

This is the authors' post-print of the manuscript titled "Extracting Reusable Design Decisions for UML-based Domain-specific Languages: A Multi-Method Study" accepted for publication by the Journal of Systems and Software. The publisher's version is available at <http://dx.doi.org/10.1016/j.jss.2015.11.037>.

(C) 2016. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>

Extracting Reusable Design Decisions for UML-based Domain-specific Languages: A Multi-Method Study

Stefan Sobernig^{a,*}, Bernhard Hoisl^a, Mark Strembeck^{a,b}

^aWU Vienna, Austria

^bSecure Business Austria (SBA) Research

Abstract

When developing domain-specific modeling languages (DSMLs), software engineers have to make a number of important design decisions on the DSML itself, or on the software-development process that is applied to develop the DSML. Thus, making *well-informed* design decisions is a critical factor in developing DSMLs. To support this decision-making process, the model-driven development community has started to collect established design practices in terms of patterns, guidelines, story-telling, and procedural models. However, most of these documentation practices do not capture the details necessary to reuse the rationale behind these decisions in other DSML projects. In this paper, we report on a three-year research effort to compile and to empirically validate a catalog of structured decision descriptions (*decision records*) for UML-based DSMLs. This catalog is based on design decisions extracted from 90 DSML projects. These projects were identified—among others—via an extensive systematic literature review (SLR) for the years 2005–2012. Based on more than 8,000 candidate publications, we finally selected 84 publications for extracting design-decision data. The extracted data were evaluated quantitatively using a frequent-item-set analysis to obtain characteristic combinations of design decisions and qualitatively to document recurring documentation issues for UML-based DSMLs. We revised the collected decision records based on this evidence and made the decision-record catalog for developing UML-based DSMLs publicly available. Furthermore, our study offers insights into UML usage (e.g. diagram types) and into the adoption of UML extension techniques (e.g. metamodel extensions, profiles).

Keywords: Domain-specific modeling, Domain-specific languages, Design decisions, UML, Model-driven development

1. Introduction

In model-driven development (MDD), a domain-specific modeling language (DSML) is a specialized modeling language tailored for graphical modeling tasks in a particular application domain, e.g. for business domains such as health-care and banking, or technical domains such as access control and system backup [1, 2, 3, 4]. MDD uses models at different abstraction levels as central development artifacts. Executable models are then derived from higher-level models through model transformations [5, 6, 7]. By raising the abstraction level in the software-development process, DSMLs aim at increasing the productivity of developers and at reducing maintenance costs [8, 2, 9].

Developing DSMLs based on the Unified Modeling Language (UML [10]) and/or on the Meta Object Facility (MOF [11]) has become a popular option in the MDD context [12, 13, 14, 15, 16]. This is because the UML/MOF infrastructure provides built-in implementation techniques for DSMLs, which allow for reusing and for extending the UML directly. As another advantage, the UML can leverage industry-grade tool support (e.g. Sparx Systems Enterprise Architect, IBM Rational Software Architect, Eclipse Model Development

Tools). UML has been subjected to scientific evaluations of its semantic foundations (see, e.g., [17, 18, 19]) and comes with standardized modeling extensions (see, e.g., [20, 21]). Moreover, the UML benefits from the systematic and continuous maintenance through the Object Management Group (OMG). Thereby, the UML and the MOF provide a rich DSML development toolkit for DSML design and implementation.

As in most software-development activities, experiences and lessons learned from developing DSMLs based on MOF/UML in a disciplined manner are particularly valuable, yet barely documented [22, 4]. Design-decision making in and on DSML development usually includes design decisions on language-model definition, constraint specification, concrete-syntax design, and platform integration [22]. Each of these concerns potentially involves multiple, interrelated design decisions.

Critical details of design decisions are rarely documented explicitly—a fact also referred to as the *capture problem* of design-rationale documentation [23, 24]. Decision details include different solutions considered before arriving at a final design decision (i.e. decision options), the decision-makers' positive and negative assessments of the considered options given a set of requirements on the DSML's design (decision drivers), and the positive and negative effects on subsequent design-decision making observed when having adopted one

*Corresponding author

Email address: stefan.sobernig@wu.ac.at (Stefan Sobernig)

or several options (decision consequences).

An important barrier to documenting design rationale in necessary detail, in DSML development and in software development in more general, is the considerable overhead that results from creating and maintaining design-rationale documentation. Other problems explored in the research on documenting design-rationale include the intrusiveness of documentation techniques, lack of incentives, and cognitive barriers in software-design processes (see, e.g., [24, 25, 26]). As a consequence, new DSML development projects cannot profit from experiences gained in prior DSML projects. Moreover, without such a body of knowledge on DSML development and evolution, existing DSMLs can become difficult to maintain.

Important goals in many fields of software-engineering (SE) research have been to limit the effort for documenting design decisions for DSML projects and to increase the quality of the documented rationale. This research has explored ways of gathering, validating, and referencing collections of *reusable generic design rationale*—such as software patterns [27]—to achieve these two goals. In our work, we build on the software-pattern tradition for capturing and for reusing SE knowledge. A software pattern documents a proven solution to a recurring engineering problem in or across different application domains, including e.g. software architecture, distributed systems, and application integration (see, e.g., [28, 29, 30]). While patterns are already useful when used in isolation, they unfold their whole potential as part of *pattern languages* [31, 32]. A pattern language documents the relations between different patterns and describes how interrelated SE problems can be solved in a proven way. In particular, a pattern language gives structure to a pattern catalog and guides the software developer in choosing a suitable selection of patterns to address a given problem.

In this spirit, in a three-year research effort, we documented recurring design decisions extracted from 90 UML-based DSMLs in terms of a decision-record catalog [33]. Similar to pattern descriptions, we document reusable DSML design decisions via a structured text format (*decision record*). Each decision record documents detailed information about a given decision-making situation (*decision point*). The catalog of decision records also describes typical combinations of decisions on developing a DSML. This way, the catalog serves as a source of reusable generic design rationale when developing DSMLs using UML.

Figure 1 shows a high-level overview of design-decision making for developing a UML-based DSML. In a decision-making situation, a decision is triggered by a decision *problem*. An exemplary problem is about how to express design constraints of a DSML. Each problem occurs in a particular *context*. For instance, one must define constraints on the DSML's language elements which cannot be expressed via graphical symbols. Decision making is affected by different *drivers* that influence the choice of a particular solution to the problem: When and where do constraints have to be checked? Are model-level checks via some graphical DSML editor needed, or instance-level checks in a particular tar-

get platform? For each problem, typically different solutions (*options*) are available; for example, defining constraints via a formal constraint language or using informal textual annotations. Each option comes with characteristic *consequences*. A consequence of formal constraints, for instance, is that they can be checked via corresponding software tools. At the same time, formal constraints are harder to change and to maintain than informal annotations.

A decision record in our catalog reflects this structure of decision-making situations by describing a recurring problem in a given context, common solution options as well as characteristic drivers and consequences. The decision-record catalog and the results of a domain analysis (see, e.g., [34]) for the DSML's application domain are then used by DSML developers to make the actual design decisions for a particular DSML. Finally, the design decisions guide the actual implementation of the DSML using UML (see Figure 1).

In this paper, we report on applying a systematic literature review (SLR) and a content analysis as methods to rigorously collect, document, and validate design decisions on UML-based DSLs. We provide examples on how the respective design decisions are documented and how they can be applied when defining a new UML-based DSML. We do not, however, discuss all design decisions and corresponding decision options identified through this study. They are documented in a companion catalog in full detail [33].¹ In other words, this paper emphasizes the research method and documents how design decisions were identified, extracted, and validated. The key contributions of this paper are:

- *Research design*: Our multi-stage and multi-method research effort took place over three years. The stages of our work include preparatory design reviews, early manual literature searches [35], a pilot study [36], construction of a quasi-gold-standard (QGS) paper corpus, an extensive automated search, and a citation-driven manual search (*snowballing*; [37, 38]). In a sequential multi-method (qualitative-quantitative) research design, we conducted an extensive systematic literature review (SLR), as well as content analysis and data mining to extract details on reusable, generic design decisions from the SLR.
- *Literature review*: The SLR was performed by adopting established guidelines [37, 39, 40]. Starting from more than 8,000 publications retrieved by automated and snowballing searches, in a two-extractor process, we selected 84 publications representing 80 UML-based DSMLs. A detailed SLR protocol is available [41].²
- *Content analysis*: Design decisions were extracted from these 84 publications using a deductive content analysis [42, 43]. This content analysis involved a coding schema (including indicators and decision rules for coders), several inter-subject iterations over the paper

¹The design-decision catalog is available at <http://epub.wu.ac.at/4312/>.

²The SLR protocol is available at <http://epub.wu.ac.at/4311/>.

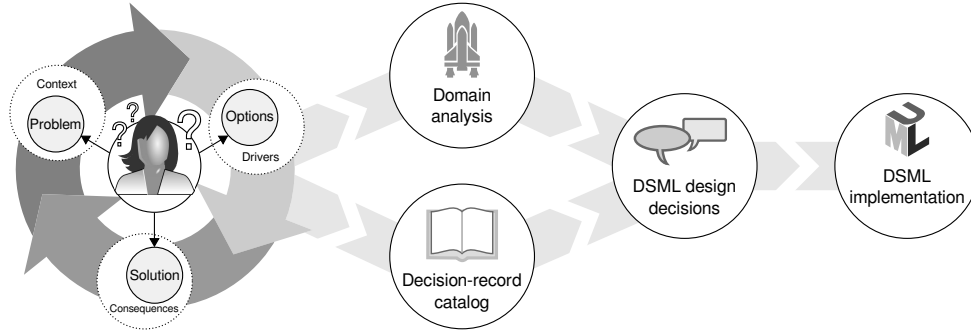


Figure 1: High-level overview of design-decision making for developing domain-specific modeling languages (DSMLs) using the Unified Modeling Language (UML).

material (segmentation, coding), and an inter-rater reliability analysis.

- *Data mining*: We performed a frequent-item-set analysis [44] to mine the design-decision data collected from the SLR for recurring decision-making patterns.

This paper is accompanied by three companion documents. First, and in line with established SLR practice, there is a detailed review protocol [41] of all steps performed to make our SLR transparent and reproducible. Second, there is a *pre-study revision* of the decision-record catalog [45] which served as input for sub-steps of the SLR (e.g. construction of a quasi-gold standard) and of the content analysis (e.g. construction of a coding schema). Third, there is a *post-study revision* [33] of this catalog which incorporates the results of our SLR as an increment to the pre-study revision.

The remainder is structured as follows: In Section 2, we justify the choice of capturing DSML design rationale via an SLR and the choice of representing design rationale in terms of structured decision records. We also motivate our research by giving two illustrative, concrete examples of using the decision-record catalog in DSML development: design-space analyses and design-process documentation. Section 3 enumerates the research questions driving the SLR study. In Section 4, we present the details of conducting and the intermediate results of the various stages of our study. The design-decision data gathered from the corresponding identified DSML designs, as well as study limitations, are described in Section 5. Important implications are discussed in Section 6. Related work on DSL development and empirical research on UML is reviewed in Section 7. Section 8 concludes the paper by reiterating over key contributions and by pointing to future work.

2. Design Rationale on DSML Development

A domain-specific language (DSL) is a software language which is specialized for addressing a given class of engineering problems, which are characteristic for an application domain. A DSL is based on abstractions aligned to this domain and provides a concrete syntax suitable for employing its abstractions effectively. A domain-specific *modeling* language

(DSML) is a DSL with a graphical concrete syntax for the primary purpose of diagrammatic modeling in a particular application domain [4, 22]. A DSML focuses on providing modeling abstractions for the problem concerns in the application domain which are independent from a given software platform, rather than on issues of implementing the domain (see, e.g., [46]). The collection of specification artifacts, which contribute to defining a DSML, is referred to as the DSML’s language model in the broader sense [22]. This language model contains a *core language-model* which describes the structural elements and their relationships using metamodeling techniques (e.g. metamodels and/or profiles in the MOF/UML). This core language-model is also known as a DSML’s abstract syntax, especially when specified using grammars. Hereafter, for brevity, we refer to a DSML’s *language model* to denote its core language-model (abstract syntax) only. As another characteristic, formal specification techniques are used in DSML development to express the structural and behavioral semantics of the DSML and its instance models [47]. A DSML is then commonly deployed as part of a model-driven development toolkit (e.g. as part of the Eclipse Modeling Framework).

For the scope of our study, we look at DSMLs which are *internal* to or *embedded* into the Unified Modeling Language version 2.x (UML 2.x; [15, 48, 49]). The language models of embedded DSMLs are defined on top of the UML 2.x or by extending the UML 2.x language model [10]. UML 2.x was rapidly adopted by DSML developers as the host modeling-language for their UML-based DSMLs. For one, this was facilitated by important modeling toolkits already providing UML 2.x support even before its final release in 2005 (e.g. Sparx Systems Enterprise Architect, Eclipse Modeling Tools). In some cases, this went hand in hand with dropping support for UML 1.x entirely. Besides, UML 2.x provides a larger base of predefined structural and behavioral diagram types (14 vs. 8 in UML 1.x), on which a DSML can be built. In addition, UML 2.x and its language architecture [50] provide for three implementation techniques for embedded DSMLs: language-model extension, language-model piggybacking, and language-model specialization. Note that a DSML development project might require a combination of these three implementation techniques.

To begin with, a DSML can be defined as a *language-model extension* [1] of UML 2.x. In this scenario, the UML 2.x metamodel is extended and modified in an additive manner. This can be achieved by first introducing new metamodel packages, which contain new or redefining metaclass definitions. These metamodel packages are defined using the Meta Object Facility (MOF) and they may reference predefined UML 2.x metamodel packages. To create an actual UML 2.x metamodel derivative, they are then merged into the UML 2.x metamodel packages [51, 52].

A second option is to have a DSML use UML 2.x as its base and add DSML-specific elements without changing the underlying structural and behavioral UML 2.x semantics (*language-model piggybacking*; [1]). This can be achieved using UML 2.x profiles [15, 48], which decorate models defined using the UML 2.x metamodel or a UML 2.x metamodel derivative.

A third DSML implementation technique applicable to UML 2.x is a variant of *language-model specialization* [1]: metamodel pruning. In metamodel pruning, an effective metamodel (i.e. the DSML language model) is extracted from an original metamodel (e.g. UML 2.x metamodel) in an automated manner [53, 54]. The extraction procedure establishes a generalization relationship between the effective (super-)metamodel and the original (sub-)metamodel. This way, DSML models defined over the smaller, effective metamodel also conform to the larger, original metamodel.

The above DSML implementation techniques for UML-based DSMLs go beyond the ones available in UML 1.x. Moreover, UML 2.x avoids known semantics pitfalls relating to these implementation techniques in UML 1.x (e.g. profiles; [50, 55]). This also determined our choice of UML 2.x.

Design rationale (DR; [23, 24]) on DSML development is the reasoning and justification of decisions made when designing, creating, and using the core artifacts of a DSML (e.g. abstract and concrete syntax, behavior specification, metamodeling infrastructure, MDD tool chain). Documenting design rationale explicitly aims at helping design-decision makers by providing and explaining past decisions (e.g. in a design-space analysis) and by improving the understanding of a DSML design during development and maintenance (e.g. as a kind of design-process documentation). Most importantly, documenting design rationale on UML-based DSML development must go beyond a mere enumeration of advantages and disadvantages specific to the three DSML implementation techniques in UML 2.x (see above). On the one hand, the choice of implementation techniques is only one decision dimension. On the other hand, the rationale behind such a choice can be specific to every DSML development project.

We distinguish between two kinds of DSML design rationale [27]: *DSML-specific* design rationale reflects reasoning required and collected explicitly during a particular design process for a single DSML. Examples of explicitly documented, specific design rationale in software-language engineering may be found in artifacts created in source-configuration management tools and development-

issue trackers and open-standards artifacts. Examples include Java Community Process documents and the recorded issue votes during the ANSI/CLISP X3J13 specification process as collected in Steele [56]. *DSML-generic* design rationale is reasoning knowledge obtained through developing multiple DSMLs, for one or several application domains. Generic design rationale is commonly found only as the implicit knowledge of experienced DSML engineers. Software patterns have been used in software-language engineering to document generic design rationale explicitly (see, e.g., [1, 57]).

An approach for managing design rationale in UML-based DSML development involves a procedure to represent design rationale, to capture it, and to put the documented design rationale to use in other DSML development projects. Representing and organizing a body of design rationale chunks deals with structuring and creating documentation items. In this paper, our emphasis is on creating documentation items on generic DSML design rationale (*decision records*) as explained in Section 2.1. Two intended scenarios for using the documented generic DSML design rationale (i.e. design-space analysis, process documentation) are presented in Section 2.2. Our approach of capturing and of transforming specific DSML rationale into generic design rationale is elaborated on in Section 4.

2.1. Representing Design Rationale

The process of developing a DSML involves different, characteristic development activities [22, 58]. From a decision-making perspective, each development activity also marks a *decision point*, i.e. a point in time at which particular design-decision problems must be addressed. In doing so, different design solutions based on their assumed or known properties for the DSML design as well as their effects on any subsequent design decisions are assessed. From the angle of design-rationale documentation, a decision point is a point in time for recording an on-going decision-making process.

In our approach, design rationale on a given decision point is captured from multiple DSML projects and represented as a reusable resource for decision making at this decision point in developing a new DSML. We refer to this reusable resource as a *decision record* (see Figure 3).³ In particular, we consider six decision points in UML-based DSML development and, therefore, six decision records (D1–D6, hereafter).

Language-model definition (D1) involves identifying the respective domain abstractions that must be represented in our DSML, after a systematic analysis and a structuring of the respective application domain. The main challenge at this decision point is how to identify and to describe these domain abstractions in order to arrive at a comprehensive and comprehensible basis for developing the DSML. *Language-model*

³Throughout the paper, we apply some notation conventions to refer to decision records and their content items such as decision options. D_i denotes a decision record corresponding to decision point i ; $O_{i,j}$ refers to decision option j at decision point i . In set notation, the O -prefix for decision options as set elements is omitted for brevity.

formalization (D2) is about how the domain definition created in D1 can be turned into a formal model. By formal model, we refer to a description that can be checked for conformance against previously defined well-formedness rules. Moreover, a formal model is amenable to processing and to manipulating by automated tools [59]. For the DSMLs studied in this paper, a formal language model is realized using well-defined metamodeling languages (MOF/UML). A metamodeling language is itself based on a well-defined and well-documented language model. The latter is referred to as the meta-metamodel, such as the CMOF for the UML metamodel. In addition, a metamodeling language provides at least one well-defined and well-documented concrete syntax to define the DSML’s language model. The DSMLs of interest in this paper can, for example, employ the CMOF diagram syntax to specify a UML metamodel extension. To remove ambiguity from a formal language model, a language model is extended to include additional *language-model constraints* (D3) to express constraints on domain abstractions, such as invariants for domain concepts, pre- and post-conditions, as well as guards. *Concrete-syntax definition* (D4) is about deciding on the “user interface” of UML-based DSMLs, looking at several options (e.g. model annotations, reuse or extend a diagrammatic syntax, mix foreign syntaxes with the UML syntax). The *behavior specification* (D5) of a DSML defines behaviors of one or more DSML language element(s). It determines how the language elements of the DSML interact to produce the behavior intended by the DSML engineers using, e.g., UML M1 behavior models or formal textual specifications. To produce executable, platform-specific model artifacts from DSML models, all DSML artifacts need to be mapped to a software platform using techniques of *platform integration* (D6; e.g. model transformations, code generators).

The order in which these development activities, and the corresponding decision points, are considered can characterize different DSML development styles [22]. Depending on the DSML development scenario some decision points may also be skipped (depending on, e.g., application domain, usage intention, and development style).

2.1.1. Decision Records

A decision record provides two or more descriptions of proven solutions to a generic and recurring problem in developing a DSML. The problem described by a decision record must not only recur, that is, be observable for many DSML development projects, but it must also have the quality of requiring an act of design-decision making. For structuring and presenting the recurring DSML design decisions as decision records, we developed and refined a document template [33, 35, 45]. Each decision record is structured into seven sections. The most important sections are identified in Figure 3, an excerpt from the actual decision record on concrete-syntax definition (D4) is depicted in Figure 2.

A decision record first describes a recurring design-*decision problem* that has been repeatedly observed for several DSML development projects. The exemplary decision record in Figure 2 gives a problem statement frequently observed when

deciding on the concrete-syntax style for a DSML: “In which representation should the domain modeler create models using the DSML?”. This problem applies to a specific *decision context*. The decision context is primarily set by one of the decision points characteristic for DSML development (D1–D6, above). In addition, a particular metamodeling toolkit (e.g. MOF/UML), the application domain modeled by a DSML, and the target software platform can contribute to establishing the decision context. To give an example: The problem statement in Figure 2 would not apply for a DSML project which is not about providing a particularly tailored or any concrete syntax to modelers. This could be because modelers are expected to work on an abstract-syntax representation only.

As its core content, a decision record lists *decision options* which describe solutions to the initial stated decision problem. The excerpt in Figure 2 shows the option MODEL ANNOTATION (O4.1) which is about realizing a tailored concrete syntax by means of model annotations. Next, a decision record documents means to select an option (or a combination of options) in terms of *decision drivers*. An exemplary driver in Figure 2 is the cognitive expressiveness of concrete-syntax styles. These drivers are likely to steer the DSML designer towards a particular option or option combination. This selection decision affects the solution spaces of subsequent decisions. For example, they can set a new decision context. To scaffold follow-up decision making, a decision record makes the DSML designer aware of known *decision consequences*. Consequences can include the need to evaluate other decision options within the same decision record or in related decision records. However, consequences can also point to follow-up decision problems not covered by the decision-record catalog alone. Such a consequence is the need for usability evaluations in Figure 2.

To provide evidence that the listed decision options are taken from observed practice, a decision record comprises DSML projects as examples of *application* of the individual options or option combinations. In Figure 2, one can see that a number of DSML projects also recorded as part of our decision-record catalog are referenced as applications. A decision record is then closed by replicating one concrete realization *sketch* of one decision option, taken from one DSML application. In Figure 2, an example of providing a combined graphical and textual concrete-syntax is given.

This document template has been derived—over multiple research stages (see Section 4)—from prior work on documenting design rationale in software engineering. In particular, it is inspired by software-pattern descriptions [32] and architectural design decisions [60]. For the content of a decision record, in particular the decision options, to qualify as generic design rationale, it must reflect *recurring* practice. To have recurrence, we require a decision option to be encountered in a certain number of DSML applications. To become included into the draft, pre-study revision of the decision records, decision options had to be accounted for by two distinct sources. For example, an option could be supported by one primary study documenting a DSML and

D4 Concrete-Syntax Definition

Problem. In which representation should the domain modeler create models using the DSML?

Context. The concrete syntax serves as the DSML's interface. Different syntax types can be defined and tailored to the need of the modeler ...

Options.

O4.1 Model annotation: Attach UML comments as concrete-syntax cues to a UML model, containing complementary domain information such as keywords, narrative statements, or formal definitions (see, e.g., [3]). ...

O4.2 Diagrammatic syntax extension: ...

Decision drivers. An overview of positive and negative links between decision drivers and available options is shown in the table below. ...

Non-diagrammatic UML notation requirements: Textual notations [1] for the UML are auxiliary representations and act as frontend syntaxes (O4.4).

Driver/Option	O4.1	O4.2	O4.3	O4.4	O4.5	O4.6	O4.7
Non-diagrammatic UML notation requirements	o	o	-	-	-	o	o
Degree of cognitive expressiveness	-	+	+/-	+/-	+/-	-	o
Disruptiveness	-	+	+	+	++	-	+/-
Degree of required modeling-tool support	++	-	+/-	+	--	++	o

Consequences.

Usability evaluation: The DSML syntax is especially important from the DSML user perspective. If a DSML is mainly used by non-programmers, a special focus on usability aspects is needed. ...

Application. In our case studies we provide a couple of different concrete syntax definitions such as UML stereotype-specific annotations for reusing symbols (P1, P3, P7, P9, P10). ...




Figure 1: Exemplary graphical and textual concrete syntax [2].

Sketch. Figure 1 shows an example of two concrete syntax definitions ...

Figure 2: Excerpt from the actual decision record on concrete-syntax definition (D4) in Hoisl et al. [33], highlighting the document sections corresponding to the key concepts in Figure 3.

one secondary study on designing UML-based DSMLs. Based on the study's results, in the post-study revision, we verified which design option has at least applications in *three* different third-party DSMLs, that is, DSMLs not developed by the authors. This way, we adopt a commonly followed rule of thumb in the software-pattern community. This rule of thumb mandates that a software-pattern description must provide at least three known uses of the pattern in existing software systems (see, e.g., [31, 32]).⁴

The catalog revised based on the study results provides six decision records—one for each decision point (D1–D6)—containing 27 decision options to describe a DSML [33].⁵ Tables 1 and 2 provide a preview of the study results, including

⁴Note, however, that the number of occurrences observed for important decision options (and decision-option sets) turned out much higher than three (see Sections 5 and 6).

⁵Note that there are actually 31 decision codes/numbers. Four of those codes/numbers serve for coding pseudo-decision options; e.g., not taking any decision. Depending on the analysis requirements, they are either ignored or included as dedicated no-option codes.

thumbnail descriptions of nine frequently adopted decision options which help characterize a majority of the DSML designs identified in our study (see also Section 6).

2.1.2. Decision-Option Sets

Making a *decision* when developing a specific UML-based DSML is about evaluating and finally adopting one or several decision options listed by a decision record in our decision-record catalog [33]. This way, a decision links to and conforms with a decision record (see also Figure 3). Design-decision making on a given DSML involves decisions at several, but not necessarily all decision points (D1–D6). Adding up all decisions results in a set of adopted decision options (an *option set*, hereafter) which represents the DSML design as product of this decision making.⁶

⁶Option sets are also an outcome of the coding step during content analysis (coding form; see Section 5.1) and a prerequisite for later data mining (see Section 3).

Table 1: Thumbnail descriptions of 12 (out of 27 total) decision options specific to the decision points 1–3 (D1–D3). 6 decision options (O1.1, O1.4, O2.2, O2.3, O3.1, O3.4; in bold font face) characterize a critical number of DSML designs, a key finding of our SLR study (see Section 6).

Problem statement	Options	Drivers
D1 How should the domain (or domain fragment) be described?	O1.1 INFORMAL TEXTUAL DESCRIPTION Use informal text to identify and to describe domain abstractions and their relationships (e.g. domain-vision statements, domain-distillation lists).	Availability of existing diagrammatic domain descriptions, intended target audience, correspondence mismatches with UML semantics, consistency preservation effort, cognitive effectiveness of a representational format
	O1.2 FORMAL TEXTUAL DESCRIPTION Use formal text to identify and to describe domain abstractions and their relationships (e.g. a grammar, a universal algebra).	
	O1.3 INFORMAL DIAGRAMMATIC MODEL Use ad hoc diagrams to identify and to describe domain abstractions and their relationships (e.g. early feature diagrams, pseudo class diagrams).	
	O1.4 FORMAL DIAGRAMMATIC MODEL Use formally defined diagrams to identify and to describe domain abstractions and their relationships (e.g. MOF and UML class diagrams, STATEMATE statecharts).	
D2 In which MOF/UML-compliant way should the domain concepts be formalized?	O2.1 M1 STRUCTURAL MODEL Implement the language model using structural UML models at level M1 (e.g. class or composite-structure diagrams).	Overlap of DSML and UML domain spaces, degree of DSML expressiveness, portability and evolution requirements, compatibility with existing artifacts
	O2.2 PROFILE RE-DEFINITION Implement the language model by creating (or by adapting existing) UML profiles (i.e. «profile» packages containing stereotype definitions).	
	O2.3 METAMODEL EXTENSION Implement the language model by creating one or several metamodel extensions (i.e. «metamodel» packages containing new metaclasses and associations).	
	O2.4 METAMODEL MODIFICATION Implement the language model by creating one or several metamodel extensions (i.e. «metamodel» packages containing redefining metaclasses and associations).	
D3 Do we have to define constraints over the language model(s)? If so, how should these constraints be expressed?	O3.1 CONSTRAINT-LANGUAGE EXPRESSION Make language-model constraints explicit using a constraint-expression language (e.g. OCL, EVL).	Constraint formalization requirements, language-model checking time, integrated language-model constraint requirements, maintainability effort, portability requirements, language model and constraints conformance
	O3.2 CODE ANNOTATION Make language-model constraints explicit using expressions in (or a specialized sub-language embedded within) a general-purpose programming language (GPL; e.g. Java).	
	O3.3 CONSTRAINING MODEL TRANSFORMATION Express language-model constraints as part of existing model (model-to-text/model-to-model; M2T/M2M) transformations, or through dedicated ones (e.g. OCL model-navigation expressions or conditional statements in ETL/EOL templates).	
	O3.4 INFORMAL TEXTUAL ANNOTATION Use informal and unstructured text annotations to capture constraint descriptions in the language model (e.g. prose text in UML comments).	

Table 2: Thumbnail descriptions of 15 (out of 27 total) decision options specific to the decision points 4–6 (D4–D6). 3 decision options (O4.1, O4.6, O6.2; in bold font face) characterize a critical number of DSML designs, a key finding of our SLR study (see Section 6).

Problem statement	Options	Drivers
D4	In which representation should the domain modeler create models using the DSML?	Non-diagrammatic UML notation requirements, degree of cognitive expressiveness, disruptiveness, degree of required modeling-tool support
D5	Do we have to define (additional) behavioral semantics for the DSML? If so, how should the additional behavior of DSML elements be defined?	Model-consistency preservation, domain characteristics, limited expressiveness, behavior verification, visualization preferences
D6	How should the DSML artifacts be mapped to (and/or integrated with) a software platform?	Targeting multiple platforms, maintainability effort of static code fragments, non-executable models
	O4.1 MODEL ANNOTATION Attach UML comments as concrete-syntax cues to a UML model, containing complementary domain information such as keywords and narrative statements.	
	O4.2 DIAGRAMMATIC SYNTAX EXTENSION Extend one or multiple UML diagram types by creating novel symbols adding to the basic UML symbol set.	
	O4.3 MIXED SYNTAX Create your DSML's concrete syntax either as a non-diagrammatic syntax (textual or tabular) or as a diagrammatic syntax not integrated with the UML's.	
	O4.4 FRONTEND-SYNTAX EXTENSION Create your DSML's concrete syntax as a non-diagrammatic (textual or tabular) one which extends a non-diagrammatic frontend syntax to the UML (e.g. HUTN).	
	O4.5 ALTERNATIVE SYNTAX Create a diagrammatic syntax extension to the UML (O4.2) and provide one or more alternative syntaxes (see O4.3 and O4.4).	
	O4.6 DIAGRAM SYMBOL REUSE Reuse built-in UML diagram symbols without modification.	
	O5.1 M1 BEHAVIORAL MODEL Specify additional behavior of language-model elements using UML behavioral models at level M1 (e.g. state machine or activity diagram).	
	O5.2 FORMAL TEXTUAL SPECIFICATION Specify the additional DSML behavior using a textual formalism (e.g. algebraic expressions or Z notation).	
	O5.3 INFORMAL TEXTUAL SPECIFICATION Specify the additional DSML behavior using an informal textual description.	
	O5.4 CONSTRAINING MODEL EXECUTION Implement behavioral constraints in a (partial) execution engine for DSML models (e.g. xMOF).	
	O6.1 INTERMEDIATE MODEL REPRESENTATION Provide for generating a second and intermediate model based on a DSML model using M2M transformations.	
	O6.2 GENERATOR TEMPLATE Create transformation templates, defined in a M2T transformation language, which turn DSML models into platform-specific textual specifications.	
	O6.3 API-BASED GENERATOR Realize the platform-specific model transformation (e.g. code generation) by instrumenting a programmatic representation of DSML models.	
	O6.4 (DIRECT) MODEL EXECUTION Use a (partial) model-execution engine to generate platform-specific instructions directly from DSML models.	
	O6.5 M2M TRANSFORMATION Perform platform integration via (multiple) endogenous M2M transformations specified via M2M transformation languages (e.g. ATL, ETL).	

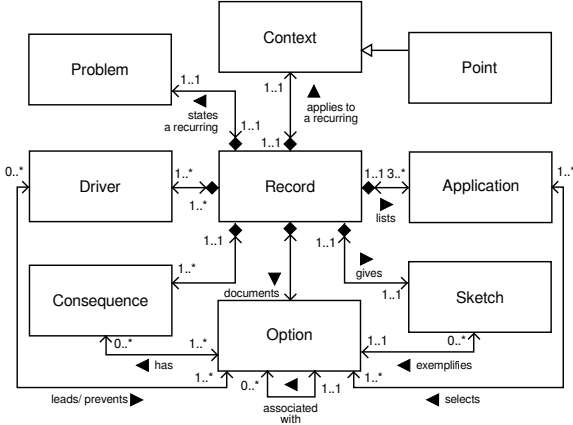


Figure 3: An overview of the 9 key concepts and their relationships: decision point, decision context, decision problem, decision record, decision option, decision driver, decision consequence, decision application, and decision sketch.

Consider the example of UML4SOA [61], one of the 80 third-party UML-based DSMLs reviewed in our study (see Section 5 and also the Appendix). UML4SOA refines the UML activity, class, and component diagrams to model service-oriented architectures (SOA). The language model of UML4SOA is defined textually (O1.1) and integrates with the UML via a UML metamodel extension (O2.3) as well as equivalent UML profile definitions for tool adoption (O2.2; see Table 1). In addition, the metamodel extension and profile definitions are accompanied by OCL constraint definitions (O3.1). The metamodel extension comes with new and resampled diagram symbols (O4.2), the profiles imply model annotations (e.g. comments containing tags; O4.1) and symbol reuse (O4.6; see Table 2). As for platform integration, UML4SOA employs intermediate model representations (O6.1) to transform extended UML activities in several steps (O6.5) into web-service orchestration specifications (BPEL) using API-based generators (e.g. the Eclipse/EMF Java API; O6.3). Based on our catalog of decision records, UML4SOA can be described as an option set containing ten options from our decision-record catalog [33]: {1.1, 2.2, 2.3, 3.1, 4.1, 4.2, 4.6, 6.1, 6.3, 6.5}.

In addition, option subsets are eligible to represent *decision-option associations*. An association between two or multiple decision options represent a possible, intentional co-occurrence of two (or more) corresponding decision options. An association denotes that two or more decision options must be considered together, without implying any particular (e.g. temporal) order of adoption. Decision-option associations can represent decision drivers and decision consequences which directly relate to decision options. Consider the proper subset {2.2, 4.1, 4.6} contained by UML4SOA’s option set above. It denotes that—as a consequence of adopting UML profiles (O2.2)—UML4SOA reuses diagram symbols of the activity diagram notation (O4.6) which become extended by model annotations (O4.1; e.g. stereotype tags, tagged val-

ues). As we will learn in Sections 5 & 6, this option subset can be frequently observed in DSMLs. Our catalog currently documents 21 such decision-option associations [33].

2.2. DR Reuse: Design-Space Analysis

In this section, we exemplify the use of our decision-record catalog. The decision records facilitate re-constructing the decision space of an existing DSML using Questions-Options-Criteria (QOC) diagrams. Note that we do not consider our catalog to be the only and authoritative source of input in this usage scenario. Rather, they complement existing material, such as pattern collections on software-language development (see, e.g., [1]).

An important area for using documented design rationale is to facilitate the creation of a systematic description of an existing design which reflects the reasoning of the designers. Here, the objective is an explicit description which is sufficiently detailed to assist in decision making and stakeholder communication during ongoing phases of an artifact’s lifecycle. Relevant phases include test creation, corrective maintenance, and authoring documentation. An established semi-formal, diagrammatic approach for describing a space of design arguments is the Questions, Options, and Criteria (QOC) notation [24, 62].

Questions deal with features of the designed artifact. In the example from Figure 4, such a feature is the representation of audit rules in the technical domain of system auditing. An audit rule determines which system events should be audited by an auditor component or subsystem. Consider file-access events by selected system users as an example. The DSML SecurityAudit provides UML-based support to model such audit rules as part of an auditing system. The DSML representation should allow domain modelers (a security auditor) to author audit rules. In QOC, *options* model possible answers to the questions. The exemplary design space shown in Figure 4 of the SecurityAudit DSML consist of three options. Clearly, the actual design space of the SecurityAudit DSML was larger. Two QOC options have correspondences to decision options from our catalog (D4: O4.2, O4.3, O4.5; see Figure 4). Depending on the design context, a QOC option might map to several decision options in our catalog or vice versa. *Criteria* indicate properties or effects of adopting given options, e.g., to satisfy certain requirements. Our catalog describes decision drivers for adopting or discarding certain decision options. The three relevant decision options (O4.2, O4.3, O4.5) describe, for example, four drivers which are reproduced in Table 3 (see also decision record D4 in [33]). The decision records enumerate positive and negative links between drivers and options. As such, our catalog offers candidate criteria to be adopted in a QOC analysis. In Figure 4, two drivers are referenced as QOC criteria (disruptiveness, modeling-tool support). The positive and negative links form the basis for the *assessment* of the QOC options to answer the QOC question. Finally, a QOC *decision* denotes the act of marking QOC options as adopted (see the solid rectangles in Figure 4). Each decision may be followed by

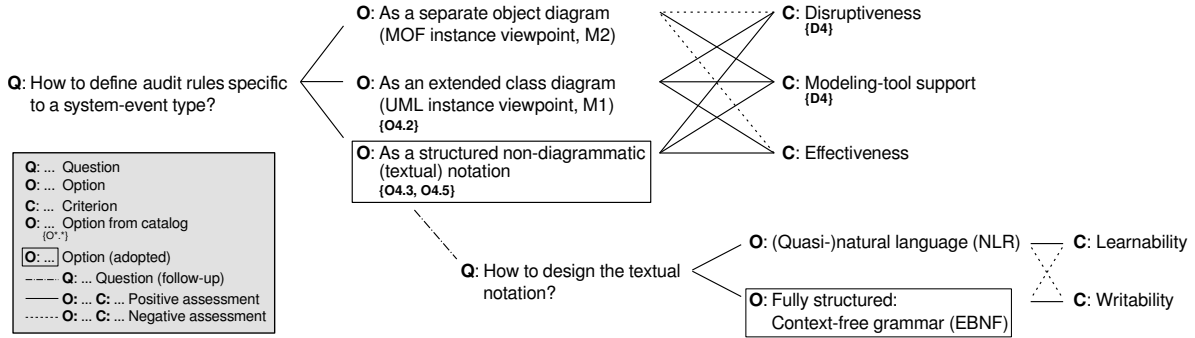


Figure 4: A partial QOC representation of a selected design-space fragment for the DSML SecurityAudit. The boxed options are the decisions made in the design of the DSML. This example follows the original QOC notation introduced in MacLean et al. [62]; advanced QOC notation elements such as arguments are omitted for brevity.

another QOC question. In our example, the choice of providing a structured, textual notation is succeeded by a question on the concrete-syntax style to be used.

Table 3: Exemplary overview of positive and negative links between criteria (or drivers as found in our catalog) and available options depicted as a rationale table. They form the basis of the assessments shown in the QOC representation in Figure 4. (+)+: (very) positive influence; o: no influence; (-)-: (very) negative influence. An option having either a (very) positive or a (very) negative influence—depending on the intended DSML’s application domain, professional background as well as prior knowledge and experience of users etc.—is denoted by (+)+/(-)-.

Driver/Option	O4.2	O4.3	O4.5
Disruptiveness	+	+	++
Modeling-tool support	-	+/-	---
Cognitive expressiveness	+	+/-	+/-
Non-diagrammatic UML notation	o	-	-

Systematically linking QOC diagrams and structured documentation of design decisions to perform design-space analyses has already been proposed. Zdun [63] integrates QOC diagrams and software patterns for systematic pattern selection.

3. Research Questions

By describing and documenting DSML designs through the lenses of decision points and decision records, we take a decision-makers’ perspective on a DSML design. On the one hand, a decision-making perspective targets the stakeholder roles relevant for DSML design-decision making (e.g. business designer, software architect, domain engineer, DSML implementer, domain modeler). On the other hand, such a perspective facilitates recording the results of decision making.

A decision-making perspective complements development-process perspectives on DSML development, such as the one proposed by Strembeck and Zdun [22]. A process perspective puts emphasis on the actual development activities and development artifacts in DSML development.

However, a decision-making perspective and the resulting design-decision documentation shifts focus on DSML development at a finer grained level of abstraction and makes decision interdependencies explicit—independent of the actual development-process style [22].

Recording and presenting repeatedly observed design decisions has the potential of facilitating the documentation of design decisions in other development contexts. Such contexts can be set by a new DSML project or by a maintenance task on an existing DSML. In Section 2.2, we make the case for a systematic reuse of previously gained process and decision knowledge (e.g. options, drivers, consequences) in design-space analyses. To this end, our study was guided by the following research question:

Research question 1: *What are the design-decision options for UML-based DSML designs reported in scientific literature?*

To answer this question, we conducted a systematic literature review (SLR) and a systematic content analysis. In its preparation, execution, and reporting, we apply established guidelines [37, 38, 39, 40]. Section 4 summarizes the SLR procedure and its results. An overview of the content analysis of the SLR paper corpus is given in Section 5. The combined SLR and content-analysis protocol can be found in Sobernig et al. [41].

Each UML-based DSML design can be described by an option set that contains a particular combination of the options described by the decision-record catalog.⁷ Therefore, at first glance, our catalog of 27 decision options defines an extensive design space of decision-option sets (see Tables 1 and 2): In a convenience view, one can generate $2^{27} - 1$ unique combinations of these 27 decision options to characterize DSML designs. This convenience view neglects any combinatorial constraints, such as the ones imposed by documented associations between decision options.

⁷Option sets are also an outcome of the coding step during content analysis (coding form; see Section 5.1) and a prerequisite for data mining (see RQ2).

We require that a DSML design must at least report one option on language-model definition (D1) and another one on language-model formalization (D2). This requirement follows from the following three assumptions: First, we only consider DSMLs that were designed in a language-model-driven development style [22], which implies at least one decision on D1. Second, at least one of the corresponding options for the extension of the UML must be chosen (D2) in order to qualify as a UML-based DSML. Third, backed by a documentation analysis based on scientific publications, we expect that these two design dimensions are mandatory in scientific reports on DSML projects. This might even be the case when they were not explicitly addressed in the actual design-decision making of the DSML designers.

Even when considering the above presence conditions on D1 and D2, our catalog still allows for expressing a vast space of distinct options sets for DSMLs: 117,964,800 unique option combinations! This number computes as follows starting from the 27 options in our catalog: $2^4 - 1$ (D1) times $2^4 - 1$ times (D2) times 2^{27-4-4} (D3–D6). In practice, however, we expect a reduced observable design-decision space across existing DSMLs. In particular, there is partial evidence suggesting that certain UML extension techniques are more commonly adopted than others (e.g. profiles [13, 64]). This leads to our second research question:

Research question 2: *What are frequently observed decision options and frequently observed combinations of decision options (option sets) in and across existing UML-based DSML designs?*

Collecting option sets that represent existing DSML designs allows for important insights beyond their mere observability. If several existing DSML designs exhibit identical option sets, then such *recurring* option sets indicate repeated decision-making practice in designing DSMLs. At the same time, repeatedly observed option sets can serve for grouping different DSML designs into families that share characteristic combinations of design options.

To address our second research question, we are thus interested in frequency patterns that occur for the collected option sets. Such frequency patterns can help in characterizing an empirically *observable* subset of the theoretically possible design space that is described by our catalog. Detecting and interpreting such decision patterns (e.g. hub decisions) has been reported as an important use case for design-rationale documentation [65].

In order to characterize the observable design-decision space for UML-based DSMLs, we mine for frequent option (sub-)sets using an analysis that is based on *frequent item sets* [44, 66]. We are interested in commonly recurring combinations of decision options which are (proper) subsets of observed option sets. Technically, we want to extract option (sub-)sets that adhere to certain constraints (i.e. minimum support, closedness, freeness, maximality; see [44, 66]). Sobernig et al. [41] provide a technical background on those concepts in the context of our study.

Such frequent option (sub-)sets can express characteristic fragments of a DSML design as well as complete DSML designs (also called “prototype option-sets”). These option sets can differ in terms of the number of options (or, proper option subsets) that they contain (size), in terms of their relatively higher or lower levels of support, and whether they are contained as-is in the base of observed option sets or not; i.e., whether they totally describe at least one DSML design alone, rather than a fragment of it. Table 4 summarizes the three kinds of characteristic option (sub-)sets that we consider in answering our second research question. Section 5 elaborates on our findings.

4. A Review-driven Approach to Capturing Generic DSML Design Rationale

To answer our research questions, we set out to distill *generic* DSML design rationale from a maximal number of DSML design documents. This way, the extracted design rationale is more likely valid beyond single DSML designs. To the best of our knowledge, no such source of generic rationale on UML-based DSML designs existed prior to our decision-record catalog. The secondary studies on DSML development, which we identified as related work, do not reflect actual design-decision making in DSML development projects; specifically not for UML-based DSMLs (see Section 7). Therefore, our emphasis was on gathering *primary* studies on UML-based DSMLs. A primary study is a piece of design documentation authored by the respective DSML developers themselves. In scientific literature, primary studies come as DSML solution proposals and/or personal experience reports [67].

Capturing design rationale (DR) behind a UML-based DSML design can be achieved in different, systematic ways [24]. First, one can recover design decisions by reviewing DSML artifacts after the fact (e.g. abstract-syntax or concrete-syntax specifications). This can either be done by the DSML developers themselves or by third-party experts in DSML development and in DR documentation (see also Section 2.2). Second, DSML designers might record their rationale themselves as a byproduct of the decision-making process. A third source are records of communication created by DSML designers, for example, language-user documentation, change/maintenance documentation, and scientific publications. Finally, if available, design-support software can be used for documenting design-decisions (e.g. IDEs including support for design-knowledge management [68]). To the best of our knowledge, however, contemporary design-support software for DSMLs does not provide DR capturing facilities.

As in other fields of software development, the *capture problem* [23, 24] often prevents DSML design rationale from becoming documented explicitly by DSML developers. In earlier, preparatory studies including manual design reviews, a snowballing study, and a pilot SLR, we found, however, that *scientific publications* of DSML developers are important primary studies for documented DSML design ratio-

Table 4: Overview of the option-set constructs considered for analyzing frequency patterns in the selected DSMLs. Each construct is defined as a set of data-mining restrictions (e.g. closedness, maximality, freeness [44]) over the space of (frequent) option sets representing the reviewed DSMLs. Details on these underlying data-mining restrictions are provided in Sobernig et al. [41].

Kind of option (sub-)set	Description
Smallest common option subset	A frequent option subset which is also a smallest (i.e. of minimal size) recurring <i>proper</i> option subset contained by observed DSML designs and/or by observed design fragments. We distinguish between two kinds of smallest common option subset: (1) option subsets specific to one decision record (D1–D6); (2) option subsets specific to two or more decision records (D1–D6).
Prototype option-set with frequent extensions	A frequent option set which represents a largest option subset (design fragment) which was also frequently found to represent complete DSML designs. This prototype option-set is frequently found extended by adding other (frequently observed) options. In this sense, it represents an <i>evolutionary</i> prototype to derive extended DSML designs.
Prototype option-set with infrequent extensions	A frequent option set which represents a largest option subset (design fragment) which was also frequently found to represent complete DSML designs. Extensions that add options to this (<i>evolutionary</i>) prototype for deriving other option sets (DSML designs) are infrequent.

nale [35, 36]. Scientific publications must not necessarily document the rationale for design decisions directly. Often, they provide a systematic overview of DSML design artifacts and references to them (e.g. abstract-syntax or concrete-syntax definitions). Moreover, DSML developers submitting to relevant scientific publication venues (e.g. SoSyM journal, MoDELS conference) are more likely to report UML-specific design decisions explicitly. They adopt best-practice examples for design documentation from OMG standard documents directly, from already published papers in the respective venues, and from secondary studies on UML-based DSML development (e.g. [69, 70, 71, 72]). Finally, for some DSMLs being research prototypes only, scientific reports are often the only source of documented design decisions.

For these reasons, we opted for a systematic literature review (SLR). The main goal of this SLR was to identify a maximum number of scientific publications which document design rationale on UML-based DSMLs as primary studies. The targeted scientific publications were required to be marked by a sufficient documentation quality. By documentation quality, we refer to the correctness and the completeness of the design documentation.

4.1. Planning & Conducting the Review

The SLR was performed in three steps (see Figure 5). First, to provide a basis for evaluation of the search procedure, we established a corpus of reference publications (viz., a quasi-gold standard, QGS; see Section 4.1.1). Based on this reference corpus, we identified the search engines for an automated publication search. We refined the corresponding search terms and the queries in several iterations. Second, we performed the actual engine-based publication search (see Section 4.1.2). Based on the bibliographical records extracted publications selected up to this point, we then performed a backward-snowballing search (see Section 4.1.3). Backward snowballing is the practice of manually identifying additional publications for selection from the reference lists (citations) of a given set of publications [37, 38].

We selected publications for inclusion and assessed their quality based on predefined criteria at each stage. The selection and quality-assessment decisions involved multiple raters per publication (i.e. the authors). Therefore, we report the inter-rater reliability (IRR) for the data extracted from the selected publications. IRR measures document patterns

of agreement and disagreement between two and more raters in their assessments of details extracted from the selected publications [73]. This section summarizes the adopted procedures at each stage (working tasks, criteria selection), the intermediate results (QGS, review, and snowballing corpora), and the intermediate evaluations (validity, reliability measurement).

4.1.1. Quasi-Gold Standard

To guide publication search, and to report the search validity, we developed a quasi-gold standard (QGS; [40, 74]). A quasi-gold standard is a collection of manually selected papers from a number of venues and outlets (e.g. journals, conference series), which are known and recognized for publishing work on the topic areas under review by a relevant community; in our case: model-driven development, UML, and DSMLs. The results of the automated literature search must include the quasi-gold standard. Otherwise, the search strategy must be adapted.

As opposed to a “gold standard” collection of papers for evaluating literature searches, a *quasi-gold standard* is specific to a set of venues/outlets over a specified time span [40]. We used the QGS to guide the main, semi-automated literature search (i.e., to select the search engines and to guide the selection of search terms) and, finally, to validate its quality attributes (i.e. sensitivity and precision).

Constructing our own QGS corpus became necessary because, at the time of designing the SLR, there was no adequate third-party publication corpus available. In 2013, the only candidate was a prior mapping study by Nascimento et al. [13] on application domains of DSLs, DSL development tooling, and research agendas on DSL engineering. However, our study aimed at a specific subset of DSLs (viz. UML-based DSMLs). Nascimento et al. [13] reviewed 170 publications on DSMLs and only 21 of these were specific to the UML. Finally, while covering an extensive publication period (1966–2011), the mapping study by Nascimento et al. [13] does not extend to 2012 as required by our study. To the best of our knowledge, there is still no adequate third-party corpus available.

The publication collections obtained from our earlier research steps [35, 36] did not qualify as a reference corpus either. This was because of their strong bias towards our research projects and our own experiences. Furthermore, the

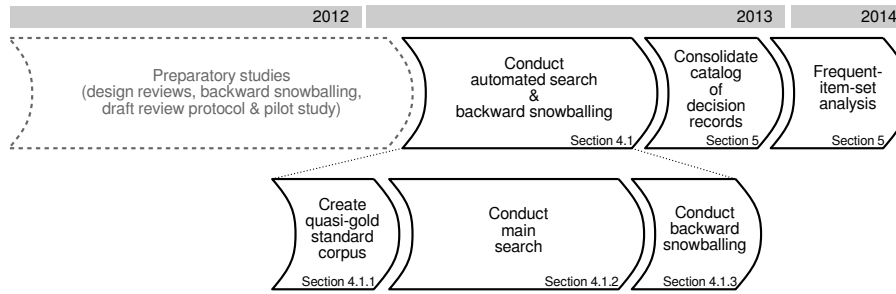


Figure 5: A timeline overview of the research stages and the corresponding sections in this paper.

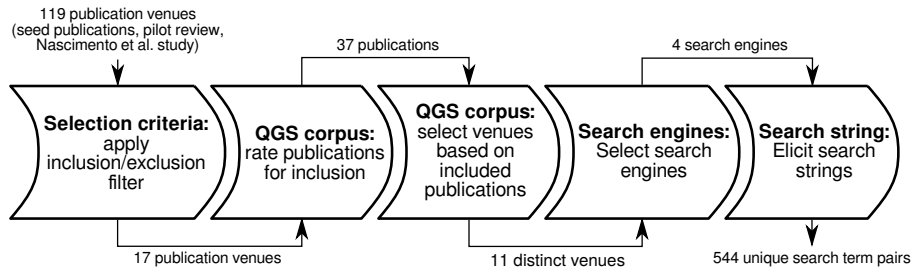


Figure 6: Sub-steps of establishing a QGS corpus and eliciting search strings for the main search.

publications from our prior research [35] were also limited to DSMLs for security-related application domains. The third-party publications collected during our pilot study were not considered as QGS candidates because of flaws in our pilot SLR design [36].

Figure 6 visualizes the five sub-steps of QGS construction. As a starting point, we considered all publications collected from the three sources above. These included our own DSML publications [35], the third-party publications from our pilot review [36], and the third-party publications returned by the mapping study [13]. The three sources accounted for 159 papers and 119 publication venues in total. The publication venues included 86 presence venues (conferences, symposia, and workshops) and 33 archival venues (journals, monographs; see also Figure 7).




Figure 7: Overview of the stepwise selection of publication venues for establishing the quasi-gold standard (QGS) corpus.

Figure 7 summarizes the intermediate results of filtering. Finally, 17 venues were selected. The seven journals included top-tier SE venues such as IEEE TSE, ACM TOSEM, and SoSyM. The ten selected conference venues comprised i.a. ICSE, OOPSLA, MoDELS, and ASE.⁸ The selected venues complied with four criteria:

1. *Time coverage*: A publication venue covers the years between and including 2005 and 2012. In 2005, the UML 2.0 specification was published. 2012 marks the year before performing the engine searches in January and February 2013.
2. *Community relevance*: An archival venue is deemed relevant by a scientific audience if it is listed with the ERA 2012 journal list [75]. The ERA list resulted from a public and international consultation process among scientists. 23 of the 31 journals were listed with ERA 2012. For a presence venue, we examined whether the venue has a regular publication history depending on the venue format (e.g. yearly, bi-yearly) during the review period. This was the case for 28 of the 75 conferences.
3. *SE focus*: The venue has a dedicated software-engineering (SE) focus. The two authors judged a venue based on publicly available SE venue lists, including SCImago Journal Rank [76] and Microsoft Academic Search [77]. 17 out of the remaining 51 venues fulfilled this condition.
4. *Content maturity*: The venue is committed to publishing mature and scientifically rigorous content. In general, we verified whether there was a peer-review procedure

Selection Criteria. In a pair session, the 119 publication venues were filtered by two authors (Hoisl, Sobernig). Fig-

⁸See Appendix A in [41] for the venues' full names.

in place. As for presence venues, 11 workshops were excluded. Two archival venues other than journals were discarded (i.e. a festschrift and a project-report monograph).

Quasi-Gold Standard Corpus. Two authors manually screened the selected 17 venues. This involved 418 journal volumes and 80 proceeding issues published between 2005–2012. In this initial iteration, candidate papers were selected by reviewing title, author-provided keywords, and abstracts (in this order). The reviewers used centrally maintained publication-history records of each venue (e.g. DBLP⁹) and the publication bases of the venue publishers (i.e. IEEE, ACM, Springer, Elsevier) to additional important metadata (e.g. abstracts).

The screening result was a collection of 83 articles (52 journal and 31 proceedings articles). Each article was rated for inclusion into the QGS corpus independently by two authors. 37 publications were positively rated by both authors and formed the final QGS corpus (24 journal and 13 proceedings articles). The two independently rating authors arrived at the same selection decision (inclusion or exclusion) for 75.9% of the 83 publications without negotiation. The conflicting selection decisions were first revisited in a joint session between the two rating authors. Then, if necessary, the third, non-involved author adjudicated on an article at issue.

Any inter-subjective rating process bears the risk of random ratings and personal bias, so that the percent agreement of 75.9% is likely to be overstated [73]. Applying standard chance-correction (Cohen’s Kappa), we established that the two rating authors achieved actual decision agreement beyond chance (inclusion/exclusion) in approx. 51% of the total expect cases beyond chance. In technical terms, Cohen’s Kappa coefficient $\widehat{\kappa}_C$ on the 83 publications and the corresponding 166 ratings amounted to 0.51 ± 0.009 .¹⁰ Kappa’s chance correction results from a worst-case assumption (random assignment), the “true”, but unknown agreement level lies between 0.51 and 0.759. According to standard Kappa benchmarks [78], this indicates an intermediate to good chance-corrected (worst-case) reliability level. Therefore, we considered the underlying selection procedure (criteria) reliable. Any decision conflicts were resolved as outlined above.

Comparatively equal shares of candidate articles published in journals and conferences, respectively, made it into the QGS corpus. Approximately 46% of the journal articles (24/52) and 42% of the proceedings articles (13/31) entered the corpus. Figure 8 depicts the number of finally included publications per venue. JOT, JSS, SoSyM, and SAFECOMP were the top four venues contributing publications to the QGS corpus. These top four accounted for $\approx 60\%$ (22/37) of the included publications. From six venues, not a single

publication was selected. The distribution over time shows that $\approx 49\%$ of the included publications (18/37) were published in two peak years: 2007 and 2011 (see Figure 9).




Figure 8: Publications found during manual screening of the 17 QGS venues and publications finally included into the QGS corpus (per venue; ordered by the decreasing number of found papers). See Appendix A in Sobernig et al. [41] for the venues’ full names.




Figure 9: The numbers of publications *found* (dashed line) and *included* (solid line) into the QGS corpus, per publication year (2005–2012).

Based on the 37 QGS publications, two subsequent steps were performed. On the one hand, the relevant search engines for the automated search were identified. On the other hand, a search string for the automated search was constructed from the QGS corpus.

Search Engines. The 37 QGS publications were published in eleven distinct venues, which are issued by five different publishers: Springer (15 papers), AITO (7), IEEE (7), Elsevier (6), and ACM (2). For these publishers, we identified four search engines: SpringerLink, IEEE Xplore, Scopus, and ACM Digital Library. These met previously defined requirements, such as full time-coverage, access to bibliographical metadata, and minimal content overlap [41].

Search String. We extracted the search string from the QGS publication corpus in a systematic and subjective manner [40]. The extracted 49 search terms were grouped into two term sets: 17 search terms specific to a (meta-)modeling technology (e.g. “uml”, “mof”) as well as 32 search terms

⁹See <http://dblp.uni-trier.de/>, last accessed: Feb 2, 2015.

¹⁰ \pm signals the leave-one-out Jackknife variance estimate $Var(\widehat{\kappa}_C)$ of Cohen’s Kappa statistic $\widehat{\kappa}_C$ to quantify the degree of statistical insecurity inherent in the data-generating process.

specific to DSL and DSML development (e.g. “profile”, “meta-model”). The resulting search string, from which the four engine-specific search expressions were derived, is documented in [41]. This final search string was developed in several iterations to maximize search sensitivity. Search sensitivity reflects the number of QGS publications retrieved by a given search string [40]. The iterations are reported in full detail in [41].

The final search string represents pairs of search terms, with each pair drawing one search term from each term set. More precisely, the terms within each set were considered alternatives (exclusive-or; e.g. “uml” *xor* “mof”). Then, to arrive at term pairs, the two sets were merged element-wise using conjunction. The final query expression, therefore, represented 544 unique term pairs (e.g. “uml” *and* “profile”). This tactic met a triple objective: First, we arrived at more precise search terms regarding our study context. Second, this tactic minimized any semantic overlap between terms related to technologies and DSML engineering, respectively. Third, all four engines supported this structured search string. We could enter the search string as one concatenated boolean expression in a disjunctive-normal form (leaving aside minor concrete-syntax differences), rather than in its complex expansion into unique term pairs.

4.1.2. Main Search

In this step, we conducted the automated publication search using the search string developed in the previous step on the four selected search engines: ACM Digital Library, IEEE Xplore, Scopus, and SpringerLink. Figure 10 provides an overview of the sub-steps of this automated search.

Engine-based Search. This step involved four activities: search execution, duplicate cleansing, validity computation, and QGS-based capping. *Search execution* yielded 5,778 search hits split into four result sets, one for each of the four search engines. When extracting the results, we applied and preserved the relevance-based sorting of search hits (as provided by each engine).

From the ACM Digital Library, 933 hits were retrieved. IEEE Xplore returned 1,845 hits. Searching Scopus and SpringerLink yielded 2,000 and 1,000 hits, respectively (see Figure 11). IEEE Xplore, Scopus, and SpringerLink cap the result sets when exporting them into a processable format. For IEEE Xplore, which enforces a capping at 2,000 hits, the restriction did not apply given a smaller number of actual hits. As for Scopus and SpringerLink, the numbers reported above correspond to the capped result sets. The actual result sets of Scopus and SpringerLink amounted to more than 16,500 and more than 8,500 hits, respectively.

The four engines demanded their own variants of the complex search string (see Section 4.1.1). This was mainly due to different concrete syntaxes of the query processors. In addition, all provided different means to limit the searches to our required range of publication years (2005–2012) and to publications having a full-text body written in English lan-

guage. The resulting four search strings are documented in the SLR protocol [41].

Duplicate detection was performed for each result set and across the four result sets. In total, we removed 1,083 hits as unwanted duplicates (see Figure 11). These duplicated hits were almost entirely caused by redundant hits between two and more result sets, rather than replicated entries within one result set alone. As a minor exception, however, Scopus contained ten duplicated hits within its result set. Duplicate detection was performed in three passes: First, we used the Document Object Identifiers (DOI) provided by the searched publication data bases for their hits. 5,232 of the 5,778 hits were equipped with a DOI. Second, for the publications lacking a DOI, we matched their publication titles character-wise and in a case-insensitive manner. Third, to overcome possible barriers of the conservative, exact matching strategy, we matched hits based on their tokenized titles and a Jaccard similarity function [79]. Barriers included the heterogeneity in the titles in the presence of punctuation, encoding artifacts, and typos. Duplicate removal was performed in a manner preserving the relevance sorting of the four result sets.

Validity computation was performed on the remaining, unique 4,695 hits (see Figure 11). For this, we computed the quasi-sensitivity of the engine-based search with respect to our QGS corpus. The overall objective was to obtain a search sensitivity at a level between 70 and 80% [40]. In other words, between 70 and 80% of the QGS publications should be contained in the collated result set. We arrived at an overall quasi-sensitivity of $\approx 75.7\%$ for the main search: 28 of 37 QGS publications were successfully retrieved across all four search engines. Scopus contributed 14, SpringerLink seven, ACM Digital Library four, and IEEE Xplore three QGS publications. Note that this validity computation was repeatedly performed, actually, over several search iterations to maximize the sensitivity.

Finally, we applied a *capping* of the search hits based on the retrieved QGS publications. We defined a cutoff at the position of the last QGS publication found in each result set. For example, the last QGS publication in the set of 853 unique ACM Digital Library hits was listed at position 781. Thus, position 781 became the cutoff position for result set of ACM Digital Library (see Figure 11). This capping strategy was justified because of the substantial quasi-sensitivity level achieved with the engine searches (see above). After capping, 2,678 hits or 46.4% of the total search hits entered the manual selection procedure.

Selection Criteria. The 2,678 papers in the cleansed, reduced, and collated result set were evaluated by two authors according to the nine selection criteria (i.e. four venue-specific and five publication-specific ones), yielding a total of 5,709 inclusion and exclusion decisions, respectively. After this selection step, 106 of the 2,678 papers ($\approx 4\%$) were included. The nine selection criteria accurately reflect the study’s objective. To give one example, the fact that we were interested in DSML designs based on the UML 2.0 led us to impose 2005 as the starting year, given that the UML 2.0 spec-




Figure 10: Sub-steps of conducting the main search.




Figure 11: Overview of cleansing (duplicate detection) and reduction steps (sensitivity cutoff) for each result set, yielding in total 2,678 papers entering the selection procedure.

ification was formally released in July 2005.

The first four criteria were specific to the publication venue of a given search hit. These four criteria corresponded to the ones already applied for selecting the QGS venues: time coverage (2005–2012), community relevance, SE focus, and content maturity (see Section 4.1.1). Testing the venue-specific criteria first allowed us to make a selection decision on papers sharing a venue at once. This helped us considerably to cope with the extensive result set.

Once satisfying all venue-specific criteria, a paper was checked for five publication-specific criteria. Most importantly, it was established whether the paper’s full-text was accessible to us, e.g. whether our institutional subscriptions to the publisher’s digital libraries covered these items, or whether an author copy could be retrieved otherwise. Only 47 papers could not be accessed at all. Based on the full-text, it was then decided whether a paper reports an actual DSML design. This decision involved several checks regarding the paper type (e.g. excluding editorial introductions or position papers) as well as the paper’s scope as either primary or secondary studies on DSMLs.

To determine whether a candidate paper and the corresponding DSML design were based on the UML 2.x, we first checked the paper’s references list for citations of the corresponding OMG UML specification documents. If missing, we reviewed the full-text for clarifying statements and the various design artifacts. If available, we also attempted to infer the UML version dependencies from the concrete-syntax elements used in diagrams. In parallel, we verified whether all necessary details of the DSML design are reported by the candidate paper. For example, we verified whether a valid specification of a UML profile and its elements is provided.

In total, approximately a quarter of the papers (i.e. 704 out of 2,678) were excluded because they did not comply with the report-type criterion. Papers not reporting on UML-based DSML designs or a secondary study on DSMLs fell into this group. 152 of the excluded publications were checked and agreed on by two authors independently from each other.¹¹ Approximately, 70% of the papers (1,868/ 2,678) did not satisfy all of the eight remaining venue- or publication-specific criteria.

Quality Assessment. The 106 papers that were included at this point were then further assessed for two quality properties: 1) duplicates and 2) erroneous DSML designs or design reports. In total, 33 papers did not pass this assessment step. Given the inherent limitations of the semi-automated duplicate detection (see above), quality assessment involved manual checking for duplicates when working on the full-text bodies and the design details. This way, we found another eight duplicates (see duplicated in Figure 12). For example,

¹¹Validation by the second extractor for false negatives was based on a 20% sample of all 704 papers processed by the first extractor and considered as either wrong report type or as containing erroneous DSML documentation artifacts. Due to this sampling, in a probabilistic projection, we risked having missed out between 0 and up to 23 publications. With one false negative found in this sample, in the worst case, we risk having missed up to 33 false negatives in the total set of 704 publications (binominal confidence interval (0, 34]; 99% confidence level). Note, however, that any false negative at this point might still have been ruled out as a true negative due to four subsequently evaluated criteria (e.g. a wrong UML version) and during quality assessment (i.e. duplicate removal). With only 73 out of 106 publications (i.e. $\approx 68.9\%$) having entered the final publication corpus (see Figure 10), we arrive at the number of up to 23 potentially missed publications in the paper corpus as false negatives under the above empirical probability: $33 * 0.689 \approx 22.7$




Figure 12: Overview of important selection decisions on the 2,678 papers entering the selection procedure, with 106 entering quality assessment, and 73 finally becoming included.

one hit turned out having two replicated entries between two or more result sets which did not share the same DOI. This way, they had escaped the semi-automated duplicate detection before.

The extent, to which we could extract design rationale from the included papers, was directly dependent on the completeness, the expressiveness, and the correctness of the design documentation by the publication authors. This challenge is a concrete instance of inappropriate rationale representation in design-rationale capturing in more general [24]. We checked the documented DSML designs for a number of well-discussed, primarily syntactic issues when it comes to designing UML-based DSMLs (see, e.g., [55, 64, 80]). To the extent concrete-syntax issues also signal semantics violations, we also considered certain classes of semantics defects. This way, we identified 25 papers whose DSML designs were reported in an inappropriate, non-extractable manner and were therefore excluded from data extraction. The quality issues are reported in full detail in the SLR protocol [41].

Extraction of Publication Data. After having completed the quality assessment, 73 papers representing 2.7% of the original search hits remained (see Figure 12). For this final publication set, we extracted or completed the publication-specific data. Besides, we coded the DSML design decision data according to our decision-record catalog (see Section 5.3).

Overall, we recorded 15 metadata items for each included paper. These items included bibliographical entries such as the publication year, paper-specific keywords, and the venue. In addition, the selection decision was noted for each hit. The included papers were further described by eight decision-mining entries: Beyond the DSML project name (if available), the DSML projects were classified according to their application domain(s) and the relevant UML diagram types. Next, the respective decision options identified for a given DSML design were recorded with respect to each of the six decision points from our decision-record catalog. Each DSML design became represented as one decision-option set (see Section 2.1.2). The final paper corpus including the 73 papers identified via the engine-based search is characterized according to the extracted publication data in Section 4.2. In

Section 5.3, we report on the extracted design-decision data in more detail.

Inter-Rater Reliability. 2,400 of the 2,678 search hits, which remained after applying the cutoff points, were reviewed by one author. 278 were rated and assessed for inclusion or for exclusion by two authors. This co-rated subset comprised all publications included by one author, all publications considered erroneous, and a randomized 20% sample of the papers excluded by one author.

For this co-rated subset, the respective rating authors achieved a percent agreement of 88.5%. That is, for more than 245 of the 278 publications two authors arrived at the same selection decision (i.e. included or excluded) independently from each other and without negotiations. Chance-corrected (worst-case) inter-rater agreement $\widehat{\kappa}_C$ amounted to $\approx 0.875 \pm 0.0004$.¹² This represents a *very good* (worst-case) agreement level according to standard benchmarks [73, 78]. Therefore, we considered the underlying selection procedure (criteria) reliable. Any conflicting selection decisions were first revisited in a joint session between the two rating authors. Then, if necessary, the third, non-involved author adjudicated on a search hit at issue.

4.1.3. Backward Snowballing

To incorporate prior work considered relevant by the authors of the 73 included papers, we performed a manual, citation-based search using the bibliographical references taken from the 73 included papers. We followed a backward-snowballing procedure as documented in [37, 38]. The procedural steps are summarized in Figure 13.

Manual Search. The snowballing procedure ended after having completed three search iterations. The initial iteration used the 73 papers included from the main search hits as a start set (see Section 4.1.2). The subsequent two iterations were triggered by newly included papers from a previous iteration. After having completed the third iteration, no new candidate papers were found. In each iteration, we first extracted the references lists from the start-set publications. In the initial iteration, we obtained 2,116 references from 73 papers. In the second iteration, we worked on 200 references taken from seven papers included during the first iteration. The third and last iteration used 21 references from one additional paper that was found during the second iteration. Across the three iterations, we reviewed a total of 2,337 references.

All 2,337 extracted references were checked by at least one reviewer, duplicate detection was performed manually. This was a consequence of the unavailability of structured and processable reference lists from all four data sources (ACM Digital Library, IEEE Xplore, Scopus, and SpringerLink). On

¹²Missing data in terms of hits only rated by one author have been considered for computing the marginal probabilities in deriving $\widehat{\kappa}_C$. \pm signals the leave-one-out Jackknife variance estimate $Var(\widehat{\kappa}_C)$ of Cohen's Kappa statistic.

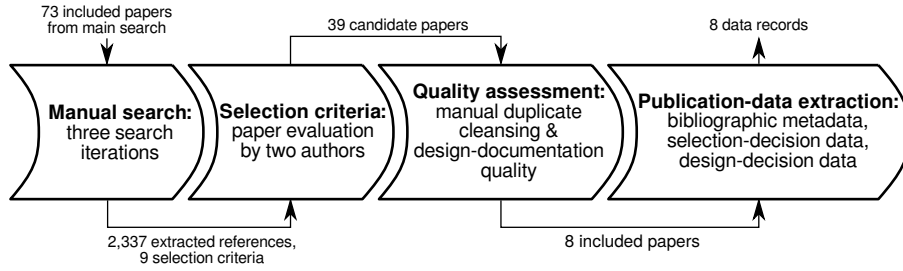


Figure 13: Sub-steps of conducting the citation-based search (backward snowballing).

top, the bibliographical reference data was widely heterogeneous (mixed bibliography styles, missing bibliographical entries such as DOIs).

Selection Criteria. We made the selection decision on the snowballing hits by applying the nine venue- and publication-specific criteria as for the main, engine-based search (see Section 4.1.2). This way, we selected another 39 publications as inclusion candidates. In addition, snowballing yielded one more QGS publication adding to the 28 QGS publications from the main search (see Section 4.1.2). We, thus, arrived at a final quasi-sensitivity of $\approx 78.4\%$ (29/37) for our review.

Approximately 7.3% of the extracted references (170/2,337) did not relate to reports on UML-based DSML designs as primary studies, but rather to other (non-UML) DSMLs, to secondary studies on DSML designs, or they were not related to DSMLs at all. 37 of these off-topic references were verified by two authors independently from each other.¹³ The majority of 91% (2,128/2,337) did not satisfy the eight remaining venue- and publication-specific criteria.

Quality Assessment. The 39 candidate publications were then assessed for duplicates and issues of design-documentation quality. In total, by comparing the snowballing hits to the cleaned automated search hits, we identified and removed 25 unwanted duplicates. For six of the remaining 14 publications, the documentation-quality assessment revealed quality issues similar to the ones during the main search. The issues included ambiguous UML metamodel extensions, faults in defining and applying the UML profiles, and syntax errors (see the SLR protocol [41]). Eight publications entered the paper corpus (see Figure 13).

¹³Validation by the second extractor for false negatives was based on a 20% sample of all 170 papers processed by the first extractor and considered as wrong report type. Due to this sampling, in a probabilistic projection, we risk having missed out between zero and up to seven publications. With one false negative found in this sample, in the worst case, we risk having missed up to 31 false negatives in the total set of 170 publications (binominal confidence interval (0, 32]; 99% confidence level). Any false negative found at this point might still have been ruled out as a true negative due to four subsequently evaluated criteria (e.g. a wrong UML version) and during quality assessment (i.e. duplicate removal). With only eight out of 39 publications (i.e. $\approx 20.5\%$) having entered the final publication corpus (see Figure 13), we arrive at the number of up to seven potentially missed publications in the paper corpus as false negatives under the above empirical probability: $31 * 0.205 \approx 6.4$.

Extraction of Publication Data. From the eight additional publications, we extracted bibliographical metadata and coded their design-decision data (see Sections 4.1.2 and 5.3, respectively). Missing (e.g. publication years) and inconsistent metadata (e.g. title formats) prevented us from extracting descriptive statistics over all backward snowballing items.

Inter-Rater Reliability. During selection and quality assessment, 2,267 of the 2,337 snowballing hits were reviewed by exactly one author. 69 hits were assessed for inclusion and for exclusion by two authors. The co-rating authors agreed on more than 80% of the rated references in each of the iterations. The corresponding chance-corrected (worst-case) agreement levels $\widehat{\kappa}_C$ for the each iteration were $\approx 0.967 \pm 0.0005$, $\approx 0.833 \pm 0.0345$, and 1, respectively.¹⁴ Therefore, we achieved a *very good* to almost *perfect* level according to standard Kappa benchmarks [73, 78] for each iteration. Therefore, we considered the underlying selection procedure (criteria) reliable. Any conflicting selection decisions were first revisited in a joint session between the two rating authors. Then, if necessary, the third, non-involved author adjudicated on a search hit at issue.

4.2. Paper Corpus

To summarize, the main, engine-based search and the snowballing searches resulted in retrieving and reviewing 5,015 publications. The main search accounted for 2,678 (53.4%), snowballing for another 2,337 publications (46.6%). From these 5,015 publications, we considered a total of 81 articles as relevant: 73 from main search plus 8 from snowballing. Recall that from 37 QGS publications, 29 had been returned by the main and snowballing searches. To complete the paper corpus, we re-considered the missing eight QGS publications for inclusion based on all 9 selection criteria. This way, we classified two QGS journal articles and one QGS conference article as relevant. We so arrived at a paper corpus of 84 publications out of a total of 5,023 reviewed publications (the complete list of all 84 publications can be found in the Appendix). The corpus was composed of 54 conference articles (64%) and 30 journal articles (36%).

¹⁴Missing data in terms of hits only rated by one author have been considered for computing the marginal probabilities in deriving $\widehat{\kappa}_C$. \pm signals the leave-one-out Jackknife variance estimate $Var(\widehat{\kappa}_C)$ of Cohen's Kappa statistic.




Figure 14: Finally selected articles per venue type (journal, conference, total) and per publication year (2005–2012). The release years of important UML revisions are plotted as vertical lines. Note that the paper corpus extends to 2013. This is due to running the search in Spring 2013 permitting results from 2013 to handle the ambiguity of “publication year” as either venue or paper-publishing year.

Figure 14 shows the timeline distribution of the 84 included articles, discriminating between archival and presence venues. The release years of important UML revisions (2.0, 2.1.2, 2.2, 2.3, and 2.4.1) in this time window are plotted as vertical lines. In general, the total number of publications on UML-based DSMLs, which are considered relevant for the scope of our study, appears to be in decline. Conference articles show a clear decrease and journal articles stagnate. 2006 and 2011 represent intermittent peaks. This general observation of a paper drop is in line with findings on UML profiles prior to 2010 by Pardillo [64].

In total, the 84 included articles were published in 29 different conference proceedings and in 13 different journals (see Table 5). The top four conference venues (MoDELS, ECMFA, SAC, and SAFECOMP; 14% of all conference venues) account for 26 publications (48% of the total 54 conference publications). The top four journal venues (SoSyM, IST, JOT, and CSI; 31% of all journal venues) account for 20 articles (67% of the total 30 journal articles). In total, more than half of all included publications (46 publications, 55%) were published in eight different venues. There are two conference outliers (MoDELS and ECMFA; 10 and 9 publications, respectively) and one journal outlier (SoSyM; 10 articles) contributing an over-proportionally high number of included articles.

Table 5: Included articles per publication venue. Journals are flagged with an asterisk (*). See Appendix A in Sobernig et al. [41] for the venues’ full names.

Frequency	Publication venue
10	MoDELS, SoSyM*
9	ECMFA
4	IST*, JOT*, SAC
3	SAFECOMP
2	COMPSAC, CSI*, EDOC, ICWE, IETSoftw*
1	CSMR, EMSOFT, ER, ESEC, FASE, HASE, ICSEA, ICSOC, ICWS, IJ-CIC*, LJSEKE*, Informatica*, ISeB*, ISSRE, JRPIT*, JSW*, OOPSLA, OTM, QSIC, SAM, SCC, SCCC, SEAA, SEFM, SEKE, SFM, SOCA*, SoMeT, SP&E*, SPLC

To map the domain coverage of the selected DSMLs, we classified every DSML project of the included publications according to the 2012 ACM Computing Classification System (CCS)¹⁵. For example, GWfM-Sec is a UML-based DSML for the modeling of security-critical, inter-organizational workflows and provide a mapping to the Web Services Choreography Description Language (WS-CDL). We assigned this article the following CCS categories: *software security engineering*, *web services*, *orchestration languages*.

In total, we applied 155 category assignments for all 84 publications (the assignments per publication can be found in the Appendix). This corresponds to a mean of 1.85 categories per publication. The 155 assignments were specific to 61 unique CCS categories. Table 6 reports the resulting frequency distribution of assigned categories. A notable number of DSMLs fall into the areas of service-oriented and embedded software systems, model verification and validation, as well as requirements analysis and security requirements. However, the frequency distribution clearly demonstrates that the paper corpus covers a very broad and diverse range of DSML application domains. This is an important achievement and prerequisite for answering our first research question (see Section 6).

5. Results: Captured Design Decisions

The paper corpus of 84 publications described the designs of 80 unique DSMLs. This is because, at closer inspection, eight of the 84 publications turned out to cover complementary details of four distinct DSML designs. In the following, we shift perspective from the 84 publications to these 80 unique DSMLs to report on the extracted design-decision data. A list of all DSMLs is provided in the Appendix.

¹⁵See <http://www.acm.org/about/class>; last accessed: Sep 9, 2015.

Table 6: Frequency of occurrence of ACM 2012 CCS categories for the corpus of 84 selected articles.

Frequency	ACM 2012 CCS categories
11	Service-oriented architectures
10	Web services
7	Embedded systems, Model verification and validation, Software development techniques
6	Security requirements
5	Model checking, Requirements analysis, Web applications
4	Data warehouses, Graphical user interfaces, Real-time systems, Software development process management, Software testing and debugging
3	Avionics, Business process modeling, Safety critical systems, Software architectures, Software evolution, Software security engineering, System on a chip, Web interfaces
2	Design patterns, Fault tree analysis, Measurement, Metrics, Orchestration languages, Reusability, Software product lines, Software safety, Transportation
1	Access control, Architecture description languages, Availability, Collaborative and social computing, Database design and models, Data mining, Distributed architectures, Electronic commerce, Engineering, Enterprise data management, Enterprise information systems, Error detection and error correction, Estimation, Hardware description languages and compilation, Hypertext languages, Operating systems security, Performance, Robustness, Scenario-based design, Semantic web description languages, Software design engineering, Software fault tolerance, Software performance, Software reliability, Systems analysis and design, Telecommunications, Trust frameworks, Ubiquitous and mobile computing, Use cases, Version control

5.1. Content Analysis

We performed a content analysis (CA) on all corpus papers using the 80 DSMLs as units of analysis. The content analysis involved human coding of the papers’ full-text according to a previously defined coding schema. This coding schema was systematically derived from the decision-record catalog by the authors. The CA purpose was a threefold test for this coding schema:

1. Test for saturation: Is each decision category reflected at least once in the corpus?
2. Test for exhaustiveness: Can each unit of coding (e.g., text fragment, specification artifact) be assigned to an already defined category?
3. Test for generalizability: Are the decision categories *general* in terms of being found for at least three units of analysis (DSMLs)?¹⁶

Hence, we applied a *directed* (deductive) content analysis using *hypothesis coding* [42]. Note that these three tests are specific to the content-analysis step of our multi-method study. Our two main research questions are exploratory (see Section 3). Running the three tests sets the context for systematically answering the research questions (esp. RQ2; see Section 6).

Coding schema. The coding schema comprised category definitions, category indicators, category examples, and decision rules. The structure and content of the (pre-study) decision-record catalog provided a framework for building the coding schema as a two-level category hierarchy. At the top level,

¹⁶The indicative threshold of three is borrowed from the software-pattern community, see Section 2.1.1 for a justification.

the decision records formed main categories (D1–Dm). Decision options became subcategories. The subcategory codes (O1.1–Om.n) were then used for material coding and coding analysis (see below). The decision records already provided all description items for the coding categories and instruction items for coders (see Table 7).

Table 7: Correspondences between content items of the decision-record catalog and the CA coding schema, which was built according to the guidelines in [43].

Decision-record catalog	Coding schema
Decision record	Main category
Decision option	Subcategory
Record name, option name	Category name
Record description, option description	Category description
Applications, sketch, option descriptions	Indicators, positive examples
Option descriptions, decision consequences	Decision rules

Indicators complement category descriptions and can signal the presence of a category of interest (decision option) in the material to the coder. Important indicators in our coding schema were *artifact types* characteristic for DSMLs and *technology projections*. The decision record on “language-model constraints” (D3 [33]), for instance, identifies kinds (and examples) of code listings. As for technology projections, this decision-record description refers to constraint-expression languages such as OCL and EVL as examples.

To illustrate the categories during coding, we included positive examples of what the category is meant to cover. Positive examples included concrete instances of artifact types (e.g. concrete code listings) and technology projections provided by the catalog (sketches). Sketches were extracted from DSMLs or secondary studies on extending the UML and on building UML-based DSMLs. The following is an exemplary sketch taken from the decision record on “language-model constraints” (D3; [33]):

“Sketch. Consider the following excerpt from P8: For a UML activity, each action can be guarded by a constraint whose conditions refer to a set of operands and checking operations. At runtime (level M0), the operations are called to evaluate whether an action should be entered, depending upon some contextual state. Constraint 1 shows a constraint-language expression (OCL) accompanied by a complementary textual annotation. [...]

Constraint 1: The operands specified in a ContextCondition are either ContextAttributes or ConstantValues:

```
context ContextCondition inv:
  self.expression.operand.oclAsType(OperandType) ->
    forAll(o |
      o.oclIsKindOf(ContextAttribute) or
      o.oclIsKindOf(ConstantValue))
```

[33, pp. 22–23]

Some categories exhibited an overlap between concepts and between indicators (examples). Consider, for example, the subtle differences between the concrete-syntax decision options (categories) DIAGRAMMATIC SYNTAX EXTENSION

(O4.2) versus MIXED SYNTAX (O4.3). To guide coding under ambiguity, the coding schema included decision helpers (“decision rules”) taken from the decision-record catalog. To draw a definitional line between the two, the decision-record description contains guidelines such as “[...] in contrast to O4.2, this option [O4.3] would define a new and domain-specific diagram type” [33, p. 24].

Another set of decision rules was derived from the description of decision associations (see Section 2.1.2). Given a unit of coding, coders are so pointed to related categories in other main categories for further consideration. For instance, once having assigned a category PROFILE RE-/DEFINITION (O2.2) to a given unit of coding, the same unit of coding is likely to contain details on specific concrete-syntax choices. Therefore, a decision rule points the coder to potentially relevant concrete-syntax categories (e.g., DIAGRAM SYMBOL REUSE; O4.6). The decision association “native stereotype definition” exemplifies such a decision rule:

“A UML profile definition (O2.2) for the language-model formalization was observed in combination with a concrete syntax specification via annotating model elements (O4.1) and reusing diagram symbols (O4.6; see, e.g., P22 or P54). [...] A stereotype inherits all semantics (abstract syntax) and the notation (concrete syntax) from its extended UML base class.” [33, p. 38]




Figure 15: Thumbnail overview of the corpus paper on the DSML BusinessActivities after segmentation. Five (out of six possible) themes are identified as colored segments. The ninth and tenth pages (465, 466) are highlighted for the coding details in Fig. 16.

Segmentation and coding were all performed using PDF copies of the selected materials, PDF reader software, and portable PDF annotations.

Segmentation. In this step, we marked relevant parts (themes) and units of coding. From preparatory work, we were aware that statements and content fragments documenting design decision do not follow the internal structure of the scientific publication (e.g., dedicated sections per design aspect). Rather, we had found them spread across the document, including appendix material and externally

referenced companion material. In addition, we had decided to consider all auxiliary design-documentation artifacts contained (e.g. diagrams) or referenced by the publication at hand, if fully accessible. Important artifacts included package diagrams as well as implementation artifacts such as metamodel, profile, and concrete-syntax specifications. Therefore, we first marked statements and content items (e.g. tables, figures, listings) having a common point of reference (themes).

We used color highlighting to identify six themes in a document (“domain analysis”, “UML2”, etc.; see Fig. 15). The available themes are defined by the decision-context descriptions of the decision records. For example, the theme “domain analysis” results from the decision record on “language-model definition” (D1) which describes the broader decision context as part of a domain-analysis step:

“Decision context. A prerequisite for DSML design is a systematic analysis and the structuring of the language domain. By applying a domain analysis method, such as domain-driven design [...], information about the selected domain is collected and evaluated (e.g. based on literature reviews, scenario analyzes, and collected expert knowledge). [...]” [33, p. 12]

Those parts of a document containing details on domain-analysis procedures, techniques, and analysis findings were marked using the theme’s color code (see the color legend in Fig. 15). In Fig. 15, five theme segments are shown color-highlighted (out of six possible ones) for the paper on the DSML BusinessActivities. As an example, the “domain analysis” segment in the themed document groups those document items which deal with analyzing the application domain (RBAC for business processes) and which define key domain abstractions and their relationships (e.g. separation and binding of duty). This way, all corpus material was split up into up to six different thematic segments (per DSML).

Within each thematic block, we then identified smaller parts (units of coding) to be assigned codes in a separate step (see paragraph on “Main coding”). Units of coding were complete phrases, phrase blocks, and content items such as tables, figures, listings, and formula blocks. Relevant phrases were identified using underline marks, other content items were marked using margin bars. See Fig. 16 for various examples of marks and units of coding. For instance, on p. 465 of the excerpt, three listings depicting constraint expressions (OCL) were marked using vertical margin bars. Each of these listings formed a separate unit to be later assigned a code. Note that the identified segments were not broken up into units of coding exhaustively. Phrases and content items deemed irrelevant for the three above tests were not considered as units of coding. The two paragraphs discussing rationale for making or limitations of given formal propositions on p. 464 in Fig. 16 are examples of such excluded content.

Segmentation as a separate working step reduced the risk of accidentally skipping relevant content during actual coding, yet it helped break down the large paper corpus into a manageable material collection for main coding. Note that segmentation step was carried as an integrated part of SLR data extraction (see Section 4.1). During data extraction,

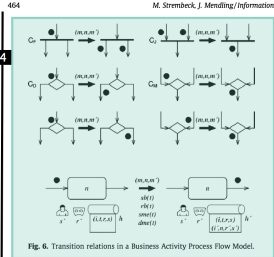


Fig. 4. Transition relations in a Business Activity Process Flow Model.

O1.4

point in time, then the process instance will deadlock. This is an analysis feature that allows us to trace back the satisfiability problem to standard techniques for deadlock detection.

O1.2 Definition 7 (Reachability Graph). Let $PFM = (N, A)$ be a Business Activity Process Flow Model with $N = \{s, c, g, u, q, c, u, c, u, u, s, t, r, e, m, e, n, t\}$. BRM is a Business Activity RBAC Model including T, P, S as a process instance, and h is the execution history of p . Then the Reachability Graph $RG = (M, TR)$ with $TR \subseteq M \times N \times M$ must comply with the following constraints:

1. The initial state is in M , i.e. $m_0 \in M$ with $T = \emptyset$.
2. If $m = (d, h) \in M$ and $n \in C_{u, s, t, r, e, m, e, n, t}$ and for all $a \in n_{u, s, t, r, e, m, e, n, t}(a) > 0$ and $m = (d, h)$ exists such that for all $a \in A(m_{u, s, t, r, e, m, e, n, t}(a) = d(a))$ for all $a \in n_{u, s, t, r, e, m, e, n, t}(a) > 0$ and for all $a \in n_{u, s, t, r, e, m, e, n, t}(a) = 1$, then $m \in M$ and $(m, n, m) \in TR$.
3. If $m = (d, h) \in M$ and $n \in C_{g, u, c, u, s, t, o, r}$ and for all $a \in n_{u, s, t, r, e, m, e, n, t}(a) > 0$, $m = (d, h)$ exists such that for all $a \in A(m_{u, s, t, r, e, m, e, n, t}(a) = d(a))$ there exists an $a \in n_{u, s, t, r, e, m, e, n, t}(a) = 1$ and there exists an $a \in n_{u, s, t, r, e, m, e, n, t}(a) = 1$, then $m \in M$ and $(m, n, m) \in TR$.
4. If $m = (d, h) \in M$ and $n \in P$ and for all $a \in n_{u, s, t, r, e, m, e, n, t}(a) > 0$ and it exists $t \in TR(p, t) \in R \subseteq S$ with $n \in TR(t, c)$ and $r \in TR(c, s)$ that for all $(t, r, p, c, s) \in h$ holds that:
 - If $n = s(t)$ then for all $t \in TR(p, s) = es(t) = es(t) = s$, and
 - If $n = r(t)$ then for all $t \in TR(p, r) = es(t) = es(t) = r$, and
 - If $n = s(m)$ then for all $t \in TR(p, s) = es(t) = es(t) = s$, and
 - If $n = r(m)$ then for all $t \in TR(p, r) = es(t) = es(t) = r$.

$m = (d, h)$ exists such that for all $a \in A(m_{u, s, t, r, e, m, e, n, t}(a) = d(a))$, there exists an $a \in n_{u, s, t, r, e, m, e, n, t}(a) = 1$, and there exists an $a \in n_{u, s, t, r, e, m, e, n, t}(a) = 1$ and $N = \{s, c, g, u, c, u, s, t, o, r\}$, $T = \{t\}$ such that $t \in TR(p, s)$, $t \in TR(p, r)$, and $m = M$ and $(m, n, m) \in TR$.

O1.2 Proposition 1 (Dynamic Correctness BRM). All BRM that comply to states M of the reachability graph are dynamically correct.

O1.2 Proof, by induction:

Base Case: In the initial state holds $T = \emptyset$ such that all four properties of definition 3 are fulfilled.

Induction: For transitions 7.2 and 7.3, BRM does not change and therefore remains correct if BRM was also correct. For transition 7.4, BRM remains correct by definition since (t, n, r, s) does not violate any constraints.

Note that Proposition 1 has implications for the so-called satisfiability problem. This is the question whether it is guaranteed that there is always a role-subject pair that allows the process to proceed [17]. However, in this paper we are concerned with modeling support for process-related RBAC models and do not address the satisfiability problem. This is because, reachability graph analysis is already an exponential problem in the general case even without considering mutual exclusion and binding constraints, but different optimization strategies can be employed [8]. What we gain though is the capability of simulating process execution including dynamically correct role and subject allocation (see also [87]). In this way, the formalization serves as the generic foundation for a thorough workflow and RBAC specification for process modeling languages.

4. UML extension for Business Activities

In this section, we extend UML activity models so that they provide modeling support for Business Activities as defined in Section 3. Therefore, we define a new package *BusinessActivities* as an extension to the UML2 metamodel. Fig. 7 shows the metamodel of the *BusinessActivities* package, including all new modeling constructs defined in this package. In particular, we introduce *BusinessActivity*, *BusinessAction*, *Subject*, *Role*, *RoleToSubjectAssignment*, and *RoleToRoleAssignment* as new modeling elements.

A *BusinessActivity* is defined as a subclass of *Activity* (from the *BasicActivities*, *CompleteActivities*, *FundamentalActivities*, and *StructuredActivities* packages, see [61]). A *BusinessAction* is defined as a subclass of *Action* (from the *CompleteActivities*, *FundamentalActivities*, and *StructuredActivities* packages, see [61]).

M. Strembeck, J. Mendling / Information and Software Technology 53 (2011) 456–483

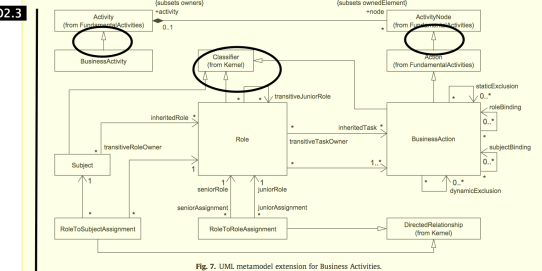


Fig. 7. UML metamodel extension for Business Activities.

4.1. Invariants for Business Activities

In accordance with Section 3, we now specify OCL invariants [58] on *BusinessActivities* and *BusinessActions* that ensure the correct semantics of models defined with our UML extension (see also [61]). In particular, these OCL invariants define the semantics of the corresponding graphical *BusinessActivity* models (see Section 4.2) and therefore make sure that the graphical models are correct and complete with respect to the formal specification from Section 3. However, for the sake of readability, this section only shows four OCL invariants as examples. The complete list of OCL invariants for the *BusinessActivity* extension is found in Appendix A.

O3.4 OCL Constraint 1. Each role may have direct and indirect/transitive junior roles. In other words, if a role r has junior roles and these junior roles have junior roles themselves, r inherits the indirect/in-transitive junior roles.

```

context Role inv:
self.roleAssignments->forall {ra}
  ra.juniorRoleAssignments->forall {ra2}
    ra2.transitiveJuniorRole->exists {ra3}
      ra3.name = ra2.juniorRole.name
  and
  ra.transitiveJuniorRole->exists {ra4}
    ra4.name = ra.juniorRole.name
  }

```

O3.4 OCL Constraint 9. To assign subjects to an instance of a *BusinessAction* and to determine the subject that executes a particular *BusinessAction*, instances included in a particular *BusinessActivity* instance, we require that each *BusinessAction* defines an attribute called *executingSubject* (and thereby each instance owns a resource *role*). Moreover, the *executingSubject* attribute must refer to a subject that is (via one of its roles) actually allowed to execute this *BusinessAction*.

```

context BusinessAction inv:
self.instanceSpecification->forall {i}
  i.roleAssignments->exists {ra}
    ra.roleAssignments->exists {ra2}
      ra2.subject.name = ra.subject.name
  and
  ra.subject.name = ra.subject.name
  or
  ra.subject.name = ra.subject.name
  and
  ra.subject.name = ra.subject.name
  }

```

O3.4 OCL Constraint 17. In a role hierarchy, mutual exclusion constraints are subject to inheritance (see Section 3). Therefore, two SME *BusinessActions* must never be assigned to the same role, neither directly nor transitively.

```

context Role inv:
self.businessActions->forall {ba1, ba2}
  ba1.roleAssignments->exists {ra1}
  ba2.roleAssignments->exists {ra2}
  and
  ra1.name = ra2.name
  and
  ba1.roleAssignments->exists {ra3}
  ba2.roleAssignments->exists {ra4}
  and
  ra3.name = ra4.name
  }

```

O3.4 OCL Constraint 19. To enforce SME constraints on *BusinessActions*, we specify that the instances of two SME *BusinessActions* must never have the same *executing subject* (see also Constraint 9 which specifies the requirement that each *BusinessAction* defines an attribute called *executingSubject*). For details on the InstanceSpecification element see [61].

Figure 16: Excerpt (two pages) from the segmented and coded corpus article on DSML BusinessActivities.

the corpus papers were divided up between the authors for extraction and segmentation. This helped avoid excessive and time-consuming iterations over the entire corpus. To establish that all authors perform segmentation in a consistent manner, segmentation was performed collaboratively in a joint coding session for papers collected during preparatory studies [35, 36].

Main coding. In the main coding phase, each unit of coding became assigned to one of the categories (decision options) of the coding schema. Category assignment was performed in line with the overall SLR extraction procedure (see Section 5). For the majority of the corpus papers, coding was conducted independently by two authors, each author being “blinded” for the assignments of his alter. A smaller share was coded/recoded by a single author. This was due to one

author having initially excluded a DSML or the underlying publication(s), for example. As a result, the overall paper corpus was effectively split into two parts: For 62 of the 80 DSMLs, coding was performed twice and independently by two authors. This part was also checked for coding consistency (see below). For the remaining 18 papers, there were assignments by a single author.

The choice of a particular category assignment was first recorded by each author for each unit of coding as a text mark. The mark’s text identified the category assigned to a given unit of coding. For example, “Fig. 7” on page 465 in Fig. 16 was handled as a single unit of coding and text-marked by “O2.3”. This text mark indicates that the package diagram documents a UML metamodel extension in a formal and diagrammatic way (METAMODEL EXTENSION). The

UML metamodel is extended by adding six sub-meta-classes (BusinessAction, Subject etc.). Note that this coding step resulted in repeatedly assigning categories to several different coding units per unit of analysis (DSML): Consider the three OCL listings in Fig. 16, each of them being marked “O3.1” (CONSTRAINT-LANGUAGE EXPRESSION). Each unit of coding could also be assigned to two or more categories. Each text mark could be accompanied with text comments, e.g., for providing some rationale for the assignment.

In a second step, the category assignments were prepared for further analysis using a summarizing coding form. The coding form is documented in the SLR protocol [41]. For every DSML, each coding author recorded whether a given category has been assigned at least once, i.e. there is at least one unit of coding having assigned a certain category (decision option). Recording and further analyzing the state of presence (absence) of one category per DSML was sufficient to answer the three tests and the overall research questions. Each row in these coding forms translates directly into the decision-option sets used for further analysis (see Section 2.1.2).

Coding Consistency. To compare the coding performed by the pairs of authors, we used summary measures of inter-rater agreement. Inter-rater agreement signals how consistent the different coders were in assigning categories (codes) of the coding schema to units of analysis (DSMLs) in the material. That is, to which extent they are in agreement on codes to assign and/or on codes to discard independently from each other, based on the coding schema, and without negotiating a joint decision. Any conflicting coding was first revisited in a joint session between the two coding authors. Then, if necessary, the third, non-involved author adjudicated on a DSML at issue.

Table 8 summarizes the percent agreement (upper bound) and the extended chance-corrected Kupper-Hafner indices (lower bound, worst case) for the six main coding categories. The two independently coding authors arrived at the same assignment of categories to units of analysis (DSMLs) for 84% (D3) through 98% (D5) of the 62 co-rated DSML designs. Applying chance correction (extended KH index \hat{C}^*) under worst-case assumptions (random assignment), we established that the two rating authors achieved actual coding agreement beyond chance in approx. 77% (D1) through 98% (D5) of the total expected cases beyond chance. In technical terms, coding consistency amounted to an extended KH index \hat{C}^* equal to and greater than 0.77 ± 0.0012^{17} for all six main coding categories (see Table 8). This includes agreement in terms of categories not assigned to DSMLs by both coders. Therefore, for the 62 co-rated DSML designs, the co-rating authors were in medium to good chance-corrected (worst-case) agreement according to standard benchmarks [73, 78]. Hence, we considered the coding

guided by the coding schema consistent. Any coding conflicts were resolved as outlined above.

Table 8: Coding consistency per main coding category (D1–D6); Kupper-Hafner (KH) Index; $\hat{\pi}^*$: Percent agreement (incl. missing ratings); \hat{C}^* : Chance-corrected KH indices (incl. missing ratings); $Var(\dots)$: leave-one-out Jackknife variance estimates.

	D1	D2	D3	D4	D5	D6
$\hat{\pi}^*$	0.85	0.96	0.84	0.94	0.98	0.87
$Var(\hat{\pi}^*)$	0.0004	0.0002	0.0005	0.0002	0.0001	0.0004
\hat{C}^*	0.77	0.95	0.80	0.92	0.98	0.84
$Var(\hat{C}^*)$	0.0012	0.0003	0.0008	0.0003	0.0002	0.0007

Test Results. Based on the coded decision data, the three tests guiding content analysis were answered as follows:

1. *Saturation: Is each decision category reflected at least once in the corpus?*
No: Three categories out of the 31 available categories were not assigned to any unit of coding, and therefore, not reflected in any DSML.
 Therefore: We annotated the corresponding decision options in the decision-record catalog to reflect this missing evidence from our content analysis. See Sections 5.5 and 6 for a discussion.
2. *Exhaustiveness: Can each unit of coding (text fragment, specification artifact) be assigned to an already defined category?*
No: The paper corpus revealed units of coding which could not be assigned to any of 26 categories of the initial coding schema.
 Therefore: The coding schema was extended accordingly, by introducing five new categories on options for behavior specification and platform integration. See Sections 5.5 and 6 for a discussion.
3. *Generalizability: Are the decision categories general in terms of being found for at least three units of analysis (DSMLs)?*
No: Seven of the finally available 31 available categories were not assigned to at least three DSMLs.
 Therefore: We proceeded with a frequency-pattern analysis for the 24 generalizable categories (see Section 5). The presentation of the decision options corresponding to the seven non-generalizable categories was revised, for example, by marking them as *candidate* options. See Sections 5.5 and 6 for a discussion.

The detailed numbers (frequencies) underlying these tests are reported in Section 5.3.

5.2. Descriptive Data

Before reviewing the design-decision data, we characterize the 80 DSMLs regarding the adopted UML diagram types and their specification size.

¹⁷± signals the leave-one-out Jackknife variance estimate $Var(\hat{C}^*)$ of the extended chance-corrected Kupper-Hafner Index \hat{C}^* .

Diagram Types. The majority of the 59 DSMLs tailor a structural UML diagram type; 38 DSMLs adopted and tailored one or more structural and no behavioral diagram types (see Table 9; the diagram types tailored by each DSML are listed in the Appendix). Class diagrams are adopted by 50 DSMLs, followed by component and package diagrams (14 DSMLs each). This observation completes the empirical picture of a preponderance of class diagrams in broader UML usage (see, e.g., [12, 81]). Profile diagrams have not been found adopted by any of the DSMLs.

With respect to behavioral diagrams, 39 DSMLs tailor at least one behavioral diagram type; 18 build on behavioral diagrams only, without building on any UML structural diagram type. Activity (20), state machine (11), and use case diagrams (8) are most frequently used, while communication and timing diagrams are not used by any DSML.

21 DSMLs provide both tailored structural and behavioral diagrams. Similar distributions between structural and behavioral diagram types were found in an earlier secondary study on UML profile usage [64].

Table 9: Overview of the structural and behavioral UML diagram-types according to Appendix A in Object Management Group [10], as adopted and/or tailored by 77 out of 80 DSMLs identified using the SLR. Three DSMLs, which target all or are unspecific about the UML diagram types, are omitted.

DSMLs (cnt.)	Diagram type (ranked by decreasing cnt.)
59	Structure diagrams
50	Class
14	Component
14	Package
7	Composite structure
5	Object
3	Deployment
0	Profile
39	Behavior diagrams
20	Activity
11	State machine
8	Use case
4	Sequence
1	Interaction overview
0	Communication
0	Timing

Specification Size. We quantified the core language-model sizes of the 80 DSMLs. Our findings below indicate that, on the one hand, our pool of DSMLs compares with prior reports on UML extension sizes. On the other hand, our DSML pool covers a greater size variety than prior work [16, 64]. Depending on the different, underlying UML implementation techniques (O2.1–O2.4), the specification size was established differently.

For 61 DSMLs defining their language models using UML profiles (O2.2), we counted the stereotype definitions and the corresponding, distinct base UML metaclasses. In this group, we find a median of 13 ± 8.9^{18} stereotype definitions

per DSML. A typical profile extends a median of 5 ± 3 distinct base metaclasses per DSML. As outliers, three DSMLs defined 40 (AspectSM), 48 (WebML), and 116 stereotypes (IEC61508), respectively. Two DSMLs extend 14 base metaclasses (UML4SOA, SafeUML). The specification sizes of these 61 DSMLs slightly differ from those reported by related work on UML profiles, but fall into a closely comparable size range. Pardillo [64] studied 39 UML profiles, with a median of 9 ± 5.9 stereotypes per profile and a median of 4 ± 1.5 extended base metaclasses per profile. Staron and Wohlin [16] cover three UML profiles containing between seven and 13 stereotype definitions.

17 DSMLs used a UML metamodel extension and/or modification (O2.3, O2.4). For these, we collected the number of newly introduced and redefined UML metaclasses. A typical DSML adds and redefines a median of 12 ± 11.9 UML metaclasses (O2.3, O2.4). One outlier includes 51 UML metaclasses (DMM/UCMM). Existing empirical work reports on UML metamodels of between 20 and 30 metaclasses [16].

For three DSMLs defining their language model using a UML class model at level M1 (O2.1), we attempted to count the number of UML classes. Given their incomplete design documentation, we failed in this attempt for two DSMLs (EM, UML-GUI). For the one remaining DSML (SECTET), we counted 20 UML classes.

5.3. Extracted Decision Data

Content analysis yielded one decision-option set per DSML (see Section 2.1.2). For brevity, we refer to decisions option by their option code (e.g. O1.4, O3.1) in the following, rather than by name (e.g. FORMAL DIAGRAMMATIC MODEL, CONSTRAINT-LANGUAGE EXPRESSION). For a complete reference on the relevant decision options, see Tables 1 and 2 in Section 2.1.

Decision Points. Each of the 80 DSMLs covered the points of language-model definition (D1) and language-model formalization (D2), respectively. This was also a minimum requirement for a DSML design to become included in our study (see Section 4.1.2). For the remaining four decision points (D3–D6), we also recorded whether or not any decision could be recovered.

Table 10 shows the frequencies per decision-option code observed for the 80 coded DSMLs. The majority of DSMLs adopted one or several decision options for D3 (language-model constraints) and D4 (concrete syntax): 48 DSMLs explicitly documented language-model constraints (D3; e.g. OCL expressions) beyond the constraints expressed directly via their formalized language models. 73 DSMLs included at least one decision on the concrete syntax (D4) of the DSML. On the contrary, only three DSMLs reported on decisions related to additional behavioral specifications (D5), and only 26 DSML designs comprised decisions related to platform integration (D6). The latter is reflected by the comparatively

¹⁸We report the variance in terms of the *median absolute deviation from*

the median (MADM) using the \pm notation along with the median value.

Table 10: The number of occurrences (abs. frequency n_{opt} , rel. frequency f_{opt}) of each of the 27 (31) decision-option codes available from the decision-record catalog [33] in the 80 reviewed DSML designs. The 27 actual decision options are described in Tables 1 and 2. Four decision-option codes represent the absence of a decision (i.e. no-option codes: O3.5, O4.7, O5.5, and O6.6).

Decision option (code)	n_{opt} (f_{opt} , %)	Decision option (code)	n_{opt} (f_{opt} , %)
Language-model definition (D1)		Concrete-syntax definition (D4)	
INFORMAL TEXTUAL DESCRIPTION (O1.1)	80 (100)	DIAGRAM SYMBOL REUSE (O4.6)	69 (86.3)
FORMAL DIAGRAMMATIC MODEL (O1.4)	23 (29.8)	MODEL ANNOTATION (O4.1)	62 (77.5)
FORMAL TEXTUAL DESCRIPTION (O1.2)	5 (7.3)	DIAGRAMMATIC SYNTAX EXTENSION (O4.2)	14 (17.5)
INFORMAL DIAGRAMMATIC MODEL (O1.3)	3 (3.8)	<i>No decision</i> (O4.7)	7 (8.8)
Language-model formalization (D2)		MIXED SYNTAX (O4.3)	3 (3.8)
PROFILE RE-/DEFINITION (O2.2)	62 (77.5)	FRONTEND-SYNTAX EXTENSION (O4.4)	1 (1.3)
METAMODEL EXTENSION (O2.3)	17 (21.3)	ALTERNATIVE SYNTAX (O4.5)	1 (1.3)
M1 STRUCTURAL MODEL (O2.1)	4 (5)	Behavior specification (D5)	
METAMODEL MODIFICATION (O2.4)	2 (2.5)	<i>No decision</i> (O5.5)	77 (96.3)
Language-model constraints (D3)		INFORMAL TEXTUAL SPECIFICATION (O5.3)	2 (2.5)
INFORMAL TEXTUAL ANNOTATION (O3.4)	35 (43.8)	M1 BEHAVIORAL MODEL (O5.1)	1 (1.3)
<i>No decision</i> (O3.5)	32 (40)	FORMAL TEXTUAL SPECIFICATION (O5.2)	1 (1.3)
CONSTRAINT-LANGUAGE EXPRESSION (O3.1)	31 (38.8)	CONSTRAINING MODEL EXECUTION (O5.4)	0 (0)
CODE ANNOTATION (O3.2)	0 (0)	Platform integration (D6)	
CONSTRAINING MODEL TRANSFORMATION (O3.3)	0 (0)	<i>No decision</i> (O6.6)	54 (67.5)
		GENERATOR TEMPLATE (O6.2)	16 (20)
		M2M TRANSFORMATION (O6.5)	9 (11.3)
		API-BASED GENERATOR (O6.3)	7 (8.8)
		INTERMEDIATE MODEL REPRESENTATION (O6.1)	4 (5)
		(DIRECT) MODEL EXECUTION (O6.4)	1 (1.3)

large numbers of no-option codes displayed in Table 10 (O5.5 and O6.6, respectively).

Language-Model Definition (D1). For identifying and describing relevant domain abstractions and their relationships, the decision-record catalog documents four decision options (O1.1–1.4). Each of the four available option codes was identified at least once. 52 DSMLs applied a single option only (see Table 10). This decision was to use a textual, natural-language representation (O1.1). The 28 remaining DSMLs are characterized by a combination of two or three D1 options. Not a single design used all four options. 23 of 28 DSMLs provide a formal diagrammatic definition (O1.4) of their language model, prior to actually formalizing (implementing) the language model on top of UML. In all 23 cases, this decision option is accompanied by a textual, natural-language definition of the language model (O1.1).

Language-Model Formalization (D2). To realize the DSML by reusing the UML, concrete UML techniques for implementing the language DSML as a UML extension, as a UML specialization, and/or by piggybacking on UML must be considered. The decision-record catalog describes four combinable options (O2.1–O2.4). Each available option code was marked at least for one DSML design (see Table 10). 75 of the 80 DSMLs applied exactly one option, the remaining five projects adopted two options. 62 of the DSML projects ($\approx 78\%$) involved UML profiles (O2.2) to realize the language model, and $\approx 21\%$ of the DSMLs (17/80) extend the UML metamodel without modifying it (O2.3). Only three DSMLs applied both, UML metamodel extensions (O2.3) and UML profiles (O2.2), 14 DSMLs used UML metamodel extensions

only. Just two DSMLs explicitly modify the UML metamodel (O2.4) when extending it (O2.3).

Language-Model Constraints (D3). Once a language model has been formalized (D2), additional structural conditions can be placed on DSML models at the level of the language model (e.g. consistency conditions on model elements). The decision-record catalog lists four options of defining language-model constraints (O3.1–O3.4). Two out of the four available options (see also Table 10) were applied in the 80 DSMLs: constraint-language expressions (O3.1 in 31 DSMLs) and textual annotations (O3.4 in 35 DSMLs). We did not find evidence for code annotations (O3.2) and translational constraining (O3.3). In 30 DSMLs, only one of these options was adopted: O3.1 13 times and O3.4 17 times. In 18 DSMLs, O3.1 and O3.4 are adopted both.

Concrete-Syntax Definition (D4). Defining the DSML’s concrete syntax requires deciding on the style of the primary modeling interface presented to the domain modeler. In addition, it must be decided whether or not, and, if yes, how to integrate the DSML concrete syntax with the diagrammatic one of UML. Our decision-record catalog documents six options (O4.1–O4.6). All six options have been found applied at least once (see also Table 10). 73 of the 80 DSMLs contained at least one decision on their concrete syntax. Many DSML projects (49) took a combination of two decision options on their concrete-syntax style; 17 DSMLs applied one and 14 DSMLs three options. The two most frequently found decision options are model annotations (O4.1: 62 DSMLs) and the unmodified reuse of existing UML diagram symbols (O4.6: 69). 48 of the 49 DSML designs, which adopt exactly two options, reflect these two options (O4.1 and O4.6).

A UML diagram-syntax extension is adopted in 14 DSMLs (O4.2; e.g. by introducing new symbols or modifying existing ones).

Behavior Specification (D5). Behavior specifications stipulate how language elements of the DSML interact to produce the intended system behavior in a platform-independent manner. Only three of the 80 DSML design involve dedicated behavioral specification artifacts. These three DSMLs document these refinements of behavioral-semantics by adopting a UML M1 model representation (O5.1), a formal textual specification (O5.2), or in an informal textual way (O5.3). Only one of these three DSMLs applies two options in a complimentary manner (O5.1 and O5.3). The remaining two picked a single option only (O5.2, O5.3). Constraining model execution (O5.4) was not found applied in any DSML.

Platform Integration (D6). Platform integration comprises support for mapping DSML models into specification artifacts, which are processable and possibly executable by a targeted software platform, in a fully or semi-automated manner. Our decision-record catalog describes five support techniques for platform integration (O6.1–O6.5). 26 DSMLs applied at least one of these five techniques and each platform-integration technique was applied at least once (see also Table 10). From the 26 DSMLs, 19 DSMLs applied a single platform-integration option only rather than a combination of two or more options. The most frequently adopted option is O6.2 (generator templates) in 20% of the DSMLs (16/80). Generator templates are also the option, which is most frequently used in isolation (with 12 of the 19 single-option DSMLs). Generation templates are followed by model-to-model (M2M) transformations (O6.5: 9 DSMLs) and API-based generators for platform-specific models (O6.3: 7 DSMLs). M2M transformations (O6.5) are found mostly in combination with at least one other D6 option (in six out of nine DSMLs). The remaining 68% of the DSMLs do not consider or at least did not document any platform-integration techniques (O6.6: 54 DSMLs).

Decision-Option Sets. A decision-option set adds up all decisions recorded for a given DSML as a set of option codes corresponding to the decision options compiled by the catalog (see Section 2.1.2). Therefore, we obtained a base of 80 decision-option sets from the 80 DSMLs (see also the Appendix). By running a frequent-item-set analysis on these 80 observed decision-option sets, one for each DSML, 53 of the 80 observed option sets were unique (non-duplicated). From these 53, 14 unique option sets represent two or more DSMLs. In our study, none of the option sets were shared by more than five DSMLs. The maximum number of decision options included in an option set is ten for the unique option sets (i.e. for option sets that were applied for exactly one DSML) and seven options for the 14 shared ones (i.e. for option sets describing two or more DSMLs).

Remember that we consider an option subset to be frequent when it is used by three or more DSMLs. This is also

referred to as a *minimum support* of three for an option subset. This follows from a threshold on backing up a software-pattern description with at least three known uses of a given pattern in existing software systems; a threshold which is commonly applied in the software-pattern community (see, e.g., [31, 32]). We found 188 of such frequent options sets; that is, option sets which are contained partly or fully in more than three observed option sets.

Smallest Common Option Subsets. In our pool of 80 DSMLs, we found two smallest common option subsets *specific to one decision point*. For language-model constraints (D3), the proper option subset {3.1, 3.4} reflects that 18 DSMLs (i.e. the option subset is said of having a support of 18) define language-model constraints using both, a constraint-expression language as well as auxiliary or complimentary textual constraint definitions in natural language. As for platform integration (D6), a second proper option subset {6.2, 6.5} (support: 3) indicates that the respective 3 DSMLs use a two-level model transformation chain (PIM-PIM-PSM): First, platform-independent models (PIM) are transformed into another PIM representation which is then transformed into a structured textual, platform-specific (PSM) representation. For example, in UML2Alloy extended UML class models (PIM) are transformed into models of an Alloy metamodel (PIM) which are finally transformed into textual Alloy definitions accepted by an Alloy model checker (PSM).

Seven smallest common option subsets contain options through two and more decision points (D1–D6). Four out of seven option subsets *specific to two or more decision points* relate language-model formalization (D2) and language-model constraining (D3). These three subsets {2.2, 3.4}, {2.2, 3.1}, and {2.2, 3.1, 3.4} show that applying one or several UML profiles is often associated with defining language-model constraints either only textually (30 DSMLs), or by using a special purpose constraint-expression language (26 DSMLs), or both (13 DSMLs). In contrast, metamodel extensions (O2.3) are found frequently combined with both constraint-definition strategies, rather than either of the two exclusively: {2.3, 3.1, 3.4}. Metamodel extensions (O2.3) are also commonly applied together with diagrammatic syntax extensions (O4.2) and M2M transformation (O6.5). Finally, 22 of the 80 DSMLs adopt UML profiles (O2.2) for realizing a language model (O1.4).

Prototype Option-Sets. For the 80 DSMLs, we extracted seven distinct prototype option-sets (see Table 11). A prototype option-set describes complete DSML designs, on the one hand. On the other hand, it is also (frequently) extended to DSML designs to include additional decision options (see Table 4). These option sets are characterized based on the underlying UML implementation techniques: UML extension, UML piggybacking, and UML specialization (see Section 2 for definitions based on [1]).

Six prototype option-sets come *with frequent extensions*. For example, the option set describing UML-PMS, a DSML for performance modeling of mobile systems building on UML

Table 11: Overview of the seven prototype option-sets (ordered by decreasing absolute support). See Tables 1 and 2 in Section 2.1 for descriptions of the option codes. Details on the exemplary DSMLs including citations are documented in the Appendix and in Hoisl et al. [33].

Prototype	Option set	Support (abs.)	Frequency (abs.)	DSMLs (ex.)
UML piggybacking plus informal constraints	{1.1, 2.2, 3.4, 4.1, 4.6}	30	5	UML-AOF, PredefinedConstraints, UML-PMS
UML piggybacking plus formal constraints	{1.1, 2.2, 3.1, 4.1, 4.6}	26	4	REMP, CUP, UML4PF
Two-level UML piggybacking	{1.1, 1.4, 2.2, 4.1, 4.6}	22	5	SPArch, MoDePeMART, RichService
UML piggybacking for domain-specific M2T system	{1.1, 2.2, 4.1, 4.6, 6.2}	15	3	DPL, WCAAUML, WS-CM
UML piggybacking plus mixed constraints	{1.1, 2.2, 3.1, 3.4, 4.1, 4.6}	13	3	ArchitecturalPrimitives, SHP, C2style
UML metamodel (“middleweight”) extension	{1.1, 2.3, 4.6}	10	4	UML2Ext, UML4SPM, MDATC
Two-level UML piggybacking plus mixed constraints	{1.1, 1.4, 2.2, 3.1, 3.4, 4.1, 4.6}	5	3	UACL, SafeUML, and IEC61508

activities, denotes four additional DSML designs (total frequency: 5). Besides, this option set is found as a large subset in the option sets of 25 additional DSMLs (support: 30). The majority of prototype option-sets (5) involve UML profiles only (O2.2) and therefore realize UML piggybacking. Just one prototype option-set—UML4SPM, an extension of UML activities and classes to model software-development processes—builds on metamodel extensions (O2.3) only and, therefore, represents a UML extension (a.k.a. middleweight extension [82]). All six prototype option-sets, that come with frequent extensions, involve at least one concrete-syntax decision option. The only platform-integration option adopted in three prototype option-sets are M2T generator templates (O6.2). These prototypes, therefore, represent the important group of domain-specific model-to-text transformation systems using UML piggybacking.

In addition, we found a seventh prototype option-set labelled “two-level UML piggybacking plus mixed constraints” which comes *with infrequent extensions* (cf. Table 4). These DSMLs are realized at two model levels: a diagrammatic language model independent from the UML and a UML profile implementing the language in the UML. The three DSMLs for this prototype are UACL, SafeUML, and IEC61508 (see Table 11). SafeUML is a domain-specific adaptation of component and class diagrams for modeling in the avionics safety domain, to give an example. This prototype option-set is, however, only extended by less than three additional DSMLs (2). Therefore, it is considered extended infrequently. These DSMLs provide for constraints expressed both via CONSTRAINT-LANGUAGE EXPRESSION (O3.1) and INFORMAL TEXTUAL ANNOTATION (O3.4) (hence: “plus mixed constraints”).

5.4. Study Limitations

We deliberately narrowed the applicability of our findings to DSMLs embedded into UML 2.x (see Section 2). Therefore, we excluded DSMLs based on UML 1.x and other, non-UML metamodeling infrastructures (e.g. Kermet, Ecore, XMF).

While this appears, at first glance, as a barrier to generalizing the recorded design decisions, the choice of UML 2.x was necessary because important design decisions taken for the UML 2.x are substantially different from those for UML 1.x; not to mention from other metamodeling infrastructures. Regarding UML 1.x, there are important lines separating the

UML 2.x and UML 1.x regarding their language architectures and the foundational semantics of the available extension techniques (e.g. profiles, package merge; see [50, 51, 55]). By focussing on UML 2.x, we aimed at increasing the internal validity of our SLR-based study design at the expense of its external validity. Note, however, that many recorded decisions can be adopted in a broader sense to be compatible with DSMLs based on other metamodeling infrastructures and DSLs generally (e.g. concrete-syntax decisions).

Based on our SLR, we applied a documentation analysis to extract design decisions from scientific publications and their companion material. We considered supporting material if reported by and available from the publication authors. A documentation analysis represents an indirect data-collection technique [83]. Therefore, information on ordering of design decisions over time (decision sequences) often remained implicit and, therefore, unrecoverable for us. Even if documented, any indirectly observed order of decision options adopted by DSML engineers might have also followed from the presentation requirements of a scientific publication (i.e. the one reporting on a DSML); an order which does not necessarily correspond to the original one during decision making. Therefore, in our research setting, we can only study option sets in terms of decision associations. For the same reason, we focused on one process style of DSML development only (i.e. language-model-driven development; see Section 3). We might have neglected design decisions characteristic for other development styles (e.g. mockup-driven DSML development; [22]).

SLR studies such as this one have the major problem of finding a representative set of relevant primary studies. In general, we closely followed established guidelines on designing and conducting SLRs available from research on evidence-based software engineering to avoid any pitfalls [37, 39, 40]. As for the search strategy, we used an extensive automated search based on four search engines by primary publishers in the software-engineering field (ACM Digital Library, IEEE Xplore, Scopus, and SpringerLink). This search yielded 5,778 hits which included the majority of reference publications (28/37) previously collected in a manual search. Backward snowballing added another reference publication. Therefore, our automated search missed eight publications that should have been found. While the missing ones were considered for inclusion in a separate step, this

still indicates that we risk having missed other relevant primary studies on UML-based DSMLs, in more general. Note that we addressed this threat right from the beginning, by building our review procedure around the principle of continuous search validation and search refinement driven by a QGS as a recommended practice [40, 74].

We applied a review and extraction process involving two independently working data extractors on each item (search hit, publication, DSML). This was to avoid limitations of an extractor-checker procedure [84]. On the flip side, this required us to control the workload per data extractor in the light of the excessive result sets when checking for false negatives. Recall that we were required to process 5,778 hits from automated search and 2,337 from snowballing. A false negative is a hit which is deemed excluded by one data extractor, in the first iteration during selection, and which becomes included after a negotiated agreement between the two extractors. This negotiated agreement is triggered by an independent inclusion decision of the second extractor. In these cases, we drew subsets of hits for the second, independent assessment (i.e. 20% random samples). As elaborated in Sections 4.1.2 and 4.1.3, due to this sampling and in a probabilistic projection (99% confidence level), we risk having missed out between zero and 22 publications as false negatives during automated search (from 704 negatives) and between zero and seven publications during snowballing (from 170 negatives), respectively.¹⁹ This threat cannot be neglected, however, we mitigated for this risk in two ways. First, we applied sampling for false-negative checking only for decisions which were based on structured and objectively evaluable decision criteria (e.g. publication year, venue names based on predefined venue lists). This way, we hope to have minimized the risk of undisclosed disagreement on negatives between the respective two data extractors. For the actual decision-data extraction, which involved more subjective assessments, no sampling was applied. For example, to establish whether a DSML's design documentation is of unacceptable quality, each DSML design was reviewed by two independent data extractors. Second, we consider the results reported (frequency counts, prototype designs) robust against the risk of false negatives. Consider, for example, the substantial support reported for the prototype designs. Each prototype design is observed for between ten and 30 DSMLs (see Table 11). It is unlikely that false negatives in the ranges above would reduce the observed support of just one or of a few prototypes, so that they would not be observable anymore.

As a result of our quality-assessment procedure, we excluded another 31 DSML candidates due to their incomplete and/or incorrect documentation of their designs. More precisely, we watched out for critical, UML-specific specification defects [55, 64, 80] which would prevent us from extracting the decision options unambiguously, based on the design artifacts available to us via the corresponding scientific publications (e.g. diagrams, structured and unstructured text on

design rationale).

At the same time, however, by excluding the DSMLs, we risk having biased our quantitative results (e.g. frequency counts and prototype designs). By making this exclusion decision, however, we strengthened the qualitative validity of our results by extracting decision options and option sets from DSMLs of acceptable documentation quality only. For the reasons put forth for potentially false negatives (see previous paragraph), we are also confident that our quantitative findings are robust against the exclusion of these DSML designs.

In our search design, we also omitted certain kinds of publications explicitly. On the one hand, we excluded publications appearing in venues not considered SE venues. This was realized by filtering venues according to publicly available third-party sources enumerating SE venues. The main objective was to put focus on DSML documentation artifacts which are likely to reflect on DSML and UML design rationale; and to avoid large numbers of publications from outside the SE domain. On the other hand, we excluded certain venue formats (e.g. workshops) and grey literature (i.e. technical reports, working papers) to increase the probability of building our analysis on documentation artifacts having a certain maturity. In addition, we targeted publications having been subjected to community-driven quality assessment (e.g. peer reviews). A similar limitation is that we restricted the search hits retrieved from the automated search based on the QGS corpus publications (see Section 4.1): from 5,778 total hits, 2,678 containing all QGS publications found during the automated search were considered for selection. This was necessary to reduce the number of publications to manageable levels for reading during selection by two independent raters. While settling us at a sufficient level of search validity (75.7% [40]), it means that we may have missed relevant publications. Note, however, that we undertook backward snowballing to make sure that we did not skip work considered relevant by the authors of the primary studies obtained as search hits.

There is a bias inherent to our SLR design in that by relying on scientific publications only, the recorded design decisions on DSMLs risk being specific to DSMLs as research-driven prototypes and proof-of-concept implementations. Design decisions during DSML development in industry might not necessarily be covered by our revised catalog and the findings in this study. The industrial case in Staron and Wohlin [16] exemplifies this. This case is based on more balanced numbers of metamodel extensions (2) and profiles (3). At a coarser-grained level, our study reports a clear preponderance of UML profiles in 80% of the reviewed DSMLs. However, we find it difficult to assess the severity of this bias. To begin with, the primary studies reviewed in this SLR did not disclose their industrial background, if any at all. Similarly, while related empirical studies on UML usage certainly document the existence of UML extensions and UML-based DSML designs (see, e.g., [12, 13]; see also Section 7), they do not discriminate between industry-driven and research-driven projects.

¹⁹Note that the two intervals cannot be simply summed up because of the backward dependency between snowballing on automated search.

For selection and data extraction, we used a process involving two independently working data extractors on each item (search hit, publication, DSML) to avoid pitfalls of an extractor-checker process in complex extraction settings and to follow closely corresponding recommendations [84]. The observed disagreement about publications and DSMLs between the extractors was rigorously documented using established IRR statistics, adding to the detailed documentation of the final extraction agreed by both extractors. To reduce some of the overhead of a two-extractor process, we had one extractor perform selection and extraction decisions on reduced data sets. This involved a first extractor reviewing all included items, but only a sample of the excluded ones by the second extractor. A final issue with respect to extraction and selection was that due to a number of constraints arising during the process, two of the three authors performed the majority of extractions. Constraints included personalized access to publisher databases and individual time constraints. This might potentially have introduced bias.

An important motivation of this SLR study—besides validation—was to remove personal bias from the decisions recorded in the initial catalog [45]. This personal bias was most likely introduced by looking at our own DSML projects at that time. However, even for the SLR study, we could not fully avoid personal bias because we risked collecting data from our own publications, if retrieved at the various SLR stages (QGS corpus construction, automated search, and backward snowballing). Therefore, we excluded our own DSML projects from the result set so that from the 80 DSMLs obtained from this SLR *none* was developed by the authors themselves. Furthermore, there was only one third-party DSML (i.e. UML4SOA in [61]) which was more familiar to us than the others, because one of our own DSMLs is built on top of it. Nevertheless, personal bias cannot be ruled out, simply because our understanding of DSML development has entered our judgments during data extraction.

5.5. Catalog Revision

An important, secondary outcome of the SLR study is a revision of the initial decision-record catalog [35, 45]. The updated and improved version included changes to the catalog’s content as well as to its presentation. Table 12 gives an overview of the changes resulting in the current decision-record catalog [33]. As a result of the SLR and our frequent-item-set analysis (see Section 5), the content items in the decision-record catalog were revised (decision points, options, and associations; see Section 6).

Important but comparatively few additions and modifications to the catalog were necessary. As shown in Table 12, the main addition content-wise was the inclusion of decision record D5 on decisions relating to *behavioral specification*. While it was necessary to cover this sub-space of decision options to fully characterize the DSMLs in our study, we also found that except for very few DSMLs, explicitly documented behavioral specifications are widely missing (see Section 5.3). However, within the limits of our study, this can be explained straightforwardly because the majority of

the reviewed DSMLs focus on structural viewpoints (i.e. UML classes) and the corresponding application domains did not require an additional behavior specification. In addition, some of the DSMLs reviewed for our study directly reuse or slightly adopt UML built-in behavior (e.g. provided by UML activities). Thus, DSMLs providing a dedicated behavior specification were documented as known-usage examples in the newly added decision record D5.

Apart from D5, only one new decision option was added. O6.5 for D6 (*platform integration*) allows for characterizing DSMLs which build on endogenous M2M transformations for platform integration, i.e. transformations between models defined over the same metamodel—rather than using M2M transformations only for creating different intermediate model representations within an MDD tool chain (which is covered by O6.1). In addition, we refined the descriptions of three decision options (O2.1, O2.2, O4.3) to cover aspects revealed by the reviewed DSMLs and to better discriminate between decision options. For example, we revised decision option O2.2 (formalizing a language model via UML profiles) to include scenarios of extending and/or redefining existing UML profiles.

Table 12: Comparative overview of the numbers of content items (e.g., decision records, decision options) included in the decision-record catalog in its pre-study version [45] and in the post-study revision [33]. *: The numbers in round brackets indicate that four decision options serve for coding pseudo-decisions only (e.g. not choosing any option); **: These entries show the numbers for third-party studies (DSMLs) compared to *all* studies (DSMLs), including our own.

	Pre-study revision	Post-study revision
Decision records	5	6
Decision options	22 (26)*	27 (31)*
Decision drivers	22	27
Decision consequences	9	13
Decision associations	11	21
Underlying resources		
Primary studies	6/ 19**	84/ 97**
DSMLs	6/ 16**	80/ 90**
Secondary studies	15	25

The resulting catalog lists 27 decision options and 21 associations between these options (e.g., dependencies; see also Table 12). In addition, the catalog provides seven prototypical solutions of existing DSMLs which represent commonly adopted combinations of decision options from 80 third-party DSMLs. Note that the catalog also records ten self-developed DSMLs summing up to a total of 90 DSMLs (see also the Appendix). 40 decision drivers and corresponding decision consequences (e.g., forward dependencies on follow-up decisions) are available to assess the decision options (e.g., in terms of rationale tables). In addition, the catalog offers application examples and implementation sketches taken from the 90 DSMLs that we examined in detail. Moreover, the decision catalog provides references to 25 secondary studies on DSML development (e.g., [47, 85]).

6. Discussion

In this three-year research project, we started from our own experiences gained from building ten UML-based DSMLs to define a preliminary decision-record catalog for such DSMLs. To extend and to validate this decision catalog, we performed this systematic literature review (SLR). Our multi-stage SLR initially returned 8,152 publications: 37 articles originated from a quasi-gold standard corpus, 5,778 from the automated and 2,337 from the backward-snowballing searches. 7,069 papers were unique, 5,023 were considered for selection and quality assessment. Subsequently, we filtered these publications based on multiple quality criteria. As a result, we identified 84 publications on 80 UML-based DSML projects (see Section 4 and also the Appendix). These 80 DSMLs are “third-party” artifacts, that is, we did not participate in developing them.

Research question 1: *What are the design-decision options for UML-based DSML designs reported in scientific literature?*

Based on the literature review and content analysis, we tested scientific literature for the presence or absence of 27 decision options. From these, 24 were reportedly adopted by at least one DSML project. We found that three options are not contained by any of the 80 DSMLs (see Section 5.3): CODE ANNOTATION, CONSTRAINING MODEL TRANSFORMATION, and CONSTRAINING MODEL EXECUTION (i.e. O3.2, O3.3, and O5.4). This follows directly from the *saturation* test for the coding schema reported in Section 5.1.

CODE ANNOTATION (O3.2) and CONSTRAINING MODEL TRANSFORMATION (O3.3) describe practices to realize and to enforce language-model constraints *indirectly*. On the one hand, this involves constraints defined as host-language expressions for a certain API (e.g. Java expressions binding the Ecore Java API; O3.2). On the other hand, such constraints could be realized in terms of assertions on a DSML model that are defined as a part of corresponding model transformations on the respective DSML models (O3.3).

Similarly, enforcing behavioral constraints (D5) via partial or prohibitive model-execution engines (CONSTRAINING MODEL EXECUTION; O5.4) was not documented for any DSML. We did, however, find these options documented in secondary studies (for O5.4 see, e.g., [86]) and applied them in our own DSML projects (for O3.3 see DSML SOFServices). As a result, our decision-record catalog lacks known-usage examples for the three decision options mentioned above. Hence, we highlighted these decision options as *candidate* options in the post-study revision of the catalog; that is, candidates requiring further evidence.

Recall that there are three basic UML implementation approaches for DSMLs: language-model extension, language-model piggybacking, and language-model specialization (see Section 2). The first two have been found employed frequently in the 80 DSMLs, most notably language-model piggybacking (see discussion of RQ2). Language-model specialization involves a metamodel derived from the vanilla UML2

metamodel systematically, e.g., through metamodel cloning and pruning [53, 54]. This approach has also been described as a “heavyweight” metamodel extension [82]. Language-model specialization is driven by the need for performing a METAMODEL MODIFICATION (O2.4), that is, the need for re-defining existing UML metaclasses, their features, and their associations. The specialization approach comes at certain maintenance and deployment costs (e.g., tracking changes in the UML metamodel, low tool interoperability [82]).

In total, we only found two DSMLs for which a METAMODEL MODIFICATION is reported: eSPEM and DMM/UCMM. The former further extends the UML2/SPEM metamodel (activities and state machines) to model software-development processes and artifact states. The latter provides a modeling frontend for automated GUI generation. DMM/UCMM adds new properties to existing UML metaclasses (e.g. Class and Property). Besides being comparatively infrequent, the respective design reports fall short in two ways:

First, they do not clarify how they realized the language-model specialization technically, i.e., managing metamodel derivation. Second, they also lack statements which document the intention and rationale for performing a METAMODEL MODIFICATION in the first place. Often, a METAMODEL MODIFICATION appears merely *accidental*. For instance, in a third DSML UCDM, existing UML metaclasses (e.g. UseCase) are associated with newly introduced metaclasses (e.g. UseCaseDescription). The metamodel definition for UCDM is underspecified regarding the ownership of association ends and, hence, regarding the choice of a METAMODEL EXTENSION (O2.3) or a METAMODEL MODIFICATION (O2.4):

(1) Both ends could be owned by the association, leaving the UML2 metamodel unchanged (O2.3); (2) one end could be owned by the association, the other one by a metaclass (O2.3 or O2.4, depending on whether the owning class is coming from the UML metamodel); or (3) both ends could be owned by their corresponding metaclasses (O2.4).²⁰

For these three reasons (infrequency, specification ambiguity, unknown metamodel management), we cannot conclude from our study that language-model specialization has been adopted.

Even when omitting the decision options not observed in the 80 third-party DSMLs, the design space described by the 24 found options amounts to a vast number of possible option combinations.²¹ Therefore, the frequency patterns identified for answering RQ2 turned out to be key.

Research question 2: *What are frequently observed decision options and frequently observed combina-*

²⁰Note that the UML has these degrees of freedom, however, to avoid ambiguities, association end ownership can always be made explicit using the dot-notation [10].

²¹See Section 3 for the rationale behind these combinatorial computations; pre-study revision (all options; 3,686,400): $2^4 - 1$ (D1) times $2^4 - 1$ times (D2) times 2^{22-4-4} (D3–D5); post-study revision (all options; 117,964,800): $2^4 - 1$ (D1) times $2^4 - 1$ times (D2) times 2^{27-4-4} (D3–D6); post-study revision (excluding unobserved options; 29,491,200): $2^4 - 1$ (D1) times $2^4 - 1$ times (D2) times 2^{25-4-4} (D3–D6).

tions of decision options (option sets) in and across existing UML-based DSML designs?

From a total of 24 options (RQ1), 16 options are frequently found in three or more projects. Among the most frequently adopted options, the following are noteworthy (ordered by frequency): DIAGRAM SYMBOL REUSE (O4.6), PROFILE RE-/DEFINITION (O2.2), MODEL ANNOTATION (O4.1), INFORMAL TEXTUAL ANNOTATION (O3.4), CONSTRAINT-LANGUAGE EXPRESSION (O3.1), FORMAL DIAGRAMMATIC MODEL (O1.4), META-MODEL EXTENSION (O2.3), and GENERATOR TEMPLATE (O6.2) (see also Table 10). We discuss their relevance below in their broader context, not necessarily in the above order.

A domain-analysis step in terms of domain engineering [87] and domain-driven design [34]) is contained by most process guidelines available for DSL and DSML development (see, e.g., [4, 22]). One challenge is to identify and to define the domain abstractions and their relationships in a manner which allows for 1) a detailed specification independent from a concrete modelling language and for 2) having the abstractions enter a UML-based language-model implementation in a seamless (e.g. instantiation-based model layers) and/or guided way (e.g., through transformations). Another challenge is making domain variability explicit, for example, in terms of feature models [87]. Guidelines and the 84 primary studies reviewed put emphasis on (semi-)structured textual analysis techniques, such as domain-vision (scoping) statements, domain-distillation lists, and feature tables. This is indicated by *all* 80 DSMLs adopting an INFORMAL TEXTUAL DESCRIPTION (O1.1) to capture their application-domain definition.

However, prior to our study, very little was known about using a FORMAL DIAGRAMMATIC MODEL (O1.4) for a domain-analysis step. A FORMAL DIAGRAMMATIC MODEL leans itself towards a seamless and guided UML implementation of the resulting language model (e.g., in the sense of model chaining through transformations). It can also serve as the definitional basis for an FORMAL TEXTUAL DESCRIPTION (O1.2; e.g., as a type graph using set theory). We found that only 23 DSMLs adopt such a FORMAL DIAGRAMMATIC MODEL (O1.4). A closer look reveals that the dominant modeling formalism in this group are E/MOF and/or UML class diagrams. Alternatives such as Ecore or entity-relationship models (while interchangeable to some extent) are not reported. Only a single DSML design (SOA) employed a type-graph notation in the sense of a graph-transformation system.

Two follow-up observations are noteworthy: First, the usage of E/MOF and/or UML class diagrams most often prepares the definition of UML profiles, rather than alternative UML language-model implementations. This is confirmed by the two “two-level UML piggybacking” prototypes (out of 7) identified in Section 5. This is typically justified by referring to profile-specific guidelines such as [69, 70, 71, 72, 88]. Second, although E/MOF and/or UML class diagrams are available, they are *not* formally linked to the corresponding UML implementation model e.g. by using model transformations, inter-model consistency constraints, and/ or instantiation-

based model chaining. This is despite the availability of suitable approaches (see, e.g., [89]). Rather, convention-based approaches dominate. A popular example is the 1:1 name mapping of language-model elements into equally named stereotypes.

The importance as well as relative pros and cons of UML profiles as implementation vehicles have been covered to some extent [69, 70, 71, 72, 88]. Basic evidence on the actual adoption level, however, has been missing or has remained partial. 62 DSMLs implement their language model using newly defined and/ or redefined UML profiles (PROFILE RE-/DEFINITION, O2.2). As a consequence, “UML piggybacking” characterizes 5 of the 7 prototype designs in Section 5. This is partly a confirmatory finding for observations of prior, but inherently limited empirical studies. For example, both Pardillo [64] and Nascimento et al. [13] characterize publication-centric trends on UML profile usage. However, they discuss UML profiles in isolation. Our findings provide empirical insights on UML profile usage in the context of implementation alternatives, which is entirely missing from previous studies: As a general observation, UML profile usage in 62 DSMLs contrasts with 17 DSMLs using UML metamodel extensions.

Furthermore, we reveal coupled design decisions such as on concrete syntax. The popularity of MODEL ANNOTATION (e.g., UML comments carrying domain-specific keywords or slot specifications; O4.1) directly follows from the UML profile preponderance: The 62 profile-based DSMLs also adopted this option. The same coupling can be found between UML profiles and the concrete-syntax option DIAGRAM SYMBOL REUSE (O4.6) in 69 DSMLs. The latter coupling partly results from the default that stereotyped elements carry on the notation defined for the extended UML metaclass, if one is defined and not overruled explicitly.

Researchers have started systematizing the use and the role of model-level consistency constraints for the UML in general [90]. We complement these findings for UML-based DSMLs. We established that 31 DSMLs adopted the CONSTRAINT-LANGUAGE EXPRESSION (O3.1) option. To be more precise, they all employ OCL expressions. This level of adoption puts new focus on the portability issues of such OCL consistency rules between different evaluation engines. Portability remains limited due to the OCL/UML language specifications leaving critical details to language and tool implementers (e.g., navigation semantics between extension and extended model elements; see [91] for an overview).

In addition, we found that constraint expressions are predominantly defined for an intra-model scope (e.g. to resolve ambiguities in the language model). Inter-model constraint expressions, including vertical constraints between different model levels (e.g. platform-independent vs. platform-specific), are the minority. Consider as an example a DSML language model defined at level M2 which must enforce consistency rules at level M0, i.e. the occurrence (instance) level of DSML models. In fact, we found that for such scenarios, DSML developers resort to an INFORMAL TEXTUAL ANNOTATION (O3.4). The use of implementation idioms (e.g., prototypical concept pattern [92]) and alternatives to metamodel-

eling based on shallow instantiation (e.g., potency and deep instantiation [46]) to work around or to overcome this limitation was not observed.

Generating textual artifacts (e.g. source code, documentation, deployment descriptors) from models in software development via model-to-text (M2T) transformation is a common MDD activity [93]. The importance for model-driven development of M2T transformations, in general, and M2T transformation systems based on generator templates, in particular, has been reported earlier [94]. For UML-based DSMLs, such evidence has been missing so far. We found that the majority of DSMLs (16) providing for platform integration (D6) adopt the GENERATOR TEMPLATE (O6.2) option. This is a mandate for DSML-specific research on M2T transformation. Examples include scalable automated change propagation from domain-specific models to text artifacts [93] and coupled evolution of generator templates under metamodel composition for DSML integration [95].

Frequency Patterns. In total, our study identified and describes 16 frequent combinations of decision options: 9 smallest common option subsets and 7 prototype option-sets (see Section 5). The latter are summarized below. To be considered frequent, an option subset was required to occur in at least three different DSMLs (see Section 3). During this pattern analysis, we learned that the reviewed DSMLs have a maximum of ten decision options per DSML (e.g. UML4SOA). For recurring option subsets, i.e. option subsets found in more than one DSML, the maximum number of options in a particular option subset was seven (e.g. UML-PMS). At the level of individual decision points, these maxima translate into frequently recurring subsets containing options from three decision records only: language-model definition (D1), language-model formalization (D2), and concrete-syntax definition (D4).




Figure 17: A feature diagram [96] representing the seven prototype option-sets found in the pool of 80 third-party DSMLs. These seven prototype option-sets include nine different decision options. Each of the seven observed prototype option-sets listed in Table 11 is one of the possible configuration of the feature space.

As a key finding, we extracted seven option subsets which form so called *prototype option-sets*. These prototype option-sets are a particularly useful way for structuring the design space described by the 24 observed decision options. First, they cover a critical share of the observed DSMLs. Second, they stress commonalities and differences in terms of deci-

sion options between these highly representative option combinations. In particular, the seven prototype option-sets characterize 30% of the observed DSMLs (24/80) in their entirety. Furthermore, they are contained as large proper subsets by 25 extended option sets; therefore reaching a total coverage of approximately 61% of the DSMLs included in our study (49/80). By looking at the common and varying decision options in the seven prototype option-sets, we find that all seven are combinations of nine decision options. These nine decision options correspond to the leaf elements of the feature diagram in Figure 17. As a result of this observation, these nine options were highlighted in our catalog. In addition, we extended the catalog to include additional reading aids based on the prototype option-sets. For example, we added a variant of the feature diagram from Figure 17 as well as thumbnail descriptions of the nine key options (see also Tables 1 and 2 in Section 2.1). Adding these content and navigation items is a concrete means to render the catalog tailorable [25, 65].

Moreover, we were able to identify ten additional associations between different decision options that are not yet included in the prototype option-sets (see also Table 12). Associations were found in terms of the two smallest common option subsets (e.g. model-transformation chaining in D6; see Section 2.1) and in terms of associations derived from the prototype option-sets. From our analysis, we also learned about the frequency at which certain associations occur. This finding also allowed us to confirm already documented associations, such as the association between UML profiles and certain concrete-syntax decisions. To reflect this evidence for the reader of the catalog, we annotated the frequently observed associations accordingly (ten in total).

Revised Catalog for DR Reuse. As an important by-product of our study, the revised decision-record catalog serves as an evidence-based and empirically validated foundation for decision making in the context of UML-based DSML projects (see Figure 18). For the *identification* of key decision problems, decision makers can consult our decision-record catalog, which documents recurring decision contexts and decision problems. The decision options and associations support the *inventory* of suitable design decisions for deriving a selection of candidate solutions. Finally, the *evaluation* of candidate solutions for their fit to domain-specific and domain-generic requirements can be performed using the documented decision drivers and decision consequences. In addition, the decision options are aligned with different styles of DSML development (abstract-syntax-first vs. concrete-syntax-first vs. extraction-based; see, e.g., [4, 22]).

The importance of the frequency patterns summarized above stems from the fact that the documented decision options provide for a huge number of possible design option-sets. Thus, without evidence-based guidance, making an informed choice from this amount of combinations becomes impractical. Likewise, performing design-space analyses or creating design-process documentation based on the decision-record catalog can become a time-consuming and

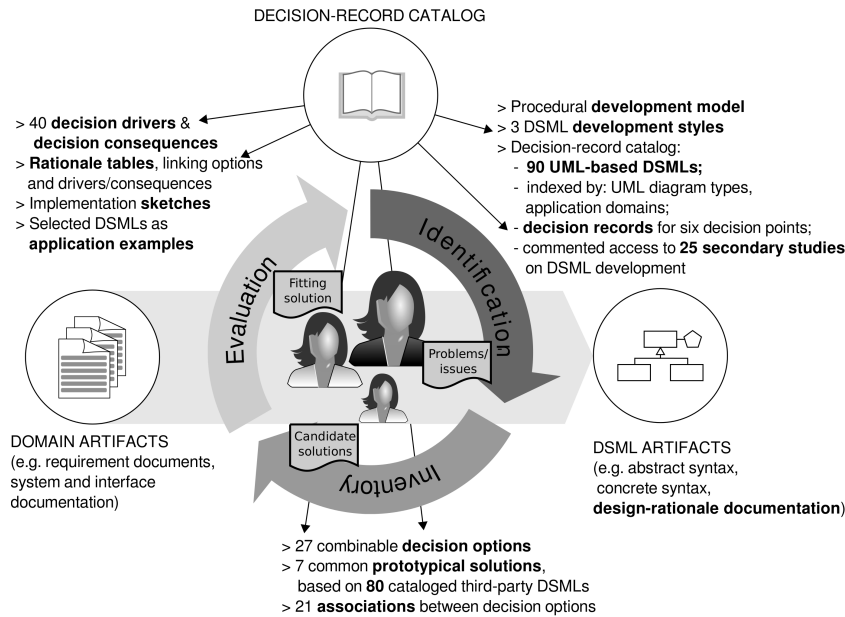


Figure 18: Supporting decision making (identification, inventory, evaluation) on developing UML-based DSMLs via an evidence-based *decision-record catalog* providing reusable design decisions [33].

tedious task (see Section 2.2). This is because of the sheer size of the documented design-decision space. Creating and maintaining design-rationale documentation for a DSML project, however, is a substantial investment, and the effectiveness of this investment is largely affected by the tailorability of the reusable documentation of generic design rationale. A *tailorable* documentation of generic design rationale [25] provides means for customization such that for each DSML project only relevant documentation items are provided to decision makers (e.g. domain engineers, modelers, language engineers, and software architects). Moreover, customization usually involves supporting visualizations (e.g. decision-flow and activity diagrams) for particular use cases [65].

7. Related Work

As explained in Section 2, the practice of documenting generic design decisions via collections of structured descriptions is an established procedure in the domain of software engineering [23, 24]. In this paper, we also documented *reusable* design decisions in a structured and *tailorable* manner [25]. The way we documented the generic DSML rationale in this paper is related to similar documentation techniques applied for software architecture design and architectural design decisions [27, 60, 68, 97].

Moreover, we identified two additional areas of related work for this paper. On the one hand, we have related work on DSL design procedures and DSL design-decision making (see Section 7.1). On the other hand, there are related contributions which have systematically gathered empirical evidence as well as generic design rationale on using the UML for domain-specific modeling (see Section 7.2).

7.1. Systematic Development of Domain-specific Languages

Several related contributions exist that describe systematic procedures for developing DSLs. Each of these approaches is based on experiences drawn from actual DSL engineering projects, and each of the related approaches provides insights into the DSL development process, or into certain aspects of DSL design, or into DSL-related design decisions. For example, Strembeck and Zdun [22] discuss different DSL development activities and describe how these activities can be combined to tailor a DSL engineering process.

In a complementary contribution, Zdun and Strembeck [4] document three main decisions to be made when applying the DSL development process from Strembeck and Zdun [22]. These decisions relate to the choices of a specific type of DSL development process, of a concrete syntax style, and of developing an external vs. an embedded DSL. To render these decision descriptions reusable, Zdun and Strembeck [4] document them in a pattern-like format. In software engineering, a pattern is a time-proven solution to a recurring design problem. A pattern description includes (at least) a “problem description”, a description of the “context” in which the respective problem occurs, and one or more (alternative) “solutions”. Typically, pattern descriptions also include different “forces” that may influence the choice of a certain solution, “consequences” that arise from a solution, as well as “known uses” of a particular solution. In this way, the description format we chose for the decision records resembles a pattern format to a certain degree. However, decision records are not identical to or variants of software patterns, since, for example, they list multiple solution propositions (decision options) rather than one.

While prior work on patterns for DSL development [4, 22] aims at describing generic procedures and decisions for DSL

development projects, our contribution in this paper provides detailed insights into design decisions for UML-based DSLs. In this way, our work complements [4, 22], as well as other DSL development approaches such as [9, 98]. This is because our work provides a systematic and in-depth documentation of the follow-up decisions that DSL engineers face after they decided to develop a UML-based DSL.²²

A number of other patterns and pattern languages exist that can be applied in DSL development and are thereby complementary to our work. This includes patterns for the design and implementation of DSLs [1], patterns for evolving frameworks into DSLs [99, 100], and approaches for pattern-based DSL development [101]. Often, DSL-related patterns do not only describe how a DSL is developed, but also why it is developed in a specific way. In addition, pattern languages also describe potential sequences in which the patterns can be applied [32]. Pattern sequences compare with our notion of option sets (see Section 2.1.2) in the sense that (ordered) option sets can represent sequences of adopted decision options. The research design of this study, however, does not allow us to assume a temporal ordering when extracting design decisions (see Section 5.4).

Mernik et al. [2] used the patterns from Spinellis [1] to conduct a survey on decision factors affecting the analysis, design, and implementation phases of DSL development. These decision factors can be considered during DSL development. For example, the decision factor *Notation* deals with the consideration whether the DSL should provide a new or an existing domain notation. For a few decision factors, Mernik et al. [2] suggest implementation guidelines. The work of Mernik et al. [2] is complementary to ours as it focuses on general issues of design-decision making and implementation, rather than on design decisions for a specific (host) language environment such as the UML.

Another group of related work reports observations from developing DSLs in (industrial) practice. For example, Luoma et al. [102] conducted a study including 23 industrial projects for the definition of DSMLs. Similar to our approach, a number of DSLs are systematically compared. However, in contrast to our paper Luoma et al. [102] provide a high-level description only and do not describe specific DSL design decisions or decision-option sets in detail. Similar to patterns, lessons learned have been used as a vehicle to preserve best practices of DSL development. For example, Wile [103] reports on twelve lessons learned from three DSL experiments. For each lesson, he introduces a respective rule of thumb and gives an overview of the experiences that are the origin of the corresponding rule. Despite Wile's lessons learned being described at a comparatively high level of abstraction, they can, in general, also be observed in our work and are hence reflected in parts of our decision-record catalog. Kelly and Pohjonen [104] present a report on worst practices found by reviewing 76 DSL development projects, and Karsai et al. [105]

propose 26 general guidelines for DSL development derived from their own experiences.

A UML-based DSL uses UML as its host language and extends the UML with domain-specific language elements and, therefore, qualifies as an embedded DSL (also: internal DSL). Related work on developing embedded DSLs includes the contributions by Günther et al. which describe a process and corresponding patterns for the development of internal DSLs on top of dynamic programming languages, such as Ruby or Python [57, 106, 107]. Other related contributions describe how to build DSLs from component building blocks that can be incrementally designed and composed (see, e.g., [108, 109]). This idea originates from approaches such as keyword-based programming [110], in which so called "keywords" serve as building blocks for DSLs. In particular, a number of (universal) keywords are suggested which are then glued together to compose DSLs. This approach was first envisioned in Landin [111] and is akin to building embedded DSLs in dynamic languages (such as Ruby, Perl, Python, or Tcl for example).

In the UML context, some authors propose approaches that define domain-specific UML extensions via UML profiles (see, e.g., [69, 70, 71, 72]). While each of these approaches is related to our work, none of them document generic design decisions for UML-based DSLs. Weisemöller and Schürr [112] give an overview of standard compliant ways to define domain-specific UML extensions, while Atkinson and Kühne [88] discuss potential issues with UML profiles and suggest a solution to address these problems. Bruck and Hussey [82] present different techniques for tailoring the UML (e.g. lightweight profile or middleweight metamodel extension). In particular, Bruck and Hussey [82] define a catalog of options and characterize different extension mechanisms accordingly. They also discuss pros and cons of using one approach or the other. However, Bruck and Hussey [82] focus on UML customization techniques in general and do not integrate design decisions in the process of DSML development (e.g. no development phases are distinguished, language-model constraints as well as platform integration are not considered).

In addition, knowledge on DSL design decisions can also be gained from analyzing toolkits for DSL development. For example, Tolvanen and Kelly [113] present a tool for the definition and usage of integrated DSMLs. Similarly, Zdun [114] presents a tool suite for textual DSL-based software and provides a discussion of architectural decisions for DSL development. However, most existing contributions have a strong focus on textual domain-specific programming languages. To the best of our knowledge, there is no report reflecting on design decisions embodied in toolkits for UML-based DSML development.

In summary, the related work on patterns, best practices, and lessons learned in DSL development have in common with our approach that all are based on experiences from actual DSL projects and contain some information on DSL design decisions and design rationale. Our work provides a systematic and detailed description of decision options for

²²Remember that each UML-based DSL is an embedded DSL and that UML-based DSLs usually have a graphical concrete syntax or a mixture of graphical and textual concrete syntaxes (see also Section 2).

building UML-based DSLs. In this way, our contribution is complementary to those other approaches and can be combined with them.

7.2. Empirical Evidence on UML and MDD Usage

We identified a number of systematic UML-related empirical studies closely related to our paper. The research methods employed in these studies include systematic reviews (5), surveys (2), a controlled experiment (1), and a case study (1). Our findings complement the empirical observations of a first group of related work (e.g. descriptive statistics on UML usage).

Dobing and Parsons [115] performed a web-based survey among 171 OMG members on the perceived importance of the use-case viewpoint offered by the UML. They found that UML class and use-case diagrams are the most widely and regularly adopted diagram types (i.e. in more than two thirds of the respondents projects), with state machine and collaboration diagrams ranking lowest. With respect to UML-based DSMLs, our study shows a slightly different picture (see Section 5). While we can confirm the preponderance of class diagrams, we found that use-case diagrams are rarely customized by means of a DSML while state machine diagrams are frequently adopted in DSMLs.

Pardillo [64] conducted a SLR for the years from 1999 to 2009 to characterize trends on UML profile usage. Most notably. The SLR documented publication trends, descriptive statistics (e.g. extended metaclasses, profile-concept usage frequencies) and definition issues of documented profile packages. The review was limited to a few handpicked, but important presence venues (ER, MoDELS/UML) including co-located workshops. An automated search in DBLP was used, yielding 63 hits. 39 of those 63 were included in the study. After a peak in 2002, Pardillo observes a decrease in publications that present a UML profile. However, the SLR study of Pardillo [64] is not directly comparable to ours. This is because of a different review period, our study's reach beyond UML profiles, and because of our more extensive pool of venues. We cover journals and conferences beyond MoDELS/UML and ER. Yet, in our study, we also observed a falling number of yearly publications on UML-based DSMLs including UML profiles (with intermittent peaks in 2006 and 2011). Conference publications are on the decline and journal publications stagnate. We can also confirm that the MoDELS conference continues to serve as the top venue for relevant publications. It ranks first in terms of included publications in our paper corpus, while the ER conference only contributes one publication (see Section 4.2).

Nascimento et al. [13] conducted a systematic mapping study on DSLs in more general and for an unbounded time range (1966–2011). They collected and reviewed 2,688 publications from 669 presence venues (i.e. conferences and workshops) and 180 archival venues (i.e. journals). The number of papers relating to (UML-based) DSMLs amounted to 163, 69 of which report on DSML creation. The data extracted from these DSML publications is unfortunately not discussed in detail in the research reported by Nascimento

et al. [13]. Only the number of DSML publications (21) referring to the usage of UML profiles is documented. To this end, we complement the broader mapping work of Nascimento et al. [13] for the later years in our review period by detailing the usage characteristics of profiles versus metamodel extensions in UML-based designs. The DSML-specific publications in their paper corpus served for establishing our QGS corpus as a third-party data source (see Section 4).

Garousi [116] performed a search-driven literature review by collecting bibliographical records on UML-related monographs from the Google Books database for a time range between 2001–2009. As a major observation, the authors state a peak year in 2005, allegedly caused by authors writing new books on the UML 2.0 specification. From 2005 onwards, however, the study finds a substantial and continued reduction in new UML books being published, down to only two books in 2008. Our study draws a comparable timeline picture for scientific publications on UML 2.0 (see Section 4.2).

Nugroho and Chaudron [117] provide empirical findings from an online questionnaire survey on the use of UML amongst 80 professional software engineers. Part of the results targeted “imperfections” in UML models; i.e. how often inconsistency, understandability, inaccuracy, and incompleteness in models led to implementation problems. The study indicates that in each of the four categories approximately 90% of the respondents reported the respective imperfection occurring sometimes, often, or very often during a project.

Hutchinson et al. [12], Whittle et al. [118] report quantitative data on MDD practice in industry based on 449 responses to an online questionnaire, supplemented by qualitative data from 22 semi-structured interviews with MDD practitioners. Their report suggests that, in practice, the UML is the most commonly adopted modeling language with 85.5% of the respondents having adopted the UML. Another key finding, however, is that DSMLs have gained an important share in adoption. 39.2% of the respondents indicated that they use a custom DSML developed in house, 21.5% make use of DSMLs bundled with their modeling tool chain. Unfortunately, their report remains silent about the share of co-adoption of the standard UML and such DSMLs. Also, the authors do not discriminate between DSMLs based on the UML and those developed independently from the UML. In an earlier study, Hutchinson et al. [14] hint at the fact that such DSMLs could qualify as DSMLs embedded within the UML (e.g. using profiles). Therefore, these studies provide motivation for our work, because they stress the need for systematic empirical research on UML-based DSML development to fill the above gap. Regarding platform integration (D6), only $\approx 32\%$ of the DSML projects considered platform-integration techniques. This observation is confirmed by our study (see Section 5.3). Hutchinson et al. [12] report a comparatively high number of projects using MDD for code generation (O6.2) as well as M2M transformation (O6.1, O6.5). Even fewer respondents indicated to rely on executable models (O6.4)—observations which are supported by our results. In addition, the reported frequency of UML diagram types matches our observations of a dominance of structural and, in particular, class diagrams

(see Table 9).

Recently, Gorschek et al. [119] reported an inquiry concerning the adoption of object-oriented design and implementation concepts (including object-oriented design models) using a web-based survey. 4,823 developers finally responded, 78.5% of which completed the survey. The only selection criterion for participants was some proficiency in an object-oriented programming language (e.g. C#, Java, Python). The authors found that approximately 50% of the respondents never (i.e. in less than 10% of the development activities) and about 70% rarely (in less than 25% of the development activities) use any diagrammatic design models. In contrast, only approximately 11% employ models in more than 75% of their development activities. In terms of demographics, this minority of software developers has received a higher level of formal education. When using a formal diagrammatic notation, the UML or a UML-inspired notation was prominently mentioned by the respondents. This large-scale survey, while being widely unspecific of the concrete type of UML usage (e.g. diagram types), just stresses that there is only very little evidence on the use of UML-based DSMLs in software-development practice and on their design details. Our study cannot fully close this particular gap given a different intention (extraction of generic design rationale) and its data sources (i.e. mainly scientific publications by the DSML developers). However, we provide a data set of DSML projects as a starting point for follow-up work on DSML usage in industry practice.

A second group of related empirical studies also documented important generic design rationale which was included in our decision-record catalog. Below, we iterate over these secondary studies to stress how we incorporated their findings on UML-specific design rationale into our catalog.

In a follow-up and partly replicated review of Pardillo [64] focusing on the data warehouse domain only, Pardillo and Cachero [15] make the case of the (unwanted) coupling of abstract-syntax and concrete-syntax decisions of a DSML when adopting UML profiles. This issue has been included into our revised decision-record catalog via option associations between D2 and D4.

The SLR of Budgen et al. [81] aims to collect and to reflect on empirical research on the UML, involving measurement, comprehension, quality, and maintenance of UML models. The review yielded 49 papers published up to 2008. The overview of papers on UML comprehension provides an important auxiliary source for our decision records on concrete-syntax design (D4). The included papers cover findings on the comprehensibility of stereotyped and non-stereotyped models for different audiences, as well as the cognitive effectiveness of different diagram types (sequence diagrams). In addition, the authors recorded the different research methods adopted. These methods included mainly small-scale controlled experiments in laboratory settings. Budgen et al. [81] also found two empirical studies on “notation extensions” in the UML, however, the respective paper category was deliberately excluded from the analysis because the authors were only concerned with contributions using standard

UML. To this end, we complement the SLR of Budgen et al. [81] by organizing papers on UML extensions explicitly, including notation (concrete-syntax) extensions. At the same time, our SLR is a primary study on mining UML-specific design rationale.

Staron and Wohlin [16] report on a confirmatory case study to test six alleged advantages and disadvantages of UML extension techniques (profiles and metamodel extensions) based on an industrial case: the UML-driven modeling toolkit Tau G2 by Telelogic. The authors triangulated their observations from data sources such as profile definitions and metamodel extensions implemented using Tau G2, requirements documents of Telelogic’s customers, a fault database, and interviews to reflect on decision-making processes. A key observation was that software products based on metamodel extensions exhibit more faults in terms of unfulfilled or violated customer requirements. In addition, metamodel extensions were found to incur substantially more development effort (three to four weeks working time) than profiles (one week). We incorporated these evidence-based findings as arguments (drivers) into decision record D2. The observations of Staron and Wohlin [16] rest upon metamodel-based extensions to activity diagrams as well as upon profile-based extensions to deployment, interaction overview, and component diagrams.

Staron et al. [120] conducted a series of four experiments on the understandability of stereotyped UML class models, all replicating a single design but varying in terms of control levels and details of the treatment (e.g. order of artifact presentation). The experiments were set in an academic environment (68 student subjects) and in an industrial environment (four professional subjects). The replicated design involved twelve tasks to be solved by the subjects. The tasks comprised counting specific model elements and identifying definition defects in stereotyped and non-stereotyped models. Non-stereotyped models contained the domain semantics of the class model as unstructured UML comments attached to model elements. The authors recorded three data points per subject: total and per-task solving time as well as the number of correctly solved tasks. In this way, the authors found that the subjects solved more tasks correctly based on stereotyped models. In addition, there was a decrease in the total and relative task-solving times for stereotyped class models. We refer to the findings from Staron et al. [120] as part of a decision driver on the understandability of concrete-syntax options (D4).

8. Conclusion

Existing approaches towards systematizing DSML development [22, 49] put forth a development-process perspective, treating DSML development as a complex flow of exploratory and iterative development activities. Key activities are language-model definition, constraint specification, concrete-syntax design, and platform integration [22]. A process perspective has immediate benefits. It establishes a shared process vocabulary (e.g., of different development

artifacts) and represents common DSML development styles via different flows between development activities (e.g., language-model-driven or mockup-driven DSML development). On the flip side, a process perspective alone abstracts from DSML development as a result of continuous decision-making by stakeholders (e.g., domain engineers, modelers, language engineers, and software architects). Such decision making is always situated in a specific context shaped by domain requirements and architecture requirements (e.g., a certain modeling technology stack such as the UML).

Therefore, our work proposes a complementary *decision-centric perspective* on DSML development. For the scope of UML-based DSMLs, our focus is on providing *tailorable* documentation [25] of *generic* design decisions [27]. This paper puts emphasis on the scientific approach: A three-year, multi-method empirical study based on a systematic literature review (SLR), a content analysis, and a frequent-item-set analysis was used to collect and to validate generic, reusable design decisions. Key findings are:

- The study revealed 24 design-decision options in 80 UML-based DSMLs reported between 2005 and 2012, covering design aspects from language-model implementation in the UML to platform integration (the complete list of DSMLs is provided in the Appendix).
- Language-model piggybacking [1] using UML profiles is the most frequently adopted DSML implementation technique (by 78% of the DSMLs).
- Approximately one third of the documented decision options (9/24) can be combined to characterize more than 60% of the reviewed DSMLs (49/80).
- Based on these 9 critical decision options, 7 different *prototype* DSML designs representing different variants of UML piggybacking and UML extension [1] were identified.
- The most frequently tailored UML diagram types are class diagrams ($\approx 63\%$ of the DSMLs) and activity diagrams (25%), respectively.
- The publication activity on UML-based DSMLs including premier journal venues (e.g. SoSyM, IST) and conference venues (e.g. MoDELS, ECMFA) is on a downward trend.

These findings were incorporated into an empirically validated and publicly available collection of reusable design decisions on UML-based DSMLs. This *decision-record catalog* [33] lists 27 *decision options* and their interdependencies for six different design-decision concerns (e.g. language-model formalization, concrete-syntax definition). We used the data on frequency as well as on commonalities and differences in DSML designs to highlight prototype designs in the catalog. This way, we provide guidance for DSML engineers and enable them to quickly find examples of decision-option combinations that were successfully applied in other DSML projects. Such guidance is particularly important because, in total, hundreds of millions of possible decision-option sets can be derived from our decision-record catalog. To the best of our knowledge, our work is the first attempt to document

design rationale on UML-based DSML development based on empirical evidence collected at a large scale.

The decision-record catalog [33] and the detailed SLR protocol [41] are publicly available. Important by-products of this research project were multiple UML-based DSMLs (see, e.g., SOFServices), techniques for integrating (otherwise independent) DSMLs (see, e.g., [95]), testing techniques for DSML artifacts (see, e.g., [121, 122]), and empirical studies to evaluate the resulting research artifacts (e.g. a controlled experiment on testing notations [123]).

In future work, we will continue to extend and to evaluate the decision-record catalog by incorporating additional third-party DSMLs. We will replicate our SLR study for publication years starting with 2013. This way, we will also seek the missing evidence for decision options deemed candidates after our SLR study (see Section 6). In parallel, we plan to add design decisions on *DSML tooling support* as a seventh decision record (D7) to cover DSML editors and generators. Furthermore, to overcome certain study limitations (see Section 5.4), we will try to engage outside DSML experts, including the authors of the 80 DSMLs identified during our SLR study. We will employ direct inquisitive (e.g. interviews) and observational research techniques (e.g. participant observations) to collect more qualitative evidence, such as on time ordering in design-decision making.

Acknowledgments

This work has partly been funded by the Austrian Research Promotion Agency (FFG) of the Austrian Federal Ministry for Transport, Innovation and Technology (BMVIT) through the Competence Centers for Excellent Technologies (COMET K1) initiative and the FIT-IT program.

References

- [1] D. Spinellis, Notable Design Patterns for Domain-specific Languages, *J. Syst. Softw.* 56 (1) (2001) 91–99.
- [2] M. Mernik, J. Heering, A. Sloane, When and How to Develop Domain-specific Languages, *ACM Comput. Surv.* 37 (4) (2005) 316–344.
- [3] S. Kelly, J. Tolvanen, *Domain-Specific Modeling: Enabling Full Code Generation*, John Wiley & Sons, 2008.
- [4] U. Zdun, M. Strembeck, Reusable Architectural Decisions for DSL Design: Foundational Decisions in DSL Development, in: *Proc. 14th Europ. Conf. Patt. Lang. Prog.*, ACM, 1–37, 2009.
- [5] B. Selic, The Pragmatics of Model-driven Development, *IEEE Softw.* 20 (5) (2003) 19–25.
- [6] S. Sendall, W. Kozaczynski, Model Transformation: The Heart and Soul of Model-driven Software Development, *IEEE Softw.* 20 (5) (2003) 42–45.
- [7] T. Mens, P. v. Gorp, A Taxonomy of Model Transformation, *Electron. Notes Theor. Comput. Sci.* 152 (2006) 125–142.
- [8] B. Langlois, C.-E. Jitka, E. Jouenne, DSL Classification, in: *Proc. 7th OOPSLA Worksh. Domain-Specific Model.*, 2007.
- [9] T. Stahl, M. Völter, *Model-Driven Software Development: Technology, Engineering, Management*, John Wiley & Sons, 2006.
- [10] Object Management Group, *OMG Unified Modeling Language (OMG UML)*, Available at: <http://www.omg.org/spec/UML>, version 2.5, formal/2015-03-01, 2015.
- [11] Object Management Group, *OMG Meta Object Facility (MOF) Core Specification*, Available at: <http://www.omg.org/spec/MOF>, version 2.5, formal/2015-06-05, 2015.

- [12] J. Hutchinson, J. Whittle, M. Rouncefield, Model-driven engineering practices in industry: Social, organizational and managerial factors that lead to success or failure, *Sci. Comput. Program.* 89, Part B (2014) 144–161.
- [13] L. Nascimento, D. L. Viana, P. A. M. S. Neto, D. A. O. Martins, V. C. Garcia, S. R. L. Meira, A Systematic Mapping Study on Domain-Specific Languages, in: *Proc. 7th Int. Conf. Softw. Eng. Adv.*, 179–187, 2012.
- [14] J. Hutchinson, J. Whittle, M. Rouncefield, S. Kristoffersen, Empirical Assessment of MDE in Industry, in: *Proc. 33rd Int. Conf. Softw. Eng.*, ACM, 471–480, 2011.
- [15] J. Pardo, C. Cachero, Domain-specific language modelling with UML profiles by decoupling abstract and concrete syntaxes, *J. Syst. Softw.* 83 (12) (2010) 2591–2606.
- [16] M. Staron, C. Wohlin, An Industrial Case Study on the Choice Between Language Customization Mechanisms, in: *Proc. 7th Int. Conf. Product-Focused Softw. Process Improv.*, vol. 4034 of *LNCS*, Springer, 177–191, 2006.
- [17] F. Mallet, C. Andre, On the Semantics of UML/MARTE Clock Constraints, in: *Proc. 12th IEEE Int. Sym. Object/Compon./Service-Oriented Real-Time Distrib. Comput.*, IEEE, 305–312, 2009.
- [18] B. Selic, On the Semantic Foundations of Standard UML 2.0, in: *Formal Methods for the Design of Real-Time Syst.*, vol. 3185 of *LNCS*, Springer, 181–199, 2004.
- [19] M. Broy, M. Cengarle, UML Formal Semantics: Lessons Learned, *Softw. Syst. Model.* 10 (4) (2011) 441–446.
- [20] Object Management Group, *OMG Systems Modeling Language (OMG SysML)*, Available at: <http://www.omg.org/spec/SysML>, version 1.3, formal/2012-06-01, 2012.
- [21] Object Management Group, *Service oriented architecture Modeling Language (SoaML) Specification*, Available at: <http://www.omg.org/spec/SoaML>, version 1.0.1, formal/2012-05-10, 2012.
- [22] M. Strembeck, U. Zdun, An Approach for the Systematic Development of Domain-Specific Languages, *Softw. Pract. Exper.* 39 (15) (2009) 1253–1292.
- [23] J. E. Burge, J. M. Carroll, R. McCall, I. Mistrik, *Rationale-Based Software Engineering*, Springer, 2008.
- [24] A. H. Dutoit, R. McCall, I. Mistrik, B. Paech, *Rationale Management in Software Engineering: Concepts and Techniques*, in: *Rationale Manag. in Softw. Eng.*, chap. 1, Springer, 1–48, 2006.
- [25] D. Falessi, L. C. Briand, G. Cantone, R. Capilla, P. Kruchten, The Value of Design Rationale Information, *ACM Trans. Softw. Eng. Methodol.* 22 (3) (2013) 21:1–21:32.
- [26] J. Horner, M. Atwood, *Effective Design Rationale: Understanding the Barriers*, in: *Rationale Manag. in Softw. Eng.*, chap. 3, Springer, 73–90, 2006.
- [27] N. Harrison, P. Avgeriou, U. Zdun, Using Patterns to Capture Architectural Decisions, *IEEE Softw.* 24 (4) (2007) 38–45.
- [28] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, M. Stal, *Pattern-Oriented Software Architecture: A System of Patterns*, John Wiley & Sons, 2000.
- [29] M. Völter, M. Kircher, U. Zdun, *Remoting Patterns: Foundations of Enterprise, Internet and Realtime Distributed Object Middleware*, John Wiley & Sons, 2005.
- [30] G. Hohpe, B. Woolf, *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*, Addison-Wesley, 2004.
- [31] J. Coplien, *Software Patterns*, SIGS Manag. Briefings., SIGS Books & Multimedia, 1996.
- [32] F. Buschmann, K. Henney, D. C. Schmidt, *Pattern-oriented Software Architecture – On Patterns and Pattern Languages*, John Wiley & Sons, 2007.
- [33] B. Hoisl, S. Sobernig, M. Strembeck, A Catalog of Reusable Design Decisions for Developing UML/MOF-based Domain-specific Modeling Languages, Available at: <http://epub.wu.ac.at/4466/>, Tech. Rep. 2014/03, WU Vienna, 2014.
- [34] E. Evans, *Domain-driven Design: Tackling Complexity in the Heart of Software*, Addison-Wesley, first edn., 2004.
- [35] B. Hoisl, S. Sobernig, S. Schefer-Wenzl, M. Strembeck, A. Baumgrass, Design Decisions for UML and MOF based Domain-specific Language Models: Some Lessons Learned, in: *Proc. 2nd Worksh. Process-based Appr. for Model-Driven Eng.*, 303–314, 2012.
- [36] E. Filtz, *Systematic Literature Review and Evaluation of DSML-Design Decisions*, Bachelor Thesis, WU Vienna, 2013.
- [37] S. Jalali, C. Wohlin, Systematic literature studies: database searches vs. backward snowballing, in: *Proc. ACM-IEEE Int. Sym. Empir. Softw. Eng. Meas.*, ACM, 29–38, 2012.
- [38] J. Webster, R. T. Watson, Analyzing the Past to Prepare for the Future: Writing a Literature Review, *MIS Quart.* 26 (2) (2002) pp. xiii–xxiii.
- [39] B. Kitchenham, *Procedures for Performing Systematic Reviews*, Joint Tech. Rep. (Keele University Tech. Rep., NICTA Tech. Rep.) TR/SE-0401, 0400011T.1, Keele University & National ICT Ltd., 2004.
- [40] H. Zhang, M. A. Babar, P. Tell, Identifying relevant studies in software engineering, *Inform. Softw. Tech.* 53 (6) (2011) 625–637.
- [41] S. Sobernig, B. Hoisl, M. Strembeck, Protocol for a Systematic Literature Review on Design Decisions for UML-based DSMLs, Available at: <http://epub.wu.ac.at/4311/>, Tech. Rep., WU Vienna, 2014/02, 2014.
- [42] J. Saldaña, *The Coding Manual for Qualitative Researchers*, Sage, 2nd edn., 2013.
- [43] M. Schreier, *Qualitative Content Analysis in Practice*, Sage, 1st edn., 2013.
- [44] C. Borgelt, Frequent Item Set Mining, *Wiley Int. Rev. Data Min. and Knowl. Disc.* 2 (6) (2012) 437–456.
- [45] B. Hoisl, S. Sobernig, S. Schefer-Wenzl, M. Strembeck, A. Baumgrass, A Catalog of Reusable Design Decisions for Developing UML- and MOF-based Domain-Specific Modeling Languages, Available at: <http://epub.wu.ac.at/3578/>, Tech. Rep. 2012/01, WU Vienna, 2012.
- [46] C. Atkinson, T. Kühne, A Tour of Language Customization Concepts, *Adv. Comput.* 70 (2007) 105–161.
- [47] E. Jackson, J. Sztipanovits, Formalizing the structural semantics of domain-specific modeling languages, *Softw. Syst. Model.* 8 (4) (2009) 451–478.
- [48] G. Giachetti, B. Marin, O. Pastor, Using UML as a Domain-Specific Modeling Language: A Proposal for Automatic Generation of UML Profiles, in: *Proc. 21st Int. Conf. Adv. Inform. Syst. Eng.*, vol. 5565 of *LNCS*, Springer, 110–124, 2009.
- [49] F. Lagarde, E. Huáscar, F. Terrier, C. André, S. Gérard, Leveraging Patterns on Domain Models to Improve UML Profile Definition, in: *Proc. 11th Int. Conf. Fund. Appr. to Softw. Eng.*, vol. 4961 of *LNCS*, Springer, 116–130, 2008.
- [50] S. Cook, Looking back at UML, *Softw. Syst. Model.* 11 (4) (2012) 471–480.
- [51] J. Dingel, Z. Diskin, A. Zito, Understanding and improving UML package merge, *Softw. Syst. Model.* 7 (4) (2008) 443–467.
- [52] X. Burgués, X. Franch, J. M. Ribó, Improving the accuracy of UML metamodel extensions by introducing induced associations, *Softw. Syst. Model.* 7 (3) (2008) 361–379.
- [53] S. Sen, N. Moha, B. Baudry, J.-M. Jézéquel, Meta-model Pruning, in: *Proc. 12th Int. Conf. Model Driven Eng. Lang. Syst.*, vol. 5795 of *LNCS*, Springer, 32–46, 2009.
- [54] A. Blouin, B. Combemale, B. Baudry, O. Beaudoux, Kompren: modeling and generating model slicers, *Softw. Syst. Model.* 14 (1) (2015) 321–337.
- [55] B. Henderson-Sellers, C. Gonzalez-Perez, Uses and Abuses of the Stereotype Mechanism in UML 1.x and 2.0, in: *Proc. 9th Int. Conf. Model Driven Eng. Lang. Syst.*, vol. 4199 of *LNCS*, Springer, 16–26, 2006.
- [56] G. L. Steele, *Common Lisp: The Language*, Digital Press, 2nd edn., 1990.
- [57] S. Günther, Development of Internal Domain-specific Languages: Design Principles and Design Patterns, in: *Proc. 18th Conf. Patt. Lang. of Prog.*, ACM, 1:1–1:25, 2011.
- [58] T. Clark, P. Sammut, J. Willans, *Applied Metamodelling: A foundation for language-driven development*, Ceteva, 2nd edn., 2008.
- [59] R. F. Paige, D. S. Kolovos, F. A. Polack, A tutorial on metamodelling for grammar researchers, *Sci. Comput. Program.* 96 (4) (2014) 396–416.
- [60] M. Shahin, P. Liang, M. R. Khayyambashi, Architectural design decision: Existing models and tools, in: *Joint Proc. 3rd Europ. Conf. Softw. Archit. and 8th Working IEEE/IFIP Conf. Softw. Archit.*, IEEE, 293–296, 2009.
- [61] P. Mayer, A. Schroeder, N. Koch, MDD4SOA: Model-Driven Service Orchestration, in: *Proc. 12th Int. Conf. Enterp. Distrib. Object Comput.*, IEEE, 203–212, 2008.

- [62] A. MacLean, R. M. Young, V. M. E. Bellotti, T. P. Moran, Questions, Options, and Criteria: Elements of Design Space Analysis, in: *Design Rationale: Concepts, Techniques, and Use*, chap. 3, Lawrence Erlbaum Associates, 53–106, 1996.
- [63] U. Zdun, Systematic pattern selection using pattern language grammars and design space analysis, *Softw. Pract. Exper.* 37 (9) (2007) 983–1016.
- [64] J. Pardillo, A Systematic Review on the Definition of UML Profiles, in: *Proc. 13th Int. Conf. Model Driven Eng. Lang. Syst.*, vol. 6394 of *LNCS*, Springer, 407–422, 2010.
- [65] P. Kruchten, P. Lago, H. van Vliet, Building Up and Reasoning About Architectural Knowledge, in: *Qual. of Softw. Archit.*, vol. 4214 of *LNCS*, Springer, 43–58, 2006.
- [66] M. Hahsler, B. Grün, K. Hornik, arules—A Computational Environment for Mining Association Rules and Frequent Item Sets, *J. Stat. Softw.* 14 (15) (2005) 1–25.
- [67] R. Wieringa, N. Maiden, N. Mead, C. Rolland, Requir. Eng. Paper Classification and Evaluation Criteria: A Proposal and a Discussion, *Requir. Eng.* 11 (1) (2005) 102–107.
- [68] A. Tang, P. Avgeriou, A. Jansen, R. Capilla, M. A. Babar, A comparative study of architecture knowledge management tools, *J. Syst. Softw.* 83 (3) (2010) 352–370.
- [69] R. Paige, J. Ostroff, P. Brooke, Principles for Modeling Language Design, *Inform. Softw. Tech.* 42 (10) (2000) 665–675.
- [70] E. Grant, K. Narayanan, H. Reza, Rigorously Defined Domain Modeling Languages, in: *Proc. 4th OOPSLA Worksh. Domain-Specific Model.*, 2004.
- [71] S. Robert, S. Gérard, F. Terrier, F. Lagarde, A Lightweight Approach for Domain-Specific Modeling Languages Design, in: *Proc. 35th EUROMICRO Conf. Softw. Eng. and Adv. Appl.*, IEEE, 155–161, 2009.
- [72] B. Selic, A Systematic Approach to Domain-Specific Language Design Using UML, in: *Proc. 10th IEEE Int. Sym. Object-Oriented Real-Time Distrib. Comput.*, IEEE, 2–9, 2007.
- [73] K. L. Gwet, *Handbook of Inter-Rater Reliability*, Advanced Analytics, LLC, 3rd edn., 2012.
- [74] B. Kitchenham, P. Brereton, A Systematic Review of Systematic Review Process Research in Software Engineering, *Inform. Softw. Tech.* 55 (12) (2013) 2049–2075.
- [75] Australian Research Council, Excellence in Research for Australia (ERA) Journal List, Available at http://www.arc.gov.au/era/era_2012/era_journal_list.htm. Last accessed: Oct 3, 2014, 2012.
- [76] Scimago Lab, SCImago Journal & Country Rank, Available at <http://www.scimagojr.com/journalrank.php>. Last accessed: Feb 2, 2015, 2013.
- [77] Microsoft Corporation, Microsoft Academic Search journal list, Available at <http://academic.research.microsoft.com/?SearchDomain=2&SubDomain=4&entitytype=4>. Last accessed: Feb 2, 2015, 2013.
- [78] J. L. Fleiss, *Statistical Methods for Rates and Proportions*, John Wiley & Sons, 1981.
- [79] F. Naumann, M. Herschel, *An Introduction to Duplicate Detection*, Synthesis Lectures on Data Management, Morgan & Claypool Publishers, 2010.
- [80] C. Atkinson, T. Kühne, B. Henderson-Sellers, Systematic Stereotype Usage, *Softw. Syst. Model.* 2 (3) (2003) 153–163.
- [81] D. Budgen, A. J. Burn, O. P. Brereton, B. A. Kitchenham, R. Pretorius, Empirical evidence about the UML: A systematic literature review, *Softw. Pract. Exper.* 41 (4) (2011) 363–392.
- [82] J. Bruck, K. Hussey, Customizing UML: Which Technique is Right for You?, Available at: http://www.eclipse.org/modeling/mdt/uml2/docs/articles/Customizing_UML2_Which_Technique_is_Right_For_You/article.html. Last accessed: Sep 9, 2015, IBM, 2008.
- [83] J. Singer, S. E. Sim, T. C. Lethbridge, Software Engineering Data Collection for Field Studies, in: *Guide to Adv. Empir. Softw. Eng.*, Springer, 9–34, 2008.
- [84] M. Turner, B. Kitchenham, D. Budgen, P. Brereton, Lessons Learnt Undertaking a Large-scale Systematic Literature Review, in: *Proc. 12th Int. Conf. Eval. and Ass. in Softw. Eng.*, British Computer Society, 110–118, 2008.
- [85] D. Moody, J. van Hilleberg, Evaluating the Visual Syntax of UML: An Analysis of the Cognitive Effectiveness of the UML Family of Diagrams, in: *Softw. Lang. Eng.*, vol. 5452 of *LNCS*, Springer, 16–34, 2009.
- [86] T. Mayerhofer, P. Langer, M. Wimmer, Towards xMOF: Executable DSMs based on fUML, in: *Proc. 12th Worksh. Domain-Specific Model.*, ACM, 1–6, 2012.
- [87] K. Czarnecki, U. W. Eisenecker, *Generative Programming: Methods, Tools, and Applications*, Addison-Wesley, 2000.
- [88] C. Atkinson, T. Kühne, Profiles in a strict metamodeling framework, *Sci. Comput. Program.* 44 (1) (2002) 5–22.
- [89] F. Lagarde, H. Espinoza, F. Terrier, S. Gérard, Improving UML Profile Design Practices by Leveraging Conceptual Domain Models, in: *Proc. 22nd IEEE/ACM Int. Conf. Autom. Softw. Eng.*, ACM, 445–448, 2007.
- [90] D. Torre, Y. Labiche, M. Genero, UML Consistency Rules: A Systematic Mapping Study, in: *Proc. 18th Int. Conf. Evaluation and Assessment in Softw. Eng.*, ACM, 6:1–6:10, 2014.
- [91] P. Langer, K. Wieland, M. Wimmer, J. Cabot, EMF Profiles: A Lightweight Extension Approach for EMF Models, *J. Object Technol.* 11 (1) (2012) 1–29.
- [92] C. Atkinson, T. Kühne, Processes and Products in a multi-level meta-modeling architecture, *Int. J. Softw. Eng. Know.* 11 (6) (2001) 761–783.
- [93] B. Ogunyomi, L. M. Rose, D. S. Kolovos, Property Access Traces for Source Incremental Model-to-Text Transformation, in: *Proc. 11th Europ. Conf. Model. Found. Appl.*, vol. 9153 of *LNCS*, Springer, 187–202, 2015.
- [94] Generative Software, Umfrage zu Verbreitung und Einsatz modellgetriebener Softwareentwicklung, Survey Rep., Generative Software GmbH and FZI Forschungszentrum Informatik, URL <http://www.mdsd-umfrage.de/mdsd-report-2010.pdf>, 2010.
- [95] B. Hoisl, S. Sobernig, M. Strembeck, Higher-Order Rewriting of Model-to-Text Templates for Integrating Domain-specific Modeling Languages, in: *Proc. 1st Int. Conf. Model-Driven Eng. and Softw. Development*, SciTePress, 49–61, 2013.
- [96] S. Apel, D. S. Batory, C. Kästner, G. Saake, *Feature-Oriented Software Product Lines: Concepts and Implementation*, Springer, 2013.
- [97] U. Heesch, P. Avgeriou, R. Hilliard, A documentation framework for architecture decisions, *J. Syst. Softw.* 85 (4) (2012) 795–820.
- [98] M. Völter, *DSL Engineering – Designing, Implementing, and Using Domain-Specific Languages*, Amazon Distribution, 2013.
- [99] D. Roberts, R. Johnson, Evolve Frameworks into Domain-Specific Languages, in: *Proc. 3rd Int. Conf. Patt. Lang. of Prog.*, 1996.
- [100] D. Roberts, R. Johnson, Patterns for Evolving Frameworks, in: *Patt. Lang. of Program Design 3*, Addison-Wesley, 471–486, 1997.
- [101] C. Schäfer, T. Kuhn, M. Trapp, A Pattern-based Approach to DSL Development, in: *Worksh. Proc. of Conf. Syst., Prog., and Appl.: Softw. for Hum.*, ACM, 39–46, 2011.
- [102] J. Luoma, S. Kelly, J. Tolvanen, Defining Domain-Specific Modeling Languages: Collected Experiences, in: *Proc. 4th OOPSLA Worksh. Domain-Specific Model.*, 2004.
- [103] D. Wile, Lessons learned from real DSL experiments, *Sci. Comput. Program.* 51 (3) (2004) 265–290.
- [104] S. Kelly, R. Pohjonen, Worst Practices for Domain-Specific Modeling, *IEEE Softw.* 26 (4) (2009) 22–29.
- [105] G. Karsai, H. K. C. Pinkernell, B. Rumpe, M. Schindler, S. Völkel, Design Guidelines for Domain Specific Languages, in: *Proc. 9th OOPSLA Worksh. Domain-Specific Model.*, 7–13, 2009.
- [106] S. Günther, M. Haupt, M. Splieth, Agile Engineering of Internal Domain-Specific Languages with Dynamic Programming Languages, in: *Proc. 5th Int. Conf. Softw. Eng. Adv.*, IEEE, 162–168, 2010.
- [107] S. Günther, T. Cleenewerck, Design Principles for Internal Domain-specific Languages: A Pattern Catalog Illustrated by Ruby, in: *Proc. 17th Conf. Patt. Lang. of Prog.*, ACM, 3:1–3:35, 2010.
- [108] N. Allen, C. Shaffer, L. Watson, Building Modeling Tools that Support Verification, Validation, and Testing for the Domain Expert, in: *Proc. 37th Winter Simul. Conf.*, IEEE, 419–426, 2005.
- [109] S. Thibault, R. Marlet, C. Consel, Domain-Specific Languages: From Design to Implementation Application to Video Device Drivers Generation, *IEEE Trans. Softw. Eng.* 25 (3) (1999) 363–377.
- [110] T. Cleenewerck, Component-based DSL Development, in: *Proc. 2nd Int. Conf. Genera. Prog. and Compon. Eng.*, Springer, 245–264, 2003.
- [111] P. Landin, The Next 700 Programming Languages, *Commun. ACM* 9 (3) (1966) 157–166.
- [112] I. Weisemöller, A. Schür, A Comparison of Standard Compliant Ways to Define Domain Specific Languages, in: *Worksh. Proc. 10th Int. Conf.*

Model Driven Eng. Lang. Syst., vol. 5002 of *LNCS*, Springer, 47–58, 2008.

- [113] J.-P. Tolvanen, S. Kelly, MetaEdit+: Defining and using integrated domain-specific modeling languages, in: Proc. 24th ACM SIGPLAN Conf. Companion on Object Oriented Prog. Syst. Lang. Appl., ACM, 819–820, 2009.
- [114] U. Zdun, A DSL Toolkit for Deferring Architectural Decisions in DSL-based Software Design, *Inform. Softw. Tech.* 52 (9) (2010) 733–748.
- [115] B. Dobing, J. Parsons, How UML is Used, *Commun. ACM* 49 (5) (2006) 109–113.
- [116] V. Garousi, Classification and trend analysis of UML books (1997–2009), *Softw. Syst. Model.* 11 (2) (2012) 273–285.
- [117] A. Nugroho, M. R. Chaudron, A Survey into the Rigor of UML Use and Its Perceived Impact on Quality and Productivity, in: Proc. 2nd Int. Sym. Empir. Softw. Eng. and Meas., ACM, 90–99, 2008.
- [118] J. Whittle, J. Hutchinson, M. Rouncefield, The State of Practice in Model-Driven Engineering, *IEEE Softw.* 31 (3) (2014) 79–85.
- [119] T. Gorschek, E. Tempero, L. Angelis, On the use of software design models in software development practice: An empirical investigation, *J. Syst. Softw.* 95 (2014) 176–193.
- [120] M. Staron, L. Kuzniarz, C. Wohlin, Empirical assessment of using stereotypes to improve comprehension of UML models: A set of experiments, *J. Syst. Softw.* 79 (5) (2006) 727–742.
- [121] S. Sobernig, B. Hoisl, M. Strembeck, Requirements-driven Testing of Domain-specific Core Language Models using Scenarios, in: Proc. 13th Int. Conf. Qual. Softw., IEEE, 163–172, 2013.
- [122] B. Hoisl, S. Sobernig, M. Strembeck, Natural-language Scenario Descriptions for Testing Core Language Models of Domain-Specific Languages, in: Proc. 2nd Int. Conf. Model-Driven Eng. Softw. Dev., SciTePress, 356–367, 2014.
- [123] B. Hoisl, S. Sobernig, M. Strembeck, Comparing Three Notations for Defining Scenario-based Model Tests: A Controlled Experiment, in: Proc. 9th Int. Conf. Qual. of Inform. and Commun. Tech., IEEE, 95–104, 2014.

Appendix

Table 13 lists all 90 DSML projects from which design decisions have been extracted—ten are our own developments (#1–#10), the rest has been retrieved via the SLR (#11–#90). The first column consecutively numbers and names each DSML (a name either specified by the authors themselves or—when no explicit name was mentioned—one chosen by us) and references corresponding publication(s). In the second column, the DSML application domain(s) are encoded according to the 2012 ACM Computing Classification System (CCS).²³ We have extracted the UML diagram type(s) tailored by a DSML as classified by the UML superstructure itself (shown in the third column of Table 13). The last column lists the decision-option set representing a DSML’s design as encoded according to our catalog.

Table 13: Application domains, tailored diagram types, and decision-option set for each of the 90 DSML projects.

DSML	Application domain(s)	Diagram type(s)	Option set
#1 ConcernActivities * ²⁴ [A 1]	Access control, Software design engineering	Activity	{1.1,2.2,2.3,3.1,3.4,4.1,4.2,4.6,5.5,6.6}
#2 BusinessActivities [A 2]	Access control, Business process modeling, Software security engineering	Activity, Class	{1.1,1.2,1.4,2.3,3.1,3.4,4.2,5.5,6.1,6.4}
#3 UML-PD * [A 3, 4]	Access control, Business process modeling, Software security engineering	Activity, Class	{1.1,1.2,1.4,2.2,2.3,3.1,3.4,4.1,4.2,4.6,5.1,5.3,6.6}
#4 UML-DEL [A 4, 5]	Access control, Business process modeling, Software security engineering	Class	{1.1,1.2,1.4,2.3,3.1,3.4,4.2,5.5,6.6}
#5 SOF [A 6]	Business process modeling, Software security engineering	Activity	{1.1,2.3,3.1,3.4,4.2,5.5,6.6}
#6 UML-PD [A 7]	Access control, Business process modeling, Software security engineering	Activity, Class	{1.1,2.3,3.1,3.4,4.7,5.5,6.6}
#7 SOFServices * [A 8, 9]	Business process modeling, Service-oriented architectures, Software security engineering, Web services	Activity, CompositeStructure	{1.1,1.2,1.4,2.2,2.3,3.1,3.3,3.4,4.1,4.6,5.5,6.1,6.3}
#8 UML-CC [A 10]	Access control, Business process modeling, Software security engineering	Class	{1.1,1.2,1.4,2.3,3.1,3.4,4.2,5.5,6.6}
#9 SecurityAudit * [A 11]	Publish-subscribe / event-based architectures, Software security engineering	* ²⁵	{1.1,2.2,2.3,3.1,3.4,4.1,4.3,4.5,4.6,5.5,6.2}
#10 MTD * [A 14]	Object oriented languages, Software architectures	Activity, Class, Object, StateMachine	{1.1,2.2,2.3,3.1,3.4,4.1,4.2,4.6,5.5,6.6}
#11 ADModel [A 15]	Business process modeling	Activity	{1.1,2.3,3.5,4.7,5.5,6.3}
#12 AspectSM * [A 16]	Robustness, Software development techniques, Software testing and debugging	StateMachine	{1.1,1.4,2.2,3.1,4.1,4.6,5.5,6.2}
#13 UML4SPM * [A 17]	Software development process management	Activity, Class	{1.1,2.3,3.5,4.6,5.5,6.6}
#14 MDATC * [A 18]	Reusability, Software development techniques, Software product lines	Activity, Package	{1.1,2.3,3.5,4.6,5.5,6.6}
#15 TLM * [A 19]	Model verification and validation, System on a chip	Class	{1.1,2.2,3.1,4.1,4.6,5.5,6.2}
#16 UPSS * [A 20]	Service-oriented architectures	Class, CompositeStructure	{1.1,1.4,2.2,3.4,4.1,4.6,5.5,6.6}
#17 BIT * [A 21]	Software testing and debugging	Class	{1.1,2.2,3.1,4.1,4.6,5.5,6.1,6.2,6.5}
#18 UML4PF * [A 22, 23]	Design patterns, Model checking, Requirements analysis, Security requirements	Class	{1.1,2.2,3.1,4.1,4.6,5.5,6.6}
#19 UP4WS * [A 24]	Service-oriented architectures, Web services	Class	{1.1,2.2,3.4,4.1,4.6,5.5,6.2}
#20 CB * [A 25]	Reusability, Software development techniques	Class, Component	{1.1,1.4,2.2,3.4,4.1,4.6,5.5,6.1,6.3,6.5}
#21 AbstractSet * [A 26]	Model verification and validation	Class, Package	{1.1,1.4,2.2,3.5,4.1,4.6,5.5,6.6}
#22 C2style * [A 27]	Architecture description languages, Systems analysis and design	Component, Sequence	{1.1,2.2,3.1,3.4,4.1,4.6,5.5,6.6}
#23 MARTE-DAM * [A 28, 29]	Embedded systems, Fault tree analysis, Real-time systems, Software fault tolerance, Transportation	Component, Sequence, StateMachine, UseCase	{1.1,1.4,2.2,3.1,4.1,4.6,5.5,6.3,6.5}

²³See <http://www.acm.org/about/class>; last accessed: Sep 9, 2015.

²⁴The DSML’s option set contains (at least) one of the seven prototype option-sets shown in Table 11.

²⁵The DSML does not tailor a UML diagram type specifically; for example, a stereotype extension of a UML element applicable in all diagram types, such as, Element (see, e.g., [A 11, 12]) or Constraint (see, e.g., [A 13]).

DSML	Application domain(s)	Diagram type(s)	Option set
#24 UMM-Local-Choreographies * [A 30]	Business process modeling, Orchestration languages	Activity	{1.1, 2.2, 3.1, 4.1, 4.6, 5.5, 6.6}
#25 RichService * [A 31]	Service-oriented architectures, Web services	Class, Component, StateMachine	{1.1, 1.4, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#26 UML-PMS * [A 32]	Performance, Ubiquitous and mobile computing	Activity	{1.1, 2.2, 3.4, 4.1, 4.6, 5.5, 6.6}
#27 SOA * [A 33]	Service-oriented architectures	Class, Component, Deployment	{1.1, 1.4, 2.2, 3.1, 4.1, 4.2, 4.6, 5.5, 6.6}
#28 SWS * [A 34]	Semantic web description languages, Web services	Activity	{1.1, 1.4, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#29 eSPEM * [A 35]	Software development process management	Activity, StateMachine	{1.1, 2.3, 2.4, 3.5, 4.6, 5.5, 6.6}
#30 RCSD * [A 36]	Transportation	Class, Object	{1.1, 2.2, 3.1, 3.4, 4.1, 4.3, 4.6, 5.2, 6.6}
#31 UML-SOA-Sec [A 37]	Business process modeling, Security requirements, Service-oriented architectures, Web services	Activity	{1.1, 2.2, 3.5, 4.1, 4.2, 4.6, 5.5, 6.6}
#32 UML2Alloy * [A 38]	Model verification and validation	Class, Package	{1.1, 1.4, 2.2, 3.1, 3.4, 4.1, 4.6, 5.5, 6.1, 6.2, 6.5}
#33 ExSAM * [A 39]	Avionics, Embedded systems, Engineering	CompositeStructure	{1.1, 1.4, 2.2, 3.4, 4.1, 4.6, 5.5, 6.6}
#34 UACL * [A 40]	Availability, Telecommunications	Class, Component	{1.1, 1.4, 2.2, 3.1, 3.4, 4.1, 4.6, 5.5, 6.6}
#35 SECTET [A 41]	Service-oriented architectures, Software security engineering, Web services	Class	{1.1, 2.1, 3.5, 4.7, 5.5, 6.2, 6.5}
#36 UML4SOA * [A 61, 43]	Service-oriented architectures	Activity, Class, Component	{1.1, 2.2, 2.3, 3.1, 4.1, 4.2, 4.6, 5.5, 6.1, 6.3, 6.5}
#37 SafeUML * [A 44]	Avionics, Software safety	Class, Package	{1.1, 1.4, 2.2, 3.1, 3.4, 4.1, 4.6, 5.5, 6.6}
#38 IStarDW * [A 45]	Data warehouses, Security requirements	Class, Package	{1.1, 2.2, 3.1, 3.4, 4.1, 4.2, 4.6, 5.5, 6.5}
#39 TestOracle * [A 46]	Software testing and debugging	StateMachine	{1.1, 2.2, 3.5, 4.1, 4.3, 4.6, 5.5, 6.2}
#40 MOCAS [A 47]	Model checking, Model verification and validation	Object	{1.1, 2.3, 3.1, 3.4, 4.7, 5.5, 6.4}
#41 CCFG [A 48]	Model verification and validation	Activity	{1.1, 2.3, 3.5, 4.2, 5.5, 6.6}
#42 TimeSeriesAnalysis * [A 49]	Data mining, Data warehouses	Class, Object	{1.1, 2.2, 3.1, 4.1, 4.2, 4.6, 5.5, 6.6}
#43 ADOM-UML [A 50]	Model verification and validation, Requirements analysis, Software design engineering	*	{1.1, 1.2, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#44 Predefined-Constraints * [A 13]	Model checking	*	{1.1, 2.2, 3.4, 4.1, 4.6, 5.5, 6.6}
#45 TAM-PM * [A 51]	Graphical user interfaces, Web interfaces	Activity, Class	{1.1, 1.4, 2.2, 3.1, 4.1, 4.6, 5.5, 6.2}
#46 SPEM4MDE [A 52]	Software development process management	Activity, StateMachine	{1.1, 2.3, 3.1, 3.4, 4.2, 5.5, 6.5}
#47 CSSL [A 53]	Collaborative and social computing	Class, StateMachine	{1.1, 2.3, 3.1, 3.4, 4.5, 5.5, 6.6}
#48 SystemC [A 54]	Embedded systems, System on a chip	CompositeStructure, StateMachine	{1.1, 2.2, 3.5, 4.1, 4.2, 4.6, 5.5, 6.6}
#49 UML2Ext * [A 55]	Requirements analysis, Software product lines	UseCase	{1.1, 2.3, 3.5, 4.6, 5.5, 6.6}
#50 HM ³ * [A 56]	Hypertext languages	Class, UseCase	{1.1, 2.2, 2.3, 3.1, 3.4, 4.1, 4.6, 5.5, 6.6}
#51 WCAAUML * [A 57]	Web applications, Web interfaces	Class, Deployment, Package	{1.1, 2.2, 3.5, 4.1, 4.6, 5.5, 6.2}
#52 IEC61508 * [A 58, 59]	Model verification and validation, Safety critical systems	Class, Package	{1.1, 1.4, 2.2, 3.1, 3.4, 4.1, 4.6, 5.5, 6.6}
#53 UCDCM * [A 60]	Use cases	UseCase	{1.1, 2.3, 3.1, 3.4, 4.4, 4.6, 5.3, 6.6}
#54 SPArch * [A 61]	Software architectures, Software development process management	Class, Component, Package	{1.1, 1.4, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#55 MoDePeMART * [A 62]	Measurement, Metrics, Software performance	Class, StateMachine	{1.1, 1.4, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#56 UPCC [A 63]	Enterprise data management, Service-oriented architectures, Web services	Class	{1.1, 2.1, 3.5, 4.7, 5.5, 6.6}

DSML	Application domain(s)	Diagram type(s)	Option set
#57 SELinux [A 64]	Access control, Operating systems security, Security requirements	Class	{1.1, 1.3, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#58 UML-GUI [A 65]	Graphical user interfaces	Class, Component	{1.1, 1.2, 2.1, 3.5, 4.7, 5.5, 6.3}
#59 SHP * [A 66]	Software security engineering	Class, Package	{1.1, 2.2, 3.1, 3.4, 4.1, 4.6, 5.5, 6.6}
#60 SMF * [A 67]	Fault tree analysis, Safety critical systems, Software safety	Class, Component, UseCase	{1.1, 1.3, 1.4, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#61 DMM/UCMM * [A 68]	Graphical user interfaces	Class, UseCase	{1.1, 1.4, 2.3, 2.4, 3.5, 4.6, 5.5, 6.6}
#62 CUP 2.0 * [A 69]	Graphical user interfaces	Activity, Class, Package	{1.1, 2.2, 3.4, 4.1, 4.2, 4.6, 5.5, 6.2}
#63 REMP * [A 70]	Embedded systems, Real-time systems, Software testing and debugging	Class, StateMachine	{1.1, 2.2, 3.1, 4.1, 4.6, 5.5, 6.6}
#64 DPL * [A 71]	Web services	Activity	{1.1, 2.2, 3.5, 4.1, 4.6, 5.5, 6.2}
#65 WebRE * [A 72]	Requirements analysis, Web applications	Activity, UseCase	{1.1, 1.4, 2.2, 3.5, 4.1, 4.2, 4.6, 5.5, 6.6}
#66 AOM-AD * [A 73]	Software development techniques	Activity	{1.1, 1.2, 2.2, 3.4, 4.1, 4.6, 5.5, 6.6}
#67 Reliability * [A 74]	Software reliability	InteractionOverview, Sequence	{1.1, 1.4, 2.2, 3.4, 4.1, 4.6, 5.5, 6.2}
#68 UML-AOF * [A 75]	Software development techniques	Class, Package	{1.1, 2.2, 3.4, 4.1, 4.6, 5.5, 6.6}
#69 CompSize [A 76]	Embedded systems, Estimation, Measurement, Metrics	Class, Component	{1.1, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#70 Architectural-Primitives * [A 77]	Design patterns, Software architectures	Component	{1.1, 2.2, 3.1, 3.4, 4.1, 4.6, 5.5, 6.6}
#71 CUP * [A 78]	Error detection and error correction, Model checking	CompositeStructure, Sequence	{1.1, 2.2, 3.1, 4.1, 4.6, 5.5, 6.6}
#72 GWFm-Sec * [A 79]	Orchestration languages, Software security engineering, Web services	Activity	{1.1, 2.2, 2.3, 3.4, 4.1, 4.6, 5.5, 6.6}
#73 SoC * [A 80]	Hardware description languages and compilation, System on a chip	Activity, Class, CompositeStructure, Deployment	{1.1, 2.2, 3.4, 4.1, 4.6, 5.5, 6.6}
#74 UMLtrust * [A 81]	Scenario-based design, Software development techniques, Trust frameworks	Class, Package, UseCase	{1.1, 2.2, 3.4, 4.1, 4.2, 4.6, 5.5, 6.6}
#75 HERM * [A 82]	Database design and models	Class	{1.1, 1.2, 1.4, 2.2, 3.1, 3.4, 4.1, 4.6, 5.5, 6.6}
#76 WebML * [A 83]	Web applications, Web interfaces	Class, Component, CompositeStructure	{1.1, 1.4, 2.2, 3.1, 4.1, 4.2, 4.6, 5.5, 6.6}
#77 ODP [A 84]	Distributed architectures	Class, Component, Object, Sequence	{1.1, 2.2, 3.5, 4.1, 4.6, 5.5, 6.6}
#78 EIS * [A 85]	Enterprise information systems	Activity, Component	{1.1, 2.2, 3.4, 4.1, 4.2, 4.6, 5.5, 6.6}
#79 SPTEExt * [A 86]	Embedded systems, Real-time systems	Activity	{1.1, 1.2, 1.3, 2.2, 3.4, 4.1, 4.6, 5.5, 6.6}
#80 CAV * [A 87]	Software architectures, Software evolution	Class	{1.1, 2.2, 3.4, 4.1, 4.6, 5.5, 6.6}
#81 SOA-NF * [A 88]	Service-oriented architectures	CompositeStructure	{1.1, 2.2, 3.4, 4.1, 4.6, 5.5, 6.2}
#82 SECRDW [A 89]	Data warehouses, Security requirements	Class, Package	{1.1, 2.3, 3.5, 4.7, 5.5, 6.6}
#83 SECDW * [A 90]	Data warehouses, Security requirements	Class, Object	{1.1, 2.2, 3.1, 3.4, 4.1, 4.2, 4.6, 5.5, 6.6}
#84 EM [A 91]	Electronic commerce, Web applications	Class, StateMachine, UseCase	{1.1, 2.1, 3.1, 3.4, 4.7, 5.1, 5.3, 6.6}
#85 WS-CM * [A 92]	Web applications, Web services	Class, StateMachine	{1.1, 2.2, 3.5, 4.1, 4.6, 5.5, 6.2}
#86 aspectJ * [A 93]	Software development techniques	Class, Package	{1.1, 2.2, 3.1, 3.4, 4.1, 4.6, 5.5, 6.2}
#87 ContextUML * [A 94]	Service-oriented architectures, Web services	Class	{1.1, 2.2, 3.5, 4.1, 4.6, 5.5, 6.2, 6.3}
#88 DifferenceMM * [A 95]	Software evolution	Class	{1.1, 2.3, 3.5, 4.6, 5.5, 6.6}
#89 Versioning * [A 12]	Software evolution, Version control	*	{1.1, 1.4, 2.2, 3.5, 4.1, 4.6, 5.5, 6.3}
#90 NFA [A 96]	Avionics, Model checking	Class	{1.1, 2.3, 3.5, 4.3, 5.5, 6.5}

DSML Paper Corpus

- [A1] M. Strembeck, U. Zdun, Modeling interdependent concern behavior using extended activity models, *J. Object Technol.* 7 (6) (2008) 143–166.
- [A2] M. Strembeck, J. Mendling, Modeling process-related RBAC models with extended UML activity models, *Inform. Softw. Tech.* 53 (5) (2011) 456–483.
- [A3] S. Schefer-Wenzl, M. Strembeck, Modeling process-related duties with extended UML activity and interaction diagrams, *Electron. Commun. EASST* 37.
- [A4] S. Schefer-Wenzl, M. Strembeck, An approach for consistent delegation in process-aware information systems, in: *Proc. 15th Int. Conf. Bus. Inform. Syst.*, Vol. 117 of LNBI, Springer, 2012, pp. 60–71.
- [A5] S. Schefer, M. Strembeck, Modeling support for delegating roles, tasks, and duties in a process-related rbac context, in: *Proc. Int. Worksh. Inform. Syst. Secur. Eng.*, Vol. 83 of LNBI, Springer, 2011, pp. 660–667.
- [A6] B. Hoisl, M. Strembeck, Modeling support for confidentiality and integrity of object flows in activity models, in: *Proc. 14th Int. Conf. Bus. Inform. Syst.*, Vol. 97 of LNBI, Springer, 2011, pp. 278–289.
- [A7] S. Schefer, Consistency checks for duties in extended UML2 activity models, in: *Proc. Int. Worksh. Secur. Aspects of Process-aware Inform. Syst.*, IEEE, 2011, pp. 680–685.
- [A8] B. Hoisl, S. Sobernig, M. Strembeck, Modeling and enforcing secure object flows in process-driven SOAs: An integrated model-driven approach, *Softw. Syst. Model.* 13 (2) (2014) 513–548.
- [A9] B. Hoisl, S. Sobernig, Integrity and confidentiality annotations for service interfaces in SoaML models, in: *Proc. Int. Worksh. Secur. Aspects of Process-aware Inform. Syst.*, IEEE, 2011, pp. 673–679.
- [A10] S. Schefer-Wenzl, M. Strembeck, Modeling context-aware RBAC models for business processes in ubiquitous computing environments, in: *Proc. 3rd Int. Conf. Mobile, Ubiquitous and Intell. Comput.*, IEEE, 2012, pp. 126–131.
- [A11] B. Hoisl, M. Strembeck, A UML extension for the model-driven specification of audit rules, in: *Proc. 2nd Int. Worksh. Inform. Syst. Secur. Eng.*, Vol. 112 of LNBI, Springer, 2012, pp. 16–30.
- [A12] P. Brosch, M. Seidl, M. Wimmer, G. Kappel, Conflict visualization for evolving UML models, *J. Object Technol.* 11 (3) (2012) 2:1–30.
- [A13] D. Costal, C. Gómez, A. Queral, R. Raventós, E. Teniente, Improving the definition of general constraints in UML, *Softw. Syst. Model.* 7 (4) (2008) 469–486.
- [A14] U. Zdun, M. Strembeck, Modeling composition in dynamic programming environments with model transformations, in: *Proc. 5th Int. Sym. Softw. Compos.*, Vol. 4089 of LNCS, Springer, 2006, pp. 178–193.
- [A15] S. Brahe, K. Østerbye, Business process modeling: Defining domain specific modeling languages by use of UML profiles, in: *Model Driven Archit. – Found. and Appl.*, Vol. 4066 of LNCS, Springer, 2006, pp. 241–255.
- [A16] S. Ali, L. C. Briand, H. Hemmati, Modeling robustness behavior using aspect-oriented modeling to support robustness testing of industrial systems, *Softw. Syst. Model.* 11 (4) (2012) 633–670.
- [A17] R. Bendraou, M.-P. Gervais, X. Blanc, UML4SPM: A UML2.0-based metamodel for software process modelling, in: *Proc. 8th Int. Conf. Model Driven Eng. Lang. Syst.*, Vol. 3713 of LNCS, Springer, 2005, pp. 17–38.
- [A18] R. Bendraou, P. Desfray, M.-P. Gervais, A. Muller, MDA tool components: A proposal for packaging know-how in model driven development, *Softw. Syst. Model.* 7 (3) (2008) 329–343.
- [A19] V. Jain, A. Kumar, P. R. Panda, A SysML profile for development and early validation of TLM 2.0 models, in: *Proc. 7th Europ. Conf. Model. Found. Appl.*, Vol. 6698 of LNCS, Springer, 2011, pp. 299–311.
- [A20] S. K. Johnson, A. W. Brown, A model-driven development approach to creating service-oriented solutions, in: *Service-Oriented Comput.*, Vol. 4294 of LNCS, Springer, 2006, pp. 624–636.
- [A21] E. L. Alves, P. D. Machado, F. Ramalho, Automatic generation of built-in contract test drivers, *Softw. Syst. Model.* (2012) 1–25.
- [A22] D. Hatebur, M. Heisel, A UML profile for requirements analysis of dependable software, in: *Proc. 29th Int. Conf. Comput. Safety, Reliab., Secur.*, Vol. 6351 of LNCS, Springer, 2010, pp. 317–331.
- [A23] D. Hatebur, M. Heisel, Making pattern- and model-based software development more rigorous, in: *Formal Methods and Softw. Eng.*, Vol. 6447 of LNCS, Springer, 2010, pp. 253–269.
- [A24] W. Dahman, J. Grabowski, UML-based specification and generation of executable web services, in: *Syst. Anal. and Model.: About Models*, Vol. 6598 of LNCS, Springer, 2011, pp. 91–107.
- [A25] H. G. Min, S. D. Kim, A technique to represent and generate components in MDA/PIM for automation, in: *Fund. Appr. to Softw. Eng.*, Vol. 3922 of LNCS, Springer, 2006, pp. 293–307.
- [A26] F. Lagarde, F. Terrier, C. André, S. Gérard, Constraints modeling for (profiled) UML models, in: *Model Driven Archit. – Found. and Appl.*, Vol. 4530 of LNCS, Springer, 2007, pp. 130–143.
- [A27] J. E. Pérez-Martínez, A. Sierra-Alonso, From analysis model to software architecture: A PIM2PIM mapping, in: *Model Driven Archit. – Found. and Appl.*, Vol. 4066 of LNCS, Springer, 2006, pp. 25–39.
- [A28] S. Bernardi, F. Flammini, S. Marrone, J. Merseguer, C. Papa, V. Vittorini, Model-driven availability evaluation of railway control systems, in: *Proc. 30th Int. Conf. Comput. Safety, Reliab., Secur.*, Vol. 6894 of LNCS, Springer, 2011, pp. 15–28.
- [A29] S. Bernardi, J. Merseguer, D. C. Petriu, A dependability profile within MARTE, *Softw. Syst. Model.* 10 (3) (2011) 313–336.
- [A30] B. Hofreiter, Extending UN/CEFACT’s modeling methodology by a UML profile for local choreographies, *Inform. Syst. and e-Bus. Manag.* 7 (2) (2009) 251–271.
- [A31] V. Ermagan, I. H. Krüger, A UML2 profile for service modeling, in: *Proc. 10th Int. Conf. Model-Driven Eng. Lang. Syst.*, Vol. 4735 of LNCS, Springer, 2007, pp. 360–374.
- [A32] V. Grassi, Performance analysis of mobile systems, in: *Formal Methods for Mobile Comput.*, Vol. 3465 of LNCS, Springer, 2005, pp. 107–154.
- [A33] L. Baresi, R. Heckel, S. Thöne, D. Varró, Style-based modeling and refinement of service-oriented architectures, *Softw. Syst. Model.* 5 (2) (2006) 187–207.
- [A34] R. Grønmo, M. I. Jaeger, H. Hoff, Transformations between UML and OWL-S, in: *Model Driven Archit. – Found. and Appl.*, Vol. 3748 of LNCS, Springer, 2005, pp. 269–283.
- [A35] R. Ellner, S. Al-Hilank, J. Drexler, M. Jung, D. Kips, M. Philippsen, eSPEM – a SPEM extension for enactable behavior modeling, in: *Proc. 6th Europ. Conf. Model. Found. Appl.*, Vol. 6138 of LNCS, Springer, 2010, pp. 116–131.
- [A36] K. Berkenkötter, U. Hannemann, Modeling the railway control domain rigorously with a UML 2.0 profile, in: *Proc. 25th Int. Conf. Comput. Safety, Reliab., Secur.*, Vol. 4166 of LNCS, Springer, 2006, pp. 398–411.
- [A37] M. Saleem, J. Jaafar, M. Hassan, Secure business process modelling of SOA applications using “UML-SOA-Sec”, *Int. J. Innov. Comput. I.* 8 (4) (2012) 2729–2746.
- [A38] K. Anastakis, B. Bordbar, G. Georg, I. Ray, On challenges of model transformation from UML to Alloy, *Softw. Syst. Model.* 9 (1) (2010) 69–86.
- [A39] R. Behjati, T. Yue, S. Nejati, L. Briand, B. Selic, Extending SysML with AADL concepts for comprehensive system architecture modeling, in: *Proc. 7th Europ. Conf. Model. Found. Appl.*, Vol. 6698 of LNCS, Springer, 2011, pp. 236–252.
- [A40] P. Salehi, A. Hamoud-Lhadj, P. Colombo, F. Khendek, M. Toeroe, A UML-based domain specific modeling language for the availability management framework, in: *Proc. 12th Int. Sym. High-Assur. Syst. Eng.*, 2010, pp. 35–44.
- [A41] M. Alam, R. Breu, M. Hafner, Model-driven security engineering for trust management in SECTET, *J. Softw.* 2 (1) (2007) 47–59.
- [A42] P. Mayer, A. Schroeder, N. Koch, MDD4SOA: Model-driven service orchestration, in: *Proc. 12th Int. Conf. Enterp. Distrib. Object Comput.*, IEEE, 2008, pp. 203–212.
- [A43] S. Gilmore, L. Gönczy, N. Koch, P. Mayer, M. Tribastone, D. Varró, Non-functional properties in the model-driven development of service-oriented systems, *Softw. Syst. Model.* 10 (3) (2011) 287–311.
- [A44] G. Zoughbi, L. Briand, Y. Labiche, Modeling safety and airworthiness (RTCA DO-178B) information: Conceptual model and UML profile, *Softw. Syst. Model.* 10 (3) (2011) 337–367.
- [A45] J. Trujillo, E. Soler, E. Fernández-Medina, M. Piattini, A UML 2.0

- profile to define security requirements for data warehouses, *Comp. Stand. Inter.* 31 (5) (2009) 969–983.
- [A46] B. P. Lamancha, M. Polo, D. Caivano, M. Piattini, G. Visaggio, Automated generation of test oracles using a model-driven approach, *Inform. Softw. Tech.* 55 (2) (2013) 301–319.
- [A47] E. Cariou, C. Ballagny, A. Feugas, F. Barbier, Contracts for model execution verification, in: *Model. Found. Appl.*, Vol. 6698 of LNCS, Springer, 2011, pp. 3–18.
- [A48] V. Garousi, L. C. Briand, Y. Labiche, Control flow analysis of UML 2.0 sequence diagrams, in: *Model Driven Archit. – Found. and Appl.*, Vol. 3748 of LNCS, Springer, 2005, pp. 160–174.
- [A49] J. Zubcoff, J. Pardillo, J. Trujillo, A UML profile for the conceptual modelling of data-mining with time-series in data warehouses, *Inform. Softw. Tech.* 51 (6) (2009) 977–992.
- [A50] I. Reinhartz-Berger, A. Sturm, Utilizing domain models for application design and validation, *Inform. Softw. Tech.* 51 (8) (2009) 1275–1289.
- [A51] J. Kryštof, An LPGM method: Platform independent modeling and development of graphical user interface, *Informatica (Ljubljana)* 34 (3) (2010) 353–367.
- [A52] S. Diaw, R. Lbath, B. Coulette, Specification and implementation of SPEM4MDE, a metamodel for MDE software processes, in: *Proc. 23rd Int. Conf. Softw. Eng. & Knowl. Eng.*, 2011, pp. 646–653.
- [A53] L. M. Bibbo, D. García, C. Pons, A domain specific language for the development of collaborative systems, in: *Proc. Int. Conf. Chilean Comp. Sci. Soc.*, IEEE, 2008, pp. 3–12.
- [A54] E. Riccobene, P. Scandurra, A. Rosti, S. Bocchio, A UML 2.0 profile for SystemC: Toward high-level SoC design, in: *Proc. 5th Int. Conf. Embedded Softw.*, ACM, 2005, pp. 138–141.
- [A55] A. Braganca, R. J. Machado, Extending UML 2.0 metamodel for complementary usages of the «extend» relationship within use case variability specification, in: *Proc. 10th Int. Conf. Softw. Product Line*, IEEE, 2006, pp. 123–130.
- [A56] P. Cáceres, V. de Castro, J. M. Vara, E. Marcos, Model transformations for hypertext modeling on web information systems, in: *Proc. 21st Annu. ACM Sym. Applied Comput.*, ACM, 2006, pp. 1232–1239.
- [A57] I.-C. Hsu, Extending UML to model web 2.0-based context-aware applications, *Softw. Pract. Exper.* 42 (10) (2012) 1211–1227.
- [A58] R. Panesar-Walawege, M. Sabetzadeh, L. Briand, A model-driven engineering approach to support the verification of compliance to safety standards, in: *Proc. 22nd Int. Sym. Softw. Reliab. Eng.*, 2011, pp. 30–39.
- [A59] R. K. Panesar-Walawege, M. Sabetzadeh, L. Briand, Using UML profiles for sector-specific tailoring of safety evidence information, in: *Proc. 30th Int. Conf. Conceptual Model.*, Springer, 2011, pp. 362–378.
- [A60] S. S. Somé, A meta-model for textual use case description, *J. Object Technol.* 8 (7) (2009) 87–106.
- [A61] F. Aoussat, M. Oussalah, M. Nacer, SPEM extension with software process architectural concepts, in: *Proc. 35th Annu. IEEE Int. Conf. Comp. Softw. and Appl.*, 2011, pp. 215–223.
- [A62] M. Bošković, W. Hasselbring, Model driven performance measurement and assessment with MoDePeMART, in: *Model Driven Eng. Lang. Syst.*, Vol. 5795 of LNCS, Springer, 2009, pp. 62–76.
- [A63] P. Liegl, C. Huemer, C. Pichler, Registry support for core component-based business document models, *Service Oriented Comput. and Appl.* 5 (3) (2011) 183–202.
- [A64] B. Agreiter, R. Breu, Model-driven configuration of SELinux policies, in: *On the Move to Meaningful Internet Syst.*, Vol. 5871 of LNCS, Springer, 2009, pp. 887–904.
- [A65] C. Song, E. Cho, C. Kim, An integrated GUI-business component modeling method for the MDD- and MVC-based hierarchical designs, *Int. J. Softw. Eng. Know.* 21 (3) (2011) 447–490.
- [A66] D. Mouheb, C. Talhi, A. Mourad, V. Lima, M. Debbabi, L. Wang, M. Pourzandi, An aspect-oriented approach for software security hardening: From design to implementation, in: *Proc. 8th Conf. New Trends in Softw. Methodol., Tools and Tech.*, IOS Press, 2009, pp. 203–222.
- [A67] M. Miguel, J. Briones, J. Silva, A. Alonso, Integration of safety analysis in model-driven software development, *IET Softw.* 2 (3) (2008) 260–280.
- [A68] A. M. R. da Cruz, J. a. P. Faria, A metamodel-based approach for automatic user interface generation, in: *Proc. 13th Int. Conf. Model Driven Eng. Lang. Syst.*, Vol. 6394 of LNCS, Springer, 2010, pp. 256–270.
- [A69] J. Bergh, K. Coninx, CUP 2.0: High-level modeling of context-sensitive interactive applications, in: *Model Driven Eng. Lang. Syst.*, Vol. 4199 of LNCS, Springer, 2006, pp. 140–154.
- [A70] M. Z. Iqbal, A. Arcuri, L. Briand, Environment modeling with UML/MARTE to support black-box system testing for real-time embedded systems: Methodology and industrial case studies, in: *Proc. 13th Int. Conf. Model Driven Eng. Lang. Syst.*, Vol. 6394 of LNCS, Springer, 2010, pp. 286–300.
- [A71] R. Barrett, L. M. Patcas, C. Pahl, J. Murphy, Model driven distribution pattern design for dynamic web service compositions, in: *Proc. 6th Int. Conf. Web Eng.*, ACM, 2006, pp. 129–136.
- [A72] N. Koch, G. Zhang, M. J. Escalona, Model transformations from requirements to web system design, in: *Proc. 6th Int. Conf. Web Eng.*, ACM, 2006, pp. 281–288.
- [A73] Z. Cui, L. Wang, X. Li, D. Xu, Modeling and integrating aspects with UML activity diagrams, in: *Proc. 24th Annu. ACM Sym. Applied Comput.*, ACM, 2009, pp. 430–437.
- [A74] G. N. Rodrigues, D. S. Rosenblum, S. Uchitel, Reliability prediction in model-driven development, in: *Model Driven Eng. Lang. Syst.*, Vol. 3713 of LNCS, Springer, 2005, pp. 339–354.
- [A75] J. U. Júnior, R. D. Penteado, V. V. de Camargo, An overview and an empirical evaluation of UML-AOF: An UML profile for aspect-oriented frameworks, in: *Proc. 25th Annu. ACM Sym. Applied Comput.*, ACM, 2010, pp. 2289–2296.
- [A76] K. Lind, R. Heldal, A model-based and automated approach to size estimation of embedded software components, in: *Proc. 14th Int. Conf. Model Driven Eng. Lang. Syst.*, Springer, 2011, pp. 334–348.
- [A77] U. Zdun, P. Avgeriou, Modeling architectural patterns using architectural primitives, in: *Proc. 20th Annu. ACM SIGPLAN Conf. Object-oriented Prog., Syst., Lang., Appl.*, ACM, 2005, pp. 133–146.
- [A78] P. Baker, P. Bristow, C. Jervis, D. King, R. Thomson, B. Mitchell, S. Burton, Detecting and resolving semantic pathologies in UML sequence diagrams, in: *Proc. 10th Europ. Softw. Eng. Conf. and the 13th ACM SIGSOFT Int. Sym. Found. of Softw. Eng.*, ACM, 2005, pp. 50–59.
- [A79] M. Hafner, R. Breu, Realizing model driven security for inter-organizational workflows with WS-CDL and UML 2.0, in: *Proc. 8th Int. Conf. Model Driven Eng. Lang. Syst.*, Vol. 3713 of LNCS, Springer, 2005, pp. 39–53.
- [A80] T. Schattkowsky, J. Hausmann, G. Engels, Using UML activities for system-on-chip design and synthesis, in: *Model Driven Eng. Lang. Syst.*, Vol. 4199 of LNCS, Springer, 2006, pp. 737–752.
- [A81] M. G. Uddin, M. Zulkernine, UMLtrust: Towards developing trust-aware software, in: *Proc. 23rd Annu. ACM Sym. Applied Comput.*, ACM, 2008, pp. 831–836.
- [A82] A. Tretiakov, S. Hartmann, Higher-order entity relationship modelling with UML, in: *Proc. 5th Int. Conf. Qual. Softw.*, IEEE, 2005, pp. 205–211.
- [A83] N. Moreno, P. Fraternali, A. Vallecillo, WebML modelling in UML, *IET Softw.* 1 (3) (2007) 67–80.
- [A84] J. Romero, A. Vallecillo, Modeling the ODP computational viewpoint with UML 2.0, in: *Proc. 9th Int. EDOC Enterp. Comput. Conf.*, 2005, pp. 169–180.
- [A85] M. Nikolaidou, N. Alexopoulou, A. Tsadimas, A. Dais, D. Anagnostopoulos, Accommodating EIS UML 2.0 profile using a standard UML modeling tool, in: *Proc. Int. Conf. Softw. Eng. Adv.*, 2007.
- [A86] C. Bartolini, A. Bertolino, G. De Angelis, G. Lipari, A UML profile and a methodology for real-time systems design, in: *Proc. 32nd EUROMICRO Conf. Softw. Eng. and Adv. Appl.*, 2006, pp. 108–117.
- [A87] I. Ivkovic, K. Kontogiannis, A framework for software architecture refactoring using model transformations and semantic annotations, in: *Proc. 10th Europ. Conf. Softw. Maint. and ReEng.*, 2006, pp. 135–144.
- [A88] H. Wada, J. Suzuki, K. Oba, Modeling non-functional aspects in service oriented architecture, in: *Proc. 3rd Int. Conf. Services Comput.*, 2006, pp. 222–229.

- [A89] E. Soler, J. Trujillo, E. Fernández-Medina, M. Piattini, Building a secure star schema in data warehouses by an extension of the relational package from CWM, *Comp. Stand. Inter.* 30 (6) (2008) 341–350.
- [A90] R. Villarroel, E. Fernández-Medina, M. Piattini, J. Trujillo, A UML 2.0/OCL extension for designing secure data warehouses, *J. Res. Pract. Inf. Tech.* 38 (1).
- [A91] A. Queralt, E. Teniente, A platform independent model for the electronic marketplace domain, *Softw. Syst. Model.* 7 (2) (2008) 219–235.
- [A92] G. M. Kapitsaki, D. A. Kateros, G. N. Prezerakos, I. S. Venieris, Model-driven development of composite context-aware web applications, *Inform. Softw. Tech.* 51 (8) (2009) 1244–1260.
- [A93] J. Evermann, A meta-level specification and profile for AspectJ in UML, *J. Object Technol.* 6 (7) (2007) 27–49.
- [A94] G. Prezerakos, N. Tselikas, G. Cortese, Model-driven composition of context-aware web services using ContextUML and aspects, in: *Proc. IEEE Int. Conf. Web Services, IEEE, 2007*, pp. 320–329.
- [A95] A. Cicchetti, D. D. Ruscio, A. Pierantonio, A metamodel independent approach to difference representation, *J. Object Technol.* 6 (9) (2007) 165–185.
- [A96] Y. Zhang, Y. Liu, L. Zhang, Z. Ma, H. Mei, Modeling and checking for non-functional attributes in extended UML class diagram, in: *Proc. 32nd Annu. IEEE Int. Conf. Comp. Softw. and Appl., IEEE, 2008*, pp. 100–107.