

Collaboro: a collaborative (meta) modeling tool

Javier Luis Cánovas Izquierdo¹ and Jordi Cabot^{1,2}

¹ Universitat Oberta de Catalunya (UOC), Barcelona, Spain

² Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain

ABSTRACT

Software development is becoming more and more collaborative, emphasizing the role of end-users in the development process to make sure the final product will satisfy customer needs. This is especially relevant when developing Domain-Specific Modeling Languages (DSMLs), which are modeling languages specifically designed to carry out the tasks of a particular domain. While end-users are actually the experts of the domain for which a DSML is developed, their participation in the DSML specification process is still rather limited nowadays. In this paper, we propose a more community-aware language development process by enabling the active participation of all community members (both developers and end-users) from the very beginning. Our proposal, called Collaboro, is based on a DSML itself enabling the representation of change proposals during the language design and the discussion (and trace back) of possible solutions, comments and decisions arisen during the collaboration. Collaboro also incorporates a metric-based recommender system to help community members to define high-quality notations for the DSMLs. We also show how Collaboro can be used at the model-level to facilitate the collaborative specification of software models. Tool support is available both as an Eclipse plug-in a web-based solution.

Subjects Programming Languages, Software Engineering

Keywords Collaborative development, Domain-specific languages, Model-driven development

INTRODUCTION

Collaboration is key in software development, it promotes a continual validation of the software to be build (*Hildenbrand et al., 2008*), thus guaranteeing that the final software will satisfy the users' needs. Furthermore, the sooner the end-users participate in the development life-cycle, the better, as several works claim (*Hatton & van Genuchten, 2012; Rooksby & Ikeya, 2012; Dullemond, van Gameren & van Solingen, 2014*). When the software artefacts being developed target a very specific and complex domain, this collaboration makes even more sense. Only the end-users have the domain knowledge required to drive the development. This is exactly the scenario we face when performing (meta) modeling tasks.

On the one hand, end-users are key when defining a Domain-Specific Modeling Language (DSML), a modeling language specifically designed to perform a task in a certain domain (*Sánchez Cuadrado & García Molina, 2007*). Clearly, to be useful, the concepts and notation of a DSML should be as close as possible to the domain concepts and representation used by the end-users in their daily practice (*Grundy et al., 2013*). Therefore, the role of domain experts during the DSML specification is vital, as noted by

Submitted 15 May 2016
Accepted 18 August 2016
Published 24 October 2016

Corresponding author
Javier Luis Cánovas Izquierdo,
jcanovasi@uoc.edu

Academic editor
Marian Gheorghie

Additional Information and
Declarations can be found on
page 31

DOI 10.7717/peerj-cs.84

© Copyright
2016 Cánovas Izquierdo and Cabot

Distributed under
Creative Commons CC-BY 4.0

OPEN ACCESS

several authors (*Kelly & Pohjonen, 2009; Mernik, Heering & Sloane, 2005; Völter, 2011; Barišić et al., 2012*). Unfortunately, nowadays, participation of end-users is still mostly restricted to the initial set of interviews to help designers analyze the domain and/or to test the language at the end (also scarcely done as reported by *Gabriel, Goulão & Amaral (2010)*), which requires the development of fully functional language toolsets (including a model editor, a parser, etc.) (*Mernik, Heering & Sloane, 2005; Cho, Gray & Syriani, 2012*). This long iteration cycle is a time-consuming and repetitive task that hinders the process performance (*Kelly & Pohjonen, 2009*) since end-users must wait until the end to see if designers correctly understood all the intricacies of the domain. On the other hand, those same end-users will then employ that modeling language (or any general-purpose (modeling) language like UML) to specify the systems to be built. Collaboration here is also key in order to enable the participation of several problem experts in the process. Recently, modeling tools have been increasingly enabling the collaborative development of models defined with either General-Purpose Languages (GPLs) or DSMLs. However, their support for asynchronous collaboration is still limited, specially when it comes to the traceability and justification of modeling decisions.

Existing project management tools such as Trac (<http://trac.edgewall.org/>) or Jira (<http://www.atlassian.com/es/software/jira/overview>) provide the environments required to develop collaboratively software systems. These tools enable the end-user participation during the process, thus allowing developers to receive feedback at any time (*Cabot & Wilson, 2009*). However, their support is usually defined at file level, meaning that discussions and change tracking are expressed in terms of lines of textual files. This is a limitation when developing or using modeling languages, where a special support to discuss at language element level (i.e., domain concepts and notation symbols) is required to address the challenges previously described and therefore promote the participation of end-users. As mentioned above, a second major problem shared by current solutions is the lack of traceability of the design decisions. The rationale behind decisions made during the language/model specification are implicit so it is not possible to understand or justify why, for instance, a certain element of the language was created with that specific syntax or given that particular type. This hampers the future evolution of the language/model.

In order to alleviate these shortcomings, we define a DSML called *Collaboro*, which enables the involvement of the community (i.e., end-users and developers) in the development of (meta) modeling tasks. It allows modeling the collaborations between community members taking place during the definition of a new DSML. *Collaboro* supports both the collaborative definition of the abstract (i.e., metamodel) and concrete (i.e., notation) syntaxes for DSMLs by providing specific constructs to enable the discussion. Also, it can be easily adapted to enable the collaborative definition of models. Thus, each community member has the chance to request changes, propose solutions and give an opinion (and vote) on those from others. We believe this discussion enriches the language definition and usage significantly, and ensures that the end result satisfies as much as possible the expectations of the end-users. Moreover, the explicit recording of these interactions provides plenty of valuable information to explain the language evolution and justify all design decisions behind it, as also proposed in requirements

engineering (Jureta, Faulkner & Schobbens, 2008). Together with the Collaboro DSML, we provide the tooling infrastructure and process guidance required to apply Collaboro in practice. In this paper we will use the term Collaboro to refer to both the DSML and the developed tooling.

The first version of Collaboro, which supported the collaborative development of textual DSMLs in an Eclipse-based environment, was presented in a previous work by (Cánovas Izquierdo & Cabot, 2013). Since then, the approach has evolved to include new features such as: (1) better support for the collaborative development of graphical DSMLs; (2) a new architecture which includes a web-based front-end, thus promoting usability and participation for end-users; (3) a metric-based recommender system, which checks the DSMLs under development to spot possible issues according to quality metrics for both the domain and the notation (relying on Moody's cognitive framework (Moody, 2009)); and (4) the development of several cases studies, which have allowed us to improve the general expressiveness and usability of our approach. Additionally, in this paper we describe how our tool could be easily adapted to support collaborative modeling.

Paper structure. The first two sections describe the proposal and approach to develop DSML collaborative. The following section shows then how our approach could be easily adapted to use any modeling language to model collaboratively. Next, the implemented tool and a case study are described. Finally, we review the related work and draw some conclusions and future work.

COLLABORATIVE (META) MODELING

While collaboration is crucial in both defining modeling languages and then using them to model concrete systems, the collaborative aspects of language development are more challenging and less studied since collaboration must cover both the definition of a new notation for the language and the specification of the language primitives themselves. Therefore, we will present first Collaboro in the context of collaborative language development and later its adaptation to cover the simpler modeling scenario. A running example, also introduced in this section, will help to illustrate the main concepts of such collaborations.

A DSML is defined through three main components (Kleppe, 2008): abstract syntax, concrete syntax, and semantics. The abstract syntax defines both the language concepts and their relationships, and also includes well-formedness rules constraining the models that can be created. Metamodeling techniques are normally used to define the abstract syntax. The concrete syntax defines a notation (textual, graphical or hybrid) for the concepts in the abstract syntax, and a translational approach is normally used to provide semantics, though most of the time it is not explicitly formalized.¹

The development of a DSML usually consists in five different phases (Mernik, Heering & Sloane, 2005): decision, analysis, design, implementation and deployment. The first three phases are mainly focused on the DSML definition whereas the implementation phase is aimed at developing the tooling support (i.e., modeling environment, parser, etc.) for the DSML. Clearly, the community around the language is a key element in the

¹ The collaborative definition of the semantics of a new DSML is out of the scope of this paper and considered as part of future work.

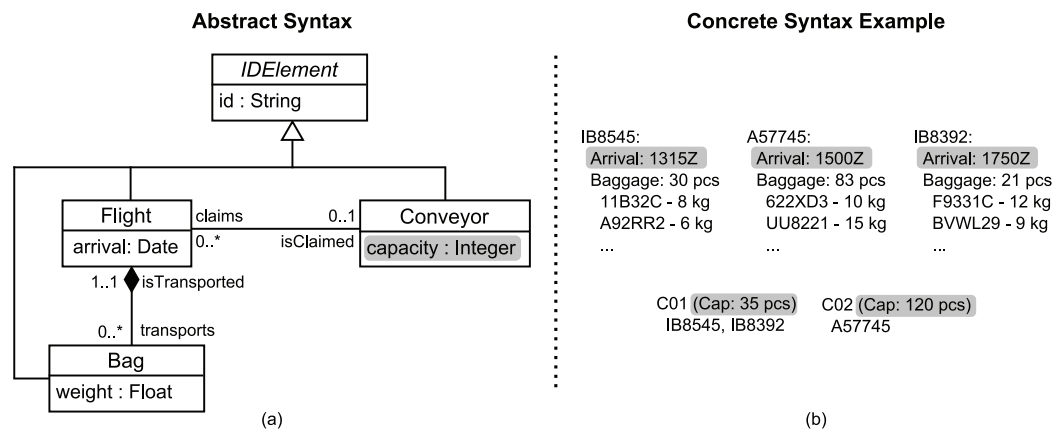


Figure 1 Abstract syntax and an example of concrete syntax of the *Baggage Claim* DSML (grey-filled boxes represent elements added after the collaboration).

process. In this paper we use the term *community* to refer to what is known as *Communities of Practice*, which is defined as *groups of people informally bound together by shared expertise and passion for a joint enterprise* (Tamburri, Lago & Vliet, 2013). In this case, the DSML community covers the group of people involved in its development, which includes both technical level users (i.e., language developers) and domain expert users (i.e., end-users of the language), where both categories can overlap.

As a running example, imagine the development of a DSML to facilitate the planification of the baggage claim service in airports. Let's assume that the airport baggage service needs to specify every morning the full daily assignment of flights to baggage claim conveyors so that operators can know well in advance how to configure the actual baggage system. For that, developers and domain experts (i.e., baggage managers) collaborate to define a DSML that serves this purpose.

Typically, domain experts are only involved at the very beginning and very end of the DSML development process. Assuming this is also the case for our example, during the analysis phase, developers would study the domain with the help of the baggage managers and decide that the DSML should include concepts such as *Flight*, *Bag* and *Conveyors* to organize the baggage delivery. Developers would design and later implement the tooling of the language, thus coming up with a textual DSML like, for instance, the one shown in Fig. 1 (both abstract and concrete syntax proposals are shown, except for the elements included in grey-filled boxes that are added later as explained in what follows). Note that the concrete syntax is illustrated by means of a sample model conforming to the abstract syntax, other possible notations could be defined.

Once the language is developed, end-users can play with it and check whether it fits their needs. Quite often, if the end-users only provided the initial input but did not closely follow how that was interpreted during the language design, they might detect problems in the DSML environment (e.g., missing concepts, wrong notation, etc.) that will trigger a new (and costly) iteration to modify the language and recreate all the associated tools. For instance, end-users could detect that the language lacks a construct to represent the capacity of conveyors, that may help them to perform a better assignment.

Developers can also overlook design constraints and recommendations that could improve the DSML quality. For instance, constructs in the abstract syntax not having a concrete syntax definition could become an issue (e.g., `arrival` attribute in `Flight` concept).

The collaboration tasks can go beyond the definition of new DSML and can cover the usage of well-known GPLs, like UML. Let's imagine for instance the collaborative definition of class diagrams in order to identify the domain of a new software artefact. In fact, we could even reuse the running example to illustrate this scenario. Thus, the definition of the abstract syntax of the previous DSML requires the collaborative creation of a UML class diagram. In this sense, end-users (i.e., domain experts) use a common language (i.e., UML) to create a new model required for a particular purpose (in this case, the definition of a DSML). As before, end-users can propose changes to the model, which can after be discussed and eventually accepted/rejected in the final version.

Our aim is to incorporate the community collaboration aspect into all DSML definition phases, making the phases of the process more participative and promoting the early detection of possible bugs or problems. As we will see, this support also enables the collaborative creation of models conforming to modeling languages. Next section will present our approach.

MAKING DSML DEVELOPMENT COLLABORATIVE

We propose a collaborative approach to develop DSMLs following the process summarized in Fig. 2. Roughly speaking, the process is as follows. Once there is an agreement to create the language, developers get the requirements from the end-users to create a preliminary version of the language to kickstart the actual collaboration process (step 1). This first version should include at least a partial abstract syntax but could also include a first concrete syntax draft (see *DSML Definition*). An initial set of sample models are also defined by the developers to facilitate an example-based discussion, usually easier for non-technical users. Sample models are rendered according to the current concrete syntax definition (see *Rendered Examples*). It is worth noting that the rendering is done on-the-fly without the burden of generating the DSML tooling since we are just showing the snapshots of the models to discuss the notation, not actually providing at this point a full modeling environment.

Now the community starts working together in order to shape the language (step 2). Community members can propose ideas or changes to the DSML, e.g., they can ask for modifications on how some concepts should be represented (both at the abstract and concrete syntax levels). These change proposals are shared in the community, who can also suggest and discuss how to improve the change proposals themselves. All community members can also suggest solutions for the requested changes and give their opinion on the solutions presented by others. At any time, rendering the sample models with the latest proposals helps members to have an idea of how a given proposal will evolve the language (if accepted). During this step, a recommender system (see *Recommender*) also checks the current DSML definition to spot possible issues according to quality metrics for DSMLs. If the recommender system detects possible improvements, it will

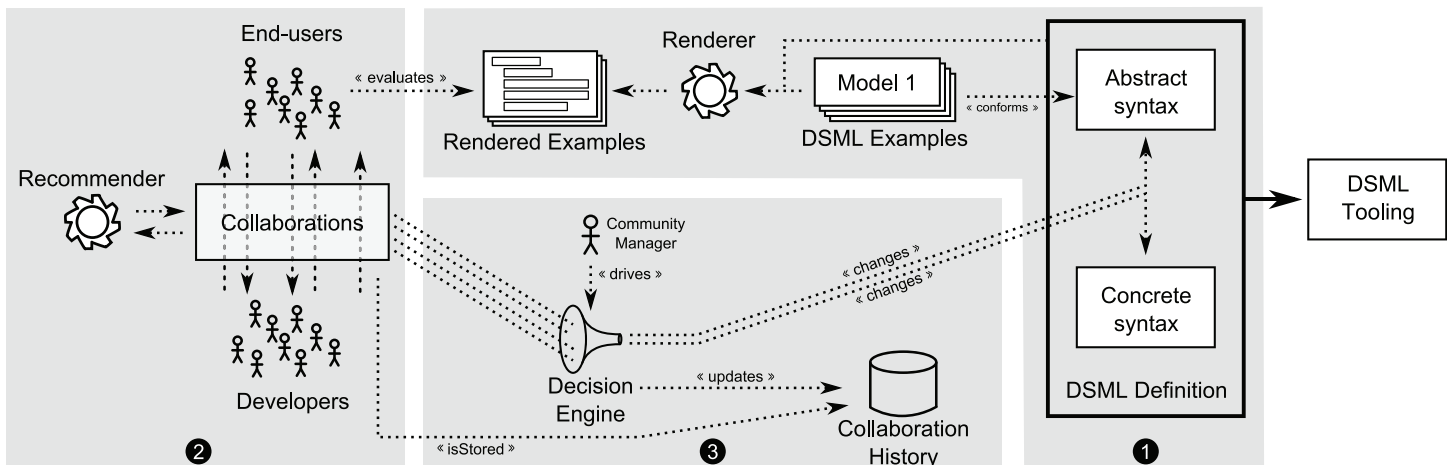


Figure 2 Collaborative development of DSMLs.

create new proposals to be also discussed by the community. All these proposals and solutions (see *Collaborations*) are eventually accepted or rejected.

Acceptance/rejection depends on whether the community reaches an agreement regarding the proposal/solution. For that, community members can vote (step 3). A decision engine (see *Decision Engine*) then takes these votes into account to calculate which collaborations are accepted/rejected by the community. The engine could follow an automatic process but a specific role of *community manager* could also be assigned to a member/s to consolidate the proposals and get a consensus on conflicting opinions (e.g., when there is no agreement between technical and business considerations). Once an agreement is reached, the contents of the solution are incorporated into the language, thus creating a new version. The process keeps iterating until no more changes are proposed. Note that these changes on the language may also have an impact on the model examples which may need to be updated to comply with the new language definition.

At the end of the collaboration, the final DSML definition is used as a starting point to implement a full-fledged DSML tooling (see *DSML Tooling*) with the confidence that it has been validated by the community (e.g., transforming/importing the DSML definition into language workbenches like Xtext or Graphical Modeling Framework (GMF)). Note that even when the language does not comply with commonly applied quality patterns, developers can be sure that it at least fulfills the end-users' needs. Moreover, all aspects of the collaboration are recorded (see *Collaboration History*), thus keeping track of every interaction and change performed in the language. Thus, at any moment, this traceability information can be queried (e.g., using standard Object Constraint Language (OCL) (*Object Management Group, 2015a*) expressions) to discover the rationale behind the elements of the language (e.g., the argumentation provided for its acceptance).

To illustrate our approach, the development of the *Baggage Claim* DSML mentioned above could have been the result of the imaginary collaboration scenario depicted in Fig. 3. After developers completed a first version of the language, the collaboration

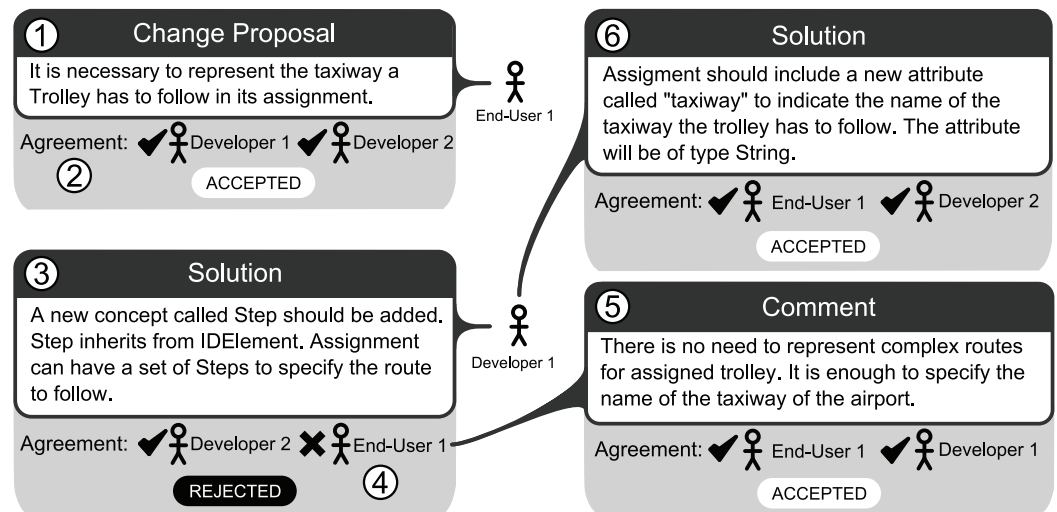


Figure 3 Example of collaboration in the *Baggage Claim DSML*.

begins with a community member detecting the need of expressing the capacity of the conveyors. Since now we are still in the definition phase, the community has the chance to discuss the best way to adapt the language to support this new information. The member that identified the problem would create a change proposal with that aim, and if the change is deemed as important by the community, other members could propose a solution/s to adapt the language. As an example, [Fig. 3](#) graphically depicts a possible collaboration scenario assuming a small community of one end-user and two developers. Each collaboration is represented as a bubble, and each step has been numbered. In the Figure, *End-User 1* proposes a language change (step 1), which is accepted by the community (step 2), and then *Developer 1* specifies a solution (step 3). The solution is rejected by *End-User 1*, including also the explanation of the rejection (step 4). As the rejection is accepted (step 5), the *Developer 1* redefines the solution, which is eventually accepted (step 6) and the changes are then incorporated into the language. The resulting changes in the abstract and concrete syntaxes are shown in grey-filled boxes in [Fig. 1](#). Clearly, it is important to make this collaboration iterations as quick as possible and with the participation of all the community members. Moreover, the discussion information itself must be preserved to justify the rationale behind each language evolution, from which design decisions can be derived.

The recommender system may also participate in the collaboration and eventually improve the DSML. After checking the DSML definition, the recommender may detect that not all the attributes in the abstract syntax have a direct representation in the concrete syntax, as it happens with the `arrival` attribute of the `Flight` concept (as we will explain later, this is the result of applying the metric called *Symbol Deficit*). Thus, the system may create a new proposal informing about the situation and then the community could vote and eventually decide whether the DSML has to be modified.

Our proposal for enabling the collaborative definition of DSMLs is built on top of the *Collaboro* DSML, a DSML for modeling the collaborations that arise in a community

working towards the development of a DSML for that community. In the next sections, we will describe how Collaboro makes the collaboration feasible by:

- Enabling the discussion about DSML elements,
- providing the metaclasses for representing collaborations and giving support to the decision-making process,
- providing a metric-based recommender that can help to develop high-quality DSMLs.

Representing the elements of a DSML

To be able to discuss about changes on the DSML to-be, we must be able to represent both its abstract syntax (i.e., the concepts of the DSML) and its concrete syntax (the notation to represent those concepts) elements. Additionally, to improve the understanding of how changes in its definition affect the DSML, we provide a mechanism to automatically render DSML examples using the concrete syntax notation under development.

Abstract syntax

The abstract syntax of a DSML is commonly defined by means of a metamodel written using a metamodeling language (e.g., Meta-Object Facility (MOF) (*Object Management Group, 2015b*) or Ecore (*Steinberg et al., 2008*)). Metamodeling languages normally offer a limited set of concepts to be used when creating DSML metamodels (like types, relationship or hierarchy). A DSML metamodel is then defined as an instantiation of this metamodeling concepts. [Figure 4A](#) shows an excerpt of the well-known Ecore metamodeling language, on which we rely to represent the abstract syntax of DSMLs.

Concrete syntax

Regarding the concrete syntax, since the notation of a DSML is also domain-specific, to promote the discussion, we need to be able to explicitly represent the notational elements proposed for the language. Thanks to this, community members will have the freedom to create a notation specially adapted to their domain, thus avoiding coupling with other existing notations (e.g., Java-based textual languages or UML-like diagrams). The type of notational elements to represent largely depends on the kind of concrete syntax envisioned (textual or graphical). Nowadays, there are some tool-specific metamodels to represent either graphical or textual concrete syntaxes (like the ones included in GMF (<http://eclipse.org/gmf-tooling>), Graphiti (<https://eclipse.org/graphiti>) and Xtext (<http://eclipse.org/Xtext>)), or to interchange model-level diagrams (*Object Management Group, 2014b*). However, a generic metamodel covering both graphical and textual syntaxes (and combinations of both) is not typically available in other tools. Therefore, we contribute in this paper our own metamodel for concrete syntaxes. [Figure 4B](#) shows an excerpt of the core elements of this notation metamodel. As can be seen, the metamodel is not exhaustive, but it is a lightweight solution that suffices to discuss about the concrete syntax elements most commonly used in the definition of graphical, textual or hybrid concrete syntaxes (offering a good trade-off

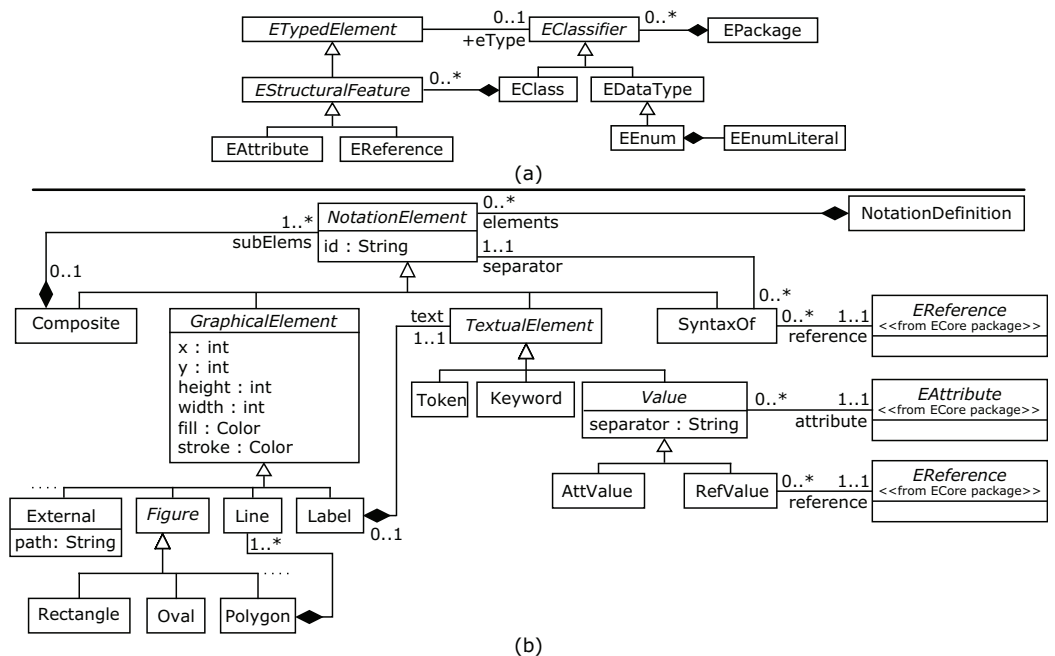


Figure 4 Excerpts of the (A) Ecore and (B) notation metamodels used to represent, respectively, the abstract and concrete syntaxes of DSMLs in Collaboro.

between expressiveness and manageability of the description in order to render and analyze/recommend changes). Note that, with this metamodel, it is possible to describe how to represent each language concept, thus facilitating keeping track of language notation changes.

Concrete syntax elements are classified following the `NotationElement` hierarchy, which includes graphical elements (`GraphicalElement` metaclass), textual elements (`TextualElement` metaclass), composite elements (`Composite` metaclass) and references to the concrete syntax of other abstract elements (`SyntaxOf` metaclass) to be used in composite patterns. The main graphical constructs are provided by the `GraphicalElement` hierarchy, which allows referring to external pictures (`External` metaclass), building figures (see `Figure` hierarchy), lines (`Line` metaclass) and labels for the DSML elements. A label (`Label` metaclass) serves as a container for a textual element. Textual elements can be defined with the `TextualElement` hierarchy, which includes tokens, keywords and values directly taken from the abstract syntax elements expressed in a textual form (`Value` metaclass). It is possible to obtain the textual representation from either an attribute (`AttValue` metaclass) by specifying the attribute to be queried (`attribute` reference), or a reference (`RefValue` metaclass) by specifying both the reference (`reference` reference) and the attribute of the referred element to be used (`attribute` reference). The attribute `separator` of the `Value` metaclass allows defining the separator for multivalued elements. The `Composite` element can be used to define complex concrete syntax structures, allowing both graphical and textual composites but also hybrids. Finally, the `SyntaxOf` metaclass allows referencing to already specified concrete syntax definitions of abstract syntax elements, thus allowing modularization and

C01 (Cap: 35 pcs)
IB8545, IB8392

...

(a)

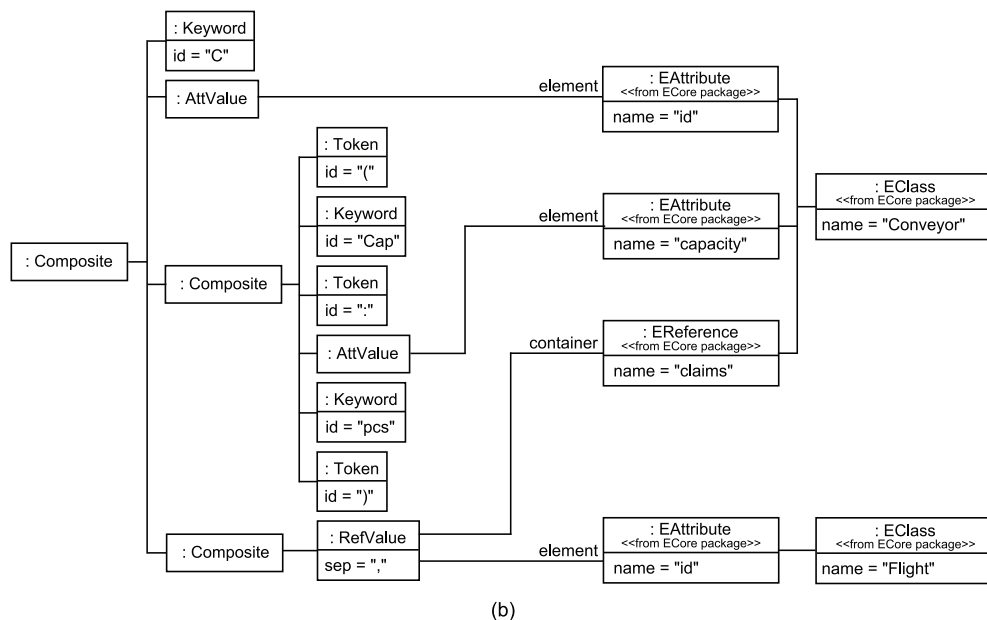


Figure 5 (A) Textual representation example of the metaclass *Conveyor* of the *Baggage Claim* DSML and (B) the corresponding notation model.

composition. The `reference` reference of the `SyntaxOf` metaclass specifies the reference to be queried to obtain the set of elements whereas the `separator` reference indicates the separator between elements.

Renderer

The current DSML notation specification plus the set of example models for the DSML (expressed as instances of the DSML abstract syntax) can be used to generate concrete visual examples that help community members get a better idea of the language being built. We refer to this generator as *renderer*. The renderer takes, as inputs: (1) the abstract and (2) concrete syntaxes of the DSML, and (3) the set of example models conforming to the abstract syntax; and returns a set of images representing the example models expressed according to the concrete syntax defined in the notation model (additional technical details about the render process will be given in the Section describing the developed tooling).

We believe the advantages of this approach is twofold. On the one hand, it is a lightweight mechanism to quickly validate the DSML without generating the DSML tooling support. On the other hand, developers and end-users participating in the collaboration can easily assess how the language looks like without the burden of dealing with the abstract and concrete syntax of DSML, which are expressed as metamodels.

Example

Figure 1A shows an example of the abstract syntax for the *Baggage Claim* DSML while Fig. 5 shows the notation model for the textual representation of the metaclass

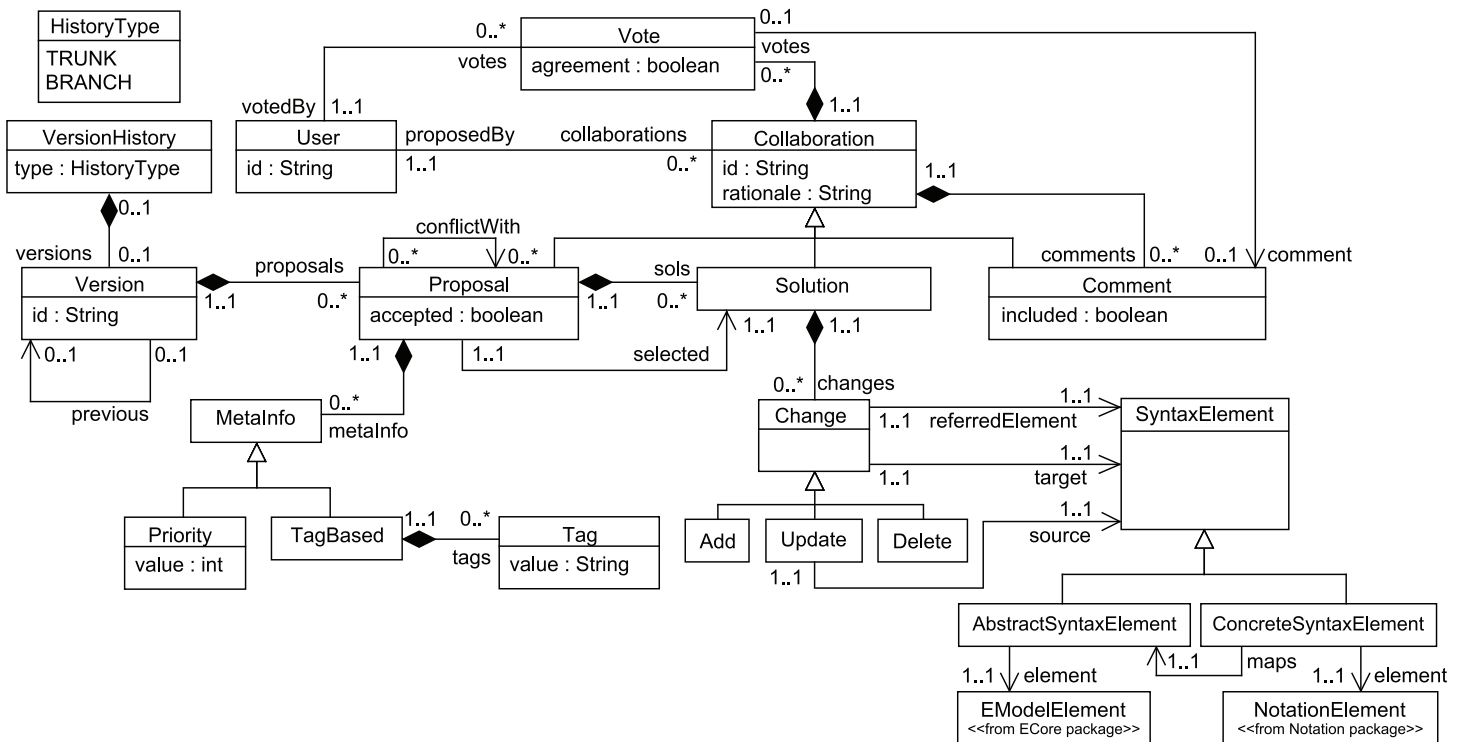


Figure 6 Core elements of the Collaboro metamodel.

Conveyor of such DSML (Fig. 5A shows a textual example and Fig. 5B shows the corresponding notation model). Note that `AttValue` and `RefValue` metaclass instances are referring to elements from the abstract syntax metamodel. Fig. 1B shows a possible renderization of a model for such language.

Representing the collaborations

The third metamodel required in our process focuses on representing the collaborations that annotate/modify the DSML elements described before. This collaboration metamodel, which is shown in Fig. 6, allows representing both static (e.g., change proposals) and dynamic (e.g., voting) aspects of the collaboration. Being the core of our collaborative approach, we refer to this metamodel as the *Collaboro metamodel*.

Static aspects

Similarly to how version control systems track source code, Collaboro also allows representing different versions of a DSML. The `VersionHistory` metaclass represents the set of versions (`Version` metaclass) through which the collaboration evolves. There is always a main version history set as *trunk* (`type` attribute in `VersionHistory` metaclass), which keeps the baseline of the collaborations about the language under development. Other version histories (similar to branches) can be forked when necessary to isolate the collaboration about concrete parts of the language. Different version histories can be merged into a new one (or the *trunk*).

Language evolution is the consequence of collaborations (`Collaboration` metaclass). Collaboro supports three types of collaborations: change proposals (`Proposal` metaclass), solutions proposals (`Solution` metaclass) and comments (`Comment` metaclass). A collaboration is proposed by a user (`proposedBy` reference) and includes an explanation (`rationale` attribute).

A change proposal describes which language feature should be modified and contains some meta information (e.g., priority or tags). Change proposals are linked to the set of solutions proposed by the community to be discussed for agreement. It is also possible to specify possible conflicts between similar proposals (e.g., the acceptance of one proposal can involve rejecting others) with the `conflictWith` reference.

Solution proposals are the answer to change proposals and describe how the language should be modified to incorporate the new features. Each solution definition involves a set of add/update/delete changes on the elements of the DSML (`Change` hierarchy). `Change` links the collaboration metamodel with the DSML under discussion (`SyntaxElement` metaclass), which can refer to the abstract syntax (`AbstractSyntaxElement` metaclass) or the concrete syntax (`ConcreteSyntaxElement` metaclass). The latter links (`maps` reference) to the abstract element to which the notation is defined. Both `AbstractSyntaxElement` and `ConcreteSyntaxElement` metaclasses have a reference linking to the element which is being changed (`element` reference). Changes in the abstract syntax are expressed in terms of the metamodeling language (i.e., `EModelElement` elements, which is the interface implemented by every element in the `Ecore` metamodel) while changes in the concrete syntax are expressed in terms of elements conforming to the notation metamodel presented before.

The `Change` metaclass has a reference to the container element affected by the change (`referredElement` reference) and the element to change (`target` reference). Thereby, in the case of `Add` and `Delete` metaclasses, `referredElement` reference refers to the element to which we want to add/delete a “child” element whereas `target` refers to the actual element to be added/deleted. In the case of the `Update` metaclass, `referredElement` reference refers to the element which contains the element to be updated (e.g., a metaclass) whereas `target` reference refers to the new version of the element being updated (e.g., a new version for an attribute). The additional `source` attribute indicates the element to be updated (e.g., the attribute which is being updated).

Dynamic aspects

During the process, community members vote collaboration elements, thus allowing to reach agreements. Votes (`Vote` metaclass) indicate whether the user (`votedBy` reference) agrees or not with a collaboration (`agreement` attribute). A vote against a collaboration usually includes a comment explaining the reason of the disagreement (`comment` reference of `Vote` metaclass). This comment can then be voted itself and if it is accepted by the community, the proponent of the voted proposal/solution should take such comment into account (the `included` attribute of `Comment` metaclass records this fact).

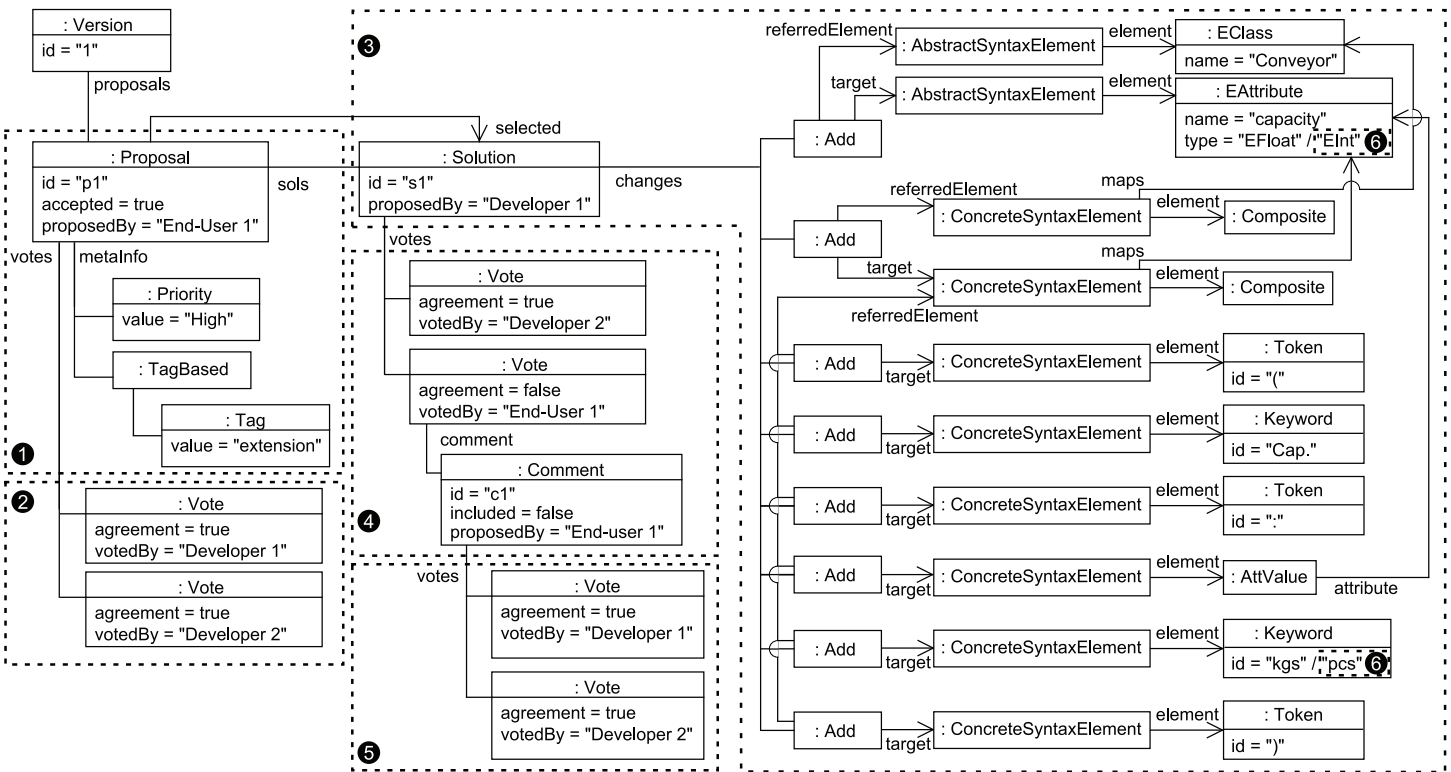


Figure 7 The collaboration model representing the collaborations arisen in the *Baggage Claim DSML*.

The acceptance of a proposal means that the community agrees that the requested change is necessary (accepted attribute). For each proposal we can have several solutions but in the end one of them will be selected (selected reference of the Proposal metaclass) and its changes applied to the DSML definition. Part of this data (like the accepted and selected properties) is automatically filled by the decision engine analyzing and resolving the collaboration.

Making decisions

Community votes are used to decide which collaborations are accepted and must be incorporated into the language. Collaboration models include all the necessary information, thus allowing the automation of the decision process (i.e., approval of change proposals and selection of solutions). A decision engine can therefore apply resolution strategies (e.g., unanimous agreement, majority agreement, etc.) to deduce (and apply) the collaborations accepted by the community. As commented before, most times it is necessary to have the role of the community manager to trigger the decision process and solve possible decision locks.

Example

As example of collaboration, we show in Fig. 7 the collaboration model which would be obtained when using Collaboro to model the example discussed previously. The figure is divided in several parts according to the collaboration steps enumerated previously. For the sake of clarity, references to User metaclass instances have been represented as

string-based attributes and the `rationale` attribute is not shown. The figure shows the collaboration as an instance of the Collaboro metamodel. At the tool level, we offer a user-friendly interface enabling all kinds of users to easily contribute and vote during the discussion process while the tool updates the collaboration model behind the scenes in response to the user events.

Part 1 of Fig. 7 shows the collaboration model just after *End-User 1* makes the request. It includes a new `Proposal` instance that is voted positively by the rest of the users and therefore accepted (see part 2). Then, a new solution is proposed by *Developer 1* (see part 3), which involves enriching the `Conveyor` metaclass with a float attribute in addition to define the concrete syntax. However, this solution is not accepted by all the community members: *End-User 1* does not agree and explains his disagreement (see part 4). Since the comment is accepted (see part 5), *Developer 1* updates the solution to incorporate the community recommendations (see part 6). Note that the elements describing the model changes in parts 3 and 6 are mutually exclusive. Moreover, the `included` attribute of the `Comment` element in part 4 will be activated as a consequence of the solution update. Once everybody agrees on the improved solution, it is selected as the final one for the proposal (see the `selected` reference).

Now the development team can modify the DSML tooling knowing that the community needs such change and agrees on how it must be done. Moreover, the rationale of the change will be tracked by the collaboration model (from which an explanation in natural language could be generated, if needed), which will allow community members to know why the `Conveyor` metaclass was changed.

Metric-based recommender

When developing DSMLs, several quality issues regarding the abstract and concrete syntaxes can be overlooked during the collaboration. While developers are maybe the main responsables for checking that the language is being developed properly, it is important to note that these issues may arise from both developers (e.g., they can forget defining how some concepts are represented in the notation) and end-users (e.g., they may miss that the notation is becoming too complicated for them to later being able to manage complex models). We propose to help both developers and end-users to develop better DSMLs by means of a recommender engine which checks the language under development to spot possible issues and improvements.

The recommender applies a set of metrics on the DSML to check its quality, in particular, to ensure that the resulting language is expressive, effective and accurate enough to be well-understood by the end-users. Metrics can target both the abstract and concrete syntaxes of a DSML. Concrete syntax metrics can in turn target either textual or graphical syntaxes. While several metrics for abstract and textual concrete syntaxes have been devised in previous works (*Cho & Gray, 2011; Aguilera, Gómez & Olivé, 2012; Power & Malloy, 2004; Črepinšek et al., 2010*), the definition and implementation of metrics for graphical concrete syntaxes is still an emerging working area. Thus, in this work, we explore how metrics for abstract and concrete

syntaxes can be implemented in our approach, but we mainly focus on those ones regarding graphical concrete syntaxes.

Abstract syntax metrics

The abstract syntax of a DSML is defined by a metamodel, as commented before. While the identification of proper DSML constructs (i.e., concepts and relationships) usually relies on the domain experts, identifying and solving design issues (e.g., creating hierarchies to promote extensibility or identifying patterns such as factoring attributes) is normally performed by the developers. Thus, to provide consistent solutions for recurring metamodel design issues, some metrics applied to abstract syntax metamodels may offer key insights on its quality.

There are currently several works providing a set of metrics for metamodels as well as for UML class diagrams that can be applied in this context (e.g., [Cho & Gray, 2011](#); [Aguilera, Gómez & Olivé, 2012](#)). As a proof of concept to evaluate the abstract syntax of DSMLs in our approach, we implemented a couple of metrics that validate hierarchical structures in metamodels (inspired by [Aguilera, Gómez & Olivé \(2012\)](#)). Thus, we consider that such structures are invalid whether either there is only one derived class or whether an inheritance is redundant (i.e., already covered by a chain of inheritance). As our approach relies on Ecore, other metrics defined for this metamodeling language could be easily plugged in by using the extension mechanism provided, as we will show afterwards.

Concrete syntax metrics

The concrete syntax of a DSML can be textual or graphical (or hybrid). As textual DSMLs are usually defined by means of a grammar-based approach, which is also the case for GPLs, existing support for evaluating the quality of GPLs could be applied (e.g., [Power & Malloy \(2004\)](#) and [Črepinšek et al. \(2010\)](#)). Apart from this GPL-related support, the current support to assess the quality of the concrete syntaxes in the DSML field is pretty limited. Thus, in this paper, we apply a unifying approach to check the quality of any DSML concrete syntax (i.e., textual and/or graphical).

With this purpose, we employ the set of metrics based on the cognitive dimensions framework ([Green, 1989](#)), later formalized by [Moody \(2009\)](#), where metrics are presented according to nine principles, namely: cognitive integration, cognitive fit, manageable complexity, perceptual discriminability, semiotic clarity, dual coding, graphic economy, visual expressiveness and semantic transparency. Several works have applied them to specific DSMLs (e.g., [Genon, Heymans & Amyot \(2011b\)](#) or [Le Pallec & Dupuy-Chessa \(2013\)](#)). Nevertheless, none of them has tried to implement such metrics in a way that can be applied generically to any DSML. As Collaboro provides the required infrastructure to represent concrete syntax at a technology-agnostic level, we propose to define a set of DSML metrics adapted from Moody's principles for designing cognitively effective notations. In the following, we present how we addressed five of the nine principles to be applied to our metamodels, and we justify why the rest of the principles were discarded.

Semiotic clarity. This principle refers to the need of having a one-to-one correspondence between notation symbols and their corresponding concepts, thus maximizing precision and promoting expressiveness. We can identify four metrics according to the possible situations that could appear: (1) *symbol deficit*, when a concept is not represented by a notation symbol (sometimes this situation could be evaluated positively as to avoid having too many symbols and losing visual expressiveness); (2) *symbol excess*, when a notation symbol does not represent any concept; (3) *symbol redundancy*, when multiple notation symbols can be used to represent a single concept; and (4) *symbol overload*, when multiple concepts are represented by the same notation symbol.

In Collaboro, these metrics can be computed by analyzing the mapping between the abstract syntax elements and the notation model elements of the DSML (i.e., analyzing the maps reference in the `ConcreteSyntaxElement` element). On the one hand, the analysis of the abstract syntax consists on a kind of flattening process where all the concepts are enriched to include the attributes and references inherited from their ancestors. The aim is to identify the DSMLs elements (i.e., concept, attribute or reference) for which a concrete syntax element has to be defined. On the other hand, the analysis of the concrete syntax focuses on the discovery of symbols. When a symbol uses multiples graphical elements to be represented (e.g., using nested `Composite` elements or `SyntaxOf` elements), they are aggregated. The result of this analysis is stored in a map that links every abstract syntax element with the corresponding concrete syntax element, thus facilitating the calculation of the previous metrics. This map will be also used in the computation of the remainder metrics.

Visual Expressiveness. This principle refers to the number of visual variables used in the notation of a DSML. Visual variables define the dimensions that can be used to create notation symbols (e.g., shape, size, color, etc.). Thus, to promote its visual expressiveness, a language should use the full range and capacities of visual variables.

To assess this principle, we define a metric which analyzes how visual variables are used in a DSML. The metric leverages on the previous map data structure and enriches it to include the main visual variables used in each symbol. According to the current support for visual variables of the notation metamodel (recall `GraphicalElement` metaclass attributes), these variables include: size (`height` and `width` attributes), color (`fill` and `stroke` attributes) and shape (subclasses of `GraphicalElement` metaclass). The metric checks the range of visual variables used in the symbols of the DSML and notifies the community when the notation should use more visual variables and/or more values of a specific visual variable to cover the full range.

Graphic Economy. This principle states that the number of notation symbols should be cognitively manageable. Note that there is not an objective rule to measure the complexity of notation elements (e.g., expert users may cognitively manage more symbols than a novice). There is the *six symbol rule* (Miller, 1956) which states that there should be no more than six notation symbols if only a single visual variable is used. We therefore devised a metric based on this rule to assess the level of graphic economy in a DSML.

Perceptual Discriminality. This principle states that different symbols should be clearly distinguishable from each other. Discriminality is primarily determined by the visual distance between symbols, that is, the number of visual variables on which they differ and the size of these differences. This principle also states that every symbol should have at least one unique value for each visual variable used (e.g., unique colors for each symbol). Thus, to assess the perceptual discriminability, we define a metric which also relies on the previous map data structure, compares each pair of symbols and calculates the visual distance between them according to the supported visual variables (i.e., number of different visual variables per pair of symbols). By default, the metric notifies the community when the average distance is lower than one, but it can be parameterized.

Dual Coding. This principle suggests that using text and graphics together conveys the information in a more effective way. Textual encoding should be then used in addition of graphical encoding to improve understanding. However, textual encoding should not be the only way to distinguish between symbols. We defined a metric that checks whether each symbol uses text and graphics elements, thus promoting dual coding. To this aim, we leverage on our notation metamodel, which allows to attach textual elements to symbols by employing `Label` elements that contain `TextualElement` elements. This metric notifies the community when more than a half of the symbols are not using both text and graphics.

The remaining four Moody's principles were not addressed due to the reasons described below.




Semiotic Transparency. This principle states that a notation symbol should suggest its meaning. This principle is difficult to evaluate as it relies on many parameters such as context and good practices in the specific domain. Furthermore, as the meaning of a representation is subjective, an automatic verification of this principle would be difficult to reach.

Complexity Management. This principle refers to the ability of the notation to represent information without overloading the human mind (e.g., providing hierarchical notations). Although this could be addressed in the notation model by providing mechanisms for modularization and hierarchical structuring, we believe that assessing this principle strongly depends on the profile and background of the DSML end-users and it is therefore hard to measure.

Cognitive Integration. This principle states that the visual notation should include explicit mechanisms to support integration of information from different diagrams. In this sense, this principle refers to the results of composing different DSMLs, which is not an scenario targeted by our approach.

Cognitive Fit. This principle promotes the fact that different representations of information are suitable for different tasks and audiences (e.g., providing different concrete syntaxes for the same abstract syntax). Like in the complexity management principle, assessing the cognitive fit of the notations of a DSML is directly related to the expertise of the different communities using the language, which is hard to measure with an automatic evaluation.

Table 1 Example of Visual Expressiveness and Perceptual Discriminality for the *Baggage Claim* DSML.

		 B	 C	VE
Shape	Polygon	Rectangle	Line	3/5
Size	H:5	H:5	H:1	2
	W:9	W:9	W:12	
Color	Fill:White	Fill:Black	Fill:White	2/49
	Stroke:Black	Stroke:Black	Stroke:Black	
PD	Visual Distance	Visual Distance	Visual Distance	
	B:2	A:2	A:2	
	C:2	C:3	B:3	

Note:

VE, Visual Expressiveness; PD, Perceptual Discriminability.

Recommending changes

The results of the previously shown metrics provide the community developing a DSML with an important feedback to address potential improvements. In Collaboro, the DSML development process incorporates a recommender that plays the role of a user in the collaboration process. This “recommender user” can check the different versions of the DSML under development according to the previously shown metrics and propose new changes identifying the weak points to be discussed in the community. Metrics can be deactivated if wished and can be given different relevance values that can also be used to sum up the results to calculate a general value assessing the quality of the DSML under development.

Example

In this section, we will show an example of the metrics regarding visual expressiveness and perceptual discriminativity for the *Baggage Claim* DSML. For the sake of illustration purposes, we describe these metrics on an alternative graphical syntax to the DSML, where the *Flight* concept is represented as a polygon with the shape of an airplane, the *Conveyor* concept is represented as a black filled-rectangle and the *claims* reference is represented as a line. The computation of these metrics are specially tailored to the visual variables supported by our notation metamodel. [Table 1](#) illustrates how these two metrics are calculated. As can be seen, visual expressiveness results assess the number of different values used for each visual variable. Thus, there are three out of five values for the shape dimension, two different values for the size dimension and two different values for the color dimension. On the other hand, the visual distance is calculated for each pair of symbols and measures the number of different visual variables between them. For instance, the black-filled rectangle differs in two visual variables (i.e., color and shape) with the airplane polygon; and all the supported visual variables with regard to the line. These results reveal a good visual expressiveness (good values for shape and size visual variables while the color range is appropriate for the number of symbols) and perceptual discriminability (visual distance is in average more than 2, where the highest value is 3) therefore validating this graphical notation proposal.

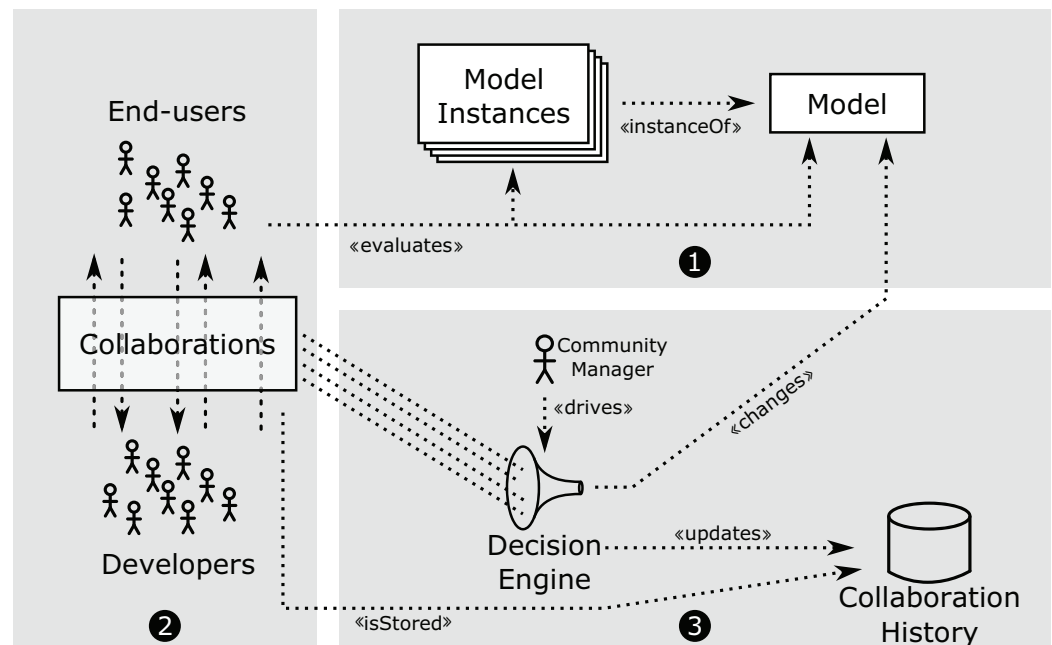


Figure 8 Collaborative modeling.

COLLABORATIVE MODELING

In this section we will show how our approach could be easily adapted to support collaborative modeling. This adaptation is depicted in Fig. 8. Unlike the Fig. 2, where we illustrated the process for the collaborative development of DSMLs, in this case the community evaluates and discuss changes about the model being developed and not the metamodel. Thus, once there is a first version of the model and a set of examples (step 1), the community discusses how to improve the models (step 2). The discussion arises changes and improvements, that have to be voted and eventually incorporated in the model (step 3). Discussion and decisions are recorded (see *Collaboration History*), thus keeping track of the modifications performed in the model.

To support this development process, the modifications to perform in the original Collaboro metamodel are very small. Figure 9 shows the new metamodel to track the collaboration. As can be seen, the only changed element is the *SyntaxElement*, which now has to refer to the main (i.e. root) metaclass of the modeling language being used to link the model elements with their metamodel definition. For instance, by default, the Figure includes the element *NamedElement* from UML, thus illustrating how Collaboro could be used for the collaborative development of UML models. Other languages could be supported following this same approach.

TOOL SUPPORT

Since the very first implementation of Collaboro was released, the tool support has evolved to integrate the full set of features described in this paper². The new architecture of the developed tool is illustrated in Fig. 10. The main functionalities of our approach are implemented by the backend (see *Collaboro Backend*), which includes specific

² The tool is available at <http://som-research.github.io/collaboro>

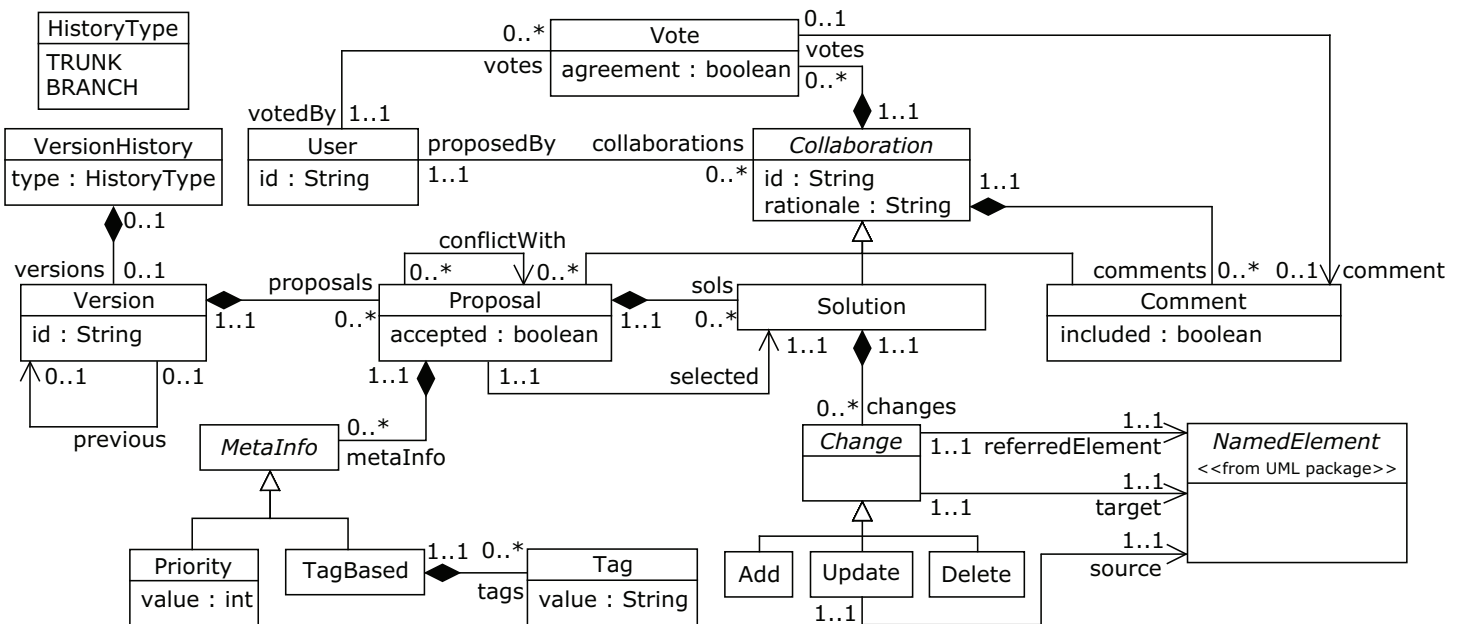


Figure 9 Core elements of the adapted Collaboro metamodel.

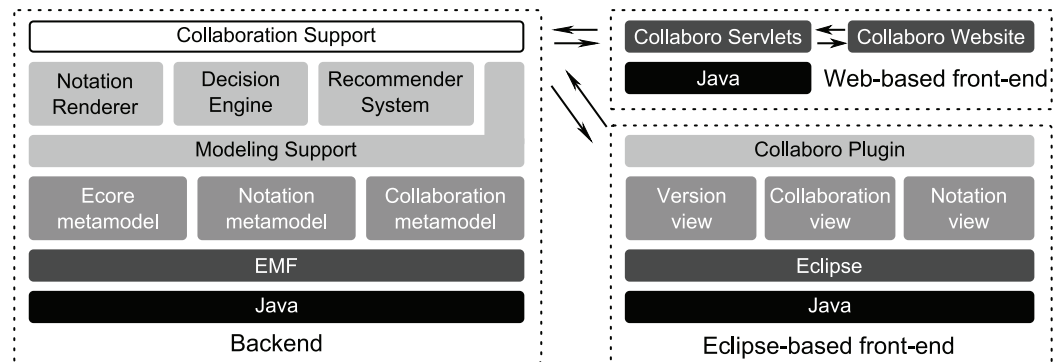


Figure 10 Architecture of Collaboro tool support.

components for modeling both the DSML elements and the collaborations (see *Modeling Support*), rendering the notation examples (see *Notation Renderer*) making decisions (see *Decision Engine*), and recommending changes (see *Recommender System*). As front-end for Collaboro, we have developed two alternatives: (1) a web-based front-end, which gives access to the collaboration infrastructure from any web browser; and (2) an Eclipse-based front-end, which extends the platform with views and editors facilitating the collaboration. Next, we describe in detail each component of this architecture.

Collaboro backend

This component provides the basic functionality to develop collaborative DSMLs as explained in this paper. Collaboro relies on the EMF framework (*Steinberg et al., 2008*) (the standard de facto modeling framework nowadays) to manage the models required during the development process. In the following, we describe the main elements of this component.

Modeling support

Collaboro provides support for managing models representing the abstract and concrete syntaxes, and the collaboration models. We implemented the metamodels described in previous sections as Ecore models (the metamodeling language used in EMF) and provided the required API. To support concurrent collaboration the tool can be configured to store the models in a CDO (<http://www.eclipse.org/cdo>) model repository.

Notation renderer

The tool incorporates a generator which automatically creates the graphical/textual representation of the DSML example models. This component enables the lightweight creation of SVG ([SVG, 2011](#)) images from notation models to help users “see” how the notation they are discussing will look like when used to define models with that DSML. The generator analyzes each example model element, locates its abstract/concrete syntax elements and interprets the concrete syntax definition to render its textual/graphical representation. `GraphicalElement` and `TextualElement` concrete syntax elements indicate the graphical or textual representation to be applied (e.g., a figure or a text field), while `Composite` and `SyntaxOf` concrete syntax elements are used for layout and composite elements. Regarding the layout of the generated graphical/textual representation, on the one hand, a block-based notation is automatically applied for textual languages, where each new `Composite` concrete syntax element (and its contained elements) is indented. On the other hand, graphical notations are rendered following a horizontal layout, thus elements are arranged from left to right as the example model is analyzed. Note that symbols acting as connectors between concepts are detected by means of the `maps` reference in `ConcreteSyntaxElement`, thus allowing the renderer to know what concepts (and their corresponding symbols) are used as source/target elements.

Decision engine

This component is responsible for updating the dynamic part of the collaboration models (recall the support for votes and decisions). The current support of the tool implements a total agreement strategy to infer community agreements from the voting information of the collaboration models.

Recommender system

This component provides the required infrastructure to calculate both abstract and concrete syntaxes metrics in order to ensure their quality. The recommender is executed on demand by the community manager. The current support of the tool implements metrics to evaluate the quality of concrete graphical syntax issues.

New metrics can be plugged in by extending the Java elements presented in [Fig. 11](#). The entry point is the `MetricFactory` class, which is created for each DSML and is responsible for providing the list of available metrics. `Metrics` have a name, a description, a dimension (e.g., each Moody’s principle), an activation, a priority level and an acceptance ratio. The acceptance ratio allows specifying the maximum number of elements of syntaxes that can be wrong (e.g., not conforming to the

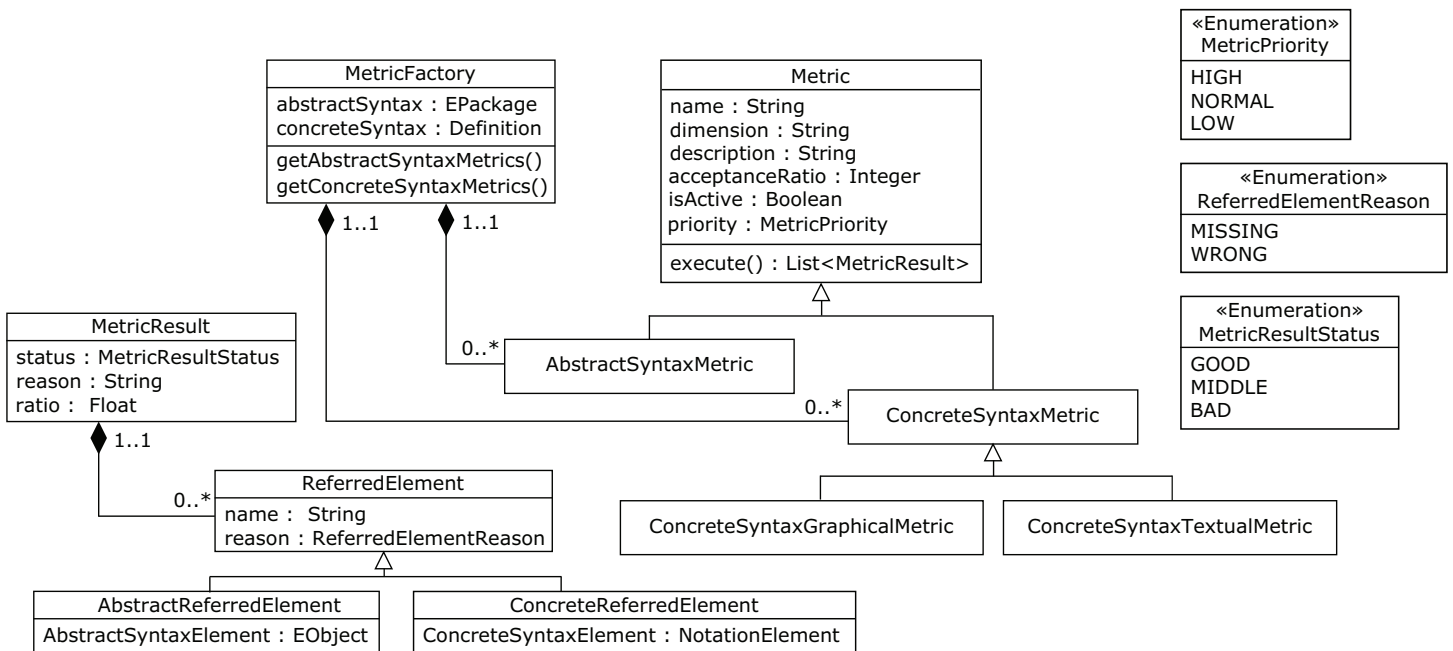
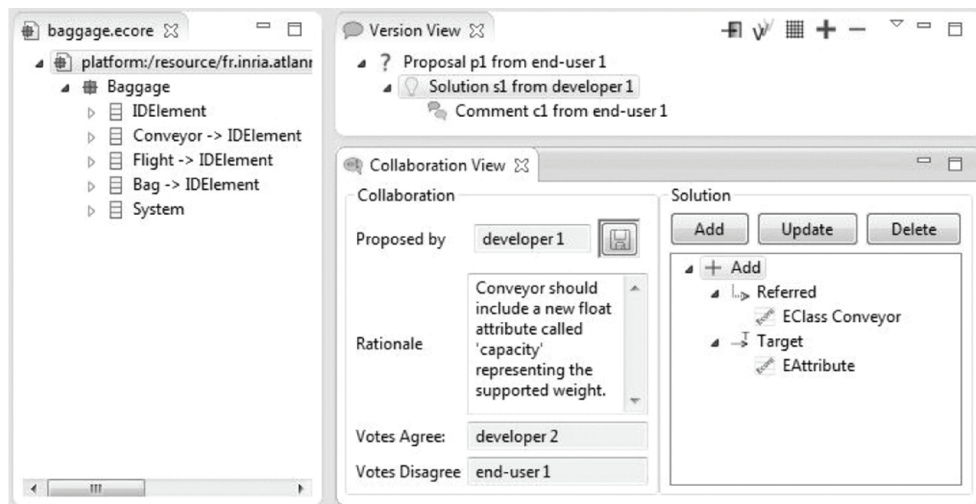


Figure 11 Core elements of the recommender engine.

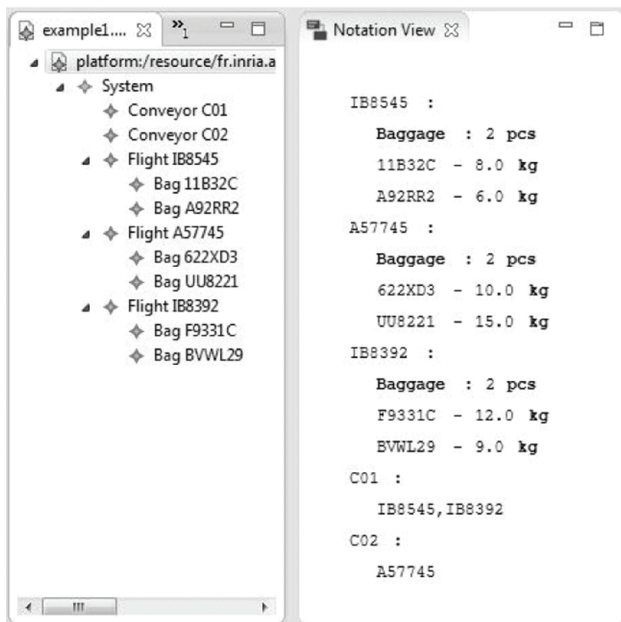
metric). Every metric also includes an `execute()` method for the recommender to compute them. This function returns a list of `MetricResults` describing the assessment of the metric. Metric results includes a status (i.e., measured in three levels), a reason describing the assessment in natural language and a ratio of fulfillment for the metric. Metric results also include a list of `ReferredElements` pointing to those abstract or concrete syntaxes elements not conforming with the metrics being calculated, thus helping developers to spot the DSML elements not satisfying each metric (if any).

Eclipse plugin

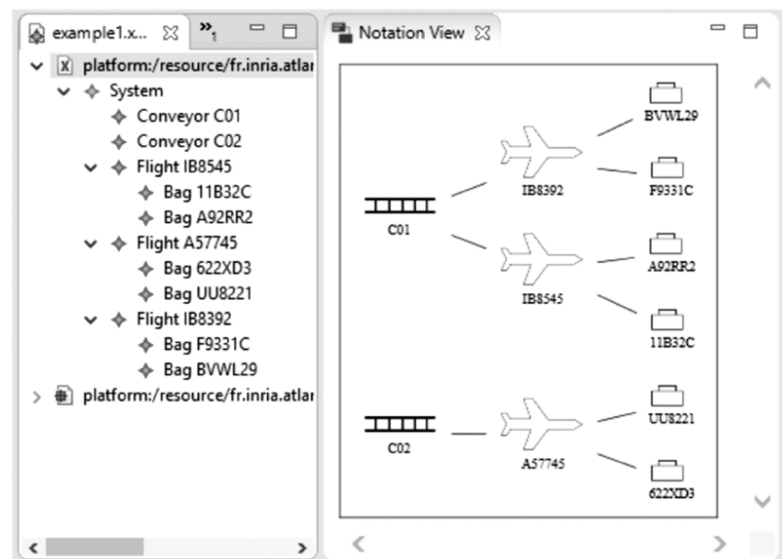
We have developed an Eclipse plugin implementing the Collaboro process and DSML. The plugin provides a set of new Eclipse views and editors to facilitate the collaboration, which can be considered a kind of concrete syntax of Collaboro itself for non-expert users. Via these editors users can propose changes (to add both new abstract syntax and concrete syntax definitions to existing abstract elements) on the Collaboro model that, once accepted, will update the abstract and concrete syntax models and link them together according to the selected solution. Figure 12A includes a snapshot of the environment showing the last step of the collaboration described in previous sections. In particular, the *Version view* lists the collaboration elements (i.e., proposals, solutions and comments) of the current version of the collaboration model. The *Collaboration View* shows the detailed information of the selected collaboration element in the Version view and a tree-based editor to indicate the changes to discuss for that element, as shown in Fig. 12A. Finally, the *Notation view* uses the notation



(a)



(b)



(c)

Figure 12 (A) Snapshot of the Collaboro Eclipse plugin. Collaboro Eclipse plugin with the *Notation view* rendering the (B) textual and (C) graphical concrete syntaxes for a model.

generator to render a full example model of the language. For instance, the Notation view in Figs. 12B and 12C show the notation (i.e., in textual and graphical forms, respectively) for an example model, which allowed detecting the missing attribute regarding the conveyor capacity.

Web-based front-end

The developed web support includes two components: (1) the server-side part, which offers a set of services to access to the main functionalities of Collaboro; and the

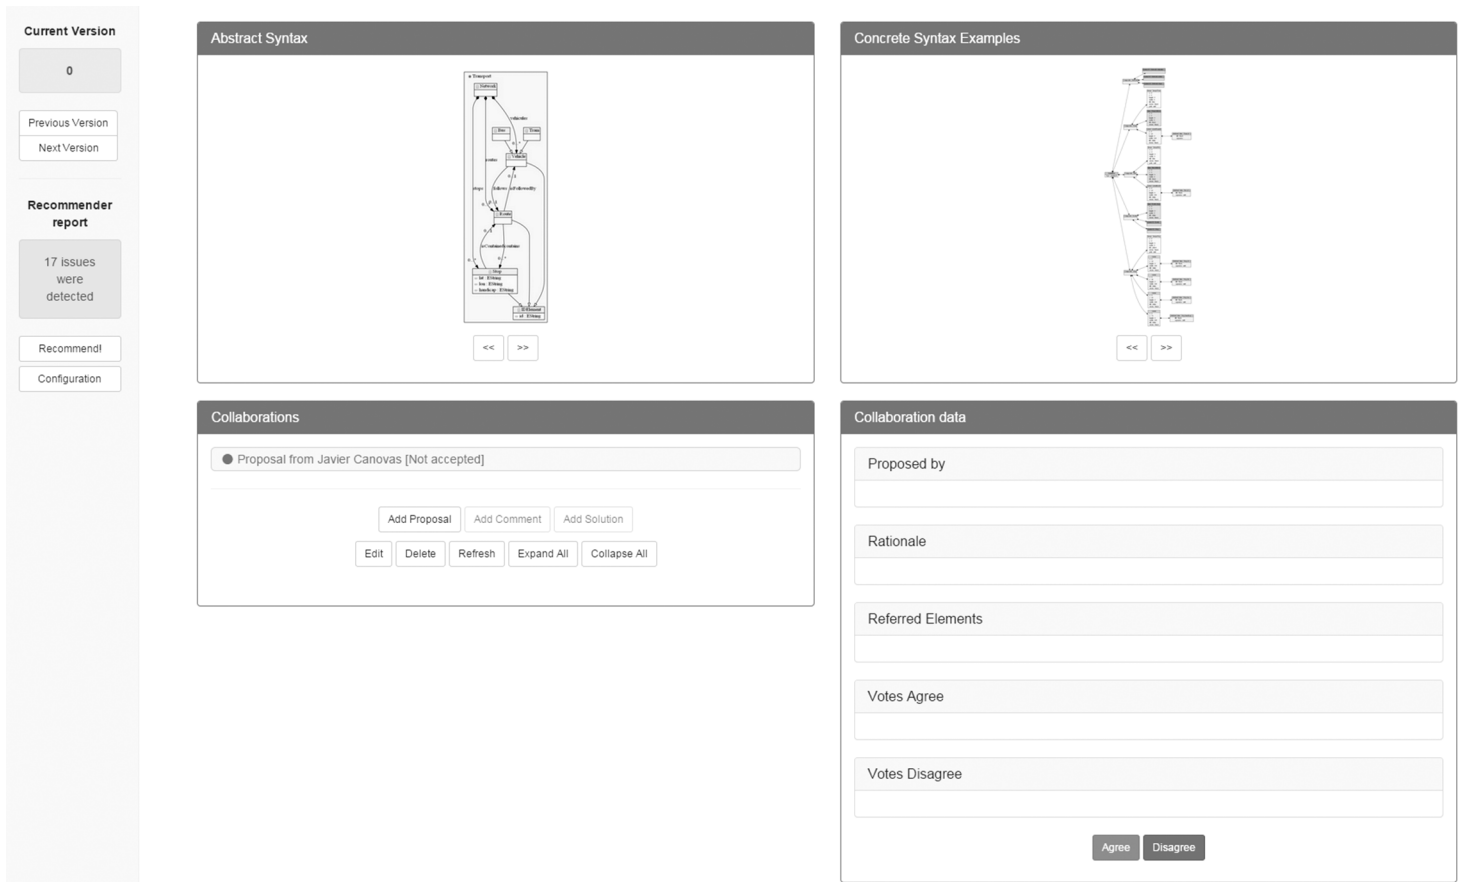


Figure 13 Snapshot of the Collaboro web client.

client-side part, which allows both end-users and developers to take part of the DSML development process from their browsers. The server-side component has been developed as a Java web application which uses a set of Servlets providing the required services. On the other hand, the client-side component has been developed as an AngularJS-enabled website.

Figures 13 and 14 show two snapshots of the developed website. As can be seen in Fig. 13, the website follows an arrangement similar to that one used in the Eclipse plugin. Thus, on top, there are two sections showing the current status of (1) the abstract syntax of the DSML on the left and (2) several model examples rendered with the concrete syntax definition of the DSML on the right (both sections are zoom-enabled). These sections include several pictures that can be navigated by the user (e.g., it is possible to evaluate the different example models rendered). At the bottom of the website, there are two more sections aimed at managing the collaborations, in particular, (1) a tree including all the collaboration elements on the left and (2) a details view on the right which shows the information of a collaboration once it is selected in the tree. Furthermore, the tree view also includes buttons to create, edit and delete collaborations.

Recommender Configuration

Graphic Economy	Active
The number of different graphical symbols should be cognitively manageable	
1 issue/s detected in the current version	
Symbol Deficit	Active
Symbol deficit occurs when there are abstract concepts that are not represented by any concrete symbol	
1 issue/s detected in the current version	
Symbol Excess	Active
Symbol excess occurs when a concrete symbols does not correspond to any abstract concept	
The current version fulfills the recommendations for this metric	
Symbol Overload	Active
Symbol overload occurs when a single concrete symbol can represent multiple abstract concepts	
The current version fulfills the recommendations for this metric	
Symbol Redundancy	Active
Symbol redundancy occurs when multiple concrete symbols can be used to represent a single abstract concept	
The current version fulfills the recommendations for this metric	
Perceptual Discriminability	Active
Different symbols should be clearly distinguishable from each other	
3 issue/s detected in the current version	

Figure 14 Snapshot of a subset of supported metrics.

The website also includes a left menu bar which allows the user to navigate through the different versions of the DSML as well as indicate some information about the recommender system status. Additionally, the user can quickly see the number of issues detected by the recommender, configure the metrics (see Fig. 14) that have to be executed and perform the metric execution to incorporate the change proposals into the collaboration.

APPLICATION SCENARIOS

In this section, we report the use of Collaboro in two types of scenarios: (1) the creation of new DSMLs, based on two different case studies; and (2) the extension of existing DSMLs, where we describe our experience in one case study. We also mention some lessons learned in the process.

Developing new DSMLs

We used Collaboro in the creation of two new DSMLs: (1) a textual DSML to define workflows and (2) three metamodels to represent code hosting platforms in the context of a modernization process. We explain each case in the following.

Creating a textual DSML

Collaboro was used in the development of a new DSML for MoDisco (<http://eclipse.org/modisco>), an Eclipse project aimed at defining a group of tools for Model-Driven Reverse Engineering (MDRE) processes. The goal of this new DSML is to facilitate the development of MDRE workflows that chain several atomic reverse engineering tasks to extract the model/s of a running system. At the moment, the only way to define a MDRE workflow is by using an interactive wizard. MoDisco users have been asking for a specific language to do the same in a more direct way, i.e., without having to go through the wizard.

Some years ago an initial attempt to create such language was finally abandoned but, to simplify the case study, we reused the metamodel that was proposed at the time to kickstart the process. Five researchers of the team followed our collaborative process to complete/improve the abstract syntax of the DSML and create from scratch a concrete syntax for it. Two of the members were part of the MoDisco development team so they took the role of developers in the process while the other three were only users of MoDisco so they adopted the role of end-users in the process. One of the members was in a different country during the collaboration so only asynchronous communication was possible.

The collaboration took two weeks and resulted in two new versions of the MDRE workflow language released. The first version was mainly focused on the polishment of the abstract syntax whereas the second one paid more attention to the concrete syntax (this was not enforced by us but it came out naturally). The collaboration regarding the abstract syntax involved changes in concepts and reference cardinalities, while regarding the concrete syntax, the community chose to go for a textual-based notation and mainly discussed around the best keywords or style to be used for that.

Defining metamodels

We have also applied Collaboro for defining a set of metamodels used in a model-driven re-engineering process (i.e., only the abstract syntax of the DSML was part of the experiment since the models were to be automatically created during the reverse engineering process). In particular, the process was intended to provide support for migrating Google Code to GitHub projects, thus requiring the corresponding Platform-Specific Models (PSM) metamodels for both platforms, plus a Platform-Independent

Models (PIM) metamodel to represent such projects at high level of abstraction (following the typical terminology defined by the Model-Driven Architecture (MDA) approach from the OMG (*Object Management Group, 2014a*)). As the developers were distributed across different geographical locations, we decided to use Collaboro to create the PSM and PIM metamodels required.

Six researchers geographically dispersed (i.e., the participants were part of three research groups, making three groups composed of 3, 2 and 1 researchers) collaborated in the definition of the metamodels. To kickstart the collaboration, one of the teams created a first version of each metamodel. As the collaboration was focused on defining only the abstract syntax of the language, there was no need for creating a notation model, and therefore the set of examples were rendered following a class-like diagram style. The collaboration took three weeks and resulted in two versions for each one of the PSM metamodels and only one version for the PIM metamodel since there the agreement was faster.

Extending an existing DSML

More recently, we were contacted by a community member of the Architecture Analysis & Design Language (AADL) (<http://www.aadl.info>), and one of the lead developers in charge of defining an extension to such language. AADL is an architecture description language used in the embedded and real-time systems field. It is a textual DSML with large abstract and concrete syntaxes. The abstract syntax contains more than 270 concepts and the concrete syntax is composed of more than 153 elements (including keywords and tokens). The language was being extended to incorporate support for behavior specification. This extension, called AADL Behavior Annex (AADL-BA) (<http://penelope.enst.fr/aadl/wiki/Projects#AADL-BA-FrontEn>), was being defined as a plugin enriching both the abstract and concrete syntaxes.

At the time, the definition of the extension was taken care by a standardization committee open to new contributions. Change proposals were informally managed by in-person voting (i.e., raising hands in a meeting) or online ballots. Later, the documentation of the change proposal was spread out in a document, presentation or online wiki documentation. As explained to us by this lead developer, this process made tracking modifications very hard in the language as well as the corresponding argumentations, and he proposed to use Collaboro to manage the development of the extension for AADL. As a first step, we created a fake AADL project so that this person could play around with the tool and assess its usefulness for the AADL community. The feedback was that the tool would be very useful for the project at hand if we were able to deal with some technical challenges linked to the current setting used by the project so far. In particular, to be able to use Collaboro for managing the AADL-BA language definition process we needed to import: (1) previous discussions stored in the wiki-based platform and (2) the current concrete syntaxes of the AADL and AADL-BA language defined in Xtext and ANTLR respectively (the abstract syntax was already defined as an EMF model so it could be directly imported into Collaboro). It was also clear that to simplify the use of the tool, we had to provide a web interface since it would be too complex for the members of

the AADL community to install an Eclipse environment just for the purpose of discussing around language issues.

In the end, time constraints prevented us to test the tool with AADL community at large (the AADL-BA committee meets at fixed dates and we did create a web-based interface but could not get a new version of the tool with all the scripts required to import the legacy data on time), but the private iterations with the AADL-BA developer and his validation and positive feedback helped us a lot to improve Collaboro and learn more about the challenges of using Collaboro as a part of an ongoing language development effort. We are still in contact with this community and we will see if we can complete the test in the future or reach out other similar standardization committees.

Lessons learned

The development of the previous case studies provided us with some useful insights on the Collaboro process that since then have been integrated in our approach. For instance, in the first and second case studies, it turned out that conflicting proposals were frequent and therefore we added a conflicting relationship information explicitly in the collaboration metamodel so that once one of them was accepted we could automatically shut down the related ones. We also noted an intensive use of comments (easier to add) in comparison with proposals and solutions. This fact together with the discussions on what should constitute a new version and when to end the discussions (e.g., what if there was unanimity but not everybody had voted, should we wait for that person? for how long?) helped us to realize the importance of an explicit community manager role in charge of making sure the collaboration is always fluid and there are no bottlenecks or deadlocks.

During the development of the three case studies, concurrent access to the models turned out to be a must as well since most of the time collaborations overlapped at some point. The experience gathered during the development of the first case study, where the collaboration was performed only in the Eclipse-based plugin, and later the requirements of the second and third case studies allowed us to provide a second front-end for the approach based on a web-client. Thus, the web-enabled support was crucial to allow all the developers to contribute and visualize how the metamodels evolved during the collaboration.

In all the case studies the notation view allowed the participants to quickly validate the concrete syntax. This is specially important since for non-technical users it is easier to discuss at the concrete syntax level than at the abstract level.

The only common complaint we got was regarding the limited support for voting (mainly raised in the first case study but also raised in the others), where participants reported that they would have preferred more options instead of just a boolean yes/no option. Note that this would have a non negligible impact on the decision algorithms that would need to be adapted to consider the new voting options. We plan to incorporate extra support to define how to make decisions, in a similar way as proposed in [Cánovas Izquierdo & Cabot \(2015\)](#).

RELATED WORK

End-user involvement is a core feature of several software development methods (such as agile-based ones). The concept of *community-driven development* of a software product was introduced by *Hess, Offenbergh & Pipek (2008)* and other authors have studied this collaboration as part of the requirement elicitation (*Mylopoulos, Chung & Yu, 1999*), ontology development (*Leenheer, 2009; Siorpaes, 2007*) and modeling phases of the software (*Hildenbrand et al., 2008; Lanubile et al., 2010; Whitehead, 2007; Rittgen, 2008*), but neither of them focuses on the DSML language design process nor they present the collaboration as a process of discussion, voting and argumentation from the beginning to the end of the language development process. End-user participation is also the core of user-centered design (*Norman & Draper, 1986*), initially focused on the design of user interfaces but lately applied to other domains (e.g., agile methodologies (*Hussain, Slany & Holzinger, 2009*) or web development (*Troyer & Leune, 1998*)). Again, none of these approaches can be directly applied to the specification of a DSML. Nevertheless, ideas from these papers have indeed influenced the Collaboro process.

Regarding specific approaches around collaboration in DSML development, some works propose to derive a first DSML definition by means of user demonstrations (*Cho, Gray & Syriani, 2012; Kuhrmann, 2011; Sánchez Cuadrado, de Lara & Guerra, 2012; López-Fernández et al., 2013*) or grammar inference techniques (*Javed et al., 2008; Liu et al., 2012*), where example models are analyzed to derive the metamodel of the language. However, these approaches do not include any discussion phase nor validation of the generated metamodel with the end-users. In this sense, our approaches could complement each other, theirs could be used to create an initial metamodel from which to trigger the refinement process based on the discussions among the different users (*Cánovas Izquierdo et al., 2013*).

Subsets of our proposal can also be linked to: i) specific tools for model versioning (e.g., AMOR repository (<http://www.modelversioning.org>) and *Altmanninger, Seidl & Wimmer (2009)*) that have already proposed a taxonomy of metamodel changes, ii) online-collaboration (*Brosch et al., 2009; Gallardo, Bravo & Redondo, 2012*) promoting synchronous collaboration among developers, iii) metamodel-centric language definition approaches (*Scheidgen, 2008; Prinz, Scheidgen & Tveit, 2007*) where the concrete syntax is considered at the same level as the abstract one and iv) collaboration protocols (*Gallardo et al., 2013*). In all cases, Collaboro extends the contributions of those tools with explicit collaboration and justification constructs, and provides as well the possibility of offline collaborations and a more formal representation of the interactions (e.g., voting system, explicit argumentation and rationale, traceability). The agreed DSML definition at the end of the Collaboro process could be then the input of the complete DSML modeling environment aimed by some of the tools mentioned above.

Regarding the recommender engine and the calculation of metrics for DSMLs, we can identify works centered on assessing the quality of both the abstract and concrete syntaxes, and the main features of the language (e.g., reusability, integrability or compatibility). There are several works providing metrics to check the quality in

metamodels (*Cho & Gray, 2011; Aguilera, Gómez & Olivé, 2012*) and in the notation used for textual DSMLs (*Power & Malloy, 2004; Črepinšek et al., 2010*). With regard to graphical DSMLs, Moody's principles (*Moody, 2009*) have emerged as the predominant theoretical paradigm. Originally based on the cognitive dimensions framework (*Blackwell et al., 2001; Green, 1989; Green & Petre, 1996*), Moody's principles address their theoretical and practical limitations. While these principles provide a framework to evaluate visual notations, other works have put them into practice by analyzing DSMLs (*Genon, Amyot & Heymans, 2011a; Genon, Heymans & Amyot, 2011b; Moody & van Hillegersberg, 2009; Le Pallec & Dupuy-Chessa, 2013*) or complement the use of Moody's principles with polls (*Figl et al., 2010*) also, thus allowing end-user feedback and involvement during the design process of a visual notation. However, the previous works are usually centered to specific DSMLs and do not provide mechanisms to be calculated to any DSML as our approach addresses. Other works such as (*Kahraman & Bilgen, 2013*) propose an evaluation framework focused on language features and therefore not particularly analyzing the quality from an end-user perspective. To the best of our knowledge, ours is the first proposal to generically assess the cognitive quality of DSMLs under development.

Finally, the representation of the collaboration rationale is related to the area of requirements negotiation, argumentation and justification approaches such as the work presented by *Jureta, Faulkner & Schobbens (2008)*. The decision algorithms proposed in those works could be integrated in our decision engine. Other decision engines such as CASLO (*Padrón, Dodero & Lanchas, 2005*) or HERMES (*Karacapilidis & Papadias, 2001*) could also be used.

CONCLUSIONS

We have presented Collaboro, a DSML to enable the participation of all members of a community in the specification of a new domain-specific language or in the creation of new models. Collaboro allows representing (and tracking) language change proposals, solutions and comments for both the abstract and concrete syntaxes of the language. This information can then be used to justify the design decisions taken during the definition or use of the modeling language. The approach provides two front-ends (i.e., Eclipse-based and web-based ones) to facilitate its usage and also incorporates a recommender system which checks the quality of the DSML under development.

Once the community reaches an agreement on the language features, our Collaboro model could be later used as input to language workbenches in order to automatically create the DSL tooling (i.e., editors, parsers, palettes, repositories, etc.) needed to start using the language in practice. For instance, it would be possible to automatically create the configuration files required for XText (for textual languages) or GMF (for graphical ones) from our notation and abstract syntax models.

As further work, we plan to extend our notation metamodel (and the renderer) to support richer concrete syntax definitions (e.g., incorporating the concept of *anchor* to specify how to represent the source/target connections for links in graphical languages).

These extensions should be defined as pluggable extensions to allow designers import them in the definition of languages where it is foreseen those additional concepts may be needed. We also find interesting to use our recommender system on existing (popular) languages as a way to assess the “quality” of such languages and, potentially, suggest changes to improve them. Furthermore, we would like to explore how to support the collaborative definition of the well-formed rules (e.g., OCL constraints) for the DSML under development. As these rules are normally expressed by using a (semi)formal textual language (like OCL), the challenge is how to discuss them in a way that non-technical experts can understand and participate. Finally, we are also exploring how to better encourage end-user participation (e.g., by applying gamification techniques) to make sure the process is as plural as possible. This could be tried as part of a new experiment in an educational setting at our institution (Universitat Oberta de Catalunya (UOC): www.uoc.edu) with the (virtual) students in our software engineering master degree.

ADDITIONAL INFORMATION AND DECLARATIONS

Funding

Javier Luis Cánovas Izquierdo benefited from two postdoctoral fellowships grants by Inria and IN3-UOC during the realization of this work. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Grant Disclosures

The following grant information was disclosed by the authors:

Inria and
IN3-UOC.

Competing Interests

The authors declare that they have no competing interests.

Author Contributions

- Javier Luis Cánovas Izquierdo conceived and designed the experiments, performed the experiments, analyzed the data, contributed reagents/materials/analysis tools, wrote the paper, prepared figures and/or tables, performed the computation work, reviewed drafts of the paper.
- Jordi Cabot conceived and designed the experiments, performed the experiments, analyzed the data, contributed reagents/materials/analysis tools, wrote the paper, prepared figures and/or tables, performed the computation work, reviewed drafts of the paper.

Data Deposition

The following information was supplied regarding data availability:

GitHub: <https://github.com/SOM-Research/collaboro>.

REFERENCES

- Aguilera D, Gómez C, Olivé A. 2012.** A method for the definition and treatment of conceptual schema quality issues. In: *International Conference on Conceptual Modeling*. Vol. 7632. Berlin, Heidelberg: Springer, 501–514.
- Altmanninger K, Seidl M, Wimmer M. 2009.** A survey on model versioning approaches. *International Journal of Web Information Systems* 5(3):271–304 DOI 10.1108/17440080910983556.
- Barišić A, Amaral V, Goulão M, Barroca B. 2012.** Evaluating the usability of domain-specific languages. In: *Formal and Practical Aspects of Domain-Specific Languages: Recent Developments*, 386–407. Available at <http://www.igi-global.com/chapter/evaluating-usability-domain-specific-languages/77793>.
- Blackwell AF, Britton C, Cox AL, Green TRG, Gurr CA, Kadoda GF, Kutar M, Loomes M, Nehaniv CL, Petre M, Roast C, Roe C, Wong A, Young RM. 2001.** Cognitive dimensions of notations: design tools for cognitive technology. In: *International Conference on Cognitive Technology, Coventry*, 325–341.
- Brosch P, Seidl M, Wieland K, Wimmer M, Langer P. 2009.** We can work it out: collaborative conflict resolution in model versioning. In: *European Conference on Computer Supported Cooperative Work*. London: Springer, 207–214.
- Cabot J, Wilson G. 2009.** Tools for teams: a survey of web-based software project portals. Available at <http://www.drdoobbs.com/tools/tools-for-teams-a-survey-of-web-based-so/220301068>.
- Cánovas Izquierdo JL, Cabot J. 2013.** Enabling the collaborative definition of DSMLs. In: *International Conference on Advanced Information Systems Engineering, Valencia*, 272–287.
- Cánovas Izquierdo JL, Cabot J. 2015.** Enabling the definition and enforcement of governance rules in open source systems. In: *International Conference on Software Engineering*. Piscataway: IEEE Press, 505–514.
- Cánovas Izquierdo JL, Cabot J, López-Fernández JJ, Sánchez Cuadrado J, Guerra E, de Lara J. 2013.** Engaging end-users in the collaborative development of domain-specific modelling languages. In: *International Conference on Cooperative Design, Visualization, and Engineering, Alcudia, Mallorca*, 101–110.
- Cho H, Gray J. 2011.** Design patterns for metamodels. In: *Conference on Systems, Programming, and Applications: Software for Humanity—Colocated Workshop*. New York: ACM, 25–32.
- Cho H, Gray J, Syriani E. 2012.** Creating visual domain-specific modeling languages from end-user demonstration. In: *International Workshop on Modeling in Software Engineering*. Piscataway: IEEE Press, 29–35.
- Črepinšek M, Kosar T, Mernik M, Cerville J, Forax R, Rouse G. 2010.** On automata and language based grammar metrics. *Computer Science and Information Systems* 7(2):309–329 DOI 10.2298/CSIS1002309C.
- Dullemond K, van Gameren B, van Solingen R. 2014.** Collaboration spaces for virtual software teams. *IEEE Software* 31(6):47–53 DOI 10.1109/MS.2014.105.
- Figl K, Derntl M, Rodríguez MC, Botturi L. 2010.** Cognitive effectiveness of visual instructional design languages. *Journal of Visual Languages & Computing* 21(6):359–373 DOI 10.1016/j.jvlc.2010.08.009.
- Gabriel P, Goulão M, Amaral V. 2010.** Do software languages engineers evaluate their languages? *Congreso Iberoamericano en Software Engineering, Cuenca*, 149–162.
- Gallardo J, Bravo C, Redondo MA. 2012.** A model-driven development method for collaborative modeling tools. *Journal of Network and Computer Applications* 35(3):1086–1105 DOI 10.1016/j.jnca.2011.12.009.

- Gallardo J, Bravo C, Redondo MA, de Lara J. 2013.** Modeling collaboration protocols for collaborative modeling tools: experiences and applications. *Journal of Visual Languages & Computing* 24(1):10–23 DOI 10.1016/j.jvlc.2012.10.006.
- Genon N, Amyot D, Heymans P. 2011a.** Analysing the cognitive effectiveness of the UCM visual notation. In: *International Workshop on System Analysis and Modeling, Oslo*, 221–240.
- Genon N, Heymans P, Amyot D. 2011b.** Analysing the cognitive effectiveness of the BPMN 2.0 visual notation. In: *International Conference on Software Language Engineering, Eindhoven*, 377–396.
- Green TRG. 1989.** Cognitive dimensions of notations. In: Sutcliffe A, Macaulay L, eds. *People and Computers*. Vol. 5. Cambridge: Cambridge University Press, 443–460.
- Green TRG, Petre M. 1996.** Usability analysis of visual programming environments: a cognitive dimensions framework. *Journal of Visual Languages & Computing* 7(2):131–174 DOI 10.1006/jvlc.1996.0009.
- Grundy JC, Hosking J, Li KN, Ali NM, Huh J, Li RL. 2013.** Generating domain-specific visual language tools from abstract visual specifications. *IEEE Transactions on Software Engineering* 39(4):487–515 DOI 10.1109/TSE.2012.33.
- Hatton L, van Genuchten M. 2012.** Early design decisions. *IEEE Software* 29(1):87–89 DOI 10.1109/MS.2012.5.
- Hess J, Offenbergs S, Pipek V. 2008.** Community driven development as participation? Involving user communities in a software design process. In: *Conference on Participatory Design*. Indianapolis: Indiana University, 31–40.
- Hildenbrand T, Rothlauf F, Geisser M, Heinzl A, Kude T. 2008.** Approaches to collaborative software development. In: *Conference on Complex, Intelligent and Software Intensive Systems*. Piscataway: IEEE, 523–528.
- Hussain Z, Slany W, Holzinger A. 2009.** Current state of agile user-centered design: a survey. In: *Symposium of the workgroup human-computer interaction and usability engineering of the austrian computer society—HCI and usability for e-Inclusion, Linz*. Vol. 5889. 416–427.
- Javed F, Mernik M, Gray J, Bryant BR. 2008.** MARS: a metamodel recovery system using grammar inference. *Information and Software Technology* 50(9–10):948–968 DOI 10.1016/j.infsof.2007.08.003.
- Jureta IJ, Faulkner S, Schobbens P-Y. 2008.** Clear justification of modeling decisions for goal-oriented requirements engineering. *Requirements Engineering* 13(2):87–115 DOI 10.1007/s00766-007-0056-y.
- Kahraman G, Bilgen S. 2013.** A framework for qualitative assessment of domain-specific languages. *Software & System Modeling* 14(4):1–22 DOI 10.1007/s10270-013-0387-8.
- Karacapilidis NI, Papadias D. 2001.** Computer supported argumentation and collaborative decision making: the HERMES system. *Information Systems* 26(4):259–277 DOI 10.1016/S0306-4379(01)00020-5.
- Kelly S, Pohjonen R. 2009.** Worst practices for domain-specific modeling. *IEEE Software* 26(4):22–29 DOI 10.1109/MS.2009.109.
- Kleppe A. 2008.** *Software language engineering: Creating domain-specific languages using metamodels*. Boston: Addison Wesley.
- Kuhrmann M. 2011.** User assistance during domain-specific language design. In: *FlexiTools Workshop, Waikiki, Honolulu*.
- Lanubile F, Ebert C, Prikladnicki R, Vizcaíno A. 2010.** Collaboration tools for global software engineering. *IEEE Software* 27(2):52–55 DOI 10.1109/MS.2010.39.

- Le Pallec X, Dupuy-Chessa S. 2013.** Support for quality metrics in metamodelling. In: *Workshop on Graphical Modeling Language Development*. New York: ACM, 23–31.
- Leenheer PD. 2009.** On community-based ontology evolution. PhD thesis, Belgium: Vrije Universiteit Brussel. Available at <https://d2so1lpcz9yn41.cloudfront.net/sites/vub/files/20090519a.pdf>.
- Liu Q, Gray J, Mernik M, Bryant BR. 2012.** Application of metamodel inference with large-scale metamodels. *International Journal of Software and Informatics* **6**(2):201–231.
- López-Fernández JJ, Sánchez Cuadrado J, Guerra E, de Lara J. 2013.** Example-driven meta-model development. *Software & Systems Modeling* **14**(4):1323–1347 DOI [10.1007/s10270-013-0392-y](https://doi.org/10.1007/s10270-013-0392-y).
- Mernik M, Heering J, Sloane AM. 2005.** When and how to develop domain-specific languages. *ACM Computing Surveys* **37**(4):316–344 DOI [10.1145/1118890](https://doi.org/10.1145/1118890).
- Miller GA. 1956.** The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review* **63**(2):81–87 DOI [10.1037/h0043158](https://doi.org/10.1037/h0043158).
- Moody DL. 2009.** The physics of notations: toward a scientific basis for constructing visual notations in software engineering. *IEEE Transactions on Software Engineering* **35**(6):756–779 DOI [10.1109/TSE.2009.67](https://doi.org/10.1109/TSE.2009.67).
- Moody DL, van Hillegersberg J. 2009.** Evaluating the visual syntax of UML: an analysis of the cognitive effectiveness of the UML family of diagrams. In: *Conference on Software Language Engineering, Toulouse*, 16–34.
- Mylopoulos J, Chung L, Yu ESK. 1999.** From object-oriented to goal-oriented requirements analysis. *Communications of the ACM* **42**(1):31–37 DOI [10.1145/291469.293165](https://doi.org/10.1145/291469.293165).
- Norman DA, Draper SW. 1986.** *User Centered System Design: New Perspectives on Human-Computer Interaction*. Hillsdale: L. Erlbaum Associates Inc.
- Object Management Group (OMG). 2014a.** Model-Driven Architecture (MDA) specification. Available at <http://www.omg.org/mda/specs.htm> (accessed 6 May 2016).
- Object Management Group (OMG). 2014b.** Object Constraint Language (OCL) specification. Version 2.4. Available at <http://www.omg.org/spec/OCL> (accessed 6 May 2016).
- Object Management Group (OMG). 2015a.** Diagram Definition (DD) specification. Version 1.1. Available at <http://www.omg.org/spec/DD> (accessed 6 May 2016).
- Object Management Group (OMG). 2015b.** Meta Object Facility Core (MOF) specification. Version 2.5. Available at <http://www.omg.org/spec/MOF/2.5> (accessed 6 May 2016).
- Padrón CL, Dodero JM, Lanchas J. 2005.** CASLO: collaborative annotation service for learning objects. *Learning Technology Newsletter* **7**(2):2–6.
- Power JF, Malloy BA. 2004.** A metrics suite for grammar-based software. *Journal of Software Maintenance and Evolution: Research and Practice* **16**(6):405–426 DOI [10.1002/smr.293](https://doi.org/10.1002/smr.293).
- Prinz A, Scheidgen M, Tveit MS. 2007.** A model-based standard for SDL. In: *International SDL Forum, Paris*, 1–18.
- Rittgen P. 2008.** COMA: a tool for collaborative modeling. In: *Forum at the International Conference on Advanced Information Systems Engineering, Montpellier*, 61–64.
- Rooksby J, Ikeya N. 2012.** Collaboration in formative design: working together. *IEEE Software* **29**(1):56–60 DOI [10.1109/MS.2011.123](https://doi.org/10.1109/MS.2011.123).
- Sánchez Cuadrado J, de Lara J, Guerra E. 2012.** Bottom-up meta-modelling: an interactive approach. In: *Conference on Model Driven Engineering Languages and Systems, Innsbruck*, 1–17.
- Sánchez Cuadrado J, García Molina J. 2007.** Building domain-specific languages for model-driven development. *IEEE software* **24**(5):48–55.

- Scheidgen M. 2008.** Textual modelling embedded into graphical modelling. In: *European Conference on Model Driven Architecture—Foundations and Applications, Berlin*. Vol. 5095, 153–168.
- Siorpaes K. 2007.** Lightweight community-driven ontology evolution. In: *International Semantic Web Conference, Busan*. Vol. 4, 951–955.
- Steinberg D, Budinsky F, Paternostro M, Merks E. 2008.** *EMF: Eclipse Modeling Framework*. Boston: Addison Wesley.
- SVG. 2011.** Scalable vector graphics 1.1. Available at <http://www.w3.org/TR/SVG/>.
- Tamburri Da, Lago P, Vliet HV. 2013.** Organizational social structures for software engineering. *ACM Computing Surveys* **46**(1):1–35 DOI [10.1145/2522968.2522971](https://doi.org/10.1145/2522968.2522971).
- Troyer OD, Leune CJ. 1998.** WSDM: a user centered design method for web sites. *Computer Networks* **30**(1–7):85–94.
- Völter M. 2011.** MD*/DSL best practices. Available at <http://voelter.de/data/pub/DSLBestPractices-2011Update.pdf> (accessed 6 May 2016).
- Whitehead J. 2007.** Collaboration in software engineering: a roadmap. In: *Workshop on the Future of Software Engineering*. Piscataway: IEEE Computer Society, 214–225.