An Assistive Approach for Learning Goal Modeling

Karenna Kung and Alicia M. Grubb
Smith College
Northampton, MA, USA
amgrubb@smith.edu

***Abstract*—Goal models, as part of a model-driven engineering approach, allow users to specify the needs and rationales of stakeholders and the dependencies between them. Goal modeling software allows users to model situations and analyze tradeoffs between various stakeholders' goals. However, building and analyzing goal models is not always intuitive for new and non-technical users. While stakeholders who routinely use goal models can receive training through their enterprises, new users and those from smaller organizations may not have the time or resources to undergo extensive training. Integrating training into goal modeling software itself would allow users to create and analyze goal models without requiring external training, while embedding explanations of goal modeling concepts into the context of the user's model. In this paper, we present our design for embedded instructional tutorials that aim to guide and instruct first-time goal modelers, while clarifying the common points of confusion.**

I. Introduction and Motivation

Goal-oriented requirements engineering (GORE) allows users to visualize and analyze the needs, constraints, and rationale s of stakeholders and dependencies between them [1]. While various approaches and notations exist for GORE (e.g., iStar [2], GRL [3], Tropos [4], KAOS [5]), their ability to connect in with other model-driven engineering approaches depends on the accuracy of the model and engagement of stakeholders. GORE tools (e.g., OpenOME [6], jUCMNav [7], piStar [8], BloomingLeaf [9]) promote the elicitation and organization of requirements, enabling communication between stakeholders and developers [1]. However, inadequate communication between companies and customers is a common cause of problems in requirements engineering [10], and stakeholders may not know how to translate their client-side knowledge of the situation into a modeling notation suitable for model-driven development. When stakeholders are unsure how to make a model, analysts must create models based on interviews with the stakeholders [11], adding an extra step with the potential for miscommunication. Further, recent work has critiqued the lack of adoption of GORE in industry [12]. Causes of this lack of adoption appear to be the complexity of tooling, lack of available training, and difficulty in understanding models [13]. While stakeholders who routinely use goal models can receive training through their enterprises, new users and those from smaller organizations may not have the time or resources to undergo extensive training.

We aim to reduce this barrier to entry for BloomingLeaf [9], the GORE tool developed by our research lab. We enhance the BloomingLeaf tool with a series of instructional modules that both explain the mechanics of using the tool and provide the conceptual understanding required to create and analyze models. Through an interactive design, we incorporate suggestions from the literature and lessons learned from an initial case study with a new user. The modules are split into three tutorials aimed to assist a user with understanding goal modeling, making a model, and developing and answering questions using simulations.

In this short paper, we first review previous approaches to assisting users with modeling in Sect. II and give an overview of our design methodology in Sect. III. We present our current implementation (i.e., the primary contribution of this work) in Sect. IV. We conclude the paper by discussing additional tradeoffs and our plan for validation in Sect. V.

II. Background and Related Work

While there are multiple texts to train individuals in the concepts of GORE (e.g., [2], [5], [14], [15]), this information is not embedded within tooling. When taught in a classroom context, instructors can provide an overview of the process and tooling for students via course notes and laboratory assignments. Practitioners and those outside of institutions where GORE is taught, require assistance in this process.

Our tool, BloomingLeaf [9], implements the Evolving Intentions framework [16], a time-based extension to Tropos [4]. The original tool provides a reference guide via a *.pdf* file available in the Documentation menu, which we refer to as the "cheatsheet" in the remainder of the paper. The cheatsheet contains a sample model and lists of the possible types of actors, intentions, links, evidence pairs, function types, EVO modes, etc. This approach is consistent with similar tools [8]. While these lists may allow users to recognize concepts that fill a similar role (e.g., goals and tasks are types of intentions), these concepts were not linked to each other and the documentation contained little information about the meanings of the different possible values.

BloomingLeaf and other GORE tools offered videos to improve the installation process and assist in tool usage. However, as the "action-breakdown-repair model" [17] informs us, videos and cheatsheets are less effective when they are not linked directly within the tool to the particular task the user is in the process of completing. Prior work found varying preferences and time savings from different formats of tutorials depending on learning style and gender [18]. Thus, our goal is to embed appropriate GORE training directly within BloomingLeaf.

TABLE I: Categorization of questions asked by the Client.

	Mechanical	Conceptual
General Modeling	-	1
Creating	7	12
Analyzing	1	6

III. Feature Elicitation and Design

We used an iterative approach to developing the interactive tutorials. We started with the cheatsheet already present in BloomingLeaf (see Sect. II) and divided it into atomic units. Next, we brainstormed an initial list of sections to be added. As tool developers, we were well versed in the processes of creating and analyzing goal models and BloomingLeaf, and much of this knowledge was tacit. To truly capture the experience of a new user, we needed their perspective. Thus, we recruited an undergraduate student (i.e., *Client*) who was completely unfamiliar with goal modeling to act as a novice user. Inspired by diary studies [19], we asked Client to document her thoughts and questions as she learned. Client was directed to run the tool on her computer, but was given no further instruction on using the tool. Over several weeks, Client recorded her questions independently, and we met with her weekly to further understand specific points of confusion and questions. We first focused on modeling and general tool usability, before working through how to analyze goal models.

We collected and analyzed Client's documented questions to inform the development of the interactive tutorials. Questions either were matched to the initial list of brainstormed sections or prompted the addition of a new section. Modules incorporated the answers to Client's relevant questions. Client asked questions about both the mechanics of how to use the tool to make the modifications she envisioned, and about the factors to be considered when making decisions about which modifications to make. Between questions that Client wrote down independently and questions that she asked while in sessions with the researchers, 27 unique questions were identified and were classified as either (1) related to overall understanding of modeling, (2) related to creating a model, or (3) related to analyzing the model. Categories (2) and (3) can be further broken down into conceptual questions (e.g., "I don't know what I want to do") and questions about the mechanics of using the tool (e.g., "I know what I want to do but I don't know how to accomplish it"). We report counts for each type of question in Table I and note that conceptual questions were more prevalent. Questions about tool mechanics are more straightforward to answer, as they can be resolved with detailed verbal explanations of a certain part of the tool and screen recordings of the tool being used. On the other hand, conceptual questions require more coordination, as understanding of a given part of the modeling process is interrelated to understanding other parts, and the order in which confusions arise may depend on the individual user.

Our work with Client reminded us that there is no standardized methodology for creating and analyzing a model.

Fig. 1: BloomingLeaf toolbar with Tutorials menu selected showing the three tutorials.

In our previous studies, new users commonly question how they know that their model is complete and may not innately know how to use analysis features to answer a question or what questions are answerable with analysis [20], [21]. While the Evolving Intentions framework proposed a methodology specific to the framework [16], it assumes users have a conceptual understanding of the types of questions that can be asked. Earlier work by Horkoff and Yu [6] did not make this assumption, instead proposing that users first explore the model using forward and backward propagation rules to evaluate the validity of the model before forming specific scenarios. Thus, through our interactive tutorials, we aim to clarify mechanical and conceptual questions, as well as help new users to propose appropriate questions for their model and interpret analysis results.

IV. BloomingLeaf Interactive Tutorials

In this section, we describe the interactive tutorials we incorporated into BloomingLeaf (available at *https://github.com/amgrubb/BloomingLeaf*).

Tutorial Structure. The interactive tutorials are divided into three topics: purpose and elicitation, model creation, and analysis (per the guidance in [6]). Each of the tutorials can be accessed through the toolbar in BloomingLeaf, see Fig. 1 for screenshot. The first tutorial, titled "Purpose and Usage" (see Fig. 1), describes motivations behind goal modeling and introduces elements of the goal modeling process to frame the user's approach. The second tutorial, titled "Build a Model", walks the user through creating a model by brainstorming and successively adding more information to their model. The third tutorial, titled "Analyze the Model", allows the user to explore relationships within the model and guides them through the process of developing questions about certain intentions and using simulated paths to answer those questions. Fig. 2 lists the modules for the first and second tutorials, with the second tutorial sub-divided into two sections. The non-linear modules in the third tutorial are shown as a flowchart in Fig. 3.

User Interaction. After selecting a tutorial, users are shown the *overview* module (see Fig. 2 and Fig. 3). Each module is encapsulated within a pop-up window. For example, we show the module for "1e. Link Intentions" in Fig. 4, with Fig. 4(a) illustrating the initial view. The main instructional section gives the essential information, including instructions for the mechanics of performing the step or the property that the user needs to identify about their model. In the example in Fig. 4(a) for the "Link intentions" module, the user is guided through the process of creating a link between two intentions.

Purpose and Usage Tutorial
0. Overview - BloomingLeaf
   0a. What is goal modeling
   0b. Sample models and insights
   0c. Evolving models
   0d. Getting started...

Build a Model Tutorial
1. Overview - Create the model
   1a. Add actor to model
   1b. Use actor inspector
   1c. Add intentions to the model
   1d. Use intention inspector
   1e. Link intentions
2. Overview - Adding evolutionary information
   2a. Set initial satisfaction value
   2b. Set evolving functions
   2c. Limit presence intervals
   2d. Change max absolute time
   2e. Set absolute time points
   2f. Set relative intention assignments
   2g. Set absolute relationship assignments
   2h. Change presence intervals
   2i. Use intermediate values table

Fig. 2: List of modules for "Purpose and Usage" and "Build a Model" tutorials.

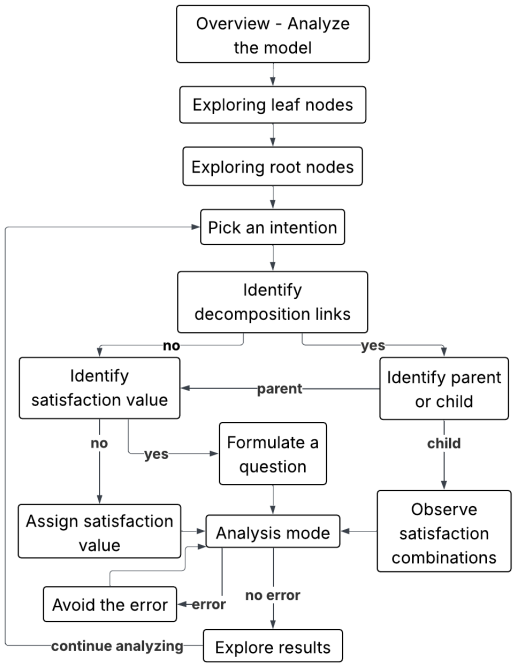

Fig. 3: Flowchart of modules in "Analyze the Model" tutorial.

Most modules also have a "Learn More" section (see Fig. 4(a) for link and Fig. 4(b) for resulting popup), which gives more information about how to decide between different options and, in many cases, screen recordings with the topic of the module applied to an example model. These expanded sections incorporate additional conceptual information and visualize mechanical operations. We evaluated having a "Learn More" section vs. having the entire module in a single window and opted for the separation because Client found the single window option overwhelming and more akin to the PDF manuals provided by most tools, whereas the "Learn More" option presented information in digestible chunks. While we do not claim the completeness of the "Learn More" sections, they allow the user to view a model in progress and focus only on the topics they have seen in previous modules.

The bottom of each module window (see Fig. 4(a)) gives the user the option to navigate to other parts of the tutorials. The "Next" and "Back" buttons navigate to the next and previous modules, respectively. Users may use the dropdown menu (bottom-center in Fig. 4(a)) to select any other module in the current tutorial by module title. Finally, above the "Next" button is a link to navigate to a different tutorial. The "Purpose and Usage" and "Analyze the Model" tutorials jump to the "Build a Model" tutorial, while the "Build a Model" tutorial jumps to the "Analyze the Model" tutorial (e.g., see "Jump to Analysis" in Fig. 4(a)). When jumping between different tutorials, the user's most recently viewed module is bookmarked within each tutorial, enabling them to return to their place. For example, if the user realized during the analysis that they had forgotten to add a link and wanted to be reminded how to do so, they could navigate within the modeling tutorial to Fig. 4(a), add the required link, and then select the analysis tutorial again, which would return the user to their previous tutorial position.

Understanding User Intent. An important aspect of goal modeling is using the model to simulate scenarios and answer trade-off questions about the model. In the "Analyze the Model" tutorial, users are first instructed to generally explore their model by setting the satisfaction values of leaf and root nodes, and then are instructed to choose an intention of interest and develop a specific question about it. Depending on the user's evaluation of the intention they chose, they are guided to ask questions either exploring the outcome of a certain choice or exploring the factors necessary for a certain outcome to occur. Researchers brainstormed a list of questions that new users might want to ask a model and were able to represent all brainstormed questions in one of two general formats: (1) "How does the [satisfaction value] of [intention A] versus [intention B] affect the model?" or (2) "Given that [intention A] is [satisfaction value], what must be true of [intention B]?". Thus, the "Analyze the Model" tutorial follows a non-linear structure that prompts the user to pursue one of these two types of analysis questions by asking them a series of questions about the intentions and links in their model. Fig. 3 shows the titles of each of the modules in the tutorial and how they are connected. Identified links in Fig. 3 (e.g., "Yes", "No") show options that replace the "Next" and "Back" buttons for the module's navigation buttons. Modules with these identified links pose questions to the user about an aspect of the model, ultimately guiding them to test satisfaction combinations of different intentions if they are asking a type (1) question, or to

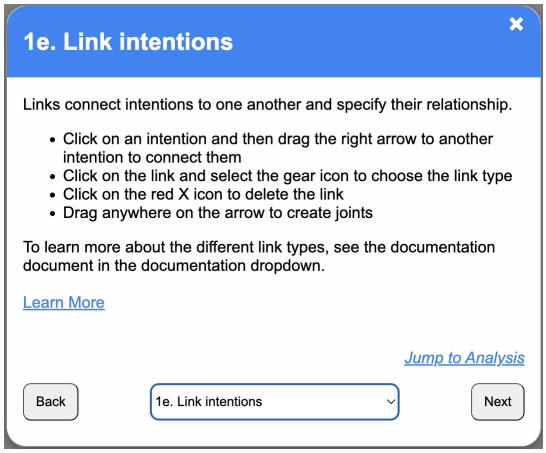

(a) Initial View

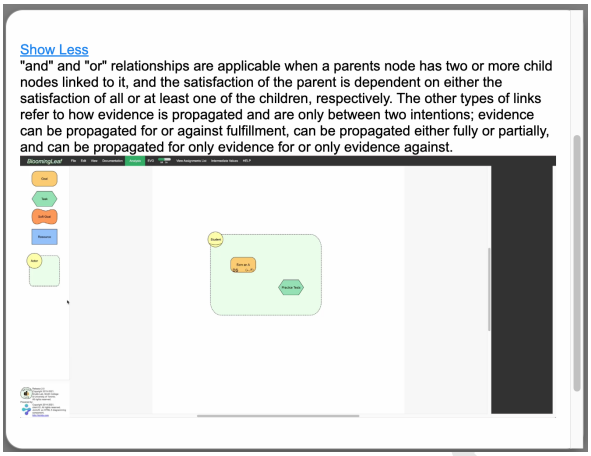

(b) "Learn More" Section

Fig. 4: Screenshots of the "1e. Link Intentions" module within the "Build a Model" tutorial.

formulate a question about possible values of another intention if they are asking a type (2) question. With this guidance, users who were previously unsure of how to begin goal modeling, can ask and answer questions of a model that they have built.

V. Discussion and On-going Work

Tradeoffs for Understandability. Our modules aim to help inexperienced users create and analyze their first models, but we acknowledge that our modules simplify the modeling process. Since the modules are presented sequentially, unless users click through all of the modules before starting, users may miss information that experienced modelers would be aware of while making early decisions.

New users experience difficulty generating questions to ask of their goal model. In the analysis tutorial in Fig. 3, we ask users to identify an intention of interest and then to formulate a question using the selected intention. This strategy depends on the user selecting an intention that results in a sensible, nontrivial question. In an alternative design, we asked the user to formulate a question first, but then could not deterministically understand the structure of the question to guide the user in how to set up the model for analysis. There is a similar tradeoff in asking users to identify questions before or after creating their model. Questions about a model and the creation of the model itself are interdependent, but this fact may not be apparent to new users. Since new users must necessarily learn about one before the other, we chose to discuss model creation first, but provide the ability for the user to jump between tutorials (see Sect. IV).

Horkoff and Yu [6] recommend that forward analysis questions start with the leaf nodes entirely satisfied or entirely unsatisfied as baselines, in comparison to their most likely values. Backward analysis may start with setting all root nodes to their most ideal satisfaction value, or their minimum accepted satisfaction value and then gradually increase their satisfaction value while checking that it remains viable [6]. Alternatively, one can start with the most obvious option to check that the model was properly built [11]. Later analyses take the form of a domain-driven analysis, in which specific questions are asked about individual nodes [6]. While experienced developers may be able to jump directly to posing a question and using the model to answer it [16], we hypothesize that new users benefit from preliminary general explorations to help them see the potential of the model and the relationships that their questions might probe. Thus, though our tool automatically performs both forward and backward analysis from any node, users are prompted to start with setting the satisfaction values of leaf and root nodes so that they can see how the simulation works before coming up with a specific question. In our ongoing work, we add a feature to assist new users in identifying root and leaf nodes, similar to the functionality of OpenOME [22].

A new user viewing the results of their simulations will see the satisfaction values but will not understand the propagation that led to those results. The modules contain information about interpreting the tool, but not the underlying algorithms for propagation. Therefore, users likely lack the intuition that experienced modelers have. However, if the new user continues to make models, they will develop this intuition through observation. Further, the purpose of the tool is to automate analysis over complex models, so the ability to use the tool to interpret models should not depend on users being able to perform analysis steps manually. In our ongoing work, we explore whether users want or need to understand the mechanics of the underlying analysis.

Future Study and Validation Plan. Thus far, we created a set of embedded tutorials for BloomingLeaf and evaluated them through our co-design process with Client, yet, the generalizability of Client's insights and efficacy of the tutorials in assisting a new user to understand goal modeling and BloomingLeaf needs to be evaluated. In a future empirical study, we aim to validate the effectiveness of the tutorials in explaining goal modeling concepts and tool mechanics for new users, as well as evaluate the completeness of the tutorials and clarify remaining points of confusion.

We proposed four research questions:

- RQ1 How do the BloomingLeaf tutorials affect users' ability to create a robust and correct goal model?
- RQ2 How do the BloomingLeaf tutorials affect users' ability to evaluate and answer questions about their model?
- RQ3 How do subjects use the BloomingLeaf tutorials?
- RQ4 To what extent are BloomingLeaf's tutorials complete?

We envision examining these research questions through a controlled experiment with subjects new to goal modeling. Previous studies in goal modeling (e.g., [23], [24]) have used standalone videos and handouts to train new users, while having a trained modeler on hand to answer questions. We compare our interactive tutorials with the video-handout model of training. Users will be split into two groups: one of the groups will be presented with the video-handout training materials and the unmodified tool, while the other group will be presented with the tool with the embedded tutorials. Using BloomingLeaf, subjects will model a chosen scenario that is of personal interest to them and for which they have domain knowledge. While subjects build their model, we measure aspects of the modeling process (e.g., time, size of the model) and how they use the tutorials, where applicable. For example, we will record which of the modules the subject viewed and whether they opened the "Learn More" section.

Once subjects have completed the modeling session, we evaluate the correctness of their model and their understanding. For example, we will interview them to check proper classification of intentions, semantics of links, placement of intentions within actors, question alignment with the model, and ability to use the model to answer specific questions they documented. Next, we evaluate the subjects' understanding of their own model and goal modeling concepts more generally. For example, we will ask the subjects to justify their decisions and answer questions about the model. By extension, we evaluate whether subjects can extrapolate to future non-guided models by asking them questions about another model not previously seen through any study materials. Finally, we ask subjects to self-report how satisfied, frustrated, or confused they were about their model and the goal modeling process in general, as well as any points of confusion throughout the process. This feedback will help identify tutorial improvements.

Summary. In this paper, we described our extension to the BloomingLeaf tool, which aims to streamline the process of learning goal modeling for new users by embedding training into the goal modeling process. The training tutorials explain the overarching purpose of goal modeling, assist users in building a model, and guide them towards an analysis question. We incorporate aspects of goal modeling methodology previously expressed in the literature, as well as points of confusion expressed by a new user. Our ongoing work focuses on gathering diverse feedback about remaining areas of confusion for new users. We plan to continue to modify and clarify the tutorials as new questions arise from users. We will evaluate the efficacy of the tutorials through a comparison of the performance and understanding of new users presented with the BloomingLeaf tool with and without these embedded tutorials.

We hope to gain feedback from the model-driven engineering (MDE) community to inform our ongoing validation, and in-turn hope to inspire other tool developers within the GORE and MDE communities to consider how their tooling is viewed by novice users. The lessons learned and methodologies from our process translate to other tools and modeling languages.

Acknowledgments. We thank Molly Daniel and Christine Dong for their contributions to this project. This material is based upon work supported by the National Science Foundation under Award No. 2104732.

References

[1] J. Horkoff *et al.*, "Goal-Oriented Requirements Engineering: An Extended Systematic Mapping Study," *Req. Eng.*, vol. 24, no. 2, 2019.
[2] F. Dalpiaz, X. Franch, and J. Horkoff, "iStar 2.0 Language Guide," *arXiv:1605.07767*, 2016.
[3] D. Amyot, "Introduction to the User Requirements Notation: Learning by Example," *Computer Networks*, vol. 42, no. 3, pp. 285–301, 2003.
[4] P. Giorgini, J. Mylopoulos, and R. Sebastiani, "Goal-oriented Requirements Analysis and Reasoning in the Tropos Methodology," *Engineering Applications of Artificial Intelligence*, vol. 18, no. 2, pp. 159–171, 2005.
[5] A. van Lamsweerde, *Requirements Engineering - From System Goals to UML Models to Software Specifications*, 1st ed. Wiley, January 2009.
[6] J. Horkoff and E. Yu, "Interactive Goal Model Analysis For Early Requirements Engineering," *Req. Eng.*, vol. 21, no. 1, pp. 29–61, 2016.
[7] G. Mussbacher, J. Whittle, and D. Amyot, "Semantic-Based Interaction Detection in Aspect-Oriented Scenarios," in *Proc. of RE'09*, 2009.
[8] J. Pimentel and J. Castro, "piStar Tool—A Pluggable Online Tool for Goal Modeling," in *Proc. of RE'18*, 2018, pp. 498–499.
[9] A. M. Grubb and M. Chechik, "BloomingLeaf: A Formal Tool for Requirements Evolution over Time," in *Proc. of RE'18*, 2018.
[10] D. M. Fernández *et al.*, "Naming the pain in requirements engineering," *Empirical Software Engineering*, vol. 22, no. 5, pp. 2298–2338, 2017.
[11] J. Horkoff and E. Yu, "Evaluating goal achievement in enterprise modeling – an interactive procedure and experiences," in *The Practice of Enterprise Modeling*, 2009, pp. 145–160.
[12] A. Mavin *et al.*, "Does Goal-Oriented Requirements Engineering Achieve Its Goal?" in *Proc. of RE'17*, 2017, pp. 174–183.
[13] M. Santos *et al.*, "On the Impact of Semantic Transparency on Understanding and Reviewing Social Goal Models," in *Proc. of RE'18*, 2018.
[14] E. Yu, P. Giorgini, N. Maiden, and J. Mylopoulos, *Social Modeling for Requirements Engineering: An Introduction*. MIT Press, October 2010.
[15] ITU-T, "Recommendation Z.151 (11/08) User Requirements Notation (URN) – Language definition," http://www.itu.int/rec/T-REC-Z.151/en., ITU-T – International Telecommunications Union, Tech. Rep., 2008.
[16] A. M. Grubb and M. Chechik, "Formal Reasoning for Analyzing Goal Models that Evolve over Time," *Req. Eng.*, vol. 26, no. 3, 2021.
[17] G. Fischer, A. C. Lemke, and R. McCall, "Towards a System Architecture Supporting Contextualized Learning," in *Proc. of AAAI'90*, 1990.
[18] D. U. Bolliger and S. Supanakorn, "Learning styles and student perceptions of the use of interactive online tutorials," *British Journal of Educational Technology*, vol. 42, no. 3, pp. 470–481, 2011.
[19] S. Ohly, S. Sonnentag, C. Niessen, and D. Zapf, "Diary Studies in Organizational Research," *Journal of Personnel Psychology*, 2010.
[20] A. M. Grubb, "Reflection on Evolutionary Decision Making with Goal Modeling via Empirical Studies," in *Proc. of RE'18*, 2018, pp. 376–381.
[21] Y. Baatartogtokh, I. Foster, and A. M. Grubb, "A Splash of Color: A Dual Dive into the Effects of EVO on Decision-making with Goal Models," *Requirements Engineering*, vol. 29, no. 3, pp. 371–402, 2024.
[22] J. Horkoff, Y. Yu, and E. Yu, "OpenOME: An Open-source Goal and Agent-Oriented Model Drawing and Analysis Tool," in *Proc. of iStar'11*, 2011, pp. 154–156.
[23] N. Cebula, L. Diao, and A. M. Grubb, "A Preliminary Investigation of the Utility of Goal Model Construction," in *Proc. of iStar'20*, 2020.
[24] Y. Baatartogtokh, K. Cook, and A. M. Grubb, "Exploring the Robustness of the Effect of EVO on Intention Valuation through Replication," in *Proc. of ICSE'25*, 2025, pp. 680–680.