



Universiteit
Leiden

The Netherlands

A relational approach to understanding interactions in interactive art

Xu, D.

Citation

Xu, D. (2025, October 7). *A relational approach to understanding interactions in interactive art*. Retrieved from <https://hdl.handle.net/1887/4266648>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4266648>

Note: To cite this publication please use the final published version (if applicable).

Chapter 3

The Development of the Relational Modelling Tool (RMT)

3.1 Introduction

In Chapter 2, we introduce a relational model designed to capture the characteristics of diverse forms of interaction in interactive art. This model analyses the actions performed by the interacting elements, be they audience members or art systems, and examines the various forms of communication between them. An intriguing aspect of the relational model lies in its equal treatment of the different types of participants—whether human, nonhuman organisms, or technical systems. This perspective not only enables us to better understand an interactive artwork from the viewpoints of different participants, but also opens creative possibilities for developing novel interactive dialogues where the participating elements take on unconventional roles.

The relational model, when applied manually, demands a strong familiarity with the model and can be both inefficient and prone to errors. To make the model more accessible and streamline the application process to some extent, we decided to translate the model into a web-based tool that can be easily learned and used by a broader group of users, including artists, researchers, and practitioners in the field of interactive art and beyond. Moreover, we also aim to implement features that allow us to explore and generate new forms of interaction using algorithmic techniques. The outcome of this process is the development of a web-based application, referred to as the Relational Modelling Tool (RMT).

RMT offers a structured and formal input area that guides users in describing an interac-

Key Components

tion using the concepts from the relational model. Given the formal description of an artwork, it automatically generates a visual graph that provides a direct overview of the interaction. The resulting descriptions can be stored in a database for later retrieval, contributing to an expanding collection of interactive dialogues. The formal structure of RMT enforces rigorous reasoning about the relationships between the actions of the elements, thereby making the description more precise. It also generates data with consistent types and structures that allow us not only to easily compare different interactions, but also to generate new interacting elements and behaviours based on the collection of interactive dialogues and draw new connections between them using different combinatorial and randomisation algorithms.

As such, we believe that RMT can not only benefit the analysis of interactive art but also facilitate the discovery and creation of new interactions. It is interesting to note that we did not initially foresee all the benefits and potential of RMT, rather, we discovered, and are still discovering, its benefits alongside its continued development.

This chapter largely builds on our earlier publication presented at the 13th EAI ArtsIT conference (Xu, Lamers, & van der Heide, 2025), which introduces the version 0.1.0 of RMT. Here, we present RMT¹'s description and visualisation components, detailing its design and development while highlighting modifications made to the original relational model. In this chapter, we focus on sharing RMT's development process and reflecting on its benefits. In doing so, we aim to highlight how this approach can advance both the research and creation of interactive art, while also inspiring the development of similar tools and applications across interactive art and other domains. The practical application of RMT in describing and visualising diverse forms of interaction in interactive art is demonstrated with concrete artwork examples in Chapter 4 and Chapter 5. Furthermore, since the initial release, we have made several key improvements to RMT based on feedback from other users during an evaluation workshop. The details of the workshop and the improvements are presented in Chapter 7.

3.2 Key Components

In this section, we focus on discussing the design and features of RMT, including some challenges encountered during its development and the adjustments made in relation to the model. RMT's development has been an iterative process. From the early stages, we have been testing the prototype to model a wide range of interactive artworks, continuously improving upon identified problems and limitations.

RMT is developed using HTML and JavaScript language, allowing easy access via a pub-

¹In the remaining of this chapter, RMT refers to the version 0.1.0.

lic webpage². The webpage contains what we call a *worksheet* that provides an interface to RMT and allows users to describe an interaction within it. Upon opening it, users can click on the *About* button to view a brief introduction to the relational model and guidance on using the tool. RMT consists of two main components: a set of formal input fields for describing an interaction, and a visualisation component that automatically generates a visual graph of the interaction based on the description. Typically, a worksheet includes a detailed description of an interactive artwork. However, if an artwork contains multiple different modes of interaction, each mode can be described on a separate worksheet.

Furthermore, RMT is connected to a cloud-hosted database using Google's Cloud Firestore service (Google, 2024). The database enables users to save a worksheet with all entered data, allowing them to revisit later or share their work for communication and collaboration. Users can also browse the artwork collection in the database to familiarise themselves with RMT or seek inspirations. A worksheet can be locked for viewing only, with all input fields disabled. Users can click on the *Unlock* button to create an editable copy. Once an unlocked worksheet is loaded, users can directly edit the existing data to refine the description or create new interactions. Additionally, the collection data³ can be used to facilitate modelling and comparing existing interactions, as well as generating new interactions, as we will elaborate later. Below, we introduce the main features of RMT in detail.

3.2.1 The description component

The description component consists of input fields that are structured based on the key concepts—element, action, and communication—and their descriptors developed in the relational model. To describe an interaction, the model specifies the types and counts of the interacting elements with shared properties and behaviours, the actions they perform, and the forms of communication resulting from each action within one element profile. We group the descriptors for each concept using different outlines or background colours and arrange them in a nested structure to highlight their hierarchical relationships. As illustrated in Figure 3.1, one element profile, indicated with a dotted outline, can contain multiple actions, shown as grey rectangles, each of which can result in multiple forms of communication, depicted as white rectangles. The input fields are either text fields or selection menus. To facilitate the modelling process, a brief explanation is displayed when users hover over a concept or descriptor text. Initially, only the input fields for an element, an action, and a communication are shown. Users can expand the input fields by adding more elements, actions, and communications.

²The latest version of RMT can be accessed via this URL: <https://modeltool.liacs.nl>.

³Currently, the database contains the descriptions of a collection of interactive artworks focusing on co-located interaction and more-than-human interaction, which are presented in Chapter 4 and Chapter 5, respectively.

Key Components

ELEMENT: #1

Type:

Count:

ACTION:

1 Action

☒ Intended ☐ Unintended

If do(es) +

Then this action

COMMUNICATION:

To:

Means: ☒ Direct ☐ Via

Config:

Count:

Access: ☒ Public ☐ Private

Effect:

Additional information:
e.g. urls or notes about the artwork.

Footnotes:

Figure 3.1: Screenshot of the input fields upon opening an empty worksheet.

The input fields largely adhere to the original model described in Chapter 2; however, several important adjustments and additions were made in RMT. Firstly, in the relational model, we previously specified the *components* for each element, which refer to the devices or apparatus constituting the element. However, this descriptor is not always effective or informative when describing human participants when they participate using their bodies. Additionally, in some artworks, elements of the same type and components may exhibit different behaviours and play distinct roles in the interaction, which requires separate analysis. Therefore, we decided to remove the descriptor component and instead ask users to indicate the role of elements after specifying their type, if such distinction is required.

Moreover, in the description of a communication, we have added the descriptor *access* to indicate the degree to which a communication is perceivable by other elements and audience, classifying it as either private or public. This distinction was first made as part of the *configuration* of a communication (see Chapter 2).

We also made some adjustments to enhance the description of an action. In the relational model, we indicated that an action can be either intended or unintended for interaction, but we did not create an explicit descriptor for this distinction. In the application of the model, this information is sometimes mentioned but not always present (Xu, Lamers, & van der Heide, 2023). During the development of RMT, we initially created only a text input field for describing what an action is. Following the strict formal structure of RMT, we soon noticed the

omission of the intention of an action. Consequently, we added a radio button after an action that requires users to indicate whether it is intended or unintended for interaction.

Another key adjustment is the replacement of the descriptor *role* for an action in the relational model with the descriptor *effect* for a communication in RMT. Originally, we use the role(s) of an action to describe “the function(s) it serves and how it relates to other actions performed by the same or different elements in the interaction”. Following the nested structure of RMT, we observed that the role(s) of an action can be further broken down into the effect(s) of the various forms of communication it creates in relation to the receiving element(s). To enhance the precision of description, we decided to use the effect of communication to specify the functions and consequences of a communication, and its impact on and relation to the receiving element. When a communication has multiple effects, users can add additional input fields to describe each effect.

The last adjustment we would like to highlight is the addition of condition for an action in the form of *If [Element] does [Actions]*. This addition allows users to specify which prior actions are required for the described action to occur, thereby clarifying the relationships between the actions. Previously, such information was implied in the role(s) of an action. However, as RMT took shape, we realised that a clear indication of the relationships among actions is necessary for users to grasp the dynamic exchange between elements. Moreover, when multiple triggering actions are involved, specifying their roles alone is insufficient to capture the connections among them.

RMT provides several options to describe the condition of an action. Users can select “self-initiated” if an action is initiated by the element itself. If an action is triggered by another action, users can specify the triggering action by selecting the corresponding element and action. In this case, the described action is designated as a reaction. An action can serve as both a triggering action and a reaction to the same or different actions. If an action is triggered by the interplay between two actions, users need to specify both actions and indicate whether they need to occur simultaneously (an *AND* relationship) or only one at a time (an *OR* relationship). With its current structure, RMT can only specify the relationships among triggering actions linearly. When there are more than two triggering actions, their interplay may be more complex. In such cases, users can add a footnote to specify such relationships. The condition primarily aims to capture the formal or logical relationships among actions. For more nuanced influences between them, users can indicate them in the effects of the resulted communications.

Furthermore, when multiple reactions are triggered by the same conditions, RMT provides options to specify the relationships between the reactions⁴. In such instances, users can

⁴This feature was not implemented in RMT v0.1.0 and was identified as a limitation in the original publication

Key Components

indicate whether the current reaction should occur concurrently with other reactions (an *AND* relationship) or whether only one reaction should be activated at a time (an *OR* relationship) by selecting the appropriate option from the dropdown menu following the “Then this action” prompt. Upon selection, input fields for specifying the associated reactions will be made available. The implementation of this feature is symmetrical to the condition specification process, thereby adhering to the same logic and constraints.

Lastly, to make the description process easier, all input fields are connected to a dropdown list of possible options. For the selection menus related to the condition of an action and describing the recipient or mediator of a communication, the options are based on present information within the worksheet and updates as more data is entered by users. For the text input fields associated with element type, action, and communication effect, the options are collected from all existing descriptions within the worksheets stored in the database. Users can click on the input field to view the attached list and type in letters or words to narrow down the suggestions. If none of the options suits the interaction under description, users can enter new descriptions that will be added to the list once they are stored in the database. We believe this feature not only eases the workload of the modelling process by reducing the need for repeated data entry but also helps users in phrasing these descriptions in a more consistent manner. Besides describing the interaction, users can also supplement additional information, such as the URL to the artwork documentation or other notes.

3.2.2 The visualisation component

The visualisation component of RMT was developed using the D3 JavaScript library (Observable, 2024). Thanks to the formal structure of the input fields, we can represent the concepts and their descriptors in the relational model using a set of visual symbols as shown in Figure 3.2 (left), and code instructions for automatically generating graphs to illustrate a described interaction. These graphs can also be downloaded as PDF files for separate usage. This feature is useful for complex descriptions where the various interplays between individually described elements are not easily perceptible. While modelling such an interaction, users can utilise the visualisation to inspect the description. Furthermore, visualising a described interaction provides users with a direct overview without necessitating a review of all the details in the input fields. As a result, users can easily compare multiple interactions and identify their similarities and differences with direct visual access.

As shown in Figure 3.2 (right), elements represented as rectangles with rounded corners are arranged horizontally in a visualisation, while their actions are depicted as standard rectangles positioned vertically along the central axis of each element. The intention of an action is




Figure 3.2: (Left) The legend describing the visual symbols representing the concepts and their descriptors for the visualisation. (Right) An example visualisation of described interaction in the artwork *Body Movies*.

indicated with different stroke styles of the rectangle. When an element count is two, two element squares are drawn on top of each other with a slight offset. If an element count exceeds two or is simply indicated as more than one, only three element squares are drawn in the same stacked and offset manner. This feature aims to provide visual cues regarding the number of elements, rather than displaying precise information. Communications are always depicted as arrows extending horizontally from an action to an element. A solid arrow line represents direct communication, while a dashed line signifies mediated communication. Configurations are registered in small circles at both ends of an arrow. The access level of a communication is indicated by different opacities: complete opacity signifies public access, while partial opacity indicates private access. Moreover, all textual information is displayed on top of the symbols.

These constructions provide a clear overview of the elements, their actions, and the resulting communications in a described interaction. Additionally, a sorting algorithm determines the display order of actions based on their conditions, ensuring that triggering actions are placed before reactions. Aside from this, the order in which the elements, actions, and communications appear follows the sequence specified in the input fields. However, because an action can be present in the conditions of multiple actions and can function both as a triggering action and as a reaction, we faced challenges in visualising the conditions without creating overlaps or messy indicators on top of the main graph.

To address this, we decided to show this information separately on the side using bars with different colours and styles (see Figure 3.2 right). Each action-condition combination is drawn along the same vertical axis and placed next to each other on the left side of the main graph. Triggering actions are indicated using red bars, drawn at the same vertical positions as the actions, while reactions are indicated with blue bars. This solution not only effectively communicates the relationships between actions without overlaying excessive visual distractions on the main graph, but also provides a system that allows for easy compositions of these relationships on its own.

3.3 Benefits of RMT

The formal structure of RMT exhibits numerous advantages over the narrative approach of the relational model when it comes to modelling and visualising interactive dialogues. Additionally, it opens up opportunities for developing features that facilitate the exploration and comparison of large collections of artworks, as well as the generation of new interactive dialogues. Below, we elaborate on these benefits of RMT in greater detail.

3.3.1 Enhance precision and practicality

Compared with the relational model, RMT significantly enhances the precision of the modelling process and the interaction description. Previously, using only the relational model, we could generate a textual description based solely on our understanding of the model. This process was more subjective and prone to errors, often resulting in descriptions that were ambiguous and incomplete. For instance, the intentions and conditions for actions were easily overlooked. In contrast, RMT provides a comprehensive and consistent scheme for each element, along with suggestions for each descriptor. These features not only guide users through the modelling process, but also aid them in formulating detailed descriptions, ensuring that all aspects of the relational model are thoroughly addressed.

In addition to generating datasets with consistent structure, the formal structure of the input fields allows us to define a visual language for depicting an interaction and automate such process. The visualisation can not only facilitate modelling and comparing different interactions by providing a direct visual impression of the descriptions, but also create easily identifiable visual patterns showing some characteristics of an interaction. Moreover, the visual language itself offers a concise and straightforward method for users to configure relationships and draw connections between elements, for instance, using the visual codes for describing conditions of actions. Its simplicity also enables users to quickly sketch interaction ideas, both digitally

and by hand. Furthermore, we can implement features that allow users to directly manipulate the visualisation, with corresponding updates made to the input fields. This integration can potentially enhance the usability of RMT, making the modelling process more intuitive.

3.3.2 Explore and compare artwork collections

Datasets of interactive dialogues generated with the formal input fields of RMT create opportunities for developing new approaches for analysis and comparison. The unified data structure not only enables straightforward comparison and robust analysis across diverse forms of interaction but also supports the application of statistical and visualisation methods to derive insights from large sets of interactive artworks, which are not always possible to achieve manually.

Specifically, we can experiment with techniques to visualise the artwork collection to create an overview of the stored interactions and the distributions of the data points across the descriptors. With a sufficiently large collection of artworks, this approach can potentially reveal patterns and trends in the evolution of interactive art. In addition, we propose the implementation of a feature that enables users to select and visualise specific artworks or custom collections, facilitating the comparison of similarities and differences between them.

Furthermore, one could imagine the construction of a *distance measure* between two artworks: a numerical expression based on similarities and differences between their interactive dialogues captured by RMT. Such a numeric construction would allow for interesting distributional analyses of a collection of artworks. For example, clusters of works could be identified based on the relative similarity of their interaction structures using standard data analysis techniques. Alternatively, multidimensional scaling techniques could then be easily applied to map a collection of works onto a multidimensional space, revealing underlying patterns of similarity, similar to the use of principal component analysis in multidimensional data analysis. At an application level, these techniques would enable automated curation of artwork collections, as well as other forms of computational and visual analysis of larger collections.

Currently, users can only access the collection by selecting an artwork's title and viewing the described interaction. We also envision implementing features that allow users to browse the data directly. Such features would allow users to select specific data points and view all artworks containing the same data points. By doing so, they can facilitate users to discover connections between artworks that might otherwise go unnoticed. Such features could also supplement the visualisation of the collection as described above.

3.3.3 Generate new forms of interaction

Beyond analytical applications, the collection of existing interactive dialogues serves as a valuable resource for imagining new interactions. Similar to the *Dadaist* cut-up technique in which a written text is cut up and randomly rearranged to create a new text (Burroughs, 1961), we can apply random and systematic generation and permutation features to explore new combinations of element types and behaviours as well as draw new connections between them⁵.

More specifically, we can leverage the existing data about elements from the artwork collections to generate new ones by randomly combining their types, actions, and communications. This feature can be particularly useful for initiating exploration and providing inspiration at the onset of the creative process. Besides generating elements from scratch, we can also provide users with partial control in this process. For example, they can specify certain aspects of an element, such as its type, and then use RMT to generate new actions and communications. Given the diversities of elements in the collection, including human, nonhuman organisms and technical systems, we believe that this feature can stimulate the discovery and imagination of new interactive behaviours and unconventional roles of elements.

Furthermore, we also envision to take advantage of RMT's formal structure to create new configurations using computational techniques. Once elements are specified, we can draw new connections between them using different combinatorial or randomisation algorithms. These connections can be realised by generating new forms of communication, such as randomly assigning the receiving elements and selecting the means, configurations, counts, and access of communications. Additionally, we can generate new combinations of triggering actions and reactions to configure new relationships between actions. Beyond randomisation, we also plan to apply algorithms that allow users to systematically explore all possible combinations and configurations. These manipulations are possible because these descriptors are confined to limited options, and we believe they can help us discover and create relationships between elements that do not yet exist.

Additionally, the standardisation of data collection allows us to leverage machine intelligence even further. For instance, we can apply data mining and machine learning algorithms to explore relations and patterns within the artwork collection and generate new forms of interaction based on these insights.

⁵The generative functions are realised in the later version of RMT, the detailed description about this is presented in Chapter 7

3.4 Current Limitations

There are a few limitations of RMT we wish to address here. One pertains to the description of action conditions. Currently, users must add footnotes to describe complex relationships involving more than two triggering actions. While this approach is effective for a human user to view a description, it is limited in processing such information using computational means, such as randomising the relationships among the triggering actions to generate new conditions. A potential solution is to use a tree-like structure that allows users to nest triggering actions within different branches and specify the orders of the branches. However, this could potentially increase the complexity of the interface and make RMT harder to use.

Balancing the level of generalisation and specification in a description remains a key challenge in the development of RMT. Presently, it has limited capacity to show information about an interaction that is not captured by the generalised formal structure of RMT. Since interactive art often seeks to create unconventional interactions, such information may be essential for understanding an artwork. For future development, we plan to implement features that allow for more open-ended entries in the input fields. However, it remains a challenge to preserve the benefits of a generalised formal structure while creating room for accommodating individual characteristics. We believe that systematic evaluations with professional users will provide valuable insights for making informed decisions regarding these concerns.

3.5 Discussion

In this chapter, we introduce a novel web-based tool-RMT-built upon our previously developed relational model, designed for modelling and visualising interactions in interactive art. The formal structure of RMT enhances the precision of the modelling process, facilitates the analysis and comparisons of different forms of interaction, and can potentially inspire the creation of new forms of interaction with data-driven approaches that were previously unattainable. Compared with the relational model, our findings highlight RMT's capacity to emphasise the relational exchange between elements as well as identify and contextualise other known interaction models⁶. RMT not only supports the research and creation of interactive art but can also be used as an educational resource in this field. While our focus remains on interactive art, we believe that RMT can also be applied in other domains of interaction design and research, and the development process can also inspire the creation and development of similar tools and applications across different fields.

Lastly, we wish to emphasise again that the development of RMT is an iterative process,

⁶This insight is shown in the practical application of RMT in Chapter 4 and 5.

Discussion

with its benefits emerging organically rather than being pre-planned. Since the creation of the first prototype, we have been continually testing the modelling capability of RMT with diverse interactive artworks. As each example emphasises on different aspects about interactions, this process has allowed us to identify limitations and further refine RMT's features, as well as resolve ambiguities inherent in the relational model. Moreover, we have also been learning about what RMT can do as we develop and enrich its features. For instance, we discovered the potential of the artwork collection and implemented the related features only after the database was implemented, despite originally only intended for storage. This is only possible as the research and development processes are integrated, the outcomes from one process continuously provide inputs and inspirations for the other.