MultiCategory: multi-model query processing meets category theory and functional programming

Uotila, Valter Johan Edvard

2021-08-20

MultiCategory: Multi-model Query Processing Meets Category Theory and Functional Programming

Valter Uotila
Jiaheng Lu
University of Helsinki
first.last@helsinki.fi

Dieter Gawlick
Zhen Hua Liu
Souripriya Das
Oracle Corporation
first.last@oracle.com

Gregory Pogossiants
SATS Technologies
gregp_21@yahoo.com

ABSTRACT

The variety of data is one of the important issues in the era of Big Data. The data are naturally organized in different formats and models, including structured data, semi-structured data, and unstructured data. Prior research has envisioned an approach to abstract multi-model data with a schema category and an instance category by using category theory. In this paper, we demonstrate a system, called `MultiCategory`, which processes multi-model queries based on category theory and functional programming. This demo is centered around four main scenarios to show a tangible system. First, we show how to build a schema category and an instance category by loading different models of data, including relational, XML, key-value, and graph data. Second, we show a few examples of query processing by using the functional programming language Haskell. Third, we demo the flexible outputs with different models of data for the same input query. Fourth, to better understand the category theoretical structure behind the queries, we offer a variety of graphical hooks to explore and visualize queries as graphs with respect to the schema category, as well as the query processing procedure with Haskell.

1 INTRODUCTION

The variety of data is one of the most important issues in modern data management systems to cope with the challenge of Big Data. In many applications, data sources are naturally organized in different formats and models, including structured data, semi-structured data, and unstructured data. To address the challenge of variety, multi-model databases have begun to emerge with a single database platform to manage multi-model data together, with a fully
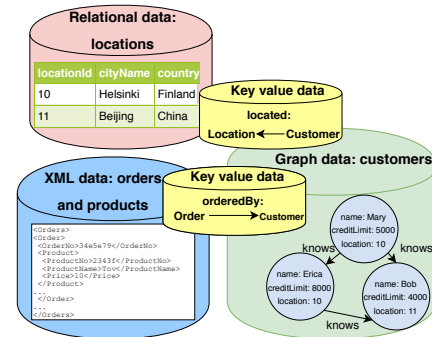
Figure 1: A multi-model data environment

integrated backend to handle the demands for performance and scalability [3].

Let us consider an example of a multi-model data environment. Figure 1 illustrates an application of E-commerce, which contains customers, a social network, and orders information with four distinct data models. The property graph data bear information about mutual relationships between the customers, i.e. who knows whom, and some customer properties such as name and credit limit. The geographic location of customers is stored in a relational table. In XML documents, each order has an ID and a sequence of ordered products, each of which includes product number, name, and price. The fourth type of data, key/value pairs, contains the relations between different data sets. In a typical application like customer-360-view, users of databases demand to analyze the information from these four different data sources together to enable a holistic analysis of customer behaviors.

Category theory was developed by mathematicians in the 1940s and has been successfully applied in many areas of science including computer science. Recent research initiatives have applied category theory for the database area. In particular, Spivak [6, 7] used a schema category and an instance functor to model relational databases. Liu et al. [2] promoted category theory to play the role of the new mathematical foundation to reason about declarative constructions and transformations between various data models.

While the previous works have envisioned the theoretical significance to model and manage data with category theory, this demonstration shows our initiative to showcase a proof-of-concept implementation of `MultiCategory`, a system to support multi-model query processing based on category theory. The core parts of the system have been coded with the functional programming language

Haskell, which is widely recognized to have a strong connection to category theory. The data storing framework of `MultiCategory` is established on the concepts of schema and instance categories [6], and the query processing structure is based on catamorphism and foldable data structures [1]. With these key properties, we can create a system that has a consistent integration with relational, hierarchical, and graph data models and we show how category theory can be used to achieve valuable perspectives for multi-model query representation and processing.

In brief, the demonstration of `MultiCategory` offers the following to the audience:

- category theoretical and functional programming oriented methods of querying and accessing multi-model data with a unified schema;
- a unified query language endowed with Haskell's lambda expressions allowing the users to submit one query to access different models of data seamlessly;
- the flexibility of output the same result with different models, which provides the users an opportunity to exploit the same data with different representations;
- to better understand the theoretical structure behind the queries, this demo also provides an interactive visualizer to understand the schema and instance categories, as well as the query processing procedure.

In our demo, attendees are welcome to compose their queries that follow the syntax of our query language to search the multi-model datasets. The source code of this system is available in GitHub [8] and the demo video can be watched online on YouTube [9].

2 PRELIMINARIES

In this section, we first review the mathematical definition of a category [4], followed by the descriptions of schema and instance categories which are influenced by [6, 7].

Definition 2.1. A category C consists of a collection of objects denoted by $Obj(C)$ and a collection of morphisms denoted by $Hom(C)$. For each morphism $f \in Hom(C)$ there exists an object $A \in Obj(C)$ that is a domain of f and an object $B \in Obj(C)$ that is a target of f. In this case we denote $f: A \rightarrow B$. We require that all the defined compositions of morphisms are included in C: if $f: A \rightarrow B \in Hom(C)$ and $g: B \rightarrow C \in Hom(C)$, then $g \circ f: A \rightarrow C \in Hom(C)$. We assume that the composition operation is associative and that for every object $A \in Obj(C)$ there exists an identity morphism $\mathrm{id}_A: A \rightarrow A$ so that $f \circ \mathrm{id}_A = f$ and $\mathrm{id}_A \circ f = f$ whenever the composition is defined.

Informally speaking, we can understand a category as a graph endowed with the composition rule.

Figure 2 constructs a unified schema category, which represents the schema information of a multi-model data environment in Figure 1. Conceptually, an object in a schema category includes two kinds of data types: (1) the first collection of data types consists of a string, integer, rational, boolean, etc., called predefined data types; and (2) the second collection of data types includes entities, such as customers and products. Morphisms are defined to be the typed functions between the data types, such as a customer is located in a certain location and an order is ordered by a customer. Furthermore, it is important to note that a schema category presents a single
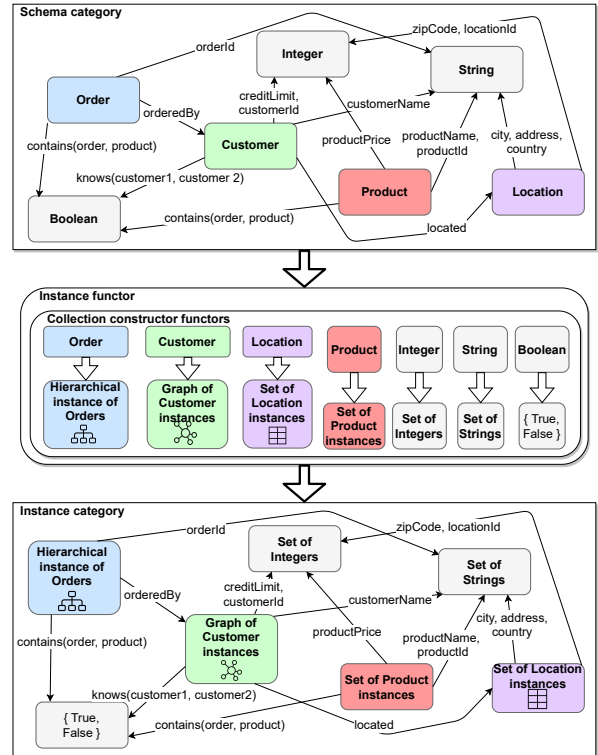

Figure 2: An example of a category theoretical construction

unified view for different models of data. Based on this view, we develop a unified query mechanism to process different models of data seamlessly.

An instance category models how the concrete data instances are stored. Each object of the schema category is mapped to the corresponding typed Haskell data structure in the instance category (see Figure 2). Each morphism in the schema category is mapped to a concrete Haskell function in the instance category. The mapping between these categories is called an *instance functor* which is defined on objects by collection constructor functors [1]. As shown in our demo, queries are formulated based on the schema category and the answers are retrieved from the instance category based on the instance functor and the collection constructor functors.

In schema and instance categories, we can follow any path to form a well-defined function between the start node and the end node of the path thanks to the composition rule in Definition 2.1. For example, there is a morphism (edge) in the instance category that gives us that the customer A makes the order B and another morphism that gives us that the order B includes the product C. Based on the composition rule of these morphisms there is a well-defined morphism that gives us that the customer A buys the product C. This compositionality property is important to guarantee the correctness of programs to traverse through multiple data models.

3 SYSTEM OVERVIEW

In this section, we provide an overview of `MultiCategory`'s architecture, query language, and query processing mechanism. For
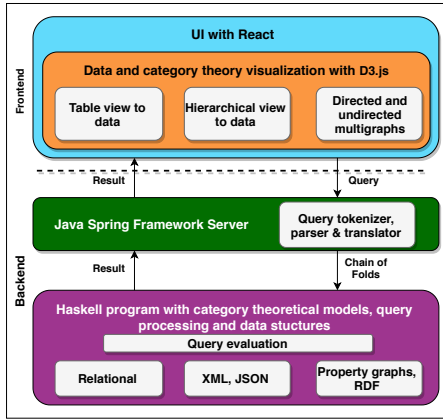
Figure 3: The architecture of `MultiCategory`

Figure 4: The workflow diagram of query processing

more details about technical solutions, a tutorial, and an installation guide you can find from `MultiCategory`'s documentation and in Github [8].

Figure 3 depicts the architecture of `MultiCategory` that consists of the frontend and the backend. In particular, the frontend creates a web interface and data visualizations for relational data, hierarchical documents, and graphs. The backend is responsible for query processing and the implementation of category theoretical constructions.

3.1 Multi-model query language

We have developed a multi-model query language that encapsulates Haskell functions and expressions. For example, the following query finds *those customers whose credit limit is greater than a threshold, say 3000.*

Example 3.1. `QUERY (\x -> if creditLimit x > 3000`
`then cons x else nil)`
`FROM customers`
`TO graph/xml/relational`

Every query block starts with **QUERY** keyword, which is followed by a function defined with Haskell's lambda notation. In Example 3.1, x is a variable that represents a customer in the graph, and the query returns the graph of the customers with credit limit > 3000. The **FROM** keyword specifies the source collection of the data. The **TO** keyword configures the data model of the result. Note that in the above example the returned model can be either *graph*, *relational* or *XML*. Figure 5 demonstrates the three different representation models for the answers of the query in Example 3.1.

Example 3.2. `LET t BE`
`QUERY (\x xs -> if elem "Book" (map productName (`
`orderProducts x)) then cons x xs else xs)`
`FROM orders TO relational IN`
`QUERY (\x -> if any (\y -> orderedBy y customers == x`
`) t then cons (customerName x, countryName(`
`located x locations)) else nil)`
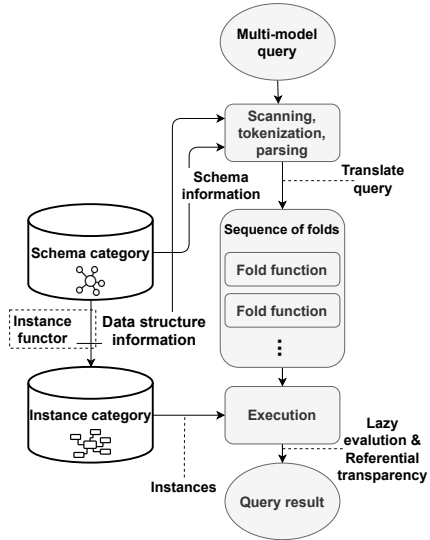`FROM customers TO algebraic graph/relational/xml`

The query in Example 3.2 returns a graph that contains names and locations of the customers who ordered a book, which involves relation, XML, and key/value data types. Figure 6 shows the result of this query. There are two QUERY clauses in this query that correspond to the two fold functions. The clauses are connected with **LET BE IN** structure which works in the same way as the corresponding mechanism in Haskell. The **LET** keyword introduces a variable (i.e. *t*) that connects fold functions together. In particular, the first **QUERY** clause finds any order which contains a book. The second **QUERY** clause finds customers who made such orders. The results contain the customer's name and location information (i.e. countryName).

3.2 Query processing mechanism

Figure 4 depicts the main workflow of query processing in `Multi-Category`. When a user inputs a query to the system, it is parsed into a sequence of *fold-functions* with respect to the schema information from the schema category. The sequence of folds is sent to the Haskell program running in the backend. The Haskell program accesses the instance category and executes the sequence of fold functions that is just pure Haskell code. Note that we do not require all data structures to be instances of Haskell's foldable type class as we can use generalizations of folds to query more complex algebraic data types. For example, we use the function `foldg` to query the algebraic graphs of the package `Algebra.Graph` [5]. Finally, the result is returned to the frontend, where it is visualized depending on the model the user defined in the query.

4 DEMONSTRATION SCENARIOS

4.1 Using schema and instance categories

We first introduce our system by inviting attendees to view the schema and instance categories of varied data sets through the graphical interface. This demo will use six different synthetic and
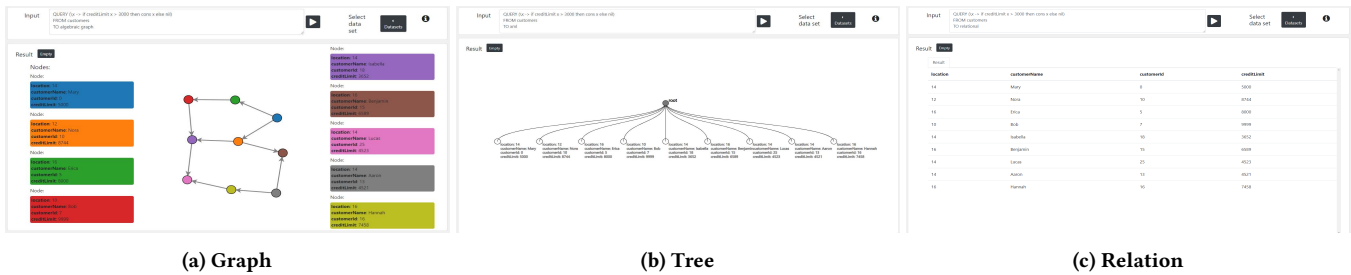
(a) Graph   (b) Tree   (c) Relation

Figure 5: Three different representations of the same query result of Example 3.1

Figure 6: The result in graph and relational models for the query in Example 3.2

real datasets. Each data set includes different models, such as relational, XML, JSON, RDF, and property graph data. Attendees can select a data set, view the related schema and instance categories, and examine the nodes and edges of any graph to find more information. For example, one may find that *customer* data type has an attribute *customerName* which is considered as a function from `Customer` to `String`.

4.2 Querying multi-model data

In `MultiCategory` we have created a collection of example queries that can query different data models together. For example, an E-commerce data set includes all the seven possible combinations of multi-model queries combining relational, XML, and graph models. This demo allows attendees to formulate their queries with guidance and observe results in different output models. Currently, the system does not support large-scale data processing due to implementational reasons but generally category theory-based frameworks scale well.

4.3 Visualizing multi-model queries

For obtaining a better understanding of the theoretical structure behind the queries, `MultiCategory` provides an automatic visualizer for the fold function-based queries. This feature visualizes queries as graphs with respect to the instance category. In particular, after the execution of a query, the query is visualized as a graph that exhibits a composition of Haskell fold functions. Every fold function has at least one lambda expression which the user can click to see the detailed information in the graphical interface.

5 COMPARISON TO EXISTING SYSTEMS

The existing multi-model databases, for example, ArangoDB and OrientDB are implemented based on a single dominant model so that they cannot be called "native" concerning all models they are supporting. In contrast, we developed `MultiCategory` in such a way that each data model is equal and each instance is stored in its native data structure. No specific transformation operations are required when the data is uploaded. A unified view is generated to accommodate different models of data together. Besides, there are very few systems that would support as many models (relational, XML, JSON, RDF, and property graphs) as `MultiCategory` supports.

6 CONCLUSION AND FUTURE WORK

`MultiCategory` is a tangible system that is applying category theory to model and query multi-model data. We implement a query language with a functional programming language. We visualize the category theoretical constructions, i.e. schema and instance categories, and query processing to show connections between theory and applications. Note that `MultiCategory` is not yet a full-fledged database system and due to its implementation it is also not a big data system.

The schema category and instance category are fixed and predefined manually. In the future, we consider automatically generating this unified category view based on input data sets. We have been developing the theoretical framework so that it can define multi-model joins and data, schema, and query transformations.

REFERENCES
[1] Torsten Grust. 1999. *Comprehending Queries*. Ph.D. Dissertation. Universitaet Konstanz, Konstanz.
[2] Zhen Hua Liu, Jiaheng Lu, Dieter Gawlick, Heli Helskyaho, Gregory Pogossiants, and Zhe Wu. 2018. Multi-model Database Management Systems - A Look Forward. In *Polystores VLDB 2018 Workshops*. 16–29.
[3] Jiaheng Lu and Irena Holubová. 2019. Multi-model Databases: A New Journey to Handle the Variety of Data. *ACM Comput. Surv.* 52, 3 (2019), 55:1–55:38.
[4] Saunders MacLane. 1971. *Categories for the working mathematician*. Springer, New York, NY. https://doi.org/10.1007/978-1-4612-9839-7
[5] Andrey Mokhov. 2017. Algebraic Graphs with Class (Functional Pearl). In *Proceedings of the 10th ACM SIGPLAN International Symposium on Haskell* (Oxford, UK) *(Haskell 2017)*. New York, NY, USA, 2–13. https://doi.org/10.1145/3122955.3122956
[6] David Spivak. 2014. Category Theory for the Sciences. (2014).
[7] David I. Spivak. 2010. Functorial Data Migration. *CoRR* abs/1009.1166 (2010). arXiv:1009.1166 http://arxiv.org/abs/1009.1166
[8] Valter Uotila. 2021. MultiCategory Documentation and System Codes. https://multicategory.github.io/, https://git.io/JvPqM. Accessed Jul. 19, 2021.
[9] Valter Uotila and Jiaheng Lu. 2021. MultiCategory demo video. https://youtu.be/uceIi91AGsg. Accessed Jul. 19, 2021.