

Facetize: An Interactive Tool for Cleaning and Transforming Datasets for Facilitating Exploratory Search

Anna Kokolaki and Yannis Tzitzikas
Institute of Computer Science, FORTH-ICS, GREECE, and
Computer Science Department, University of Crete, GREECE
kokolaki@ics.forth.gr, tzitzik@ics.forth.gr

December 27, 2018

Abstract

There is a plethora of datasets in various formats which are usually stored in files, hosted in catalogs, or accessed through SPARQL endpoints. In most cases, these datasets cannot be straightforwardly explored by end users, for satisfying recall-oriented information needs. To fill this gap, in this paper we present the design and implementation of *Facetize* an editor that allows users to transform (in an interactive manner) datasets, either static (i.e. stored in files), or dynamic (i.e. being the results of SPARQL queries), to datasets that can be directly explored effectively by themselves or other users. The latter (exploration) is achieved through the familiar interaction paradigm of *Faceted Search* (and *Preference-enriched Faceted Search*). Specifically in this paper we describe the requirements, we introduce the required set of transformations, and then we detail the functionality and the implementation of the editor *Facetize* that realizes these transformations. The supported operations cover a wide range of tasks (selection, visibility, deletions, edits, definition of hierarchies, intervals, derived attributes, and others) and *Facetize* enables the user to carry them out in a user-friendly and guided manner, without presupposing any technical background (regarding data representation or query languages). Finally we present the results of an evaluation with users. To the best of your knowledge, this is the first editor for this kind of tasks.

Keywords: Exploratory Search, Data Exploration, Data Transformation, Data Cleaning, Faceted Search

1 Introduction

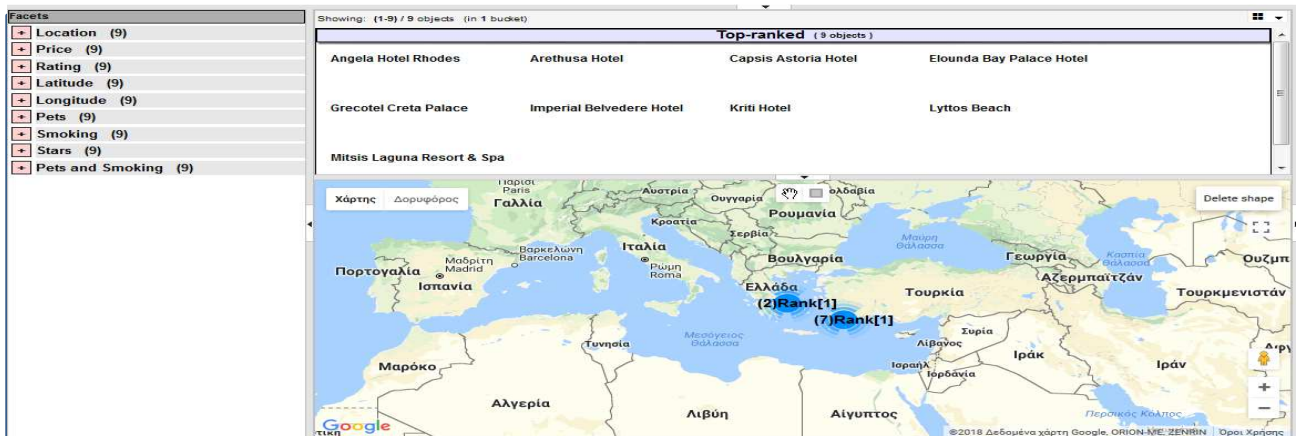
Although there is a tendency to publish data as Linked Data, the majority of datasets is still represented in simple data file formats such as CSV and TSV. Such datasets cannot be easily exploited by end users. A user can view directly these CSV files using a text editor, or spreadsheet application, or on the other extreme, one should build dedicated applications for offering a more user friendly exploration service and experience. To tackle this difficulty, our objective is to provide a general purpose solution that enables users to directly explore such datasets and/or to configure the way they are explored, for better supporting recall-oriented information needs. To this end, we rely on a quite familiar access method, specifically on *Faceted Search*, which is the defacto standard in e-commerce and booking applications. However, a straightforward loading of such datasets in a faceted search system will not always result to a satisfying solution, since additional tasks are usually required. This includes deciding (a) the parts of the dataset that should be explorable, (b) the attributes that should be visible and their order, (c) the transformations and/or enrichments that should be done, (d) the groupings (hierarchical or not) of the values that should be made, (e) the addition of derived attributes, and others. For this reason in this paper we present the design and implementation of an editor, called *Facetize*, that allows the user to carry out these tasks in a user-friendly and guided manner, without the need to write any script or use any programming language.

To grasp the idea, Figure 1(a) shows a CSV file containing information about 9 hotels each described with 10 properties. Now suppose that one would like from this dataset to produce an explorable dataset, as sketched in Figure 1(b), with the following specific requirements:

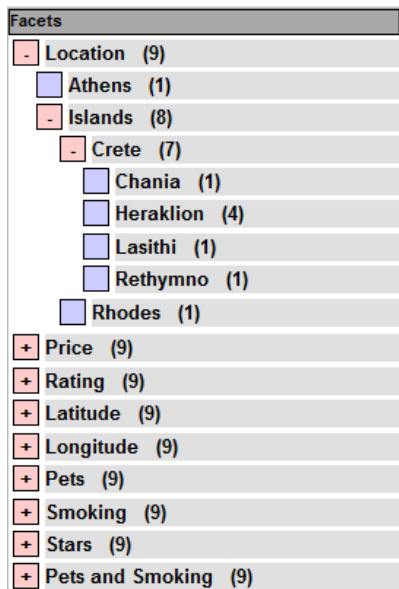
- (1) The dataset must contain hotels, only located in Greece (so one row should be deleted).
- (2) The dataset should be enriched with a new hotel with values: *Mitsis Laguna Resort & Spa, Heraklion, 25.371359, 35.307237, 5,385,115,8.7, allowed, not allowed*
- (3) The properties with name *Longitude* and *Latitude*, must be marked as properties that contain geographic information for longitude and latitude respectively, for enabling an exploration system to show the location of each hotel on a map.
- (4) The hotel entities should take as names the value of the property *Name*.
- (5) The value *Iraklio* in property *Location*, must be replaced with the value *Heraklion*, in all entities that contain that value.

1	Name, Location, Longitude, Latitude, Stars, Rooms, Price, Rating, Smoking, Pets
2	Imperial Belvedere Hotel, Heraklion, 25.401855, 35.3042, 4, 330, 106, 8.6, allowed, allowed
3	CitizenM Paris Gare de Lyon, Paris, 2.371776, 48.843585, 4, 300, 94, 8.9, allowed, not allowed
4	Arethusa Hotel, Athens, 23.733043, 37.975271, 3, 87, 50, 8.2, allowed, not allowed
5	Angela Hotel Rhodes, Rhodes, 28.220095, 36.450108, 3, 220, 22, 8.0, allowed, not allowed
6	Lyttos Beach, Iraklio, 25.3523, 35.3352, 4, 250, 156, 9.0, allowed, not allowed
7	Kriti Hotel, Chania, 24.025017, 35.515731, 3, 84, 69, 9.1, not allowed, not allowed
8	Elounda Bay Palace Hotel, Lasithi, 25.723416, 35.261350, 5, 200, 264, 8.8, allowed, allowed
9	Grecotel Creta Palace, Rethymno, 24.518737, 35.36092, 5, 300, 157, 8.5, allowed, not allowed
10	Capsis Astoria Hotel, Iraklio, 25.136372, 35.338397, 4, 312, 59, 8.1, allowed, allowed


(a)




(b)



(c)



(d)



(e)

Figure 1: Indicative Scenario

- (6) The property *Rooms* should not appear in list of facets: it should either be hidden or deleted.
- (7) A new property with name *Pets and Smoking* should be created, with values as shown in the bottom part of Figure 1(e), and each hotel should be associated with the right value.
- (8) The values in property *Location* should be organized hierarchically as shown in Figure 1(c).
- (9) The values in property *Price*, must be organized in interval hierarchies, as shown in Figure 1(d).
- (10) The order of the facets should be as shown in Figure 1.

With the approach that we describe in this paper, and the tool *Facetize*, a naive user can make all these transformation in an easy manner and produce an output file that is directly loadable to a system that supports Faceted Search (or Preference Enriched Faceted Search), like *Hippalus*, which the users can directly explore using a GUI as shown in Figure 1(b). The overall process is sketched in Figure 2.

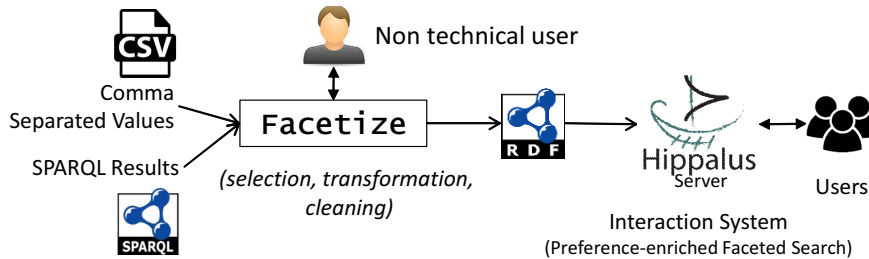


Figure 2: The process

In a nutshell, the key contributions of this paper are: (a) we identify the basic requirements for this kind of tasks, (b) we present a set of operations (for selecting, transforming, cleaning or enriching a dataset), (c) we describe the design and implementation of *Facetize* that supports these operations over static datasets (i.e. stored in files) as well as dynamic (i.e. being the results of SPARQL queries), and (d) we report the results of a task-based evaluation with users. In comparison to related systems, *OpenRefine* is probably the more relevant system. One distinctive feature of *Facetize* is that it supports the creation of hierarchies and numeric intervals and these features are crucial for managing the complexity of large datasets, and producing datasets that can be easily explored. Moreover, *Facetize* can fetch data directly from SPARQL endpoints, making it appropriate for dynamic datasets.

The rest of this paper is organized as follows: Section 2 describes background information and related work, Section 3 identifies the requirements, Section 4 presents the *Facetize* editor, Section 5 presents the results of the evaluation, and finally Section 6 concludes and identifies issues for further work and research.

2 Background and Related Work

2.1 Datasets

Although there is a tendency to publish data as Linked Data, e.g. see [10], a publishing method that facilitates data linking, aggregation and integration, and is beneficial for the preservation of data (e.g. see [19]), a big percentage of datasets is represented only in simple data file formats such as CSV. For instance, in *Govdata.de*, one of the most popular open data portals, almost 26.76% of its files are in CSV format¹. In *Data.gov.uk*, the official open data portal of the UK Government, 18.70% of them are in CSV format. In the open data portal of Italy, *Dati.gov.it*, around 29.61% files in CSV format, while in the european data portal.eu (EDP) around 11.68% of the files are CSV files.

2.2 Faceted Search

Faceted Exploration (or *Faceted Search*) is a widely used interaction scheme for Exploratory Search. It is the de facto query paradigm in e-commerce [13, 15]. In a short (and rather informal) way we could define it as a *session-based interactive method for query formulation (commonly over a multidimensional information space) through simple clicks that offers an overview of the result set (groups and count information), never leading to empty results sets*. Faceted search has been generalized also for RDF datasets [18].

¹From open data formats statistics of the main catalogs (November 2016)

2.3 PFS: Preference-enriched Faceted Search

Preference-enriched Faceted Search [20], for short PFS, is an extension of Faceted Search that supports preferences. PFS offers actions that allow the user to order facets, values, and objects using *best*, *worst*, *prefer to* actions (i.e. relative preferences), *around to* actions (over a specific value), or actions that order them lexicographically, or based on their values or count values. Furthermore, the user is able to *compose* object related preference actions, using *Priority*, *Pareto*, *Pareto Optimal* (i.e. skyline) and other. The distinctive features of PFS is that it allows expressing preferences over attributes whose values can be hierarchically organized (and/or multi-valued), it supports preference inheritance, and it offers scope-based rules for resolving automatically the conflicts that may arise. As a result the user is able to restrict his current focus by using the faceted interaction scheme (hard restrictions) that lead to non-empty results, and rank according to preference the objects of his focus. Recently, PFS has been used in various domains, e.g. for offering a flexible process for the identification of fish species [17], as a Voting Advice Application [16] and it has been expanded with geographic anchors for being appropriate for the exploration of datasets that contain also geographic information [8]. Applications of the model in the context of spoken dialogue systems are also emerging [11].

In our work we decided to use Hippalus which is a publicly accessible web system that implements the PFS interaction model. The information base that feeds Hippalus is represented in RDF/S² using a schema adequate for representing objects described according to dimensions with hierarchically organized values.

2.4 Related Work

The existence of anomalies in real-world data motivates the development and application of data cleansing methods. This is a hot topic, e.g. see [1] for a recent overview, and there are several tools for such tasks including, DataPreparator³, Potter's Wheel [12], WinPure [6], OpenRefine⁴, Karma[5], DataX-Former [2], Katara[3], Data Wrangler [7] and others. Existing data cleansing approaches mostly focus on the transformation of data, the elimination of duplicates, syntax and lexical errors detection. They also offer methods for create and edit attributes using a number of functions. In general, the users can see the transformations in each step of editing and they can export the data after applying the transformations in various file formats, such as Excel, CSV, RDF and others.

We could say that the the available data cleaning solutions and tools ([1], [9]) belong to one or more of the following four categories:

- *Rule-based detection algorithms* that can be embedded into frameworks and the user can specify a collection of rules that clean data will obey and the tool will find any violations.
- *Pattern enforcement and transformation tools* such as OpenRefine⁵, Data Wrangler [7], Trifacta⁶, Katara[3], and DataX-Former [2]. These tools discover patterns in the data, either syntactic (e.g., OpenRefine and Trifacta) or semantic (e.g., Katara), and use these to detect errors (cells that do not conform with the patterns). They can also be used to change data representation and expose additional patterns.
- *Quantitative error detection algorithms* that expose outliers, and glitches in the data.
- *Record linkage and de-duplication algorithms* for detecting duplicate data records, such as the Data Tamer system [14]. These tools perform entity consolidation when multiple records have data for the same entity. Conflicting values for the same attribute can be found, indicating possible errors.

3 Requirements

We dichotomize requirements to transformation requirements and to workbench-related requirements. They are presented in a more structured way, later in Table 1 of Section 5.1.

3.1 Transformation Requirements

Even from the running example in the introductory section it is evident that for preparing a dataset appropriate for exploration one should be able to define the facets that should be visible, and their order. The user should be able to specify the type of the terms of a facet, e.g. identifier, integer, float, string, longitude and latitude. By defining the type of a facet as identifier, the entities will take as names their corresponding values in this facet, and all values in this facet should be distinct. Moreover it should be possible to define new facets whose terms are derived by applying

²<http://www.w3.org/TR/rdf-schema/>

³<http://www.datapreparator.com/>

⁴openrefine.org

⁵<http://openrefine.org/>

⁶<http://www.trifacta.com>

functions over the terms of other facets. For avoiding cluttering the GUI, and aiding the exploration, it should be possible to define hierarchical groupings of terms, as well as intervals to numerical values. Finally, it should be able to delete all rows with a specific value or a condition, to replace each distinct value in a facet with a new one in each row of the dataset, and the user should also be able to add, edit and delete individual rows.

3.2 Workbench-related Requirements

The notion of project, should be supported, allowing the user to create, open, edit and save the changes on disc. Moreover the system should be able to keep the history of transformations that have been applied and offer to the user the ability to undo the desired ones (and redo). Due to data dynamicity, the user should be able to open a project and change the input dataset, i.e. by giving a more recent version of the dataset, or the file with the hierarchical information. This is very important for datasets that change over time, since we would not like to lose the transformations that have been defined for a past version of the dataset. Finally, it should be possible to export the transformed dataset in RDF according to a schema that is compliant with Hippalus, for enabling the straightforward loading by Hippalus.

4 The Tool Facetize

Supported Transformations. In brief, Facetize supports all transformation requirements described in §3. As regard hierarchies, the user is able to move values in existing hierarchies, and add intermediate terms. In case of string values, the user is able to use functions to define expressions for creating groups of values with same prefix and values that start with the same range of letters for aiding the exploration. For numerical facets, the specification of intervals is supported and the user can define linear intervals, logarithmic intervals, and intervals with specific bounds. For example, Figure 3.a shows the list of facets of the running example described in the introductory section. Notice that facet *Location* contains 6 distinct values. By left clicking on the term *Chania*, and selecting the option *Add Parent* from the menu, a (bootstrap.js) modal dialogue appears with a text field for adding the parent value (see Fig. 3.b). If for example the user inserts the value *Crete*, the hierarchy Chania/Crete will be created in facet *Location* (as shown in Fig. 3.c).

In terms of the data error types, Facetize assists the user to detect and remove them. Specifically, in case of *Rule*, *Pattern violations* and *Outliers* errors in the dataset, the user can edit the values that cause the inconsistencies and replace them with other values. In case of *Duplicates*, the user can delete or edit duplicate rows.

The Notion of Project. Facetize supports the notion of project, allowing the user to gradually configure the desired presentation and also to update (refresh) the underlying dataset without losing the transformations that have been defined. Each project has a name, an input dataset and a configuration of the transformations that should be applied for producing an output dataset that is suitable for exploration by Hippalus. In brief, Facetize supports three types of project: (a) Project with Single File Dataset, (b) Project with Multiple Files, (c) Project with SPARQL Query.


(a) Project with Single File Dataset The first row in a dataset file is the header row. Each column in this row must contain a Property name (the name of the corresponding facet). Hierarchies can be specified (a) internally, (b) externally, or (c) by both methods. According to method (a), the user can specify the complete path from the leaf term to its root, e.g. Mazda/Japanese/Asian. According to method (b), the hierarchy is specified through a configuration file that contains hierarchical information about some or all the values of the tabular data file, i.e. each line contains hierarchical information about one term using the following syntax: name + path from the specific term to the root of the hierarchy separated with slashes (/), e.g. Mazda/Japanese/Asian/Manufacturer, where Manufacturer is the name of the facet that the mentioned hierarchy belongs. The order of rows is not important and the import of a configuration file in each project is optional, since the user can create the desired hierarchies through the editor.

(b) Project with Multiple Files. Facetize supports input from multiple CSV or TSV files, located at a specified folder. This structure is convenient for enriching the dataset with more dimensions. The files that such project can contain are:

Object Id File: The object id file is specified explicitly by the user and is the one that holds all the ids of the objects of this project. It is a CSV/TSV file, with a single column, containing the names of the entities.

Dimension (Property) Files. These files have 2 columns: the first is for the object id while the second for the value that the specified entity has for this property. The hierarchies can be expressed anywhere in the file (at the beginning, end, or mixed) and they can be straightforwardly represented: we just have to put the child at the first column and its parent at the second column. The translator detect that a property is hierarchical when the name in the first column is not a valid object id.

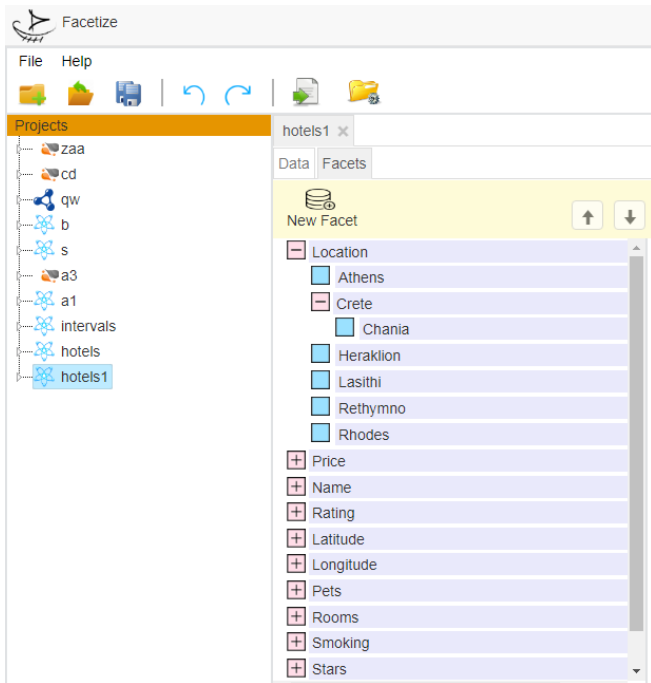
(c) Project with SPARQL Query. For exploiting the wealth of Linked Data and the available SPARQL endpoints, Facetize allows the user to specify the address of a SPARQL endpoint and the SPARQL SELECT query to be sent,



(a)



(b)



(c)

Figure 3: (a) Click on a term and open menu, (b) Set parent *Crete* for term *Chania* in the hierarchy, (c) Hierarchy Chania/*Crete* has been created

e.g. the query `SELECT ?Speed ?Price ?Weight` will lead the tool to create 3 properties: Speed, Price and Weight. Moreover the user can save queries and endpoints in the favourites list, for re-running them in the future.

Project Refreshing. Every time the user applies a transformation to the dataset of a Facetize project, that transformation is saved in JSON format in a file in the project's folder. For each transformation, its type, as well as all related information, are saved, for enabling reapplying the transformation in the future.

Implementation. Facetize is a Web Application, based on the architecture of client-server. The presentation layer (*front end*) is implemented using HTML, CSS for rendering elements appropriately on screen, and JavaScript libraries such as `Bootstrap.js` to create responsive content. The data access layer (*back end*), is implemented using Java Servlets technologies and Jena Semantic Web framework⁷. To run the editor the user should have installed at his computer Java Runtime Environment (JRE) and a Java Servlet Container.

Table 1: Features of data cleansing systems

Category	Feature	DataPreparator	Potters Wheel	WinPure Clean	Karma	OpenRefine	Facetize
Import/Export	Import/Export text files, EXCEL Import/Export Data from databases Import/Export RDF files Export R2RML/JSON files Execute SPARQL queries to retrieve data from a source	✓ ✓		✓ ✓	✓ ✓	✓ ✓	✓ ✓ ✓
Data cleaning: Rows	Deletion of rows Removal of duplicate rows Addition of rows Removal of empty rows			✓	✓	✓ ✓ ✓	✓ ✓ ✓ ✓
Data cleaning: Values	Deletion of records containing missing values Character removal, text replacement, date conversion Value editing Impute missing values Creation of value hierarchies Creation of intervals for arithmetic values Correction of bad values using string similarity metrics Clustering of the values in a property that contain specific characters (e.g. same prefix)	✓ ✓ ✓		✓	✓ ✓	✓ ✓ ✓	✓ ✓ ✓ ✓ ✓ ✓ ✓
Data cleaning: Columns	Delete/move attributes Filtering Renaming of Columns Set column types Creation of new columns using expression Split of columns Deletion of columns Rearrangement of columns Hiding of columns Addition/merging of columns Split records into columns	✓	✓ ✓	✓	✓ ✓ ✓ ✓	✓ ✓ ✓ ✓	✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓
Editor Facilities	Undo/redo transformations Display the distinct property values and their frequency Ability to reapply same range of transformations		✓			✓ ✓	✓ ✓ ✓
Other Export Functions	Export macro or a C program, or a Perl program Extraction of Entities		✓		✓		

⁷<http://jena.apache.org/>

5 Evaluation

5.1 Comparison with Related Tools

Table 1, shows the features of each data cleansing system that was mentioned in §2.4, as well as of Facetize. We observe that one distinctive feature of Facetize is that it supports the creation of hierarchies and numeric intervals and these features are crucial for managing the complexity of large datasets, and producing datasets that can be easily explored. Moreover, Facetize can fetch data directly from SPARQL endpoints, making it appropriate for dynamic datasets. Finally, users have the ability to reapply the same range of transformations to an existing project and this is very convenient in case they want to refresh the dataset of a project (with a new version of the dataset) and would like to apply the same transformations.

We tried to carry out scenario of the introductory section, using the tools mentioned in Table 1, i.e. DataPreparator, Potters Wheel, WinPure Clean, Karma, OpenRefine, Facetize. We wanted to check what transformations were feasible to perform and how much time it took to make them. Of course none of these tools (apart from Facetize) can produce an RDF file that is directly loaded to Hippalus, however we wanted to test the rest aspects of the scenario. With DataPreparator we managed to perform requirements 5 and 6 (as numbered in the introductory section) in 1 (one) minute. With Potters Wheel we managed to perform requirements 6, 7 in 2 (two) minutes. Using WinPure Clean we did not manage to satisfy any of the requirements. With Karma we managed to perform requirements 2, 5, 7 in 3 (three) minutes. Finally with OpenRefine we managed to perform requirements 1, 2, 5, 6, 7, 10 in 6 (six) minutes, i.e. it is the tool that covered most of the requirements.

5.2 Task-based Evaluation with Users

We conducted a task-based evaluation with users for getting general and specific feedback, and for testing the usability of Facetize and user satisfaction. We used the scenario described in Section 1 with the dataset of hotels of Figure 1(a). We prepared a simple text tutorial of 45 slides, a questionnaire and a file with the description of the aforementioned scenario. We invited by email various persons to participate in the evaluation voluntarily. The users were asked to carry out the tasks and to fill the questionnaire. It was stated to them clearly, that they should not rush up. The participation to this evaluation was optional. Twenty persons (20), eventually participated. The number was sufficient for our purposes, since according to [4] 20 evaluators are enough for getting more than 95% of the usability problems of a user interface. In numbers, the participants were 11 (55%) female and 9 (45%) male, with ages ranging from 18 to 64 years. As regards occupation and skills, users have studied Computer Science. In detail, 5 (25%) were undergraduate students, 12 (60%) of them postgraduate students and 3 (15%) computer engineers and researchers.

We used a questionnaire that users had to fill it in, and then send it back to us by email. The questionnaire is shown below, enriched with the results of the survey in the form of percentages written in bold.

- *How many mistakes did you make ?* No mistake (**20%**), 1-3 mistakes (**45%**), 3-6 mistakes (**15%**), 6-9 mistakes (**5%**), More than 9 mistakes (**15%**)
- *How much time you spent for carrying out the task?* 6-7 minutes (**0%**), 7-10 minutes (**20%**), 10-15 minutes (**25%**), 15-20 minutes (**45%**), More than 20 minutes (**10%**)
- *Have you ever used data transforming systems like Facetize?* Yes, I have used systems like Facetize. (**25%**), No, I have not used any system like Facetize. (**75%**)
- *What was the final outcome of the scenario?*
I successfully completed the entire scenario, and the data are displayed exactly as it was requested. (**70%**), I did the whole script, except for some requirements that I did not perform successfully and did not receive all the correct results from Hippalus. List the numbers (in the order in which they appear) of the requirements you did not successfully execute. (**20%**), I quitted the task in the requirement with number..... (fill in the number of requirement you were trying to perform or whatever you did). (**10%**)
- *How would you rate Facetize as a data transformation system?* Very Useful (**40%**), Useful (**60%**), Little Useful (**0%**), Not Useful (**0%**)

The results are quite satisfactory, since 14 (70%) users managed to complete the task successfully, 4 (20%) executed the scenario, except from some requirements that they probably did not perform successfully, and only 2 (10%) quitted the task. For those users who did not get the right results, or quitted the process, they did not understand the system and used it wrongly. As regards the overall rating, 12 (60%) users rated the system *Useful*, while (40%) *Very Useful*. Here we include only two plots from the analysis of the results. In Figure 4 we can see the "Dedicated time-Success percentage" where we observe that that users that spent 10-15 minutes were the most successful. From the "Errors-Success percentage" we can see that the users who made 1-3 errors were the most successful; 35% success, those who made no error had 20% success, those with 3-6 errors had 10% success, and finally those who made more than 9 errors had 5% success.




Figure 4: Top: Dedicated time-Success percentage. Bottom: Errors-Success percentage

6 Concluding Remarks

To conclude, there is a demand for systems that help users with no particular technical background to clean and apply transformations on plain datasets for turning them easily explorable. However, a straightforward loading of such datasets in a faceted search system will not always result to a satisfying solution, since additional tasks are usually required. This includes deciding (a) the parts of the dataset that should be explorable, (b) the attributes that should be visible and their order, (c) the transformations and/or enrichments that should be done, (d) the groupings (hierarchical or not) of the values that should be made, (e) the addition of derived attributes, and others. We presented the design and implementation of an editor, called *Facetize*, that allows the user to carry out these tasks in a user-friendly and guided manner, without presupposing any technical background about the data representation language or the query language. *Facetize* can be construed as an authoring environment for setting up exploration services over static files (represented in various formats) or results produced by querying SPARQL endpoints. It is worth noting that the authoring environment, supports the notion of project, allowing the user to gradually configure the desired presentation and also to update (refresh) the underlying dataset without losing the transformations that they have been defined. We evaluated the *Facetize* prototype with users over real data and the results of the evaluation were positive. Overall the combination of *Facetize* with *Hippalus* offer a handy approach for supporting exploratory search over static and dynamic datasets.

We plan to make *Facetize* publicly available soon. Directions that are worth further work and research include methods for predicting the missing values. Such functionality is helpful in cases where there are a lot of missing values that we would like to fill in the dataset at hand, and various models could be employed for that purpose (e.g. dependence tree, Naive Bayes model, etc). Another direction is to devise methods that can spot possible errors or misspellings by exploiting string similarity-based clustering methods.

References

- [1] Ziawasch Abedjan, Xu Chu, Dong Deng, Raul Castro Fernandez, Ihab F Ilyas, Mourad Ouzzani, Paolo Papotti, Michael Stonebraker, and Nan Tang. Detecting data errors: Where are we and what needs to be done? *Proceedings of the VLDB Endowment*, 9(12):993–1004, 2016.
- [2] Ziawasch Abedjan, John Morcos, Ihab F Ilyas, Mourad Ouzzani, Paolo Papotti, and Michael Stonebraker. Dataformer: A robust transformation discovery system. In *Data Engineering (ICDE), 2016 IEEE 32nd International Conference on*, pages 1134–1145. IEEE, 2016.
- [3] Xu Chu, John Morcos, Ihab F Ilyas, Mourad Ouzzani, Paolo Papotti, Nan Tang, and Yin Ye. Katara: A data cleaning system powered by knowledge bases and crowdsourcing. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 1247–1261. ACM, 2015.
- [4] Laura Faulkner. Beyond the five-user assumption: Benefits of increased sample sizes in usability testing. *Behavior Research Methods, Instruments, & Computers*, 35(3):379–383, 2003.
- [5] Shubham Gupta, Pedro Szekely, Craig A Knoblock, Aman Goel, Mohsen Taheriyani, and Maria Muslea. Karma: A system for mapping structured sources into the semantic web. In *Extended Semantic Web Conference*, pages 430–434. Springer, 2012.

- [6] Hamed Ibrahim Housien, Zhang Zuping, and Zainab Qays Abdulhadi. A comparison study of data scrubbing algorithms and frameworks in data warehousing. *International Journal of Computer Applications*, 68(25), 2013.
- [7] Sean Kandel, Andreas Paepcke, Joseph Hellerstein, and Jeffrey Heer. Wrangler: Interactive visual specification of data transformation scripts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3363–3372. ACM, 2011.
- [8] Panagiotis Lionakis and Yannis Tzitzikas. Pfsgeo: Preference-enriched faceted search for geographical data. In *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*, pages 125–143. Springer, 2017.
- [9] Jonathan I Maletic and Andrian Marcus. Data cleansing. In *Data Mining and Knowledge Discovery Handbook*, pages 21–36. Springer, 2005.
- [10] Michalis Mountantonakis and Yannis Tzitzikas. High performance methods for linked open data connectivity analytics. *Information (2078-2489)*, 9(6), 2018.
- [11] Alexandros Papangelis, Panagiotis Papadakos, Margarita Kotti, Yannis Stylianou, Yannis Tzitzikas, and Dimitris Plexousakis. Ld-sds: Towards an expressive spoken dialogue system based on linked-data. In *Search Oriented Conversational AI, SCAI 17 Workshop (co-located with ICTIR 17)*, 2017.
- [12] Vijayshankar Raman and Joseph M Hellerstein. Potter's wheel: An interactive data cleaning system. In *VLDB*, volume 1, pages 381–390, 2001.
- [13] Giovanni Maria Sacco and Yannis Tzitzikas. *Dynamic taxonomies and faceted search: theory, practice, and experience*, volume 25. Springer Science & Business Media, 2009.
- [14] Michael Stonebraker, Daniel Bruckner, Ihab F Ilyas, George Beskaes, Mitch Cherniack, Stanley B Zdonik, Alexander Pagan, and Shan Xu. Data curation at scale: The data tamer system. In *CIDR*, 2013.
- [15] Daniel Tunkelang. Faceted search. *Synthesis lectures on information concepts, retrieval, and services*, 1(1):1–80, 2009.
- [16] Y. Tzitzikas and E. Dimitrakis. Preference-enriched faceted search for voting aid applications. *IEEE Transactions on Emerging Topics in Computing*, PP(99):1–1, 2016.
- [17] Yannis Tzitzikas, Nicolas Bailly, Panagiotis Papadakos, Nikos Minadakis, and George Nikitakis. Using preference-enriched faceted search for species identification. *International Journal of Metadata, Semantics and Ontologies*, 11(3):165–179, 2016.
- [18] Yannis Tzitzikas, Nikos Manolis, and Panagiotis Papadakos. Faceted exploration of RDF/S datasets: a survey. *Journal of Intelligent Information Systems*, 2016.
- [19] Yannis Tzitzikas and Yannis Marketakis. *Cinderella's Stick - A Fairy Tale for Digital Preservation*. Springer, 2018.
- [20] Yannis Tzitzikas and Panagiotis Papadakos. Interactive exploration of multidimensional and hierarchical information spaces with real-time preference elicitation. *Fundamenta Informaticae*, 20:1–42, 2012.