

Towards Semantically Enhanced Data Understanding

Markus Schröder

Smart Data & Knowledge Services Dept., DFKI GmbH
Computer Science Dept., TU Kaiserslautern
Kaiserslautern, Germany
markus.schroeder@dfki.de

Jörn Hees

Smart Data & Knowledge Services Dept., DFKI GmbH
Computer Science Dept., TU Kaiserslautern
Kaiserslautern, Germany
joern.hees@dfki.de

Christian Jilek

Smart Data & Knowledge Services Dept., DFKI GmbH
Computer Science Dept., TU Kaiserslautern
Kaiserslautern, Germany
christian.jilek@dfki.de

Andreas Dengel

Smart Data & Knowledge Services Dept., DFKI GmbH
Computer Science Dept., TU Kaiserslautern
Kaiserslautern, Germany
andreas.dengel@dfki.de

ABSTRACT

In the field of machine learning, data understanding is the practice of getting initial insights in unknown datasets. Such knowledge-intensive tasks require a lot of documentation, which is necessary for data scientists to grasp the meaning of the data. Usually, documentation is separate from the data in various external documents, diagrams, spreadsheets and tools which causes considerable look up overhead. Moreover, other supporting applications are not able to consume and utilize such unstructured data.

That is why we propose a methodology that uses a single semantic model that interlinks data with its documentation. Hence, data scientists are able to directly look up the connected information about the data by simply following links. Equally, they can browse the documentation which always refers to the data. Furthermore, the model can be used by other approaches providing additional support, like searching, comparing, integrating or visualizing data. To showcase our approach we also demonstrate an early prototype.

KEYWORDS

Data Understanding, Data Mining, CRISP-DM, Semantic Web, Data Catalog, Deep Linking, RDF

Markus Schröder, Christian Jilek, Jörn Hees, and Andreas Dengel. 2018. Towards Semantically Enhanced Data Understanding. In *preprint*, 4 pages.

1 INTRODUCTION

The purpose of data mining is to discover patterns in datasets using various methods from multiple research fields. In order to guide data scientists in this complicated task, several process models for data mining exist [9]. They all have in common that they emphasise on understanding the data before analysing it. For instance, the popular Cross-Industry Standard Process Model for Data Mining (CRISP-DM) [3] includes a “Data Understanding” phase in order to guide analysts in getting initial findings from data, which in particular means collecting, describing, exploring and verifying datasets.

Usually, analysts ask many questions when facing unfamiliar data in the understanding phase. For example, they would like to

know what quality issues exist, what data dependencies are known, which values occur most frequently, what columns exist, how data can be accessed, etc. Answering those questions immediately and specifically in conjunction with the corresponding data facilitates its comprehension. However, it requires that data scientists have to read or compose various reports describing the data. The containing information helps them to grasp the contents of the data and to ease the incremental, repeating and time-consuming process of understanding unknown datasets.

Nevertheless, data scientists are confronted with several typical challenges in this process. To illustrate this, we think of an imaginary data scientist who has to analyse data from some foreign company. After the analyst begins to understand the project’s objective and goals, the unknown dataset is investigated.

The heterogeneous data sources are stored using different technologies, each forcing the usage of other tools to access and browse them. During exploration, the analyst notices that some parts of the datasets are not documented. A couple of notations for tables, columns and values are incomprehensible because they contain technical terms, abbreviations, or identifiers. After several iterations with his contact person in the other company, all necessary information is gathered and written down in a separate document. But the data scientist also investigates the datasets on his own using various tools and visualizations [8] that produce multimedia results, like reports, lists and diagrams.

For some datasets, the company provides several documents that explain, among many other things, the meaning of the tables’ columns. But the employee needs a lot of time to find the corresponding documents and the pages explaining a certain table. During the project, this information has to be looked up repetitively which proves to be very time-consuming, all because data and its documentation are physically and logically separated. Additionally, because of the large amount of different kinds of documentation, the analyst loses track of which information belongs to which dataset content. The data scientist would like to use helping tools, but the unstructured reports are rather difficult for algorithms to process.

In conclusion, the scenario shows the main problem: data and its documentation are not logically “connected”. As a result, this hinders data scientists in freely moving within and between those two information spaces and consequently causes considerable overhead in the data understanding process.

2 APPROACH

To ease the understanding of data, we would like to support data scientists with an easy way to look up corresponding documentation about unknown data. In the following, we will step-by-step explain how our proposed methodology enables this.

A first solution would be to include the documentation into the data source. In fact, some information systems allow to add textual comments for their data elements (e.g. PostgreSQL supports a COMMENT command [5]). Another approach in this regard is CCSV [10] which combines content and context (in form of semantic data) in regular CSV files. However, such practices have several limitations and drawbacks. The data source has to be modified to include comments and they only reside in the system where they are created. As a consequence, in order to make them available for others, the whole data source has to be shared. Often, the additional information is limited by size and restricted to a certain format (mostly textual).

Instead of locally stored annotations, we would like to save them externally. This has the benefit that they are not limited by the data source’s capabilities. If annotations are modelled separately, they must refer to the piece of data they annotate. Usually, analysts use labels in their reports to implicitly address data elements like tables or columns, but this causes the look up overhead which was discussed earlier. Instead, we would like to have a simple resolvable hyperlink into tabular data. Actually, such a deep link should be able to refer to a specific part of a table. That way, a look up becomes as easy as browsing a web site. This approach is comparable with URI fragment identifiers¹ which refer to subordinate resources of primary ones as known from HTML. In the context of tabular data, fragment identifiers have been proposed for CSV files². We already demonstrated [12] that such deep links are a convenient way to address certain locations within desktop resources.

With this method, data scientists are now able to use hyperlinks in their reports to refer to the described data. But, as we have already noticed, reports in free text form are rather difficult to process. That is why our approach suggests that the documentation is (a) modelled in a machine-processible form or that (b) the documents’ contents is at least similarly addressable. The latter means that we are able to refer to parts of documents about data (like reports, presentations or spreadsheets). With this approach it is now possible to interlink data with its documentation. We utilize the Resource Description Framework³ (RDF) which allows to link resources (identified by hyperlinks) using triples. As a result, the look up of associated documentation reduces to a simple traversal in a semantic graph.

Although we create another separate model, this approach does not duplicate the datasets. Often, legacy data is fully converted to RDF before it is further processed in an RDF graph. For instance, csv2rdf⁴ suggests how tabular data should be converted. Other approaches bypass this by generating RDF on-demand during runtime, similarly as we do. For example, D2RQ [2] focuses on making relational databases queryable by SPARQL⁵ without copying the

database. However, the application initially needs a mapping which defines how database contents is mapped to RDF. In comparison, our approach only requires that the data source is addressable in order to make statements about it. Because of such referencing, the data stays where it is and details are only requested when necessary.

Having a semantic representation of the data and its documentation yields additional modelling possibilities. The first one is cross-referencing, which allows the reciprocal linking of two or more datasets or documents. This is helpful when information is spread across various sources, for instance, when a CSV file and a database table stores information about the same entity which is also mentioned in a reporting document. In addition, ontologies [6] can be used to further model the data’s domain. The same way, we are able to link the legacy data to resources from DBpedia [1], which contains structured information from Wikipedia. As an example, a column containing motorcar identifiers could link to DBpedia’s car entity⁶ to clarify its content. Such approaches are used in Semantic Data Mining [4]: a research field that incorporates domain knowledge into usual data mining processes.

The modelled RDF data is queryable with SPARQL which already covers many use cases. To give an example: now data scientists are able to list columns from several CSV files that have only one distinct value. Other supporting approaches can make use of the semantic graph, too. Search engines that index all datasets and documentation retrieve semantic entities based on a keyword-based search. In this context, thanks to the modelling, semantic search [7] also becomes possible. Besides, data visualization can make use of the additional meta data. Based on the data properties, approaches can automatically suggest more appropriate depictions using the right diagram type. In this regard, it always remains possible to generate artefacts like reports or plots from the semantic data. Especially reporting documents are well-known and understood exchange mediums in the data science community.

To sum up, our contribution is the combination of mentioned solutions to formulate a methodology that uses a semantic model which interlinks addressable data with its documentation. That way, data scientists interact with one model allowing for an easy look up of required information to understand data. Such models are commonly classified as data catalogues⁷ which store meta data about dataset contents like tables and columns. We describe our specialized model as a “semantic catalogue of data made addressable” to emphasise that our catalogue semantically enriches referred data.

¹<https://tools.ietf.org/html/rfc3986#section-3.5>

²<https://tools.ietf.org/html/rfc7111>

³<https://www.w3.org/TR/rdf11-primer>

⁴<https://www.w3.org/TR/csv2rdf>

⁵<https://www.w3.org/TR/sparql11-query>

⁶<http://dbpedia.org/resource/Car>

⁷Also known as or comparable with database catalogue, data dictionary/directory or meta data repository.

Figure 1 exemplarily illustrates our approach and the relationships of the various solutions. Data ① and documentation ② are interlinked using an RDF graph ③. With our specialized tool ④ data scientists can browse and annotate table contents as well as documents. Deep links ⑤ are used to identify and refer to specific data and documentation parts. The semantic graph interconnects and further enriches them with class types, literals and resources from external knowledge bases ⑥. SPARQL ⑦ enables querying the graph while other approaches ⑧ are now able to consume and contribute to it.

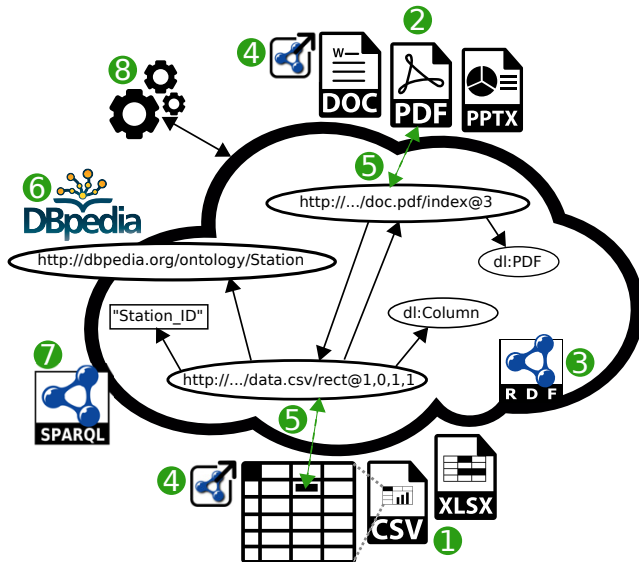


Figure 1: ① Data, ② Documentation, ③ RDF Graph, ④ Tool for ⑤ Deep Linking, ⑥ External Knowledge Base, ⑦ Querying, ⑧ Supporting Approaches.

There are also few remaining challenges. Our approach provides that a specialized tool (see next section) is used to semantically browse and annotate datasets. To aid data scientists in becoming familiar with the proposed methodology, we will make use of our previously developed approaches [11, 13].

3 IMPLEMENTATION

To showcase our approach we demonstrate an early prototype which is an extension of DeepLinker [12]. For better demonstration, an online version is available⁸.

Our locally running server application allows to browse a table-based file (CSV or Excel) in a web browser. The application responds with an HTML representation of the resource along with links pointing deeper into the table. As a result, instead of just stopping at the resource's content (surface link), users can traverse further into them. The resulting deep links allow users to point to any cell or region which enables to make statements about them in form of RDF triples.

⁸<http://www.dfki.uni-kl.de/~mschroeder/demo/deeplinker-data-understanding>

For demonstration purposes, we randomly chose a dataset about vegetation plots⁹ provided by the Environmental Conservation Online System (ECOS) via the U.S. Government open data platform¹⁰. Figure 2 shows a screen shot of our tool which browses the file CWEM_2013_Plots.csv¹¹ while cell H1 is selected. In general, the page shows the accessed link (top) and a simple form to add and list RDF triples (highlighted in gray). Below, an excerpt of the table together with the selected cell is rendered. Note that the column and row headers as well as each cell is again a deep link.

	A	B	C	D	E	F	G	H	I	J	K	L	
1	Author	Station_ID	Observat	Observat	plot	Lead	plot	Contr	Effort	Floristic	Bryophyt	Lichen	NVC
2	137-06-0RRV013I29-	29-	Forbes	BOYLE:	Very	High	Not	Not	Taxodiun	Bald-			
3	137-06-0RRV013I29-	29-	Forbes	STANTO	Very	High	Not	Not	Taxodiun	Bald-			
4	137-06-0RRV013I29-	29-	Forbes	STANTO	Very	High	Not	Not	Taxodiun	Bald-			
5	137-06-0POC016I30-	30-	Forbes	STANTO	Very	High	Not	Not	lex	Little			
6	137-06-0POC016I30-	30-	Forbes	STANTO	Very	High	Not	Not	Pinus	Pond			
7	137-06-0POC016I30-	30-	Forbes	STANTO	Very	High	Not	Not	Pinus	Pond			
8	137-06-0ALL030E01-	01-	Forbes	STANTO	Very	High	Not	Not	Pinus	Pond			
9	137-06-0ALL030C01-	01-	Forbes	STANTO	Very	High	Not	Not	Pinus	Pond			
10	137-06-0ALL030A01-	01-	Forbes	STANTO	Very	High	Not	Not	Pinus	Pond			
11	137-06-0PLD010I02-	02-	Forbes	STANTO	Very	High	Not	Not	Juncus	Black			

Figure 2: A page which shows a selected cell (H1) of a CSV resource together with a form to add and list RDF triples about this cell (highlighted in gray).

Since we would like to demonstrate how analysts would interact with documented data, we implemented some automated analyses that scan the data and generate RDF statements. Having columns and their values located in the table, it counts the total, distinct, blank, and empty values. Descriptive statistics, such as minimum, average, standard deviation and maximum, are created for the values' lengths. A histogram for distinct values is generated showing their frequency distribution.

Every insight mentioned above is formulated as one or more RDF statements. The resolvable deep links are used at the triples' *subject* and *object* positions and identify and refer to the table's contents. We plan to utilize other existing approaches that generate analogous results of analysis, like [15] or [14], by adjusting them to generate similar RDF data. Besides automatically generated statements, data scientists may add further annotations while exploring the data. For that, DeepLinker offers a form to add RDF statements about the currently browsed resource by providing a simple vocabulary (classes and properties). An even easier RDF data entry method could be provided by using the spreadsheet metaphor [13].

⁹<https://ecos.fws.gov/ServCat/Reference/Profile/44947>

¹⁰<https://www.data.gov>

¹¹<https://ecos.fws.gov/ServCat/DownloadFile/125253>


Since an RDF graph is rather difficult for data scientists to read and understand, our tool renders a familiar reporting document based on it (Figure 3). A template language¹² is used that queries the semantic graph and outputs HTML. The pages heavily uses deep link anchors in order to directly refer to the data. Describing tool tips are shown when mouse hovers.

Data Understanding Report ?

Data Collection Report ?

This report is about test.csv, a file, spreadsheet file, and csv file.

Data Description Report ?

The file is a -delimited file, spreadsheet file, and csv file. It is interpreted as spreadsheet and csv. The table has 61 rows and 9 columns, in total 549 cells. We identified the following 9 columns.

Column Observation Start Date

The column has 60 cells that contain values.

From 60 values there are ...

- 20 are distinct values.
- 0 are blank values.
- 0 are empty values.

Distinct Value Distribution

Representative Count

 21-May-2013	6
 08-May-2013	4
 07.08.13 00:00	3
 18.06.13 00:00	3
 29.04.13 00:00	3

Value Length Distribution

The length of the values have the following distribution:

11 / 12.8 ± 1.4820968814488364 / 14

Figure 3: The report’s HTML page which is generated based on the analysed data.

4 CONCLUSION AND OUTLOOK

In this paper, we identified several challenges in the process of understanding data. Our proposed solution suggests the usage of a semantic model that interlinks data with its documentation, both of which were made addressable before via our deep linking approach. We argued that this enables a simple look up of expressive information about unknown data which facilitates its comprehension. Several additional benefits of our approach were discussed. Some were also demonstrated in an online¹³ prototype.

In the future, we plan to conduct a user study to measure the impact of our methodology in real world scenarios. Looking at other data science tasks, we would like to use the acquired semantic graph to also enhance business understanding, data preparation, modelling and evaluation. We are confident that our approach enables new possibilities to assist in Semantic Data Mining, a research field that incorporates domain knowledge into usual data mining processes.

REFERENCES

- [1] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. In *The semantic web*. Springer, 722–735.
- [2] Christian Bizer and Andy Seaborne. 2004. D2RQ-treating non-RDF databases as virtual RDF graphs. In *Proceedings of the 3rd international semantic web conference (ISWC2004)*.
- [3] Pete Chapman, Julian Clinton, Thomas Khabaza, Thomas Reinartz, and Rüdiger Wirth. 1999. The CRISP-DM process model. *The CRISP-DM Consortium* 310 (1999).
- [4] Dejing Dou, Hao Wang, and Haishan Liu. 2015. Semantic data mining: A survey of ontology-based approaches. *Proceedings of the 2015 IEEE 9th International Conference on Semantic Computing, IEEE ICSC 2015* (2015), 244–251. <https://doi.org/10.1109/ICOSC.2015.7050814>
- [5] Korry Douglas and Susan Douglas. 2003. *PostgreSQL: a comprehensive guide to building, programming, and administering PostgreSQL databases*. SAMS publishing.
- [6] Thomas R Gruber. 1993. A translation approach to portable ontology specifications. *Knowledge acquisition* 5, 2 (1993), 199–220.
- [7] Ramanathan Guha, Rob McCool, and Eric Miller. 2003. Semantic search. In *Proceedings of the 12th international conference on World Wide Web*. ACM, 700–709.
- [8] Jiawei Han, Micheline Kamber, and Jian Pei. 2011. *Data Mining Concepts and Techniques*. Morgan Kaufmann.
- [9] Lukasz A Kurgan and Petr Musilek. 2006. A survey of Knowledge Discovery and Data Mining process models. *The Knowledge Engineering Review* 21, 1 (2006), 1–24.
- [10] Henrique Santos, Vasco Furtado, Paulo Pinheiro, and Deborah L McGuinness. 2017. Contextual data collection for smart cities. *Computing Research Repository eprint Journal (CoRR)* abs/1704.01802 (2017). <https://arxiv.org/abs/1704.01802>
- [11] Markus Schröder, Jörn Hees, Ansgar Bernardi, Steffen Stadtmüller, Daniel Ewert, and Peter Klotz. 2018. Simplified SPARQL REST API - CRUD on JSON Object Graphs via URI Paths. *The Semantic Web: ESWC 2018 Satellite Events, Revised Selected Papers*. (to be published).
- [12] Markus Schröder, Christian Jilek, and Andreas Dengel. 2018. Deep Linking Desktop Resources. *The Semantic Web: ESWC 2018 Satellite Events, Revised Selected Papers*. (to be published).
- [13] Markus Schröder, Christian Jilek, Jörn Hees, Sven Hertling, and Andreas Dengel. 2018. An Easy & Collaborative RDF Data Entry Method using the Spreadsheet Metaphor. *Computing Research Repository eprint Journal (CoRR)* abs/1804.04175 (2018). <https://arxiv.org/abs/1804.04175>
- [14] Dina Sukhobok, Nikolay Nikolov, and Dumitru Roman. 2017. Tabular Data Anomaly Patterns. In *2017 International Conference on Big Data Innovations and Applications (Innovate-Data)*. 25–34. <https://doi.org/10.1109/Innovate-Data.2017.10>
- [15] Juha Vesanto and Jaakko Hollmén. 2002. An Automated Report Generation Tool for the Data Understanding Phase. In *Hybrid Information Systems*. 611–626.

¹²<https://freemarker.apache.org/>

¹³<http://www.dfki.uni-kl.de/~mschroeder/demo/deeplinker-data-understanding>