



HAL
open science

An Attempt to Use Ontologies for Document Image Analysis

Bart Lamiroy, Shabai Zheng

► **To cite this version:**

Bart Lamiroy, Shabai Zheng. An Attempt to Use Ontologies for Document Image Analysis. RISE 2014 6ème Atelier Recherche d'Information SEmantique, ARIA, Mar 2014, Nancy, France. hal-00959487v1

HAL Id: hal-00959487

<https://inria.hal.science/hal-00959487v1>

Submitted on 14 Mar 2014 (v1), last revised 15 Mar 2014 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An Attempt to Use Ontologies for Document Image Analysis

Bart Lamiroy¹ and Shabai Zheng²

¹ Université de Lorraine – LORIA (UMR 7503)

Campus Scientifique – BP 239, 54506 Vandoeuvre-lès-Nancy CEDEX, France
`Bart.Lamiroy@loria.fr`

² Université de Lorraine – École Nationale Supérieure des Mines de Nancy
Campus ARTEM – CS 14 234, 92 Rue Sergent Blandan, 54042 Nancy, France

Abstract. This paper presents exploratory work on the use of semantics in Document Image Analysis. It is different than existing semantics-aware approaches in the sense that it approaches the problem from a very domain specific angle, and tries to incorporate an open model based on a reduced ontology. As presented here, it consists of enhancing an existing platform for Document Image Analysis benchmarking using off-the-shelf tools. The platform on which it is based hosts a wide variety of image interpretation algorithms as well as a wide range of benchmarking data. These data are stored in a relational database, as well as their type definition, the association between data and algorithms, *etc.* This work tries to provide an experimental indication whether ontologies and automated reasoning can provide new or alternative ways to extract relations among different stored facts, or infer dependencies between various user-defined types, based on their interactions with algorithms and other types of data.

1 Introduction

The aim of this paper is to report preliminary experimental setups related to introducing open semantics in Document Image Analysis (DIA). The topics described relate to a benchmarking platform, specifically conceived for DIA research, and stem from the need of intelligently handling a specific kind of semantics, as will be made clear later. The DAE Platform[1], on which this work is based, provides an experimental environment in which users can upload algorithms, store relevant data, conduct experiments and compare their results. This system is conceived around a flexible and open data model that links together document images, algorithms, user-defined interpretations and user information. All is stored in a relational database. Among these, the most important is a wide range of data related to document analysis algorithms that can generally be divided into user-defined data types as inputs of algorithms and automatic interpretations as results. However, the flexibility of this model in terms of interpretations and data types is one of its main weaknesses. One may define a type without considering that it may have other equivalent types in the system defined by others. Moreover, sometimes existing

types may form super types or subtypes of the one the user intended to create. The goal of this research consists in developing a semantic model of the knowledge stored in the database, which will allow automatic reasoning engines to infer dependencies between various user-defined types, based on their interactions with algorithms and other data types. One important step in this direction is to represent the underlying data model as an ontology.

The next section will explain in further details the goals of the benchmarking platform, the specific need for semantics it has, and how it differs from other existing platforms. In section 3 we explain the techniques used to integrate ontologies into the model. Section 4 provides some experimental results.

2 Context

2.1 DAE-platform

The DAE-Platform's [1] primary goal is to serve as a reference repository for reproducible experimental research in Document Image Analysis, and as such offers access to reference data, state-of-the-art algorithms, and benchmark annotations. It enables registered users to upload new reference material, to define, extend or correct the interpretation of stored data and aims to support experimental research through an open and evolving framework that can be extended by community driven contributions.




Fig. 1. Architecture of the DAE Platform

Previous accounts of the features and the potential uses of the platform [2,3] focused on its architecture, implementation choices, and its potential to impact experimental research, especially with respect to reproducibility and accountability. The platform has been implemented using a classical LAMP architecture [4] but also integrates a full WSDL interface [5], allowing it to be integrated in a fully distributed service-oriented architecture. The DIA reference data and annotations themselves are stored in an Oracle relational database back-end. An independent cluster of servers is used to execute stored algorithms. Fig 1 shows its structure.

2.2 Data model

Besides the technical architecture of the server, as depicted in Fig 1, the DAE Platform is an implementation of a data model [6]. The data model is based on the following claims: all data is typed; users can define new types;

- Data can be attached to specific parts of a document image (not obliged)
- Both data and algorithms are modelled; algorithms transform data from one type into data of another type;
- Full provenance of data history is recorded.

The result is that the platform can be used as a benchmarking repository. Since all provenance is recorded, one can leverage it to compare results from algorithms on identical data, and compare annotations by different users.

A simplified representation of the data model is represented in Fig 2. It consists of three key elements: `algorithms`, `data_items` and `algorithm_runs`.




Fig. 2. The DAE Data Model (simplified)

The underlying reasoning is that data is transformed by algorithms. `data_items` are instances of data and are related to algorithms by explicit `algorithm_runs`. New `data_items` thus produced are stored in the

database with the exact information of how they were obtained. There are also pre-defined types of `data_items`: `page_images`, `page_elements`, `page_element_properties`, `files` and `datasets`. In this document we focus on the first three.

- `page_image` is an image file representing a physical page at a given resolution and with a given image quality. It is perfectly possible to have multiple `page_images` representing the same physical page, as shown in Fig 2;
- `page_element` is an area of a `page_image`, defined in as unconstrained a way as possible, either by a bounding box or a pixel map
- `page_element_property` are any kind of user-defined property or interpretation attached to a `page_element`.

3 Semantics and Ontologies

3.1 A Specific Kind of Semantics

As already mentioned, the flexibility of DAE platform causes polysemy. Unlike other image annotation initiatives, our main goal is not to try and relate visual representations to higher level known or fixed ontologies [7]. Although we will eventually benefit from connecting our ontology with it, it is currently not within the scope of this paper.

Our main concern comes from the fact that different users very likely operate in different interpretation contexts, which inevitably may result in several types for the same visual element. We have already shown that this is a normal, and necessary "feature" in machine perception and interpretation related topics [8]. The problem is, the system cannot automatically deduce or indicate the relation among these interpretations. Our goal is to try and investigate if an appropriate ontology, capturing the semantics of our benchmarking data can help resolve this.

3.2 Example

In order to illustrate our point, let us develop the example depicted in Fig. 3. It represents the platform with various data objects and algorithms, created by different users (researchers).

We assume that one of the stored algorithms was contributed by researcher A, who defined its input as being of type X and its output of type Y. Another researcher B successfully execute this algorithm with, as input, a specific data item *i* whose type was previously defined as being Z.

It is straightforward to infer that:

$$\begin{cases} \text{Type}(Z) \cap \text{Type}(X) \neq \emptyset \\ i \in \text{Type}(Z) \cap \text{Type}(X) \end{cases}$$

The fact that types *Z* and *X* are considered being different may come from either having different users using different name labels for the same semantics (*e.g.* "text area" and "text zone"), or it may result




Fig. 3. Various Data Used by Different Users




Fig. 4. Different Data Used by the Same Algorithm

from having both users using slightly different interpretation contexts with overlapping, yet not fully identical semantics. This situation is neither far-fetched, nor unique. Since the researches in domain of document analysis can be very specific (e.g. concentrate in translating a file from a particular form to another) and the domain itself is becoming more and more complex, some researchers may work on quite specialised sub-problems, that may eventually prove to be instances of a more generic problem. In order to detect this, we need to extract the dependency between equivalent user-defined types.

3.3 Semantics of the data model

In the DAE Data Model, algorithms are declared having precise `data_types` for their input and output values (called `data`). However, an algorithm can have any number of executions (`algorithm_runs`), associating specific `data_items`. In other words, specific `data` (algorithm inputs/outputs) may possess various instances at run-time. The problem discussed above lies in the fact that, a given `data_item` can be sometimes applied in executions of different algorithms, which means it serves as `data` of potentially different types. Conversely, this means that once we find all the `data_types` of a `data_item`, we can easily trace to every algorithms where this `data_item` is involved. Our study is based on this point. Therefore, one important issue is to establish the link from specific `data` to each of its instance `data_items`.

Given the architecture of the DAE platform, we could achieve this using SQL queries. However, this solution is difficult to generalise: every time when we want to check a particular `data_item`, we need to modify the query statement (e.g. `data_item`'s id). Since the number of `data_items` is enormous, using SQL complicated, tedious and error prone. Introducing semantic querying provides a more readable presentation for complicate hierarchical relationships between information. Furthermore, DL reasoners will allow us to infer a transitive chain of a series of relations between `data_type` and `data_item`.

3.4 D2R

To achieve this, we will be using D2R [9]. D2R Server is a tool for publishing the content of relational databases on the Semantic Web. Database content is mapped to RDF by a declarative mapping which specifies how resources are identified and how property values are generated from database content. D2R Server uses a specific language³ to express mappings between application-specific database schemas and RDFS schemas or OWL. They specify how resources are identified and how property values are generated from database content. The central object in D2RQ is the *ClassMap*. A *ClassMap* represents a mapping from a set of entities described within the database, to a class or a group of similar classes of resources. Each *ClassMap* has a set of *PropertyBridges*, which specify how resource descriptions are created. Property values can be created directly from database values or by employing patterns or translation tables. D2RQ supports conditional mappings on *ClassMap* and *PropertyBridge* levels, mapping of multiple columns to the same property and the handling of highly normalised table structures where instance data is spread over multiple tables. It also includes a tool that automatically generates a D2RQ mapping from the database table structure, generating the appropriate RDF for each database, using table names as class names and column names as property names.

4 Experiments and Analysis

4.1 Matching data and data_items

As mentioned above, given a specific algorithm execution allows to relate a `data_item` to `data` properties. We know that an algorithm is defined as: $(output_1, output_2, \dots, output_n) = \text{algorithm}(input_1, input_2, \dots, input_m)$ Where $input_i$ ($i = 1 \dots m$) is defined as `datai`, who possess a type `data_typei`. In the same way, $output_j$ ($j = 1 \dots n$) is defined as `dataj` of `data_typej`. These data are associated directly with the algorithm in a strict order. This position information is available in our provenance data from the DAE platform. Fig. 5 we see the logical order of the inference. We have implemented this inference chain in Protégé⁴. Fig. 7 shows that we are capable of correctly inferring new types on data items.

³ D2RQ Mapping Language: <http://d2rq.org/d2rq-language>

⁴ <http://protege.stanford.edu/>




Fig. 5. Inference of Data Types from Algorithm Executions

Up to this point, we can conclude that if all relevant properties are properly defined(extracted), the desired deduction set by inference chain will be established automatically on the whole model after running the DL reasoning engine. However, after the implementation and tests, several issues left to be resolved:

- Complexity: a chain of 4 properties slows down the reasoning;
- Limited by the fact that the enforcement of argument positioning requires reification of the relations (our example is limited to two inputs/outputs only);
- We failed to realize the extraction by ordering: Because the D2RQ mapping language cannot yet function normally on conditional mapping.

Besides its inconveniences and technical limitations, this semantic model functions as expected, so the applied methods proved to be correct. Moreover, it proposes an important usage of the ontology modeling: verification of complex relations in relational database.

4.2 Simplification




Fig. 6. Simplified Inference of Data Types from Algorithm Executions

In order to avoid all the problem occurred in the first approach and to provide a functional model, we decided to replace the ordering information by storing a direct link between `data_item` and `data_type` in our provenance database. Every time when an algorithm runs, we record each of the `data_items` used as its inputs/outputs, along with the information

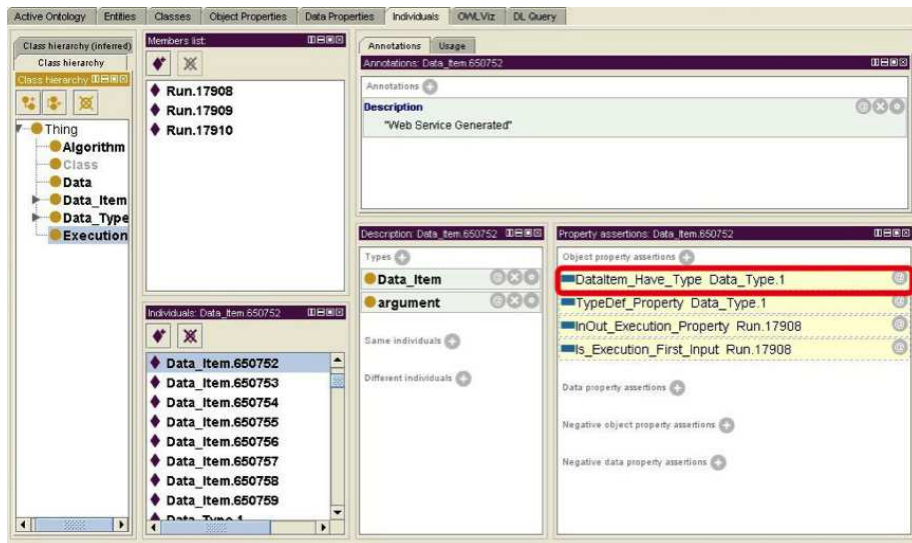



Fig. 7. Results of Type Inference from Algorithm Executions using Protégé

of this algorithm in the database. Besides, the corresponding data are also recorded. The resulting semantics is shown in Fig. 6.

This model provides the solution in a most simplified way with only three classes and is more efficient than the one in the first approach.

5 Conclusion

In this work, we have demonstrated that the use of semantics opens new perspectives of looking at document image annotation and benchmarking repositories, since it allow to infer and extract new kinds of relations between data. This can be used as a basis of trying to uncover whether different algorithms or different experimental setups share similar interpretation contexts, and thus discover new relations between various interpretations of the same data. Our approach offers advantages compared to the previously existing tool. Relations can be more easily created, modified and integrated using the appropriate an ontologies. In the domain of document analysis research, the ontology modelling is helpful in inferring the dependency between existing algorithms developed by individual researchers. Moreover, it can be considered as an efficient way to verify existing relationships between tables in the database.

Acknowledgements

Part of this work was based on previous, unpublished work on configuring a D2RQ server, by students A. Hombourger, P. Lauffenburger and J. Pruvost. Without their technical help, this work would never have been possible.

References

1. Lamiroy, B., Lopresti, D., Korth, H., Hefin, J.: How Carefully Designed Open Resource Sharing Can Help and Expand Document Analysis Research. In Gady Agam, C.V.G., ed.: Document Recognition and Retrieval XVIII - DRR 2011. Volume 7874., San Francisco, United States, SPIE, SPIE (January 2011)
2. Lamiroy, B., Lopresti, D.: An Open Architecture for End-to-End Document Analysis Benchmarking. In: 11th International Conference on Document Analysis and Recognition - ICDAR 2011, Beijing, China, International Association for Pattern Recognition, IEEE Computer Society (September 2011) 42–47
3. Lamiroy, B., Lopresti, D.: A Platform for Storing, Visualizing, and Interpreting Collections of Noisy Documents. In: Fourth Workshop on Analytics for Noisy Unstructured Text Data - AND'10. ACM International Conference Proceeding Series, Toronto, Canada, IAPR, ACM (October 2010)
4. Wikipedia: Lamp (software bundle) — Wikipedia, the free encyclopedia (2008) [Online; accessed 01-Sep-2008].

5. Chinnici, R., Moreau, J.J., Ryman, A., Weerawarana, S.: Web services description language (wsdl) version 2.0 part 1: Core language. World Wide Web Consortium, Recommendation REC-wsdl20-20070626 (June 2007)
6. Korth, H.F., Song, D., Heflin, J.: Metadata for structured document datasets. In: Proceedings of the 8th IAPR International Workshop on Document Analysis Systems. 547–550
7. Halaschek-Wiener, C., Golbeck, J., Schain, A., Grove, M., Parsia, B., Hendler, J.: Photostuff - an image annotation tool for the semantic web. In: 4th International Semantic Web Conference - Poster Paper. (2005)
8. Lamiroy, B.: Sur les limites de la perception artificielle et de l'interprétation. Hdr, Université de Lorraine (December 2013)
9. Bizer, C., Cyganiak, R.: D2r server-publishing relational databases on the semantic web. 5th international Semantic Web conference (2006)