

# A Scientific Workflow Platform for Generic and Scalable Object Recognition on Medical Images

Manuel Möller, Christopher Tuot, Michael Sintek

German Research Center for Artificial Intelligence (DFKI) GmbH  
Kaiserslautern, Germany

`manuel.moeller@dfki.de`, `christopher.tuot@dfki.de`, `michael.sintek@dfki.de`

**Abstract.** In the research project THESEUS MEDICO<sup>1</sup> we aim at a system combining medical image information with semantic background knowledge from ontologies to give clinicians fully cross-modal access to biomedical image repositories. Therefore joint efforts have to be made in more than one dimension: Object detection processes have to be specified in which an abstraction is performed starting from low-level image features across landmark detection utilizing abstract domain knowledge up to high-level object recognition. We propose a system based on a client-server extension of the scientific workflow platform Kepler that assists the collaboration of medical experts and computer scientists during development and parameter learning.

## 1 Introduction

Rapid advances in imaging technology have dramatically increased the amount of medical image data generated every day.<sup>2</sup> This data is produced in different modalities—such as CT, MRI, Ultrasound and PET/CT—which each have their own characteristics.

In recent years great improvements have been made in object recognition on medical images [1]. However, they have not yet resulted in more flexible or generic image understanding techniques. Instead, the analysis methods are still very object specific and difficult to scale across different applications (*cf.* [2]). Consequently, current image search is still dependent on the manual and subjective association of keywords to images. On the other side there are declarative representations of medical knowledge in formal ontologies like the *Foundational Model of Anatomy* (FMA) [3].

Our approach aims to integrate existing pattern recognition algorithms with formal ontologies. The benefit of this integration is twofold: Firstly, ontologies are

---

<sup>1</sup> This research has been supported in part by the THESEUS Program in the MEDICO Project, which is funded by the German Federal Ministry of Economics and Technology under the grant number 01MQ07016. The responsibility for this publication lies with the authors.

<sup>2</sup> For example, University Hospital of Erlangen, Germany has approx. 150,000 medical examinations producing 13 TB per year.

used to improve the image analysis through checking the output from the pattern recognition algorithms for their medical consistency. Secondly, the mapping of images to concepts in an ontology allows semantic and cross-modal information retrieval. We propose to use a scientific workflow engine to make the complexity of our approach tractable and aid the collaboration of medical experts and computer scientists.

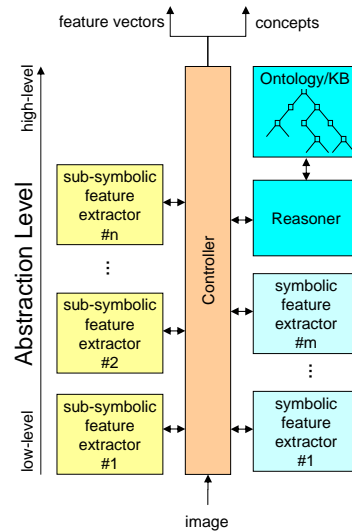
## 2 Materials and Methods

As presented in [4] we seek a deep fusion of sub-symbolic pattern recognition algorithms and domain knowledge represented in formal ontologies. We aim to eventually allow a fully automatic image segmentation and semantic annotation.

The adoption of domain knowledge from ontologies like the FMA (we rely on the OWL [5] version by [6]) has great potentials. It allows iterative processes of interaction between purely sub-symbolic pattern recognition algorithms and reasoners like Pellet [7] based on these ontologies.

To give an example: In the first step a couple of pattern matching algorithms is applied to a particular image. They output a list of detected objects. The domain knowledge can be used in conjunction with a reasoner to rule out results like a ventricle and a knee joint next to each other on the same image since this combination is anatomical nonsense.

On the other hand, the adoption of formal ontologies entails a higher complexity of the object recognition processes which makes them harder to develop, adjust and maintain. Fig. 1 shows our proposal for a system that has sub-symbolic feature extractors on the left side which operate on bit vector representations of the images. They are complemented by high-level object recognition algorithms and a reasoner interfacing medical ontologies and knowledge bases to provide the system with domain knowledge. The input in form of plain images is fed in from below. The output is twofold: low-level features are returned to be saved in a database for content based image retrieval (CBIR) based on *visual similarity* like in [8]. Concepts are returned which map image regions to corresponding concepts in the ontology to support retrieval based on *semantic similarity*. This technique allows to answer a query for the “heart” with images from various modalities which can look different from each other but all contain information about the heart.



**Fig. 1.** Integration of Feature Extraction and Background Knowledge

When a medical expert performs object recognition for an image, *e. g.*, to detect lymph nodes in a CT volume data set, he follows an implicit process consisting of several steps. Throughout this document we will use the term “process” to describe sequences of steps that have to be performed to reach a certain goal while a “workflow” is a formalization of such a process to be processed by a computer. Our architecture aims to make such workflows explicit and implement them (to a reasonable extent) in the controller component. It coordinates the feature extractors, interacts with the domain knowledge base (represented by reasoner and ontology) and directs data streams.

Therefore we base our framework on the open-source scientific workflow platform Kepler [9]. It is based on Ptolemy II, a system for heterogeneous, concurrent modeling and design developed at UC Berkeley. In Kepler, workflows can be designed following a 100% visual programming principle which allows medical experts with non or slight programming skills to design workflows and execute them remotely on a high-performance Grid-computing system.

### 3 Results

Fig. 2 shows a screenshot of our client/server extension to the Kepler system with a sample workflow for the detection of lymph nodes. Starting from the image reader the *Generic Shape Detector* performs a first preprocessing of the image. The output of this actor is passed to the *Body/Non-Body Border Detection* to separate the image into body and non-body regions. In parallel the *Anatomical Shape Detector* tries to select those shapes which are likely to belong to anatomical structures. Based on these shapes, a *Vertebrae Detection* is performed that tries to detect a vertebra as a landmark. Based on this data the reasoner searches the anatomical ontology for possible locations of lymph nodes. A resulting list of image region is subsequently processed by the *Lymph Node Detector Grid*. The results of this grid are checked for anatomical consistency. Anatomically unlikely or impossible location combinations are passed back to the *Lymph Node Detector Grid* to be checked again.

### 4 Discussion

#### Visual Representation

We believe that a formal and visual representation of the processes has benefits for all involved persons. As discussed above, the knowledge transfer from the medical experts into computers is a complicated problem. Our system addresses this fact by aiding the collaboration of medical experts and computer scientists. The clear visual representation is of substantial help when it comes to modeling complex processes. Misunderstandings between both groups can be detected easier, since the workflow is presented visually right right after it is modeled and even before it is executed for the first time.

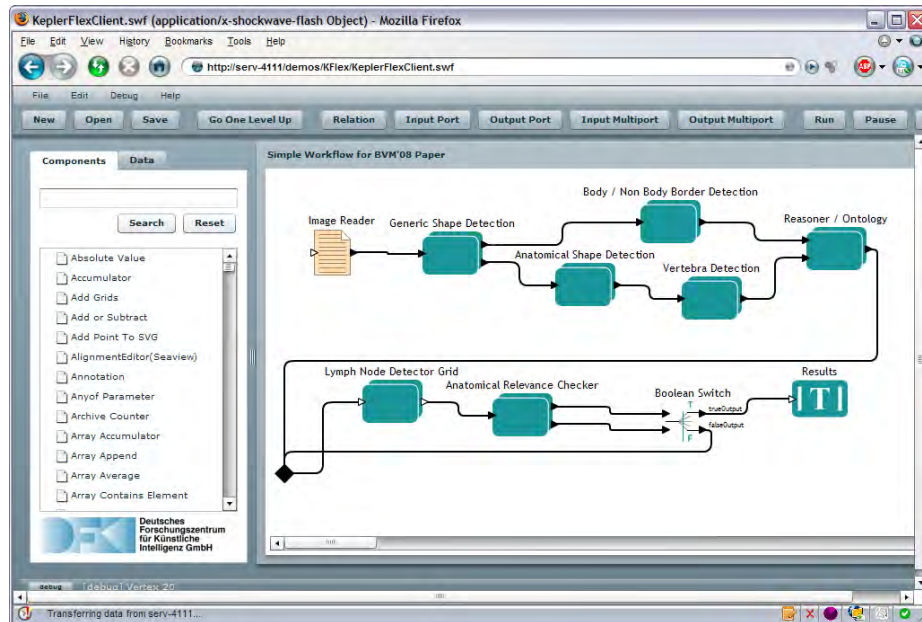


Fig. 2. Screenshot of the Interface

## Separation of Workflows from Application Logic

To our knowledge, all approaches which integrate symbolic and sub-symbolic techniques fuse them into monolithic applications. These approaches work well for particular tasks (*e. g.*, [10]) but make maintenance and reusability difficult and do not scale for other object recognition tasks. Our proposed workflows have a greater *maintainability* since they are no longer hard-coded into the application code but are modifiable remotely, via a graphical user interface. Additionally, this architecture supports the distribution and coordination of a complex architecture of remote services independent of hardware platform and physical location. This is a great benefit, since it allows to perform data intensive calculations near to where the data is stored. The Kepler platform has also built-in support for the distribution of multiple tasks over a battery of equal webservices. This feature allows to parallelize long running computing across several machines easily.

## Support for Parameter Learning

For medical image analysis often object recognition algorithms are used which incorporate inductive learning. They require a set of labeled samples compiled by medical experts who annotate image regions with particular descriptions. These descriptions can be used to train the inductive algorithm. Usually, the training phase of the classifier includes tweaking various parameters. This parameter adjustment has to be repeated several times, is time consuming and requires human interaction. Apart from the object recognition workflows our system also

supports the visual modeling of the processes which train and evaluate the classifiers. It even allows parameter tweaking for workflows containing multiple data preprocessing steps with additional parameters.

## Current State and Next Steps

Currently, our implementation of the system architecture is based on an extension of the search engine Lucene<sup>3</sup> and the Lucene Image Retrieval Image Retrieval (LIRE) library.<sup>4</sup> Among our next steps are the integration of a reasonable number of existing object recognition algorithms and the definition of workflows which connect these algorithms to proceed towards a generic image understanding.

## References

1. Tu Z, Zhou XS, Bogoni L, Barbu A, Comaniciu D. Probabilistic 3D Polyp Detection in CT Images: The Role of Sample Alignment. IEEE Computer Society Conference on Computer Vision and Pattern Recognition. 2006 May;2:1544–1551.
2. Chan AB, Moreno PJ, Vasconcelos N. Using Statistics to Search and Annotate Pictures: an Evaluation of Semantic Image Annotation and Retrieval on Large Databases. Proceedings of Joint Statistical Meetings (JSM). 2006; Available from: <http://www.svcl.ucsd.edu/publications/conference/2006/jsm06/jsm2006-imgnote.pdf>.
3. Rosse C, Mejino JL. A Reference Ontology for Bioinformatics: The Foundational Model of Anatomy. In: Journal of Biomedical Informatics. vol. 36; 2003. p. 478–500. Available from: <http://sigpubs.biostr.washington.edu/archive/00000135/>.
4. Möller M, Sintek M, Buitelaar P, Mukherjee S, Zhou XS, Freund J. Medical Image Understanding through the Integration of Cross-Modal Object Recognition with Formal Domain Knowledge. Proceedings of Healthinf 2008. 2008;.
5. McGuinness DL, van Harmelen F. OWL Web Ontology Language Overview; 2004. Available from: <http://www.w3.org/TR/owl-features/>.
6. Noy NF, Rubin DL. Translating the Foundational Model of Anatomy into OWL. In: Stanford Medical Informatics Technical Report; 2007. Available from: <http://smi.stanford.edu/smi-web/reports/SMI-2007-1234.pdf>.
7. Sirin E, Parsia B. Pellet: An OWL DL Reasoner. In: Haarslev V, Möller R, editors. Description Logics. vol. 104 of CEUR Workshop Proceedings. CEUR-WS.org; 2004. Available from: <http://dblp.uni-trier.de/db/conf/dlog/dlog2004.html#SirinP04>.
8. Deselaers T, Weyand T, Keysers D, Macherey W, Ney H. FIRE in ImageCLEF 2005: Combining Content-based Image Retrieval with Textual Information Retrieval. Working Notes of the CLEF Workshop. 2005 September; Available from: <http://www-i6.informatik.rwth-aachen.de/~deselaers/files/deselaers-fireinclef.05.pdf>.
9. Ludäscher B, Altintas I, Berkley C, Higgins D, Jaeger-Frank E, Jones M, et al. Scientific Workflow Management and the Kepler System. In: Concurrency and Computation: Practice & Experience. vol. 18; 2006. p. 1039–1065.

<sup>3</sup> <http://lucene.apache.org/>

<sup>4</sup> <http://www.semanticmetadata.net/lire/>

10. Mechouche A, Golbreich C, Gibaud B. Towards an Hybrid System Using an Ontology Enriched by Rules for the Semantic Annotation of Brain MRI Images. In: Marchiori M, Pan JZ, de Sainte Marie C, editors. Lecture Notes in Computer Science. vol. 4524; 2007. p. 219–228.