



Combining and Uniting Business Intelligence with Semantic Technologies


Acronym: CUBIST

Project No: 257403

Small or Medium-scale Focused Research Project

FP7-ICT-2009-5

Duration: 2010/10/01-2013/09/30



cubist

Your Business Intelligence

Project Fact Sheet

Abstract: n/a

Type	Other
Document ID:	CUBIST D5.1.1
Workpackage:	WP5
Leading partner:	SAP
Author(s):	Frithjof Dau (SAP)
Dissemination level:	PU
Status:	Final
Date:	21 December 2010
Version:	1.0



<Confidential>



Versioning and contribution history

Version	Description	Contributors
0.1	draft	Frithjof Dau (SAP)
1.0	Final version	Frithjof Dau (SAP)



<Confidential>



1	PROJECT FACT SHEET	4
1.1	INTRODUCTION	4
1.2	APPENDIX: FACT SHEETS	5



1 Project Fact Sheet

1.1 Introduction

In this document, the project fact sheets of CUBIST are introduced.

The CUBIST consortium decided to provide two versions of the fact sheet:

- A one-page factsheet about CUBIST.
- A two-page fact sheet about CUBIST, which extends the one-page fact sheet by including research challenges and expected impact.

Both versions can be downloaded from the CUBIST website (www.cubist-project.eu, see Deliverable D.5.1.3). The sheets are found under “Publications / General Information”.

Below in the next two figures, screenshots of the fact sheets are provided.

Combining and Uniting Business Intelligence with Semantic Technologies

At a glance

Main Objectives

The constantly growing amounts of data and an emerging trend of incorporating unstructured data into analytics is bringing new challenges to Business Intelligence (BI). Contemporary BI solutions fall short in the following aspects: Firstly, they focus only on structured data and disregard the increasing amount of information hidden in unstructured data. Next, BI users are dealing with increasingly complex analyses, but the complexity of BI tools becomes the biggest barrier for their success.

CUBIST is an EU-funded research project which envisions to leverage BI to a new level of precise and user-friendly analytics of data. CUBIST follows a best-of-breed approach that combines essential features of Semantic Technologies, Business Intelligence and Visual Analytics. It aims to

- support Federation of data from unstructured and structured sources,
- persist the data in an Information Warehouse: an approach based on a BI enabled triple store,
- provide novel kinds of Visual Analytics based on meaningful diagrammatic representations.

Use Cases

The research results will be demonstrated in three innovative use cases:

Biomedical Informatics: CUBIST will combine and explore the semantic representation of spatial temporal biomedical data from two biomedical atlases and three gene expression databases.

Control centre operations: In mission control rooms in space control centres very large volumes of data are obtained from heterogeneous sources. CUBIST targets to reveal hitherto undiscovered information and to provide support for taking better decisions.

Market and competitive intelligence: This job market use case combines information from crawled job advertisements and an existing firmographic database which enables insights about recruitment activities.

Expected Impact

CUBIST is based on Semantic Technologies, particularly on RDF for data representation and triple stores for federating and persisting the data, and Formal Concept Analysis (FCA) for conceptually clustering the data. These clusters will be used for the Visual Analytics. We expect the following impact:

Semantic Technologies: CUBIST aims to bring semantic technologies to a level where they can be successfully applied in industrial settings using huge data sets, comparable to established technologies such as relational databases and BI.

Business Intelligence: Incorporating unstructured data will be very important for future BI systems. CUBIST takes an essential step in this direction. From a technological perspective, CUBIST will propose new architectures which can be used for future BI systems. From a business perspective, CUBIST will help overcome the barrier of complexity of current BI tools and apply BI functionalities to new scenarios.

Visual Analytics: From a technological perspective, CUBIST will help using FCA in real business settings. From a user perspective, using FCA has a solid mathematical foundation and a close link to the human perception of concepts, thus FCA will drive visual analytics to a new level of theoretically precise and humanly comprehensible visualizations.

www.cubist-project.eu

Partners: SAP AG, Oxford, Sheffield Hallam University, Centre Recherche S.A., Heriot-Watt University, Semantic Solutions Services, Innovantage

Core Information: Call: FP7-ICT-2009-5, Objective: 6.3, Intelligent Information Management, Project ID: CUBIST-252483, Funding Scheme: Smp, Duration: 01/01/2010-01/2013, Total Cost: €4,357,834, EC Contribution: €1,029,139

Project Coordinator: Dr. Ingrid Isik, SAP AG, SAP Research Center Dresden, Tel: +49 (0)911 4081 6162, Ingrid.Isik@sap.com

Logos: SAP, ontotext, Sheffield Hallam University, spaceapplications, innovantage, HERIOT WATT UNIVERSITY

Fig 1: CUBIST fact sheet, short version



<Confidential>





Combining and Uniting Business Intelligence with Semantic Technologies

Main Objectives

The constantly growing amounts of data and an emerging trend of incorporating unstructured data into analytics is bringing new challenges to Business Intelligence (BI). Contemporary BI solutions fall short in the following aspects: Firstly, they focus only on structured data and disregard the increasing amount of information hidden in unstructured data. Secondly, BI users are dealing with increasingly complex analyses, but the complexity of BI tools becomes the biggest barrier for their success.

CUBIST is an EU funded research project with a visionary approach that leverages BI to a new level of precise, meaningful and user-friendly analytics of data by following a best-of-breed approach that combines essential features of Semantic Technologies, Business Intelligence and Visual Analytics. CUBIST aims to

- support federation of data from unstructured and structured sources
- persist the federated data in an Information Warehouse; an approach based on a BI enabled triple store.
- provide novel ways of applying Visual Analytics based on meaningful diagrammatic representations

	Classical Business Intelligence	Semantic Business Intelligence
Output		
User Interaction	Restricted queries / analytics	Free-libs and visual queries / analytics
Store	Data Warehouse	Triple Store
Gathering Information	ETL	Semantic ETL
Data sources	Databases	Databases, Forums, blogs, Office-docs

Use Cases

The research results will be demonstrated in three innovative use cases:

Biomedical Informatics: CUBIST will combine and explore the semantic representation of spatial temporal biomedical data from two biomedical atlases and three gene expression databases.

Control centre operations: In mission control rooms in space control centres very large volumes of data are obtained from heterogeneous sources, including structured and unstructured data. CUBIST will provide online support for taking better decisions, reveal hitherto undiscovered information and provide supportive evidence in debriefing and decision making processes related to the organisation of space control centre operations.

www.cubist-project.eu

At a glance

Partners

SAP AG
Olivier
Sheffield Hallam University
Centric Technology SA
Heriot-Watt University
Open Applications Services
Innovantage

Core Information

Call: FP7-ICT-2009-5
Objective: 4.3. Intelligent Information Management
Project ID: CUBIST-237403
Funding Scheme: Stsp
Duration: 30/2010-09/2013
Total Cost: € 4.897.004
EC Contribution: € 3.029.339

Project Coordinator

Dr. Patrick Doherty
SAP AG
SAP Research Center Dresden
Tel: +49 (0)513 6111 6133
patrick.doherty@sap.com








Market and competitive intelligence: This is a job market use case which combines information from job advertisements crawled by CUBIST and an existing firmographic database. CUBIST will enable market intelligence (insights about who is recruiting, and where, and when and how they recruit) and competitive intelligence to help employers track and better understand the recruitment activity of their competitors.

Research Challenges and Expected Impact

CUBIST is based on Semantic Technologies, particularly on RDF for data representation and triple stores for federating and persisting the data, and Formal Concept Analysis (FCA) for conceptually clustering the data and organizing the clusters into hierarchical relationships. These clusters will be used for the Visual Analytics. This leads to the following challenges and expected outcomes:

Semantic Extract, Transform and Load (ETL): Source data in unstructured sources might be noisy, inconsistent and incomplete. Extracted data from unstructured sources has to be brought into relationship with extracted data from structured sources. CUBIST targets semantically enriched lineage information, error detection and identity resolution within extracted data, and a semantic ETL component that provides SPARQL (the query language for triple stores) endpoints for various data sources.

Query Language: SPARQL lacks complex aggregate functionality, reporting functions and rollup/cube expressivity. In alignment with the efforts of W3C, CUBIST will extend SPARQL by needed OLAP functionalities.

Performance and scalability of the Triple Store: Current state-of-the-art implementations of triple stores are for tens of billions of triples, a magnitude lower than the data volumes for a state-of-the-art data warehouse. CUBIST will significantly multiply the number of triples the triple store can deal with.

FCA and Triple Stores: Most FCA applications are stand-alone solutions. In CUBIST, a layer within the warehouse will integrate the triple store with the FCA-based visual analytics.

Scalability of FCA: Existing FCA solutions do not scale to large amounts of data. CUBIST will investigate high-performance FCA algorithms and tools, including parallel processing algorithms for multi-core architectures.

Visual Analytics: Current FCA visualization tools have been designed without very large data sets in mind, interlinked with best practices from known BI visualizations. CUBIST will scrutinize novel approaches for FCA-based visualisations which allow for in depicting, navigating through and visually querying the data.

To summarize, for the core fields Semantic Technologies, Business Intelligence and Visual Analytics, we expect the following impact:

Semantic Technologies: CUBIST aims to bring Semantic Technologies to a level where they can be successfully applied in industrial settings using huge data sets, comparable to established technologies such as relational databases and BI.

Business Intelligence: It is expected that incorporating unstructured data will be very important for future BI systems. CUBIST takes an essential step in this direction. From a technological perspective, CUBIST will have an impact on the architecture of future BI systems. From a business perspective, CUBIST will help overcome the barrier of complexity of BI tools and interfaces and apply BI functionalities to new business scenarios and new user groups.

Visual Analytics: From a technological perspective, using FCA for representing and navigating the large amounts of data will open up FCA to new horizons in terms of applicability in real business settings. From a user perspective, using FCA has a solid mathematical foundation and a close link to the human perception of concepts, thus FCA will drive Visual Analytics to a new level of theoretically precise and humanly comprehensible visualizations.

www.cubist-project.eu

Fig 2: CUBIST fact sheet, extended version

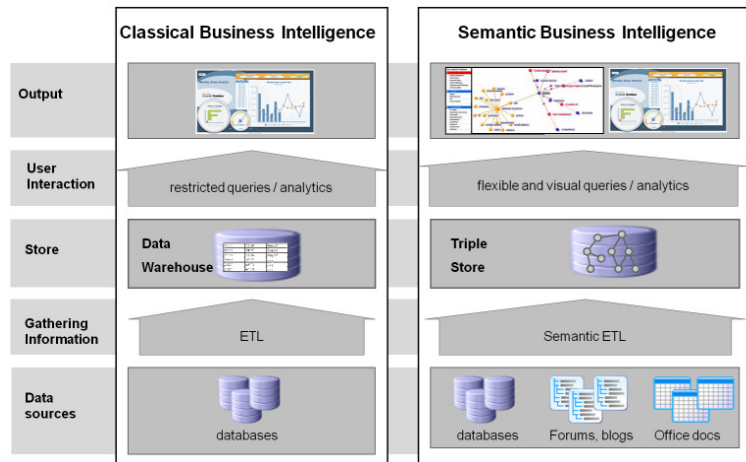
1.2 Appendix: Fact Sheets

On the next three pages, the fact sheets are attached in its original layout. The word-documents of the sheets can be provided on request.

Combining and Uniting Business Intelligence with Semantic Technologies

Main Objectives

The constantly growing amounts of data and an emerging trend of incorporating unstructured data into analytics is bringing new challenges to Business Intelligence (BI). Contemporary BI solutions fall short in the following aspects: Firstly, they focus only on structured data and disregard the increasing amount of information hidden in unstructured data. Next, BI users are dealing with increasingly complex analyses, but the complexity of BI tools becomes the biggest barrier for their success.



CUBIST is an EU-funded research project which envisions to leverage BI to a new level of precise and user-friendly analytics of data. CUBIST follows a best-of-breed approach that combines essential features of Semantic Technologies, Business Intelligence and Visual Analytics. It aims to

- support federation of data from unstructured and structured sources,
- persist the data in an Information Warehouse; an approach based on a BI enabled triple store,
- provide novel kinds of Visual Analytics based on meaningful diagrammatic representations.

Use Cases

The research results will be demonstrated in three innovative use cases:

Biomedical informatics: CUBIST will combine and explore the semantic representation of spatial temporal biomedical data from two biomedical atlases and three gene expression databases.

Control centre operations: In mission control rooms in space control centres very large volumes of data are obtained from heterogeneous sources. CUBIST targets to reveal hitherto undiscovered information and to provide support for taking better decisions.

Market and competitive intelligence: This job market use case combines information from crawled job advertisements and an existing firmographic database which enables insights about recruitment activities.

Expected Impact

CUBIST is based on Semantic Technologies, particularly on RDF for data representation and triple stores for federating and persisting the data, and Formal Concept Analysis (FCA) for conceptually clustering the data. These clusters will be used for the Visual Analytics. We expect the following impact:

Semantic Technologies: CUBIST aims to bring semantic technologies to a level where they can be successfully applied in industrial settings using huge data sets, comparable to established technologies such as relational databases and BI.

Business Intelligence: Incorporating unstructured data will be very important for future BI systems. CUBIST takes an essential step in this direction. From a technological perspective, CUBIST will propose new architectures which can be used for future BI systems. From a business perspective, CUBIST will help overcome the barrier of complexity of current BI tools and apply BI functionalities to new scenarios.

Visual Analytics: From a technological perspective, CUBIST will help using FCA in real business settings. From a user perspective, using FCA has a solid mathematical foundation and a close link to the human perception of concepts, thus FCA will drive visual analytics to a new level of theoretically precise and humanly comprehensible visualizations.

At a glance

Partners

SAP AG
 Ontotext
 Sheffield Hallam University
 Centrale Recherche S.A.
 Heriot-Watt University
 Space Applications Services
 Innovantage

Core Information

Call: FP7-ICT-2009-5
 Objective: 4.3: Intelligent Information Management
 Project ID: CUBIST-257403
 Funding Scheme: Strep
 Duration: 10/2010-09/2013
 Total Cost: € 4.357.834
 EC Contribution: € 3.029.193

Project Coordinator

Dr. Frithjof Dau
 SAP AG
 SAP Research Center Dresden
 Tel: +49 (0)351 4811 6152
 frithjof.dau[at]sap.com



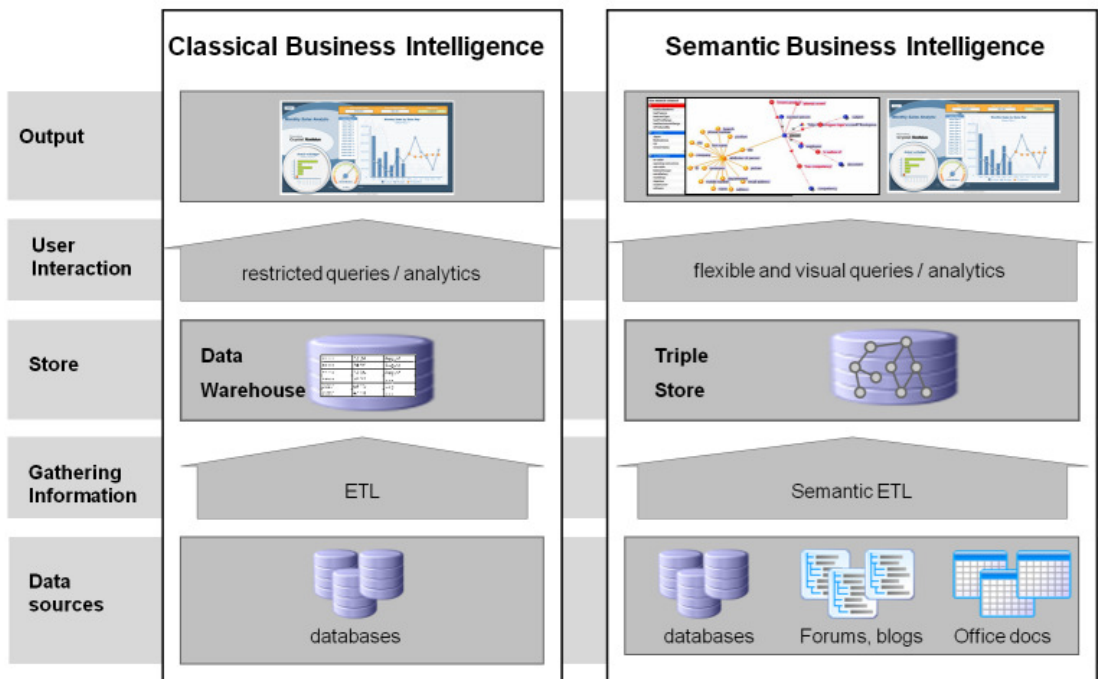
Combining and Uniting Business Intelligence with Semantic Technologies

Main Objectives

The constantly growing amounts of data and an emerging trend of incorporating unstructured data into analytics is bringing new challenges to Business Intelligence (BI). Contemporary BI solutions fall short in the following aspects: Firstly, they focus only on structured data and disregard the increasing amount of information hidden in unstructured data. Secondly, BI users are dealing with increasingly complex analyses, but the complexity of BI tools becomes the biggest barrier for their success.

CUBIST is an EU funded research project with a visionary approach that leverages BI to a new level of precise, meaningful and user-friendly analytics of data by following a best-of-breed approach that combines essential features of Semantic Technologies, Business Intelligence and Visual Analytics. CUBIST aims to

- support federation of data from unstructured and structured sources
- persist the federated data in an Information Warehouse; an approach based on a BI enabled triple store
- provide novel ways of applying Visual Analytics based on meaningful diagrammatic representations



Use Cases

The research results will be demonstrated in three innovative use cases:

Biomedical informatics: CUBIST will combine and explore the semantic representation of spatial temporal biomedical data from two biomedical atlases and three gene expression databases.

Control centre operations: In mission control rooms in space control centres very large volumes of data are obtained from heterogeneous sources, including structured and unstructured data. CUBIST will provide online support for taking better decisions, reveal hitherto undiscovered information and provide supportive evidence in debriefing and decision making processes related to the organisation of space control centre operations.

At a glance

Partners

SAP AG
 Ontotext
 Sheffield Hallam University
 Centrale Recherche S.A.
 Heriot-Watt University
 Space Applications Services
 Innovantage

Core Information

Call: FP7-ICT-2009-5
 Objective: 4.3: Intelligent Information Management
 Project ID: CUBIST-257403
 Funding Scheme: Strep
 Duration: 10/2010-09/2013
 Total Cost: € 4.357.834
 EC Contribution: € 3.029.193

Project Coordinator

Dr. Frithjof Dau
 SAP AG
 SAP Research Center Dresden
 Tel: +49 (0)351 4811 6152
 frithjof.dau[at]sap.com



Market and competitive intelligence: This is a job market use case which combines information from job advertisements crawled by CUBIST and an existing firmographic database. CUBIST will enable market intelligence (insights about who is recruiting, and where and when and how they recruit) and competitive intelligence to help employers track and better understand the recruitment activity of their competitors.

Research Challenges and Expected Impact

CUBIST is based on Semantic Technologies, particularly on RDF for data representation and triple stores for federating and persisting the data, and Formal Concept Analysis (FCA) for conceptually clustering the data and organizing the clusters into hierarchical relationships. These clusters will be used for the Visual Analytics. This leads to the following challenges and expected outcomes:

Semantic Extract, Transform and Load (ETL): Source data in unstructured sources might be noisy, inconsistent and incomplete. Extracted data from unstructured sources has to be brought into relationship with extracted data from structured sources. CUBIST targets semantically enriched lineage information, error detection and identity resolution within extracted data, and a semantic ETL component that provides SPARQL (the query language for triple stores) endpoints for various data sources.

Query Language: SPARQL lacks complex aggregate functionality, reporting functions and rollup/cube expressivity. In alignment with the efforts of W3C, CUBIST will extend SPARQL by needed OLAP functionalities.

Performance and scalability of the Triple Store: Current state-of-the-art implementations of triple stores are for tens of billions of triples, a magnitude lower than the data volumes for a state-of-the-art data warehouse. CUBIST will significantly multiply the number of triples the triple store can deal with.

FCA and Triple Stores: Most FCA applications are stand-alone solutions. In CUBIST, a layer within the warehouse will integrate the triple store with the FCA-based visual analytics.

Scalability of FCA: Existing FCA solutions do not scale to large amounts of data. CUBIST will investigate high-performance FCA algorithms and tools, including parallel processing algorithms for multi-core architectures.

Visual Analytics: Current FCA visualization tools have been designed without very large data sets in mind. Interlinked with best practices from known BI visualizations, CUBIST will scrutinize novel approaches for FCA-based visualisations which allow for in depicting, navigating through and visually querying the data.

To summarize, for the core fields Semantic Technologies, Business Intelligence and Visual Analytics, we expect the following impact:

Semantic Technologies: CUBIST aims to bring Semantic Technologies to a level where they can be successfully applied in industrial settings using huge data sets, comparable to established technologies such as relational databases and BI.

Business Intelligence: It is expected that incorporating unstructured data will be very important for future BI systems. CUBIST takes an essential step in this direction. From a technological perspective, CUBIST will have an impact on the architecture of future BI systems. From a business perspective, CUBIST will help overcome the barrier of complexity of BI tools and interfaces and apply BI functionalities to new business scenarios and new user groups.

Visual Analytics: From a technological perspective, using FCA for representing and navigating the large amounts of data will open up FCA to new horizons in terms of applicability in real business settings. From a user perspective, using FCA has a solid mathematical foundation and a close link to the human perception of concepts, thus FCA will drive Visual Analytics to a new level of theoretically precise and humanly comprehensible visualizations.
