

DAVIDE SOTTARA

PRESENT ADDRESS

2700 56th St NW
Unit 1 Apt 3
Rochester, MN 55901
(480) 628-7805
davide.sottara@asu.edu
skype: davide.sottara

PERMANENT ADDRESS

P.zza VII Novembre 1944
Bologna, Italy 40131
+39 (051) 551079
dsotty@gmail.com

SUMMARY

Enterprise Knowledge Architect with experience in Applied Clinical Informatics, Research and Development. Skills and Background include Artificial Intelligence, Software Engineering, Knowledge Representation, Standards, Interoperability and Clinical Application Development.

EDUCATION

University di Bologna (Italy), School of Engineering,
Department of Computer Science, Electronics and Systems:

- PhD in Computer Science, Electronics and Telecommunications Engineering, 2010
- Laurea Magistrale (Masters Degree) in Computer Science Engineering, 2006, with honors
- Laurea (Bachelor Degree) in Computer Science Engineering, 2003, with honors

LANGUAGES

- Italian (native)
- English
- German and Spanish (basic understanding)

EXPERIENCE

Knowledge Architect Mayo Clinic Fall 2012-Present
Rochester, MN

Enterprise Knowledge Architecture.
Knowledge Elicitation, Management, Representation, Delivery, Execution. Cognitive Support Applications.
Data Quality, Interoperability, Semantic Data Enrichment, Standards Compliance.

CTO Healthcare Service Platform Consortium July 2015-August 2016
US

Service Oriented Architectures.
Knowledge Architectures, Interoperability, Cognitive Support Application Development.

Clinical Assistant Professor Biomedical Informatics Department, ASU Fall 2014-Present
Scottsdale, AZ

Clinical Decision Support Systems, Clinical Applications, Automated reasoning systems.

Assistant Professor Biomedical Informatics Department, ASU Fall 2012-Fall 2014
Scottsdale, AZ

Clinical Decision Support Systems.
Standards for Clinical Knowledge Elicitation, Representation and Interchange.

Contractor	Soadex, Inc. Boston, MA	Jan 2011 - Jan 2012
	DOD-funded KMR-2 project. Personalized, agent-based clinical decision support systems. Hybrid systems combining rule-based decision support and predictive modeling techniques	
Post Doc Grant	Dept. of Computer Science (DEIS) University of Bologna, Italy	Jun 2011 - May 2012
	Hybrid decision support systems based on signal processing and interpretation techniques. Risk estimation models for geriatric patients.	
Consultant	SPES s.c.p.a. / ENEA Fabriano (AN) / Bologna, Italy	Jul 2009 - Sep 2010
	Development, Integration and Testing of a remote data management system (“PLT Line”) with an environmental decision support system.	
Internship	Univ. Newcastle / JBoss, Newcastle upon Tyne (UK)	Feb 2009 - Jun 2009
	Integration of “imperfect” reasoning techniques in the open source rule engine JBoss Drools	
Internship	ENEA Bologna, Italy	Feb 2009 - Jun 2009
	Integration of “imperfect” reasoning techniques in the open source rule engine JBoss Drools	

PROJECTS AND ACTIVITIES

Enterprise Knowledge Management and Delivery Platform Architecture (KMDP) (Mayo Clinic)

- **Knowledge Management Metadata Model / Knowledge Asset Catalog. (Mayo Clinic, with J.Shellum, K.Pavek et al.)**

I have optimized the information model used to manage and catalog knowledge assets (valuable pieces of clinical knowledge curated for delivery and consumption by specific audiences), aligning it with HL7 standards and grounding it into OWL ontologies used in library sciences. I have helped the Clinic’s Knowledge Management team define a roadmap for the implementation of the model at scale. The model is also informing the design of the Mayo Clinic Knowledge Asset Catalog, which indexes and categorizes the assets so that they can be searched and browsed by subject matter experts and other interested parties.

- **KMDP APIs / API4KP (Mayo Clinic, OMG).**

I am designing a portfolio of KMDP APIs and their underlying service model. The APIs serve some of the KMDP capabilities for consumption by clinical applications. The APIs cover the storage and retrieval of knowledge content from dedicated repositories, the syntactic and semantic manipulation of knowledge assets, the construction of knowledge bases and the integration of automated reasoning components. A first version of the API specification has been submitted to OMG, in response to their API4KP Request for Proposals, and accepted in late 2017. Currently, I am refining the specification and working on a reference implementation.

- **Clinical Concept Glossary (2017, Mayo Clinic, with S.Demuth, J.Shellum, D.Kor, R.Chaudry et al.)**

To organize the Clinic’s Knowledge-based data enrichment framework, the notion of a ‘glossary’ has been introduced, to assert and catalog the clinical semantics of each data transformation module. I have helped review a number of preexisting related tools and technologies, and designed the initial architecture of the Glossary. This preliminary work resulted in a technical report, and the idea is being considered for a joint Mayo/Epic co-development partnership. Currently, I am also developing a proof of concept implementation to better inform the technical requirements.

Standards Based Representation of Clinical Knowledge (Mayo Clinic)

- **Sharable CDS (Mayo Clinic, with J.Shellum et al.)**

I am investigating the use of standards for the expression of platform-independent representation of CDS logic, with focus on delivery to different audiences. The work focuses on the relationship between CDS endorsed by the enterprise, its implementation in the Epic EHR, and its diffusion among partners in the Mayo Clinic Health System, and its potential commercialization. I have produced a preliminary report that compares the use of current standards (including but not limited to RuleML, SWRL, SBVR and HeD/KNART), and I am leading a pilot project that uses a combination of KNART, DMN and CQL for the reverse engineering of CDS implemented in Epic.

- **Care Process Models. (Mayo Clinic, with R.Nishimura, J.Shellum et al.)**

Mayo Clinic formalizes and delivers internal guidelines for diagnosis, prevention and treatment of a variety of conditions in the form of Care Process Models (CPM). My work consists in the elicitation of clinical knowledge from a subject matter expert, its formalization into CPMs by means of standards-compliant languages, and the actual modelling of a few strategically valuable CPMs (Atrial Fibrillation, Pulmonary Embolism). The CPMs are considered valuable knowledge assets, and are on a roadmap to be consumed by the Mayo Expert Advisory applications to provide cognitive support to clinicians at the point of care .

- **Business Process Management for Healthcare. (OMG)**

I have been an active participant in a multidisciplinary, multi-actor initiative, under the umbrella of OMG, whose goal is the evaluation of the gaps and opportunities for the adoption of cross-industry knowledge representation standards (BPMN, CMMN, DMN) to healthcare processes. The workgroup is preparing a white paper that describes methods and best practices for the modelling of clinical processes.

Clinical Applications Development

- **Mayo Expert Advisory (MEA) v3.0 Discovery. (Mayo Clinic)**

This discovery project is developing a standards compliant, knowledge-driven cognitive support application for the delivery of Mayo Clinic Care Process Models at the point of care. MEA3D aims to (i) provide an innovative user experience for clinicians, (ii) demonstrate the benefit of knowledge-based systems and their underlying delivery platforms. I have led the initial architectural design, and I am specifically in charge of the integration patterns and protocols with the rest of the Mayo Clinic platforms. The initial prototype received very positive feedback, and the project has been encouraged to grow by means of both internal Mayo Clinic funding and external strategic partnerships.

- **Consent 2 Share. (ASU, with A.Grando)** I have advised Dr.Grando and mentored two BMI student on the technical aspects of the Consent2Share open source platform, used in the MyDataChoice project. The activity allowed to deploy an enhanced instance of Consent2Share at ASU, and is leading towards the improvement of the precision of the system in the identification of sensitive patient data that may be subject to consent policies.

PAST PROJECTS and ACTIVITIES

- **EASE Project (2014-15, Mayo Clinic, with D. Larson, J. Lovely, M. Burton et al.)**

A large scale (ca. 80 people) project aimed at providing clinicians with cognitive support, monitoring, decision support, and population based analytics for colorectal surgery patients, using knowledge-based architectures. I have contributed to the design and implementation of the data enrichment pipelines, the domain specific clinical terminologies and ontologies, semantic data models, clinical decision support engines, provenance models, and integration within the Mayo Clinic IT infrastructure. While generally not considered successful, the project resulted in the release of open source software, a number of technical reports, and the initiation of the KMDP and MEA3D discovery projects

- **PGx Guideline Modelling (2015, Mayo Clinic, with R.Freimuth)**
Exploratory formalization of pharmacogenomics guidelines with rules and process models. The projects resulted in a conference paper.
- **Rule/Data Model/Terminology integration (2015, Mayo Clinic, with H.Solbrig and K.Peterson)**
Integration of conceptual data models, information models and terminology systems with CDS rules and other clinical logic, both at authoring time and runtime. The project resulted in the release of open source software.
- **Comparison of Decision Modelling Authoring Tools. (2016, Mayo Clinic)**
I have conducted a specific survey and comparison of knowledge authoring tools specifically designed for the OMG Decision Modelling Notation (DMN) standard, with a secondary investigation on process standards such as BPMN and CMNN. The project resulted in a technical report, and influenced the acquisition of one of the tools.
- **Knowledge Management and Delivery Platform Enterprise Architecture. (2016, Mayo Clinic, with K.Toussaint)**
The Mayo Clinic IT infrastructure is based on four enterprise-wide platforms - Identity Management, API Management, Unified Data and the nascent Knowledge Management and Delivery (KMDP). My work consisted in the use of a TOGAF-compliant framework for the definition of core KMDP business and technical capabilities, and their mapping to ongoing strategic projects at Mayo. I have used the capability framework to evaluate existing IT systems for knowledge, content and terminology management (KCMS, TERMS), as well as several new candidates for acquisition. I have also consulted on the interfacing of the KMDP with other platforms, with particular focus on the UDP. The activity resulted in a report, an architectural specification and a business capability model.
- **“Piper” project. (2016, ASU, with Drs. Greenes and Burton)**
I have overseen the design and architecture of a context-driven clinical application framework, funded by the Piper grant money awarded for 2016. This framework allows to develop applications that maintain their internal state aligned with a clinical workflow, trying to provide better cognitive support (and thus better experience) to their users. The framework uses concepts from the Work Domain Ontology, and is designed to evolve into an HSPC compliant architecture. The outcome of the project has seeded, among others, Dr.Greenes’ HII-C and VIP program, and influenced the design of the Mayo Clinic MEA3D application.
- **Knowledge Enrichment Framework (2016, Mayo Clinic, with M.Burton and A.Knopp)**
I have contributed the design a DMN-oriented framework for the formalization, orchestration and application of (clinical) knowledge to the validation, normalization, classification and augmentation of clinical data. Specifically, I have recommended the use of standards for knowledge representation, and provided formal definitions the various kinds of enrichment. This framework, described in a number of technical reports and presented at the AMIA conference, is going to be implemented in the Mayo Clinic Unified Data Platform.
- **HSPC Service Oriented Architecture (2016, with HSPC).**
The Healthcare Service Platform Consortium efforts have led to the identification of a number of candidate standards and reference implementations for clinical application Service Oriented Architectures. I contributed to the early stages of investigation.
- **(Clinical) Work Domain Ontology (2016/17 joint Mayo Clinic / ASU, with M.Burton and S.Arabandi).**
I have worked on a modular, large scale OWL ontology of clinical activities, environments, roles and information. I have helped identify suitable standard components, created bridge ontologies to mediate them, designed an architecture for the organization of the knowledge artifacts and provided a preliminary recommendation for the management of the ontology using the Mayo Clinic ontology and terminology environments. This ontology has been a foundational asset in several Mayo Clinic and ASU projects, including the SSS and Piper projects mentioned in this list.
- **Enterprise URI strategy. (2017, Mayo Clinic)**
I have provided a recommendation for the adoption of semantic web compatible, standard

identifiers across the enterprise. This strategy, and a (web) service model built around it, is being piloted in the MEA3D project.

- **Source System Semantics (SSS) (2017, Mayo Clinic, with M.Burton, A.Knopp et al.)**

As the Mayo Clinic is implementing Epic, a large number of other commercial and homegrown systems are being decommissioned. The Source System Semantics (SSS) project aims to categorize and document the meaning of the data that is exported out of those applications. The semantics to be preserved comes from the data itself, the business logic of the applications and the inferences made by clinicians interacting with the applications. I have focused on the processes and the tools that use concepts from the Work Domain Ontology to describe the semantics of the data in a way that is both consumable by subject matter experts, but also computable to search and query for specific information. The project resulted in several technical reports, a methodology adopted by the Mayo Clinic staff and a growing knowledge base about applications and data semantics.

OPEN SOURCE PROJECTS CONTRIBUTIONS

JBoss Drools, 5.x and 6.x series

I have been involved with the Drools community since 2008. Initially a user and community member, since 2009 I have become a committer, contributing code to the project. I have been working specifically on dynamic data models for rule execution, but I have also contributed a variety of fixes and improvements. The contributions, necessary for most of my other project, became the core of a minor release (version 5.6). I am currently considered a honorary core team member and work closely with Mark Proctor as a scientific advisor to the team. I am continuing to contribute to the 6.x series and participate in the design and implementation of the new features.

Drools-PMML

Started in 2011, this project is aimed at the integration of predictive models and rule-based reasoning. It leverages the PMML standard and allows to deploy models such as neural networks or decision trees in a rule-based system. This allows to combine the native, qualitative nature of the business rules with the quantitative approach of the predictive models. The project has become a sub-module of the official release since Drools 6.1.

Drools-Chance

This module extends the rule language and the business rule engine, adding a generic and customizable framework for non-boolean reasoning. The initial implementation is focused on fuzzy logic, but the proposed extension can also accommodate some forms of probabilistic reasoning.

Drools-Shapes

This module is focused on the integration of semantic web technologies with the production rule engine. It consumes OWL ontologies and generates an enhanced set of Java classes and interfaces based on the concepts and properties defined in the original ontology. The classes are compatible with major persistence and serialization frameworks (JPA, JAXB, Jena). The module also generates classification rules and provides helpers to convert XSD schemas into equivalent OWL representations.

Drools-MAS

Drools MAS is a Drools-based experimental framework for the creation of intelligent agents compatible with the FIPA standard. The behavior of the agents is completely specified using a combination of built-in and user-defined rules. It supports dynamic knowledge bases, where the rules can be replaced at runtime, allowing for adaptive behaviors.

Socratic Grid

Socratic Grid, developed during the KMR2 project, is a collection of open source services and utilities that aim at the creation of a platform for Clinical Decision Support. I have contributed the CDS components, in the form of three instances of Drools-MAS agents: a general purpose CDS agent, an agent managing the communications between patients and providers and a resource optimization

agent that plans and schedules appointments and visits.

Health eDecisions editor

The editor, developed for the SHARPC2B project, is an ontology-driven, visual editor compatible with the HeD standard. It is designed as a web application and can be deployed in a cloud-based environment. It offers a technical view, where expressions can be composed using a visual approach, and a subject-matter expert mode, which assists the user providing preconfigured templates for most common use cases. Both the domain model, the expression language and the assisted mode templates are pluggable, making the editor potentially reusable in a variety of contexts.

Other contributions

I occasionally participate in, but do not lead, the following other open source projects:

- **Empire** : a JPA implementation that relies on triple stores rather than relational databases for persistence
- **MVEL** : a lightweight expression language and interpreter
- **JBoss Roaster** : an API for code generation
- **JBoss Tohu** : a Drools-based tool for dynamic, adaptive questionnaires
- **OntoMaven** : Tools for the management of OWL ontologies

Code available from <https://github.com/sotty>

Computer Skills

As a software engineer, computer scientist and software developer, I have expertise across the entire life-cycle of an application, from requirements elicitation to design, development, testing and deployment, under different architectural styles and principles.

- **Service Oriented Architectures:** REST, OpenAPI/Swagger, SOAP
- **Knowledge Based Systems:** Rule Engines, Process Engines, Description Logic (OWL) Reasoners, Predictive Models and Data Mining, Terminology Systems; including respective APIs.
- **Model Driven Architectures:** UML, Code Generation frameworks.
- **Databases:** SQL, Triple Stores.
- **Programming Languages:** Java, Matlab, C/C++/C#; XML, RDF, JSON and YAML-related technologies.
- **Software Engineering Tools:** Maven, Git/GitHub, TFS
- **HW/SW Administration:** Assembly and Configuration of Personal Computers; Microsoft Windows and Linux/Ubuntu Management, Administration and Installation.

Knowledge Representation

I am familiar with the following standards for knowledge interoperability.

- **Business Modeling Languages:** BPMN, CMMN, DMN, UML
- **Predictive Models:** PMML
- **Expression Languages:** CQL/ELM, FHIRPath, FEEL
- **Semantic Web:** OWL, RDF, SKOS, SPARQL, SWRL
- **Rule Languages:** RuleML, RIF, SBVR, HeD/KNART, Arden Syntax
- **Healthcare Data Models:** HL7 RIM, FHIR
- **Healthcare Terminology Systems:** SNOMED, RxNORM, LOINC, ICDx
- **General purpose languages:** Prolog

Standards Development Organizations

I am involved in business- and healthcare- oriented working groups in the following organizations:

- **Health Level 7 (HL7)** : Health eDecisions / KNART, FHIR
- **Object Management Group (OMG)** : API4KP, Health BPM
- **RuleML**

Professional Affiliations

I am affiliated with the following organizations:

- **Professional Order of Engineers for the province of Bologna (Italy)**
- **Object Management Group (OMG)**
- **RuleML**

Honors and Awards

SPINNER Technology Transfer Grant Beneficiary (Januray 2007 - June 2007)

Teaching

I have been the lead instructor for the following graduate-level (MS,PhD) classes in the ASU Biomedical Informatics Program:

- **Knowledge Management and Engineering (online) : 2017-**
- **Clinical Applications Development 2015**
- **Foundations of Biomedical Informatics II : 2013-2015**
- **Problem Solving In Biomedical Informatics (Applied Software Engineering) : 2012-2014**
- **Reading and Conference (individualized study) : 2012-**
- **Applied Project / Dissertation : 2012-**

Other Academic Activities

- **Journal Article Reviewer: JBI, ENVSOFT**
- **Program Committee member / Reviewer: AMIA, RuleML, Jurix**
- **Co-chair of several sessions at the RuleML conferences (2008-2014)**
- **Chair of the Healthcare Day at Decision Camp 2013 and 2014**
- **Mentoring of MS and PhD students**