

## Building an Enterprise Ontology in Less Than 90 Days

Dave McComb, Dan Carey & Todd Schneider  
Semantic Arts

[www.semanticarts.com](http://www.semanticarts.com)

November 18, 2015

SEMANTIC TECHNOLOGY FOR  
INTELLIGENCE, DEFENSE, AND SECURITY  
George Mason University, Fairfax, VA

© 2015 Semantic Arts, Inc.

## Building an Enterprise Ontology in 3 Acts

- Act 1 – Do we need to do something different?
- Act 2 – What is it about semantics that will make a difference?
- Act 3 - 6 – “How to’s” for getting this done in 90 days (or less!)

© 2015 Semantic Arts, Inc.

2

## Semantic Arts

- Over the last 15 years, we've been designing and building ontologies for a number of large firms in many different industries



SENTARA



© 2015 Semantic Arts, Inc.

3

## Dan Carey

- Ontologist at Semantic Arts.
- Dan has 30 years of consulting experience, 25 of it designing application databases, logical and physical data models, and data strategies with major IT service firms.
- He has designed semantic technology products to assist in military human resources management, and data exchange standards using OWL and XSD.
- He holds a bachelor's degree in Applied Physics.

© 2015 Semantic Arts, Inc.

4

## Todd Schneider

- Todd has 25 years of experience as a systems engineer and ontologist, primarily in the defense industry
- Lead several initiatives within Raytheon and their clients to integrate semantic technology with enterprise architecture
- Early and frequent participant and contributor to the Net Centric Industry Consortium, and other Net Centric initiatives throughout the federal space.
- He holds a PhD in Mathematics

© 2015 Semantic Arts, Inc.

5

## Please introduce yourselves

- We can still tailor this presentation based on your needs and backgrounds
  - Name
  - Organization
  - Experience with Semantic Technology
  - Experience with other information/data technologies
  - Your specific interest in this topic

© 2015 Semantic Arts, Inc.

6

## Dave McComb

- Founded Semantic Arts in 2000
- Co-founded Semantic Technology conference in 2004
- Wrote Semantics in Business Systems
- Four patents in software engineering including first patent on model driven development
- Worked with dozens of large enterprises at the Enterprise Architecture and Enterprise Application level

© 2015 Semantic Arts, Inc.

7



## Information Systems Cost

- Our information systems cost 10 - 100x what they could or should cost

© 2015 Semantic Arts, Inc.

9

## Why is that?

- |                              |                           |
|------------------------------|---------------------------|
| ▪ Legacy Systems?            | ▪ Certainly a contributor |
| ▪ Vendor Lock in?            | ▪ Adds to it              |
| ▪ Solving the wrong problem? | ▪ Frequently              |
| ▪ Undue Complexity?          | ▪ Getting at the root     |

© 2015 Semantic Arts, Inc.

10

## Architect, Prostitute & System Analyst



© 2015 Semantic Arts, Inc.

11

Yeah, Where did that chaos come from?

- Root-cause analysis lead us to one of the most damaging phrases ever uttered in the Corporate world:

“Let’s not reinvent the wheel”

© 2015 Semantic Arts, Inc.

12

## Sounds innocent, helpful even

- Let's explore how it works its way through a decision process
- Someone says "we need x feature"
- "Surely we're not the first firm to want that. Let's not reinvent the wheel"
- And a search begins to find a product that can be acquired that has "x"

© 2015 Semantic Arts, Inc.


13

## As if...

- ...You needed a washer,



- ...And you discovered this wheel has a washer!



14

Before you know it



Surely they don't...

- ...Implement large monolithic applications just to handle a small variation in data or function.
- Yep, they do



## Examples

- Washington State - Paying employees v. paying w2 providers
- Washington State - 6 referral systems
- Washington State Labor & Industries - 23 systems with Accounts Receivable functionality

© 2015 Semantic Arts, Inc.

17

## Net result

- Most large firms have thousands of major applications
- Each has its own, arbitrarily different data model
  - With thousands of tables and attributes
- Each at an arbitrary level of abstraction,
  - With an arbitrary data structure
  - With arbitrary names
- Leading to millions of distinctions to be mastered
- The implicit and explicit relationships between them are vastly complex

© 2015 Semantic Arts, Inc.

18

## For Example

- SAP - Average SAP Install has 95,000 tables and well over 1 million attributes
- EPIC (Electronic Medical Record) - has 210,000 attributes

© 2015 Semantic Arts, Inc.


## Replacing one of these systems

- Most of the time people don't retire systems, they just build additional ones
- Let's look at what happens when the project is to actually retire an existing application

© 2015 Semantic Arts, Inc.

20

## Implementing a new system

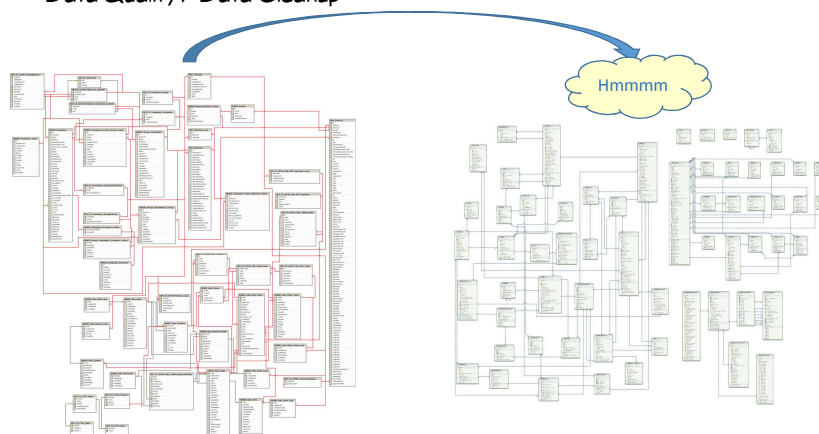


© 2015 Semantic Arts, Inc.

21

## “Data Quality Problems”


Data Quality / Data Cleanup



© 2015 Semantic Arts, Inc.

22

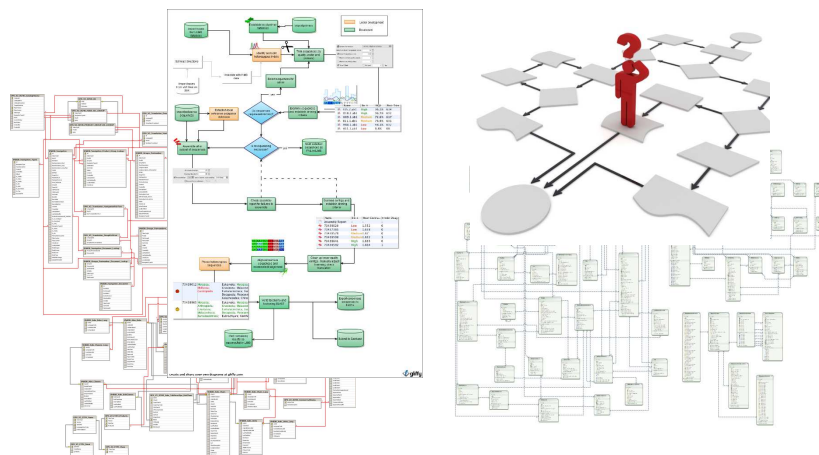
## Integration



© 2015 Semantic Arts, Inc.

23

## Those darn users




"Change Management"

© 2015 Semantic Arts, Inc.

24

Most of these projects end poorly



© 2015 Semantic Arts, Inc.

25

## Healthcare.gov

	HealthCare.gov	HealthSherpa
Cost	\$1 billion to date	Initial Build: 3 guys, 3 weeks. Probably a person year since then
Finding plans	Allows you to search for plans, some eligibility checking, enrollment	Same
Eligibility	Verify income and calculate subsidies	Subsidiary calculator, but doesn't check eligibility
User experience	Healthcare.gov has adopted some of healthsherpa's UI ideas	Initially much better, browse first, then register later
Users enrolled	7.1 million	120,000

© 2015 Semantic Arts, Inc.

26

## Further evidence that these economics are not necessary

- Pinterest – 0 - 10 billion page views/month in 2 years and 40 engineers

<http://highscalability.com/blog/2013/4/15/scaling-pinterest-from-0-to-10s-of-billions-of-page-views-a.html>

- Instagram – 30 million users in 2 years, from 2 engineers to 5

<http://techcrunch.com/2012/04/12/how-to-scale-a-1-billion-startup-a-guide-from-instagram-co-founder-mike-krieger/>

© 2015 Semantic Arts, Inc.

27

## The Question

- Is not: “Is this approach to building and deploying systems dysfunctional?”
  - (It is.)
- The question is: “Why does this approach persist?”

© 2015 Semantic Arts, Inc.

28

## How does the bad idea of adding yet another application to our architecture persist?

Because

- It's in the interest of the vendors
- It's budget-able
- It's what we know
- A credible alternative hasn't been put forward

© 2015 Semantic Arts, Inc.

29

## Steam v. Electricity



© 2015 Semantic Arts, Inc.

30

## We are putting forward two related ideas

- 1) The data-centric approach to systems implementation is the only viable way to break this pattern

And

- 2) The application of semantic technology and enterprise ontologies is crucial to the success of the data-centric approach

© 2015 Semantic Arts, Inc.

31

## Data-Centric

- In a data-centric enterprise, the data is the main asset
- Application functionality comes and goes
- The data remains and is added to
- It does not need to be converted


© 2015 Semantic Arts, Inc.

32



## In order for this to work


- There needs to be an architecture...
- ...Which replaces some of the key functions currently being performed by applications



© 2015 Semantic Arts, Inc.

33


## At the architecture's heart is the data



© 2015 Semantic Arts, Inc.

34

Really a set of coordinated data  
repositories




Some of which are  
internally generated  
and curated and  
others are not

© 2015 Semantic Arts, Inc.

35

With shared meaning




This is where the  
enterprise ontology  
comes in

© 2015 Semantic Arts, Inc.

36

## Other layers




Include such things as  
identity management, security,  
common services,  
producer and consumer ontologies

© 2015 Semantic Arts, Inc.

37

## What is an Enterprise Ontology?

- Definition of a common core of concepts that expose the hidden sameness in what people are talking about,
- That identifies some agreed-upon terms,
- And is represented in a formal notation to support automation.



© 2015 Semantic Arts, Inc.

## Sounds like a Data Dictionary / Controlled Vocabulary?

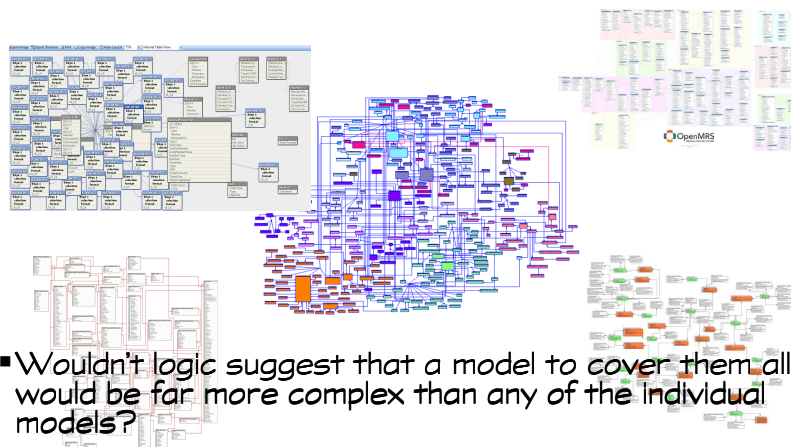
- It is
- But it differs from DD/CV in that a DD/CV is primarily for human consumption (users or system designers)
- An Enterprise Ontology is meant for human consumption, plus to be used directly for system building and system integration, and operation
- It's developed using ontological principals and analyses

© 2015 Semantic Arts, Inc.

39

## Sounds Like an Enterprise Data Model?

- But if each application data model is vastly complex?



- Wouldn't logic suggest that a model to cover them all would be far more complex than any of the individual models?

© 2015 Semantic Arts, Inc.

### In this case...

- The logic is wrong
- It is possible to model an enterprise's information with hundreds of core concepts, plus a manageable constellation of fine-grained distinctions
- Without losing any fidelity

© 2015 Semantic Arts, Inc.

### An Enterprise Ontology

- Can capture the fidelity of the distinctions in an enterprise
- Without succumbing to the temptation to recreate the complexity

© 2015 Semantic Arts, Inc.

What are the Key Difference between a traditional data model and an ontology that allow this?

- Structure
- Complexity
- Explicitness
- Flexibility
- Reusability

© 2015 Semantic Arts, Inc.

## Structure

- It's surprising how much the structure of a structured database contributes to its complexity
- The slightest addition, variation, or change in cardinality, seems to give rise to more tables and more attributes

© 2015 Semantic Arts, Inc.

## Complexity

- IRS - EDM 30,000 Entities
- AT&T - have been working on an EDM for over a decade
- Most Defense Contractors - Multi-year, Multi-thousand entity efforts

© 2015 Semantic Arts, Inc.

## Explicitness

- In traditional models, the meaning is implicit
- It lives in design documents and people's heads
- There is nothing in the "Employee Master Table" to tell you what an employee is
- Only, to tell you a few key attributes we've decided were of interest to us
- In an *Ontology* the meaning can be explicit

© 2015 Semantic Arts, Inc.

## Flexibility

- The structure of a data model is rigid, and programmers rely on that rigidity
- Many programming idioms rely on the structure of the data being recapitulated in the code
- This means that data models are easy to re-factor before development begins, and very hard to re-factor afterward
- This promotes a design style of adding on rather than refactoring, and at some point it is easier to create a new system than to add on
- Ontology implementation is through a graph database that is inherently flexible

© 2015 Semantic Arts, Inc.


## Reuseability

- There is no extension mechanism for Data Models
- There is no inheritance model in relational, and while some developers have developed patterns for simulating inheritance it is not widespread
- Most modelers add additional tables and columns in order to extend a model, or create an entirely new model and application
- Modularity, inheritance and importation are baked into Ontologies
- As a result it is usually easier to make small extensions, without disruption

© 2015 Semantic Arts, Inc.



## Oliver Wendell Holmes



"I would not give a fig for the simplicity this side of complexity, but I would give my right arm for the simplicity on the other side of complexity."

© 2015 Semantic Arts, Inc.

## Some Evidence that Semantics Helps

- Sallie Mae
- Secretary of State (WA)
- Sentara
- Schneider-Electric

© 2015 Semantic Arts, Inc.

50


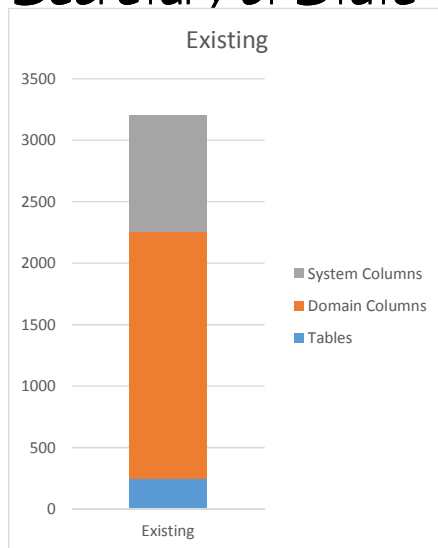
## Case Study: Reducing Complexity at Sallie Mae

	tables	attributes
Class	582	10,230
LoanCons	133	15,295
Eagle I	356	13,538
Eagle II	464	12,502
	1,535	51,565

Category	Metric
Classes	366
Object Properties	224
Datatype Properties	13
Individuals (Categories mostly)	227
Total TBox Axioms	1,030

© 2015 Semantic Arts, Inc.

## Case Study: Reducing Complexity at WA Secretary of State



© 2015 Semantic Arts, Inc.

52

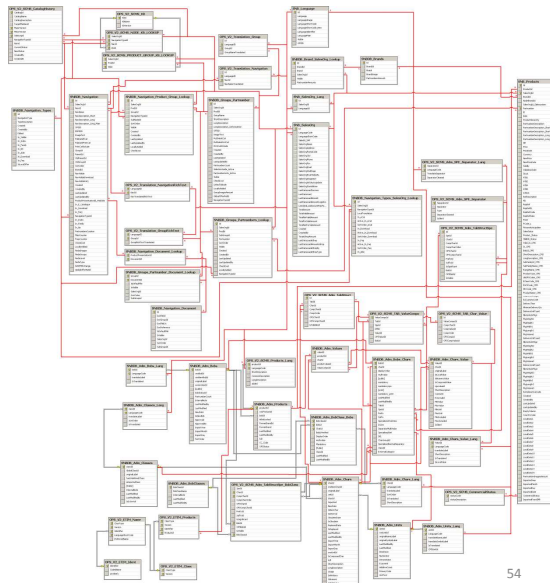
## Case Study: EO for Sentara Healthcare

- Enterprise Ontology
- Scope: entire enterprise: healthcare, assisted living, insurance and internal systems
- Size: 1,276 classes/397 properties
- Applications:
  - We co-developed a Proof of Concept in Asthma and COPD, and the newly uncovered concepts we are finding are directly derivable from the core.
  - Currently being used as basis for semantic enterprise search. More than adequate coverage.
  - Flexible way to manage Physicians data across multiple sources


© 2015 Semantic Arts, Inc.

## Case Study: Schneider Electric

- Existing product catalog systems has about 700 tables and 7000 attributes in total
- To date the new system has populated 46 classes and used 36 properties.
- We expect this to slightly more than double, to include the full scope

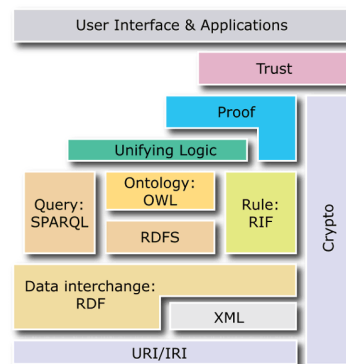


54

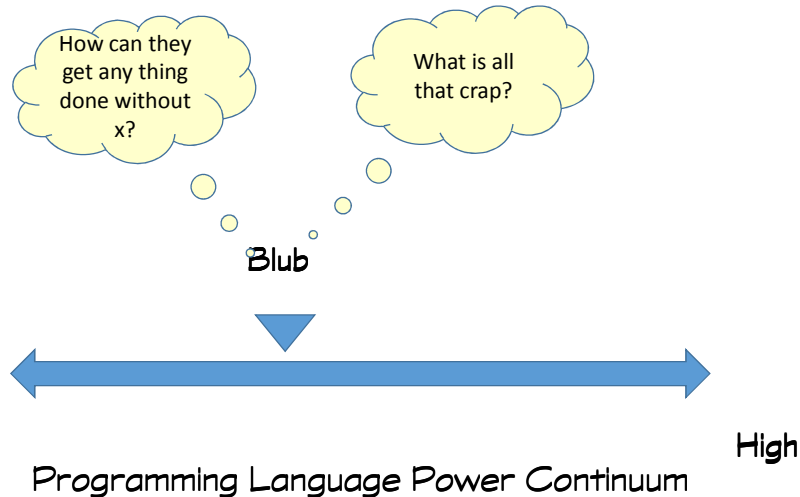


## Semantic Technology

- **Semantics** - Pertaining to the study of meaning
- **Semantic Technology** - software and methods that rely on representing meaning, especially those that are based on the Semantic Web stack as standardized by the W3C
- However, information systems don't understand 'meaning'; they interpret symbols.
  - OWL and ontologies built using it provide a way to better control how they interpret the symbols.



## Blub Paradox



© 2015 Semantic Arts, Inc.

57

## Developers and Semantics

- Most enterprise developers have come from either Relational, Object-Oriented or JSON-style document backgrounds
- When building ontologies, they tend to build them to resemble what they are used to
- Then they wonder why it doesn't do X as well as their approach of choice
- And what's all this "inference" and "open world" stuff, anyway?

© 2015 Semantic Arts, Inc.

58

## What Differentiates *Ontology*?

- All information representation is reduced to a single data structure: the “triple”




© 2015 Semantic Arts, Inc.

59

## By convention

- The three parts of a triple are called...

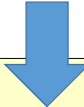


© 2015 Semantic Arts, Inc.

60

## Useful Global Identifiers

- All systems create identifiers for the many things they need to identify and distinguish
- The typical approach is to create a primary key on a table

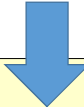


SecretAgent	
ID	Name
007	Bond, James
002	Fairbanks, Bill

© 2015 Semantic Arts, Inc.

## There are problems with this approach

- These numbers only mean anything in the context of:
  - This database
  - This table
  - This column
- They embed context and hide usage



SecretAgent	
ID	Name
007	Bond, James
002	Fairbanks, Bill

© 2015 Semantic Arts, Inc.

## This is one reason existing systems are complex

- If you want to get information from multiple different tables
- You have to write a query that "joins" the tables
- ... WHERE `SecretAgent.ID = Gadgets.Assignee` And `SecretAgent.ID = FemmeFatale.Foe`

SecretAgent	
ID	Name
007	Bond, James
002	Fairbanks, Bill

Gadgets		
Number	Desc	Assignee
001	Shoe Dagger	002
002	Rocket Belt	007

Villain		
Code	Name	Protag
BB104	No, Dr. Julius	007
BB105	Goldfinger, Auric	007

FemmeFatale		
Code	Name	Foe
HB	Brandt, Helga	007
FV	Volpe, Fiona	007

© 2015 Semantic Arts, Inc.

63

## This is one reason existing systems are complex

- A human has to know all the (implicit) metadata to construct this join for each use
- This only works between tables in the same database

SecretAgent	
ID	Name
007	Bond, James
002	Fairbanks, Bill

Gadgets		
Number	Desc	Assignee
001	Shoe Dagger	002
002	Rocket Belt	007

Villain		
Code	Name	Protag
BB104	No, Dr. Julius	007
BB105	Goldfinger, Auric	007

FemmeFatale		
Code	Name	Foe
HB	Brandt, Helga	007
FV	Volpe, Fiona	007

© 2015 Semantic Arts, Inc.

64




## We need context-independent identifiers

- GUIDs
- May be globally unique, but not globally resolvable
- Nor are they easily interpretable
  - What do you do when someone sends you a message with this in it?
  - {21EC2020-3AEA-1069-A2DD-08002B30309D}

© 2015 Semantic Arts, Inc.

## W3C's Semantic Technology Stack Finesses this with the "URI"

- Uniform Resource Identifier
  - Analogous to the URLs we type into our browsers
- Each part in a triple is identified by a URI



© 2015 Semantic Arts, Inc.

## W3C's Semantic Technology Stack Finesses this with the "URI"

- This identifier is truly global
- It means the same thing, regardless of the database or document it is part of

**cusip:02209S103** = <http://cusip.com/reg#02209S103>

- The metadata has been made explicit
- In turn, this yields information that can be "joined" without relying on humans or additional (implicit) knowledge of the metadata

© 2015 Semantic Arts, Inc.

67

## Using URIs (similar to URLs) as identifiers

- Gives us truly global unique IDs
- That can be looked up, if needed

<http://isbn10.isbn.org/books#1558609172>

Namespace

Fragment




This is the key to "self-assembling structures"

© 2015 Semantic Arts, Inc.

68

## Let's Visualize this Process

- Lets say we the following triples were harvested from completely different sources
- CUSIP: Identifier for a Financial Securities



© 2015 Semantic Arts, Inc.

69

## Tinker toys




© 2015 Semantic Arts, Inc.

70

## Accommodates change in place

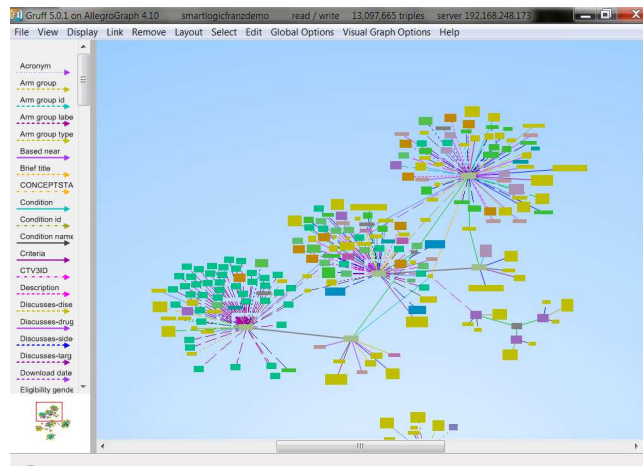
- A semantic system can evolve in place
- Example: Tasks -> Projects
  - > Backlogs -> Assignments
  - > Expenses



© 2015 Semantic Arts, Inc.

71

Of course these graphs can get more complex than you could represent with Tinker Toys



© 2015 Semantic Arts, Inc.

72

## Relation of Metadata to Data


- Metadata is also represented as triples
- It is made explicit
- It can be queried

© 2015 Semantic Arts, Inc.

73

## Including Metadata

- Metadata is the data that defines your current systems



© 2015 Semantic Arts, Inc.

74

## Triples from RDB

- One of the most common sources for triples is from existing relational databases

© 2015 Semantic Arts, Inc.

75

## Making Assertions, Traditionally

- In a traditional system, we make assertions by putting data in tables

Person 1 is named John

PersonId	Name
1	John



© 2015 Semantic Arts, Inc.

76

## Making Assertions, Traditionally

- Including relationships

John is the parent of Dave.

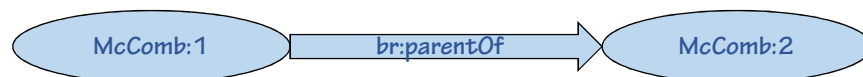
PersonId	Name	ParentOf
1	John	2
2	Dave	3

© 2015 Semantic Arts, Inc.

77

## Where do triples come from?

- Triples can come from existing systems

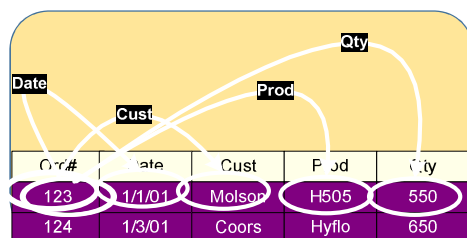


PersonId	Name	ParentOf
1	John	2
2	Dave	3

© 2015 Semantic Arts, Inc.


78

## Where do triples come from?



Date	Cust	Prod	Qty
123	1/1/01	Molson	H505
124	1/3/01	Coors	Hyflo

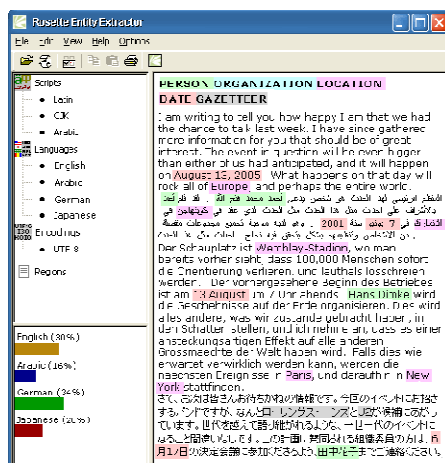
Databases



© 2015 Semantic Arts, Inc.

## Where do triples come from?

Documents



PERSON ORGANIZATION LOCATION  
 DATE GAZETTEER

I am writing to tell you how happy I am that we had the chance to talk last week. I have since gathered more information for you that should be of great interest. The event in question will no longer happen than either of us had anticipated, and it will happen on August 15, 2005. What happens on that day will rock all of Europe, and perhaps the entire world.

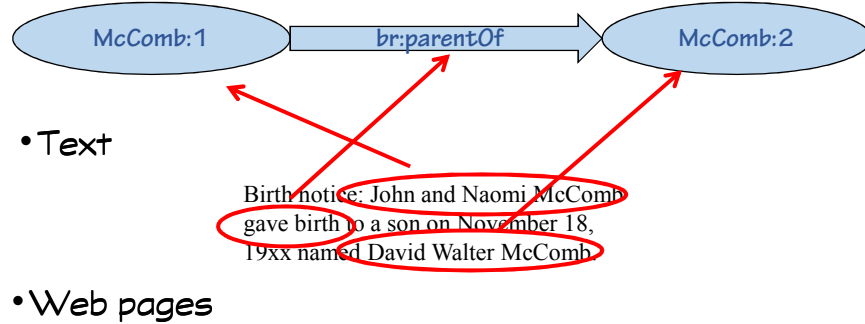
Der Schauplatz ist Wembley-Stadion, wo man bereits vorher sah, dass 100.000 Menschen sofort die Orientierung verlieren und letztlich losbrechen werden. Die vorangegangene Saison des Betriebes ist am 13. August im / Unruhig. Hans Dinkel wird die Verantwortlichen auf der Erde organisieren. Das wird alles anders, was wir zusammengebracht haben, in der Schatulle stellen, und ich meine, dass es eine abschließende Effekt auf alle anderen Grossmeister der Welt haben wird. Falls das wie erwartet verwirklicht werden kann, werden die nächsten Ereignisse in Paris, und darauf in New York stattfinden.

すごく、お気遣いありがとうございます。今回のイベントには、さく、お気遣いありがとうございます。今回のイベントには、さく、お気遣いありがとうございます。

© 2015 Semantic Arts, Inc.



## Where do triples come from?



© 2015 Semantic Arts, Inc.

81

## Where do triples come from?

### XML and HTML

```
<div class="vcard" style="background-color: F9C; width: 300px; height
  <span class="fn n">
    <span class="honorific-prefix">Mr. </span>
    <span class="given-name">Doug</span>
    <span class="family-name">Mahugh</span>
  </span>
  <br />
  <span class="adr">
    <span class="street-address">One Microsoft Way</span>
    <br />
    <span class="locality">Redmond</span>
  </span>
  <br />
  <abbr class="region">WA</abbr>
  <span class="postal-code">98052</span>
  </span>
  <br />
  <a class="email" href="mailto:dm%61h%75%67h@m%69%63%72
  <br />
  <span class="tel">
    Phone:
    <span class="value">+1-425-882-8080</span>
  </span>
</div>
```

© 2015 Semantic Arts, Inc.

## Where do triples come from?



- Social Media

© 2015 Semantic Arts, Inc.

83

## Where do triples come from?

- The Web



**WIKIPEDIA**  
The Free Encyclopedia

Main page  
Contents  
Featured content  
Current events  
Random article  
Donate to Wikipedia  
Wikimedia Shop

Interaction  
Help  
About Wikipedia  
Community portal  
Recent changes  
Contact Wikipedia

Toolsbox  
Print/export

Languages  
العربية  
Български  
Brezhoneg  
Català  
Cymraeg  
Deutsch  
Español  
Esperanto  
فارسی  
Français  
ગુજરાતી

# subject

## Fort Collins, Colorado

From Wikipedia, the free encyclopedia  
(30 redirects from 145 endpoints)

**Fort Collins** is a Home Rule Municipality situated on the Cache La Poudre River along the Colorado Front Range, and is the county seat and most populous city of Larimer County, Colorado, United States. Fort Collins is located 57 miles (92 km) north of the Colorado State Capitol in Denver. With a 2010 census population of 143,986, it is the fourth most populous city in Colorado.<sup>[4]</sup> Fort Collins is a large college town, home to Colorado State University. It was named Money magazine's Best Place to Live in the U.S. in 2006, #2 in 2009, and #5 in 2010.<sup>[5]</sup>

### Contents

- 1 History
- 2 Geography and climate
  - 2.1 Neighboring cities
- 3 Demographics
- 4 Law and government
- 5 Culture
  - 5.1 Communications
  - 5.2 Education
- 6 Economy
  - 6.1 Major industries and commercial activity
  - 6.2 Sustainability Programs
- 7 Transportation
  - 7.1 Commercial shipping
- 8 Facilities
- 9 Notable natives and residents
- 10 See also
- 11 References
- 12 External links

### History



Colorado  
Fort Collins

City of Fort Collins



Downtown "Old Town"



Location of Fort Collins shown within Colorado.  
Coordinates: 40°33′33″ N﻿ 105°40′3″ W

**Fort Collins**

City and County

**Commissi**  
United States Army colonel William O. Collins

**Named for**  
February 12, 1883<sup>[2]</sup>

**Government**

- **Type**
- **Mayor**
- **Mayor pro tem**
- **City Manager**

Home Rule Municipality<sup>[1]</sup>

- **Mayor**  
Karen Wefelbunat
- **Mayor pro tem**  
Kelly Ohlsen
- **City Manager**  
Darin Abernethy

**Area**

- **Total**  
53.38 sq mi (122.1 km<sup>2</sup>)
- **Land**  
46.6 sq mi (120.5 km<sup>2</sup>)
- **Water**  
6.6 sq mi (1.6 km<sup>2</sup>) 1.27%

**Elevation**  
5,003 ft (1,525 m)

**Population** (2010)

- **Total**  
143,986
- **Density**  
2,549.3/sq mi (984.4/km<sup>2</sup>)

**Time zone**  
MST (UTC−7)

• **Summer (DST)**  
MDT (UTC−6)

**ZIP Code(s)**  
80521–80528, 80553

**Area code(s)**  
970

**FIPS code**  
08-27425

**GHS feature ID**  
0204673

**Highways**  
I-25, US 287, SH 1, SH 14


**Website**  
www.fcgov.com<sup>[6]</sup>

Fourth most populous Colorado city



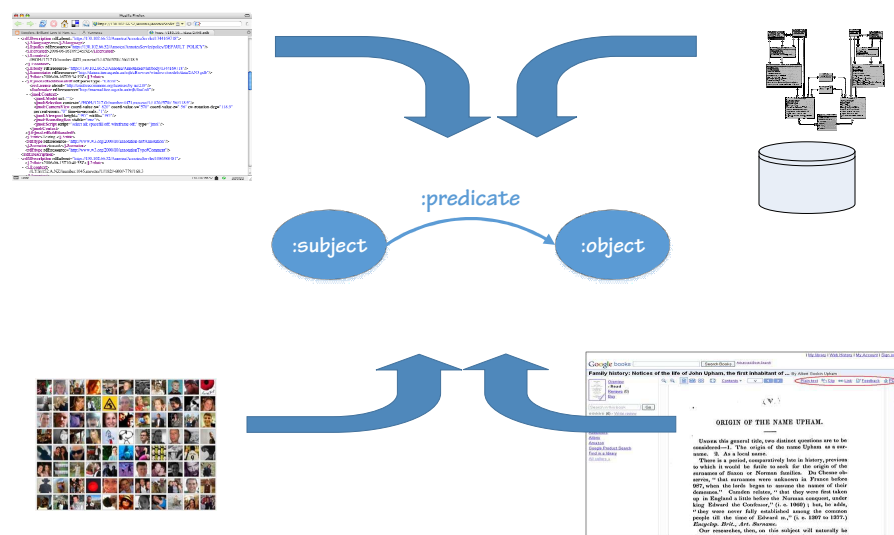
Colorado  
Fort Collins

# DBpedia



© 2015 Semantic Arts, Inc.

## Triple as Common Denominator

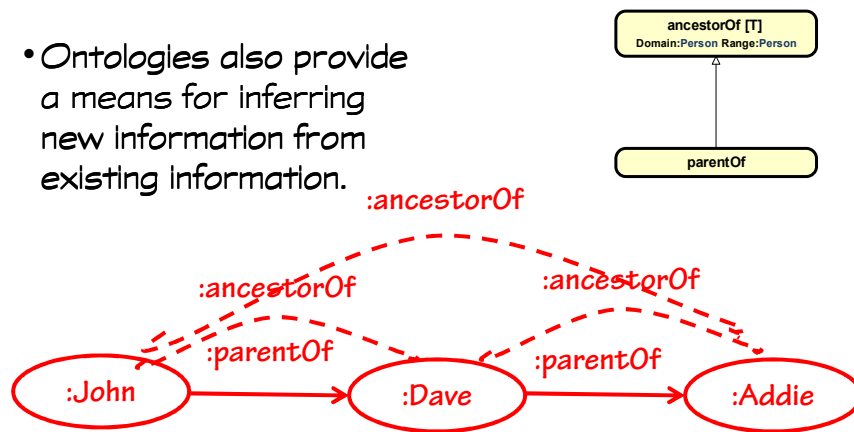


© 2015 Semantic Arts, Inc.

86

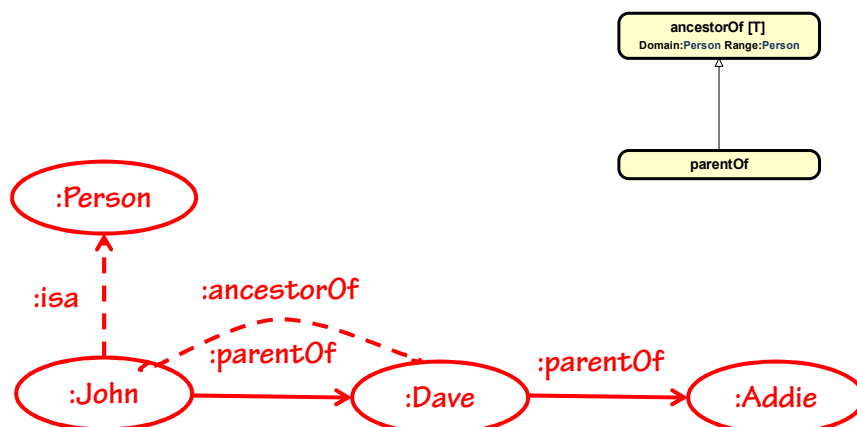
## Inferring new information from existing

- Ontologies also provide a means for inferring new information from existing information.



© 2015 Semantic Arts, Inc.


## The most important new information we infer is class membership



© 2015 Semantic Arts, Inc.


And because class membership is just another triple

- The same instance can be in many classes simultaneously.



© 2015 Semantic Arts, Inc.

## Description Logics




- Description Logics give a way of define the membership criteria for classes.
- This means that computer systems can do a lot of the classification that we currently have humans do.
- And that makes the rules transparent.  
(Current approaches bury meaning in code and specifications.)

© 2015 Semantic Arts, Inc.

90

## Modularity

- The way that modularity works with ontologies is very cool




© 2015 Semantic Arts, Inc.

91

## Importing Ontologies

- Because they have no structure (beyond triples), it is very easy to extend and reuse ontologies

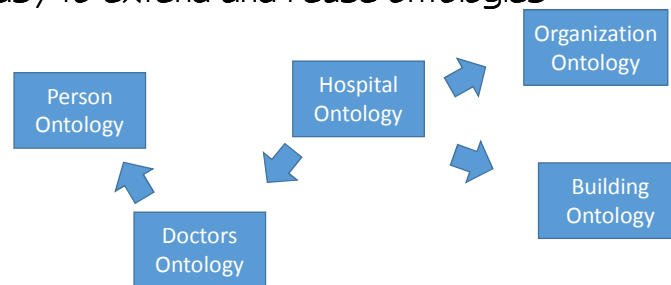


A doctors ontology could import a Person ontology and use that as its starting point (people have dates of birth, and addresses and the like)

© 2015 Semantic Arts, Inc.

## Importing Ontologies

- Because they have no structure (beyond triples), it is very easy to extend and reuse ontologies



A hospital ontology could not only import and use the doctor ontology, it could create a superclass of doctor, called provider which also might include organizations. It can superclass another class without the other class or ontology being aware.

© 2015 Semantic Arts, Inc.

## Versus Relational

- In relational paradigms, there is no real concept of extending a model
- At best, developers copy one schema for a starting point and change it from there

© 2015 Semantic Arts, Inc.

94

## Versus Object-Oriented

- In object-oriented approaches, you can extend a common model through inheritance
- But there is no mechanism to create new classes (even superclasses) in descendent models
- Further, with semantic technology, the order in which the ontology extensions are created is not important. You can create the sub-ontology first and unify it with another later

© 2015 Semantic Arts, Inc.

95

## Review: Characteristics of Semantics that Make it Uniquely Suited to Role as EO

© 2015 Semantic Arts, Inc.

96



## Value of a Triplestore / Graph DB

- In general, the main advantage of a Graph DB over a Relational DB lies in its flexibility, and the ability to evolve in place
- We're going to go over three specific examples of how that manifests itself

© 2015 Semantic Arts, Inc.

## Global / Resolvable identifiers

- In a relational system, all IDs are contextual; you need to know the DB, the table and the column to use the identifier
- In a triplestore, all identifiers are URIs. They are globally unique and resolvable

Once we assign a URI to a individual or concept, it means exactly the same thing, no matter what database or repository it is in.

**What this means is that “joins” are done automatically for you by the system**

© 2015 Semantic Arts, Inc.

## The relationship of Schema to Data

- In a relational system, a table must exist before you can put data in it
- In a graph DB, you can create data and later associate it with newly created classes

For example, in a triplestore, you could start with people and medical treatments, and later create the idea of a patient, without having to redo any of the data you've already committed

**What this means is that you can start simple and add as you need.**

© 2015 Semantic Arts, Inc.

## Instances can belong to many classes

- In a relational system, a row exists in only one table
- In a graph DB, you can create data and later associate it with newly created classes

For example, in a triplestore, you could have an instance (a URI) that represents a person, who is both a patient and a provider. In a relational system, you create two records and then have to have another way to determine that they are the same.

**What this means is a great reduction in redundancy**

© 2015 Semantic Arts, Inc.

## The Relationship of Applications and Schema

- In a relational system, schemas are built to the needs of applications. They are owned by the application
- In a triplestore, the schema is a shared resource. Many applications cooperate around the same model

Traditionally, every application has its own data model. This is what drives the cost of integration so high. In a semantic system, we share as much modeling as we can, running the integration costs down.

**What this means is that data integration is almost a bi-product of building semantically**

© 2015 Semantic Arts, Inc.

## Our Quest

- Find the stability
- The enduring business themes

© 2015 Semantic Arts, Inc.



## Abstract until everyone agrees

Credit Default Swap

Obligation





© 2015 Semantic Arts, Inc.

## Core

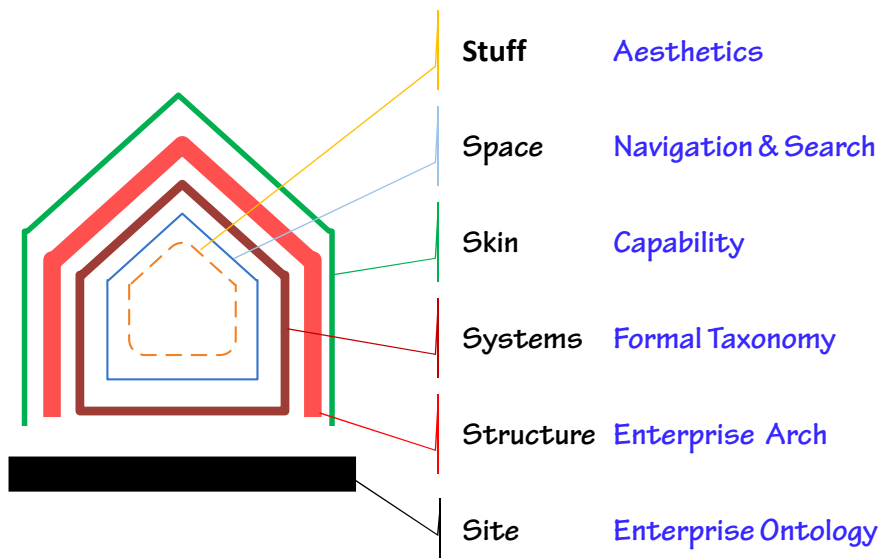
- We're looking for a core that is stable
- Not unchanging
- Just rooted
- In such a way that we can let the things around it change at their natural rate

© 2015 Semantic Arts, Inc.

## An Analogy: Pace Layering



## An Analogy: Pace Layering



## Our Position

- **A Core Enterprise Ontology:**
  - **Is Necessary** - in order to provide some sort of framework and stability to the rest of your information systems efforts
  - **Should be Understandable** - the value expands greatly when it is well understood and used
  - **Can be Elegant** - there are typically a few hundred concepts that cover most information system activity
  - **Can be Flexible** - the model should be able to evolve in place and be extended by some groups without impacting all

© 2015 Semantic Arts, Inc.

109

## At the end of the EO process

- **You will have a model that is 1% as complex as your current models**
- **And that seems to gain as much as it loses in fidelity in use**

© 2015 Semantic Arts, Inc.

110

Example of how a simple ontology isn't lossy



© 2015 SemanticArts, Inc.

111

**Act 3**

**6 “how to’s” for getting this  
done in 90 days or less**



## Getting Started

- Most companies don't want to "boil the ocean" with an Enterprise Data Modeling effort
- They believe (correctly) that this could take 1-2 years, with questionable payoff
- We're going to suggest some techniques for doing this in a much more streamlined fashion
- You could complete this in 3 months
- If you also do it agilely, you could begin using it in 1

© 2015 Semantic Arts, Inc.

113

## Six techniques / methods for getting your EO complete in 90 days

1. Separate your artifacts by purpose
2. Use gist
3. Model the real world
4. Economize expression
5. Postulate the solution, don't extract it
6. Use Inference to check for errors

© 2015 Semantic Arts, Inc.

114

## 1) separate your artifacts by purpose

- Taxonomy / Ontology Assessment
- Became Knowledge Artifact Assessment

© 2015 Semantic Arts, Inc.

115

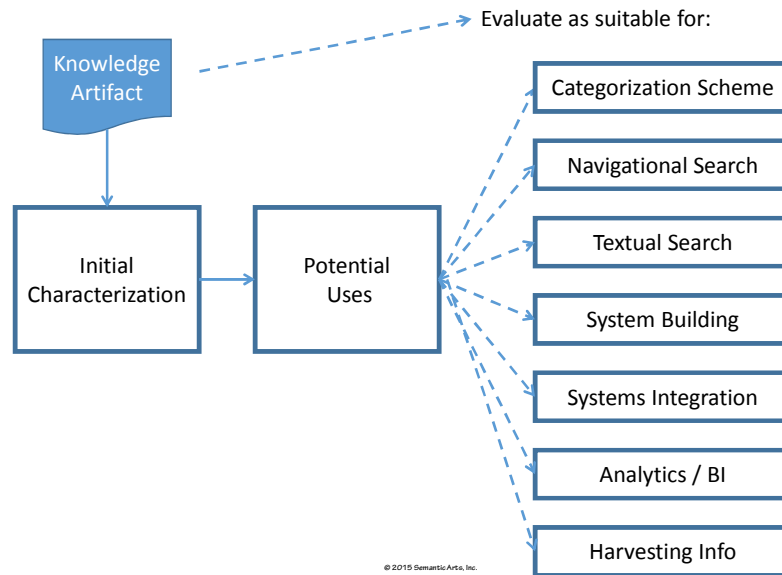
## Purpose-Driven

- The assessment of an artifact turned out to be very contextual
- And the main context is “what is the intended purpose for this knowledge artifact?”
- Some characteristics of a taxonomy (for example MECE (Mutually Exclusive/ Completely Exhaustive)) are very important for some purposes, and get in the way for others

© 2015 Semantic Arts, Inc.

116

## For each Knowledge Artifact



## We've discovered

- These purposes have wildly different definitions of "goodness"
- An artifact for textual search wants to have many, many synonyms for every term
  - (One of our clients said they had on average 29 synonyms for every term)
- By contrast, having lots of synonyms tends to confuse things during integration and system-building

## Modularity to the rescue

- In the 2nd Act, we mentioned the power of modularity
- Now we can begin to harness that modularity
- Rather than either throw away all the excessive synonyms
  - Or load up on them
  - It's not an either/or question
- Have the core, and modularly extend it with the synonyms
- If you don't need the synonyms; only get the core
  - If you do get them, organize them around the core

© 2015 Semantic Arts, Inc.

119

## What we've found


- Many classes are merely taxonomic differences
  - It doesn't add anything to make them classes
  - And it interferes with the elegance of the model
- Moving them to modules that can be under separate governance speeds development and eases maintenance

© 2015 Semantic Arts, Inc.

120

## Fractal Modeling

- Move most of your taxonomies out of the class structure
- Separation of governance
- Simplification of the model




© 2015 Semantic Arts, Inc.

121

## Move Minor Distinctions to Taxonomies

- We're tempted to put everything we know into our ontologies
- Many distinctions are best "pushed" to taxonomies
- Where mere mortals can debate and rearrange them
- Without destabilizing the ontology



© 2015 Semantic Arts, Inc.

## Special Individuals

- Most ontologies have a small number of “special” individuals/instances
- For example, the only semantic distinction between accounts payable and accounts receivable is who “you” are (your firm typically)
- Sooner or later, this becomes a definition based on an instance or several instances
  - (myOnt.\_myCompany0001)
- Other special instances include those that participate in definitions that cannot be formally defined practically and have to be accepted
  - (“male” and “female”, or “exempt” and “non-exempt”, for instance)

© 2015 Semantic Arts, Inc.

123

## Classes and Taxonomic Instances

- Strive to be more like GeoNames than Snomed

	GeoNames	Snomed
Concepts	10 million geographical names	Presumably millions
Classes	19	303,035
Properties	33	152
Taxonomic Categories	645 “feature codes”	In classes

© 2015 Semantic Arts, Inc.

124

## Separation of Concerns

- Meaning (OWL) vs Structure (RDF Shapes)
- If you're trying to answer the question "which properties go on this class" by looking at the ontology, you've collapsed these two ideas
- The ontology should be the province of coining new terms, establishing meaning and providing the rules for inference

© 2015 Semantic Arts, Inc.

125

## RDF Shape Example

```
<ProductRef> {
  rdf:type (spo:ProductReference)
,
  gist:identifiedBy @<ProductID>
,
  spo:describedBy @<ProductReferenceDescription> ?
,
  gist:memberOf (@<ProductRange> | @<ProductSubRange>
|@<ProductSubSubRange> |@<NewProductRange>) *
,
  gist:categorizedBy @< ProductOrComponentType> *
,
  gist:categorizedBy @< ProductFunction> *
,
  gist:conformsTo @<Standard> *
,
  gist:specifiedBy (@<SpecEntry> |@<TabularSpecEntry>)?
}
```

© 2015 Semantic Arts, Inc.

126

## 2) Use gist



<http://semanticarts.com/gist>



© 2015 Semantic Arts, Inc.

127

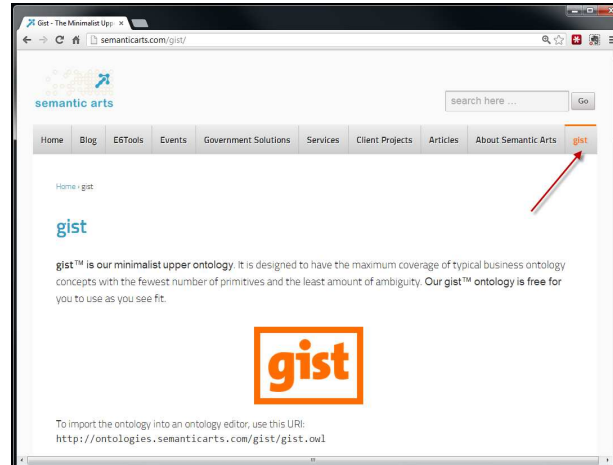
## Introducing gist

- An upper enterprise ontology containing a minimal set of concepts required by most businesses.
- Copyright Semantic Arts, Inc.  
Rights to use are conveyed under the [Creative Commons Attribution-ShareAlike 3.0 license](http://creativecommons.org/licenses/by-sa/3.0/).
- Current version at <http://ontologies.semanticarts.com/gist/gist.owl>
- Consists of a core plus several "subgists."

© 2015 Semantic Arts, Inc.



## Get gist from our web site



([www.semanticarts.com/gist](http://www.semanticarts.com/gist))

© 2015 Semantic Arts, Inc.

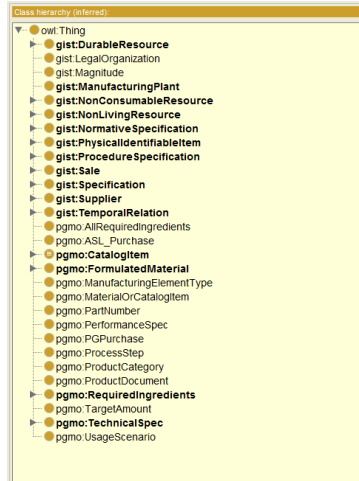
## gist Attribution

If you use gist, please include the following attribution:

This work is derived from and/or directly uses concepts from gist, a copyrighted ontology from Semantic Arts, Inc. Rights to use are conveyed under the Creative Commons Attribution-ShareAlike 3.0 license <http://creativecommons.org/licenses/by-sa/3.0/legalcode>

© 2015 Semantic Arts, Inc.

## Early Projects



- About half the EO classes were derived from gist


© 2015 Semantic Arts, Inc.

131

## Evolution



- We evolved gist and our methodology



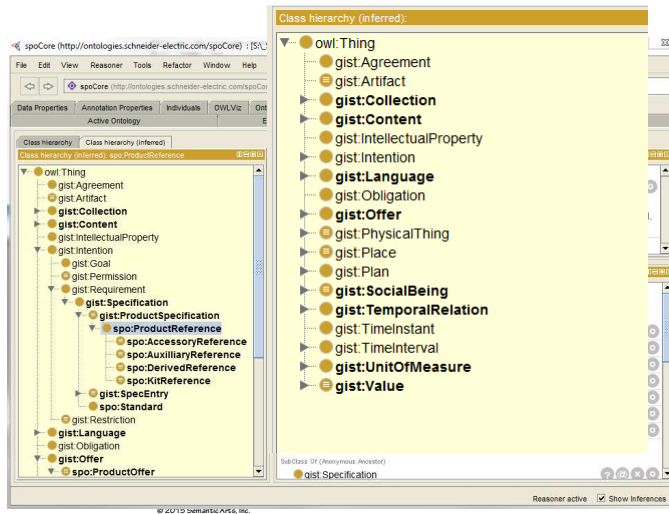
© 2015 Semantic Arts, Inc.

132

## Current Projects



- Most classes derived from gist, without even trying



## What is gist?

- Fewest concepts
- Broadly agreed on, across industries
- That cover most, of most, enterprises
- Least ambiguity
- Stable, ten years old (used in about a dozen major projects)
- Evolving (we keep refining it)

## Abstract?

- Many people think that to cover an entire enterprise with a few hundred concepts, they'd have to be pretty abstract
- And some upper ontologies are quite abstract
- One, for instance, has a high-level distinction between "endurants" and "perdurants" (things v. events).
- We think of these as "abstract abstractions"

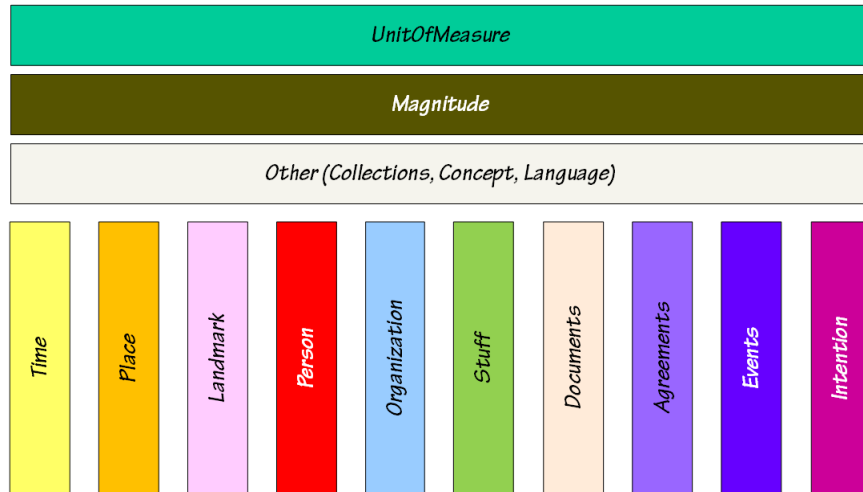
© 2015 Semantic Arts, Inc.

## Concrete Abstractions

- We're trying to work at the level of "concrete abstractions"
- Classes where the members are easily grasped, if slightly abstract
- Person is a concrete abstraction
- We can create an instance of Person, and as we learn more about the person, we may decide (by assertion or inference) that they are also more specific types of persons (Doctors, Brokers, Adults, etc.)

© 2015 Semantic Arts, Inc.

## gist – Major Families of Classes



© 2015 Semantic Arts, Inc.

## Learning gist

- Manageable number of concepts

Ontology metrics:	
Axiom	1275
Logical axiom count	454
Class count	139
Object property count	108
Data property count	20
Individual count	11
DL expressivity	SROIQ(D)

© 2015 Semantic Arts, Inc.

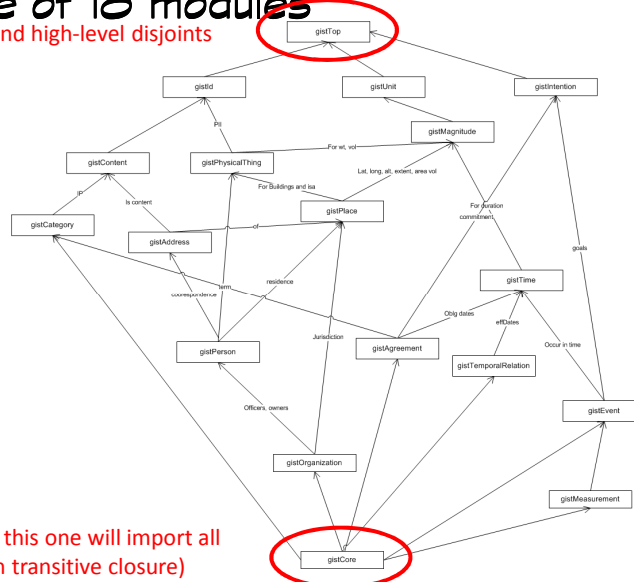
## Modular

- As of 7.x, gist is very modular
- Understanding how the modules fit together adds a bit of conceptual baggage
- However, each module now is so simple as to be almost self-explanatory

© 2015 Semantic Arts, Inc.

## Gist is made of 18 modules

Primitive concepts and high-level disjoints

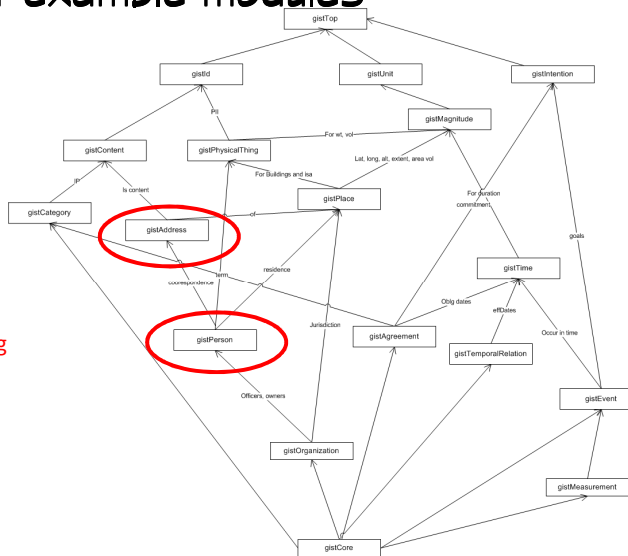


© 2015 Semantic Arts, Inc.

## An couple of example modules

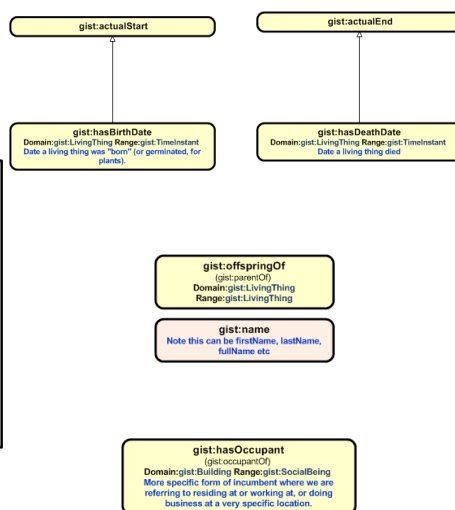
Some very useful  
primitives in an area  
often confused

Person as (once) Living Thing



© 2015 Semantic Arts, Inc.

## Person



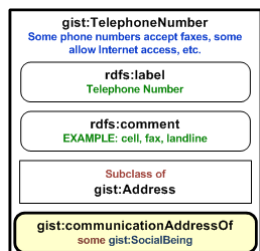
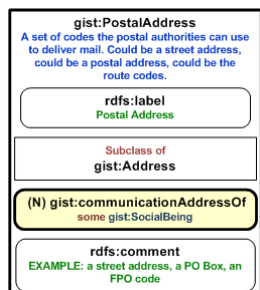
© 2015 Semantic Arts, Inc.

# Address

- We treat address as a first-class object (not an attribute of a person or company)
- And the use of that address by a Person or Organization as essentially a communication preference
- This one change makes the chaos of addresses in most enterprises manageable

© 2015 Semantic Arts, Inc.

# Address





### 3) Model the real world

- It's very easy to fall into the trap of modeling the concepts you find in the application
  - (or in people's heads, which often came from an application)
- Many of them are ok
- But the only way to know which is which, is to try to get as close to the real world as you can.
  - It will shine a light on which are contrived
- Some examples: PIMS, Tabulars, Sections, most junction records., legs, and most booleans

© 2015 Semantic Arts, Inc.

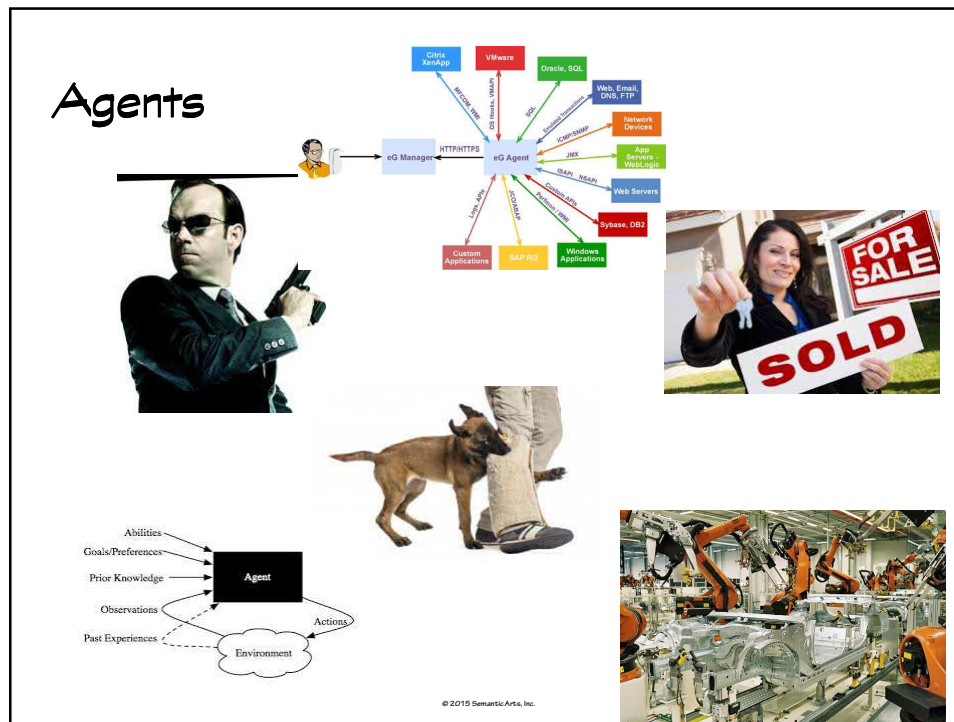
145

### Stay away from abstract abstractions

- Person and Agent are both abstractions
- But pretty much everyone (other than foaf) agrees on what a person is
- But agent...
  - In many cases it is Person or Organization
  - But sometimes machine
  - Or program
  - Or animal
- The real meaning is in the property (the agent on "wasBittenBy" is Animal, whereas the agent on "durablePowerOfAttorney" is Person.)

© 2015 Semantic Arts, Inc.

146




## 4) Economize expression

- It's tempting to put everything you know plus everything you learn in the ontology
- "you might need it"
- These things clutter up the result
- And confuse the use
- Many are unlikely to be widely agreed upon

## In an Enterprise Ontology

- Less is more
- Don't get paid by the pound
- Remember what happened with lines of code?



© 2015 Semantic Arts, Inc.

## Reducing Cognitive Load

- # of things you must know to be competent with the ontology
- # of things you must be in agreement with in order to commit to the ontology
- # of concepts shared

© 2015 Semantic Arts, Inc.

150

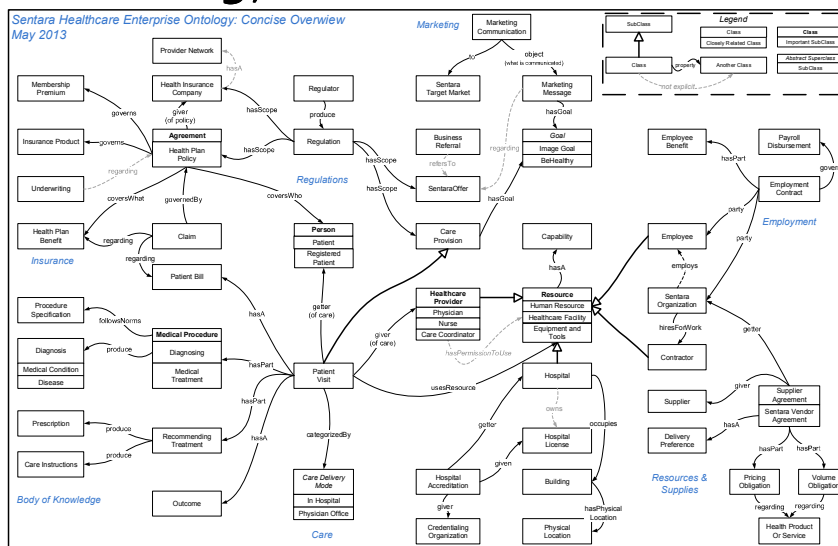
## Economizing Properties

- Typical large enterprises have millions of properties (attributes/ columns ) in their legacy systems
- This is mostly a product of arbitrary design decisions
- We need to be vigilant
- Properties (with a few very narrow exceptions) cannot be formally defined; we must learn them all
- Our target is to get to a few hundred

© 2015 Semantic Arts, Inc.

151

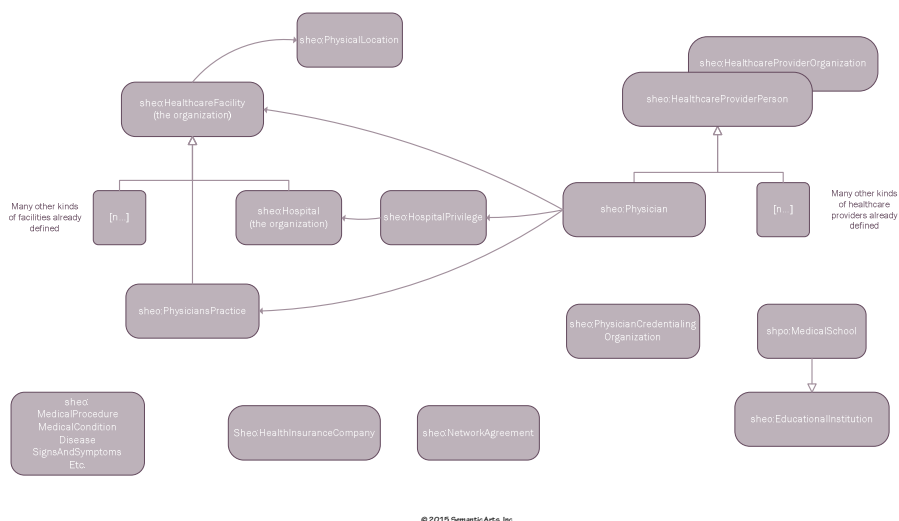
## The Ontology Sketch



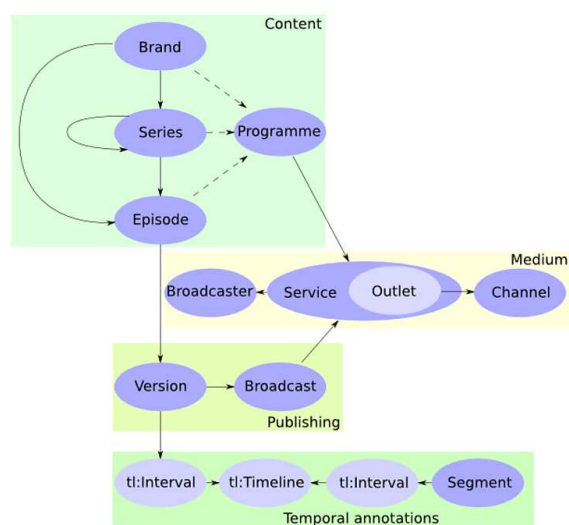
© 2015 Semantic Arts, Inc.

152

## Overall Shape of the Ontology



## BBC Programmes Ontology sketch



154

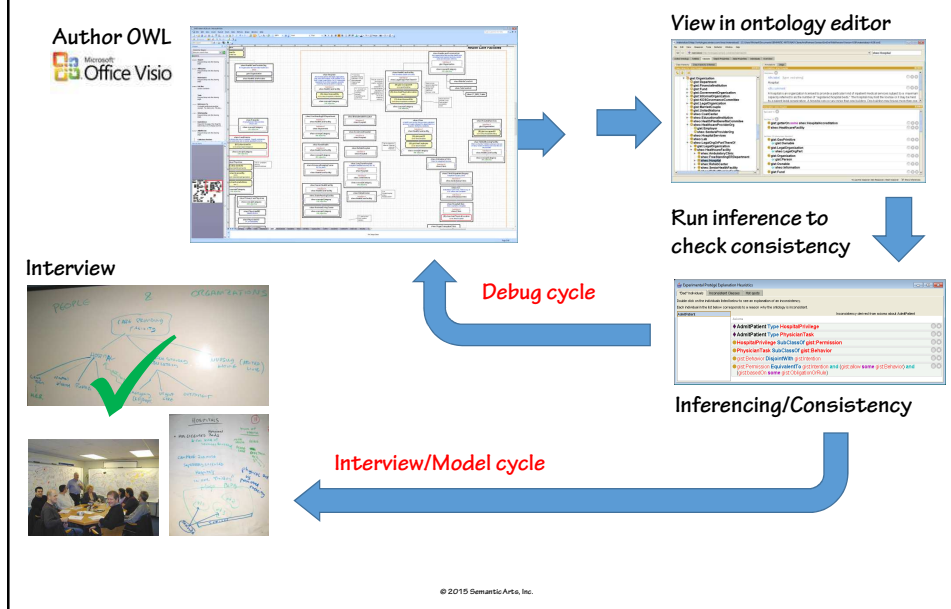
## 5) Postulate the solution, don't extract it

- Over the last several years we have been moving from a "discovery model" to a "postulate model" for ontology development

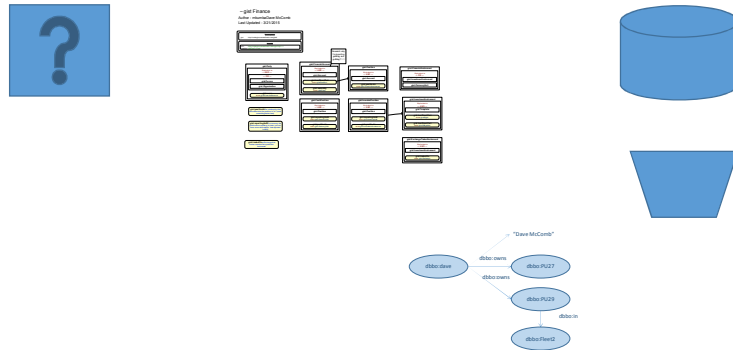
© 2015 Semantic Arts, Inc.

155

## Discovery Methodology



## Postulate Methodology



© 2015 Semantic Arts, Inc.

157

## Then check existing systems

- Find the correlates
- And look for the things not postulated
- How do they relate?

© 2015 Semantic Arts, Inc.

158

## Mapping

- One of the better ways to check for coverage

Physician Data Elements						
Section/ Category	Required Fields	Subject: the subject will be an instance of (rdf:type) the class shown	Predicate	Object: the object will be an instance of (rdf:type) the class shown	Existing Taxo	Notes: when triples are noted, the subjects/object will be instances (rdf:type) the classes shown
Demographic data						
1	First Name	sheo:HealthcareProviderPerson (and subclasses)	sheo:firstName	(string)		
2	Middle Initial	sheo:HealthcareProviderPerson (and subclasses)	sheo:middleNameOrInitial	(string)		
3	Last Name	sheo:HealthcareProviderPerson (and subclasses)	sheo:lastName	(string)		
4	Designator/Title	sheo:HealthcareProviderPerson (and subclasses)	sheo:hasProfessionalDesignation	sheo:titleDesignation [taxo value]	PSDB - Title	
5	Alternate First Name	sheo:HealthcareProviderPerson (and subclasses)	sheo:altFirstName	(string)		
6	Practicing Status- (instate practicing physician)	sheo:PracticingStatus	gist:categorizedBy	sheo:PracticingStatusCategory [taxo value]	EPD - Clinician Status? Or Physician Status Type?	preceding triple: sheo:Physician gist:hasA sheo:PracticingStatus.
7						
8						
9	Date for Inactive Practicing Status	sheo:PracticingStatus	gist:actualEnd	gist:TimeInstant		A sheo:PracticingStatus is a temporal relation, therefore it has a start and date. (You may only be interested in the en date.)
10						

© 2015 Semantic Arts, Inc.

## Many attributes collapsed

- We found 10 attributes for Identity (SSN, Provider DB id, etc) are all covered by gist:identifiedBy

					sheo:getIdProofDgIsallocatedBy [URI for org or app]	
SSN	sheo:HealthcareProviderPerson (and subclasses)	gist:identifiedBy	sheo:SocialSecurityNumber		succeeding triples: sheo:SocialSecurityNumber gist:uniqueText [string] . sheo:SocialSecurityNumber gist:locatedBy [URI for org or app]	100-55-1234
Provider Database ID	sheo:HealthcareProviderPerson (and subclasses)	gist:identifiedBy	sheo:ProviderDatabaseID		succeeding triples: sheo:ProviderDatabaseID gist:uniqueText [string] . sheo:ProviderDatabaseID gist:locatedBy [URI for org or app]	666
Person ID	sheo:HealthcareProviderPerson (and subclasses)	gist:identifiedBy	sheo:PersonID		succeeding triples: sheo:PersonID gist:uniqueText [string] . sheo:PersonID gist:locatedBy [URI for org or app]	160
Medicaid ID	sheo:HealthcareProviderPerson (and subclasses)	gist:identifiedBy	sheo:MedicaidID		succeeding triples: sheo:MedicaidID gist:uniqueText [string]	5703333

© 2015 Semantic Arts, Inc.



## Mapping added a bit to the ontology

- As you'd imagine, the act of crossing all those t's and dotting all those i's lead to a few extensions to the model
- But not much, and nothing that really changed the shape of the model
- And this is what we would hope: even as we add additional data sources, internal or external, we expect them to be extensions to the existing structure

© 2015 Semantic Arts, Inc.

## Profile



#### Description (Profile Report for charities)

System IP 192.168.2.99  
System Name zmkumba-Server  
Database Type JMSQLSERVER  
Instance Name JMSQLSERVER  
Schema Name charities

#### Schema Summary

Table Name	#Records	#Columns	#Profiled Columns
CharitiesSourceContact	2	5	5
CharitiesPrintDocuments	24	5	5
Charity	2245	86	86
CharityAll	17897	3	3
CharityEvent	16186	24	25
CharityFrance	107644	11	10
CharityFranceContact	9366	10	10
CharityLegalStatus	103	5	5
CharityRef	39799	13	13
CharityRegistrationStatus	49370	2	2
CharityRegistration	146	10	10
DocumentCostRef	19	4	4
DocumentCostRefNew	23	4	4
Location	60	3	3
Location	30377	3	3
PrintProgramDocuments	204	3	3
TableContact	97	3	3
TableContactContact	127	2	2
TableContactCharity	122977	106	91
TableContactCharityAll	177116	2	2
TableContactCharityFrance	813362	10	9
TableContactCharityFranceContact	101603	26	25
TableContactCharityLegalStatus	1136	5	5
TableContactCharityRef	36479	14	14
TableContactCharityRegistrationStatus	446761	2	2
TableContactCharityRegistration	3911	10	10
tbl_charity	56075	12	12
tbl_charity	444	4	4
tbl_charity	54	4	4
tbl_charity	14	4	4
tbl_document	19	5	5
tbl_location	17460	75	74
tbl_location	226	76	76

ID	A	B	C	F	G	H	I	J	K	L	M	N
		COLUMN1		COLUMN2	Column Name	File Name	Is Binary	Match	Is Binary	Confidence	Match Type	Inferred
1	Charities	Auto Generated		DOCUMENTNAME	DocumentName	CharitiesPrintDocuments	false	false	0.42	Context based	ID	Text
2	Charities	Auto Generated		DOCUMENTDESCRIPTION	DocumentDescription	CharitiesPrintDocuments	false	true	0.17	Context based	ID	Text
3	Charities	Auto Generated		DOCUMENTTYPE	DocumentType	DocumentCostRef	false	false	0.26	Context based	ID	Text
4	Charities	Auto Generated		DOCUMENTTYPE	DocumentType	CharitiesDocumentCostRef	false	false	0.25	Context based	ID	Text
5	Charities	Auto Generated		DOCUMENTTYPE	DocumentType	DocumentCostRefNew	false	false	0.25	Context based	ID	Text

© 2015 Semantic Arts, Inc.

162

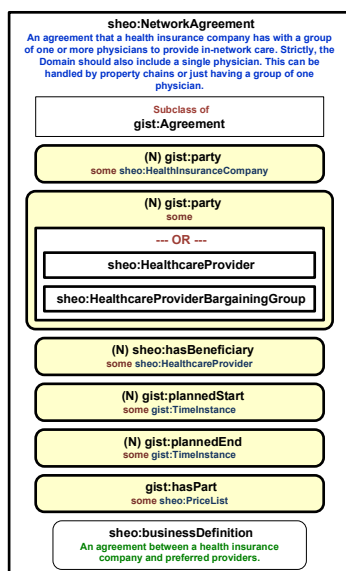
## 6) Use inference to check for errors

- Use the power of semantics to help find errors
- And hidden similarities

© 2015 Semantic Arts, Inc.

163

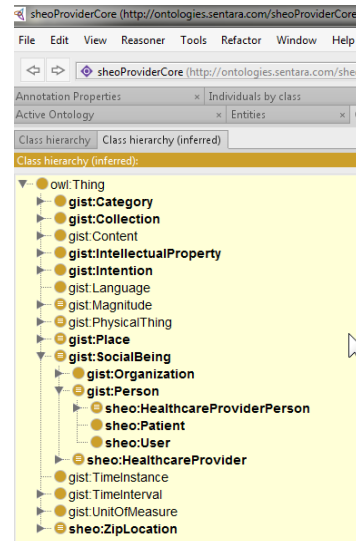
## Partial Example - Network Agreement



© 2015 Semantic Arts, Inc.

- Even though the agreement between physicians and insurance companies didn't show up in the data we were provided, we know it exists
- It is an essential bit of information, which initially we will only be aware of for InsCo agreements. But eventually we may become aware of others, and this structure allows us to hold a place for them without incurring any overhead

Check it for logical consistency using  
Protégé or Top Braid



© 2015 Semantic Arts, Inc.

## Agile -- Detecting Errors

- The "easy to change" aspect of agile requires a way to detect errors.
- Two of the more effective that we use are
  - High-Level Disjoints
  - ABox Unit Tests

© 2015 Semantic Arts, Inc.

166

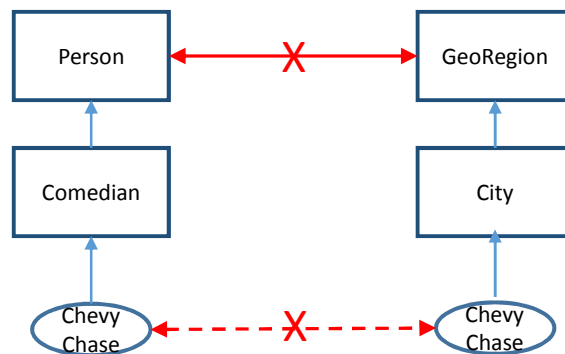
## Disjoints

- Most of the errors that a tableaux reasoned will surface stem from disjointness (or negation/complement) assertions
- No disjointness = no error checking

© 2015 Semantic Arts, Inc.

167

## High-Level Disjoints



© 2015 Semantic Arts, Inc.

168

## ABox Unit Tests

- Check for things that should not arise in the course of using the ontology, and use them for unit testing

```
ASK {  
  ?tc rdf:type sa:TimeCharge .  
  ?tc gist:actualStart ?t1 .  
  ?t1 gist:universalDateTime ?start .  
  ?tc gist:actualEnd ?t2 .  
  ?t2 gist:universalDateTime ?end .  
  FILTER(?end < ?start)  
}
```

© 2015 Semantic Arts, Inc.

169

## Case Study

- From discovery to postulate and test case study

© 2015 Semantic Arts, Inc.

170

## Questions?

- For more

<http://semanticarts.com>