

Trust, -- but Verify!

Emergence, Trust and Quality in Intelligent Systems

Vipul Kashyap
Clinical Informatics R&D, Partners Healthcare System, Inc.
vkashyap1@partners.org

1 Introduction

Trust, --but Verify!, a well known political dictum was enunciated by the late Ronald Reagan (United States President, 1980-88) in his attempts to break down barriers and integrate the communist (political, economic and nuclear) systems into a capitalist world¹. Nowhere is this dictum more appropriate than in the context of attempts to enable information integration and service composition for creation and deployment of intelligent systems. A representative domain that presents complex and interesting challenges is that of biomedical information and research, and is also a prime candidate for application of semantics-based approaches.

Biomedical and biological research has been transformed from a cottage industry, marked by scarce, expensive data generated manually, to a large-scale data-rich industry, marked by factory scale sequencing. Biomedicine is fast becoming an information-based science with data/information playing a big role across the “research flow”, e.g., Genomics → Transcriptomics → Proteomics → Metabolomics → Final Products/Results. Scientific Data Integration has been identified as one of the most daunting challenges at the interface between computer science and biology [1], and has been seen as restraining rapid progress in biomedical research [2]. Biomedical data integration, poses a unique set of challenges [1,3,14]:

- **Diversity of Information Objects, including Data Types and Queries:**
 - Sequences, complex phenotypic and disease-relevant data, graphs, 3D structures, images
 - Similarity-based queries (e.g., sequence similarity), classification based queries (e.g., Papers about Gene X), what-if hypotheses generating queries (what if Gene X was suppressed, will Protein Z exist?)
- **Diversity of information based computations and operations:**
 - Experimental Plans and Protocols involving complex repetitive computations involving data retrieval, fusion and analytics.
 - Data Curation and Annotation Tasks.
 - Hypothesis validation across multiple scenarios, Federated search/query processing.
- **Semantic Heterogeneity:**
 - Multiple Controlled Vocabularies and Ontologies: Integration, Interoperation and Composition
 - Semi-automatic creation and verification of Mappings
 - Mappings and complex relationships across vocabularies and ontologies
 - Mappings, Annotations of data objects to concepts in controlled vocabularies and ontologies
- **Dynamic and evolving nature of Biomedical Research**
 - Evolution of schemas, ontologies and vocabularies and their impact on underlying mappings or annotations.
 - Evolution of data objects and their impact on associated mapping and annotations
 - Uncertainty and inconsistency of data
 - Support for pro-active data mining and hypothesis generation

The Semantic Web effort [4] attempts to address some of these above problems by explicitly capturing the *semantics* of information and services in a machine readable format [9,10]. In the meanwhile, there has been a realization that over and above the scale of data/information explosion, *semantics* itself is a moving target that keeps evolving with time and usage. This calls into question the formal top-down and classical logic-based approaches being proposed by various researchers working on the Semantic Web. An alternative and potentially better suited approach that is

¹ The author is agnostic towards the efficacy of capitalistic and communist systems and the advisability of integrating them into a common framework.

bottom-up and statistical in nature investigates the phenomena of *emergence* in the context of information and service semantics.[8,11,12,13,14] in response to the above challenges².

The critical question that arises is *how does one guide this process?, i.e., How does one make sure that the emergent processes do not generate spurious and flawed information?* As the emergent semantics process flow (discussed later) goes through multiple iterations, it generates various artifacts (viz., taxonomies, ontologies, mappings, etc.). The *quality* of the data/information these artifacts represent can be used to guide these emergent processes. However, in the initial stages the process is likely to be in a bootstrapping mode, *we need an apriori measure of information quality*. We would like to argue here, that an *apriori trust* in the component or source of the artifact is the only way this is possible. However as more semantics begins to be created and evolved by the system, this apriori measure needs to be modulated by something that measures the internal consistency and validity of the artifacts generated, i.e., we need to verify the apriori measures, hence the dictum, *Trust – but Verify!*

2 An Emergent Semantics Framework for Biomedical Informatics

An *emergent semantics-based* information infrastructure [14], illustrated in **Figure 1**, would be a pro-active platform where people and applications collaborate for creation of dynamic *semantics* reflecting the current state of knowledge in biomedical research. Some interesting properties of this infrastructure are:

- **Self-description:** The infrastructure shall enable self-description of biological data and content, the focus of XML-based markup languages (e.g., BioPAX [5], OWL [10]) and ontologies/vocabularies (e.g., GeneOntology [6], UMLS® Semantic Network [7]).
- **Self-genesis:** The infrastructure shall proactively analyze data and content flowing through it, to create models, ontologies and concepts to capture semantics, enabling *bootstrapping* of prevalent meanings.
- **Auto-emergence:** The infrastructure shall monitor interactions between people and applications to capture and describe new meanings/knowledge that *emerge*, or existing meanings/knowledge that *evolve* from these interactions. New biomedical hypotheses may be generated pro-actively based on new emerging knowledge.
- **Self-organization:** The infrastructure shall (re-)organize itself to perform new tasks, in response to new requests for information and services; or due to emergence of new meanings. Two types of self-organization are:
 - **Self-interoperation/integration:** The infrastructure shall integrate or interoperate between existing meanings, in response to new requests for information and services. This may lead to creation of data annotations or mappings between concepts in ontologies.
 - **Self-provisioning:** The infrastructure shall, monitor the data retrieval and computation operations invoked and pro-actively create experimental plans, workflows, protocols and models.

² This proposition itself is worthy of debate, but this is not the focus of this essay.

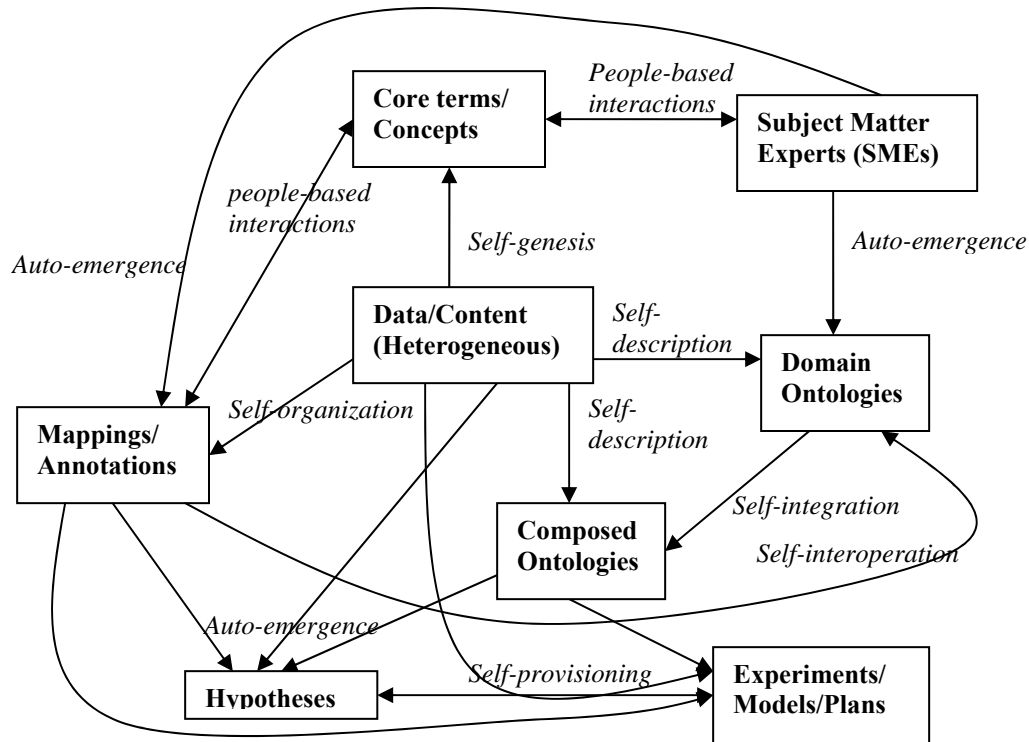


Figure 1: A Strawman Emergent Semantics Infrastructure/Process

3 Emergence, Information Quality and Trust

The collection of artifacts or components that are generated by the emergent semantics process (Figure 1) may be classified into two broad categories: “source” components present at the beginning of the process and artifacts “derived” by the process:

- Some “source” components are:
 - Data/Content Sources
 - Subject Matter Experts (SMEs)
- Some “derived” artifacts are:
 - Domain Ontologies: Single and Composed
 - Annotations and Mappings
 - Hypotheses
 - Experiments/Models/Plans

It is very important that the artifacts generated be of good quality. “Quality” may refer to both information/data quality of domain ontologies, annotations, mappings and data sources as well as quality of service related to software components that may be developed to realize various plans, experiments, etc. Quality metrics for the artifacts discussed above fall into two broad categories:

- Intra-component/artifact quality measures
 - Measures that evaluate the quality based on some internal consistency criteria
 - Measures that are assigned apriori to an artifact
- Inter-component quality measures
 - Measures that are derived from quality measures assigned to other components related to it

In the initial bootstrapping stages, where very little information is available about the various components, we have to depend on apriori measures of quality. Furthermore, trust metrics are highly correlated with these quality measures at least in the initial stages. Finally, quality metrics over time should be adjusted to reflect quality measures of external components in a consistent manner. *Trust – but Verify!* One possible approach to generate a joint measure or metric to reflect this interrelationship was discussed in [15]. We propose pragmatic and workable definitions (and not philosophical ones) for the notion of trust and other concepts it depends on. A proposed taxonomy/flow chart of definitions is illustrated in the Figure 2.

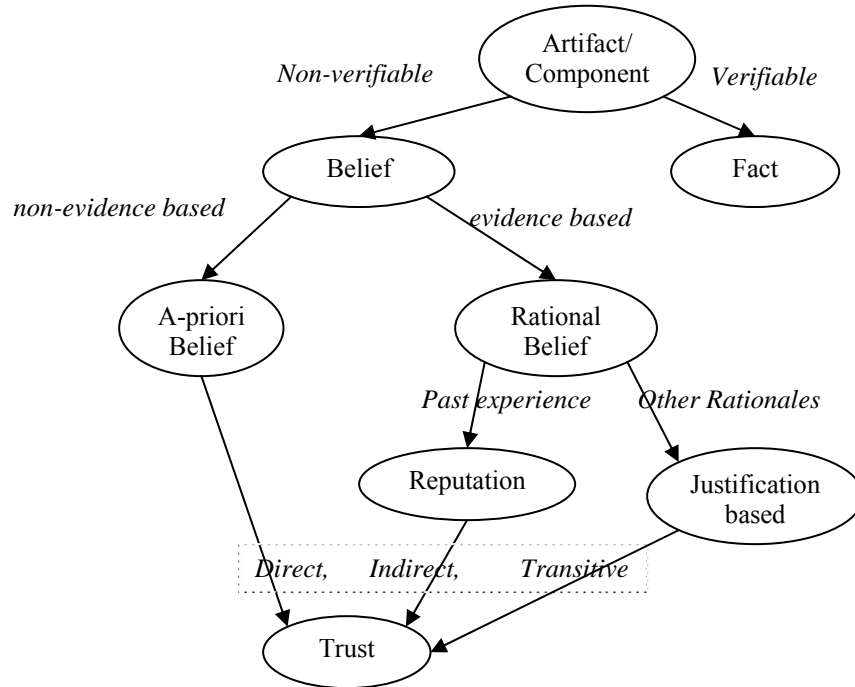


Figure 2: A Taxonomy of Fact, Belief and Trust

Verifiability of the data/information captured in an artifact or consumed/produced by the service provided by a component; forms the central axis of the taxonomy (**Figure 2**) along which various notions of trust can be organized.

- **Fact:** If data/information can be verified by some means, it is denoted as a “fact”. Verifiability might either be based on established theories (e.g., $3 + 2 = 5$), or on the basis of “trust” in an authority (for e.g., the US Government website currently states that George W. Bush is the current US President). If the information stored in the data cannot be verified by some means, then we need to have “belief” in that piece of information.
- **Belief:** When we assume the validity of some piece of data/information or the quality of a service without objective verification or recourse to a trusted authority, it is denoted as “belief”. According to Webster’s dictionary, belief is *a state or habit of mind in which trust or confidence is placed in some person or thing*, a notion intimately connected to trust.
- **A-priori Belief:** When we assume the validity of some data/information in an apriori manner, without an underlying rationale (bordering on faith), it is denoted as “a-priori belief”. This notion appears biblical, but an analysis of how intelligent systems are used or deployed may reveal that this notion is used quite often.
- **Rational Belief:** When the validity of some data/information is assumed on the basis of some *rationale* or *justification*, this is denoted as rational belief. An example is that of **reputation** of an intelligent system which is based on *past experiences* with the system. In our approach, reputation based belief is relevant in the context of subject matter experts (SMEs) who are a “component” of our approach. These past experiences themselves may be measured in an objective and empirical fashion. For instance one may use measures such as **consistency** and **completeness** or **quality of service** typically used in the data quality literature to characterize the past results received from an information system.
- **Trust:** The notion of trust is based on the continuum of belief notions that are defined above. So one may have a-priori trust in some information/data or service. This a-priori trust could be very useful in the initial bootstrapping stages and could be modulated with more objective measures, such as consistency, completeness and quality of service. Finally, trust could either be *direct* (I trust this information) or *indirect/transitive* (I trust this information because my friends trust it) [16]. In the latter case, trust is assumed to decay in proportion to the length of the “trust chain”.

The trust/quality metric may be viewed as a multi-dimensional measure. One set of dimensions could describe objective (or verifiable) measures, and another could comprise of subjective (or apriori) measures. These dimensions could be combined in two ways:

- **Multi-dimensional comparison:** The dimensions may be prioritized, say according a higher priority to the objective dimensions. Thus a comparison function may first choose compare on the objective dimensions. In the case where the objective dimensions are equal or within a statistical error range, the subjective dimensions could be compared.
- **Composition into a single measure:** The other approach is to combine the values of dimensions into a single common measure by assigning weights to the various dimensions. For instance, one could combine a subjective measure like **Apriori Belief** and an objective measure like **Reputation** as follows:

$$\text{Quality} = \alpha * \text{AprioriBelief} + \beta * \text{Reputation}$$

The weights may reflect application requirements. For instance, in the absence of any other information, we may want to set $\alpha=1$, $\beta=0$. Furthermore, **Apriori Belief** itself might have varying levels of subjectivity and **Reputation** might be characterizable in terms of information/data quality factors such as consistency and completeness; or quality of service factors.

$$\text{Quality} = \alpha_1 * \text{DirectTrust} + \alpha_2 * \text{IndirectTrust} + \beta_1 * \text{Consistency} + \beta_2 * \text{Completeness} + \beta_3 * \text{QualityOfService}$$

Of course, more experiments need to be done to refine these empirical equations and understand the true nature of interrelationship between trust and quality. Till then, we would have to *Trust – but Verify!*

References

1. Digital Biology: The Emerging Paradigm, November 6-7, 2003, NIH Natcher Conference Center, Bethesda, MD
2. NIH Roadmap: Accelerating Medical Discovery to improve Health, Bio-Informatics and Computational Biology, <http://nihroadmap.nih.gov/bioinformatics/index.asp>
3. Workshop on Data Management for Molecular and Cell Biology, February 2-3, 2003, Lister Hill Center, NLM, NIH, Bethesda, MD, <http://pueblo.lbl.gov/~olken/wdmbio/>
4. T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web: A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities," Scientific American, May 2001.
5. BioPAX: Biological Pathways Exchange, <http://www.biopax.org>
6. Gene Ontology, <http://www.geneontology.org>
7. V. Kashyap and A. Borgida, "Representing the UMLS® Semantic Network using OWL (Or "What's in a Semantic Web Link?")", Proceedings of the 2nd International Semantic Web Conference (ISWC), October 2003, Sanibel Island, Florida.
8. C. Behrens and V. Kashyap, "The "Emergent" Semantic Web: A Consensus approach for Deriving Semantic Knowledge on the Web", Real World Semantic Web Applications, Frontiers in Artificial Intelligence and Applications, Vol 92.
9. Resource Description Framework (RDF), <http://www.w3.org/RDF>
10. OWL Web Ontology Language Overview, <http://www.w3.org/TR/owl-features>
11. A. Mädche, F. Nack, S. Santini, L. Steels. **Emergent Semantics**. IEEE Intelligent Systems, Trends & Controversies, 17(1), Jan/Feb 2002, pp. 78 -86
12. K. Aberer et al, "[Emergent Semantics Systems](#)", Proceedings of the International Conference on Semantics of a Networked World: Semantics for Grid Databases, Paris, June 2004
13. K. Aberer, P. Cudré-Mauroux, M. Hauswirth, "[The Chatty Web: Emergent Semantics through Gossiping](#)", The Twelfth International World Wide Web Conference, WWW2003, Budapest, Hungary, 20-24 May 2003.
14. V. Kashyap, "Emergent Semantics: An organizing principle for Biomedical Informatics and Knowledge Management", AIS SIGSEMIS Bulletin 1(2), Semantic Web Challenges for Knowledge Management, July 2004
15. S Boll, C Cappiello, M Gertz, V Kashyap, K Sattler and C Zeigler, Do you Trust in Data Quality?, Working Group on Trust and Data Quality, Dagstuhl Workshop on Data Quality, November 2003, <http://sirius.cs.ucdavis.edu/Dagstuhl03/wgs/03362.KashyapVipul.Other.ppt>
16. M. Richardson, R. Agrawal and P. Domingos, "Trust Management for the Semantic Web", Proceedings of the 2nd International Semantic Web Conference (ISWC), October 2003, Sanibel Island, Florida.