Jul 1st, 12:00 AM

Addressing Cultural and Institutional Barriers to Data and Model Interoperability

P. Dibner

David Arctur

Dibner, P. and Arctur, David, "Addressing Cultural and Institutional Barriers to Data and Model Interoperability" (2008). *International Congress on Environmental Modelling and Software*. 138.
https://scholarsarchive.byu.edu/iemssconference/2008/all/138

Addressing Cultural and Institutional Barriers to Data and Model Interoperability

P. Dibner[a], D. Arctur[b]

[a]*OGC Interoperability Institute, Inc., Los Altos, California USA, pdibner@ogcii.org;*
[b]*OGC Interoperability Institute, Inc., Austin, Texas USA, darctur@ogcii.org*

Abstract: Despite recent technological advances, proliferation of online databases and community data collection and modeling efforts, the environmental observatory and modeling communities remain fragmented due to the lack of summaries of available observations data, differences in information models and metadata, lack of common data discovery and access protocols tuned to model requirements, and significant semantic differences in data description. However, purely technical solutions for interoperability are insufficient for establishing a shared interoperable infrastructure. The establishment of any technology in a population of potential adopters takes place within a context of social and economic processes. A community cannot become instantaneously aware of all the available data and software resources available to it, nor can consensus arise uniformly about which information models, data access mechanisms, and encodings are most appropriate. The theory of technology diffusion describes the process by which innovations are accepted by successive subgroups within a population, and suggests a series of activities to encourage adoption. The related phenomena of increasing returns and path dependence can facilitate the spread of interoperable software services and harmonized data models, but may also lead to lock-in of suboptimal solutions. Fortunately, the low cost, consensus-driven development and minimally invasive nature of interoperable web service interfaces, and the abstract foundations of well-conceived data models minimize this risk, and also further promote acceptance of these technologies.

Keywords: Interoperability; Semantics; Data sharing; Information models; Technology diffusion; Increasing returns.

1. INTRODUCTION

The holy grail of current-day integrative scientific studies is seamless interoperability across arbitrary data sets from arbitrary domains of knowledge. The purely technical aspects of this challenge have been the topic of years of effort in a broad range of studies in the domains of information science and data communications, resulting in various solutions that have been called technical or syntactic. Such solutions consist of suites or libraries of function calls, messages, or service operations that can be implemented by software components distributed throughout the Internet, and whose parameters and results are expressed using well-defined and well-known data structures.

Equal in importance to technical interoperability is semantic or data interoperability, wherein the requests and results passed among the components of a distributed software system carry content that is understood in some broad sense to convey meaning, to carry concrete values for concepts that are interpreted equivalently by all the elements of the system, and also by the system's human users. Technology for expressing such notions has also been developed and is the topic of much current research and a growing range of applications; it falls under the domain of knowledge representation, is expressed as ontologies (controlled vocabularies and, in brief, the definitions of their elements), and forms a fundamental component of the emerging semantic world-wide web.

In addition to current-day requirements for integration of information contained in widely distributed data sources is a growing awareness of the need to integrate data *processing* capabilities within and across domains of knowledge, or put another way, a need to enable interoperability among software *models*. The term "model" is overloaded even within the context of data processing and computer science. For purposes of this discussion, a model is a software program, representing some physical process or phenomenon, that requires a set of parameters as input, possibly at many points in time during its execution, and generates a result or set of results. We specifically exclude from this definition the term "information model" or, equivalently, "data model," which refer to the conceptual view and semantic content of elements within a domain of knowledge.

Technologies for structural and semantic aspects of these issues are still evolving, but have achieved a certain degree of maturity. Yet, they are only the first requirement for a truly interoperable and universally accessible solution to data sharing and integration. The ultimate goal of universal interoperability for data resources and models depends upon propagation of knowledge about the various solutions, acceptance and awareness of the advantages that it confers, agreement about which interoperability technologies are appropriate and desirable, concurrence about the representation and specific content of semantic descriptions, and generation of motivation and development resources for implementation and adoption. These are all social rather than technical processes, and subject to many factors, including limitations and outright barriers that characterize their cultural and institutional milieus.

Coming to grips with these issues in a generic, intellectually defensible, and broadly comprehensible manner requires expertise and experience both within a comprehensive sample of the domains of integrative scientific research, and also well beyond them, touching on topics that embrace economics, psychology, business administration, and a range of other esoteric domains, each potentially a life's work in its own right. We offer here a few perspectives from the standpoint of our particular experience in data modeling, environmental and biological science, and the process of standards development and advocacy from within the context of one international standards body, the Open Geospatial Consortium (OGC).

2.   PRELUDE: INTEROPERABLE SOLUTIONS AS A SINGLE ENTITY

Much of the following discussion refers to the uptake of interoperable technologies as if they were one monolithic entity, or treats one such technology as a proxy for all of them. In point of fact, even if we restrict our attention to the OGC alone, we encounter a substantial and growing number of specifications and standards, of a variety of types. The best known, and probably still the most widely represented in deployed systems, are data services, which allow any compliant client application to request and obtain information maintained by the archive and supporting software behind the service interface. A significant and growing number of data models have also flowed from OGC efforts, starting with the foundational Geography Markup Language (GML) [OGC, 2007c] and proceeding to a series of domain-specific application schemas. Most of these data models are specified in terms of their structural or syntactic encodings in the XML Schema language, but the semantic concepts underlying them are also part of the specifications and are well understood within their respective information communities. OGC specifications also include catalogues and associated methods for discovery of resources, as well as an entire set of documents that support sensing devices and human-mediated data collection programs. Specifications that can provide interoperable interfaces to computational models and workflows include the Sensor Planning Service (SPS) [OGC, 2007b], which is more a generic tasking interface than a sensor-specific interface protocol, and the Web Processing Service (WPS) [OGC, 2007a].

For all this, from the standpoint of potential adopters, there is no deep distinction between interoperable geospatial services, data and metadata encodings and meanings, and interoperable modeling paradigms. Essentially the same considerations apply to adoption of any or all of these technologies, and they amount in sum to a substantial component of the "whole product," a concept introduced by Moore [1991] and revisited below.

3.   TECHNOLOGY DIFFUSION AND THE PROCESS OF ADOPTION

Any examination of factors that delay or prevent adoption of interoperable technologies would be incomplete without consideration of patterns and processes understood to govern or describe technology adoption in the absence of any particular cultural or institutional barriers. Study of these phenomena has increased during the past several decades, and they now constitute one of the pillars of marketing strategy.

3.1   Technology Diffusion: Background

The adoption of new technology is not and can never be instantaneously universal. There is a dynamic to the spread of information about an innovation, the way it garners the attention of interested parties, and the degree to which it provokes a sense of necessity within members of the target audience.

We consider this process in the light of the theory of diffusion of innovations throughout a population of potential adopters, as originally described by Rogers (see, e.g., Rogers [2003]), and applied and popularized by Moore [1991]. In brief, the theory describes a Technology Adoption Life Cycle in which five different groups within a target population successively adopt innovative technologies. In general, an innovation must be accepted by one group before the next group starts to consider it seriously. The first group, the Innovators, is typically a small minority of the target population. Innovators are technologists, enthusiastic about exploring new technology on its own merits. The next group are the Early Adopters, still a minority, but visionaries who see new technology as a way to establish a competitive advantage, or more generally, produce a major step forward for the sake of their organization or its mission. Their purchasing decisions are based more on their imagination and vision than on established references, and they are willing to take risks in order to gain the benefits that they foresee. The Early Majority are the next group to adopt a technology, comprise a large proportion of the target population, and are key to wholesale acceptance and operational deployment of innovations. Individuals who belong to this group are driven by pragmatic concerns such as remaining abreast of proven solutions and not falling behind beneficial practices that competing organizations have embraced. They are interested only in solutions that are stable and well-supported, and they base purchasing decisions on references from sources that they trust, which necessarily depend on the accumulation of time and experience, and are not available until later in the product's lifecycle. The final two groups are the Late Majority, who adopt technology only after it has been incorporated into standard operations, and the Laggards, who resist adoption. These latter two groups are not of concern in this discussion except to note that they are the final adopters (or non-adopters), in any population.

In the case of a discontinuous innovation, i.e., one that involves a substantial change in strategy, employee behaviour, and new capabilities with new implications for an entire business, there is a particularly difficult gap that new technology must cross in order to succeed and be accepted by the entire population. This "chasm" is the divide between the Early Adopters and the Early Majority. It is hard to traverse because the Early Majority depend on relationships, proven reliability, and references, and none of these are available early in a product's lifecycle. Experience at this point lies with the Early Adopters, but their criteria are not of the sort valued and trusted by the pragmatist. The theory goes on to suggest a four-step process that maximizes the opportunity for "crossing the chasm", i.e., enabling the widespread adoption of the new technology by the Early Majority. We discuss these below in the context of interoperable geospatial technologies.

3.2   Applicability

This theory was conceived in the context of commerce and marketing of technical products by individual organizations. The case we are considering here involves the global spread of open interoperability standards for data and models, and engagement in technical development as well as operational deployment of these models by many cooperating as well as competing organizational entities. Many of the same considerations apply, but the

cases are really quite different and merit further exploration, as well as examination of the recent history of the development of standards-compliant implementation and uptake in actual deployments.

3.3   Process for "Crossing the Chasm" in the Arena of Geospatial Interoperability

We return to the four-step process prescribed by Moore [1991] in further evaluating the condition and progress of interoperable solutions in the geospatial domain, with particular reference to the process and guidelines for developing such solutions in the context of programs conducted by the Open Geospatial Consortium. Recall that the aim of this process is to engender the adoption of innovative technology by the pragmatic Early Majority, who insist upon reliable and well-referenced solutions before accepting them as a part of their organizations' infrastructure. Moore poses the response to this problem using the analogy of a military operation.

Target the Point of Attack

The first step in this process is to "target the point of attack," or in more conventional terms, to identify a market segment that truly requires and will measurably benefit from the innovation. This step is vital, particularly in the context of commercial organizations. Experience has shown that it is impractical and generally not productive to attempt to promote adoption by the entire target population. An effort of this magnitude would be diluted by the sheer number of individuals that the program was attempting to reach, and would be a fatal flaw within a group that depends heavily on well-established relationships and references. A coordinated effort in one market segment avoids this pitfall.

In the geospatial domain, the collective pain of what might be considered a single market segment engendered the creation of the precursors to the OGC and ultimately the Consortium itself. These were the government agencies, including the Earth Observation agencies, who were stewards of increasingly large volumes of substantial imagery datasets, and who needed to share their data with other government, academic, or commercial organizations. Replicating such large datasets was often prohibitively expensive, and converting formats to allow integration of these with other sources of information was not only costly, but a significant engineering enterprise in its own right, with its own technical challenges, requirements for specialized development, and other sources of delay. Even as computing power and storage capacity increased, so did the capability and practice of collecting satellite and aerial imagery, as well as the types of imagery and other specialized data with spatial and temporal attributes. The OGC started as an early coalition of Innovators and Early Adopters, who sought and ultimately developed workable solutions to the interoperability problem. The need for these innovations was such that the early efforts at development and experimentation were also blessed with an audience from within the Early Majority among the agency population, awaiting a set of technologies that were sufficiently stable to deploy as operational systems.

The success in this market segment, such as it was, provided the basic essentials for propagation of interoperable geoprocessing capabilities to additional agencies and institutions: a well-defined technical solution, a suite of increasingly reliable commercial products, and a growing set of operational deployments, already tested in one segment of the population of potential adopters. This has become a beachhead for further propagation into the Early Majority within other organizations, including city and county governments and the community of professionals involved with environmental modeling and analysis.

Assemble the Invasion Force

This second step of the Moore paradigm entails making sure that the product, when presented to the customer, provides a complete solution, or what Moore calls a "whole product." In practical terms, this means not only that the product is functional and reliable

as an isolated entity, but that other services and capabilities essential to support it are also ready and available to potential adopters.

In the context of interoperable data services, this could include documentation and training, a community of developers and consultants, and an entire suite of products that can be compared and assessed on the road to implementation as operational deployments.

In the context of interoperable data models, it means not only that the model is supported by documentation and a community of implementers, but that the paradigms supported by the data model are truly representative of the target population of adopters. In order for an information model to be useful in the context of interoperable web services, there must also be ancillary technology such as client software that understands the model sufficiently to render or analyze the data, format conversion tools for dataset translation or on-the-fly conversion so that legacy data available via interoperable services, catalog services to enable discovery.

In the context of OGC technology, this step is far from complete, and its maturity varies considerably across the spectrum of services and information models that comprise interoperable geospatial technology. However, the "invasion force" is sufficiently well assembled to gain acceptance and credibility in some domains, representing the leading edge of the "assault." In particular, the OGC data services, including the Web Map Service (WMS) [OGC, 2006], Web Coverage Service (WCS) [OGC, 2007d], and Web Feature Service (WFS) [OGC, 2005] have been incorporated into numerous commercial and open-source products and enjoy increasingly broad use in the geospatial analysis community. Well-defined and broadly accepted information models compatible with the WFS are also emerging in increasing numbers, but still await broad deployment. Approaches to interoperability of simulations and computational models are still undergoing experiment and development. However, their ultimate adoption will be facilitated by the technologies that are already enjoying success in the marketplace, and upon which their technical underpinnings are substantially based.

Define the Battle

This step refers to the need to introduce a product in the context of other competing solutions. If the goal is interoperability of geospatial data and services, then the battle is already defined. The "competition," such as it is, consists of datasets and models that do not interoperate. The costs of non-interoperability are typically well articulated by outreach organizations and individuals who promote interoperable solutions in their own organizations. Thus the real challenge is to elucidate these factors in the context of actual communities that are just starting to encounter the need for interoperable solutions to their data sharing and service operations.

Launch the Invasion

The final step in Moore's paradigm is to initiate the actual marketing and sales campaign, with channels to the customer appropriately selected to reach and match the requirements of the decision makers. The invasion in this case has been in process for several years, and it is ongoing. The initial success of the WMS, WFS, and WCS has already been mentioned. These are in effect the leading edge of the wave of web services in a growing and maturing architecture.

4.   INCREASING RETURNS AND PATH DEPENDENCE

Technology markets are different from traditional commodity- or resource-based economic models like those formulated in the nineteenth century. Such conventional models are based on the phenomenon of diminishing returns, i.e., that each increment of investment is accompanied by a reduction in the value gained or profits realized as return on the incremental investment. It is particularly easy to see how this model applies when the amount of a desired resource is limited: as the resource becomes scarcer or more difficult to obtain, its cost increases and thus reduces the profitability of products or processes based upon it. Thus the dynamic is characterized by negative feedback, a stabilizing force. In the

case of competing enterprises based upon different resources, a random increase in the market share of one enterprise increases the costs associated with it, and its competitors become incrementally more attractive by comparison. The classical result is a stable equilibrium in which a combination of the solutions persists in the market, in proportions dependent solely on the marginal costs and returns of the various options, and independent of the history or timing of random fluctuations in market dominance by any of the competitors.

The situation is very different if returns increase rather than decrease on marginal investment. Far from acting as a stabilizing force, economic activity generates positive feedback. In the case of competing enterprises, a slight advantage gained by one is amplified, giving it a greater advantage and increased market share. The end result is a situation where the entire market is dominated by one of the competitors. Moreover, the selection of which one is very sensitive to the history of fluctuations in market share or other advantages, and not determined exclusively by the cost-value propositions offered by the competing enterprises. Therefore an increasing-returns market may lead to broad adoption of suboptimal products that are difficult to displace.

The world of technology is frequently characterized by an increasing-returns dynamic. This is particularly true of the software industry, especially in the domain of application programming interfaces (APIs). Software developed to use a particular set of APIs promotes the acquisition of systems that support those APIs, which therefore assume a greater market share and so become an increasingly desirable platform upon which to base new software. Suites of compatible protocols that enable interoperability among diverse services, data sources, and computational installations are of course prime candidates for triggering this behavior in the marketplace.

To the extent that OGC standards are appearing in operational deployments, this would seem to be a benefit to the goal of promoting interoperability, but it is important to consider and somehow address the danger of locking in suboptimal or even poor solutions. We consider this issue below.

5. DISCUSSION AND ANALYSIS

It is almost legend that some corners of the research community are concerned that adherence to standards may have negative impacts on creativity, or may not satisfy requirements in some way essential to scientific progress. But in our experience, this legend ultimately holds little substance. Every community with which we have engaged includes a contingent of informatics experts who are well versed in the characterization of data in their own domain, who are also seeking, if not already knowledgeable about, the proper approach to integrate with and make use of standard services and data paradigms. However, scientific communities too have their populations of Early Adopters and Early Majority, and as in the commercial domain, the latter are more concerned with maintaining a reliable and productive research facility than with investing resources in promising but still risky propositions for enhancing productivity.

Even when the relevant interoperability technologies have achieved a level of stability and generality that would be acceptable to the Early Majority population segment, barriers remain. They may include extensive investment and training in existing or historical practices, legacy datasets, priorities for proximal goals that have little requirement for or benefit from interoperability, and others, all in combination with limited and previously dedicated resources.

OGC standards were designed be compatible with existing installations and on existing operations, and we posit that this would be a feature of any suite of interoperable service protocols or data models designed for widespread adoption. Many of the requirements that lead to wide acceptance are already built into OGC standards because they were developed by representatives from many domains, whether technical, scientific, or commercial, and intended to be accessible to very broad implementation. It was a given from the outset that they must work with all major software and hardware platforms, they must be neutral with respect to data storage formats, management and custodial practices, and they must also be

reasonably straightforward to install and operate.

The protocols themselves therefore tend to be minimal, and to make few assumptions about the installations that they support. They can in general be implemented as façades for proprietary or more complex systems for data retrieval or analysis, and they need not be exclusive: for a given service they can act as one among many interfaces to the worldwide web.

The same is true of the data models promoted by OGC technology. They are based upon the unadorned, abstract ISO Feature model as specified in the ISO 19101 [2003] and ISO 19109 [2006] specifications, and are not inherently dependent either on GML encoding or OGC service protocols. The utility of data models harmonized to one another and compatible with OGC technology extends beyond OGC-compliant installations.

If indeed these characteristics are essential for any suite of technologies that enable very broad-based deployment of interoperable services, then not only do they present a low barrier for implementation by technologists and acceptance by pragmatists within their respective domains, but they also serve as a counterweight to the lock-in phenomenon that can characterize increasing-returns markets. Service installations should not represent a huge investment, and should not be prohibitively expensive to replace if need be.

6.   CONCLUSIONS AND RECOMMENDATIONS

The primary conclusion of this assessment is that data and model interoperability appear poised to undergo an adoption process similar to that of the simpler geospatial data services and products in other highly technical domains. The picture is a bit confused and substantially blurred because of the great diversity of potential adopter populations and the urgency with which they need to integrate their data assets and modeling capabilities with each other and with those of other communities. Additional factors include different rates of uptake for the broad array of service and data models that have been defined, and differences between the populations of end users, service implementers, and modelers.

An increasing returns market and associated path dependency seem to favor development of interoperable service, data, and software modeling in a consistent direction, but raise the specter of inflexibility and broad acceptance of suboptimal solutions. Service definitions that impose minimal restrictions upon existing or new implementations, and carefully conceived and broadly vetted data models both mitigate these potential negatives and facilitate acceptance within communities wary of impacts upon their operations or resources.

The recommendations that flow from these observations are aimed at promoting the further spread of interoperable data services and computational models, and they fall into two main categories.

First, engage in activities that lower the barriers to technology diffusion. In particular, focus on issues that are appropriate to the current stage of development of the standard of concern. For technologies that are still experimental, seek involvement and feedback from the Innovators and Early Adopters in target communities of interest, including your own organization. As the technologies mature, leverage and participate in efforts to refine the product offerings, whether through implementation, development of tools or documentation, or participation in beta tests for commercial or open-source developers.

In terms of communication with hesitant parties, it is important both to maintain the mindset and communicate the truth of the fact that supporting effective interoperability standards is an opportunity, not a restriction. Data standards must arise within the community that understands and uses the underlying models.

The second general recommendation is to leverage the phenomenon of the increasing returns market.

First and foremost, develop, adapt, or incorporate existing standards-compliant products in your operations. To the extent that experimental or pilot implementations are available, use them, and do not overlook the substantial and growing framework of geographically extensive and increasingly high resolution datasets with spatial and temporal content, and the network of OGC services that provide access to them. Early in the history of the OGC Web Services (OWS) initiatives, there were only a few service implementations that supported broad datasets such as the USGS transportation networks, hydrological features, political boundaries and similar assets, and satellite imagery were available, but they added immense value to experiments and demonstrations. Today, many more standards-compliant services are available, and access to them is made all the easier due to the wide availability of software tools, both commercial and open-source, that support the same standards.

A small or partial solution is in general preferable to no solution. If you are able to publish a portion of your data using an interoperable service, or are able to access only one of many interoperable web-based resources, do so.

For data modeling efforts, engage with the broadest possible segment of the relevant community. It is not necessary for a community to produce only one data model. Sometimes there are different, even incompatible models produced by different segments of the same community. But usually, even though the models cannot be mapped successfully to one another, there is an overarching, often simpler construct to which the mutually incompatible models can be successfully mapped. Regardless, consider implementing a data service that supports the model you prefer. At a minimum, you will be making it accessible to the market of potential users, which promotes experimentation and allows broader assessment.

REFERENCES

Arthur, W. Brian, *Increasing Returns and Path Dependence in the Economy*, U. Michigan Press, 201 pp., Ann Arbor, 1994.

ISO 19101:2003, *Geographic Information – Reference Model*, 2003.

ISO 19109:2006, *Geographic Information – Rules for application schemas*, 2006.

Moore, Geoffrey A., *Crossing the Chasm*, Harper Business, 227 pp., New York, 1991.

OGC, *OpenGIS® Web Feature Service (WFS) Implementation Standard*, OGC Document 04-094, 2005.

OGC, *OpenGIS® Web Map Server (WMS) Implementation Specification*, OGC Document 06-042, 2006.

OGC, *OpenGIS® Web Processing Service (WPS) Implementation Specification*, OGC Document 05-007r7, 2007a.

OGC, *OpenGIS® Sensor Planning Service (SPS) Implementation Specification*, OGC Document 07-014r3, 2007b.

OGC, *OpenGIS® Geography Markup Language (GML) Encoding Standard*, OGC Document 07-036, 2007c.

OGC, *OpenGIS® Web Coverage Service (WCS) Implementation Standard*, OGC Document 07-067r5, 2007d.

Rogers, Everett M., *Diffusion of Innovations*, 5th ed., Free Press, 518 pp., New York, 2003.