

Expert Committee Meeting

Saturday, May 29, 2004
Madison, Wisconsin

Present:

Hans Jorgen Marker (Danish Data Services), Vice Chair; Atle Alvheim (Norwegian Social Science Data Services [NSD]); Pat Doyle (U.S. Census Bureau, Demographic Surveys Division) - by telephone; Ilona Einowski (University of California, Berkeley, UCDATA Archive); Arofan Gregory (AEON Consulting); Reto Hadorn (Swiss Data Archive, SIDOS); Carol Haney (SPSS Inc.); Pascal Heus (World Bank); James Jacobs (University of California, San Diego); Ryan Johnson (Washington State University); Mari Kleemola (Finnish Data Services); Julie Linden (Yale University, Social Science Libraries & Information Services); Marc G. Maynard (University of Connecticut, Roper Center); Meinhard Moschner (Zentralarchiv fuer Empirische Sozialforschung-ZA); Ron Nakao (Stanford University); Rob O'Reilly (Emory University); Jostein Ryssevik (Nesstar Ltd.); Janet M. Eisenhauer Smith (University of Wisconsin, Madison); Ken Miller (UK Data Archive); Wendy L. Thomas (University of Minnesota, Minnesota Population Center); Mary Vardigan (ICPSR); Joachim Wackerow (Zentrum fuer Umfragen, Methoden und Analysen-ZUMA).

Also attending as observers:

Bill Block (University of Minnesota); Dan Gillman (U.S. Bureau of Labor Statistics); Ann Green, Steering Committee (Yale University); Sanda Ionescu (ICPSR); I-Lin Kuo (ICPSR); Walter Piovesan (Simon Fraser University); Richard Rockwell, Steering Committee (Roper Center); Marion Wittenburg (NIWI/Steinmetz Archive).

Annual Report

After introductions and welcoming comments from Hans Jorgen Marker, Vice Chair of the Committee, and Ann Green, IASSIST President, the Committee briefly discussed the FY04 Annual Report for the Alliance. The DDI Alliance formally began operations on July 1, 2003. While the first fiscal year is not yet complete and thus there is no final budget report, the Alliance will finish out the year with a positive balance that can be used for meetings of working groups and other activities in the upcoming year. There were 25 members in FY04, with the prospect of some new members in FY05.

Upcoming Meetings

In discussing a possible fall 2004 meeting of the Alliance, the possibility of having separate European and American meetings was raised, but this was deemed too difficult to execute. It was decided that the Alliance should attempt to support fall meetings of the Working Groups whose activities are most central at the time. There has already been an informal request from the Structural Reform Working Group to meet in the fall. The Comparative Data group may also be meeting in conjunction with the MetaDater initiative meeting in Cologne in the fall.

The date for the next meeting of the full Expert Committee was scheduled for Sunday, May 22, 2005, in Edinburgh, in advance of the IASSIST conference to be held that week.

Communication Mechanisms

It was generally agreed that ezboard was not working as a general communications vehicle for the Alliance and that we should not invest in the bulletin board technology any longer. Most members prefer to communicate via the email lists that have been established for the Working Groups and for the Expert Committee as a whole. These lists have archives that can be consulted for message threads. It was also agreed that we should continue using the Expert Committee Web page of the DDI Alliance site for file sharing. This site can be password-protected if files of a sensitive nature are to be added. Telephone conference calls have been working fairly well, although this mode of communication can be difficult for participants who are not native English speakers.

Reports of Working Groups

Complex Files Working Group

The Complex Files Working Group had prepared a proposal, which was brought to the Expert Committee for discussion. Pat Doyle described the proposal in detail and walked the Committee through the examples. This proposal is intended to document a system of files as opposed to one file at a time. The proposal can document, for instance, a time series or a structure of person-level and household-level files used in tandem. Another example would be a three-wave longitudinal study in which data are collected at various points in time and then accumulate into a group of related files. The proposal is purely to document these structures independent of applications.

The Complex Files proposal recommends the addition of a new section of the DDI (Section 6) called File Group (similar in concept to Variable Group and Category Group). File Group would be repeatable and would provide the functionality to:

- Identify the set of files that make up the system
- Indicate how they can be used together through a linkage or a join (optional)
- Identify the files generated as a result of the join

A new Section 6 was developed rather than using Section 3 (File Description) in order to provide for an added DDI layer that points to other files. The proposal is mechanistically based and would fit into the modular structure for the new Version 3.0 that is being envisioned. We do, however, need to clean up Section 3 as we move forward.

A comment was made that this is really a revolutionary proposal in that it creates a structure that sits above other DDI instances. We need to take relevant examples and see if they work with this proposal. This is basically a relational system that is being described, and it needs to be generic enough to work in any situation. It is possible that we may need some semantic relationship controlled vocabularies.

We need to decide if the proposal covers comparative data and all types of longitudinal files also. In general, we need to make the distinction that the Complex Files proposal is intended to address the matching of cases, while the Comparative Data area refers instead to the matching and harmonization of variables. Both levels need to be coordinated in the DDI.

Julie Linden volunteered to test the Complex Files proposal with aggregate/tabular data, and Jostein Ryssevik will test it with Nesstar. Pat will act as Architect of the proposal and in that capacity will work with the SRG to populate a spreadsheet detailing the relationships inherent in the proposal elements and attributes. [Note: Pat Doyle died shortly after the meeting. Janet Eisenhauer Smith subsequently agreed to act as Architect of the proposal. -MBV]

Comparative Data/Families of Datasets Working Group

Meinhard Moschner provided a summary of the work of the nine-member group to date. The group distributed an initial brainstorming paper outlining the scope of their substantive concern, which is linking elements across studies or instances and over space and time. The group needs to look not only at the projects designed as comparative studies but also at potential families of studies that may not exist physically but only logically.

Indicating comparability is not easy because methods or measures are not always equivalent. Deviations are unavoidable and also necessary sometimes. We need to document the deviations at the study and variable levels, comment on the reasons for the deviation and on harmonization procedures, and provide for potential trend analysis. Harmonization can sometimes lead to the loss of information -- for example, if one has to collapse categories - and this needs to be documented. In the case of potential trends, we need to be able to link variables and questions across studies and to provide enough information to show that harmonization is possible. This may require a new DDI level - a collection or family level or something like a collection variable group to describe loose trends. The group will prepare a common data model for discussion.

The MetaDater Project, also concerned with comparative data, is in the process of preparing a metadata model, which will be designed to be compatible with the DDI model. There is now an internal vision of the

MetaDater model, and the project is aiming to have a more complete model for an expert workshop in the fall.

It was pointed out that W3C standards to integrate references (XInclude or XLink) may be useful in the data model to describe comparative data. We also need some means of formal description to construct new variables for purposes of harmonization. This could perhaps be a subset of MathML.

ISO-11179 establishes comparability even if there is no actual physical data collection. It lifts variables up to a higher level of abstraction and links to concepts. If we atomize variables into their component parts, as in the ISO-11179 model, do we lose the study context? Keeping links to studies that gave rise to the variables is vitally important. Also, comparability is hard to establish. The DDI needs to keep its information simple and descriptive and provide the information on which researchers can base comparability decisions.

We need to make sure that the DDI enables what comparative researchers really want to do and is useful for them. We might involve researchers from the Comparative Study of Electoral Systems or the Luxembourg Income Study to ensure that we are meeting the needs of the community. Comparative research is affected by methodology, sample design, and many other factors.

Focusing on comparative data issues allows us to compare potentially comparative variables after which an application can capture these new variable groups or relationships into a knowledge base. However, in the process of harmonization, we need to ensure a strict division of labor between the DDI and applications.

It is possible that we should treat harmonization and comparability separately. Should we take the stance that the elements we provide in the DDI are what the researcher needs to know, or should we come up with a measure or index of comparability?

In the Madiera project, the goal is for researchers to make the actual decisions about comparability but to identify the factors that influence their comparability rankings. It's also important for the researcher to feed back into a system to say that he or she performed a certain harmonization.

We need to have tags to indicate "these measures were designed to be comparable". The results of harmonization itself are in effect a new dataset.

We also need to know that we are measuring the same concept, which is where ISO 11179 again becomes relevant. The DDI already has a concept element, which can point to a vocabulary outside of the DDI; this could be an ISO-11179 repository.

The DDI may need more controlled vocabularies, but the Committee was cautioned not to embed controlled vocabularies into the specification and to keep the DDI XML independent of vocabularies.

Structural Reform Working Group

Wendy Thomas reported for the SRG, which works in parallel with the other working groups and is tasked with maintaining consistency in design across the proposals of the substantive content groups. To this end, the group created a diagram of the data life cycle with a modular structure for review by the Committee to ensure that there was agreement on what the DDI is designed to document. It was noted that there is a chapter of a MetaNet report on the life cycle of statistical data that we should also consult. The SRG also did a mapping of the current DDI tags to the life-cycle model.

This [life-cycle model](#) helps to determine what is in and out of scope for the DDI. The diagram starts at the study design stage and continues on to the archiving of a dataset and beyond, with DDI embedded in the process throughout.

The modules of the life cycle are:

- Study/Survey Design - With Concept sitting in a tier above
- Data Collection - With Data Collection Process above
- Data Processing - With Physical Encoding and Logical Encoding above
- Data Dissemination - With Archiving above
- Data Discovery
- Data Analysis

Running left to right through the life cycle are actions that are part of Data Use: Study Discovery, Detailed Discovery, and Data Access.

The current DDI specification is comprehensive enough for a single survey, but in general it represents the tail end of the life cycle.

The sense of the Committee was that the new life-cycle model was appropriate for the DDI. It was noted also that having a model that spans the life cycle of statistical information fits into the new vision of SPSS. SPSS created SPSS Dimensions, which was geared toward data collection and was principally a tool for market research. But now the goal is to have a robust and sophisticated suite of tools that span the data life cycle and that handle large datasets.

It was pointed out that the word "survey" is ambiguous and that the DDI could encapsulate a number of instruments. A question was raised regarding what exactly the digital object being described was - a dataset or a study or some other entity. We need to be clear about our definitions of these terms.

Related to this, there is a preservation metadata model that maps to the OAIS model. Could another model, such as METS, be used to "wrap" the DDI?

In October through December, the SRG will be working on the data model for Version 3.0 with the goal of having people comment by the end of January 2005. This data model will either be in the form of a spreadsheet or a UML model.

A question was raised about whether the DDI is intended to document both the conceptual and the physical. The DDI started with a physical object - a social science codebook -- that was documented and would be preserved. However, we have now separated the physical and the logical, which moves us away from the traditional codebook structure where we started. We need to be able to preserve the conceptual structure archivally without worrying about the physical form.

What we have with this life-cycle model is a set of modules and the DDI instance is a way of combining the modules from different places. We need to also think about versioning across the life cycle.

Aggregate Data, Time, and Geography Working Group

Ilona Einowski reported for this group, which is in the process of obtaining background information to move forward. Wendy Thomas has sent information on the current aggregate model and will also send information on what the NHGIS project has learned in using the nCubes model for that project.

We need to think about how aggregate and tabular data are different. The cube specification needs to be improved in Version 3.0. We want to be able to say that a dimension in one cube is the same as in another cube, and currently there is no way to do that. Right now we have to artificially locate them under the same study, but they could come from different sources.

In terms of time and geography, the group is working on identifying problems with the current specification. The Madiera project completed a review of the existing geographic elements. Atle Alveim will determine whether the final report can be circulated.

We need suggestions for Version 3.0 in terms of geography, and we need to determine how the DDI relates and should relate to other geographic standards. Ilona will fill in as committee chair for Margaret Low while she is on leave. Julie Linden will look at the mapping from DDI to FGDC.

Instrument Documentation Working Group

It is still an open question how much emphasis the DDI should place on documenting survey instruments. Should this be separate or a part of the DDI? We need to compile a list of potential tags related to instrument documentation that are not currently in the DDI. This can be partially facilitated through ICPSR's collaboration with the Survey Research Operations group at the University of Michigan's Institute for Social Research. This group has created a Blaise documentation program that produces an XML codebook, and ICPSR is currently mapping the XML tags to the DDI. ICPSR will also compile a list of what is missing from the DDI in its current form.

Usability and Outreach Working Group

The ICPSR Web site (www.ddialliance.org) provides a lot of information on how to use the DDI in specific situations, given various sources of information. We need to continue to solicit information on what others are doing and describe the different projects on the site. We also need to find out what people who are not using the DDI need to know in order to understand the value of using it. Providing good examples is extremely important. We could also use outreach materials directed to different audiences.

A useful resource would be to show how to create a Dublin Core record using the DDI. We are currently working on a DDI to MARC conversion.

Outreach to grad students is a potentially promising area now that data integrity issues are so prominent. We need to convince students to document their data to protect themselves.

The view was expressed that the DDI is basically selling itself now and is in fact an easy sell. It is expanding outside the original committee. Transport for London, for example, is going to use it. We need to reach funding agencies as well. We could try to stipulate that projects need to use the DDI to get their funding.

Potential New Working Groups

Other Working Groups have been proposed to address issues of:

- Qualitative Data
- Longitudinal Data
- Historical Data
- Language

We will investigate further and determine the need for separate groups.

Persistent URLs

Joachim Wackerow followed up on his email discussion regarding this issue and suggested that the Alliance may want to make a recommendation on how to identify codebooks in a unique and persistent way, perhaps using URNs, which could be mapped to URLs through file resolvers. However, there is currently no widespread accepted resolution system.

Structured URNs are becoming more common. We should look at OASIS. A question was raised about whether the persistent identifiers should be at the study or the codebook level and whether there should be a central registry. For uniform naming conventions we don't need a registry. We should have a publicly available naming scheme to which we all adhere.

We still have the problem of duplicate holdings. If the Internet domain is the first part of a structured URN, this helps to solve the problem. The SRG considers this issue in its purview and will look into this further. We need to include all archives in such a discussion.

Open Access Protocol

Joachim suggested that the DDI might benefit from having a central interface and repository for search and retrieval of DDI files and a standardized transmission protocol to exchange files.

The industry standard at this time appears to be Simple Open Access Protocol, or SOAP, which is being developed by the W3C as part of Web Services. However, the archives often align themselves with the library community, which uses the Open Archives Initiative Protocol for Metadata Harvesting, or OAI-PMH. SOAP is not part of the OAI specification. If we are interested in developing a registry, we might look at SDMX, which is already going that route. Developing and maintaining a registry has high overhead, so piggybacking on an existing structure would ease that burden.

The UK Data Archive is looking at becoming either OAI or Z39.50 compliant to meet the requirements of its funders. A drawback of OAI is that metadata harvesters can present the results as their own.

We need to investigate the SOAP and OAI protocols further before making a decision.

Procedures Manual

The Manual prepared by the SRG clarifies the process for changing the specification that is outlined in the DDI Alliance Bylaws and distinguishes between proposals for major and minor changes. It sets out a process involving a spreadsheet, which details the relationships between elements and attributes in a proposal; this makes things easier for the Working Groups, who do not then have to write XML. A Working Group can either have a member fill out the spreadsheet or can work with the SRG to build the document.

We are using the Complex Files proposal as a prototype and will be following it through the processes stipulated in the Manual. The proposal currently does not have a corresponding spreadsheet, which would complete Part 1 of the process, but Pat Doyle as the Architect of the proposal will work with the SRG to develop one.

For substantive content groups with overlapping interests, there should be broad discussion of developing proposal before the proposals are formalized and before the Expert Committee has to vote. This will ensure that we don't work at cross-purposes or take radically different strategies. Working Groups should feel free to start to work with the SRG as early in the process as possible. This can be done informally.

Timeline

The main point about the Timeline, which is now published on the DDI site, is that the final date for proposals to be submitted to become part of Version 3.0, which is planned for January 1, 2006, is March 1, 2005. This deadline is intended to provide adequate time for a proposal to make its way through the specified channels.