Submitted to INOSA 2012, http://inosa.icst.org/2012/show/home

Boundary Objects for Online Knowledge Management

Dino Karabeg[1] and Benjamin Johansson[1]

[1] Institute for Informatics, University of Oslo,
POB 1080 Blindern
0316 Oslo, Norway
{dino, benjamij}@ifi.uio.no

Abstract. Knowledge Federation complements the work in technical communities such as Semantic Web and Global Sensemaking by developing and putting into real-world use those desirable knowledge-work patterns and practices that are in principle enabled by technology. The Domain Map Object and the Value Matrix Object, which are described here, facilitate this task by serving as boundary objects: On the system development side, they provide building blocks in terms of which desirable functionality can be implemented; on the technology and the theory side, they provide a stable set of functions that need to be implemented, roughly analogous to sorting and matching in conventional computing. As NoSQL databases, those two objects respectively serve as repositories for information about a domain, and for information about the value of the resources in a domain.

Keywords: social engineering, information overload, global issues, knowledge federation, systemic innovation, socio-technical system design.

> *"As long as the centuries continue to unfold, the number of books will grow continually, and one can predict that a time will come when it will be almost as difficult to learn anything from books as from the direct study of the whole universe. It will be almost as convenient to search for some bit of truth concealed in nature as it will be to find it hidden away in an immense multitude of bound volumes."*
>
> Denis Diderot

1 Introduction

Diderot already predicted that information overload would grow to paradoxical proportions [1]; the persistence of 'global issues' also points that we have difficulty putting our knowledge and minds together. The question that motivates this work is "*What should we do* to reverse this trend?"

We readily notice that much of our challenge is in the way in which this question is answered: If we shall only diagnose the problem and tell what is to be done, our

solution may well fall prey to the very problem it is intended to solve—vanish in the jungle of information and never be implemented. We therefore *implement* what we consider to be key elements of the solution, albeit still as small research prototypes; and we invite the participants of INOSA 2012 to join us in developing them further.

To see what in principle needs to be done, we submit, is easy: Technology (Semantic Web, Topic Maps, DebateGraph...) has been developed that can revolutionize knowledge work; but this technology tends to be simply added to practices that have evolved based on entirely different technology [2], and in a society with different problems and needs. The critical task now is not just adoption of technology, but a complete redesign of the conventional practices (in public informing, science, education, governance...) to include coordination of activities, organization of resources and creation of socially relevant insights.

We consider the following to belong to the critical path to solution.

1.1 Design Epistemology

By *epistemology* we mean the assumptions on which knowledge work is based. The conventional *epistemology*—that the purpose of knowledge work is to provide an objective reality picture— has been challenged [3].

We therefore submit it as good practice to state one's *epistemological* position up front, as one would declare variables at the beginning of a computer program, or define *x* and *y* at the beginning of a theorem. And we declare that our position is *design epistemology* [4].

Design epistemology considers knowledge and knowledge work to be core components of various social systems, and stipulates that the core task is to adapt them to those roles, so that everything else in the society may fulfill their roles. The *design epistemology* changes our self-perception in knowledge work from 'objective observers' to 'designers.' A colloquial explanation is that *design epistemology* fosters the kind of attitude one manifests when stopping the car one is driving to change the wheel that has a flat tire.

A characteristic technique—roughly analogous to the experiment in conventional science—is a *prototype* (a model implemented and tested in practice).

Under *design epistemology,* the work on recreating knowledge-work practices to better serve their various systemic functions becomes a new notion of 'basic research.'

1.1 Knowledge Federation

Who will recreate knowledge-work practices? And in what way? Knowledge Federation community has self-organized to provide answers to these questions.

The Knowledge Federation community complements the work in various technical IT communities, by developing and putting into real-world practice the desirable or remedial ways of working and systemic patterns, which are in principle enabled by new technology. Knowledge Federation distinguishes itself from conventional socio-technical design communities and projects (which would for ex. develop a system for sharing information among hospitals, industry and patients) by focusing on systemic 'basic research' in knowledge work, and developing completely new ways in which its key components, including public informing, research, education and governance, might function.

The Knowledge Federation community originated at TMRA07 conference, where several of us noticed that we were already working on this new agenda. Our first workshop, at Inter-University Centre Dubrovnik in 2008, brought together a small group of knowledge media researchers to begin a community and map the relevant tools and approaches [5]. Subsequently, having identified that its core task is to 'develop self-organization in knowledge work,' through a series of workshops, Knowledge Federation evolved to become a heterogeneous community, or 'a federation of knowledge workers and other stakeholders' representing a suitable combination of backgrounds, and using itself as sandbox to co-create and test systemic solutions [6]. At the recent workshop in Barcelona journalists, knowledge media workers and other stakeholders came together to 'co-create an innovation ecosystem for good journalism' (the workshop title) [7].

By enabling systemic innovation, Knowledge Federation enables also the following *natural* way of developing information technology: Instead of marketing new technology to power the old ways of working, which breeds overload, we first develop new patterns in knowledge work and in knowledge-work organization that we wish to support, and *then* produce the technology needed to implement those patterns.

A core task of Knowledge Federation is to change the knowledge-work practices. A *prototype* strategy, currently under implementation, is called The Game-Changing Game.

1.2 Boundary Objects

The two boundary objects that are reported on here are *prototypes* fulfilling different functions in four distinct contexts.

In the context of socio-technical system design, they serve as building blocks, providing exactly the kind of functionality that is needed for creating and supporting remedial knowledge-work practices. The specific designs are drawn from the experiences in Knowledge Federation so far—by identifying that a certain set of functions is needed in pretty much all on-going socio-technical design projects.

In the context of an application, they serve as affordances, which (1) make good practice such as organizing the resources in a domain, natural and easy or (2) enable the creation of a knowledge-work ecology (reputation management, reward system) that supports good practice. As objects, they 'hide implementation and export function' in a similar way as the automobile cabin 'exports' the steering wheel and the gas pedal (the functions needed for driving) and 'hides' the details of the engine.

In the context of technology development, the proposed boundary objects play the role of requirements specification, pointing at a number of functional patterns that need to be implemented.

Their third role—and the one in which they appear in the work reported on in this article—is to serve as 'boundary objects' [8] for communication between the mentioned two domains or communities. They allow the Knowledge Federation (socio-technical design) community to tell the technical communities what functional *patterns* are needed. Those *patterns* can then serve an analogous role as 'problems' such as sorting and matching had in traditional computer science—invite the creation of suitable procedures, technical tools and their continuous improvement. But the communication is also enabled in the opposite direction—the technical researchers are invited to comment on the design of those objects and propose improvements, based on the technology and insights developed in their circle.

In the remaining sections we define the Domain Map Object and the Value Matrix Object in terms of the *patterns* they enable, describe a prototype implementation, discuss two applications that are under development, and introduce a DebateGraph-based platform created to enable evaluation and further improvement through a cross-community dialog. The concluding section will point at further research.

2 Domain Map Object

A Domain Map Object (DMO) is a knowledge management object assigned to a community of interest, which serves as an online representation of its domain of interest, providing access to resources and ways to organize them and create high-level, community-authored views.

The Domain Map Object implements all the functions that the community members need to keep their domain organized in a way that minimizes overload, and makes relevant insights and other resources available to the larger community. Useful metaphors for understanding its role may be a filing cabinet (a suitable places for organizing and keeping knowledge resources) and system of geographical maps, with different views (physical, political, climatological...) and levels of detail [9].

2.1 Patterns Enabled by the Domain Map Object

The following knowledge-work patterns are enabled by the Domain Map Object:

- **Domain mapping**. This pattern allows the community members to create a representation of their domain. Periodically, the community 'goes off-line' and 'redraws its map.' A domain map is composed through so-called 'structural abstraction'—by reasoning about the domain, and representing not only the knowledge that exists, but also the knowledge that naturally belongs to the domain but does not yet exist. A rudimentary implementation is a taxonomy. Visual [9, 10] and other media techniques here find natural application. A domain map offers, and maintains, multiple views of the domain, reflecting distinct purposes and levels of detail [11]. Analogy with geographical maps is useful where we have physical, political, climatological and other maps, at a variety of scales. Another useful metaphor is the one of a mountain – which allows us to come out of information jungle and get an overview [12].
- **Putting resources on the map**.
- **Putting questions on the map** is a way to initiate collaboration with 'whoever out there may know how to solve this problem.'
- **Federation** enables the community to co-create 'their messages to the world.' Disagreements are part of the community's internal creative process. At the same time, the function the community has within the larger community (specifically an academic discipline within the society) is to *federate* (synthesize, agree upon, express in an accessible way, make known to people and communities who may need them) essential high-level insights about its domain of interest. Analogy with political federation is useful: If we think of conventional published articles as independent countries, and a Wikipedia article as a single state without autonomy, then *federation* stands for the rich space of possibilities between those two extremes.
- **Registering interest** allows users to be notified when new resources are placed at chosen areas of map.
- **Locating resources relevant to a subject** is made possible through a position on the map.
- **Exporting resources**. The owner community makes specific resources (typically conveniently expressed federated views) available to journalists, general public and specific communities of interest.
- **Importing resources**. The Domain Map Object has provisions for allowing other communities to make suitable resources available to the owner community.

2.2 Processes Enabled by the Domain Map Object

The Domain Map Object will facilitate or enable the following processes:

- **Structuring and presenting a domain**—allowing a community of interest to create the way its members, and the rest of the world, see their domain.
- **Organizing resources**, by making them visible and available on the map.
- **Directing research**, by posting questions, and showing areas where more research is needed.
- **Life-long / flexible learning** by providing a map for orientation.

3 Value Matrix Object

The Value Matrix Object (VMO) is an object associated with a resource (such as a document or a person), which accumulates and makes available all information that could be relevant for estimating the value of that resource for any specific query. The VMO may be visualized as a matrix whose columns are criteria, and whose rows are various ways of evaluating those criteria.

The key insight is that the value information needs to be *federated*: Decisions about what is good and useful and what is not should not be made centrally (which in a political metaphor would resemble an authoritarian rule rather than a federation); rather, all information that may be relevant to such decisions is collected and made available through suitable functions or queries. A reason is that in knowledge-work sometimes (but of course not always!) the 'underdog' turns out be the 'winner' (brings vital new insights).

Another key idea is that by *federating* the information about *all* contributions of a community member, undervalued contributions such as knowledge organization and evaluation can be made visible and adequately rewarded.

3.1 Patterns Enabled by the Value Matrix Object

- **Customizing the VMO**—the community chooses the values it wants to support.
- **Assigning a value matrix to a resource** is a refinement of 'putting a resource on a map;' the resource's visible or 'existence on the map' is relative to its value.
- **Recording evaluations**—the information about the value of resources is itself valuable information; contributions to to it are recorded and rewarded; the identity and the value (matrix) of the community member doing the evaluation is a relevant part of the evaluation record.
- **Recording events**—a challenge here is to keep track of *every* event (quotations, visits etc.) that may later be relevant for evaluating the resource.
- **Querying the VMO**—a challenge is to provide functions and algorithms (i.e. affordances) for taking advantage of the accumulated information.

3.2 Processes Enabled by the Value Matrix Object

The Value Matrix Object facilitates or enables the following processes:

- **Precise query**. If we think information-theoretically about information overload, then it becomes clear that if a query should distinguish a resource or a handful of 'right' resources from a large number of competing ones, sufficient amount of information must exist in the system for making such distinctions. A role of the VMO is to collect and provide that information.
- **Creative reputation management**—borrowing an idea from computer games: We construct 'games' to reward all contributions, and create good systemic ecology. The VMO enables a promotion committee to obtain accurate profile of a candidate, across multiple dimensions.
- **Creative query strategies**—one may, for ex. recognize Christos Papadimitriou as an especially insightful community member, and ask to see articles about a subject that Papadimitriou ranked as outstanding; this will motivate Papadimitriou to spend time on annotation.
- **Creative remuneration**—to stimulate investment into quality rather than quantity, for ex. in journalism.
- **Crowdsourcing without overloading**—everyone can be given a voice and a chance to contribute; students can learn actively, by extracting meaning and organizing resources; good grades increase visibility.

4 Initial Prototype

We here present a partial implementation of the Domain Map Object and the Value Matrix Object, which is currently under design as part of a MS thesis. Its purpose is to serve as a research prototype—to identify and exemplify the design issues, and make an initial step towards solution.

4.1 Domain Map Object Prototype

The Domain Map Object in the prototype only provides a container for topics and access to knowledge resources. In other words, it serves as a map where users of the prototype can place new resources within topics as well as browsing topics to find resources.

In addition to acting as a map, the DMO presents a few functions that let users interact with the map. The first of these allows any user to add and place new knowledge resources to the map. For simplicity, we let all knowledge resources in the presented example be scientific research articles.

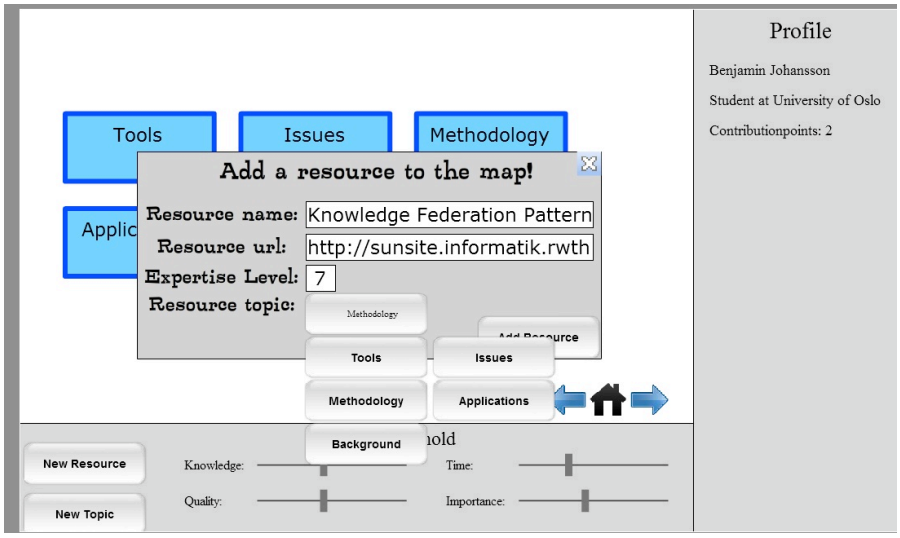

Fig. 1. A screenshot of the DMO user interface, showing the 'add new resource' dialog.

In the screenshot (Fig. 1) we can see what happens when a user activates the "add new resource" function. The dialogue window lets the user input information about the resource before submitting it. In this prototype the information required is the resource's name, it's url that points to the physical location of the resource, the recommended expertise needed to understand the resource content, and lastly the topic the resource should be assigned to.

Similarly to resources, users also have the possibility to add new topics to the domain map.

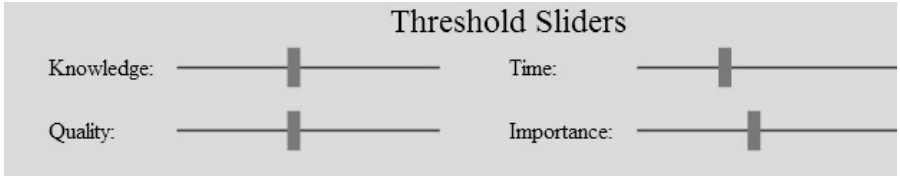

Fig. 2. A screenshot of the DMO 'threshold sliders.'

Lastly, we have the control panel functions. The control panel allows the users to specify which resources are shown within a topic by dragging a set of sliders up or down. Fig. 2 shows the 4 sliders that are currently available in the prototype. The first slider, named "knowledge", operates in the interval [1, 10] and controls the lower bound of personal expertise required to understand the content of the resources. For example, a university professor might set this slider to "7" in order to only display the

resources which have an expertise level of 7 and above, because he is not interested in introductory articles in the topic he is browsing. It is possible to expand this slider to take into account both upper and lower bound so that users can find resource with an expertise level between, for example, 3 and 7.

Next, we have the "time" slider. This allows users to exclude resources that were added or updated after a given time. Currently, the slider has 4 levels; lifetime, 6 months, this month and this week. Obviously, manipulating this slider allows the user to only see recent resources or all resource if he wishes.

The last two sliders are called "quality" and "importance". They both operate in the interval [1, 10] and allow the user to specify the level of quality or importance respectively a resource needs to have in order to be displayed. In other words, a user can choose to only show the very highest level quality resource, or everything that is not mediocre or poor.

Naturally, all sliders can be combined. For example, a user might want to look at all high quality resource that were added in the last month and require a high level of expertise to be understood.

4.2 Value Matrix Object Prototype

The role of the Value Matrix Object in the prototype is to collect information about the value of the resource the value matrix is attached to. From an endless range of possibilities for data collection we have chosen a handful. The most important dimensions are quality, importance and relevance, which are collected through user ratings. Obviously, a more advanced prototype would have automatic ways of collecting these data as well.

The ratings that have been collected and stored in the value matrix are used by a reputation computation engine that computes a reputation score for each of the mentioned dimensions. These scores are the basis for the sliders that the domain map provides, i.e. a resource with the quality score of 4 will not be displayed on the domain map if the user has set the slider threshold to above this number. Currently, the prototype uses an average of ratings algorithm to compute reputation scores, but it is preferable to use a more advanced scheme, such as a Bayesian system, in the future.

Relevance does not have its own slider, because a resource should only be relevant or not relevant to a topic. The prototype treats resources that have a relevance reputation score of less than 0.5 as not relevant, and thus will not appear on the map within that topic and vice versa.

The minor dimensions on the prototype are time, knowledge level and keywords. The value matrix keeps track of the time its resource was added to the map as well as the time when the resource was last interacted with. This way, resources that were added at the beginning of the domain map's lifetime, but are still interacted with, will still appear even as the time slider is adjusted to only show recent resources.

When a user adds a new resource to the map he also sets a knowledge level on that resource. In the prototype, this is the only time information about the knowledge level is collected. However, in future versions we might allow other users to voice their opinion about a resource's knowledge level, in a similar way in which quality, importance and relevance are presently handled.

We use the concept of keywords in the context of the value matrix to create a loose relationship between resources. It enables the use of "keyword search" through the domain map's control panel to find resources that have the same subset of keywords, and by extension talk about the same general subject. For example, a user might search for the keyword "reputation system" and as a result get back all resources that contain that keyword. It is also possible to use several keywords in a single query, and then the resources that have the largest subset of the query keywords are ranked highest.

In our present prototype, a resource's keywords are added manually by users. In future versions it would be preferable to have an automatic way of generating or extracting keywords as well. Since our resources are scientific articles and they often contain a list of keywords, it is possible to extract those and add them to the resource's value matrix. Alternatively, the keywords could be generated from the resource content itself.

In addition to attaching a value matrix to resources, also users have their own value matrix. This side of the VMO concept is not elaborated in this prototype. A user's value matrix's responsibility is to collect information about the user's behavior. For example, adding new resource's to the domain map, or contributing to the federation of knowledge through interacting with already existing resource's impacts the user's value matrix positively. Naturally, deliberately flooding the domain map with resources or giving a large number of ratings within a short time span is unnatural, unwanted behavior and will impact the value matrix negatively.

A user's value matrix is used to compute a reputation score for the user. This reputation score will in turn impact the weight of the ratings the user assigns to resources. In other words, a user whose interest lies in contributing in a good way to knowledge organization will have a higher influence on the value of resource than a user who tries to deliberately damage the system.

5 Examples of Application

Knowledge work that is oblivious of (the ineffectiveness of) its own practices is a contradiction in terms; we predict that self-organization will be recognized as a core task of academic and other communities of interest, and offer the Domain Map Object and the Value Matrix Object as an initial *prototype* of a suitable toolkit.

In addition, our boundary object enable completely new ways of organizing knowledge work, as the following two examples might illustrate.

5.1 Good Journalism Transdiscipline

A discipline joins together researchers and students with similar background to study subjects within the domain of the discipline; a *transdiscipline* joins together knowledge workers from a variety of fields, to work on an issue of contemporary interest.

A discipline inherits its body of knowledge through education; a *transdiscipline* creates its body of knowledge by organizing relevant resources from existing disciplines, with the help of the Domain Map Object.

The Good Journalism Transdiscipline has been initiated at the Knowledge Federation workshop in Barcelona, in November 2011. The insight that motivates this work is that public informing is too important as systemic component to be abandoned to spontaneous and commercial interests. The Good Journalism Transdiscipline is a community consisting of journalists, knowledge media workers, collective intelligence experts, IT companies and other stakeholders, organized to continuously co-create a public informing that can enable contemporary democracies handle contemporary issues.

A prototype Domain Map object belonging to this transdiscipline is implemented by using DebateGraph [13].

5.2 Knowledge Federation Course

A globally federated Knowledge Federation course is being prepared, to be offered once a year through Inter University Centre Dubrovnik, beginning September 30, 2012. This course will be co-created or *federated* by an international team of experts, and offered to students worldwide.

To the course expert team, the course domain map provides the tools that are needed to organize the resources; to the students, it provides access to resources and orientation.

The students learn actively, by contributing resources to the course domain map. The associated value matrix objects ensure that the best ones remain visible. All work —including the student projects, and grading—contributes to the knowledge-work ecosystem.

We predict that as technology makes many jobs and professions obsolete, life-long education will become the rule rather than exception. The flexible course model developed in this project will not only enable students to learn exactly what they need, but also to learn while doing useful work, and contributing. Grades and degrees

may ultimately become be conferred by more advanced versions of our boundary objects.

6 Boundary Object Dialog

The described object prototypes are made available online, and a dialog space for their evaluation and further development is created the on the DebateGraph platform [14].

Combined with the objects, this dialog completes 'the boundary' enabling two-way communication between the Knowledge Federation (system innovation) and technical communities that can help implement or improve the initial toolkit.

7 Conclusion

The information overload and the global issues are consequences of the fact that our knowledge work is conceived as a collection of document-producing disciplines or professions, not as a 'global mind' capable of fulfilling its core functions. Hence the 'flat tire' metaphor applies—our situation calls for stopping and taking care of structural problems.

We predict that the solutions will be found within an approach to IT innovation where suitable socio-technical systems—capable of performing the core functions—are designed first, and then used to develop specifications for technical components that are required. We submit the work described here—the Domain Map Object, the Value Matrix Object and the Boundary Object Dialog—as a step in that direction.

References

1. Diderot, D., *Rameau's Nephew and Other Works*. Doubleday & Co. pp. 277-307 (1956).
2. Karabeg, D., Information Age Paradox. Blog post at Knowledge Federation community Website. Doubleday & Co. pp. 277-307 (1956).
3. Toulmin, S., Return to Reason. Harvard University Press, 2001.
4. Karabeg, D., Design Is an Alternative to Tradition. Proc. EAD06, Bremmen (2005).
5. Karabeg, D.,Park, J. (Eds.): Proceedings of the First International Workshop on Knowledge Federation, Dubrovnik, Croatia, October 2010. CEUR Workshop Proceedings Vol. 822, http://ceur-ws.org/Vol-822/ (2011).
6. Karabeg, D.: Knowledge Federation – An Enabler of systemic Innovation. Proc. Triple Helix IX Conference, Stanford University, July 2011. http://knowledgefederation.net/Articles/DK-3H9.pdf (2011).
7. Karabeg, D., Herrera, R., Price, D. and Sangüesa, R., An Innovation Ecosystem for Good Journalism. Submitted to the workshop "New Political Communication and EU."

8. Star, S., Griesemer, J.: "Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39". Social Studies of Science 19 (3): 387–420. (1989).
9 Quaggiotto, M..: Images of Knowledge. Interfaces for Knowledge Access in an Epistemic Transition. Proc. of the First International Workshop on Knowledge Federation, Dubrovnik, Croatia, October 2010. CEUR Workshop Proceedings Vol. 822. http://ceur-ws.org/Vol-822/MQ.pdf (2011).
10. Visual Insight: Blog, http://visualinsight.net/portfolio.html.
11. Guescini, R., Karabeg, D., Nordeng, T.: A Case for Polyscopic Structuring of Information. . In Maicher, L., Park, J (Eds.): Charting the Topic Maps Research and Applications Landscape. LNAI 3873 (2005).
12. Karabeg, D.: Knowledge = Mountain. Five-minute speech transcript, TMRA 2007. http://www.informatik.uni-leipzig.de/~tmra/2007/slides/karabeg_TMRA2007_os.pdf (2007).
13. Knowledge Federation: Good Journalism Domain Map. Domain map implemented as DebateGraph map, http://debategraph.org/Stream.aspx?nid=127659&iv=05 (2011).
14. Knowledge Federation, Boundary Object Dialog. http://DebateGraph.org/BoundaryObjectDialog (2012).