

UNIVERSITY OF



Communications of the UO

December 2004

Table of Contents

Crossing the Line between Mapping and Map Making	1
Crossing the Line Between Mapping and Map Making	1
Artistic approaches to cartography	1
A London Free Map	6
A London Free Map	6
Step 1 - Collect GPS data	6
Step 2 - Normalise and Synthesise Data into Real Shapes	7
Step 3 - Annotate Shapes with Real-World Semantics	7
Step 4a - Annotate Completed Map with Civic and Historic Information	8
Step 4b - Find Partner Organisations to Contribute to, and Back, the Project	8
Step 5 - Give Everything Away	8
The Freemasons of the Future	9
The Freemasons of the Future	9
Who Are The Freemasons of the Future?	9
Mapping the Semantic Web	10
Technical Developments	10
'Total Information Awareness' & 'Consensual Reality'	11
Technical Examples	12
Ontology Examples	13
Free Association	14
Some Responses	15
Glossary and Links	16
Footnotes	17
The Epistomat: Generating Consensus About RDF Ontologies and Rules	18
The Epistomat: Generating Consensus About RDF Ontologies and Rules	18
Abstract	18
Keywords	18
Contents	18
Introduction	18
Generating Interfaces with RDFS and OWL	19
A Game of Nomic	20
A Voting Model for Ontologies and Rules	21
Ontomation and the Semantic Self-Defense	22
Generating Consensus	24
The Epistemological Congress	25
Conclusions	25
Acknowledgements	26
References	26
Lucrative Media	27
Lucrative Media: ownership and authorship of public space.	27
Credits	33
Credits	33

Crossing the Line Between Mapping and Map Making

Artistic approaches to cartography

'Mapmaking is *not* a universal expression of individual existence (like something we might call mapping), but an unusual function of specifiable social circumstances.'[1]

In his essay 'The Fine Line Between Mapping and Map Making', Dennis Wood draws on cognitive science to inform his definition of 'mapping', a process he sees as common to all humans; analysing sensory data and building a mental representation that we can use to 'get around and do stuff'. He sees this process as both universal and individual in that our 'mapping' process is made unique by our varied abilities, disabilities, and subjective perceptions. Map making then, is an attempt to externalise, or transcribe some aspects of this experience using shared languages, rituals and patterns that are culturally specific.

Artists working with cartography, experimenting with its established techniques, often find themselves working on Wood's 'fine line between mapping and map making'; between a sensory and psychological reaction to a terrain and the production of an inscription that attempts to communicate that experience to others. However, beyond this simple description which could almost be said of any plastic art, cartography has some specific qualities and potentials that increasing numbers of artists are identifying and using in their work.

This fashion is partly due to the scientism of cartography, how its representational forms and spatial metaphors are so well established, widely understood and accepted as fact. While many are suspicious of contemporary art and regard it with reverence or disdain, leaving it to qualified experts to produce and interpret, maps are still a vernacular form that everyone is expected to be able to read and write. So for artists who wish to communicate beyond a community of experts, cartography offers a formal structure that will probably be legible to most people, and that has an implicit utility. This is not to say that any map, artistic or not is necessarily useful for getting from A to B, but that it at least asks the reader to situate themselves, the map author and the things represented by it in relation to each other. This relationship may be spatial or conceptual, but at the very least, maps demand this much.

On a more basic level, thinking about what the map does, what its functions are, and how they can be used by artists for their own ends (about which more will be said later), Michel Serres' concept of the 'quasi-object' provides a useful model:

'This quasi-object that is a marker of the subject is an astonishing constructor of intersubjectivity. We know, through it, how and when we are subjects and when and how we are no longer subjects. 'We'; what does that mean? We are precisely the fluctuating moving back and forth of the 'I', The 'I' in the game is a token exchanged and this passing, this network of passes, these vicariances of subjects weave the collection.'[2]

Serres uses the example of the rugby ball as an extreme form of 'quasi-object' which he calls an 'operator'. As the 'operator', in this case a the rugby ball, is passed from player to player, it takes on the indeterminate function of a joker in a deck of cards or a blank domino in a set; an object that can change the course of play and the roles of the players. As the ball is passed between the subjects and groups of subjects (the teams in this case), the players take turns at being the 'I'. In this shifting, convoluted dynamic, Serres describes the subjectivity of each player being woven by the passage of the 'operator'. He then turns to less extreme 'quasi-objects', that are not as indeterminate as operators, but which are created by customs, languages, and rituals to solidify social relations. Steven Brown extrapolates on this

idea using the example of how the sword formalises combat, or cutlery stabilises eating arrangements:

'Objects make stable human relations possible, but it is the putting of objects into circulation as quasi-objects which realises the human collective as such... The quasi-object does not so much reverse play as formalise it, often giving rise to concrete, highly deterministic social practices.'[3]

Knives, forks and spoons provide a set of tools for the establishment and development of manners, just as certain weapons determine and change how violence is ritualised and performed. This is how Serres would view the map, as a formalisation of human relations, a representation within which each player becomes a subject. This is the use of the map as a communicative tool, as successive players engage with the map, each locates themselves in the representation, the 'I' is shifted from person to person, and between person and group.

The 'highly deterministic social practices' that Steven Brown mentions refer to the way these tools can fix relationships and representations in a framework that is not easily malleable or reconfigurable by their users. The instances of artist-made maps I will discuss below focus on attempts to destabilise this deterministic function of map making, opening up maps to the desires and needs of those who communicate through them.



Snapshot of the Downhill Map of Bristol, 2004, Heath Bunting. ([http://locate.irational.org/bristol_map/.](http://locate.irational.org/bristol_map/))

Artist Heath Bunting's Downhill Map of Bristol is a map for skateboarders, cyclists and walkers. Rolling through the streets of Bristol, Bunting and the Bristol Skate Survey group personally surveyed private paths that connect public ways, road surface quality, altitudes and the location of fruit trees: detail that would be unavailable to most map makers. The web-based software Bunting wrote to enable them to map the city allows for a great deal of flexibility in how the map is eventually rendered, as well as allowing the survey group to plot and annotate points on the map, extending it by connecting points representing road junctions with a simple point-and-click drawing interface. The use of technology in the practice of the survey is very minimal: the web site allows the surveyor to print out little strips of out-of-copyright maps showing unsurveyed spaces that then enable the group to skate around drawing alterations, additions and notes that can then be added to the map via a web form.

There have been many projects, artistic and commercial, that provide a similar set of functions. Most prominent is The **Urban Tapestries** (<http://urbantapestries.net/>) project which has developed a platform for 'collaborative mapping' based on the use of handheld computers, Global Positioning System and 3G mobile phones. This project uses the concept of 'annotation' of points on an existing map with photos, video and text. Many artistic mapping projects[4], work with similar ideas of augmenting existing maps and spaces with local, user-contributed information.

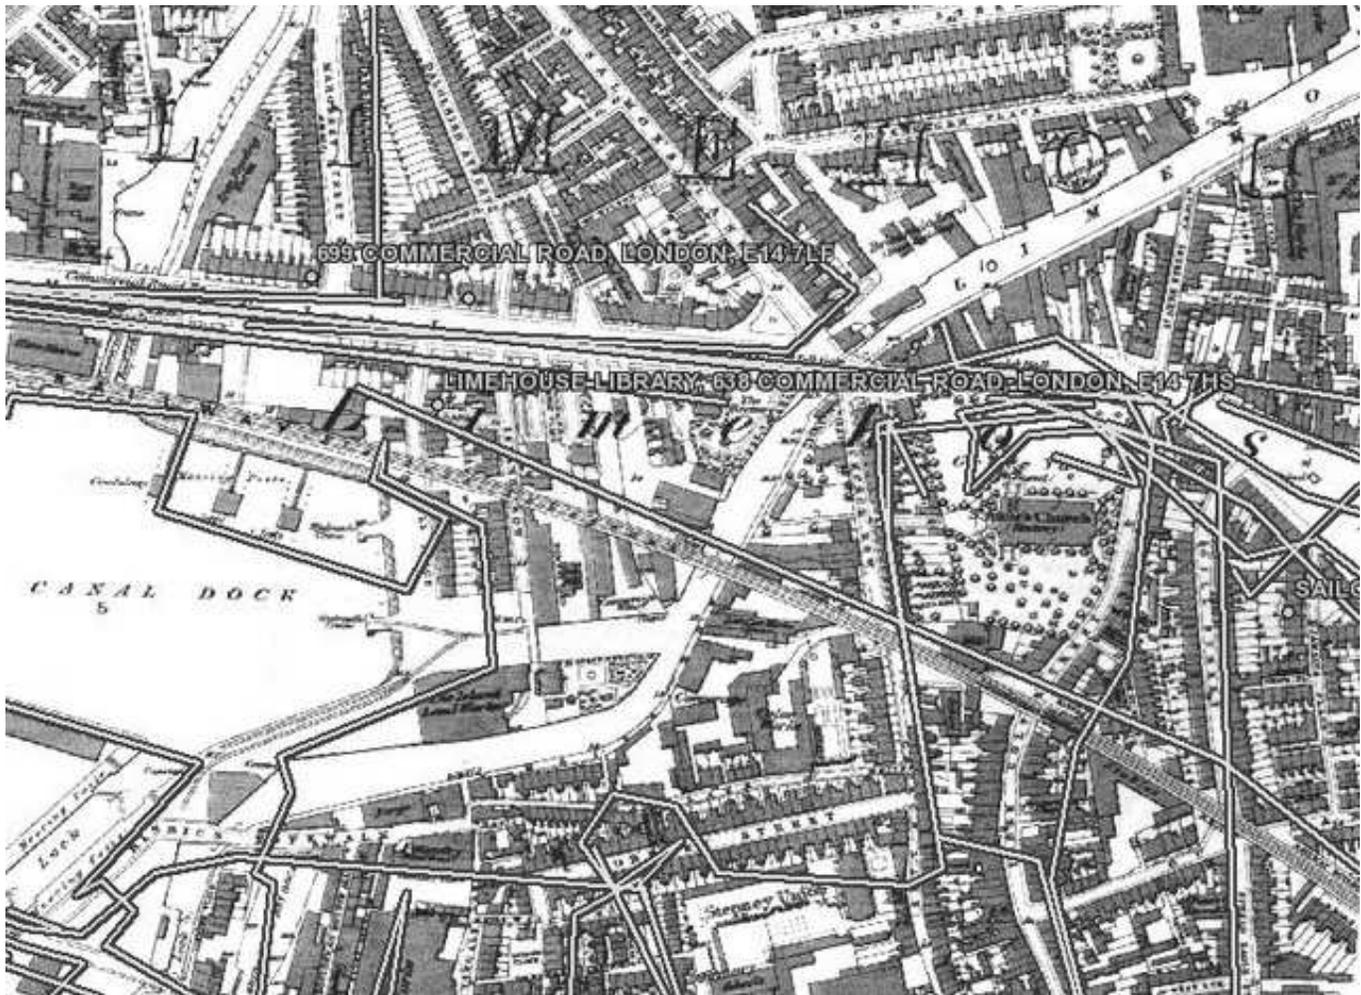
The distinctive factor of Bunting's map is the reliance on the survey group to become more than just 'users' of the map by taking part in the process that builds and extends the map itself. Bunting's Downhill Map of Bristol is a functional outcome of that mapping process that provides users of human powered vehicles in Bristol a unique navigational tool, providing details about road surface and gradient that open up the possibility of constructing a functional online bicycle and skateboard route planning service[5] that no official mapping agency would be able to offer.

Returning to the idea of opening up the map to the needs and desires of its users, this ability to reuse and extend the base-map itself, and for the users to help shape the functional operations, not just the appearance or 'content' of the map, addresses far more basic concerns than the 'annotation' based projects that use preexisting maps as the basis for historical or personal narrative layers.

This distinction is far more poignant in the UK where all geodata is restricted by crown copyright and owned by the Ordnance Survey, a quasi-commercial agency that licenses geodata collected using public funds back to the public for extortionate sums. In this context, the ability to create a map, use it, publish it, and to give away its underlying data is a political statement in itself. Also, the process of the survey, observing the terrain and representing it as a map, engaging sensory experience and the impulse to transcribe that experience, comes closer to Dennis Wood's 'fine line between mapping and map-making' than the single-point annotation process which some skeptics have jokingly called 'Red Dot Fever'. This pejorative label not only comments on the virulent proliferation of these 'annotation' projects, but seen in terms of what the map does, or as Wood asks 'what its discourse function is', 'Red Dot Fever' projects do little more to open up use of the map than allow the user to add their own 'you are here' symbols.

This can be seen as a step in the right direction, but in many cases, particularly in the case of projects such as Urban Tapestries, where the use of the map is dependent on the use of high-tech, expensive consumer equipment and relies on sponsorship from 3g networks such as Orange, equipment manufacturers HP, as well as proprietary geodata donated by the Ordnance Survey, the addition of a layer of user-contributed narrative begins to seem more like a solicitation for 'content provision': asking people to help convince consumers to buy this equipment and these services. This project might destabilise the deterministic functions of maps and map making, but given the technical and economic dependencies it demands of its 'users', Urban Tapestries appears to be little more than thinly veiled advertising and R&D for equipment manufacturers and service providers.[6]

For the Downhill Map of Bristol, Bunting addresses this issue by using cheap or free technology and plans to make both the software, and the data underlying the maps he is making available to the public under open licenses. Each user of the map, given these conditions, is capable of setting up their own mapping system and using it for their own ends whether artistic, cartographic, both or something else entirely. The difference in this case is that the aims, methods and the outcomes of the project are malleable on all levels - the practice does not necessarily become the kind of 'deterministic social practice' Steven Brown talks about in relation to objects such as maps.



Jo Walsh and Schuyler Erle, snapshot from the [London Free Map](http://uo.space.frot.org/?node=LondonFreeMap) (<http://uo.space.frot.org/?node=LondonFreeMap>)

In a similar vein, Jo Walsh and Schuyler Erle from the University of Openness' Faculty of Cartography and Steve Coast from the department of Physics at University College London have been working on London Free Map and Openstreetmap.org. These projects encourage participants to walk, drive, cycle or skate through city streets with GPS units and then use the 'traces' of points representing latitude and longitude that these devices generate to help create publicly licensed geodata. They are also working on adapting existing open source software[7] to enable people to annotate these maps in a very flexible way. The example above shows straight lines representing GPS traces made wandering the streets of Limehouse, East London, an area undergoing a huge urban regeneration process in preparation for the proposed Olympic games in 2012. The labelled points are the locations and names of approved planning permission applications made to the London Borough of Tower Hamlets in the last five years, automatically retrieved from the council's website and plotted onto a scanned out-of- copyright historical map of the area from 1916. This is just one potential use of the London Free Map, a way of visualising physical and historical changes to a space undergoing a huge social and economic upheaval. The potential for further uses and the development of new, as yet unimagined maps from this project seems evident.

Seen in these terms, Bunting's Downhill Map of Bristol and the London Free Map projects bring the practices of cartography and art together as part of a larger movement that begins to dispense with the deterministic aspects of art, cartography and other practices that generate or exchange intellectual property in this way. This movement, that has been loosely referred to as 'Free Culture'[8], encompasses Open Source software, 'copyleft' licensing structures and other practices that attempt to enable access to and ownership of their ideas, data, and methods by the people, or 'public' who use them. It is this level of access to the map, through which each map reader potentially has the ability to add to, reconfigure, or completely restructure the map they are using that finally crosses Dennis Wood's 'fine line between mapping and map-making'. The subjective understanding of the mapped terrain that each map-reader has, but is usually compelled to abandon to the norms of scientific cartography, or to the whims of the artist/cartographer, potentially becomes the starting point for a re-working of that map based on, their experiences, needs and desires.

[1] Dennis Wood, 'The Fine Line Between Mapping and Mapmaking', *Cartographica*, 30.4 (1993), p. 50.

[2] Michel Serres, *The Parasite*, (Baltimore: John Hopkins, 1981), p. 227. Quoted in Brown, 'Michel Serres: Myth, Mediation and the Logic of the Parasite'.

[3] Steven Brown, 'Michel Serres: Myth, Mediation and the Logic of the Parasite' (Loughborough University unpublished draft, 2000), see: <http://devpsy.lboro.ac.uk/psygroup/sb/Serres.htm> (04/02/2003).

[4]A short list is available here: <http://uo.space.frot.org/?node=RedDotFever> (25/11/04).

[5]This service is actually under development at this address:
http://locate.irational.org/cgi-bin/bristol_map/bristol_map.pl (25/11/04).

[6]On an artistic level, soliciting participation in this narrative layer on top of a fixed base-map achieves a similar effect for the artist as for the corporations who sponsor them. By slightly shifting the determinism of the map, allowing it to be re-centred and contributed to in a 'red dot' mode, the activity and creativity of anonymous 'contributors' bolsters the credibility and value of the artist's work. The artist can use this reputation value to seek remuneration in various academic and commercial contexts. The destabilisation of one deterministic practice (that of cartography) can be seen as leading to another (that of Art).

[7] They are using the Open Guides system (see <http://london.openguides.org>), a 'wiki' type website that any visitor can edit, but enhanced specifically for spatial references, with latitude and longitude markers for each item on the site.

[8]This term has been popularised by Intellectual Property lawyer Lawrence Lessig, who has recently published a book titled 'Free Culture: How Big Media Uses Technology and the Law to Lock Down Culture and Control Creativity'. It is free to read and reuse and is available for download here: <http://www.free-culture.cc/>

A London Free Map

We're making free-of-copyright street level maps of London, starting in the East End and working outwards. We plan to rapidly develop evocative prototypes, working through Steps 1-3 below, and repeating, working outwards.

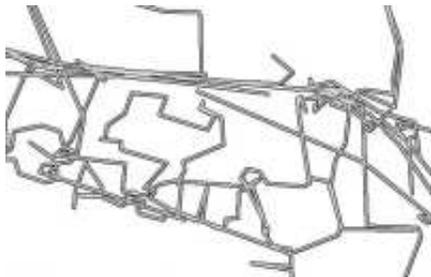
This effort is in collaboration with [Openstreetmap](http://openstreetmap.org/) who've been working on a similar plan for a few months, and with whom we hope to share data, hosting and algorithms, and keep the design language-neutral. To get involved, please [join the mailing list](http://bat.vr.ucl.ac.uk/cgi-bin/mailman/listinfo/openstreetmap) (<http://bat.vr.ucl.ac.uk/cgi-bin/mailman/listinfo/openstreetmap>).

On [GIS Day, Nov 17th 2004](http://gisday.com/) we gave a short talk about Our Plan at [dorkbotlondon](http://dorkbot.org/dorkbotlondon/) - see [FreeMapsAtDorkbot](http://dorkbot.org/dorkbotlondon/), and the [landsat-based demo showing tracks, consume nodes and old maps](http://locative.us/freemap/map.cgi?x=0&layer=landsat) that we showed.

Step 1 - Collect GPS data

This part is both easy and hard. We walk, or cycle, the streets of London, with GPS units, recording our traces and waypointing useful and easily identifiable points.

We're building a simple web-based interface where people can upload GPS tracks and points in [GPX](http://www.topografix.com/gpx.asp) format, a simple and common XML format. We have two GPS units to lend, and are looking for resources to acquire more.



Any kind of previously collected GPS traces - netstumpers, geocachers? - are of use and interest to us. If you can offer these, we'll provide an upload facility.

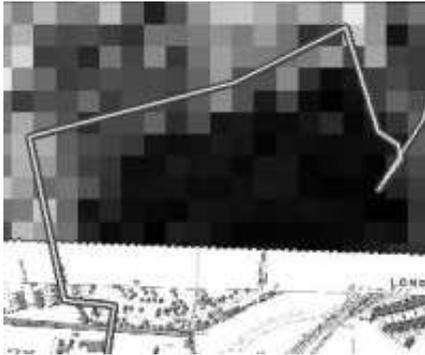
There is a *window of opportunity* several times a day - currently around 7-9pm and 4-6am - when there are 3 satellites directly overhead, thus accuracy is better and it is possible at all to trace areas with tall buildings and/or narrow streets. The calendar server we were using is currently down; we'll build our own if necessary.

Raw GPX tracks will probably be stored on a shared WebDAV server, perhaps hosted at UCL, for anyone to grab and re-process.

We'll store the collected points and segments in a simple [PostGIS](http://postgis.refrations.net/) database.

Step 2 - Normalise and Synthesise Data into Real Shapes

This is the hard part, code-wise. We start with a big collections of points, which belong to arbitrary segments. We'll write algorithms which decide when several very close lines are actually the same line and thus need to be averaged together. We'll write algorithms which decide where intersecting lines are junctions, and the street needs to be split up into segments. We'll decide when small gaps in lines probably indicate a path running between them. We'll have techniques for turning collections of lines into polygons - which describe parks, or squares, for example.



Technically, this will involve a lot of hard thought and frequent mistakes in a quiet room.

- Line identification/simplification
 - co-linearity
 - **Douglas-Peucker** (http://geometryalgorithms.com/Archive/algorithm_0205/)
- Line extension/connection
- Line deduplication

Step 3 - Annotate Shapes with Real-World Semantics

This is the demanding part, effort-wise: *Ground Truthing* and improving on raw street data. Streets and street segments need to be labelled with their names; some of this will be possible online, from local Knowledge or by cross-referencing with old maps.

Some street semantics are crucial for navigation: e.g., "This road can be driven along but not walked", "This path can be walked but not driven along", "This road is one-way in a north-to-south direction along this segment", "This road is a major artery or a minor road".

An ideal goal is to label street segments with street numbers, so we can do proper free address geolocation. It may be possible to conduct a "civic survey", offering explanatory leaflets to homeowners in exchange for divulging their postcode.

To complete this long task properly, we need a portable annotation device, ideally with wireless connectivity. We no longer need to be going out with a GPS at this stage. An assortment of iPaks which we can lend to trusted collaborators would be perfect.

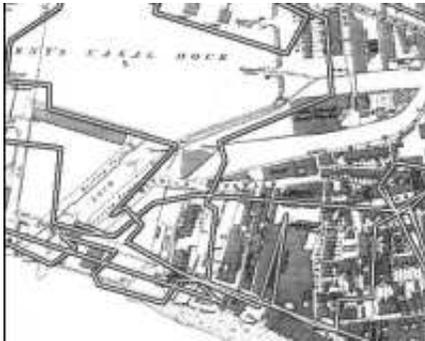
We'll use **UMN Mapserver** (<http://mapserver.gis.umn.edu/>) to build a first-cut annotation interface on top of PostGIS.

We need to develop, or find, some kind of standards-looking, good enough MapOntology.

Step 4a - Annotate Completed Map with Civic and Historic Information

We're looking for interested parties with deep local knowledge and existential commitment. Prime candidates in the East End are the [East London Postcard](http://www.eastlondonpostcard.co.uk/) group, who are taking historical/contemporary comparative photographs, and are interested in doing online maps of walks to accompany their own publications.

Annotations don't need to be contained in our own repository; they can live at another specialised service, like the [Open Guide To London](http://london.openguides.org/), that we can talk to.



There are old maps, out of copyright, in circulation. Some are very high detail, but a lot of East London has changed beyond recognition since they were made. Still, they are of essential social interest and value.

Step 4b - Find Partner Organisations to Contribute to, and Back, the Project

[London Green Map](http://londongreenmap.org/) have a great collection of data, mostly points, describing green spaces and sustainable public resources like recycling stations, which is crying out for open publishing and inclusion in a London Free Map. We need to get to the stage where we have a plausible local prototype before starting to negotiate with them.

Local authority Planning and Building, and Street Management, departments have a lot to gain from a collaboratively built and maintained, free of copyright, body of geographic data describing their locality. It may be possible to suggest a join application for an Open Source GIS project with Hackney and/or Tower Hamlets in the future.

Step 5 - Give Everything Away

Shapefiles, GML over Web Feature Service, RDF - any format we can think of or that people ask for export in. In addition to a Mapserver based web representation which people can call on as a base layer in their own sites and overlay points on top of. All the software we develop for London Free Map will be freely available under the GPL. We'll hopefully use [Here Be Dragons](http://herebedragons.org.uk/) as a 'geodata portal' for giving things away.

The Freemasons of the Future



Who Are The Freemasons of the Future?

The first and most salient fact about the Freemasons of the Future is that they do not exist. However they are sentient beings trying to struggle into existence. From their perspective they are involved in a life and death struggle to ensure their own past, some of which we perceive as the present. Located in the distant future after time travel has become commonplace they endeavour to sojourn into what they regard as history to create the conditions which they consider as necessary for their own existence. But this does not mean they are going to be successful.

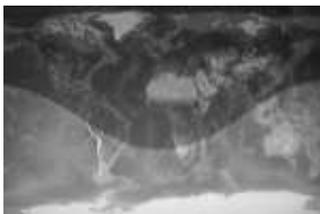
They justify their existence upon their existence, which they consider self-evident. The current conditions of life are to be extrapolated into the future until every quanta of human activity has been commodified: from genetic engineering to nanotechnology, the search for profitability is projected in every conceivable way. A world where sentient scraps of human biology exist as islands of wetware within the framework of microscopic cathedrals of computerised electronics. The distinction between human, animal and machine is dissolved as these products of bio-engineering are installed to fulfil operative functions within a nauseous system developed to do nothing but to manifest continuously expanding value. Whether we can regard such creatures as our offspring, or whether they are simply genetically engineered mutant beings created out of this or that strand of DNA perhaps is beside the point. This is the nightmare world to which we are heading, and which would provide the sort of massive bio-computer needed by the Freemasons of the Future to realise their greatest desire: unequivocal existence.

Faced with this onslaught which we can see around us, as all barriers to genetic engineering are torn down one by one, how can we respond? The class struggle now manifests itself in dimensions which have recently been invaded by the process of industrialisation.

From the industrialisation of the imagination, through television, to the industrialisation of knowledge through the internet, the information age continues to build on the 'achievements' of the Age of Steam, the Age of Petrol and the Atomic Age. The current episode we are living through is rattling asunder as the ripples of the QuantumTimeBomb[1] penetrate the deepest recesses of human activity.

-- Harry Potter[2]

extract from the announcement of the LimehouseTimeTravelRally [3]



Mapping the Semantic Web

Talking about the Web as a single entity has become a formally inadequate description of the increasing electrical interconnectedness of devices, processes, information and indices. Inadequate because these words imply a coherence that is not evident in the use of many incompatible formats, private networks, and non-indexed sections of network. This incoherent, frayed mess of networks are like an expanding and obscure territory for which there are no maps, or at least, no maps with standard keys, scales or control co-ordinates. In some ways 'surfing' or 'browsing' are increasingly appropriate metaphors for the superficial and indiscriminate ways our browsers allow us to use the Web. These limited researches are almost entirely dependent on the indices of one of the major search engines which has become the limit of the network, everything else is uncharted, unconnected and therefore largely inaccessible.

;; "The Semantic Web is an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation." -- Tim Berners-Lee, James Hendler, Ora Lassila, The Semantic Web, Scientific American, May 2001.

By attempting to develop an extensible and syntactically coherent language to describe networks and information resources, the Semantic Web Project promises (or threatens) to help make these maps. Using computer readable data formats and programmable agents with which to collect and categorise them, the object is to produce a schema from which to build a local description of local data formats, network topographies and information resources. This local description, fitting into the logical framework of the Semantic Web, can then be transposed into other contexts, linked to similar or related descriptions of other resources and networks, understood and used by human and software agents; put on the map.

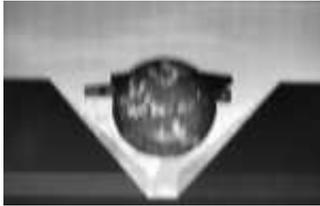


Technical Developments

In the mid 90's Ramanathan V. Guha went to work for Apple, where he developed a metadata format called Meta Content Framework which described websites, files, filesystems and relationships between them. The intention was that using Apple's 'Hotsauce' browser, users could fly through a 3 dimensional representation of that content. However, it was only when Guha moved to Netscape in '97 and Extensible Markup Language (XML) became a common standard for the exchange of structured, computer-readable data that his ideas about representing semantic associations between bits of data began to gain influence.

At that time the World Wide Web Consortium (W3C), the international web standards body founded by Tim Berners-Lee, began a general-purpose metadata project, loosely termed the 'Semantic Web', to develop ways of representing data on the web. Based on the Resource Description Framework (RDF), the basic idea is that resources are named with Uniform Resource Locators (URLs or web addresses) and described by the links between them using machine readable XML for syntax. The framework is general enough that it is not limited to describing data on the web, crucially it can also be used to describe and interrelate things in the world: people, places, objects, abstract concepts.

If we can assign a URL to a physical object, person, or idea, we can link other URLs to it, which in turn refer to people, objects, ideas or other links. Someone (or something) looking at this association can then make inferences about what is being represented from its associations, which can be further described and qualified by more links. The 'name-spaces', or vocabularies used in these descriptions can also be seen as nodes in this semantic network and linked to, extended, re-written and re-defined, so these representations are always contingent and non-originary. There is no start or end point, and no point of observation that can be outside them, just new nodes in the network.



'Total Information Awareness' & 'Consensual Reality'

The totalising invective of the Semantic Web project was very evident in one of its main predecessors, the CYC corporations' proprietary 'common sense' knowledgebase. This 'big AI' project was Guha's first job out of university, and involved the collation of a huge database of so-called 'common sense' statements. These statements were machine-readable so that software agents would be able to search through and make inferences based on them. A typical example is a CYC-based search engine that could respond to the question 'what is bravery?' by looking through its knowledgebase, finding an assertion that a property of 'brave' is 'danger', finding another saying that rock climbing is dangerous, and then retrieving a picture of a rock climber.

The notion of collating all 'common sense', (or 'consensual reality' as Cycorp sometimes put it), as a basis for artificial intelligence is a genuinely totalising and largely discredited idea. This problem, and the fact that the format of the knowledgebase and the modes and methods used to describe its contents were fixed, prescribed by Cyc's designers and their proprietary legal structures frustrated Guha, and gave him and his collaborators the impetus to attempt to formulate a more malleable framework, without this dubious premise.

The development of the Semantic Web; a machine-readable representation of everything, and its relationship to everything else, does sound like a step towards 'total information awareness'. It is true that the enriched and extensible vocabularies that the Semantic Web uses to describe relationships will expedite morally dubious activities such as surveillance, unsolicited direct marketing and military operations. These technologies will refine existing authoritarian systems for associating and describing things and people (consumer profiling systems for example) which are usually imposed without negotiation or consent, and by virtue of their limited interest in the person as a 'consumer', these representations currently remain very unsophisticated.

However, the extensibility of the Semantic Web, the fact that the person doing the describing can define the terms, the 'vocabulary' of that description suggests a less totalising, more heterogeneous 'information awareness'. This is both promising and potentially dangerous. Augmented by many more layers of information and description, volunteered by the person being represented, the 'consumer profile' becomes infinitely more insidious and detailed. At the same time, the greater sophistication of the Semantic Web's descriptive language enables someone to consciously and deliberately allow or deny access to specific data that they produce. Using cryptography, and 'friend of a friend' testimonial systems (sometimes called 'trust' networks) at least offers some degree of control over and awareness

of the data being exchanged about us. On a more structural level, the development of many divergent, even antagonistic descriptions of the world and the people in it moves away from the idea of any imposed 'consensual reality' and suggests a mode of representation that can be multiply subjective.



Technical Examples

RDF was developed as an open framework from philosophical inquiries by W3C about creating universal categorising systems, with the understanding that such a framework can never be comprehensive, hence the ability to add and modify the vocabularies used to describe and categorise things. These vocabularies are grouped into machine-readable XML documents called 'namespaces'.

The Semantic Web's use of the RDF common framework allows the data used in each description to be fully distributed in terms of storage and authorship. Not only can groups collate and share their own data, but also automate the aggregation and inclusion of publicly accessible data sources such as company profits, IMF trade data, names and connections between regulatory board members etc.

RDF's more widely-known derivative is Rich Site Summary (RSS), a format often used to syndicate news stories and blog postings between websites. Both RDF and RSS are machine readable web standards for expressing metadata (data about data) but whereas RSS has a predetermined and fixed vocabulary specifically for reading news, RDF is an extensible common framework for vocabularies, and their namespaces.

Using the framework of RDF you can create an ordered list about a category of things (a namespace). For example Foaf Corp namespace which came about as a vocabulary to convert the <http://theyrule.net> project into a Semantic Web-compatible format started with the original vocabulary below:

- fc (foaf corp)
 - fc: Company
 - fc: Committee
 - fc: Board
 - fc: Member
 - fc: Stock code
 - fc: Filings

and then in June 2003 the MCC (Mapping Contemporary Capitalism) project proposed the following additions extensions:

- fc: Owns - internal, external.
- fc: Shareholders - list of shareholders, number of shares on each market, percentage of shares.
- fc: Company employs - (this is a crude category which will display multiple categories: business management, investment banking, marketing, personnel etc).
- fc: Company is funding - (this data may be unavailable but we can draw many inferences from its patchiness).

- fc: Company affiliation - company member affiliation (e.g. Gate Foundation).
- fc: Company's geographical locations.

http://chinabone.lth.bclub.org.uk/albums/flight_carto/DSCN2732.thumb.jpg

Ontology Examples

"An ontology defines the terms used to describe and represent an area of knowledge. Ontologies are used by people, databases, and applications that need to share domain information (a domain is just a specific subject area or area of knowledge, like medicine, tool manufacturing, real estate, automobile repair, financial management, etc.). Ontologies include computer-usable definitions of basic concepts in the domain and the relationships among them (note that here and throughout this document, definition is not used in the technical sense understood by logicians). They encode knowledge in a domain and also knowledge that spans domains. In this way, they make that knowledge reusable " Quote: <http://www.w3.org/TR/2003/WD-webont-req-20030203/#onto-def>

Once web content has been formatted using an RDF vocabulary from a namespace, such as !FoafCorp, then it becomes possible to infer meaning from the associations between the things it describes. To make those inferences, the Semantic Web uses 'Web Ontology Language (OWL)', a language for asking logical questions about metadata, to ask questions about the assertions in RDF documents.

A set of OWL ontology code could include a namespace, an initial set of URL's to visit and then call on a number of logical declarations. Because the semantic web deals with web content, it is inherently distributed, so expect OWL ontologies to also be distributed. One consequence of this is that OWL generally makes an 'open world assumption' allowing it to move across networks, finding new bits of RDF metadata, new assertions and new questions, and adding them to the initial ontology.

Owl would be employed in the form of a 'bot, spider or scutter', a set of code sent out onto the web to gather and interpret RDF data. For example, Ed Dumbill's 'FOAFBot' sits on an IRC channel, listening for snippets of the conversation that it is programmed to understand:

```
<edd> foafbot, edd's name
<foafbot> edd's name is 'Edd Dumbill', according to Dan Brickley,
Anon35, Niel Bornstein, Jo Walsh, Dave Beckett, Edd Dumbill,
Matt Biddulph, Paul Ford
```

The FOAFbot is invoked when edd calls its name in the IRC channel, and then responds to his command 'edd's name' by searching through the statements in the FOAF files of Dan Brickley, Anon35, Niel Bornstien etc.. and inferring from those that the nickname 'edd' refers to 'Edd Dumbill'. It can retrieve any information about edd that is available in the statements in those FOAF files, such as links to pictures of 'edd' or lists of the people that edd says he knows in his FOAF file.[3] This simple functionality can then be re-used by other bots, built on and re-purposed to create hugely complex and nuanced systems of distributed information storage and retrieval.



Free Association

Prior to the invention of the printing press, there was no such thing as an index. Books copied by hand would have different pagination, so the idea of correlating specific sections in the book with certain ideas, and collating them in an index at the back never occurred. Similarly, without standardised grammar, spellings or spacings between words, hand-written script tended to run into long, unbroken lines of letters that needed to be read out and understood aurally for meaning to emerge. The visual comprehension of words on a page without a spoken and heard intermediate stage was again, a development of the printing press. These two developments made possible access to, and use of information with formerly unimaginable speed and sophistication. The Semantic Web promises a similar acceleration and transformation in our relationship with information. The vision of computers and people, working in 'co-operation' as Berners-Lee puts it, casts aside superficial metaphors of 'pages' to be 'explored' or 'navigated' and instead suggests the Web as a growing network of prosthetic comprehension, and potentially, a treacherous one.

;"The third wave of network attacks is semantic attacks: attacks that target the way we, as humans, assign meaning to content. " -- Bruce Schneier, Semantic Attacks, The Third Wave of Network Attacks, Crypto-gram newsletter, October 2000. [5]

Although here Bruce Schneider is talking about the immanent threat of a catastrophic hacker assault on computer security systems, he could just as well be referring to the standard operation of certain search engines. Although Google currently maintains a fairly clean track-record with regards to how it indexes, ranks and displays its search results, the potential for massaging and manipulating those operations is huge. Dependence on a single system of information association, particularly an unaccountable commercial system whose ownership may change at any moment, makes our use of the Web very vulnerable to abuse.

The enclosure of a potential 'information commons' by an anarchistic elite of corporate/state bodies is well underway. Alongside this enclosure, strong and vibrant hobbyist movements are flourishing. Free Software activists, Free Hardware geeks and Free Networkers, natives of the information commons are continuing to fiddle, peeking under the bonnet of their technologies, creating and manipulating their information environment as they see fit. Bearing in mind the problematic heritage of the Semantic Web project, there is still potential in its use and development as a part of the Free Information movement. The three strands of this movement mentioned above share the !SemWeb's dubious origins, but are pursuing a difficult and tortuous course that avoids a deterministic return to authoritarian and profit-driven exploitation. As it is, these movements are disparate, unconnected, resembling the state of the Net itself; an incoherent mess of networks. Worse, the connections between these networks are almost always proprietary at some point. Downloading your Free Software, it will almost certainly be passing over a proprietary network, and somewhere in that transaction, there is a dependency on the permission and profit-margin of a corporation, a media owner, an ISP, the DNS system. You might not even have found out about the software if Google hadn't permitted it to be indexed and returned in your search results.

Without the associations and indices that allow access to information, that information is inaccessible, valueless. As the density and quality of Semantic Web meta-content grows, that meta-content will become an extremely valuable asset in itself. To protect the integrity and trustworthiness of their meta-content, Semantic Web developers and meta-content producers will need to co-operate with, and adopt similar legal defence strategies to the Free Software groups: asserting the intellectual property rights of an author to allow their works to be maintained in the public domain.

But here is the most treacherous part. Asserting intellectual property rights over associations, vocabularies, descriptions, the relationships between things in the world, as much as data on the Web, is premised on the assumption that this kind of information must be seen as property. As the Semantic Web stretches over more and more areas of knowledge production, histories, identities, interpersonal relationships, and language, this assumption feeds nauseous system of self-industrialisation and commodification, the process by which we are transforming ourselves into Fortress Europe.



Some Responses

Surely it is complete political irresponsibility to launch an attack on people who you openly admit do not really exist at a time of extreme political volatility when there are plenty of real tangible issues around preventing war, famine and environmental destruction to be confronted?

On the contrary, it would be a complete abdication of political responsibility to try and deal with these issues with the autistic attitude of neo-classical leftism. It would be a symptom of intellectual bankruptcy not to use and develop precisely the scientific developments of the twentieth century to give birth to the scientific socialism of the twenty-first.

That's as maybe, but it seems to be another product of deracinated intellectualism faced with an irreversible decomposition of class politics in Europe and North America.

We must recognise that the modern neo-liberal world is as much a product of the liberalism of classical leftism as it is of the development of corporate power. Modern corporate expansionism has used the experience of the Grand Coulee Dam Project, with all its Stalinist hangers on, in its invasion and disruption of already traumatised communities in those parts of the world only partially penetrated by capitalist relations.

Electrification, David Stone Martin, 1940

“The town that I was born in, called Mason City, was really primarily a construction site for the construction of this extraordinary dam. And after it was built, the town really. But then it was renamed Electric City at some point afterwards, and it really is not -- it certainly is not a city. It's not even a town. It's kind of a ghost town without a town. It does not exist.” Interview with Dennis Oppenheim by S. Boettger, July 12, 1995 artarchives.si.edu/oralhist/oppenh95.htm

In the U.S. the New Deal was the point at which the Bolshevik Counter-revolution met corporate America. However, there have been dramatic changes . . .

...erm aren't you going off at a tangent here?

I would only expect such a comment from someone who was used to going round in circles. I would just like to clarify what I mean by scientific socialism. This expression has rightly gained a lot of opprobrium from its connection with some of the excesses of Bolshevism, Social Democracy and whole menagerie of vultures who pick over the bleached bones Frederick Engels. This is clearly of no interest to revolutionaries. We raise the question of electricity because it is explicitly tied up with the success of the Bolshevik counter-revolution. Their control of electricity generation was essential to their defeat of the factory committees in the early months of 1918. It makes more sense to consider the Bolsheviks as implementing Francis Bacon's vision of the New Atlantis than the overthrow of bourgeois society. They were a sort Anglicanism with electricity replacing God as the central cog of their party machine. Indeed Lenin conflates materialism with acceptance of an absolute Newtonian space-time and relativity theory was condemned as much by the 'Soviet' Union as by Nazi Germany. I do not raise this to indulge in some sort of liberal speculation - the nature of Bolshevism emerged from the social tasks which embodied: the development of capitalist social relations - and in particular to move from the formal to the real extraction of surplus-value - in the specific context of the twentieth century. In these terms we can see Lysenkoism as a programme to expand productivity by manipulating environmental circumstances of plant life in order to increase the relative surplus value extracted from the peasantry as opposed to the restrictions placed on peasant consumption of the food they grew in order to increase the amount of absolute surplus value.

Glossary and Links

- Semantic Web - The Web of data with meaning in the sense that a computer program can learn enough about what the data means to process it.
- RDF - Resource Description Framework - Designed for expressing metadata about things in the form of 'triples', using vocabularies that are published on the web. see Mute Map Vocabularies (above). An introductory (business-oriented) slideshow by Tim Berners-Lee has some interesting visualisations and talks about using an 'RDF Integration Bus' like the Mute Map Infomesh for applications. <http://www.w3.org/2003/Talks/03-pcforum-tbl/slide15-4.html>
- W3C RDF primer: <http://www.w3.org/TR/rdf-primer/>
- History of RDF by Tim Bray: <http://www.tbray.org/ongoing/When/200x/2003/05/21/RDFNet>
- RSS - Rich Site Summary - An RDF vocabulary and RDF/XML format for distributing news, increasingly popular with websites, many newsreaders available for example: <http://amphetadesk.com> for windows, <http://www.netnewswire.com> for mac. There are also many RSS aggregation services like <http://syndic8.com>. Easy to write 'crawlers' and 'scrapers' can convert HTML, email, irc, nntp etc... to RSS format.
- FOAF - FriendOfAFriend - A vocabulary for describing people and networks of people in RDF. <http://rdfweb.org/foaf>.
- Friends of Corporate Friends (FOAFCorp) : <http://rdfweb.org/foaf/corp>.
- !FoafNaut, a visual tool for navigating the foaf network done in SVG : <http://foafnaut.org>.
- A bug in FOAF, explaining the difficulties of modelling groups of people: http://rdfweb.org/issues/show_bug.cgi?id=8.
- OWL - Web Ontology Language - A language (expressed in RDF) that allows us to apply logical and taxonomic constraints to RDF data and the things expressed in RDF vocabularies. Still in development.

<http://www.w3.org/TR/2002/WD-owl-guide-20021104/>

- SVG - Scalable Vector Graphics - An XML format for describing vector graphics. with SMIL and javascript it can do Flash-like things; it also does lovely scalable static images.
- An SVG orgchart demo : <http://swordfish.rdfweb.org/discovery/2003/03/6deg.svg>
- Carto.net, cartography and SVG : <http://www.carto.net> .
- SVG London tube map : <http://space.frot.org/rdf/tubemap.svg>
- XML (Extensible Markup Language) - A simplified successor to SGML. W3C's generic language for creating new markup languages. Markup languages (such as HTML) are used to represent documents with a nested, treelike structure. XML is a product of W3C and a trademark of MIT.
- Scutter, spider, bot : in the Semantic Web context this would be a set of code containing logical instructions, that is then sent to a number of URI's to apply the code to RDF data it finds at these addresses.
- Namespace - repository for Semantic Web vocabulary
- URI - Uniform Resource Identifier. The generic set of all names/addresses that are short strings that refer to resources.
- URL - Uniform Resource Locator. An informal term (no longer used in technical specifications) associated with popular URI schemes: http, ftp, mailto, etc.
- W3C (World Wide Web Consortium) - A neutral meeting point of those to whom the Web is important, with the mission of leading the Web to its full potential.

If you are interested in taking part in Semantic Web or cartography projects, you are welcome to join the University of Openess Faculty of Cartography: <http://uo.theops.net/FacultyCartography>, or find more information at one of the key resources listed below:

- RDFweb and FOAF development: <http://rdfweb.org/>
- Geowanking - An important mapping list <http://lists.burri.to/mailman/listinfo/geowanking>
- !TheMuteMap - Semantic Web/ SVG development space <http://themutemap.3d.openmute.org>
- The Locative Media Lab: <http://locative.org/>

Footnotes

[1] The Quantum Time Bomb is an expression to refer to the whole range of anomalies which will occur when a Quantum Computer is linked to the internet. Perhaps this has already happened on August 14th, when much of North America experienced a power cut.

[2] For more information about this author, see http://www.wikipedia.org/wiki/User:Harry_Potter

[3] The Voodoo Science Club and the London Psychogeographical Society Historification Committee, Friday 22nd August 2003, Announced overnight cycle trip from Limehouse, London, to the Cave of the Illuminati, Royston, Herts. <http://uo.theops.net/LimehouseTimeTravelRally>.

[4] Download source09/12/2003 code and find more information about Edd Dumbill's FOAFbot at <http://usefulinc.com/foaf/foafbot/>.

[5] See <http://www.schneier.com/crypto-gram-0010.html>.

The Epistomat: Generating Consensus About RDF Ontologies and Rules

Abstract

What happens when you take an ontology-driven interface generator, and feed it the ontology for OWL itself? A simple recursive technique allows for bootstrapping a world model and programmatically generating interfaces to it. Such a model can then serve as the platform for the evolution of collective consensus about ontologies and statements, by means of a curious game of Nomic.

Keywords

Ontology driven interface, epistemological agents, bot arbitration, knowledge representation

Contents

- [Introduction](#)
- [Generating Interfaces with RDFS and OWL](#)
- [A Game of Nomic](#)
- [A Voting Model for Ontologies and Rules](#)
- [Ontomatic and Semantic Self-defense](#)
- [Generating Consensus](#)
- [The Epistemological Congress](#)
- [Conclusions](#)

Introduction

As semantic web applications become more pragmatic, questions about ontology merging arise with more frequency. How will a user of the system be able to refine their worldview, accept certain statements from some proposed ontologies, but not others? How will they be able to collaborate and maintain their ontological commitments, which may need to change with their world view over time?

We describe a set of application layers, built over a HTTP-based RDF/XML application framework. There is a small 'ontomatic' at the core; a client interface providing awareness of RDF Schemas and a subset of OWL semantics, which is used to generate several interfaces to the application. Each of these interfaces can then be used to query and manipulate the model.

Following the example of Nomic, a game of self-amendment, the ontomatic is employed as a basis upon which groups can fashion consensus about collected statements and ontologies by means of metastatement, e.g. voting, acceptance states, and commentary. We will look at how this emergent 'epistomat' can be used to model individual and group world views, bootstrap the consensus process, resist common forms of 'semantic attack', and facilitate the generation of meaningfully networked views of the Web.

Generating Interfaces with RDFS and OWL

By definition, ontologies specify the way that people interface with information. An ontology outlines the kinds of statements we can make about objects in a particular domain, directing and constraining the ways we interact conceptually with those objects. One can easily take this a step further, and imagine building user interfaces by simply mapping form fields or prompts to the kinds of statements an ontology suggests we can make about objects of a given type. Interfaces for creating and amending statements about those objects can then programmatically emerge from ontological metadata about the RDF classes in question.

From this approach, we can easily imagine an interface to any RDF model, given an OWL ontology, or rather a selection of annotated classes and properties from different ontologies, which provide a 'view' into the model. This view can then be used to automatically generate interfaces to the model in a web browser, an 'instant message' conversational interface, a GUI description language such as XUL, or other potential targets. We will refer to the type of system that provides such a user interface to an underlying RDF model as an 'ontomatic'.

The ontomatic itself can be bootstrapped from a small 'core model': a tiny subset of RDF and RDFS semantics sufficient to describe the basic features of OWL; a kind of RDF-Tiny. An open installation with data from the wild, subject to a collaborative process, would likely want a more complete awareness of these schemas, though we may not want to see every detail in the interface, or use them all in inference augmentation attempts.

We then seed the ontomatic with knowledge of basic common ontologies - FOAF, spacenamespace, iCal, Dublin Core, ISO vocabularies for countries and world cities, as well as the McC vocabulary describing organisational networks and structures, which we are in the process of creating. We are currently using the ontomatic to manage and publish this last ontology collaboratively, using the web interface. This model will later form the basis of a 'microtheory', or 'worldview', that a person or group or bot can subscribe to, in order to filter new statements, or ask certain questions about the world.

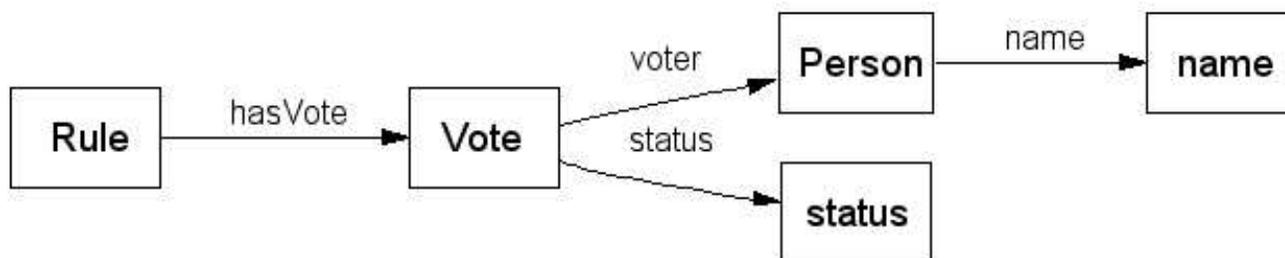
Thus we build up aspects of ontologies that we like, or that are meaningful to us, and connect them together. The same interface allows us to edit and add subject-verb-object statements, and the drop-down lists for suggestions of new properties are populated by simple reasoning on behalf of the model - for example, if a verb 'isInCountry' has a domain of iso:Country, then we should populate the list with the labels of all things of type iso:Country, and so on. Additional hints about interface styling or preferences can be provided by further annotation of the given model's ontologies - and, once bootstrapped with a core model, these embellishments can be added via the ontomatic itself.

Next, we can provide an *stateful conversational interface* to the ontomatic, by way of a generic 'ontobot', accessible via Internet Relay Chat or an Instant Message protocol. During process startup, the bot consults an annotated ontology at a URI; this can be any ontology as long as it contains an rdf:label for each class and property - this is used as short-hand for the bot. In this application, the bot consults the ontology at a URI which maps to a method on the RDF/XML HTTP API; alterations made to the model can thereby be instantly reflected in the interface to it.

as 'Pure Nomic', in which no rules are immutable, and play opens with as few rules as possibly enable a functioning system.

To this end, we collaborated on the development of a Pure Nomic bot for IRC, which keeps track of proposed rules, and allows voting and commentary on them. The object is to propose rules or emendations to rules and get them accepted; there may be a point-scoring and 'winning' threshold (as in traditional Nomic), but our original version was more of an experiment in participatory rule-based democracy amongst a small group, and a search for rules that would reinforce the continuation of the game (as Pure Nomic games tend to stabilize and 'end' after a while).

In Nomic, one votes on rules containing descriptions of the game, or of other things in the world, and which may refer to each other in ways like "this rule supersedes rule 6 and adds this coda."

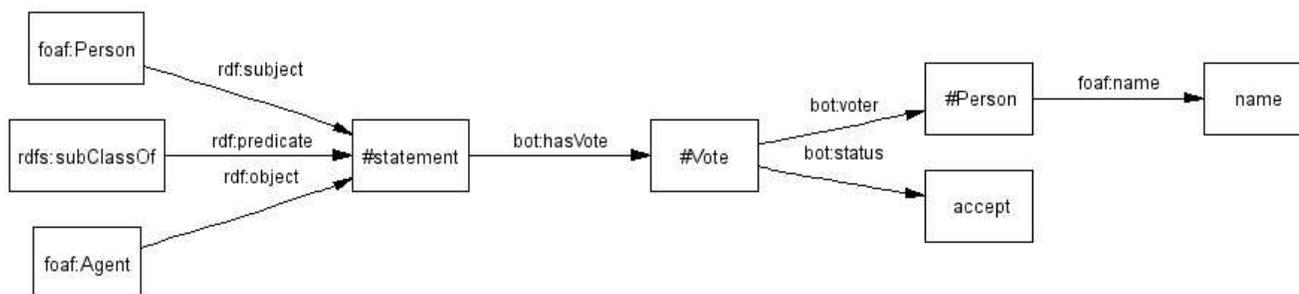


A simple graph model for a Nomic vote.

Consider, then, if we are voting not on full text rules but on statements expressed in OWL, or on inference rules expressed in a higher level language.

If we choose to vote directly on statements made in an OWL ontology, voting on Classes and their Properties, we need a way to consider them in collections, to be able to "agree with all statements with this subject" without having to vote through one by one, losing a sense of the bigger picture. We may wish to define closures based on common patterns in a model or a domain; we may wish to use the OWL schema itself to suggest entailments and implications of statements. Equally, we need to provide useful interfaces for ontology traversal and graph visualisation closely tied to this voting process.

A Voting Model for Ontologies and Rules



An RDF graph model suggesting an implementation of votes on statements using reification. The interface allows for acceptance or rejection of statements in batches, where appropriate.

Each vote within a model essentially corresponds to a statement of belief about another statement or subgraph, a confirmation or refutation of it. Our game of Nomic eventually postulated seven states which an agent can position themselves in relation to a question, ranging from complete agreement, general agreement, partial agreement, neutrality, partial disagreement, general disagreement and complete disagreement. A combination of these states can be used to compute a threshold for acceptance or refutation of a particular statement.

What happens if a statement is contained in one subgraph that you agree with, and in another that you don't? In this event, a participant would have to propose an intersection - or have a game heuristic that provided an intersection - as a new rule, and it would then be that player's responsibility to convince other players of its consistency. Ideally, these proposals should be expressed in some logical form with an unambiguous interpretation.

We can then filter our view of the ontology to only include those elements about which everyone in the group is in complete accord; or filter to only include statements confirmed by one person - essentially a microtheory reflecting a person's system of beliefs - or the intersection of any set of people, perhaps drawing those sets from enumerated FOAF networks. Thus the application is looking more like an *epistomatic* than an ontomatic.

We can also develop a voting system, based on opinions about queries as expressed in UL-1, a small RDF query and rule language developed by the authors with this application in mind. UL-1 is a second-generation RDF query language sharing some characteristics with sesame's **SeRQL** (<http://sesame.aidministrato.nl/publications/users/ch05.html>), with *composability* as a principle design aim; it should be easy to read and understand in our Nomic application. UL-1 statements are stored in the model as strings, and can be called up and 'enacted'.

UL-1 can express assertions of OWL statements, and also express 'constructs', or sub-graphs of the model. This way, we can vote on whole sub-graphs, or make refinements of them as queries expressed in UL-1. These queries are then run for demonstration in a voting interface, providing a visualisation of the model. Results are returned as a graph, usually serialised in RDF/XML.

We have not essayed an implementation, but strongly feel a Nomic voting system would be perfect for RuleML, or a similar rule language with an RDF/XML graph representation. We can run the rule subject to vote in-memory and demonstrate the new graph returned by its production. Votes which are enacted and which have entailments immediately become applied to the model. We posit being able to describe the mechanics of the game - "This statement can only have the property 'approved' if more than 75% of votes are in agreement and none in more than slight disagreement" - which would immediately propagate into the rules in play.

Immediately, we might try this as an interesting way to generate consensus on a collaborative ontology - providing some discussion mechanism through IRC, Jabber networks, email lists, web discussion forums, and so on - and as a side effect, to be able to generate a query interface into that, or any other, ontology.

Ontomation and the Semantic Self-Defense

The ontomatic allows a group of users to add and manage OWL ontologies interactively, or to import them from URIs. There is a third and interesting option; that of running a 'scutter', an RDF spidering bot, collecting triples from the wild. A near-future related option would be to accept a stream of triples from an RDF aggregator providing a scuttering service; however, these RDF aggregators provide scattered data

that may well have already been subjected to a 'smushing', or some other inference augmentation process.

Naturally, we don't want to take everything our scutter comes across on trust, and immediately attempt to augment our model with it. There are many ways a reasoner or a scutter can be subjected to 'semantic attack': by misuse of commonly implemented OWL rules, e.g. false inverseFunctionalProperties to generate false sets of joined statements; or of false subclassOf and subPropertyOf assertions, to create lots of spurious attributes in the interface; or the far more trivial attack of simply asserting a lot of factual lies or gobbledegook data in RDF.

We have several strategies for maintaining the conceptual integrity of a model, whilst opening it up to free input 'from the wild', based on the workings of the ontomatic and the epistomat. Though developed to deal with the particular pressures of a scattered data store, they apply equally to an internal-use RDF-based system with a reasoning component and human access to predicates which affect the operation of rules.

One important tactic is to keep a 'ground model' of every node and arc in an RDF graph encountered by the interface. In our implementation, the storage consists of a simple table containing every node, and a second table of triples alongside context, timestamp, and some search keys. Every statement we ever hear is recorded in it and never removed; no smushing or other normalising changes to the model based on the operation of OWL rules is contained in this model. Retractions are simply recorded in kind as a metastatement. Our candidate statements for OWL constraints and other rules that we will accept is also expressed within this 'ground' repository. It turns out this representation proves useful for versioning, and for applications which expect frequent changes in information available from many single sources, such as a FOAF network, or a news aggregator. When we want to generate an RDF model based on an ontology, or when we want to generate an ontology based on rules about it or votes regarding it, we can trivially extract a sub-graph for the present or any time in the past, and then augment the selected sub-graph with the then-accepted inference rules.

Running a scutter on top of the data store, or an 'infomesh' behind the ontomatic, we maintain two separate RDF models based on the ground store - one representing the ontologies from which the interface rules are constructed, and another containing the general 'world model' of the application. Naturally, all triples with predicates in the RDF, RDFS, or OWL namespaces, and all statements containing their subjects or objects should go in the ontological model, while all other statements are placed into one or more general models. It is these latter general models against which most application queries will be performed in the course of ordinary use.

In essence, we maintain thereby a separation comparable to that of 'mutable' versus 'immutable' rules in traditional Nomic. For an internal or critical system, such as a contacts management interface, we can even go so far as to keep the ontology of the world model completely constrained, and possibly even maintain it not in the store but in an external file at a URI. Access to instance data is managed by authentication, and there is provenance for every statement and any subsequent retraction or restatement.

In this event, we would want to be able to gather assertions made in OWL and RDFS but not statements about the OWL and RDFS vocabularies themselves - a OWL-DL-style restrictive approach to what we're prepared to infer. The new OWL statements have to be subjected to some kind of approval process before they can become real assertions that we are happy to repeat. Each individual subject URI in the model has a status, as status of terms is marked in the FOAF vocabulary, with an option to just not use or publish them.

We consider the following strategies for arbitrating RDF statements, including OWL statements which, if used to augment our model, may effect interesting changes on our world.

- We can moderate the triples, either simply via a 'super-user' role and on/off switch, or in a more nuanced and granular way, via the consensus process enacted in a game of Nomic (as it were) about the OWL ontology. The voting process can have different degrees of granularity - we posit a starting seven - and different rules about what constitutes a valid or invalid statement or rule with implications. These rules can then themselves be managed through the epistomat's interfaces.
- We could choose to accept statements about ontologies only from the URI that each Class or Property lists in its `rdfs:isDefinedBy` property. We could also choose to accept, or emphasize in an approval process, new predicates that are defined in well-known namespaces that are still under development, such as FOAF. Though this would not protect us from deliberate attack, it goes some way to avoiding harm from careless violence while exercising some liberality about what new schema information we accept.
- We could implement special handling policies for known issues that are likely to cause trouble to reasoners, in addition to otherwise being liberal or conservative. For example, new statements about `InverseFunctionalProperties` could cause quite drastic changes to our model, and could lead us to a lot of false inferences if misapplied; perhaps there should always be an approval process in front of them. A Nomic-style epistomat might require a rule dependent on such special case statements to have quantitatively more assent before being added to the model proper.
- We could look upwards one layer, and consider implementing a web of trust with a PGP-style signing system. Trust would extend to the origin of the statements - documents should be implicitly traceable to some agent. This seems to be the logical next step in the development of the Friend of a Friend ontology and its related toolset. We can envisage a connection-based trust metric like the one used on [Advogato](http://www.advogato.org/) (<http://www.advogato.org/>), which could apply directly to statement provenance, and also used to 'overlay' a completed model of votes on ontologies and rules. We can offer trust in certain kinds of RDF statements at a finer grain, offering precepts like 'trust foaf:agent X with wot:pgpKey Y to make statements (definedBy ontologies) (hasStatus status) (toModel model)'.
- One can easily imagine building this process into a scutter that traverses a FOAF network, mapping it against an arbitrary trust metric, to derive measures of more or less trusted sources of information across different domains of information.

The final step we envision involves building a partially or totally automated aggregator of public opinion, by querying distributed models via a simple HTTP-based UL-1 or other query protocol, for triples related to trust in statements, and building an ontology model out of what we can scutter thereby.

Generating Consensus

If we are to perform arbitrary acts of 'buy-in' to statements about ontologies that we find on the 'Net, or we are told by people we don't know, we need a strategy for generating consistent epistemologies within models. The need to protect ourselves from semantic attack was an important motivator for our base assumption: we are keeping one enormous 'ground' model for the ontomatic which consists of every RDF statement that has ever been asserted to us, with provenance for it, but with no local 'smushing' or inference enhancement having been performed, and use this basis to generate 'sub-models' according to different views determined by ontological commitments.

To phrase this in terms of logic programming theory, we can state that the contents of this triplestore, including the OWL meta modelling statements, and all the predicates available to any query language and reasoner we run on top of it, represent the 'Herbrand universe' of our shared world model - all the symbols, ground, variable, or functional, available to us to combine. The potential actuation of all avail-

able inference rules over the triplestore is the Herbrand base.

Of course, many or even most statements in the Herbrand base will be unprovable except by the wildest configurations. An "implementation" of the Herbrand base will be a set of ontological constraints and inference rules that when applied will express a consistent if possible not coherent worldview. A new model, augmented and compressed, generated from this implementation at a point in time, we'll refer to specifically as a "worldview", mostly to avoid conflating the logical and the RDF-related senses of the term "model".

Our game of Nomic comes in very useful now, as a source of meta-statements from which to generate candidate worldviews. We have some facility in generating different kinds of consent, which can be expressed in the rules of the game itself - rules about what percentage of votes have to be in a certain state for a rule to be accepted or rejected, for example. We can attempt to use one or several of the Nomic bot's native ideas of consensus to bootstrap the discussion - a provisional consensus worldview. By looking at the histories of individual voters, we can extrapolate a personal epistemology given the context for a query. Similarly, we can look at how different members or sub-groups of a group will cluster together over arbitrarily chosen key issues, and use this to generate sets of statements localised to the context of a given sub-group.

The Epistemological Congress

By augmenting the ontomatic with a model of consensus and principled hesitance, we can have it generate ontologies and rule sets that represent a particular world view - that of any group of one or more individuals in the system, including treating a group consensus view in the same way as an individual, allowing for a blunt federation of this system. We can use each ontology to generate a complete 'worldview' based on the ground model, possibly selecting only the view of classes and properties that it describes, and from there derive an full inference-augmented model, according to our worldview's prescriptions.

An actual storage strategy for such a system would be dependent on a less experimental use case than our own. For example, it will likely make sense to generate and incrementally maintain an augmented model for each individual, and perhaps for the general consensus according to the current rule state of the system. Meanwhile, other cross-sections of opinion could be generated afresh in response to a request, unless that request is frequent. One might liken this to the operation of indexes and views in SQL databases.

Note, we still don't have any way of *proving* that each world view will be consistent, nor do we yet have a metric for establishing 'relative' consistency between them. Our aspiration is that each worldview generated be at least sound, but probably incomplete. Further questions about correctness based on a process of consensus are absolutely fascinating in and of themselves, but more philosophical than the scope of this paper allows.

Conclusions

We believe that consensus negotiation of ontologies and rules - between business partners, or public service providers, to be briefly practical - is a critical and underexplored area. Despite the potential of topic-based trust mechanisms and collaborative reasoning systems, we believe that overt human feedback about belief is indispensable especially at this stage of the development of the semantic web.

Nomic is a way of gathering those statements which is 'fun', and indicates the potential for an 'epistomat', a small but fully self-describing system, given a sufficiently composable rule language that operates directly on the application logic. Using the ontology composed by the Nomic game, we can try out different 'perceptions' - we hesitate to say 'belief systems' - represented as differently augmented graphs inferred from the same underlying base graph.

We also offer the consensus system provided by Nomic as an extra means of protecting a knowledge base from semantic attack, either from bad instance data or bad rules. We are very curious to see what graphs and tendencies emerge as these expressions of group consensus are aggregated both within and across larger organizations, and then across the Web at large.

Acknowledgements

The authors would like to thank the members of the #bots community on irc.perl.org for their support, participation, and feedback.

References

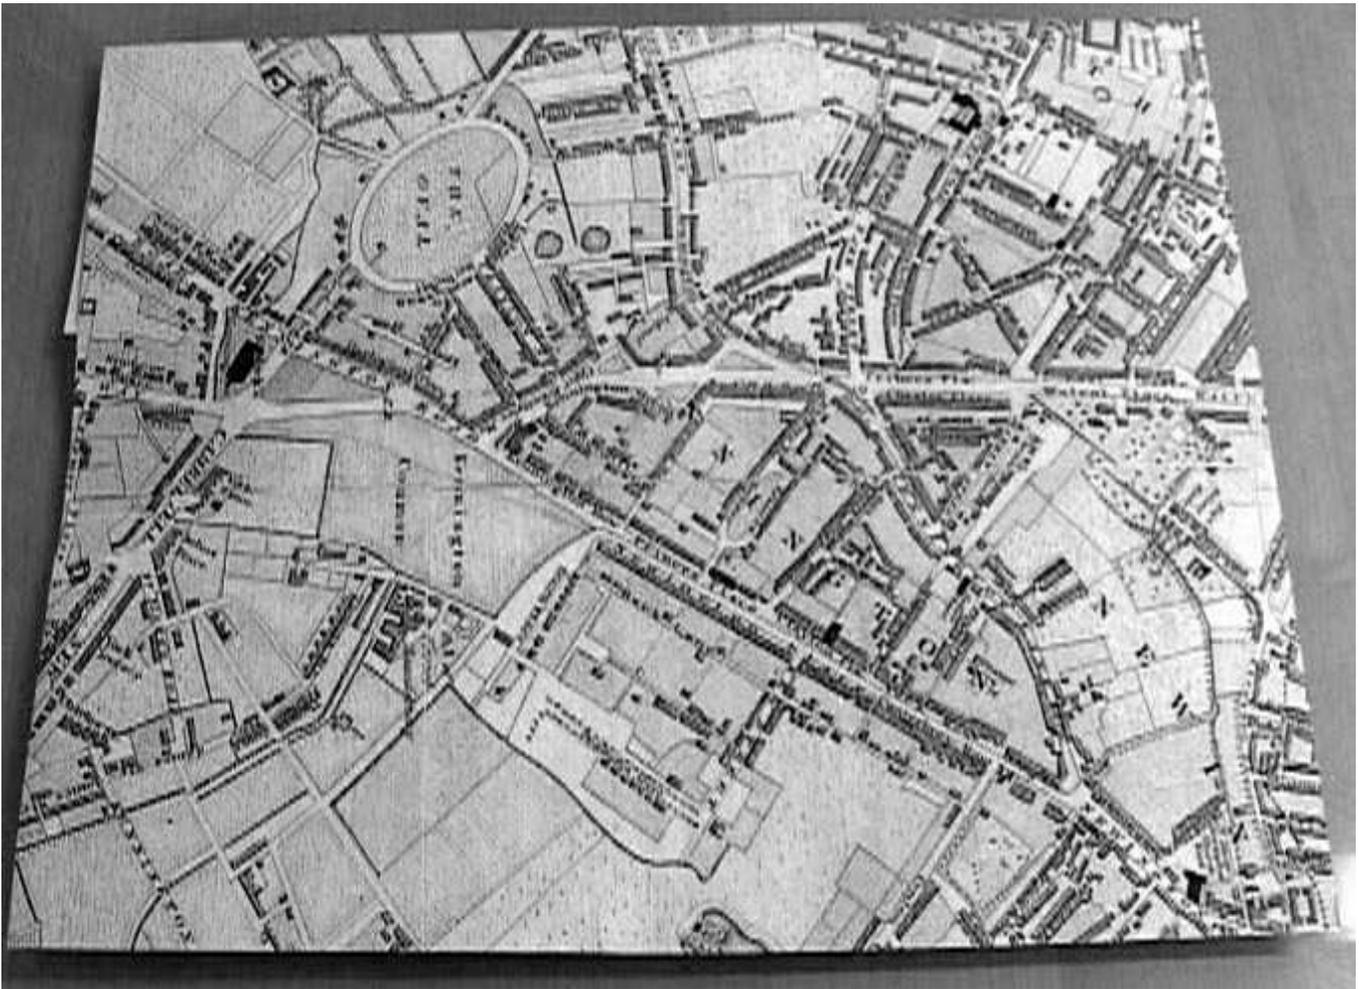
- [An RDF/XML web API](http://mutemap.openmute.org/doc/infomesh-paper.html) (http://mutemap.openmute.org/doc/infomesh-paper.html) for an ontomatic.
- An [ontological model that powers a registration / social documentation interface](http://mutemap.openmute.org/dmz_ont.rdf) (http://mutemap.openmute.org/dmz_ont.rdf), including members of the FOAF and Geo vocabularies.
- A [basic RDF model](http://mutemap.openmute.org/onto.rdf) (http://mutemap.openmute.org/onto.rdf) used to bootstrap the ontomatic.
- The [Joseki RDF web API](http://www.joseki.org/protocol.html) (http://www.joseki.org/protocol.html)
- [Squish2SQL](http://rdfweb.org/2002/02/java/squish2sql/intro.html) (http://rdfweb.org/2002/02/java/squish2sql/intro.html), in Libby Miller's [Inkling](http://swordfish.rdfweb.org/rdfquery/) (http://swordfish.rdfweb.org/rdfquery/)
- [Redland](http://www.redland.opensource.ac.uk/) (http://www.redland.opensource.ac.uk/), an RDF application framework
- [Nomic](http://www.earlham.edu/%7Epeters/nomic.htm) (http://www.earlham.edu/%7Epeters/nomic.htm), a game of self-amendment, by Peter Suber
- [SeRQL](http://sesame.aidministrator.nl/publications/users/ch05.html) (http://sesame.aidministrator.nl/publications/users/ch05.html), a query language for sesame, an RDF Schema-based repository and querying facility.
- [Advogato](http://www.advogato.org/) (http://www.advogato.org/), an Open Source advocacy and discussion site.

Lucrative Media: ownership and authorship of public space.



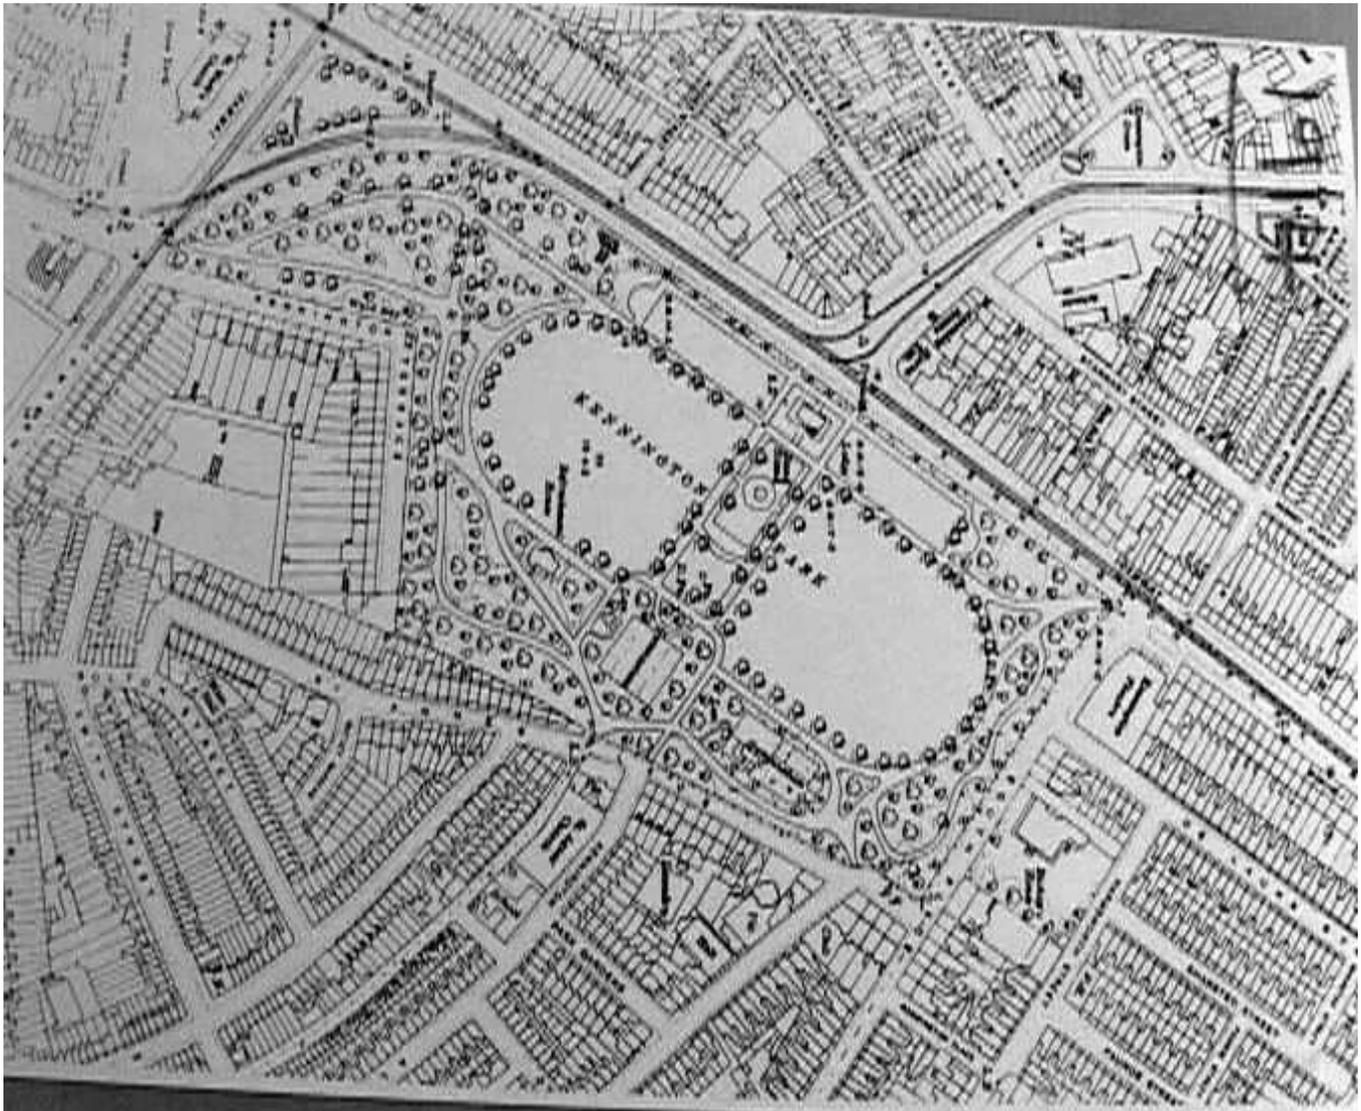
The Great Chartist's Rally on Kennington Common, 10th April 1848

During the first Cartographic Congress, London 2003, Stefan Szczelkun gave a socio-historical tour of Kennington Park ex-common: a site of popular protest by Chartists, Republicans, Methodists and dissenting groups of all kinds since before the Norman conquest of Britain.



Pre-1854 map of Kennington Common, from the collection of Stefan Szczelkun.

This map of the park pre-enclosure, in the early 19th Century shows an ad-hoc boundary, crossed by a road, asymmetrical; a coloured-in drawing of the negative spaces between private houses, roads and public thoroughfares. Non-commodified, non-landscaped, open to use and interpretation by anyone, the park was identified as a threat to authority, to be brought under spatial and symbolic control.



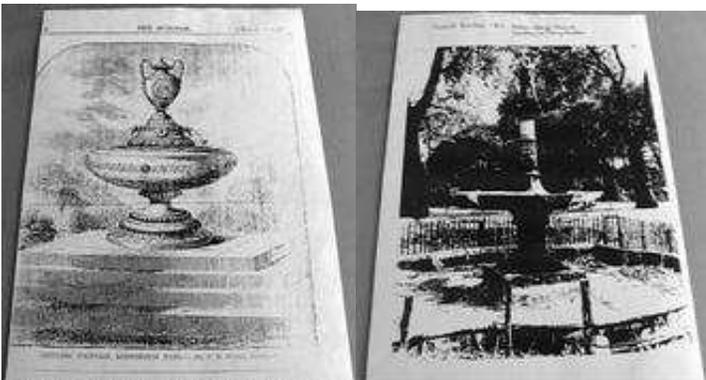
1914 Map of Kennington Royal Park, post-enclosure, from the collection of Stefan Szczelkun.

In 1852 the King enclosed the common. It became a Royal Park, patrolled by royal guards. The park was 'landscaped' and trees were planted in straight lines forming a border, with victorian monuments and fountains scattered around, seemingly at random. What was formerly a negative space of non-ownership between roads and houses became an aesthetized space for specified leisure experiences.



Kennington Park today, taken during the Cartographic Congress tour of Kennington Park ex-common.

The park is not so different today. Legislation such as the Criminal Justice Act 1998 was designed to outlaw gatherings of more than 10 people, especially where music is involved, and more recent anti-terror or 'public nuisance' bills propose even more restricted use of these notionally 'public' spaces.



Kennington Park's now-dilapidated 'Victoriana'; monuments to nothing, from Stefan Szczelkun's collection.

Stefan Sczcelkun's detailed research (recorded and transcribed here: <http://bak.spc.org/kenningtonpark/>) made a compelling hypothesis, that the victorian monuments scattered through the park were not, in fact, scattered at random. Rather, he suggested, these peculiarly un-attributed objects, without plaques or dedications were placed strategically a short distance from sites of particular historical significance or psychic intensity. Like the magician's trick to distract the audience's attention by flourishing his wand while pulling the rabbit out of his trousers with the other hand, anonymous chunks of victoriana, now fallen into disrepair, were placed just far enough from the sites of public executions and mass rallies to misdirect attention and focus from those emotive and resonant sites.

Observing the changes on the map, the modifications of historical use and interpretation of the space only begins to describe what the enclosure of Kennington Park ex-common meant. This continual process of privatisation and commodification of public space was not only about spatial control. The enclosure was as much an attack on the interpretation of that space; the ability of the people who used the park to collectively imagine new uses for it and stories about it.

It is this act, the enclosure and commodification of the interpretation of public space that many artists have identified and tried to challenge when working with concepts of 'locative media'. Locative media, as it is often described, engages and represents our interpretations of spaces, often using cartography, or technologies such as Global Positioning Systems and other location-sensing devices to augment the use of digital media with an indexical link to a location. This practice is also often coupled to the use of new 'open' licences, through which these media can be shared 'freely' in the public domain. As the argument goes, we may not be able to directly change the material relations of the ownership and control of space, immediately returning it to the status of a 'commons', but if the imagination and interpretation of the space can be made public, perhaps collective action, and even change of the material ownership will follow.

The use of these 'open' licences for 'locative media' and associated practices follow the development of Free and Open Source Software as a parallel development model to proprietary software production. There have been many attempts to construct similar licencing agreements for other areas of cultural production. The foremost of these licences is the *Creative Commons* (<http://creativecommons.org>) (<http://creativecommons.org/>), which offers licence users a range of options as to which rights they wish to reserve for the re-use of their work: that they require attribution, insist on solely non-commercial re-use, or that the user shares derivative works under the same licence.

The proposition made by many locative media projects, that interpretations, annotations and mediated subjective experiences of a location can create a new public domain of information-augmented space is compelling. However, authorship and ownership may be synonymous in terms of intellectual property, but not so spatially. When we are talking about a 'commons' of intellectual property, a 'creative commons', does it mean the same thing as when talk about the history, use and interpretation of somewhere like Kennington park?

Clearly not. The *Creative Commons* are a reaction to antiquated legal definitions of intellectual property, providing more sophisticated and complex descriptions to enable 'humans' (via lawyers and machines), to stake their claims to their informational products more efficiently. On the net, where the informational products of a million authors are exchanged, re-used and transformed in a tumultuous stream of value, these subtle legal descriptions are a way of facilitating financial or reputation-based remuneration for immaterial labour. The 'commons' on the other hand, are a space of non-ownership, free for use and interpretation by all.

It is appropriate that this acceleration and transformation of value should be initiated by the 'creative' industries. In the recent history of urban regeneration, the use of artists as the vanguard of gentrification is now a well established process¹. Investigation into the affective labour and enthusiastic self-exploitation of artists by Anthony Davies and Simon Ford describes the artist as a willing scab on the picket lines of precarious, de-unionised information workers.² Worse, this facility for self-exploitation is often developed and honed through the artist's mediation and representation of the 'public' with which they purport to engage, and in many cases, simply use as objects in some self-aggrandising aesthetic arrangement. Those few that are conscious enough to recognise their role in this grotesque process seek impunity by applying 'creative commons' or related licences to their artworks, an act which in itself has no bearing on the process of value creation and exchange in the art and property markets.

The *Creative Commons* licences may introduce greater efficiency in describing ownership of intellectual property and in the regulation of immaterial labour, but while opening up markets and proposing matrices of accounting for the slightest transactions of affect, speech and thought, they do nothing to address the labour relations in which that property is produced. For example, these licences make the assumption that all exchanges of intellectual property should be mediated by the threat of litigation, while ignoring the fact that most people either have no recourse to a fair legal system, or are understandably reluctant to engage with legal processes that in most cases dispense justice for the wealthy and influential. Untroubled by such social or political concerns, the *Creative Commons* mundanely facilitates the limitless expansion of value creation into all spheres of human activity.

This conception of a 'creative commons', spatialised by 'locative media' seems less an aspirational return to the free interpretation and imagination of public spaces, than an ever more precise delineation of their ownership and authorship: an endless series of micro-enclosures, minute plots of intellectual territory, staked out by law.

¹See David Panos, *Create Creative Clusters*, Mute 28, Summer/Autumn 2004. <http://www.meta-mute.com>

²See Anthony Davies and Simon Ford, *The Surge to Merge Culture with Economy*, Copenhagen Free University, 29 September 2001 - <http://www.copenhagenfreeuniversity.dk/AD01.html> (<http://www.copenhagenfreeuniversity.dk/AD01.html>) (05/12/04)

Credits

This journal was assembled by researchers from the University of Openess Faculty of Cartography over two days in response to talks and presentations given at Verbindingen / Jonctions 8, 3-19 Dec. 2004, Brussels organised wonderfully by www.constantvzw.com/vj8

The booklet layout was generated from html sources exclusively using the following free software: pstools, html2ps and ghostscript.