



Peter Arbuthnott, Quakr, *peter@quakr.co.uk*
Katie Portwin, Quakr, *katie@quakr.co.uk*
David Sant, Quakr, *dave@quakr.co.uk*

Building the World, one photo at a time

Abstract

Is it possible to build a virtual world from user-contributed photographic metadata?

The Giants' Approach: control everything.

Google Earth and Microsoft's Virtual Earth have 3D elements, but they have both chosen to control the data that creates those elements. Microsoft's photo-realistic cities are 'built' from especially commissioned photos.

An Alternative: we the people.

But the data is already out there. Amateur photographers have readily adopted the concept of location metadata; there are currently 0.5 million geotagged photographs on Flickr. Meanwhile, digital cameras embed Exif metadata (some already include GPS data), and geo-standards including KML and GML are in widespread use.

The view from 30,000 feet.

Where are the overlaps between what the standards advocate, what cameras can provide, and what people actually do in practice?

Data (1): wish list.

Is the data sufficient? What data does a 3D rendering engine ideally need to position photos in 3D space? Are the geo-standards sufficient to describe that data set? If not, how might they be extended? Could users realistically contribute that data set? What tools would help them?

Data (2): making do.

Finally, to what extent can an open-world-minded application make sense of the sparse, ambiguous data which is available right now?

Our application.

This paper describes the experience of developing a 3D virtual world, based on publicly available images and existing geo-metadata, including Flickr. We make recommendations for extending geo-metadata standards to facilitate such applications, and we document our D.I.Y. efforts at constructing a geo-camera of the future, capable of capturing the metadata we advocate.

Our application sits in the heart of the ubiquitous web, at the intersection between increasingly sophisticated gadgets, metadata standards, and the surprising extent to which users are interested in both.

Part 1: Laying the foundations

The Giants' Approach

There are currently two high profile projects aiming to provide a 3D world experience: Google Earth and Microsoft Virtual Earth. 3D objects in Google Earth are built with manual labour; Google SketchUp enables users to "build and modify 3D models quickly and easily". Key landmarks have been prioritised and built by Google.

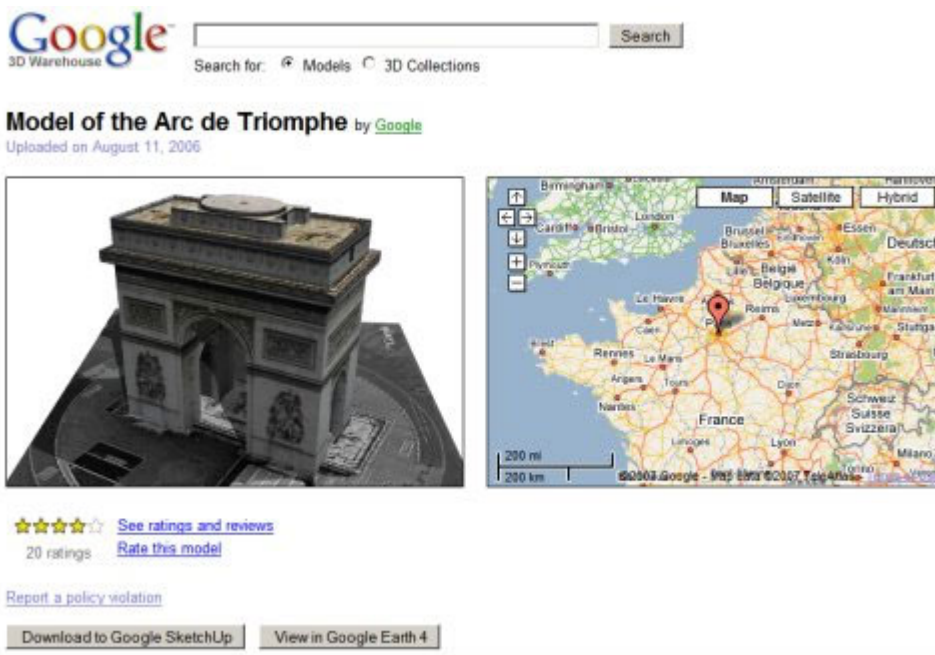


Figure 1. Google 3D Warehouse model for Arc de Triomphe.

For Virtual Earth, Microsoft have gone a step further and retained complete control over construction. Their world is built from photographs gathered by photographers, commissioned “to take millions of pictures of urban landscapes from planes, vans and motorbikes” [Swinford06]. These photographs are then stitched together and photogrammetry algorithms used to determine the structure of the 3D world.



Figure 2. A screenshot of Microsoft Virtual Earth.

Both of these approaches suffer from several drawbacks. First, in the Microsoft approach, one organisation is both data provider and data consumer, isolated from other data which might help them build their world. Second, the geographical coverage of the 3D elements in each application is limited. At the time of writing, whilst Denver may

look fabulous [Weinberg07], Oxford, UK has no 3D representation in either world. Finally, these applications have a significant installation overhead, which includes the .NET framework for Virtual Earth.

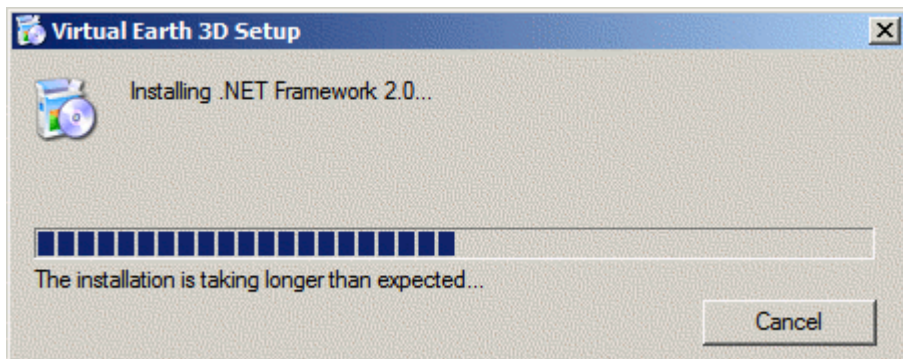


Figure 3. An example of the installation problems encountered.

An Alternative

Geotagging

Meanwhile, web users are rapidly creating their own records of what is located where. They are doing so by geotagging their photos, blogs, and even sound recordings. With the use of a GPS device, users can specify the exact latitude and longitude which best describes their resource, and (perhaps surprisingly) many people do. Flickr, which is currently one of the leading photo sharing and archiving websites, contains over 475,000 photos tagged with `geo:lat`. Furthermore, since the launch of the Flickr Map in August 2006, users have been able to geolocate their photos without using a GPS device. At the time of writing the map contains almost 14 million geolocated photos.

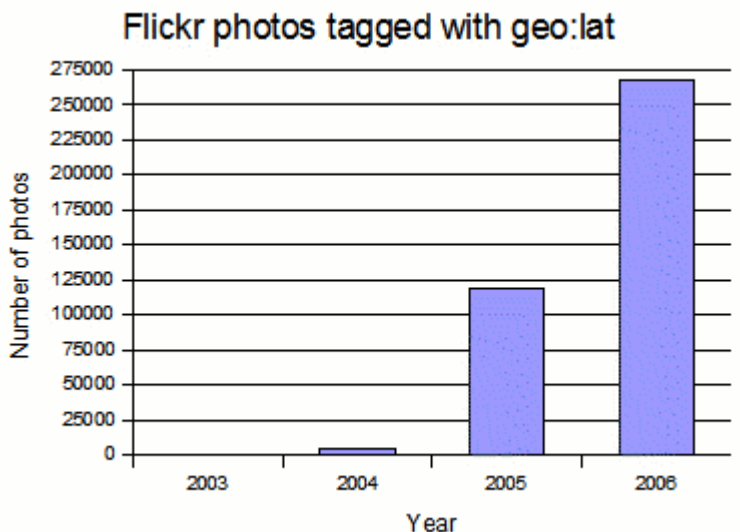


Figure 4. The growth of photographs on Flickr tagged with `geo:lat`. (Source of data: Flickr date-limited searches)

This ever-growing body of photographs and metadata is made available via the Flickr API. Flickr is not the sole data provider in this field. As of April 2007, ProgrammableWeb lists 17 photo-related APIs.

Building a navigable 3D world from user contributed photographs is therefore an exciting possibility.

Flat Mashups

The enthusiasm for photo-related mashups is second only to those based on mapping[ProgrammableWeb07], and at the time of writing, at least 84 fall into both categories. Many use the Google Maps API and Flickr API to

position photo thumbnails or markers onto a map, (e.g. loc.alize.us, Zoomr) or Google Earth, (e.g. Panoramio, FlickrFly). A recurring problem is that by presenting significantly different images side by side, the mashups lose the ability to represent what it is *actually like* to be at that point. Thus, thumbnails of photographs taken from the top of the Empire State Building looking down are displayed in exactly the same way as photographs taken from the bottom looking up.



Figure 5. A demonstration of the limitations of typical mashups. (Parte de vest a Manhattanului cu sediul ONU.. © gigi4791 / cc by-nc-sa/2.0 and Edificio del Estado del Impero © júbilo-haku / cc by-nc-nd/2.0)

To know what it is like to *be* at that point in 3D space, it is necessary to position the surrounding photos in 3D space in a meaningful way. Quakr is an attempt to do just that.

Metadata for photos in 3D

So, what photographic metadata would be necessary to achieve this ambition? Is it measurable by the average photographer?

Metadata Wish List

What metadata is necessary? To establish a photograph at a point in space, the latitude and longitude must be supplemented with altitude.



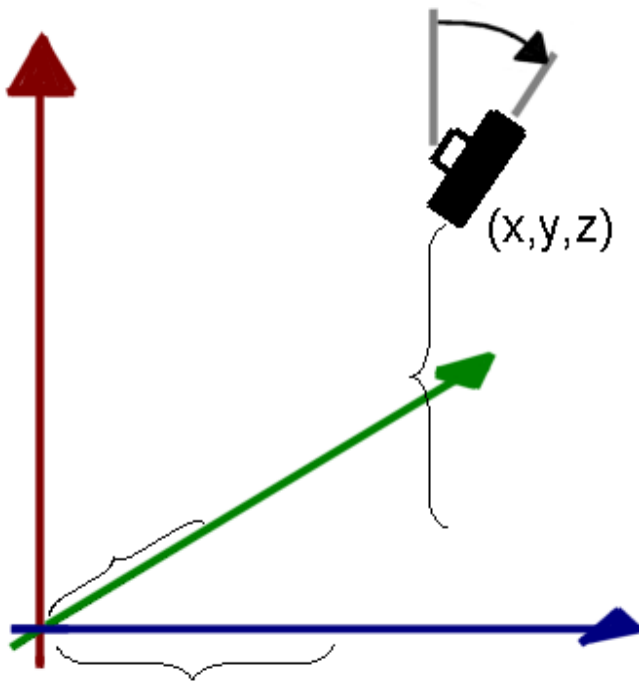


Figure 6. Positioning the photograph in space.

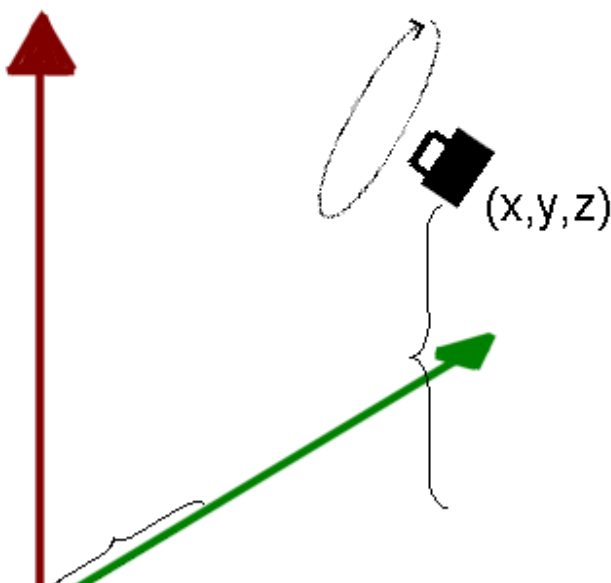
A further vector is required to establish the direction in which the camera is pointing. This vector could be described in a number of ways, for example with reference to another point in space (x_2, y_2, z_2) . In practice it may be more feasible to measure the vector as a result of three rotations around the three axes.



latitude
 longitude
 altitude
 rotation(y) – NSEW?



latitude
 longitude
 altitude
 rotation(y) – NSEW?
 rotation(z) - tilt?



latitude geo:lat
 longitude geo:long
 altitude geo:alt
 rotation(y) geo:dir ge:heading
 rotation(z) ge:tilt
 rotation(x) ge:roll

Figure 7. The three rotations required to describe the photograph's vector.

The rotations are commonly known in flight dynamics as pitch, roll and yaw (rotations around the z, x and y axes, respectively).




Figure 8. Flight dynamics representation of the three rotations. (Image:Flight dynamics.jpg © ZeroOne / Wikipedia / cc by/ 2.5)

Metadata and Practical Considerations

Is the metadata measurable by the average photographer? In common language, the rotations are described as the answers to three questions:

1. What is the compass reading for the direction of the camera?
2. How tilted up to the sky or down to the ground is the camera?
3. Is the camera rotated to be portrait, landscape, or is it some other rotation?

Metadata and Motivation

Is it reasonable to expect the users to measure and store this metadata? Why would a user start tagging with `geo:dir`? Why do users tag at all? Four basic motivations for tagging can be identified (there may be more):

1. Personal benefit – users tag to organise their data for later retrieval, (e.g. “birthday”, “cape coast”).
2. Collaboration – users tag to guide others, perhaps based on an emergent taxonomy, (e.g. “xtech2007”, “nikon”).
3. Applications – users tag to see their resource being used, perhaps using machine tags, (e.g. `geo:lat=51.75425500`).
4. Automatic – a device creates metadata automatically whether the user is interested or not (e.g. Exif information)

The complex and structured tagging necessary to build a 3D world arises from categories 3 and 4 and is already at a useful level.

Metadata: making do

In reality, there are often gaps in the metadata wish-list. This need not be fatal to the use of a photo in the 3D world. At the time of writing, the Quakr order of preference for metadata is: latitude and longitude, compass bearing, tilt, roll, altitude, zoom, date, and then all other tagging.

The more relevant metadata is present, the better the representation of the photograph in the Quakr 3D world. An application which employs this philosophy of *tolerant encouragement* will improve its performance as users become more interested, camera technology becomes more advanced and metadata quality increases.

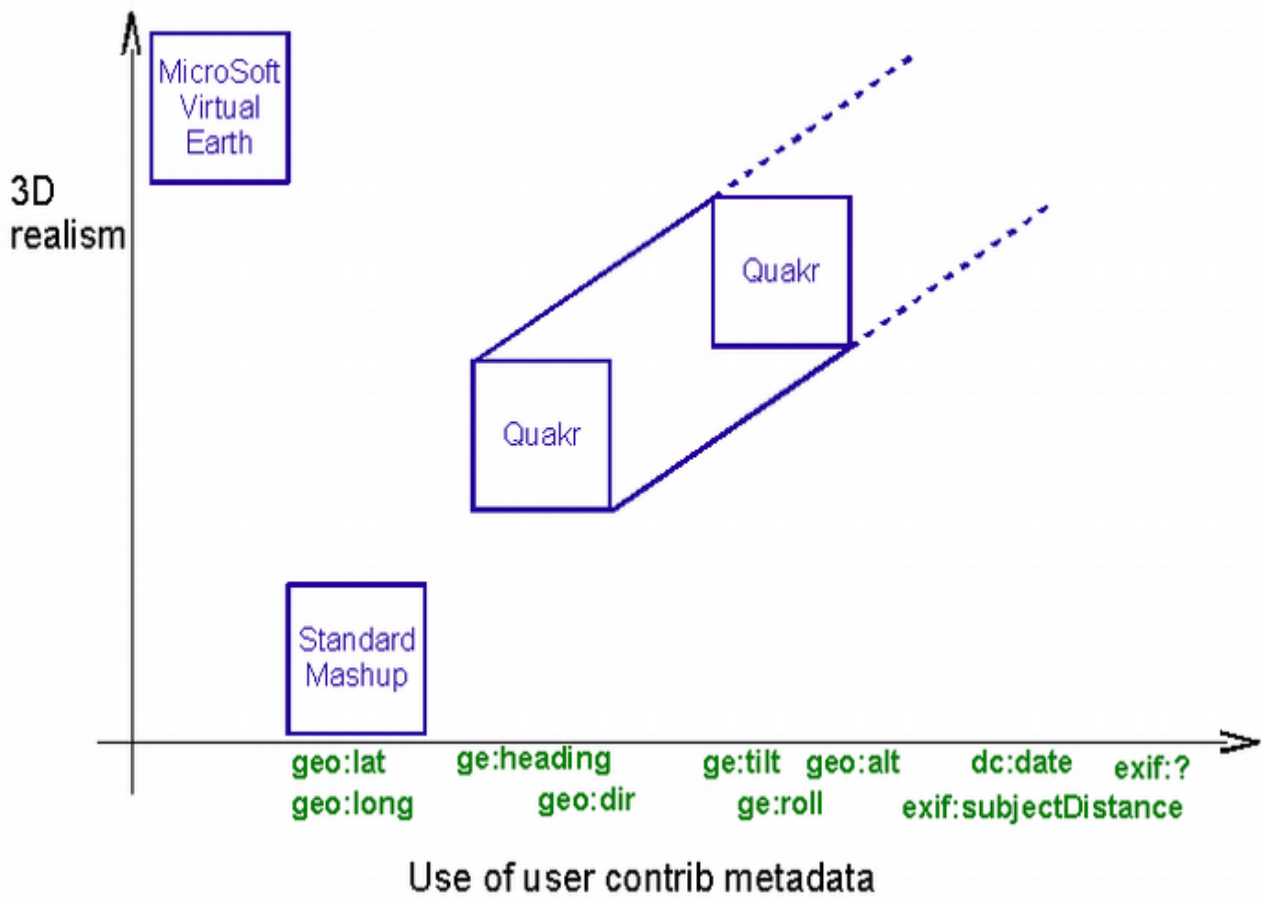


Figure 9. Use of user contributed metadata and 3D realism.

Aims, Architecture and Constraints

Quakr consists of a client consuming multiple XML resources and combining them in such a way to present a novel experience to the user. This is made possible via a Web 2.0 architecture.

The following five constraints were defined:

1. Quakr should run in a browser.
2. Quakr should require little or (preferably) no installation.
3. Quakr should be a thin client, consuming raw resources.
4. Quakr should add context information to enhance the 3D experience.
5. Quakr should adapt to an ever changing semantic web.

Part 2: Bricks and Mortar

Quakr development effort focused on three areas: generating suitable test data, choosing an appropriate taxonomy, and building the Quakr application.

Generating test data

Existing Capture Devices

As discussed above, the metadata wish list includes first, the standard geotagging measurements (latitude, longitude and altitude), and second, the three rotational measurements (compass bearing, tilt, roll).

The former can be captured with specialist GPS units or on-board GPS integrated into other devices such as mobile phones – several are available on the consumer market at the time of writing, e.g. the HP 6945 Phone [GPSPHones07].

The latter set present a greater challenge. However, the new Nintendo Wii [Nintendo07] controller purports to know which way up it is, as does the Sony PS3 [Sony07] SIXAXIS controller. The Apple MacBook [Apple07] has an accelerometer which detects "freefall". A recently released Nokia phone has "tilt" [Nokia07] built in and an alternative has an integral digital compass [Nokia07b]. These are all interesting developments; a device capable of automatic capture of rich location metadata suitable for Quakr may shortly enter the mass consumer market.

In the interim, the Quakr team built a "D.I.Y." solution.

D.I.Y. Efforts



Figure 10. The 7D Tiltometer – Protractors and Plumb Bobs.

The internally named "7D-Tiltometer" is a low-tech implementation which satisfies the Quakr requirements for a passive capture device. It consists of an HP R927 camera, a hand-held GPS device, 2 protractors, a magnetic compass, string, adhesive putty and double-sided-sticky-tape.



Figure 11. Attaching the GPS Device.

Photographs were first taken in Broad Street, Oxford and Cornmarket Street, Oxford, with one project participant operating the camera, and a second participant logging data from the attached devices.



Figure 12. Using the 7D Tiltometer.

Despite the difficulties inherent in utilising the 7D-Tiltometer, the results were effective.

Metadata Storage

The test data set of captured photographs was transferred to Flickr to make it available to the semantic web. Each was tagged with altitude, latitude, longitude, compass bearing, tilt and roll. As discussed elsewhere in this paper, there are a number of emergent taxonomies in this field. As an open-world-minded application, Quakr supports many of these as listed in the table below, and the list will increase over time.

Table 1. Preferred and supported tagging .

| Dimension | preferred | also support | acknowledged but not support |
|-----------------|-----------|--|------------------------------|
| altitude | geo:alt | | |
| latitude | geo:lat | | |
| longitude | geo:long | geo:lon | |
| compass bearing | geo:dir | ge:head ge:heading kml:heading kml:dir exif:GPSImgDirection quakr:compass | |

| | | | |
|-----------------------|---------------|---------------------|------------------|
| tilt | ge:tilt | kml:tilt quakr:tilt | |
| roll (orientation) | ge:roll | kml:roll quakr:roll | exif:orientation |
| timestamp | exif:dateTime | dc:date | |
| zoom | | quakr:zoom | |

1. Use the semantically correct term
2. Use whatever everyone else is using

For example, for altitude, latitude and longitude, `geo:` satisfies both of these principles. However for tilt, the situation is less satisfactory: `ge:tilt`, based on KML [GE07], explicitly denotes an application, not a concept. However, the size of the existing data set and user group justified the adoption of `ge:tilt` for the test data set.

Building the “Quakr Viewr” Application

XML APIs

The initial API call retrieves all the metadata for photos within a “Bounding Box” [Flickr06] defined in latitude and longitude.

A truncated example is shown below;

```
Request: http://api.flickr.com/services/rest/?method=flickr.photos.search
&api_key=xxx&min_taken_date=1970-01-01 00:00:00&bbox=
-1.3,51.754,-1.2,51.755&tags=quakr&extras=machine_tags
```

Response:

```
<?xml version="1.0" encoding="utf-8" ?>
<rsp stat="ok">
  <photos page="1" pages="1" perpage="250" total="23">
    <photo id="369020058" owner="85097477@N00" secret="9782e09dd0"
      server="167" farm="1" title="DSC02762" ispublic="1"
      isfriend="0" isfamily="0" machine_tags="" />
    <photo id="364490136" owner="56861038@N00" secret="bcf40fbd4c"
      server="167" farm="1" title="Bodlean roof, Broad Street"
      ispublic="1" isfriend="0" isfamily="0"
      machine_tags="geo:lon=125490000 geo:long=125490000
      geo:lat=5175475000 geo:alt=71 geo:heading=150 ge:roll=0
      ge:tilt=120" />
    <photo id="364490580" owner="56861038@N00" secret="426a5c9ec2"
      server="164" farm="1" title="martyr's cross, broad street"
      ispublic="1" isfriend="0" isfamily="0"
      machine_tags="geo:lon=125771667 geo:long=125771667
      geo:lat=5175430000 geo:alt=70 geo:heading=260 ge:roll=0
      ge:tilt=0" />
    ...
    <photo id="337675634" owner="47685763@N00" secret="5e6f23a168"
      server="133" farm="1" title="Balliol gate" ispublic="1"
      isfriend="0" isfamily="0" machine_tags="" />
  </photos>
</rsp>
```

Implementation investigations

Bearing in mind the first constraint (“Quakr should run in a browser”) several browser extensions were considered.

Initial investigations evaluated the use of a Java Applet. An alpha demonstration identified a fundamental problem. Creating a 3D world using a Java Applet required a deep knowledge of 3D Maths or a well-documented open source library, neither of which were available to the project team. Furthermore, it would be necessary to build an interface to allow users to move around the world, as these do not come as standard with the JAI graphics library.

The next implementation solution investigated was VRML/X3D. Displaying X3D or VRML content in a standard browser requires a plugin, which proved to be problematic to install – breaking the second constraint. There is not, at this time, a stable, browser-compatible plugin for this technology.

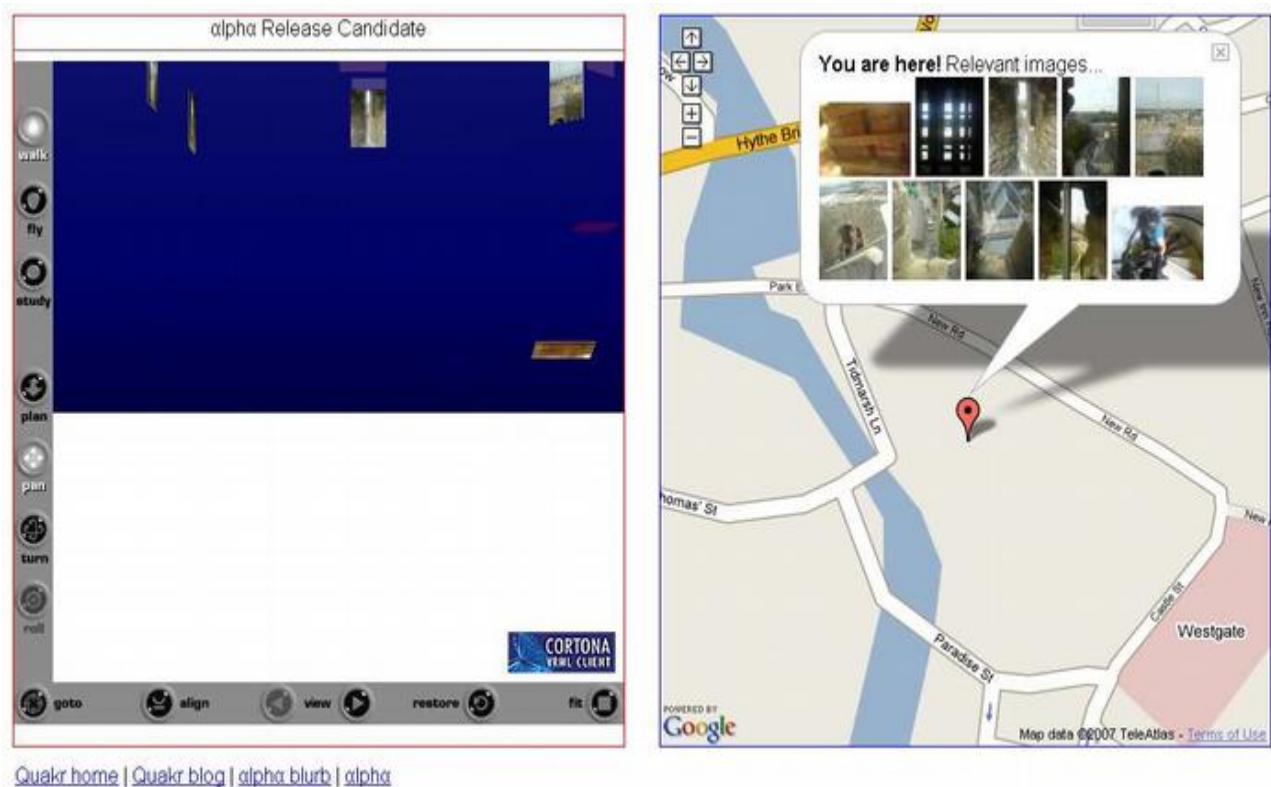


Figure 13. An Alpha Quakr Viewr in VRML Plugin form (the Google Maps display is for guidance) – 3D is present, but the plugin is unstable and has display problems.

Flash is a well-supported and well used application development environment, primarily focussed on delivering to the browser. Whilst Flash represented an effective replacement for the Java Applet, it still required 3D Maths coding and the creation of a navigation interface. The Sandy Flash 3D library, an open source library developed in Actionscript 2.0, was identified as a potential solution.

Flash and Sandy met the five constraints. A Quakr Viewr built in Flash would run in a browser, in a plugin with a minimum installation overhead. Actionscript allows communication with XML APIs, enabling Quakr to operate as a thin client, be extensible and provide the capability to integrate contextual information from a variety of sources.

Application Development

The team reviewed additional sources of contextual information to enhance the virtual world. The Google Maps API presented the most obvious source of mapping data, with other mapping applications (e.g. Yahoo!, MultiMap) also meriting consideration. In order to return the actual Google Map "tiles" relevant for any latitude and longitude request, an amount of algorithmic interpretation was required.

The application was coded in Actionscript. The following screen captures demonstrate the Alpha Release 0.21 version of Quakr. A live demo can be viewed at <http://viewr.quakr.co.uk/>

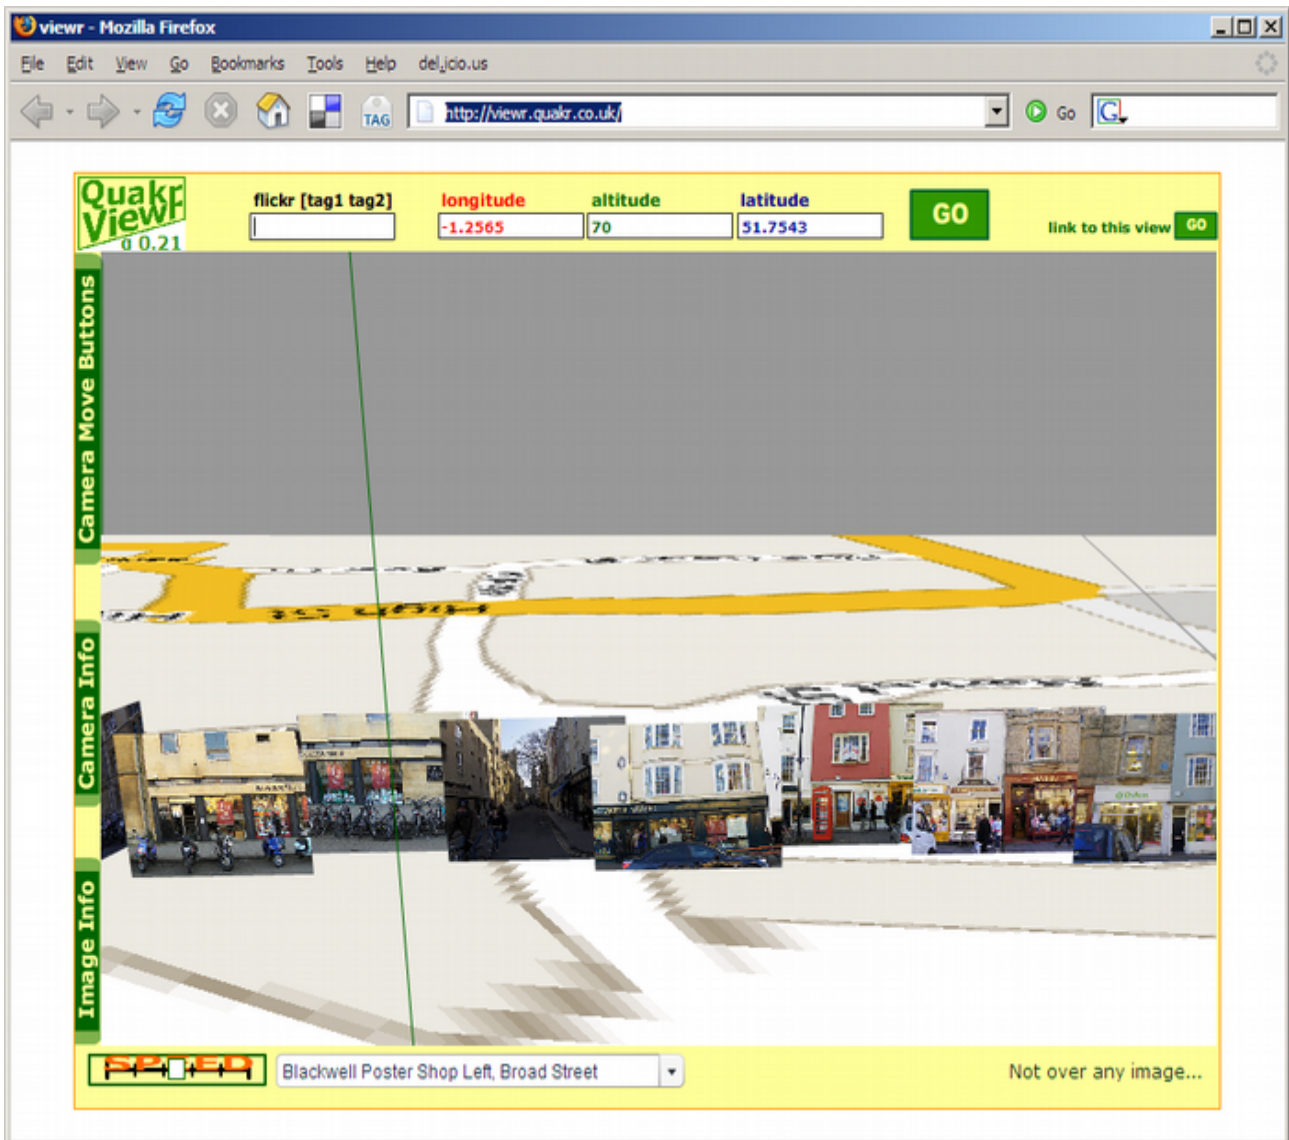


Figure 14. The Quakr Viewr Application – alpha0.21 – looking across Broad Street, Oxford (Google Maps is integrated).



Figure 15. The Quakr View Application – alpha0.21 – looking down Cornmarket Street, Oxford.

Part 3: What is a “Brick”, anyway?

Quakr exhibits certain metadata problems which are associated with semantic web applications in general, with those which consume tagging-based, community-contributed metadata in particular. The final part of this paper details this approach, and examines some of the repercussions.

Philosophy of data tolerance

Emergent taxonomies

Quakr relies on metadata created by tagging. Users type words into a free-text field.

“Tagging: describing web content using whatever words seem right” [Beckett06]

The advantage of tagging is that it is quick, simple and accessible to non-experts, and so results in a great volume of metadata. In the past year, Flickr users have added a million tags each week [Dubinko06].

However, the flexibility of tagging is also a weakness. By contrast with a formal taxonomy, tags are 'coarse-grained, informal and vague' [Wu06].

“People will in general use the minimum amount of convention to solve their co-ordination problem.”
[Halpin06]

The tagging versus taxonomies debate has been described as “evolution versus intelligent design” [WWW06]. Meaning created by tagging is “evolutionary”, “organic”. Naming conventions may exhibit a degree of “crystallization” [Cattuto06], taxonomies may “emerge” [Merholz04]. But clients which consume tags must be able to withstand unhygienic and evolving metadata.

Data Reuse

Quakr relies on metadata created for other purposes. When creating geotags some users were aiming at other applications, some may have been guided by standard vocabularies. However, by focusing on the areas of overlap and agreement between ontologies, Quakr is able to adopt existing metadata.

This data-tolerance approach can be summarised with two maxims:

1. Tolerate unhygienic data
2. Recycle everything

Costs of data-tolerance

This philosophy also has its costs. Applications which follow it suffer from metadata issues including semantic ambiguity, sparseness and accuracy. Some of these issues are examined below, with practical examples from Quakr.

Synonyms and Homonyms

An emerging taxonomy rarely exhibits a one-to-one mapping between terms and meanings. Multiple terms exist for the same concept, and vice versa. In the case of `ge:heading` and `geo:dir` (both meaning, approximately, compass bearing) synonyms exist in spite of users' valiant adherence to emerging taxonomies, probably because of the the presence of two competing published standards – KML [GE07] and GEO [W3C03]. On Flickr, in March 2007, there are 2,612 occurrences of `ge:head(ing)` and 4,362 of `geo:dir`.

In the case of `geo:lon` versus `geo:long`, (both meaning WGS84 longitude,) users are guided by a standard – 'long' from the GEO vocabulary. However, at the time of writing 26,237 photographs are tagged on Flickr with the correct term, `geo:long`, while 129,026 photographs are tagged with the more popular `geo:lon`. Perhaps users found that 3 character “lon” matched “lat” more memorably, or found the shorter term more practical. Perhaps this is a case not of taxonomy “emergence”, but rather “taxonomy drift”; the community of taggers has taken a standard term and refined it for practical use.

Quakr follows the philosophy of data tolerance: where synonyms exist, both are supported. Quakr understands both `geo:dir` and `ge:heading`, and has no preference regarding `geo:lon` or `geo:long`.

Semantic ambiguity

Tags in an emerging taxonomy tend to have no accompanying `<rdfs:comment>` definition. Sometimes users rely on obviousness of meaning, which works well for tags such as “Eiffel tower”, but badly for “zoom” (a close up photograph, or a fast car?). Occasionally, published documentation exists but it is often controlled by application developers, (for example, the KML specification [GE07].) The result is ambiguity.

What does 'Tilt' mean? In March 2007, 10,646 photos were tagged with `ge:tilt` on Flickr. The Quakr concept of 'tilt' matches the everyday meaning of that term, as demonstrated in Figure 16 - measured as the angle of the object from the normal to the Earth's surface. On this definition, a camera in a standard upright position has a tilt of “0”, which seems appropriate.

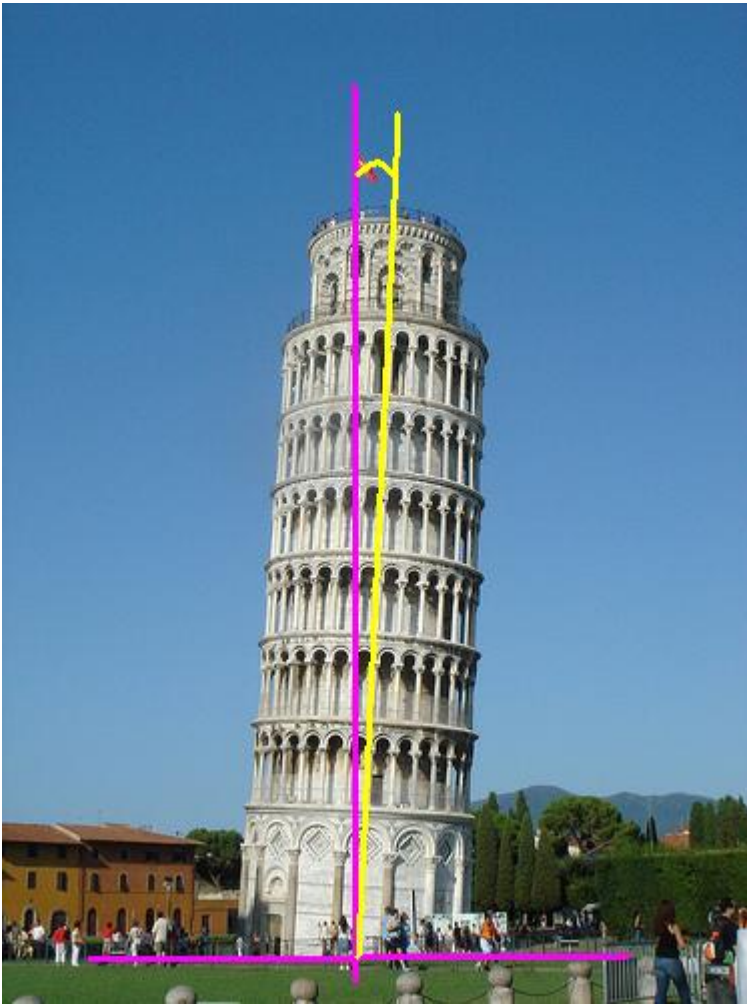


Figure 16. Common-sense concept of “tilt”. (The Leaning Tower of Pisa ©?natsu? cc by-nc-sa/2.0)




Figure 17. Common-sense concept of “tilt”, applied to holding a camera.

However, KML defines tilt as;

“Angle between the direction of the LookAt position and the normal to the surface of the earth” [GE07].

This means that tilt of "0" denotes pointing down at the ground.

Quakr is able to use the wealth of existing 'tilt' data, but the subtle difference in meaning of this term, and the discrepancy between the everyday and the documented meaning, are a source of confusion.

A second ambiguity arises: **What do the terms geo:lat and geo:lon mean?** The answer seems initially uncontroversial – they describe location. However, photographs have two locations – the location of the subject of the photograph, and the location of the photographer. Which is meant when a user adds these tags? Users who hold a GPS device in one hand and a camera in the other are measuring the latter. Users who use Flickr Map tool may be thinking of the former – or more likely, not thinking about this ambiguity at all. This is a case where a term is well defined in a general sense, but where subtleties of its meaning in the domain of images are still emerging.

Sparseness

The Quakr ontology contains at least 6 terms; Quakr performs best with photographs with a complete set. However, tag based metadata suffers from the lack of formal ontologies – we can not enforce `minCardinality` on Flickr users. Even if such constraints were possible, it is not clear that excluding sparse data would be desirable.

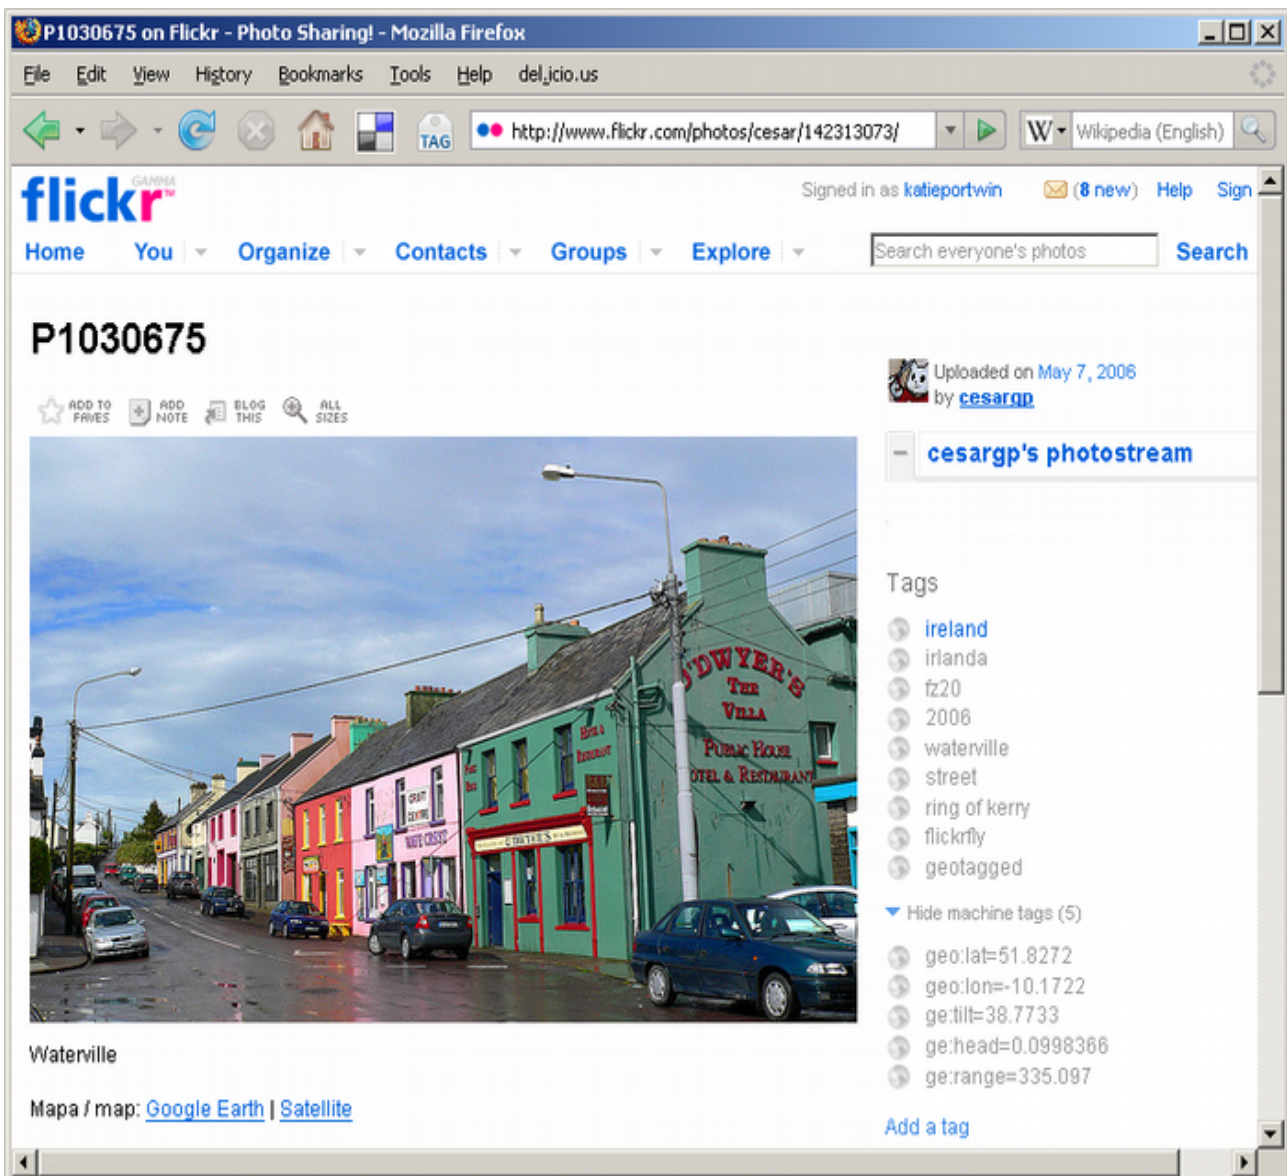


Figure 18. Example image with a “sparse” (in Quakr terms) but useful metadata set. (P1030675 © cesargp / cc by-nc-sa/ 2.0)

Two heuristics enable Quakr to tolerate sparse data sets: defaulting and enhancement.

Defaulting

The concepts of tilt and roll are amenable to defaulting. Their values are not likely to be evenly distributed - most photographs are taken with the camera upright and pointing straight forward. A tilt of 0 (`ge:tilt=90`) and a roll of 0 are reasonable defaults. There are of course cases when this heuristic fails to achieve good results - for example, when dealing with aerial photographs. However, the cost of occasional error is outweighed by the benefit of access to a larger data set.

Enhancement

Quakr enhances sparse data by taking advantage of XML Web Services to provide the missing values. Altitude is amenable to enhancement. EarthTools.org provides a RESTful [Fielding00] XML interface to NGIA/NASA satellite geospatial data [USGS07]; one of the many services returns the altitude of the earth's surface for a given latitude/longitude point [Stott05]. Assuming that most people take photographs standing on the surface of the

earth, Quakr is able to make an estimate of altitude. Again, the assumption breaks down in some cases – for example, photographs taken from helicopters, in mines etc.

Sample response from EarthTools (<http://www.earthtools.org/height/51.75425000/-1.25735000>)

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<height xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation=
  "http://www.earthtools.org/height.xsd">
  <version>1.0</version>
  <location>
    <latitude>51.75425000</latitude>
    <longitude>-1.25735000</longitude>
  </location>
  <meters>70</meters>
  <feet>229.7</feet>
</height>
```

Inaccuracy

Accuracy in the Quakr 3D world depends on the accuracy of metadata which underpins it. The accuracy of consumer grade GPS devices can be as low as 15 metres [Garmin07]. Although “Selective Availability” [NOAA00] was discontinued in 2000, technical factors conspire to reduce GPS accuracy [WP07]. The accuracy of data created with Flickr Map has also been widely criticised [Torniai07] (although alternatives such as GmIF look promising). In addition to this domain-specific problem, Quakr suffers from a general issue inherent tagging based metadata – human error. Users hand-typing tags may miss a decimal point, or confuse latitude with longitude. Without a formal ontology to range-check values, many geotags contain inaccurate or even invalid data.

Abundance

The strength of the Quakr approach over proprietary efforts is that data is contributed by a large community. A single Quakr 3D street may be composed of many different users' images. However, the difficulty with this approach is that Quakr has no control over data abundance at a given location. There is likely to be an *image jungle* at the Eiffel Tower, and an *image desert* in the Channel Tunnel.

Quakr uses two methods to disperse image jungles. The first of recruits the help of the user. Quakr users may customise their world by giving additional tags to restrict the image search. For example, “evening”, or “snow” give interesting results for Oxford city centre locations.

The second approach is to make selections automatically, developing heuristics based on co-occurring tags and other metadata. For example, the heuristic layer selects the most recent of two available images for the same location based on Exif date stamp. There is considerable scope for creativity in these heuristics, and it is likely to be an area of future development.

The Power of Heuristics

Quakr implements the data-tolerant approach with a layer of heuristics, some of which are controversial, all of which are in ongoing development. Heuristics map multiple overlapping ontologies, some stable, some at the early stages of “emerging”, into the Quakr ontology. Heuristics make assumptions about default values, insert enhancements into sparse data, and prune data jungles.

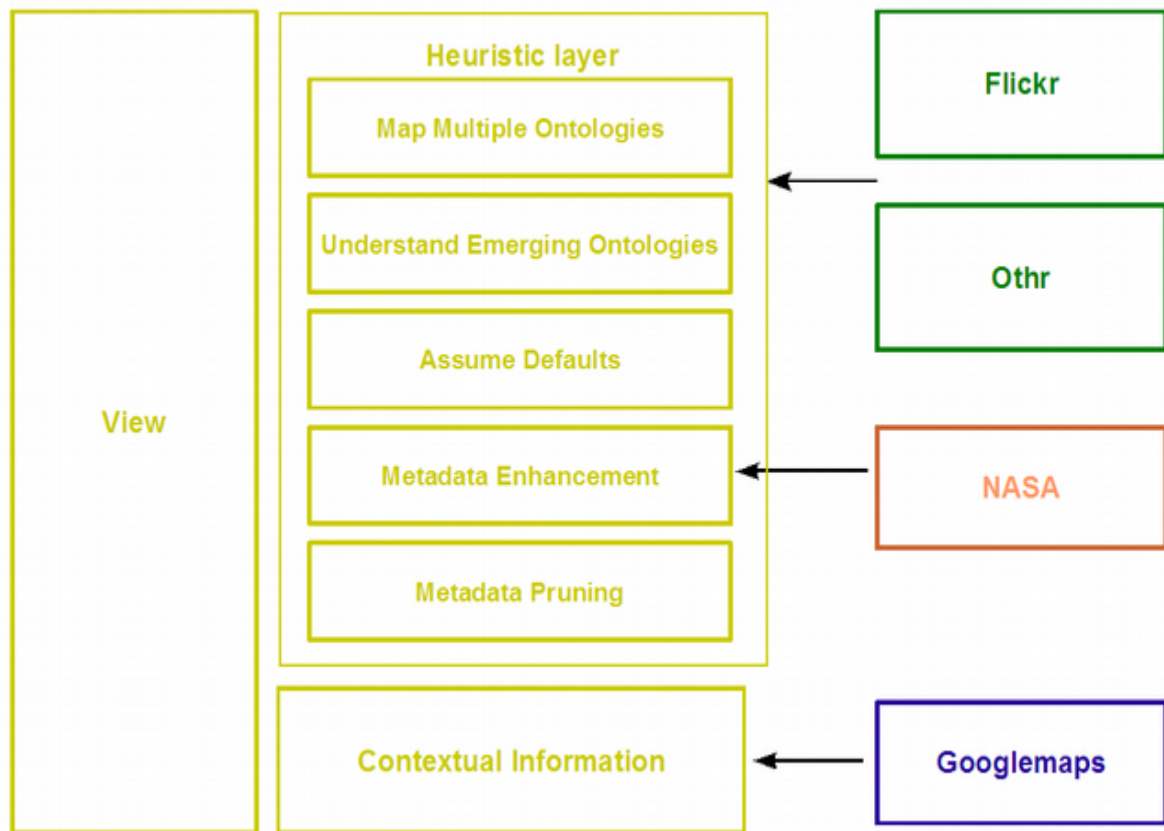


Figure 19. Quakr Heuristic layer.

Conclusion, Future Work

It seems likely that the popularity of geotagging will continue its exponential growth, and interest in the more subtle location metadata concepts will expand. Unlike “the Giants”, Quakr has everything to gain from this metadata explosion.

In Quakr, image deserts will shrink; in fact, dealing with image jungles is likely to become the significant challenge.

Other areas of anticipated development include:

1. Heuristics to deal with newly emerging taxonomies.
2. Enhancement of sparse data, as new location-focused web services come on-stream.
3. Enhancement of sparse data based on analysing co-occurring tags, descriptions or other metadata, (“..looking down at..”, “..standing on..”)
4. Further contextual information as it becomes available. So far, “context” is a visual map on the “floor” of the Quakr 3D world. As a next step, the Quakr team are interested in locating *sound* clips.

Links

- [Quakr homepage](#)
- [Quakr blog](#)
- [Quakr application](#)

- How to

Bibliography

- [Apple07] Apple, “MacBook”, 2007. <http://www.apple.com/macbook/macbook.html>
- [Beckett06] D. Beckett “Semantics Through the Tag”, Proceedings of XTech 2006: Building Web 2.0, 2006, Amsterdam. <http://xtech06.usefulinc.com/schedule/paper/135>
- [Cattuto06] C. Cattuto et al, “Collaborative Tagging and Semiotic Dynamics”, Proc Natl Acad Sci USA, 2007, 104(5), p1461-1464. <http://arxiv.org/abs/cs.CY/0605015>
- [Dubinko06] M. Dubinko M. et al, “Visualizing tags over time”, Proceedings of the 15th international conference on World Wide Web, 2006, Edinburgh. <http://doi.acm.org/10.1145/1135777.1135810>
- [Fielding00] R. Fielding “Architectural Styles and the Design of Network-based Software Architectures, Chapter 5”, 2000. http://www.ics.uci.edu/~fielding/pubs/dissertation/rest_arch_style.htm
- [Flickr06] Flickr, “flickr.photos.search, bbox”, 2006. <http://www.flickr.com/services/api/flickr.photos.search.html>
- [Garmin07] Garmin, “How accurate is GPS?“, 2007. <http://www.garmin.com/aboutGPS/>
- [GE07] Google, “KML 2.1 Reference”, 2007. http://earth.google.com/kml/kml_tags_21.html#heading
- [GPSPhones07] Amazon, “HP 6945 Phone”, 2007. http://www.amazon.com/Hewlett-Packard-HP6945-Phone-Unlocked/dp/B000ILDSLO/ref=pd_bbs_1/102-4047085-6525713?ie=UTF8&s=wireless&qid=1175787570&sr=8-1
- [Halpin06] H. Halpin “Identity, Reference, Meaning and the Web” Proceedings of the Identity, Reference, and the Web Workshop at the World Wide Web Conference, 2006. Edinburgh. <http://www.ibiblio.org/hhalpin/irw2006/hhalpin.html>
- [Merholz04] P. Merholz, “Metadata for the Masses”, 2004. <http://www.adaptivepath.com/publications/essays/archives/000361.php>
- [Nintendo07] Nintendo, “The Wii”, 2006. <http://wii.nintendo.com/>
- [Nokia07] Nokia, “Phone with Tilt”, 2007. http://www.allaboutsymbian.com/features/item/Accelerometers_Tilt_Sensors_on_Symbian_S60_3rd_Edition_smartphones.php
- [Nokia07b] Nokia, “Phone with Digital Compass”, 2007. <http://europe.nokia.com/A4145041>
- [NOAA00] National Oceanic & Atmospheric Administration, US Department of Commerce, “THE WHITE HOUSE Office of the Press Secretary, STATEMENT BY THE PRESIDENT REGARDING THE UNITED STATES' DECISION TO STOP DEGRADING GLOBAL POSITIONING SYSTEM ACCURACY”, 2000. http://www.ngs.noaa.gov/FGCS/info/sans_SA/docs/statement.html
- [ProgrammableWeb07] Programmable Web, “200 Photo Mashups”, 2007. <http://blog.programmableweb.com/2007/02/11/200-photo-mashups/>
- [Sony07] Sony, “Sony Playstation 3”, 2007. <http://uk.playstation.com/ps3/>
- [Stott05] J. Stott, “Earthtools.org, Web Services: Elevation/Height Above Sea Level”, 2005. <http://www.earthtools.org/webservices.htm#height>
- [Swinford06] S. Swinford “Zoom in: the world on your laptop in 3-D” The Sunday Times, November 26, 2006. <http://www.timesonline.co.uk/tol/news/uk/article650133.ece>
- [Torniai07] C. Torniai, S. Battle, S. Cayzer, “Sharing, Discovering and Browsing Photo Collections through RDF geo-metadata”, Proceedings of the 3rd Italian Semantic Web Workshop; Semantic Web Applications and Perspectives SWAP 2006, 2006, Pisa. <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-201/07.pdf>
- [USGS07] United States Geological Survey (USGS), “Shuttle Radar Topography Mission (SRTM)”, 2007. <http://srtm.usgs.gov/>
- [W3C03] W3C Semantic Web Interest Group, “Basic Geo (WGS84 lat/long) Vocabulary”, 2003. <http://www.w3.org/2003/01/geo/>
- [Weinberg07] N. Weinberg “Google Earth vs. Windows Live Maps 3D: Round Two”, 2007. <http://google.blognewschannel.com/archives/2007/02/15/google-earth-vs-windows-live-maps-3d-round-two/>

- [WP07] Wikipedia Community, “Global Positioning System“, 2007. http://en.wikipedia.org/wiki/Global_Positioning_System#Accuracy_and_error_sources
- [Wu06] X. Wu et al. “Exploring social annotations for the semantic web”, Proceedings of the 15th international conference on World Wide Web, 2006, Edinburgh. <http://doi.acm.org/10.1145/1135777.1135839>
- [WWW06] Panel Discussion ,“Meaning on the web: evolution vs intelligent design?” Proceedings of the 15th international conference on World Wide Web, 2006, Edinburgh. <http://doi.acm.org/10.1145/1135777.1135886>