



City Research Online

City St George's, University of London

Citation: Walker, R., Slingsby, A., Dykes, J., Xu, K., Wood, J., Nguyen, P., Stephens, D., Wong, W. & Zheng, Y. (2013). An Extensible Framework for Provenance in Human Terrain Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics*, 19(12), pp. 2139-2148. doi: 10.1109/tvcg.2013.132

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/2619/>

Link to published version: <https://doi.org/10.1109/tvcg.2013.132>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

An Extensible Framework for Provenance in Human Terrain Visual Analytics

Rick Walker, Aidan Slingsby, Jason Dykes, Kai Xu, Jo Wood, *Member, IEEE*,
Phong H. Nguyen, Derek Stephens, B.L.William Wong, and Yongjun Zheng




Fig. 1. Graphical summaries of bookmarks are used to record and browse the analytical process, here ordered (row-by-row) in the sequence in which they were bookmarked. Each can be used to access the live data, enabling analysts to revisit parts of the analytical process and helping verify past interpretations. A legend describing the encodings is provided in Fig. 6.

Abstract—We describe and demonstrate an extensible framework that supports data exploration and provenance in the context of Human Terrain Analysis (HTA). Working closely with defence analysts we extract requirements and a list of features that characterise data analysed at the end of the HTA chain. From these, we select an appropriate non-classified data source with analogous features, and model it as a set of facets. We develop ProveML, an XML-based extension of the Open Provenance Model, using these facets and augment it with the structures necessary to record the provenance of data, analytical process and interpretations. Through an iterative process, we develop and refine a prototype system for Human Terrain Visual Analytics (HTVA), and demonstrate means of storing, browsing and recalling analytical provenance and process through analytic bookmarks in ProveML. We show how these bookmarks can be combined to form narratives that link back to the live data. Throughout the process, we demonstrate that through structured workshops, rapid prototyping and structured communication with intelligence analysts we are able to establish requirements, and design schema, techniques and tools that meet the requirements of the intelligence community. We use the needs and reactions of defence analysts in defining and steering the methods to validate the framework.

Index Terms—Human terrain analysis, provenance, framework, bookmarks, narratives.

1 INTRODUCTION

Understanding the ‘human terrain’ is now considered essential to enable effective operations in regions in which military and intelligence resources are deployed. This landscape of social, political and economic organisation, conflicting beliefs and values and various interactions between local populations and military operations is complex and dynamic. Obtaining the information required involves diverse teams at many levels (strategic, tactical and operational) including intelligence from troops on the ground, many of whom have distinct socio-cultural skills and expertise. First-hand human terrain information is collected

through observation, interviews, and surveys – current practice is described in the handbooks published by the U.S. Army [15, 16]. Information about the human terrain is thus diverse. It is also inconsistent in quality, source and structure and frequently contains conflicting views.

Intelligence analysts use this diverse, incomplete and conflicting socio-cultural information to undertake human terrain analysis (HTA) [7] – the process through which what is known about the human terrain is studied, built into knowledge that describes the situation and used to understand and evaluate the effects and interactions associated with a military presence [51]. Doing so helps predict and assess potential threats and likely effects of intervention [37]. Distinct populations, their characteristics, key individuals and allegiances are identified and assessments made about how populations, groups or individuals are likely to behave. Analysts are interested in identifying recruiting methods, locations, targets, important dates, means and methods of transportation and communication, power structures and dependencies within supply networks, where tensions and disagreements occur and suchlike [37]. Some authors interpret this as cultural or socio-cultural intelligence (e.g. [37]) – enabling analysts to draw on the broader methods developed in sociology, cultural studies and anthropology whilst focusing on turning insights developed through these methods into actionable intelligence. In practice, HTA analysts seek to understand and predict the amount of danger that military per-

- Rick Walker, Kai Xu, Phong H. Nguyen, B.L.William Wong and Yongjun Zheng are with Middlesex University. E-mail: {r.walker, k.xu, p.nguyen, w.wong, y.zheng}@mdx.ac.uk.
- Aidan Slingsby, Jason Dykes and Jo Wood are at the giCentre, City University, London. E-mail: {a.slingsby, j.dykes, j.d.wood}@city.ac.uk.
- Derek Stephens is with Loughborough University. E-mail: d.p.stephens@lboro.ac.uk.

Manuscript received 31 March 2013; accepted 1 August 2013; posted online 13 October 2013; mailed on 4 October 2013.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

sonnel will expect to face, or the extent of cooperation they can expect from the local populations.

De Vries [7] argues that one of the most difficult aspects of analysing human terrain information is to turn data into actionable intelligence because many office-based intelligence analysts seldom experience the socio-cultural environment they are studying [37]. Moore [31] develops ideas of critical thinking [38] in the context of intelligence analysis through a number of case studies. These case studies show where problems have occurred and the kinds of mistakes that could have been avoided if analysts were more critical in their thinking. Patton [37] summarises a list of critical thinking questions that focus on the socio-cultural issues in analysing the human terrain. Understanding the human terrain is an ongoing task that requires comprehensive background knowledge so that trends can be judged and information assessed [30], with critical thinking a key component [31]. The process by which this knowledge is acquired must ensure that analysts are aware of data provenance whilst supporting critical thinking. It must involve recording and storing the analytic process so that intelligence can be justified, processes analysed and replayed, and gaps in information identified. Analytic tools and means of capturing and using data and analytical provenance are key here.

We tackle the challenges of undertaking rapid and meaningful analysis to develop insights that can be used as actionable intelligence through Human Terrain Visual Analytics (HTVA). Our work with defence analysts (DAs) applies visual approaches to HTA, where field reports from diverse sources of varying provenance, currency and quality are used to assess and compile a description of the human terrain. The objective is to use cutting edge ideas from visualization and visual analytics [26, 53] to develop techniques that are meaningful in HTA and can be taken through to delivery by the DAs.

A human-centred approach involving workshops with DAs allowed us to establish that compiling human terrain from such diverse sources of data requires an appreciation of the reliability of sources, how these sources relate to each other and how the sources were assessed in order to come up with interpretations. We use an iterative development process to design and deliver prototype HTVA tools to assist in interpreting these data, and to maintain a record of how these were used to arrive at interpretations. Our contributions are an *extensible framework for HTVA* that involves:

- i. the identification of *characteristics of data* used at the analytical end of the HTA chain;
- ii. a *schema* (ProveML) for modelling and storing the data, analytical provenance and interpretation;
- iii. a *prototype system* for performing HTVA;
- iv. a means of storing, browsing and recalling analytical provenance and process through *analytical bookmarks*;
- v. a means of forming *narratives* from analysts' findings, that link to the original live data.

The work is grounded [23] through a structured human-centred process [28] in which the kind of iterative approach described by Munzner [34] guards against threats to validity. This enables us to design schema, techniques and tools that link directly back to the requirements of the intelligence community and meet their needs.

2 APPROACH

Two full-day workshops held in a neutral space allowed us to establish current practice and perceived needs from the application domain and acquire feedback on progress and additional requirements in response to prototype solutions. Requirements for the HTVA prototype, data and provenance framework were developed in the light of this knowledge. Design and development used rapid, iterative agile methods that accommodated feedback on progress and enabled suggestions for enhancements and additional needs to be prioritised and implemented.

Four DAs participated in the *requirements workshop* along with eight academic members of the research team with expertise in visualization design, analytic provenance and information management.

The workshop was collaborative, with DAs and academics working together in mixed groups to share and discuss ideas and build common understanding [28].

Having shared preliminary information through 'fast forward' style presentations with a single indicative graphic, we established context through a series of presentations by DAs on their current work and objectives. The emphasis here was on describing current systems and perspectives on human terrain analysis. Participants recorded descriptions of the *good* and *bad* aspects of the current situation along with *aspirations* on colour-coded Post-it notes. DA contributions were marked accordingly and each participant was asked to highlight a 'favourite' statement that they deemed most important. We emphasised the value of all contributions and the need for inclusivity by asking each individual to read out their preferred characteristic [11]. This resulted in 78 descriptions of the current situation: 18 'good', 26 'bad' and 13 'aspirations' with 19 identified as 'favourites' including indicative aspirations such as: "*provide good and rich summaries for which details can be accessed*", "*a larger suite of tools / techniques that compliment each other well*" and "*using geospatial techniques to map micro narratives – showing how they evolve*". Both DAs and academics recorded these descriptions with ratios of 7:11, 11:17 and 11:21 from DA and academic participants respectively.

The DAs operate with classified data, meaning that our work needed to focus on accessible data sets with analogous characteristics. Data characteristics and analytical needs were established in small groups each containing a DA. A series of candidate data sets that might be used to develop novel tools and techniques for HTVA were considered in the discussions.

Given the shared context established, much discussion continued over lunch. This was followed by a *visualization awareness workshop* [27] showcasing selected academic and design work to demonstrate the variety of approaches that might be considered. Themed sections on stories, text, cartography, interaction and fluidity, transformation and user-centred visualization were used to stimulate 'like'/'dislike' responses captured once again on Post-it notes during the presentation. Of the 216 responses collected, 83 were from DAs. Having identified possible visual approaches, participants were then asked to extract generic and specific ideas from these visual stimuli that might be applied to address the *good*, *bad* and *aspirational* characteristics through which we described the current situation. We captured 51 of these ideas to be used as *guiding principles* in design, 33 coming from DAs. They were prioritised through a hierarchical process in which pairs of participants selected the most significant of their characteristics before pairing with other groups and repeating the process iteratively. This gave us an indication of the significance of each guiding principle and levels of agreement.

Having considered these various visualization stimuli, we worked individually on scenarios for visualization in the HTVA context. We established 20 of these relatively specific foci for development work. Our DA colleagues then took the initiative – an interesting move that emphasized the participatory nature of the exercise – by suggesting we work on scenarios based on the impending London Olympics in group discussion. Four groups, each containing and to an extent led by a DA, developed user stories, considered relevant datasets, discussed alternative scenarios and sketched solutions based on perspectives relating to hotels, the economy, socio-political issues and transportation. A rapporteur subsequently communicated key elements of the discussion to the wider group with sketches and flip charts.

A *feedback workshop* held at the same venue was used to steer development work by establishing DA reactions to prototype solutions developed on the basis of the established requirements. A structured process was used to evaluate the HTVA proposals developed involving: an introduction in light of the specific requirements being addressed; a demonstration of progress; the presentation of direct questions to inform subsequent development; open discussion about possibilities and directions. The uncertain nature of the proposed designs was emphasized to invite critique and suggested revisions. This was followed by open discussion about data, possible scenarios and views on progress and priorities. Reactions and ideas captured by transcribing

ing the meeting in full enabled us to identify popular features, those that required more detailed or auxiliary data and those that were out-of-scope. We were then able to prioritise additions and development.

3 REQUIREMENTS AND INFLUENCES

All information recorded at the initial workshop was transcribed, resulting in 402 Post-it notes and various sketches. The DAs reported negative issues with their current setup as including the dominance of single views rather than linked views or transforming views; the lack of uncertainty or provenance visualization; the difficulty in mapping and identifying nuances; and unseen bias in data. Positives included the well established workflow, strong understanding of analytical techniques and some recent tools that provide a platform for analysis. Responses that were indicative of more than one captured view were flagged as resonating.

Aspirations included the need for techniques that showed networks in their spatial context, mapping ‘micro-narratives’ to see how situations evolve, the concept of a “*Digital Recce*” – a summary of what is known about unstructured information relating to a particular theme – and an emphasis on “*Rapid 70%+ methods to give you a good look at a problem in a short time*”.

Examples of guiding principles that scored well in our prioritization exercise included: ‘lead generators’ – “*you can delve into the data to find new things (multiple questions)*”; ‘visualization in context’ – “*seeing not just the one thing I’ve selected, but how it fits into the whole set*”; ‘comparison (and exploration) over time, over space, over additional attributes’ – especially when comparison involves coordinates of the plane (superposition, juxtaposition [19]); ‘dynamic visualizations’ – “*movement of information between states (e.g. network and geo) would improve analytical understanding*”

As noted by [44], domain experts can be more forward in suggesting solutions than discussing problems. Given the design emphasis of the visualization guidelines exercise many of the scenarios generated were less applied and more focussed on tool development than HTVA. Nevertheless, we captured a series of questions that might be answered through the proposed designs including: *Who is the major influence here (now)?*; *What is the current situation here (now)?*; *Why is this group becoming more prevalent here (now)?*

Responses were coded to identify key themes but given the mass of detailed information these were deemed too general to guide design effectively. Alternatively a series of *design rules* were established to relate development of software components and the languages and processes used to link them to the range of requirements established at the workshop. This enabled us to make systematic use of the knowledge acquired and ensure that it influenced our work effectively. Having considered the responses and themes, their individual and collective scope and varying degrees of detail we established a principle whereby solutions should address:

- i. *at least 5 characteristics* captured, with an emphasis on DA responses and those not yet addressed;
- ii. *at least 5 aspirations*, again emphasizing those from DAs, those yet to be addressed and those highlighted as *favourites*;
- iii. *at least 3 guiding principles* that were from DAs and achieved *at least round 2 status* in the prioritization exercise.

In the case of visualization design ideas we agreed to additionally consider the influential aspects of the visualization examples and emphasize reactions originating from DAs, those identified as resonating, and those that had not yet been influential in existing designs. The project team complied with this approach by digesting the full transcribed data set and recording characteristics, aspirations, guidelines and visualization examples used as influences in their design ideas for others to see. Whilst the numbers used here are somewhat arbitrary they were determined following consideration of the data. Most importantly, this agreed practice ensured that designers engaged with requirements and that different designs tackled HTVA from different perspectives. It enabled us to justify design decisions and gave us a good working ba-

sis for grounding them in requirements established through the workshops.

4 DATA

The group work on data and analytical needs from the requirements workshop revealed *data* to consist of unstructured or semi-structured reports representing multiple, often conflicting, viewpoints through micro-narratives that are full of nuance. These time-stamped reports use a standard format and clear language (no metaphors or sarcasm). They are about places, have a named author and include links and cross references. There may be a time lag between data collection and the appearance of a report in the system. An analyst may receive thousands of these reports on a daily basis.

Various forms of uncertainty are associated with these data. This is implicit in the reports, but made explicit through subjective measures of credibility applied to sources. Data are incomplete – with known and unknown gaps – likely to involve collection bias and potentially prone to deliberate spoofing. Translation and transcription may involve additional uncertainty and bias. Automated techniques such as named entity extraction or sentiment analysis may be used to derive additional data from reports, but also introduce uncertainty. Spatial and temporal uncertainties exist but were considered “well understood”. Uncertainty associated with the analytic process is not formally documented, but is assessed informally.

Analytical needs included capabilities for triaging (grouping), selecting and linking reports. Multiple aspects of reports need to be considered – e.g. people or places – in light of multiple themes. The ability to look for bias, explore data gaps and restructure reports using alternative schema were highlighted. This work is frequently undertaken in the context of other data sets. There was a request for ‘schemas within schemas’ through which data can be viewed and transformed and different perspectives encoded.

We used these characteristics to select an appropriate surrogate from publicly available data sources. Standard social media sources such as Twitter or Facebook are not appropriate – the language being too informal, the text lacking structure, the message length being limited and the quality (and quantity) of information being relatively low. Internet review sites offer a useful alternative. We selected Qype [42] – Europe’s largest user-generated local reviews site – as our source through the ‘Qype for Developers’ API [43].

Qype has a number of similarities with HTA structured reports. Every review is a text-based report with information about the author, time of review and rating given. Each review links to the place reviewed, each place has a geographical location and is assigned to one or more hierarchical categories (e.g. “Restaurant”, “Indian Restaurant”). There may be a time lag between the event reported (visit to the place) and the preparation and posting of the review. Whilst some reviewers follow a consistent structure, this structure is not the same across reviewers. There is often lack of consensus amongst reviewers since they vary in their opinions. The quality of an establishment may also change over time and its context may change (e.g. new ownership) in ways that may or may not be picked up in the reviews.

There are also some important differences. Qype reviews describe places, whereas HTA reports tend to describe people, groups of people and organisations. Geographical information in real intelligence reports may be richer than simple point locations – for example areas or multiple points along a patrol route. They also are written by a very diverse set of people, including social scientists, regional studies experts and military sources.

Despite these differences, the DAs agreed that Qype represents an appropriate surrogate data source for our work. Through the Qype API we extracted 2669 reviews from 1551 reviewers of 1000 places across 959 categories.

5 FACETS AND SCHEMA

In order to ensure that our work with the Qype review data remains applicable to the HTA context, and that we are able to provide a general framework for HTVA, we abstracted key characteristics of the data that matched the properties of typical HTA reporting data as described




Fig. 2. ProveML representations of (a) facets of the Qype dataset and their dependencies represented using Open Provenance Model (OPM) notation; (b) an insight about an insight in which an analyst builds on one of his or her previous insights (in this case, concluding that three authors are in fact the same person); (c) a *state* (section 7.1) as a *document collection* with a 'state' **themed document** containing a URL to visual summary of the state and a text description of the visual encoding; and (d) a *bookmark* (section 7.2) which is like a *state*, but with a 'bookmark' **theme** and analyst comment in the **document**. We omit the Process in (b), (c) and (d) for clarity of representation in these diagrams.

by DAs in our workshops. The DAs informed us that it was important to be able to analyse different aspects of HTA data. We define and identify *facets* that do this. These facets reflect both data properties and the organisational emphasis given to the data in the visual analytic views of them. The concept of the facet as a means of providing a perspective on data is used elsewhere in information visualization and visual analytic design [3, 10, 48, 56]. Such abstractions can be seen as a generalisation of other approaches used to separate visual analytic data into space, time and object classes (e.g. [1, 39]).

We identified five facets as most significant in HTVA through our analysis of the data generated at the requirements workshop: **place**, **time**, **author**, **theme** and **document**. These abstract entities were adequate for mapping Qype data to the characteristics of HTA data supplied by the DAs. They do not specify how each is to be represented, but rather what kinds of relationships might exist between them. So for example, **place** can equally be described through geocoded point values, areal units, vague geographic place names or trajectories, whilst **time** might be a single point in time, a range or collection of overlapping periods and could relate to an number of event types [2]. A **document** might be text-based, an image, an audio recording or any other artifact that can be related to one or more other facets.

Each facet may be related to a collection of any other, so for example a **document** may have an **author**, refer to a specific point in **time** and **place** and be categorised by **theme**. Importantly we allow **themes** (and **themes** alone) to have a reflexive relationship so that we can build thematic hierarchies as well as distinguish analysts' annotations from primary and secondary sources. Our schema design is extensible: further facets may be added in order to reflect key relationships and visual emphasis of any systems that might use them.

5.1 ProveML

The Open Provenance Model [33] provides an abstract provenance layer that can be used for the exchange of provenance information between systems. Provenance is represented as a directed acyclic graph showing past execution of a process – the relationship of provenance to workflow is like that of logbook to recipe [32]. OPM defines three types of node:

- *Artifact* – an immutable piece of state
- *Process* – acts on artifacts, produces new artifacts
- *Agent* – controls or enables a process

These nodes are linked by edges representing dependencies between nodes in a semantic manner. For example, a Process and an Agent can be linked by a *wasControlledBy* dependency. *Roles* can be used to add additional semantic information to these dependencies.

Modelling our facets as OPM relations gives us the structure for Qype shown in Fig. 2a. **Place**, **document**, **time** and **theme** are all artifacts, an **author** is an agent and each **document** is produced as a result of a document writing process controlled by the **author**. The

structure of connections between facets is specific to the dataset in question – for example, **theme** could be linked **document** instead of to **place**.

We developed ProveML, a schema for provenance, by extending OPM and implemented it as an XML schema to allow easy transformation between representations. In addition to storing the provenance of documents in Qype, we support analyst-generated documents. Such documents may refer to multiple places, times, themes or authors and we distinguish them from source documents by linking them to the special 'Insight' **theme**. This also allows us to record the provenance of insights about previous insights across different analysts, who themselves act as authors in the system as shown in Fig. 2b.

To handle other specialised requirements of provenance capture for intelligence analysis, ProveML has two additional constructs: explicit modelling of uncertainty and collections.

5.2 Storing Uncertainty

There are many general categorisations of uncertainty in the literature: Gershon [18] discusses a taxonomy of imperfection; Thomson *et al.* [54] present a typology of uncertainty for geospatially referenced information for intelligence analysts; Skeels *et al.* [46] derive a classification from commonalities in uncertainty areas uncovered through qualitative interviews; and Zuk and Carpendale [59] extend Thomson's typology to support reasoning uncertainty. In our workshop session described in section 4, we identified three specific types of uncertainty as being of interest to our DAs in their HTA.

- source uncertainty*: how reliable is this source? Deliberate misinformation or spoofing is included in this category.
- collection bias*: there will be more information about some things than others. Is some characteristic of the data significant, or is it a result of the data collection methodology?
- process related uncertainty*: uncertainty that is introduced through attempts to automate extraction of information. For example, how accurate is the assessment of sentiment in the text?

We can store metrics for these types of uncertainty as annotations on our ProveML graph: any node or edge can have one or more types of uncertainty associated with it. We established simple numerical measures of these different types of uncertainty – for example, generating a measure of *source uncertainty* for all facets – but the framework can equally support more sophisticated numerical or ordinal measures as required. By attaching uncertainty to the graph structure in this fashion, there is potential for modelling uncertainty propagation [5].

5.3 Collections

Insights are rarely formed from a single **document**, **time**, **place** or **theme**. In fact, certain decisions (such as increasing surveillance on an individual) require multiple pieces of independent supporting evidence. ProveML can store insights drawn from multiple artifacts, but

this does not reflect the process involved in selecting those documents: elements of provenance are lost if this is not recorded. For this purpose, we define *Collections* in ProveML.

A Collection is a list of artifacts of one type (which each have their own provenance graph) together with the query that was used to select them and the timestamp at which the query was performed. This makes the collection reproducible: assuming that artifacts are tagged with a creation timestamp, the same query can be performed using the data available at the time. If new data is added, the change in Collection contents can be tracked. This may affect the degree of certainty of insights that rely on that Collection elsewhere in the provenance graph. It is also valuable in the context of an audit trail: an analyst may need to justify the conclusions they arrived at given the data available at the time those conclusions were drawn.

6 VISUAL ANALYTICS

Having established important characteristics of the data used in HTA, identified a suitable surrogate, modelled it appropriately and achieved a strong sense of the kind of support required by DAs, we are able to design visual analytics techniques to help support and record analytical work that encourages critical thinking in the context of data provenance as described in the introduction to this paper.

Visual analytics techniques have the potential to give overviews of data, establish trends and anomalies, and help generate insights [26, 53]. Design that is informed by *analysts' needs*, *analysts' feedback* and best practice is crucial to the successful achievement of these aims. Instructive feedback from the requirements workshop influenced our designs and our systematic use of the *design rules* described in section 2 ensured that our solutions addressed requirements, as we built the HTVA prototype that implements these designs (Fig. 3). Users' needs change as they are exposed to new designs and techniques because they draw attention to previously unconsidered possibilities [28]. Demonstrating progress using the HTVA prototype allowed us to acquire helpful feedback during the design process. It enabled the DAs to explore new avenues and all of us to discuss subsequent design possibilities openly. We used the Processing environment [17] to rapidly prototype ideas, eventually developing and combining those that received positive feedback into prototype tools robust enough to demonstrate to our DAs to elicit further feedback at our workshops.

To help determine best practice, we also looked at what others had done with similar data. OpinionSeer [58] uses a combination of scatterplots and radial visualization to support the analysis of review data from TripAdvisor.com. Wu *et al.* [58] use the same data source to demonstrate modelling of uncertainty propagation in an analysis. Chen *et al.* [4] use text analysis and coordinated multiple views in their analysis of reviews of a book from amazon.com and Oelke *et al.* [36] extract discriminating features from reviews on amazon to generate summary reviews and correlation plots.

6.1 Overview

Our design was guided by the *aspirations* and *guiding principles* identified in the requirements workshop. The idea of a '*Digital Recce*' coupled with the request for '*rich graphical overviews*' demonstrated a strong interest in broad synoptic graphics that show the distribution and range of the data. The '*rapid and dynamic filtering*' guiding principle showed the value DAs attached to exploratory techniques in which different aspects of data can be rapidly and visually explored. '*Intuitive*' and '*aesthetically-pleasing*' were also considered important characteristics to which we designed.

The way in which we abstracted data into *facets* (section 5) also informed the design. This abstraction allowed us to use the surrogate dataset in a way that directly maps onto the DAs' data. In the same way, basing our design around facets enabled the DAs to map the visual analysis techniques to their context – an important way in which we ensure our work remains grounded in the application domain despite our inevitable distance from the classified defence data. Each facet has attributes: name and location for **author** and **place**; text, rating and sentiment for **document** (we used Thewall *et al.*'s [52] measures of sentiment as an example here, where negative and positive val-




Fig. 4. Summary graphics showing the distribution of values for each attribute in each of the facets of (from left to right) of **author**, **place**, **document** and **theme** using dot maps for spatial data, bar charts of counts for categorical data and frequency histograms for continuous numerical data, including time.

ues indicate the degree of negative and positive sentiment); and name and level for **theme**. We chose not to make reference to the **time** facet in our HTVA prototype, instead addressing identified analytical needs by making it an attribute of the **document** facet through which documents can be selected. Despite this emphasis, the application reads and writes the **time** facet in ProveML accordingly, demonstrating the flexible way in which applications can be developed to use our faceted data.

In line with these requirements – and others from the workshop – the HTVA prototype allows graphics to be generated rapidly. Analysts can select attributes from the various facets to be plotted using a number of techniques or ‘visualization types’ that emphasise different aspects of the data. Each visualization type uses one of the facets as its *subject* and allows visual variables to be mapped as per section 6.3. For example, the scatterplot in Fig. 3 (left) has **place** as its subject – here dots in this visualization type represent places. Faceted overviews that show the distribution of attribute values for each facet type and allow them to be filtered are described in section 6.2.

6.2 Faceted overviews

Faceted overviews as shown on the right hand side of the screenshots in Fig. 3 reveal the characteristics of each of the facets. These are used to show global distributions of values for each attribute for the **author**, **place**, **document** and **theme** facet types in Fig 4. The summary graphic type for each attribute depends on the attribute type: a map for location information, a bar chart for categorical data and a frequency histogram for continuous quantitative data, including time. Each of these summary graphics is zoomable with data being appropriately rebinned on-the-fly. In addition to their other attributes, each facet type has a measure of source uncertainty, helping address the lack of uncertainty visualization in existing systems identified in the requirements workshop. The facet overview is displayed on the right of the screen in the HTVA prototype (Fig. 3) and the user can choose which of these to display and rapidly move between faceted overview.

6.3 Constructing a graphic

To facilitate the flexible and interactive construction of graphics [49], we allow the analyst to map data attributes from any of the facets to visual variables. This mechanism is based on the concept of *configuring layouts to address research questions* as implemented in HiDE [47], but we extend this work by allowing attributes from *different facets* to be used; for example, in Figure 3 (right), we size **place** facets by the average rating of **documents** associated with each **place**. The elements plotted (the subject) are **authors**, **places**, **documents** or




Fig. 3. Screenshots of the HTVA prototype with visual summaries of the attributes for a particular facet on the right and mappings between visual variables and facet attributes at the bottom of each of the three screen shots. *Left*: scatterplot of **places** showing the relationship between various **documents**’ attributes (count: x , average sentiment: y , standard deviation of sentiment: *size* and average source uncertainty: *colour*). *Middle*: density map of **document** sentiment (purple=positive; orange=negative), constrained to a **document** time range; *Right*: Dot map of **places**, coloured by their average **document** rating and sized by the standard deviation of this (a rudimentary indication of the degree to which views about a **place** are contradictory).




Fig. 5. Maps of **place** sentiment within different **document** time ranges are produced instantly as the DA interactively drags a selected time period on the **document** facet overview. Creating and dragging selection boxes in each of these views has similar effects in terms of view coordination.

themes. This provides a quick and flexible means to construct graphics relevant to the analytical task in hand.

The initial HTVA prototype offers three visualization types: a scatterplot, a density map and a dot map. These can be seen in Fig. 3. Each visualization type has different set of visual variables to which data variables can be mapped: dots in the scatterplots can be positioned in x and y , sized, and coloured by any attribute from any facet type; dots on the dot maps can be positioned, sized and coloured; and the colour on the density map is assigned to an attribute.

6.4 Filtering

Overviews such as those we describe in section 6.2 are useful means of *filtering* [9, 11]. We facilitate filtering by allowing the analyst to draw an area or range on the map, barchart or histogram. Fig. 5 shows that overall sentiment differs between two different time ranges, as indicated on the time histogram. These filters can be interactively dragged around and the main view updates in real time. This helps study changes over space, time and attribute space and conforms to the ‘multiple views with dynamic filtering’ guiding principle deemed so important in our requirements exercise.

7 ANALYTICAL PROVENANCE

Analytical provenance is the *process* by which the analyst achieves analytical outcomes. Feedback from the requirements workshop suggests that this is currently poorly supported, yet important for HTVA

because interpretations need to be audited, justified, revisited in the light of new information, and tackled by multiple analysts.

Gotz and Zhou [20] note that analytical provenance can be collected at different semantic levels. Low-level *events* (e.g. mouseclicks and keystrokes [6]) and higher-level *actions* (e.g. filtering and sorting [20, 24]) can be captured automatically, but these have low semantic content. Higher-level *sub-tasks* and *tasks* with higher semantic content cannot easily be captured automatically because they require analyst input. Our discussions with DAs suggested systems that offer optional comments at any stage as a good compromise between automatic capture and high semantic content – an approach we adopted.

Allowing analysts to review their history of automatically-captured *actions* and *events* may help them recall what they did, why they did it, identify gaps and retrospectively annotate [8, 40, 45]. Such *actions* and *events* can also be used in discussion forum contexts [22] or be presented as a narrative [13]. There are many examples in which restoring the graphical state used by the analyst is an important aspect of the system [14, 22, 47, 49, 55]. We use a similar approach, in which we can browse, reorder and restore past analytical states as a way of contributing to provenance by supporting, understanding and analysing analytic process and developing narratives. States are stored with annotation in ProveML as described in section 7.1.

The obtrusive nature of asking analysts for detailed semantic information during the analytical process prompted Gotz and Zhou [20] to investigate ways to partially process this. They produced a taxonomy of *actions* and suggest how patterns of these can be used to infer higher level semantic goals. This useful approach is something we will investigate in future to make provenance capture even less intrusive.

We presented the HTVA prototype to the DAs with some initial ideas on capturing and presenting analytical provenance at our feedback workshop. The reception was positive. A considerable number of very specific ideas were generated in a lively session where the emphases on dynamism, interaction and provenance were clear. The iterative, human-centred approaches employed here have resulted in ongoing feedback that steers and validates the approaches we present.

To accommodate and support analytic provenance we modelled the VAST Challenge 2011 Mini Challenge 1 data set [21] using ProveML. The rationale here was to show how the HTVA methods that we were developing could be applied to a *known* problem relevant to the HTA domain. The synthesised data consist of geo-coded microblog messages that contain more structure and less low quality information than standard social media sources. Our DA colleagues supported this decision, which also enables us to demonstrate the flexibility of the framework in terms of encoding different data sets.




Fig. 6. Graphical summary of a bookmark that provides information about any filtering used, the characteristics of the filtered subset of data and the visual encoding.

7.1 Capturing and browsing analytical process using states

We capture the analytical process by storing a series of *states*, each of which is automatically and non-obtrusively captured, when an *event* or *action* takes place, i.e. every time any aspect of the visualization is changed or data filtered. Each *state* contains enough information for the software to restore the visualization state from the live data, including information about how the data were filtered. These states are snapshots – we do not explicitly record the *operators* [47] used to move between states, which gives us the flexibility to reorder these states later – a key DA requirement (section 7.4).

Within the HTVA prototype, stored states can be restored from live data and live data accessed from stored states. Since visual representations are usually software specific, we use the lower fidelity representations of a bitmap thumbnail image and a high-level description of the visual encoding in an approach that reflects that used in HiVE [47]. States are modelled in ProveML as collections associated with a ‘state’ **themed document** containing a URL to a visual summary (section 7.3) of the state and a text description of the visual encoding. This configuration is shown schematically in Fig. 2c.

7.2 Bookmarks

Automatically-captured *states* contain little semantic information [20]. Although the DAs considered that there was a need to capture the entirety of the analytical process, they asserted the importance of ‘flagging’ particular moments that the analyst considers significant and attaching a comment or insight to this.

To address this, the HTVA prototype allows the analyst to make semantically-rich *bookmarks* at any time during the analytical process or afterwards when browsing the state history. Bookmarks are annotated with comments or statements of insight. Like states, bookmarks can be restored, allowing the analyst to revisit them. They are modelled in ProveML in the same way as states, but by additionally appending the comment to the text of the **document** and adding a ‘bookmark’ **theme**. This allows significant states to be recorded as bookmarks with limited effort on the part of the DA. A schematic representation is shown in Fig. 2d.

7.3 Graphical summaries of states and bookmarks

To help analysts handle the potentially large number of states and bookmarks that will result from analysis, we designed graphical summaries shown in Fig. 6 that aim to succinctly summarise important aspects of the state or bookmark – a common approach ([22, 25]). It comprises a thumbnail of the graphic constructed by the analyst, overlaid with a description of the visual encoding, how it has been filtered and characteristics of the filtered subset. Hue indicates facet type, as used throughout this paper.

Once again, the high-level description of the visual encoding avoids implementation-specific detail through an approach that has similarities with HiVE [47]. The visual encoding description is necessary because although the thumbnail may serve as visual cue, it is unlikely that enough details can be visually-resolved to determine what it shows.

The visualization type is coloured by the hue corresponding to the facet type that is the subject of the graphic (**document** in Fig. 6). Beneath this, the mappings between available visual variables and data attributes are listed, coloured by attributes’ facet types. Available visual variables depend on the visualization type – note how these differ between the visualization types in Fig. 1.

Rather than providing full details of the filter used, we supply a list of all the attributes that were involved in the filtering, coloured by their facet type, at the bottom of the summary. Characteristics of the resulting filtered set of facets are also displayed. Counts of **authors**, **documents**, **places** and **themes** are shown, along with the average *source uncertainty* of each facet type mapped to lightness. In Fig. 6, the source uncertainty for **themes** (light colour) is very high compared to the others, indicating that perhaps interpretations made on the basis of **theme** are less reliable than for other facet types – potentially important information in assessing the confidence one might place in the information encoded in a state or bookmark. In addition, a summary of the **place** location is provided as a *standard ellipse* [12] which summarises the extent and location of the **places**. The ellipse’s centre is located at the mean, while the long axis shows the direction of greatest dispersion and has a length of twice the standard deviation of positions along that axis. The length of the short axis indicates dispersion in the perpendicular direction. Finally, the mean and standard deviation of **document** time are shown as a rectangle. The spatial and temporal views are scaled to the maximum for all the states and bookmarks.

Within the software, these graphical summaries link directly to ‘live’ data that can be restored and revisited when the state or bookmark summary graphics are clicked.

7.4 Ordering and reorganising bookmarks

During or after the analytic process, the analyst is able to browse the graphical summaries of the various states and convert any of them to a bookmark by adding appropriate semantic information.

States and bookmarks can be browsed in the order in which they were generated, as in Fig. 7 (left). But as indicated by the graphical summaries shown, characteristics of states and bookmarks can be expressed quantitatively. Quantitative measures such as the number of **authors** in the filtered selection, the spatial extent of the **place** locations or the average **document** time can be used to sort bookmarks, offering useful alternative ways of browsing insights generated. For example, in Fig. 7 (right, ordered by average **document** time) the order of the events represented by the bookmarks can be seen, and Fig. 8 (arranged spatially by **place**) indicates that bookmarks are centred on just four distinct **place** locations and those on the periphery are the most spatially focussed.

As well as ordering by quantitative measures of bookmark characteristics, DAs expressed considerable interest in being able to drag bookmarks around the screen: “*you almost want to favourite the key decision points – maybe visually rearrange them*”. This would enable analysts to use the screen space to organise their ideas: “*I can drag stuff in and say that this group of documents support my hypothesis... this stuff refutes my hypothesis...*”. We also detected interest in using the screen coordinates more explicitly: “*...put stuff above or below the line – this would be really cool*”. There is certainly evidence from the sensemaking [41] literature that letting analysts drag bookmarks around in a ‘sandbox’ [57] helps organise their ideas and helps externalise some of the details about how insights relate to each other [35]. Recent work also suggests that enforcing a classification or ordering pattern can actually undermine an analyst’s comprehension of a data set [29]. We use the idea of a sandbox to allow analysts to organise states or bookmarks into ordered sequences or groups and subsequently store these configurations.

7.5 Forming narratives

The results of HTA are interpretations that lead to recommendations that help inform decision making. Feedback from the workshops indicated that analysts are regularly asked to give verbal briefings to decision-makers higher up in the chain of command with little advance warning. These are often short – in the range of 30-120 seconds




Fig. 7. Bookmarks created during an analysis of the VAST Challenge 2011 Mini Challenge 1 data set[21], which presents a fictional scenario of epidemic outbreak and spread in a city using geotagged Twitter-style messages. *Left*: Ordered (row by row) by the order in which the bookmarks were created by the analyst. *Right*: Ordered by the average time at which the documents were written.




Fig. 8. The bookmarks shown in Fig. 7, but placed on a map at the mean locations of the bookmarks' filtered subsets of places. Standard ellipses [12] summarise the spatial extent of these places. Some bookmarks are tightly focussed on a particular area, indicating a localised event, while others appear to be more general observations.




Fig. 9. A narrative has been formed by dragging bookmarks into an ordered group and attaching narrative text. Each bookmark can be linked back to the live data.

– but the time available may change with little notice. Addressing this scenario was an important requirement for our DAs.

Narratives are ordered sets of facts, assumptions, assertions and interpretations used to communicate a message – in our case, using states and bookmarks obtained from HTVA. It is unlikely that these would be presented in the same order in which they were found by the analyst. Using the sandbox functionality described above, we support analysts in creating narratives by arranging states or bookmarks into sequences or groups and then attaching a narrative message to each group, as illustrated by the simple examples in Fig. 9.

In ProveML, narratives are modelled by linking bookmark documents to a narrative document – that is, a document with a ‘narrative’ theme. The ordering of bookmarks within the narrative is encoded as the role for each bookmark, and we also annotate the link between each bookmark and the narrative with information about the spatial position of the bookmark in the sandbox. We can thus record the analyst’s use of screen space, which may have specific semantic meaning.

A narrative may pass through different levels of summarisation, dependent on the target audience and the time allocation. For example, an aide to a general might require a five minute extended summary, from which a smaller number of findings could be extracted as a new narrative for the final briefing. By representing these narratives using ProveML, an audit trail is maintained throughout the analysis, and in after-event analysis it can easily be determined what information was presented at each level from analyst to decision-maker.

By sharing narratives rather than the entire process, we also allow for different levels of security classification: while narratives for senior officers at strategic level would include (and link to) all data, those aimed at a tactical or operational level may simply provide appropriate findings through the graphical summaries without exposing the underlying evidence.

8 REACTION

Feedback that validates the approaches presented here has also been provided through ongoing communications with two DA colleagues, selected as key players on the basis of background and experience from the earlier workshops. Whilst this has been informal and collected during structured discussion at our feedback workshops, responses have been very positive on the whole: for example on the elegance of the abstraction of facets and the rapid update of the interface through which these can be explored, or in relation to our means of capturing and recording insights and the way that our approaches

enable analysts to show rigour and process in decision-making. When demonstrating the approaches we were informed: “[there are] no dead ends. All usable, all relevant and of value.”

The HTVA prototype evidently demonstrated potential for VA in HTA: “there is stuff in here that we cannot do, or haven’t thought about needing to do but potentially should be doing or just don’t know how to do it.”

The flexibility of the underlying XML representation of ProveML was noted: “I like that it’s XML based because that makes it extremely portable - you can just apply an XSLT and spit out some RDF”

ProveML was considered suitable for providing an audit trail, as a means of transferring knowledge and reasoning and to enable analysts to record, re-use and analyse their work: “Knowledge goes with the analyst. How can we make it persist? ... Using something like this so that a new analyst can go in and not lose the benefit of all that tacit knowledge would be really good.”; “Very interesting. You might want to go back through your reasoning and analyse your own analysis. OK, so I have written this assessment report, how did I come up with that assessment and how confident am I in my analysis?”

The HTVA framework was selected from a series of funded technological innovations for presentation to a broad audience at a recent DSTL showcase event. One of our DA contacts described the approach as “one of the most exploitable and operationally relevant projects on show”. Whilst we have not deployed methods in the defence domain we have established through this kind of evidence and our close work with DA colleagues that there is a strong likelihood that the methods used here are suitable for delivery to support informed and effective HTVA through which meaningful critical analysis can be undertaken, recorded and communicated to support HTA.

9 CONCLUSIONS AND FUTURE WORK

The approaches presented here show potential for VA in HTA. They are designed to support DAs in their critical thinking in the context of uncertain, unstructured data in which situations change and viewpoints vary as they make sense of the human terrain. Reactions from our DA colleagues indicate that they have considerable potential for so doing.

We identified important characteristics of the reports used by DAs in their HTA. These are about places, people and organisations, written by a diverse set of individuals, including social scientists, regional studies experts and other military sources. They are characterised by a lack of consistency, variable levels of uncertainty in the content, and lack of consensus. We found Qype web-review data to be a suitable surrogate. Our abstraction of data characteristics into facets allows us to use this surrogate data in ways that are meaningful to DAs and it ensures that techniques are transferable. This was further demonstrated through our use of the VAST challenge microblog data set.

In line with the DAs’ requirements, aspirations and guiding principles, the HTVA prototype allows the data to be visually explored through different facet types and allows informative interactive graphics to be quickly and flexibly generated by mapping data attributes to appropriate visual variables. Interactive faceted selection and filtering can be undertaken on-the-fly in real time and helps address the DAs’ requirement of exploring the data through space, time and other attributes. These sophisticated faceted filtering and multi-view representations form the basis of our visual approach to HTVA – providing the kind of ‘Digital Recce’ requested for HTA through which uncertainty can be considered, as data that represent conflicting views and originate from different sources are assessed, and insights achieved. We have learned plenty about restaurant reviews and reviewers [42] and revisited events in Vastopolis [21] and are confident that the facets used in our framework are relevant to the tasks to which our DAs apply themselves, but for security reasons we are unable to confirm that insights can be achieved from the confidential data used in HTA.

We address the DAs’ requirement to record the analytical process and allow key aspects of process to be bookmarked, in ways they find useful. The summary graphics of analytical bookmarks can be browsed in flexible ways that allow analysts to revisit process and to verify their own past interpretations and those made by others, perhaps in the light of new information or a change in context. This mechanism

also helps analysts rearrange bookmarks into narratives, designed for audiences at different levels. This enables analysts to brief decision-makers higher up in the chain of command, linking the narrative to the live data views encapsulated by the bookmarks. It also opens up possibilities for communicating intelligence *down the chain*, giving those with the socio-cultural expertise that enables them to collect the data opportunities to see some of the benefits of their work and communicate with analysts. This offers potential for addressing the shared experience issue raised by De Vries [7] with possible benefits for the HT information supply chain.

ProveML allows us to encode the faceted information and provenance relating to both data and process, transform data and share it between systems in an extensible way. For example, sets of narratives, that may contain conflicting or supporting perspectives, can be encoded and replayed later.

In combination, these techniques constitute an effective and extensible framework for analysis and provenance in HTA that meets many of the requirements and addresses many of the issues expressed and raised by DAs during our human-centred design.

Our flexible faceted design should support deployment of our framework in a defence context. By adjusting the facets to more closely match the real data, our tools and techniques should be applicable in the same fashion with minimal adjustment, and the feedback from our DA colleagues indicates that our methods for recording data and analytic provenance, bookmarking and constructing narratives are meaningful contributions to the process of HTVA.

The current prototypes have served their purpose reasonably well in providing ample evidence of possibilities for HTVA. They use adequate, plausible representations to address established requirements. If we were to develop prototypes into deployable software we would work on a number of design issues. The design of graphical summaries used for states and bookmarks would be key here. These are information rich and whilst likely to be learnable may not be optimal or intuitive – issues that we could test using formal evaluation methods. As narratives grow in size, limitations of screen space may require a reduction in size of the thumbnails. This gives scope for multi-scale representations. Bookmarks and their underlying collections of facets could also be summarised in other ways, for example as Document Cards [50]. Developing software for deployment would take advantage of our user-centred design process but would require us to augment it with quantitative studies of usability and utility – activity that is beyond the scope of our current work and not required to address our current objectives.

In the future we do aspire, as do Gotz and Zhou [20], to capture semantically-rich analytical provenance automatically and unobtrusively, and develop techniques that infer semantic information from analysts’ actions. Developing means of inference from the way in which analysts select and spatially organise bookmarks may offer interesting insights into analyst intent that enhance the HTVA process.

ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council and the UK Ministry of Defence (MOD) Defence Science and Technology Laboratory (Dstl) through their Data Intensive Systems (DaISy) programme. We would like to thank Captain Ashley Jeyes, Royal Corps of Signals, British Army, whose dissertation on the BSc Information Management and Computing degree in the Department of Information Science at Loughborough University “Can Information Visualization Enhance the Military’s Understanding of a Population (Human Terrain Analysis)?” helped inspire this work.

REFERENCES

- [1] N. Andrienko and G. Andrienko. Spatial generalization and aggregation of massive movement data. *IEEE TVCG*, 17(2):205–219, 2011.
- [2] K. Beard, H. Deese, and N. R. Pettigrew. A framework for visualization and exploration of events. *Information Visualization*, 7(2):133–151, 2008.

- [3] N. Cao, J. Sun, Y.-R. Lin, D. Gotz, S. Liu, and H. Qu. Facetatlas: Multifaceted visualization for rich text corpora. *IEEE TVCG*, 16(6):1172–1181, 2010.
- [4] C. Chen, F. Ibekwe-SanJuan, E. SanJuan, and C. Weaver. Visual Analysis of Conflicting Opinions. *IEEE VAST 2006*, pages 59–66, 2006.
- [5] C. Correa, Y.-H. Chan, and K.-L. Ma. A framework for uncertainty-aware visual analytics. In *IEEE VAST 2009*, pages 51–58, 2009.
- [6] P. Cowley, J. Haack, R. Littlefield, and E. Hampson. Glass box: capturing, archiving, and retrieving workstation activities. In *ACM Workshop on Continuous Archival and Retrieval of Personal Experience*, pages 13–18, New York, New York, USA, 2006. ACM Press.
- [7] A. de Vries. The human terrain of counterinsurgency operations: Developing the military mindset and social science support. In *The Cornwallis Group XV: Analysis For The Future of Afghanistan*, National Defense University, Washington DC, 2010.
- [8] M. Derthick and S. F. Roth. Enhancing data exploration with a branching history of user operations. *Knowledge-Based Systems*, 14(1-2):65–74, 2001.
- [9] M. Dork, S. Carpendale, C. Collins, and C. Williamson. Visgets: Coordinated visualizations for web-based information exploration and discovery. *IEEE TVCG*, 14(6):1205–1212, 2008.
- [10] M. Dork, N. Riche, G. Ramos, and S. Dumais. PivotPaths: strolling through faceted information spaces. *IEEE TVCG*, 18(12):2709–2718, 2012.
- [11] J. Dykes, J. Wood, and A. Slingsby. Rethinking map legends with visualization. *IEEE TVCG*, 16(6):890 – 899, 2010.
- [12] D. Ebdon. *Statistics in Geography*. John Wiley & Sons, Limited, 1977.
- [13] R. Eccles, T. Kapler, R. Harper, and W. Wright. Stories in GeoTime. *IEEE VAST 2007*, 7(1):19–26, 2007.
- [14] R. Eccles, T. Kapler, R. Harper, and W. Wright. Stories in geotime. In *IEEE VAST 2007*, pages 19–26, 2007.
- [15] N. Finney. *Human Terrain Team Handbook*. U. S. Army, 2008.
- [16] N. Finney. *Commander's Guide: Employing a Human Terrain Team in Operation Enduring Freedom and Operation Iraqi Freedom - Tactics, Techniques, and Procedures*. Center for Army Lessons Learned, 2009.
- [17] Fry, B. and Reas, C. Processing - (<http://processing.org>), 2013.
- [18] N. Gershon. Visualization of an imperfect world. *Computer Graphics and Applications*, 18(4):43–45, 1998.
- [19] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, 2011.
- [20] D. Gotz and M. X. Zhou. Characterizing users' visual analytic activity for insight provenance. *Information Visualization*, 8(1):42–55, 2009.
- [21] G. Grinstein, K. Cook, P. Havig, K. Liggett, B. Nebesh, M. Whiting, K. Whitley, and S. Knoecni. VAST 2011 Challenge: Cyber Security and Epidemic. *IEEE VAST 2011*, pages 299–301, 2011.
- [22] J. Heer, F. B. Viégas, and M. Wattenberg. Voyagers and Voyeurs: Supporting Asynchronous Collaborative Visualization. *Communications of the ACM*, 52(1):87–97, 2009.
- [23] P. Isenberg, T. Zuk, C. Collins, and S. Carpendale. Grounded evaluation of information visualizations. In *BELIV 2008*, page 6. ACM, 2008.
- [24] T. J. Jankun-Kelly, K.-L. Ma, and M. Gertz. A model and framework for visualization exploration. *IEEE TVCG*, 13(2):357–369, 2007.
- [25] S. Kaasten, S. Greenberg, and C. Edwards. How people recognize previously seen web pages from titles, urls and thumbnails. In *People and Computers XVI (Proc. HCI)*, pages 247–265, 2001.
- [26] D. A. Keim, J. Kohlhammer, G. Ellis, and F. Mansmann. *Mastering The Information Age-Solving Problems with Visual Analytics*. 2010.
- [27] L. C. Koh, A. Slingsby, J. Dykes, and T. S. Kam. Developing and applying a User-Centered model for the design and implementation of information visualization tools. In *2011 15th International Conference on Information Visualisation (IV)*, pages 90–95. IEEE, 2011.
- [28] D. Lloyd and J. Dykes. Human-centered approaches in geovisualization design: Investigating multiple methods through a long-term case study. *IEEE TVCG*, 17(12):2498–2507, 2011.
- [29] L. A. McNamara and N. Orlando-Gay. Reading, sorting, marking, shuffling. In *BELIV 2012*, 2012.
- [30] Ministry of Defence. *British Army Field Manual Volume 1 Part 10 Countering Insurgency*. Army Code 71876, 2009.
- [31] D. T. Moore. *Critical Thinking and Intelligence Analysis*. National Defense Intelligence College, 2007.
- [32] L. Moreau. The foundations for provenance on the web. *Found. Trends Web Sci.*, 2(2-3):99–241, 2010.
- [33] L. Moreau, B. Clifford, J. Freire, J. Futrelle, Y. Gil, P. Groth, N. Kwasnikowska, S. Miles, P. Missier, J. Myers, et al. The open provenance model core specification (v1. 1). *Future Generation Computer Systems*, 27(6):743–756, 2011.
- [34] T. Munzner. A nested model for visualization design and validation. *IEEE TVCG*, 15(6):921–928, 2009.
- [35] C. North, A. Endert, C. Andrews, and G. Fink. The visualization pipeline is broken, FODAVA-10-18. 2009.
- [36] D. Oelke, M. C. Hao, C. Rohrdantz, D. A. Keim, U. Dayal, L.-E. Haug, and H. x. r. Janetzko. Visual opinion analysis of customer feedback data. *IEEE VAST 2009*, pages 187–194, 2009.
- [37] K. Patton. *Sociocultural Intelligence: A New Discipline in Intelligence Studies*. Continuum, July 2010.
- [38] R. Paul and L. Elder. *The Miniature Guide to Critical Thinking-Concepts and Tools*. Foundation for Critical Thinking, Oct. 2009.
- [39] D. J. Peuquet. It's about time: A conceptual framework for the representation of temporal dynamics in geographic information systems. *Annals of the Association of American Geographers*, 84(3):441–461, 1994.
- [40] W. Pike, J. Bruce, B. Baddeley, D. Best, L. Franklin, R. May, D. Rice, R. Riensche, and K. Younkin. The Scalable Reasoning System: Lightweight visualization for distributed analytics. *Information Visualization*, 8(1):71–84, 2009.
- [41] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proc. International Conference on Intelligence Analysis*, 2005.
- [42] Qype GmbH. Qype - <http://www.qype.co.uk>, 2013.
- [43] Qype GmbH. Qype for developers – <http://www.qype.co.uk/developers/api>, undated.
- [44] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE TVCG*, 18(12):2431–2440, 2012.
- [45] Y. B. Shrinivasan and J. J. van Wijk. Supporting the Analytical Reasoning Process in Information Visualization. In *ACM Conference on Human Factors in Computing Systems*, pages 1237–1246, New York, New York, USA, 2008. ACM Press.
- [46] M. Skeels, B. Lee, G. Smith, and G. G. Robertson. Revealing uncertainty for information visualization. *Information Visualization*, 9(1):70–81, 2010.
- [47] A. Slingsby, J. Dykes, and J. Wood. Configuring hierarchical layouts to address research questions. *IEEE TVCG*, 15(6):977 – 984, November 2009.
- [48] G. Smith, M. Czerwinski, B. R. Meyers, G. Robertson, and D. Tan. Facemap: A scalable search and browse visualization. *IEEE TVCG*, 12(5):797–804, 2006.
- [49] C. Stolte and P. Hanrahan. Polaris: A system for query, analysis and visualization of multi-dimensional relational databases. *IEEE TVCG*, 8:52–65, 2002.
- [50] H. Strobel, D. Oelke, C. Rohrdantz, A. Stoffel, D. A. Keim, and O. Deussen. Document cards: A top trumps visualization for documents. *IEEE TVCG*, 15(6):1145–1152, 2009.
- [51] The Development, Concepts and Doctrine Centre. Joint Doctrine Publication 04 Understanding (JDP 04). Technical report, Ministry of Defence, Shrivenham, Swindon, Wiltshire SN6 9RF, December 2010.
- [52] M. Thelwall, K. Buckley, G. Paltoglou, D. Cai, and A. Kappas. Sentiment in short strength detection informal text. *J. Am. Soc. Inf. Sci. Technol.*, 61(12):2544–2558, Dec. 2010.
- [53] J. J. Thomas and K. A. Cook. *Illuminating the path: The research and development agenda for visual analytics*. IEEE Comp. Soc. Press, 2005.
- [54] J. Thomson, E. Hetzler, A. MacEachren, M. Gahegan, and M. Pavel. A typology for visualizing uncertainty. In *Proc. SPIE*, volume 5669, pages 146–157, 2005.
- [55] F. Viegas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon. ManyEyes: a site for visualization at internet scale. *IEEE TVCG*, 13(6):1121–1128, 2007.
- [56] M. Weiland and R. Dachsel. Facet folders: flexible filter hierarchies with faceted metadata. In *Proc CHI'08*, pages 3735–3740. ACM, 2008.
- [57] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, B. Cort, and O. I. Inc. The sandbox for analysis concepts and methods. *ACM CHI*, 2006.
- [58] Y. Wu, F. Wei, S. Liu, N. Au, W. Cui, H. Zhou, and H. Qu. OpinionSeer: Interactive Visualization of Hotel Customer Feedback. *IEEE TVCG*, 16(6):1109–1118, 2010.
- [59] T. Zuk and S. Carpendale. Visualization of Uncertainty and Reasoning. In *SG '07: Proc. 8th Int. Symp. on Smart Graphics*. Springer-Verlag, 2007.