Tasking the Tweeters: Obtaining Actionable Information from Human Sensors

Alun Preece[a], Will Webberley[a], Dave Braines[b]

[a]School of Computer Science and Informatics, Cardiff University, Cardiff, UK
[b]Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, UK

ABSTRACT

Social media sources such as Twitter have proven to be a valuable medium for obtaining real-time information on breaking events, as well as a tool for campaigning. When tweeters can be characterised in terms of location (e.g., because they geotag their updates, or mention known places) or topic (e.g., because they refer to thematic terms in an ontology or lexicon) their posts can provide actionable information. Such information can be obtained in a passive mode, by collecting data from Twitter's APIs, but even greater value can be gained from an active mode of operation, by engaging with particular tweeters and asking for clarifications or amplifications. Doing so requires knowledge of individual tweeters as "sensing assets". In this paper we show how the use of social media as a kind of sensor can be accommodated within an existing framework for sensor-task matching, by extending existing ontologies of sensors and mission tasks, and accounting for variable information quality. An integrated approach allows tweeters to be "accessed" and "tasked" in the same way as physical sensors (unmanned aerial and ground systems) and, indeed, combined with these more traditional kinds of source. We illustrate the approach using a number of case studies, including field trials (obtaining eyewitness reports from the scene of organised protests) and synthetic experiments (crowdsourced situational awareness).

Keywords: social media; intelligence, surveillance, reconnaissance; controlled natural language

1. INTRODUCTION

There has been an upsurge of interest in recent years in viewing social media streams as sources of actionable information for situational awareness. It has been observed that certain forms of social media can serve as a human-based sensor network.[1] Twitter stands out as particularly useful in this respect due to its real-time characteristics and follower-based model, often making it the first place users turn to when reporting breaking news and events as they happen.[2] Indeed, we have seen Twitter become highly significant as a channel for public reporting on high-profile crimes in their immediate aftermath, including the Boston Marathon bombing in the US and the Lee Rigby murder in the UK, leading to issues for the police and other authorities.[3] The utility of Twitter in this context has led to the creation of a number of platforms designed to obtain situational awareness and event detection from Twitter data, including Twitcident,[4] Apollo,[5] ReDites[6] and Sentinel.[7]

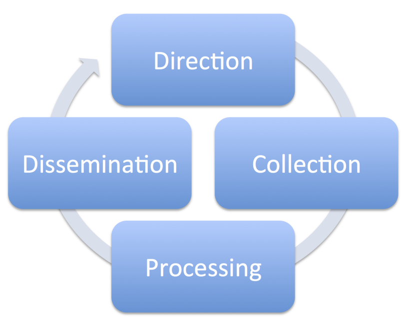

Figure 1. The direction-collection-processing-dissemination (DCPD) cycle

Send correspondence to A.Preece, E-mail: PreeceAD@cardiff.ac.uk, Tel: +44 (0) 29 2087 4653

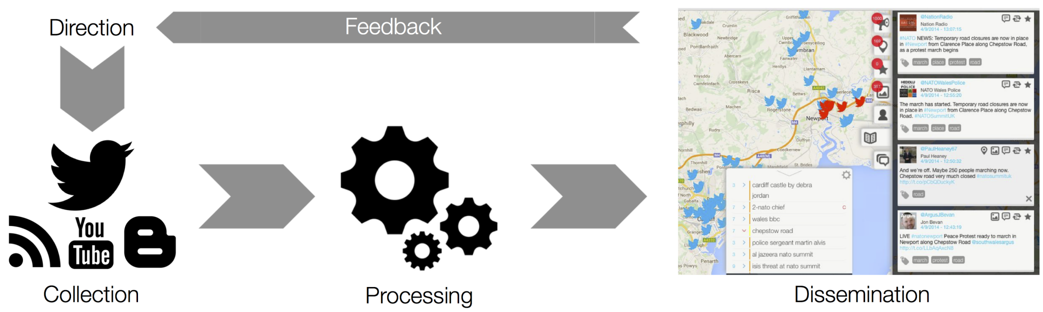

Figure 2. A generic social media processing pipeline mapped to DCPD steps

Management of sensing assets for intelligence, surveillance and reconnaissance (ISR) traditionally follows a well-known cycle, referred to in the UK as DCPD: direction, collection, processing and dissemination.[8] The cycle is shown in Figure 1. In the US variant of the DCPD cycle (called TCPED: tasking, collection, processing, exploitation and dissemination), direction is referred to as 'tasking' and the processing step is divided into two parts, 'processing' and 'exploitation', where the former is essentially pre-processing to put data into a usable form, and the latter involves putting the information into the context of a particular decision.

Viewing a social media stream such as Twitter as a human-based sensing system means that we can map its use to the DCPD cycle as follows:

Direction involves establishing what data to collect from the available streams. In principle with Twitter it is feasible to collect all available data, though the cost of this is substantial, as are the computational resources required to handle that volume of data. Most approaches to using Twitter data manage data volumes by using Twitter's streaming API* to collect tweets that match a set of search terms or are posted within a given geospatial region. Sometimes, as with systems such as Social Sensor[9] and Sentinel,[7] search parameters are chosen by a subject-matter expert; in other cases such as Twitcident[4] tweets are obtained based on terms identified in other feeds such as from the emergency services or a newswire service.

Collection depends on the mechanisms available for the form of social media of interest. As noted above, Twitter is particularly convenient in this respect, offering a number of APIs for streaming, searching and sampling. Other forms of social media, such as comment threads on YouTube, are also obtainable via APIs. Certain other platforms such as Facebook, however, are more restrictive in their access policies and offer only limited potential for open source data collection. Commonly, systems that process Twitter data perform some filtering after collection, to remove "noisy" tweets that can otherwise skew subsequent analysis. For example, geospatial collections are often dominated by tweets about particular global celebrities or wishes of "Happy birthday".

Processing can take a variety of forms including probabilistic techniques,[10] natural language processing,[4] sentiment identification[11] and event-detection.[12] Often several techniques are applied in conjunction, using a modular pipeline architecture.[6,7] The processing step is shown in Figure 2 in the context of a generic social media processing pipeline. The objective of this step is generally to provide semantic enrichment of the data, to make it useful in situation understanding. This process often involves detecting trends or clusters in the tweets or other media fragments. In this it resembles the traditional information fusion process, moving from low-level signal to higher-level information products.[13]

Dissemination of the results of the processing step may involve visualisation (as in the same screenshot fragment — from the Sentinel tool[7] — in Figure 2), summarisation, alerting an analyst via a notification, depending on the user's preferences and means of access (for example, via a computer or mobile device). Common visualisations of Twitter data include social network graphs, topic maps, timelines and geospatial plots, corresponding to who, what, when and where kinds of question. Dissemination will typically provide

*https://dev.twitter.com/streaming/overview the user with a means to further query the results. Sometimes these queries will be answerable using information already generated as a result of processing; often however, the further queries will require directions to collect additional or alternative data, which is why DCPD is a closed-loop process.

Framing the use of social media sources in terms of the DCPD process allows us to consider such sources as "just another kind of sensor" in multisource ISR. In our previous work we've addressed the problem of dynamic ISR asset assignment to mission tasks,[14] most recently in a human-in-the-loop context using knowledge representations based on a controlled natural language[15] that is both human-consumable and machine-processable.[16] The primary motivation for this paper is to extend our previous work in sensing asset modelling to consider the features of human sensors, specifically tweeters, and to revisit our approach to modelling mission tasks to consider the requirements for direction and collection of data from Twitter. Before examining these modelling requirements in Sections 3 and 4, we introduce a pilot study that was conducted in July 2014 focussing on the use of Twitter data for situational awareness of a large protest march in the UK. The results of this experiment illustrate many of the issues in using social media data as a source of intelligence.

2. A PILOT STUDY

The results reported here were captured as part of a pilot study to assess the usefulness of Twitter data in obtaining situational awareness in relation to a major event with potential for public order disruption. On July 26, 2014, protests were held in a number of cities across the UK in relation to Israeli incursions into Gaza. As part of this national campaign, well over a thousand people marched through the centre of Cardiff in South Wales. Automated data collection was performed from Twitter using the Sentinel tool[7] using Twitter's streaming API with three sets of search parameters: (1) to collect all geotagged tweets originating in the South Wales area, including Cardiff city, (2) to collect tweets containing "topical" search terms such as "Cardiff Gaza" and "Cardiff protest", and (3) to collect tweets mentioning locations near the expected route of the march such as "Cardiff City Hall", "Cardiff Saint Mary Street" and "Queen Street". In addition to this automated collection, a member of the team attended the scene and used standard Twitter search tools on a mobile device (iPhone) to identify "key" tweets as events unfolded.

The chart in Figure 3 shows the overall volume of tweets collected by Sentinel in relation to the July 26 protest in Cardiff, from 14:00 BST on July 26 to midnight on July 27. The charts are annotated with some key events during the relevant period. Tweet volumes are shown for the following: (i) tweets explicitly mentioning "Cardiff" and "protest" or "Gaza"; (ii) geotagged tweets originating in the Cardiff area and mentioning "protest" or "Gaza"; (iii) tweets mentioning "protest" or "Gaza" (but not Cardiff-specific); and (iv) tweets mentioning "police" or "swpolice" — the latter being the name of the South Wales police force official Twitter feed, *@swpolice*. The chart is labelled with a number of key events:

- The start of the march, around 14:45 BST on July 26. During the march, protesters were subjected to verbal and physical abuse when they passed a number of bars on St Mary Street and Mill Lane, around 15:15 BST. On Mill Lane, one of the marchers retaliated, leading to a fight outside a bar on the street. The march ended around 15:40 BST.
- On the evening of July 26, around 21:00 BST, a video was posted on YouTube of the Mill Lane violence. This video is now marked as private and no longer available.
- Around noon on July 27 the BBC ran a news story alleging that the violence resulted from "poor policing", linking to a second YouTube video[†].

The graph shows far more tweets mentioning Gaza protests in general than the Cardiff protest in particular, largely due to there being a national programme of marches in the UK on July 26. Many of these tweets were about the event happening concurrently in London. Tweets referring to the Cardiff march exhibited considerable variance in estimating the size of the crowd, from a few hundred to "thousands". The BBC reported the official estimate at 1,500. There was a significant increase in the volume of Cardiff-specific tweets as the marchers

[†]http://www.bbc.co.uk/news/uk-wales-south-east-wales-28509791

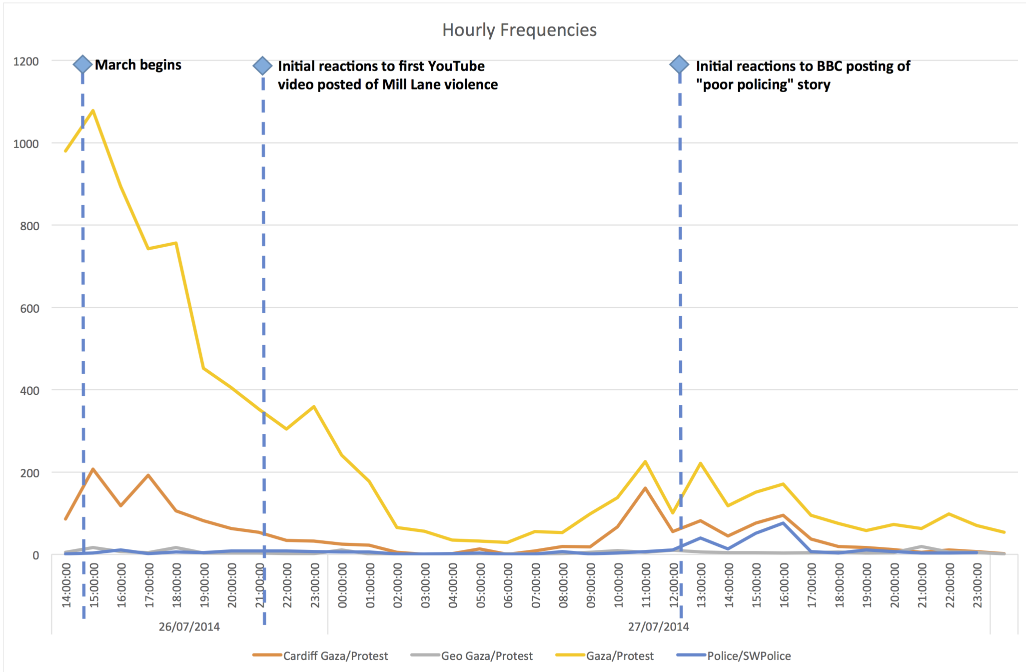

Figure 3. Timeline of the July 26, 2014 protest and its aftermath

assembled and set off. This was followed by a dip at the end of the march, then a recovery in volume as people tweeted after the event. A negligible portion of tweets mentioning Gaza or the protest were geotagged. This is consistent with few Twitter users in general geotagging their tweets. There was no significant volume of tweets on the Mill Lane violence at the time. However, the events were detectable from Twitter data as we discuss below. Few tweets mentioned "police" or "swpolice" until Sunday afternoon, following mainstream media reports of the violence. Nevertheless, there were some interesting "small signals" relating to the @swpolice account, highlighted below.

Contemporaneous tweeting of the Mill Lane events by one of the marchers included the following: "Glasses thrown from Walkabout Cardiff Customers at #gazaj26 marchers. It's slowing nobody down." (15:13 BST) and "More glasses thrown by customers of bars on Cardiff's Mill Lane at peaceful marchers. Thrown at women and children. Animals. #gazaj26" (15:18 BST). "Walkabout" in the first tweet is the name of a bar on St Mary Street in Cardiff. The march progressed down that street before turning onto Mill Lane. Both these tweets use the *#gazaj26* hashtag chosen for the UK-wide day of protests.

There were no tweets posted from the @swpolice Twitter feed during July 26–27. However, a number of the collected tweets were directed at the official account, including this one at 15:57 BST on the 26th: "A mass #FreePalestine protest in Cardiff city centre and where were the @swpolice?" In response to this tweet, the poster was contacted for comment by the WalesOnline local news service on July 27.

Immediately post-event, there was a low level of concerned tweeting regarding the violence; however, this gathered momentum when the first YouTube video appeared on the evening of July 26. By 23:00 BST, a number of retweets of the link to the video contained the phrase "riot in mill street cardiff". By the afternoon of July 27, the focus of the story had become alleged "poor policing" by South Wales police. After the BBC ran the story on July 27, a link to the article was retweeted 46 times in one hour, and subsequent public reaction revealed a variety of opposing views: "Saw a video of the Cardiff protest. Drunken idiots chucking pints at protesters kicked everything off :(" (12:09 BST, 27 July) and "Drunken idiots attack Gaza demonstration…Not drunken idiots, just good Cardiff boys standing up to Muslim Scum !" (20:38 BST, 27 July).

It is also worth noting that, concurrent with the start of the march, there was a much smaller-scale second protest occurring on Queen Street near to the route of the main protest. We could find no contemporaneous tweeting of this event; however, later tweets on July 26 mentioned it, with images: "Shahada flag in Cardiff + jew-bashing = wonder where those lads who joined ISIS got their inspiration? @swpolice http://pic.twitter.com/9XrZzgCHK6" (16:47 PM, 26 July) and "#Gaza Important march in Cardiff today even though ISIS held a gathering by Aneurin Bevan's statue!" (20:18 BST, 26 July). The mention of ISIS in the first tweet refers to the case of three Cardiff youths who reportedly joined the jihadist group now known as Islamic State in Syria in early 2014[‡]. This tweet is also directed at the @swpolice account. The statue of Aneurin Bevan mentioned in the second tweet is a well-known location for public protests in Cardiff.

In conclusion, this pilot study highlights the value of Twitter-derived data in providing situational awareness. A combination of search terms and a geospatial area of interest proved effective for focussed data collection. Relatively simple term frequency analysis was able to identify useful small signals in the data, including (at the time of the violence) "mill lane". Peaks caused by retweeting of common phrases such as "riot in mill street cardiff" are informative — in this case not of the violence itself, but of growing public awareness of it, which as we saw can then tip over into wider concern over official handling of the event. We are not claiming here that social media posting is representative of the views of the population at large. Our purpose is not to conduct surveys[17] but rather to detect activity. We accept that social media is prone to the propagation of misinformation.[18] Indeed, one could argue that calling the small-scale incident on Mill Lane a "riot" is an example of misinformation or at least exaggeration. But this case also shows why detection of peaks around misinformation is important to authorities in terms of situational awareness and the need for responsive action.

3. MODELLING TWEETS AND TWEETERS

As noted at the end of Section 1, our recent work has involved the use of a controlled natural language (CNL) for knowledge and information representation. Our CNL, ITA Controlled English (CE),[19] is approximately as expressive as the W3C's Web Ontology Language (OWL),[20] and features a rule language with similar capabilities to the Semantic Web Rule Language.[21] Our motivation for adopting a CNL is to improve human-computer collaboration by using a single representation shared by the machine and human users. Previously, we have shown how its use can support soft and hard information fusion, including the tasking of physical sensors.[22]

For illustration, a sample CE model definition is shown below.

```
conceptualise a ~ twitter account ~ A that
  is an online identity and
  is a temporal thing and
  has the value L as ~ location ~ and
  has the value NT as ~ number of tweets ~ and
  has the web image PP as ~ profile picture ~ and
  has the value NT as ~ number of tweets ~ and
  has the value NFR as ~ number of friends ~ and
  has the value NFO as ~ number of followers ~.
```

A `conceptualise` sentence defines a new concept in a CE model (ontology). New terms in the model — concepts, properties and relationships — are introduced between the tilde (~) symbols. The example defines the concept `twitter account` as being a child of the parent concepts `online identity` and `temporal thing`, and having properties such as `location`, `number of tweets`, `profile picture`, etc. The property definitions include the type of the value: either a literal value (e.g., for `number of tweets`) or a concept type (e.g., `web image` for the `profile picture` relationship).

Instances (facts) are defined using a similar syntax. The example below shows an instance of the concept `journalist`. (This example was chosen because the individual is a public figure and the BBC publicly lists the professional Twitter accounts of its journalists.)

```
there is a journalist named 'Paul Heaney' that
  uses the twitter account 'paulheaney67' and
  works for the media organization 'bbc'.
```

‡http://www.bbc.co.uk/news/uk-wales-south-east-wales-28116575

This instance is named `Paul Heaney` and has a `uses` relationship with an instance of the concept `twitter account` (as defined above). The `twitter account` is named `paulheaney67`. The `journalist` instance `Paul Heaney` also has a `works for` relationship with an instance of the concept `media organization`, named `bbc`. The role of CE is to have extensible models with whatever concepts, properties and relationships are needed. So `works for` is just one relationship that we chose to model, but there can be any of these. As we will show later, the model can be extended at run time also.

Modelling tweeters allows us to process Twitter data such as the set collected during the pilot study in Section 2 and to build profiles of individual human "sensors" and connections between them, including the following elements:

1. information derived from an individual tweeter's Twitter profile including the facts listed in the `twitter account` concept definition above, including their stated location where available;
2. a set of accounts with which an individual tweeter frequently interacts, via public "@" messages;
3. a set of accounts an individual tweeter talks about, via "@" mentions;
4. a set of accounts that influence an individual tweeter in terms of their retweeting of posts from those accounts;
5. a collection of recently-posted photos from an individual tweeter's account;
6. a set of term names derived from an ontology of concepts relating to crime and social disorder which are mentioned in the tweeter's recent posts (either directly by concept name or by a synonym);
7. a set of locations mentioned in the tweeter's recent tweets, which we attempt to contextualise in terms of whether the tweets indicate travel to or from those locations.

Items 2–4 indicate the tweeter's social network in terms of who they interact with and are influenced by. These relationships allow us to construct a wider network graph of related tweeters. Location information from their profile (item 1) and mentioned in their recent tweets (item 7) is referenced against a gazetteer of place names and, where matches are found, the names are converted from literal (text) values to the names of instances of the geospatial place concept. For each term name in item 6 found in one or more recent posts by the tweeter, we process the text of the tweet(s) using a simple sentiment classifier[§] in an attempt to determine the context in which the term was used: positively, negatively, or neutrally.

All of the above data extracted from the collected tweets is stored as CE facts in a knowledge base. As an illustration of how the information can be used, we built a prototype app which allows a user to look up details of an individual tweeter or tweet. Figure 4 shows two screenshots from this app, which was created as an extension to our Moira (Mobile Intelligence Reporting Agent) application that offers a Siri-like conversational natural language interface for common ISR tasks.[23] The figure shows elements 1, 5, 6 and 7 from the above list. Information on the individual's social network is omitted here.

Use of a CNL for knowledge representation allows us to rapidly extend our models at run time during a situation understanding exercise, and thereby enrich our knowledge and fact base. For example, a number of anti-war protests in South Wales in the summer of 2014 were linked to the area hosting the NATO Summit in September 2014. We ran an exercise to perform real-time situation awareness of local community reaction to the Summit. During this exercise it became important to extend our model of tweeters with information on their *stance* in relation to the Summit being held in the area: whether they were in favour of it, or opposed to it. Use of CE allowed us to make this extension in a way that maintained full engagement with colleagues who were subject matter experts but had no training in knowledge representation, because the formal machine-processable representation of the `stance` concept was readable and understandable to them. An example of the use of the `stance` concept via the Moira conversational interface is shown in Figure 5. (The example query is in relation to an individual who is a local political public figure.)

[§]http://www.datumbox.com

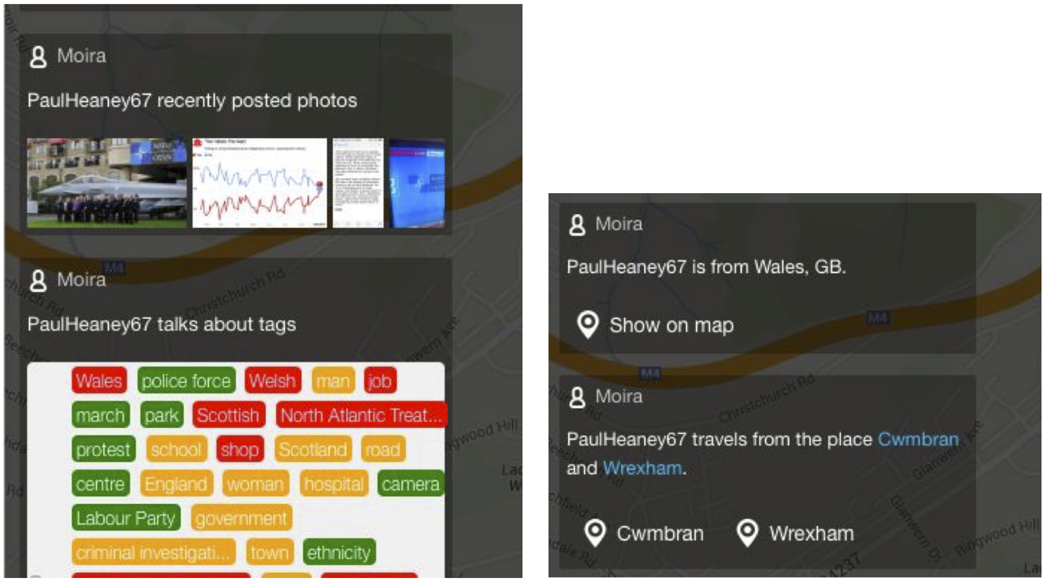

Figure 4. An example Moira query, showing some of the elements of the tweeter model

4. TASKING TWEETERS

In our previous work on assigning sensing assets to ISR mission tasks,[16] we characterised a `task` in CE as follows:

```
conceptualise the task T
  ~ requires ~ the intelligence capability IC and
  ~ is looking for ~ the detectable thing DT and
  ~ operates in ~ the spatial area SA and
  ~ operates during ~ the time period TP and
  ~ is ranked with ~ the task priority PR.
```

Here, instances of the `intelligence capability` concept include modalities such as *detect* and *localize*, while instances of `detectable thing` are drawn from an ISR ontology. In terms of directing the collection of Twitter data — i.e. the *direction* stage of the DCPD loop — the important elements here are the `detectable thing` and `spatial area` concepts. As we have seen in Section 2, a set of search terms that combines topics (for example, "protest" and "march") with spatial names (for example, "Cardiff") can be highly effective in obtaining useful

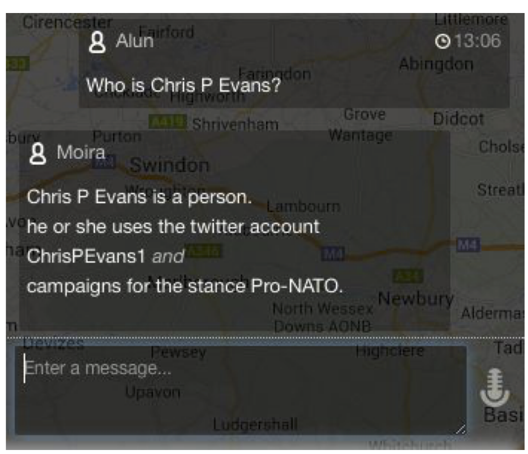

Figure 5. Use of the `stance` relationship in a conversation with Moira

tweets. A spatial area defined by geospatial coordinates can be applied to a Twitter search directly, though the volume of geotagged tweets is likely to be relatively small in comparison to those that are not geotagged, so a combination of terms and coordinates is always to be recommended. Thus, we can derive the specification of a Twitter data collection using the streaming API automatically from our existing ISR task representation.

In terms of the processing stage of the DCPD loop, the `intelligence capability` plays a role in what services to employ. For example, if the requirement is event detection, a variety of existing event detection and tracking algorithms exist.[6,12] If the requirement is localisation, then it may be possible to derive location data from the tweets (e.g., if some of them are geotagged) or by attempting to locate the tweeter.[24] Here then, the task representation tells a system which components to use in the pipeline illustrated in Figure 2.

The `time period` element of our CE `task` definition establishes temporal bounds on Twitter data collection. The `task priority` may be useful where resources for data collection are limited (given the high volumes of Twitter data this is certainly possible) and choices must be made on which collections are most important.

Experience from our own pilot study and studies of other incidents such as the Boston Marathon bombing and Lee Rigby murder[3] suggest that particular tweeters often become key sources of valuable information as events unfold, due to their being in a position to observe events, and their skills in Twitter use. The tweeter who was part of the protest march and reported the violence as it was happening ("Glasses thrown from Walkabout Cardiff Customers at #gazaj26 marchers. . . ", "More glasses thrown by customers of bars on Cardiff's Mill Lane at peaceful marchers. . . #gazaj26") is a good example of such an individual. This kind of tweeter tends to be identifiable through a combination of characteristics, drawing on elements of our model from the previous section:

- their use of spatial terms or geotagging: they want the online audience to know where they are, so take care to reveal their location over a series of tweets ("Walkabout Cardiff", "Cardiff's Mill Lane");
- their use of particular terms or hashtags: they want their tweets to be found and seen so take care to use words and tags that others are also using ("march", "#gazaj26");
- analysis of their social network: who they retweet, talk to and talk about; and who talks to, talks about and retweets them.

Our model addresses the first of these (location information) under items 1 and 7 listed in Section 3. The second (terms) is covered as item 6 (and can be extended to hashtags also). The third characteristic (social network analysis) derives from items 2–4. All of these characteristics, as well as being identifying features, are proxies for information quality: their location puts them in a "position to know", their use of words and tags relates to accuracy, and their social network characteristics point to influence and trust. (We will return to questions of information quality in Section 5.)

Going further in terms of Twitter data processing, we are experimenting with a bag-of-words natural language processing approach to fact extraction from tweet text, building on techniques developed and tested in a crowdsourcing study we performed in earlier work.[23] The goal here is to extract information in the form of CE facts so that it can be used in further processing and inference. Currently we offer the user an opportunity to confirm that the extracted information is correct before it is committed to the knowledge base — an example exchange using the Moira agent is shown in Figure 6. Our interface allows the user to select any tweet for attempted interpretation by the agent.

The example here is one that we created on a private account, during a piece of fieldwork in connection with the NATO Summit in September 2014. This mode of interaction would be intended for use only with small numbers of significant tweets (the ones in relation to the Mill Lane violence being good examples) where there is value to be gained from automating the process of rapidly adding the equivalent CE facts to the knowledge base. A use case for doing so would be if the user — for example, a police analyst — wanted to quickly disseminate the information to relevant patrols, or to automatically task sensing assets — for example, CCTV cameras — to obtain further information on the event reported in the tweet.

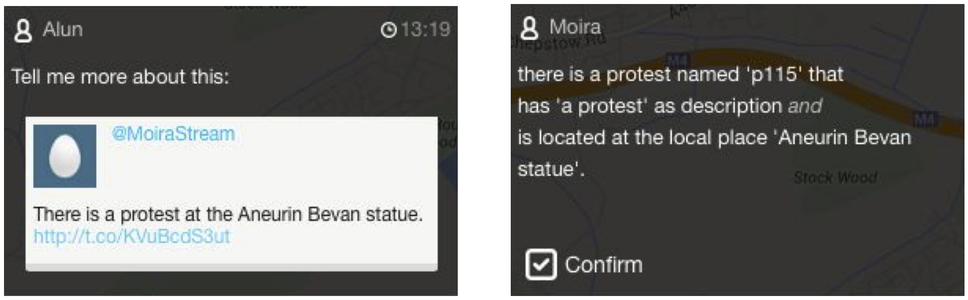

Figure 6. An example of fact extraction from tweet text using Moira

5. DISCUSSION AND CONCLUSION

In this paper we have shown how the use of streamed social media can be incorporated into the traditional ISR asset management cycle, and specifically how its use can be aligned with previous work in the automated assignment of sensing assets to ISR tasks. This facilitates a greater degree of automation in the use of social media streams as an additional source of intelligence. We believe our approach is compatible with a number of existing systems for social media monitoring as described in the literature.[2,4,5,7,9]

A significant issue in the use of all open source intelligence is the potential for misinformation.[18] There is work being done to mitigate these risks[5] but we also observe that patterns of misinformation flow are often valuable in terms of situational awareness: for example, rumouring is often a form of coordinated activity that needs countering.[7]

We have focussed here on the use of social media for text-based soft information, but often there is even more value in the imagery data attached to tweets in the form of photos and videos. The application of image processing to these sources to extract key features — particularly common symbols and objects such as weapons, as well as face recognition — offers considerable potential. In effect, social media is a source of both hard and soft data, leading to interesting challenges in data fusion.[13,22]

Some of our recent work has focussed on viewing ISR information pipelines as bidirectional chains, where humans and machines work in collaboration.[23] We are currently conducting experiments with human subjects using the Moira agent to facilitate crowdsourced situational understanding. The experiments include combining the use of human and physical sensors, dealing with issues of incomplete and conflicting information, and maximising the value of scarce resources by using relevancy criteria.

Acknowledgements

This research was sponsored by the US Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defence or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

Development of the Sentinel platform was funded by the European Commission under the project "Tackling Radicalisation in Dispersed Societies (TaRDiS)", and the ESRC via the project "After Woolwich: Social Reactions on Social Media" (ES/L008181/1). Cardiff University provided funding for the pilot study examining community impacts of the NATO Summit.

We thank Kieran Evans and David Rogers (Cardiff University) for setting up the data collection pipeline for the pilot study in Section 2 and assistance with the data analysis. We thank Darren Shaw (IBM Emerging Technology Services, UK) for creating the tweeter locator service in Section 3. Valuable insights on policing and community reaction to events such as the ones featured in our pilot study were provided by Martin Innes, Colin Roberts and Sarah Tucker (Cardiff Universities Police Science Institute, http://www.upsi.org.uk).

REFERENCES

1. M. Srivastava, T. Abdelzaher, and B. Szymanski, "Human-centric sensing," *Phil. Trans. R. Soc. A* **370**(1958), pp. 176–197, 2012.
2. M. Osborne, S. Moran, R. McCreadie, A. Von Lunen, M. D. Sykora, E. Cano, N. Ireson, C. MacDonald, I. Ounis, Y. He, T. Jackson, F. Ciravegna, and A. O'Brien, "Facebook, Twitter and Google Plus for breaking news: Is there a winner?," in *Proceedings of the AAAI Conference on Web and Social Media*, 2014.
3. M. Innes, C. Roberts, and D. Rogers, "Critical timing," *Police Professional* **January**, pp. 17–18, 2014.
4. F. Abel, C. Hauff, G.-J. Houben, R. Stronkman, and K. Tao, "Semantics + filtering + search = Twitcident. exploring information in social web streams," in *Proceedings of the 23rd ACM Conference on Hypertext and Social Media*, pp. 285–294, 2012.
5. D. Wang, M. T. Amin, S. Li, T. Abdelzaher, L. Kaplan, S. Gu, C. Pan, H. Liu, C. C. Aggarwal, R. Ganti, X. Wang, P. Mohapatra, B. Szymanski, and H. Le, "Using humans as sensors: An estimation-theoretic perspective," in *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks, IPSN '14*, pp. 35–46, 2014.
6. M. Osborne, S. Moran, R. McCreadie, A. Von Lunen, M. D. Sykora, E. Cano, N. Ireson, C. MacDonald, I. Ounis, Y. He, T. Jackson, F. Ciravegna, and A. O'Brien, "Real-time detection, tracking, and monitoring of automatically discovered events in social media," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pp. 37–42, 2014.
7. C. Roberts, M. Innes, A. Preece, and I. Spasić, "Soft facts and spontaneous community mobilisation: the role of rumour after major crime events," in *Data for Good: How big and open data can be used for the common good*, P. Baeck, ed., pp. 37–43, Nesta, 2015.
8. G. Pearson and T. Pham, "The challenge of sensor information processing and delivery within network and information science research," in *Proc Defense Transformation and Net-Centric Systems 2008 (SPIE Vol 6981)*, SPIE, 2008.
9. L. M. Aiello, G. Petkos, C. Martin, D. Corney, S. Papadopoulos, R. Skraba, A. Gker, I. Kompatsiaris, and A. Jaimes, "Sensing trending topics in Twitter," *IEEE Transactions on Multimedia* **15**(6), pp. 1268–1282, 2013.
10. T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes Twitter users: Real-time event detection by social sensors," in *Proceedings of the 19th International Conference on World Wide Web*, pp. 851–860, 2010.
11. R. Socher, A. Perelygin, J. Y. Wu, J. Chuang, C. D. Manning, A. Y. Ng, and C. Potts, "Recursive deep models for semantic compositionality over a sentiment treebank," in *Proc Empirical Methods in Natural Language Processing (EMNLP 2013)*, pp. 1631–1642, 2013.
12. K. N. Vavliakis, A. L. Symeonidis, and P. A. Mitkas, "Event identification in web social media through named entity recognition and topic modeling," *Data and Knowledge Engineering* **88**, pp. 1–24, 2013.
13. J. Llinas, C. Bowman, G. Rogova, A. Steinberg, E. Waltz, and F. White, "Revisiting the JDL data fusion model II," in *Proc Seventh International Conference on Information Fusion (FUSION 2004)*, pp. 1218–1230, 2004.
14. A. Preece, T. Norman, G. de Mel, D. Pizzocaro, M. Sensoy, and T. Pham, "Agilely assigning sensing assets to mission tasks in a coalition context," *IEEE Intelligent Systems* **Jan/Feb**, pp. 57–63, 2013.
15. T. Kuhn, "A survey and classification of controlled natural languages," *Computational Linguistics* **40**, pp. 121–170, 2014.
16. A. Preece, D. Pizzocaro, D. Braines, and D. Mott, "Tasking and sharing sensing assets using controlled natural language," in *Proc Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR III (SPIE Vol 8389)*, SPIE, 2012.
17. D. Gayo-Avello, "A meta-analysis of state-of-the-art electoral prediction from twitter data," *Social Science Computer Review* **31**(6), pp. 649–679, 2013.
18. F. Jin, W. Wang, L. Zhao, E. Dougherty, Y. Cao, C.-T. Lu, and N. Ramakrishnan, "Misinformation propagation in the age of twitter," *IEEE Computer* **47**(12), pp. 90–94, 2014.
19. D. Mott, "Summary of ITA Controlled English," 2010.
20. "OWL 2 Web Ontology Language document overview (second edition)." World Wide Web Consortium, Dec. 2012.

21. "SWRL: A semantic web rule language combining OWL and RuleML." World Wide Web Consortium, May 2004.
22. A. Preece, D. Pizzocaro, D. Braines, D. Mott, G. de Mel, and T. Pham, "Integrating hard and soft information sources for D2D using controlled natural language," in *Proc 15th International Conference on Information Fusion*, 2012.
23. A. Preece, C. Gwilliams, C. Parizas, D. Pizzocaro, J. Z. Bakdash, and D. Braines, "Conversational sensing," in *Proc Next-Generation Analyst II (SPIE Vol 9122)*, SPIE, 2014.
24. J. Mahmud, J. Nichols, and C. Drews, "Home location identification of twitter users," *ACM Transactions on Intelligent Systems and Technology* **5**(3), pp. 47:1–47:21, 2014.