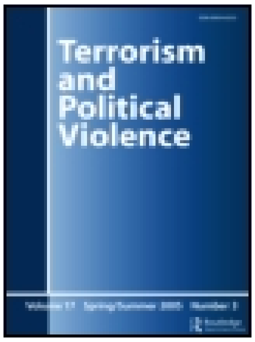

Terrorism and Political Violence

ISSN: 0954-6553 (Print) 1556-1836 (Online) Journal homepage: http://www.tandfonline.com/loi/ftpv20

Natural Language Understanding and Multimodal Discourse Analysis for Interpreting Extremist Communications and the Re-Use of These Materials Online

Peter Wignell, Kevin Chai, Sabine Tan, Kay O'Halloran & Rebecca Lange

TERRORISM AND POLITICAL VIOLENCE
https://doi.org/10.1080/09546553.2018.1520703

Routledge
Taylor & Francis Group

Check for updates

Natural Language Understanding and Multimodal Discourse Analysis for Interpreting Extremist Communications and the Re-Use of These Materials Online

Peter Wignell, Kevin Chai, Sabine Tan, Kay O'Halloran and Rebecca Lange

Faculty of Humanities, Curtin University, Perth, Australia

ABSTRACT
This paper reports on a study that is part of a project which aims to develop a multimodal analytical approach for big data analytics, initially in the context of violent extremism. The findings reported here tested the application of natural language processing models to the text of a sample of articles from the online magazines *Dabiq* and *Rumiyah*, produced by the Islamic extremist organisation ISIS. For comparison, text of articles found by reverse image search software which re-used the lead images from the original articles in text which either reported on or opposed extremist activities was also analysed. The aim was to explore what insights the natural language processing models could provide to distinguish between texts produced as propaganda to incite violent extremism and texts which either reported on or opposed violent extremism. The results showed that some valuable insights can be gained from such an approach and that these results could be improved through integrating automated analyses with a theoretical approach with analysed language and images in their immediate and social contexts. Such an approach will inform the interpretation of results and will be used in training software so that stronger results can be achieved in the future.

KEYWORDS
data analytics; Islamic State/ISIS; multimodal discourse analysis; natural language understanding; social semiotics

Introduction and background

This paper reports on work that is part of a larger project, which aims to develop a multimodal analytical approach (see theoretical framework below) for big data analytics, initially in the context of violent extremism. By applying and advancing methods from multimodal analysis, the larger project addresses online violent extremist discourse, where the most salient feature is not individual uses of text or image, but the ways in which the text is often made an integral part of the image and vice versa[1] and, "similar to conventional social multimedia, the messages exchanged in 'extremosphere' typically consist of text and visual content."[2] In order to more fully understand how online violent extremist discourses operate as forces for radicalisation, the issue of the meanings arising from the integration of language and images thus needs to be addressed.

There is a growing body of work on the use of images by Islamic violent extremist groups across mainstream and social media platforms[3] and on the re-use of images across platforms.[4] Some of this work has also examined how images and text combine in material

CONTACT Kay O'Halloran kay.ohalloran@curtin.edu.au Curtin University, Perth 6845, Australia
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/ftpv.

produced by Islamic violent extremist organisations.[5] Analysis of the text in which the images appear is less abundant.[6] Recontextualisation of images in relation to the texts in which they reappear appears not to have been examined.

This study aims to work towards remedying this imbalance by focusing on the text that accompanies the images and in which the images are embedded. The trial aims to identify patterns of variation in text in terms of co-contextualising (similar meanings) and re-contextualising relations (changed meanings) by analysing the text associated with images in their original context in materials produced by the Islamic extremist organisation which refers to itself as Islamic State (hereinafter referred to as ISIS) and the text associated with images when they are recontextualised. The follow-up phases of the larger project will involve investigating the results of the automated analysis of images and the integration of automated text and image analyses.

ISIS has produced prolific amounts of propaganda material. Likewise, ISIS's activities have generated widespread coverage in both mainstream and social media. An important issue that arises in relation to the analysis of this material is that such analysis is usually based on either a detailed examination of a relatively small amount of data or on a relatively superficial analysis of larger quantities of data. This paper attempts to reconcile these approaches by adopting a method which uses readily accessible natural language processing software to analyse a quantity of text material and applies a consistent theoretical approach to the interpretation of the results of this analysis, using this to evaluate the results and interpretations of the analysis as they stand and offering suggestions for possibly improving the results of such larger-scale analyses.

Material produced by ISIS can be assumed to be propaganda in support of violent extremism and as such provides a baseline for identifying violent extremist and pro-violent extremist discourse. Not all text produced by ISIS is violent extremist in nature. For example, *Dabiq* magazine contains articles about subjects such as provision of public services to citizens of the "caliphate" and on building the apparatus required to run a "state." These texts are, however, parts of the ISIS agenda and violent extremism is integral to their implementation.[7] Material which reports on and discusses ISIS activities represents views ranging from support, more or less straight reporting on ISIS activities, to active antagonism. Almost all of this material is multimodal in nature, consisting of combinations of text and images.

Empirical analysis and testing of observations are required to map discourse patterns and trends in order to identify similarities and differences between these sources. However, the sheer volume of material represents a serious obstacle for analysts. It is simply not possible to manually analyse such a volume of material. One way around this issue is to develop tools for the automated analysis of text, images, and how they combine to make meaning. Large-scale analyses of language and images tend to consider the potential meanings of these modalities separately, whereas in actual instances of use, the meanings arise from a complex integration of the contributions made by images and language.

This study represents a first step: testing the possibilities of using automated text analysis to identify features and combinations of features of language which might be used to help distinguish between material which actively promotes terrorism, material which does not promote violent extremism, material which might inadvertently promote violent extremism, and material related to Islam which is unrelated to violent extremism. A further aim is to investigate the limits of readily available automated text analysis tools and to suggest possible modifications to improve performance. Previous studies using natural language processing to investigate violent Islamic extremist language have concentrated on material produced by violent extremist groups.[8] These studies have been lexically focused, based on frequencies, collocations, and relationships among words. Where this current study differs is that its intention is to build a stratified analytical and interpretive model which incorporates context on several scales and connects context to patterns of linguistic choices in texts produced by ISIS and in texts which report on material produced by ISIS.

Theoretical framework

While the focus of this study is on language, the larger project focuses on meanings which arise from the combination of language and images in text. This approach is referred to as multimodal discourse analysis, which is the study of meanings arising from the integration of language, images, and other resources in texts, interactions, and events. Multimodal discourse analysis has emerged as an interdisciplinary field of research that provides new and developing frameworks for analysing how language and images combine to communicate meaning.[9] The theoretical approach is based on social semiotics, which studies how sign systems are used to create meaning in context.[10] The specific approach to social semiotics and multimodal discourse analysis used here is referred to as Systemic Functional Multimodal Discourse Analysis.[11] This approach is derived from the application of Michael Halliday's systemic functional theory, where language, images, and other resources are viewed as resources for making meaning.[12] Such meaning-making resources are formalised as networks of options from which choices are made. Specifying those options and showing how they are made in particular artefacts gives a powerful, contextually-motivated method for characterising the communicative effects when texts and images are combined. Systemic functional theory, when applied to the analysis of language, is referred to as Systemic Functional Linguistics.[13]

A foundational principle of a systemic functional approach is that resources for making meaning are organised according to the functions which they have evolved to serve. These functions are themselves organised into three fundamental strands of meaning, which are referred to as metafunctions.[14] These are: a) ideational meaning, which relates to meanings about things like people, objects, actions, and places in the material world and how events are connected to each other logically in text; b) interpersonal meaning, which is concerned with meanings about the social nature of communication, for example different communicative choices made by people according to their relative social status or power in an interaction; and c) textual meaning, which connects text to its context and organises messages into coherent forms. The messages in any communicative situation are characterised along these dimensions. That is, every act of meaningful communication will be about something, involve some kind of relationship among the participants, and be organised so that it works as cohesive and coherent text. The systems operate across different scales. For example, they involve grammatical choices at the level of the sentence and choices which operate cohesively over larger units of text. That is, the metafunctions are also realised through choices at different ranks and at different scales (see Figure 1).[15] With regard to textual meaning, the flow of information in, say, a newspaper article can be viewed on several scales. For example, the relationship between the headline and introductory paragraph of an article to the rest of the article is the same as the relationship between the topic sentence of a paragraph to the rest of the paragraph: each sets the scene for what is developed in what follows.

In the model shown in Figure 1, what is referred to as register is made accessible through configurations of three variables: field, tenor, and mode.[16] Field refers to meanings about what is happening, what kind of human activity the text is about, and is realised primarily through choices from ideational meaning. Field addresses questions about "what." Tenor refers to the kinds of social relationships a text is constructing, such as the relative status of participants in a text or activity, and is realised primarily through choices from interpersonal meaning. Tenor addresses questions about "who." Mode is realised primarily through choices from textual meaning and refers to the part language is playing in the construction of a text, how information flows through a text, and how a text is organised in relation to its context and addresses questions about "how."

Register links language to its contexts and is realised in language through choices from systems at lower ranks: grammatical and discourse-semantic systems. Meanings can be realised in the grammar, for example, in sentences and in groups of words (e.g., a noun and all the words modifying that noun). Discourse semantics refers to cohesive links which carry meanings across and through passages of text larger than a clause or sentence (e.g., taxonomic relations among words in a text or in using pronouns to track participants in a text). For example, to find the identity of *she* in a text, it is necessary to go beyond the sentence *she* is in. Patterns of choices from grammar and discourse semantics allow a text's field, tenor, and mode[17] to be deduced. Texts which exhibit recurring patterns of similar choices can be regarded as belonging to the same register. For example, the register of a casual conversation between friends will be different from the register of a hard news report. In short, meanings at higher ranks are made accessible through patterns of choices from lower ranks. From the patterns of these combinations of choices, the ideological position of a text can also be inferred.

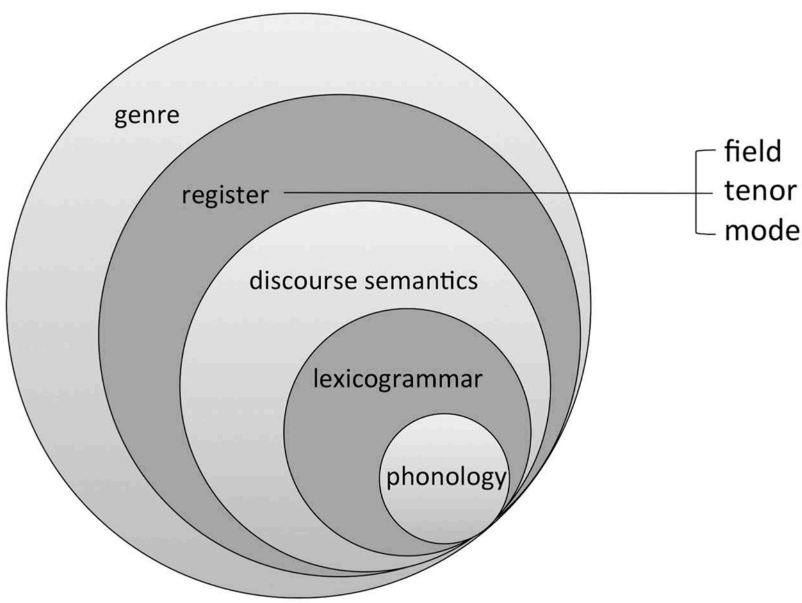

Figure 1. Stratified model of language and context (From Martin, 2014, p. 14).

These systems, when combined, produce an integrated model of language and how it relates to its contexts, both its immediate context and to its broader social and cultural context. Approaches derived from the same theoretical model have been applied to the analysis of images.[18] It is hypothesised that a stratified metafunctional approach, if it could be even partly automated, could provide a major breakthrough in the analysis of discourse patterns in very large data sets.

Connections between systemic functional multimodal discourse analysis and natural language and image processing are developing and have made notable advances over the past decade,[19] but are not yet at a stage where they can be applied to automated analysis of large data sets. For this reason, it was decided to trial only the automated analysis of language, use readily available and accessible natural language processing tools, and interpret the results through a systemic functional lens.

Method

Islamic violent extremist discourse is most clearly exemplified and articulated in material produced by ISIS. ISIS has clearly articulated its agenda, beliefs, goals, and strategies in its online magazines *Dabiq* and *Rumiyah*. These magazines thus represent a data set of material in support of and promoting violent extremism which has been produced by an Islamic extremist organisation. The initial data set of Islamic extremist material used in this trial consisted of the 15 issues of *Dabiq* magazine and 12 issues of *Rumiyah* magazine (*Rumiyah* replaced *Dabiq* in September, 2016). This data set consisted of 519 articles: 290 articles in *Dabiq* and 229 articles in *Rumiyah*. These articles had been classified into 21 article types, ten of which are common to both magazines, six are found in *Dabiq* but not in *Rumiyah*, and five are found in *Rumiyah* but not *Dabiq*.[20] From this initial set, one typical example of each of the article types was selected for analysis. The articles all promote ISIS's agenda through texts which report on ISIS's military action, incite lone-wolf terror attacks, encourage recruitment and migration, and provide religious "justification" of ISIS's agenda. While these articles differ somewhat in content and style, they all have one thing in common: they promote and incite violent extremism.

From each article the lead image, when there was one, was chosen. Where there was no clear lead image either one typical image was chosen or, in the case of articles such as infographics, where the whole text is an image, the whole text was used. These images were then used to identify online texts which re-used the images from the *Dabiq* and *Rumiyah* articles in different contexts.

A reverse image search using TinEye (https://www.tineye.com/) was conducted on the 21 images to find examples of their redistribution after publication in *Dabiq* or *Rumiyah*. This search provided all instances that TinEye could find of re-use of the images. One example of a recontextualisation of the image from each article type was selected. The recontextualisations covered a range of media, such as mainstream news services, political commentary/analysis sites, blogs, and tweets. The text of each example article from *Dabiq* and *Rumiyah* and the text of each example of a recontextualisation was extracted for automated analysis. Table 1 shows the sources of the original articles from *Dabiq* and *Rumiyah* and the type of site and URL of where the recontextualisations chosen for analysis were found. In some cases, comments are made about the type of content on the site.

Table 1. Sources of original articles and articles in *Dabiq* and *Rumiyah* and sources of articles with recontextualised images.

	Article Type	Magazine and Issue	Site type and URL of Recontextualisation
1	Cover/Table of Contents	*Rumiyah*, 2	Twitter http://linkedin.com/in/rita-katz
2	Cover	*Dabiq*, 7	Clarion Project: Political commentary https://clarionproject.org/islamic-state-isis-isil-propaganda-magazine-dabiq-50/
3	Table of Contents	*Dabiq*, 2	Intermedia Education Project: 'information/propaganda site: negative towards Islam http://www.ieproject.org/projects/dabiq2.html
4	Foreword	*Dabiq*, 8	Sydney Morning Herald: mainstream newspaper http://www.smh.com.au/federal-politics/political-news/terror-plot-teenagers-linked-to-top-islamic-state-recruiter-abu-khalid-alkambodi-20150419-1mock5.html
5	Hikmah (Wisdom)	*Dabiq*, 5	Tomfernandez28s Blog: blog https://tomfernandez28.com/2015/11/27/german-police-chief-isis-hiding-among-refugees-entering-europe/
6	Among the Believers are Men	*Dabiq*, 7	Daily Mail: mainstream newspaper http://www.dailymail.co.uk/news/article-3435947/Second-member-ISIS-Beatles-QPR-fan-west-London-Alexanda-Kotey-32-identified-spies-group-four.html
7	To/From Our Sisters/ For Women	*Dabiq*, 9	Business Insider: mainstream magazine https://www.businessinsider.com.au/isis-is-recruiting-in-the-most-perverse-way-imaginable-2015-8?r=US&IR=T
8	From the Pages of History	*Dabiq*, 15	Fox News: mainstream news http://www.foxnews.com/world/2016/01/21/new-issue-isis-magazine-dabiq-calls-for-war-on-muslims.html?intcmp=trending
9	Interview	*Dabiq*, 12	Site Intelligence Group: news/political commentary: focus on Islamic extremism https://news.siteintelgroup.com/Jihadist-News/dabiq-12-is-reveals-bomb-used-on-russian-airliner-execution-of-chinese-and-norwegian-hostages-cantlie-returns.html
10	Near Enemy Issues	*Dabiq*, 6	Daily Mail: mainstream newspaper http://www.dailymail.co.uk/news/article-3300562/Al-Qaeda-leader-calls-new-9-11-strikes-against-praises-Palestinian-knife-attacks-Israelis-new-audio-message-fanatics.html
11	In the Words of the Enemy	*Dabiq*, 9	Washington Post: mainstream newspaper https://www.washingtonpost.com/world/asia_pacific/in-qatar-us-taliban-talks-remain-on-the-line/2013/06/22/f45e381e-db7d-11e2-b418-9dfa095e125d_story.html?utm_term=.39df9821ec1b
12	Feature Articles	*Dabiq*, 15	The Counter Jihad Report: political commentary: focus on Islamic extremism https://counterjihadreport.com/category/dr-sebastian-gorka/
13	Far Enemy Captives	*Dabiq*, 3	Crime Files Network: "library" of crime reports http://crimefiles.net/category/religious-crimes/page/2/
14	John Cantlie	*Dabiq*, 9	New York Daily News: mainstream newspaper http://www.nydailynews.com/news/world/isis-nukes-allowed-consolidate-expert-article-1.1958855
15	ISIS Reports	*Dabiq*, 5	Federation for Defense of Democracy's Long War Journal: Political commentary/news site: focus on Islamic extremism https://www.longwarjournal.org/archives/2014/10/islamic_state_releases_picture_2.php
16	Advertisements	*Rumiyah*, 6	No recontextualisation found
17	Eulogy/Obituary	*Rumiyah*, 1	Boston Globe: mainstream newspaper http://www.bostonglobe.com/news/world/2016/08/31/slain-figure-was-powerful-leader-with-multiple-roles/M4dFMxLF1BNAHDAHQMgDoM/story.html
18	Infographic Religious	*Rumiyah*, 1	Cracked: humour, "oddities" "library" site http://www.cracked.com/personal-experiences-2425-how-isiss-new-magazine-inspired-ohio-attack.html
19	Infographic Military	*Rumiyah*, 11	Twitter https://twitter.com/hashtag/tehranshooting?lang=en
20	Procedural	*Rumiyah*, 3	Business Insider: mainstream magazine https://www.businessinsider.com.au/isis-calls-for-vehicle-knife-attacks-2016-11?utm_content=allverticals&utm_medium=referral&utm_source=hearst&r=US&IR=T
21	Last Page Message	*Dabiq*, 7	Anti-Defamation League: "liberal" civil rights, news commentary http://austin.adl.org/oren-segal-director-of-center-on-extremism-visits-austin/

The complete text of 40 sample articles (one from each article type in *Dabiq* and *Rumiyah*, with the exception of advertisements, for which no recontextualisations were found) and the text from each of the articles circulated on the internet with recontextualised images was analysed using the IBM Watson Natural Language Understanding (NLU) online demonstration models (https://natural-language-understanding-demo.ng.bluemix.net/). The Watson NLU tool performs a number of automated Natural Language Processing (NLP) tasks which include sentiment analysis opinion mining,[21] emotion detection,[22] keyword extraction,[23] named entity recognition,[24] text categorisation,[25] concept extraction/discovery,[26] and semantic role labelling.[27] Many of these NLP techniques employ machine learning algorithms that are also used for dealing with nontextual data such as image classification, object detection, speech recognition, and data visualisation (e.g., representing the relationship between terms) but can be slightly customised to perform better with the textual data.

These natural language processing tasks have been relabelled in the IBM Watson NLU tool from their academic terms to *sentiment, emotion, keywords, entities, categories, concept, semantic roles, metadata*, and *relations*. (The NLU models *metadata* and *relations* are not available in the online demonstration.) A general description of the NLU analytical models are provided online, but the criteria, rationale, and algorithms are not provided; that is, the algorithms are a "black box." In what follows, we provide a description of the results obtained when these models are applied for text analysis, and how they relate to the SFL model of language.

Figure 2 shows screenshots of the analysis of *categories, concepts, entities*, and *keywords* carried out with the IBM Watson Natural Language Understanding (NLU) online demonstration models for the sample of analysed texts. The results of *emotion* and *sentiment* analysis are shown in Figure 3.

Categories analyse text according to a hierarchy with a maximum of five levels. The demo, however, identifies three categories for each sample of text. *Categories* assign the text to a field (and sub-fields) of human activity and provide a score of between 0 and 1 for relevance. Figure 2(a) shows a three-level categorisation hierarchy. Interpreted through a systemic functional lens, *categories* correlate closely with the register variable of field. *Categories* situate the text within a domain of human activity and show in broad terms what the text is about. In the example in Figure 2(a), the three categories represent three intersecting domains of human activity. It is intersections and combinations such as these that are likely to prove useful in helping to identify violent extremist discourse.

Concept is used to "identify high-level concepts that aren't necessarily directly referenced in the text" (https://console.bluemix.net/docs/services/natural-language-understanding/index.html#about). Figure 2(b) shows *concept*s for the same text used in Figure 2(a). *Concept* also relates to domains of human activity but is more specific as it identifies particular organisations, places, and individuals related to the domains shown in *categories*. Identification as a *concept* does not depend on something being specifically found in the text or on the frequency of something being found. For example, the word *Islam* appears as a *concept* but is not specifically found in the text (although *Islamic* is found). *Caliphate* appears in the text once and is identified as a *concept*, whereas *jihad* appears in the text eight times but does not appear as a *concept*.

Entities refers to individual people, specific places, and specific organisations in a text. It mainly identifies proper nouns in a text and classifies them according to type of *entity* (see

Hierarchy	Score
/ religion and spirituality / islam	0.47
/ law, govt and politics / government	0.36
/ society / unrest and war	0.27

(a): *Categories* hierarchy for text of 'In the Words of the Enemy', *Dabiq* Issue 9, pp. 60-64

Concept	Score
Al-Qaeda	0.98
Osama bin Laden	0.85
Syria	0.69
Taliban	0.68
Iraq	0.57
Ayman al-Zawahiri	0.56
Caliphate	0.55
Islam	0.53

(b): *Concepts* for the text of 'In the Words of the Enemy', *Dabiq* Issue 9, pp. 60-64

Name	Type	Score
Islamic State	Organization	0.86
crusaders	Organization	0.56
United States	Location	0.49
Syria	Location	0.40
ENEMY	Organization	0.32

(c): Top 5 *entities* from 'In the Words of the Enemy', *Dabiq*, Issue 9, pp. 60-64

Text	Relevance
Islamic State	0.94
al qaeda	0.79
Free Syrian Army	0.62
jihād claimants	0.57
United States	0.55

(d): Top 5 *keywords* from 'In the Words of the Enemy', *Dabiq*, Issue 9, pp. 60-64

Figure 2. Screenshots of analysis of categories, concepts, entities, and keywords performed with IBM Watson Natural Language Understanding (NLU) online demonstration models.

Figure 2(c)). *Entities* frequently misclassifies the type of *entity*, especially when the name of the *entity* is named in Arabic.

Keywords for the most part appears to identify other nouns in the text, although there is some overlap with *entities*. In Figure 2(c) and (d), the Islamic State and the United States

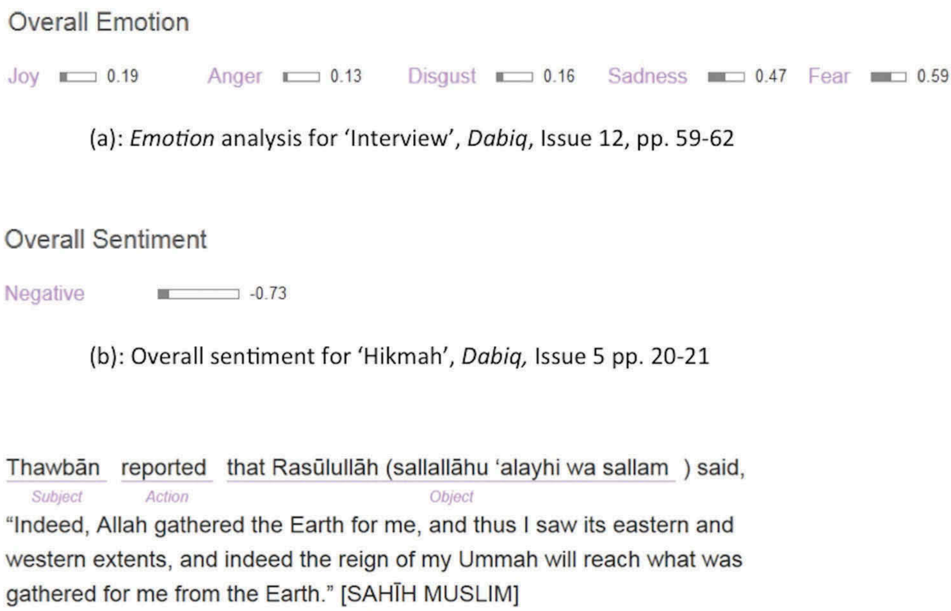

(a): *Emotion* analysis for ‘Interview’, *Dabiq*, Issue 12, pp. 59-62

(b): Overall sentiment for ‘Hikmah’, *Dabiq,* Issue 5 pp. 20-21

(c): *Semantic roles* in the first sentence of ‘Hikmah’, *Dabiq,* Issue 5 pp. 20-21

Figure 3. Screenshots of analysis of emotion, sentiment, and semantic roles analysis.

appear as *entities* and also appear *as keywords*. The analysis of keywords frequently overlaps with *entities*, with many *entities* also appearing as *keywords* but less so the other way around as *entities* focuses on proper nouns. In addition, occasionally *keywords* identifies a group of words, such as “explosive Islamic movement” or noun and verb combinations such as “Allah humiliates kufr.” However, these are rare. *Keywords* are given a weighting of between 0 and 1 according to how relevant they are considered to be. Lists of keywords can be very long. In this trial, only the top 5 weighted keywords for each text are used. Both *entities* and *keywords* relate to ideational meaning.

Emotion is described as being able to “analyze emotion conveyed by specific target phrases or by the document as a whole” (https://console.bluemix.net/docs/services/natural-language-understanding/index.html#about). The option for analysing the whole text assigns a weighting of between 0 and 1 for each of five emotions: joy, anger, disgust, sadness, and fear. Figure 3(a) shows the *emotion* analysis for the text of Example 9 in Table 1. According to the weightings, the principal emotions realised in the text are sadness and fear. There is also an option in the demonstration models for what is referred to as targeted emotion. However, this requires phrases to be entered individually. Whether it is possible to automate this option for large stretches of text has not yet been explored.

Sentiment analyses positive or negative dispositions in the text. It does this either by allocating a negative or positive value of between –1 and 1 to the whole text or by applying targeted sentiment to particular participant and process combinations in the text. In the demo, targeted sentiment requires individual wordings to be analysed one at a time. For this trial, sentiment analysis is applied only to the whole text (see Figure 3(b)). Figure 3(b) shows a weighting for negative emotion of –0.73, but does not provide information about who is being negative or what or who they are negatively disposed towards. Like *emotion, sentiment* relates to interpersonal meaning. The option to analyse for *semantic roles* was not used in this trial. An example of the semantic roles analysis for the first sentence of the text used in Figure 3(b) is shown in Figure 3(c). If an analysis of semantic roles could be automated and applied using a grammatical analysis that was part of a unified theoretical approach, it could prove useful in not only providing information about who is doing what to whom but also in interpreting the direction of emotion and sentiment.

Findings

The following section discusses the results of the automated analysis of all the texts from both sets of data. Following this, observations about the value of the analyses and suggestions for improving such analyses are outlined.

The results of the analyses were tabulated and the information in these tables was converted into visualisations which display patterns in the data in forms more accessible than in the tables. Observations drawn from the summaries of the analyses are offered as suggestions for further exploration with the aim of identifying what in these automated analyses might be worth pursuing further and how the automated analyses might be improved.

Categories

Figure 4 displays the results for *categories* for the original texts from *Dabiq* and *Rumiyah* and for the recontextualisations in a Sankey diagram. The diagram shows the article types from *Dabiq* and *Rumiyah* vertically in the centre. The two most delicate levels of the three-level category hierarchy are shown for a) the original articles (on the left) and b) for the recontextualisations (on the right). A screenshot of the interactive diagram is shown here. The interactive diagram shows both the relative values of each category in each set of texts and the relationship of categories to the article types in *Dabiq* and *Rumiyah*. This relationship is potentially important as a correlation between categories and article types could prove to be a strong indicator of what type of text is most likely to most strongly function as ISIS propaganda. This is even more likely to be the case once a correlation between features of images and features of text is firmly established.

Figure 4 shows that there is one *category* that is found in almost all texts in both sets of data. The *category* Religion and spirituality/Islam is identified in 19 of the 20 texts from *Dabiq* and *Rumiyah* and in 16 of the 20 examples of recontextualisation. The odd example from the ISIS data is from a field report of military activity from Issue 5 of *Dabiq*. These reports, while still glorifying ISIS, tend to be written by people in the field rather than by ISIS media people. It is not clear why Religion and spirituality/Islam are not identified as a *category* in the four recontextualised texts, although all use words such as ISIS, jihad, Islamic, ISIS terrorists, terrorist attacks, or *Rumiyah*. Perhaps if all five levels of the *category* hierarchy were shown, this *category* would appear.

That said, the category Religion and spirituality/Islam shows promise in distinguishing between text which is possibly related to Islamic violent extremism (either for or against) and text which is not related. It would, however, also include text from mainstream, non-violent, moderate Islamic sources. If used in combination with the *category* society/unrest and war, this would help to distinguish a) Islamic extremist discourse from other non-extremist

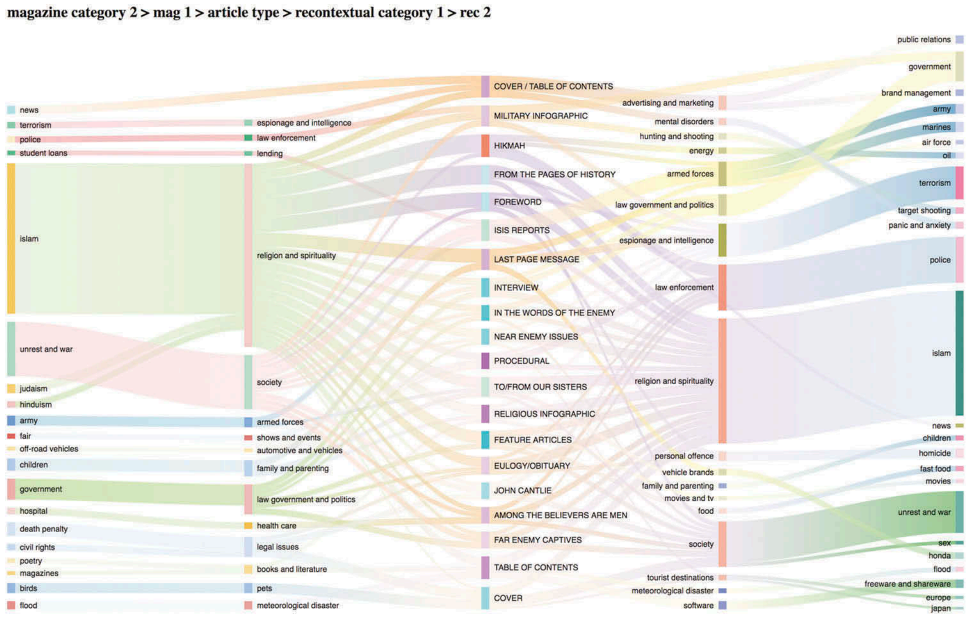

Figure 4. Screenshot of Sankey diagram showing results of categories analysis of *Dabiq* and *Rumiyah* texts and texts containing recontextualised images from those texts (https://curtinic.github.io/mma-violent-extremism/sankey_nlp_categories.html).

Islamic discourse; and b) to help identify discourse in the field of Islamic extremism, both promoting/supporting and resisting/reporting on. It is also possible that *categories* such as police and terrorism, which are found in the recontextualisations but either not found or only minimally found in the texts produced by ISIS, can be used to aid in distinguishing between material produced by ISIS and material produced in response to ISIS material.

Concepts

The Sankey diagrams in Figures 5 and 6 display the results of the automated analysis according to *concept*. Figure 5 displays *concepts* identified in *Dabiq* and *Rumiyah* and Figure 6 identifies concepts identified in the recontextualisations. The figures show the distribution of *concepts* across article types.

In contrast, Figure 7(a) and 7(b) show word clouds which aggregate the concepts and show their relative proportions across all of the texts. This provides a general pattern for the proportions of concepts across the whole sample of texts.

The analysis of *concepts* appears to duplicate to a large extent the information found in *categories* except that it appears to breaks *concepts* down into more detail. What the Sankey diagrams do show is that the *concepts* are not evenly distributed. Some are related to a single article. For example, the concept marriage relates only to the article “To Our Sisters,” while the concepts Islam, Prophets of Islam, Muhammad, and Caliphate relate to a range of article types. The pattern is different for the recontextualisations. *Concepts* are more thinly spread over texts with fewer and smaller clusters of *concepts* associated with many texts. *Concepts* which relate to a large range of article types are more restricted to

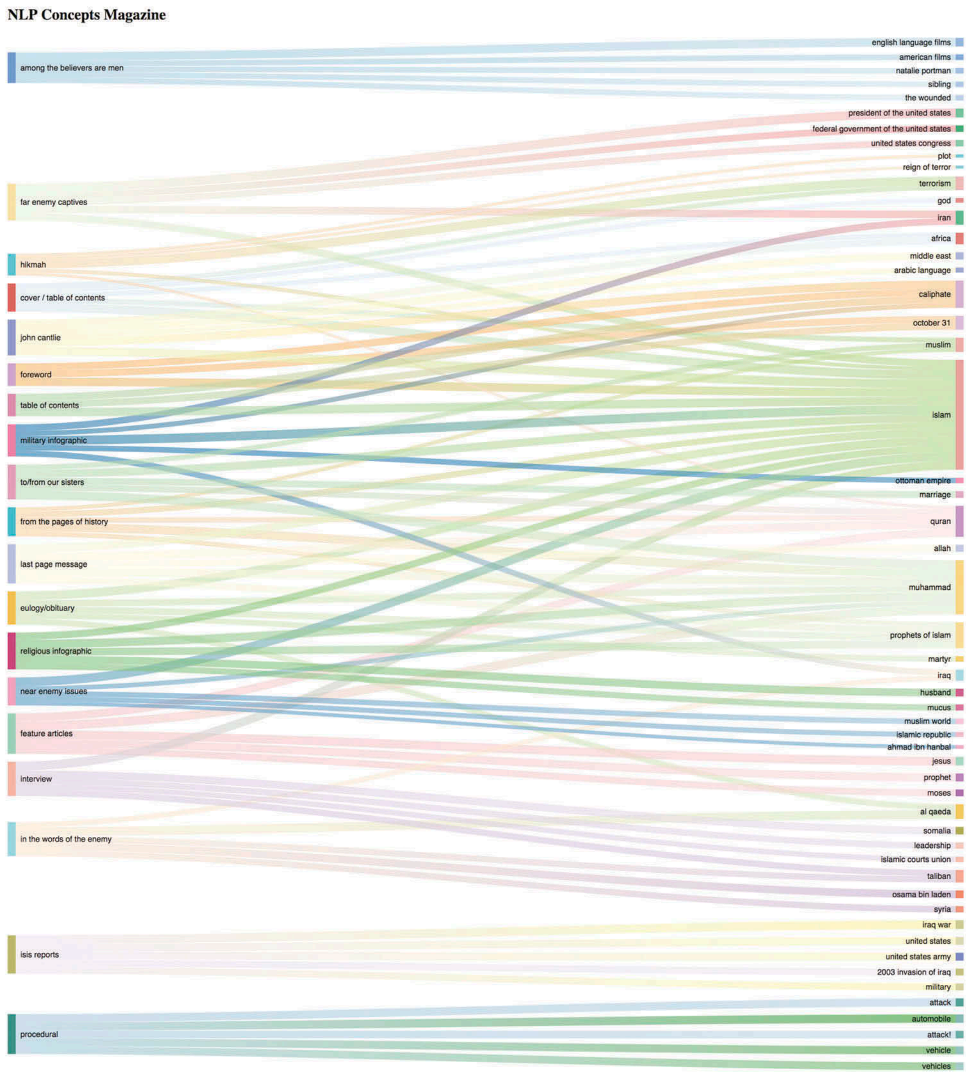

Figure 5. Screenshot of Sankey diagram showing results of analysis for concepts for *Dabiq* and *Rumiyah* texts (https://curtinic.github.io/mma-violent-extremism/sankey_nlp_concepts_mag.html).

the *Dabiq* and *Rumiyah* texts and therefore could prove useful in helping to distinguish between texts produced by ISIS and texts produced in response to those texts.

The word clouds also show some striking differences. For example, a most noticeable feature found in the word clouds is that the word cloud for the texts from *Dabiq* and *Rumiyah* foreground religion, the United States, ISIS's caliphate, and other Islamic extremist organisations (Taliban), while the concepts from the recontextualisations foreground Islam, terrorism, and attack. One aspect worth further investigation is that many *concepts* appear to be closely associated with particular article types. When mapped against *categories*, the *concepts* most closely related to the most prominent categories in Figure 2 can be used to help differentiate between texts that promote violent extremism from other texts.

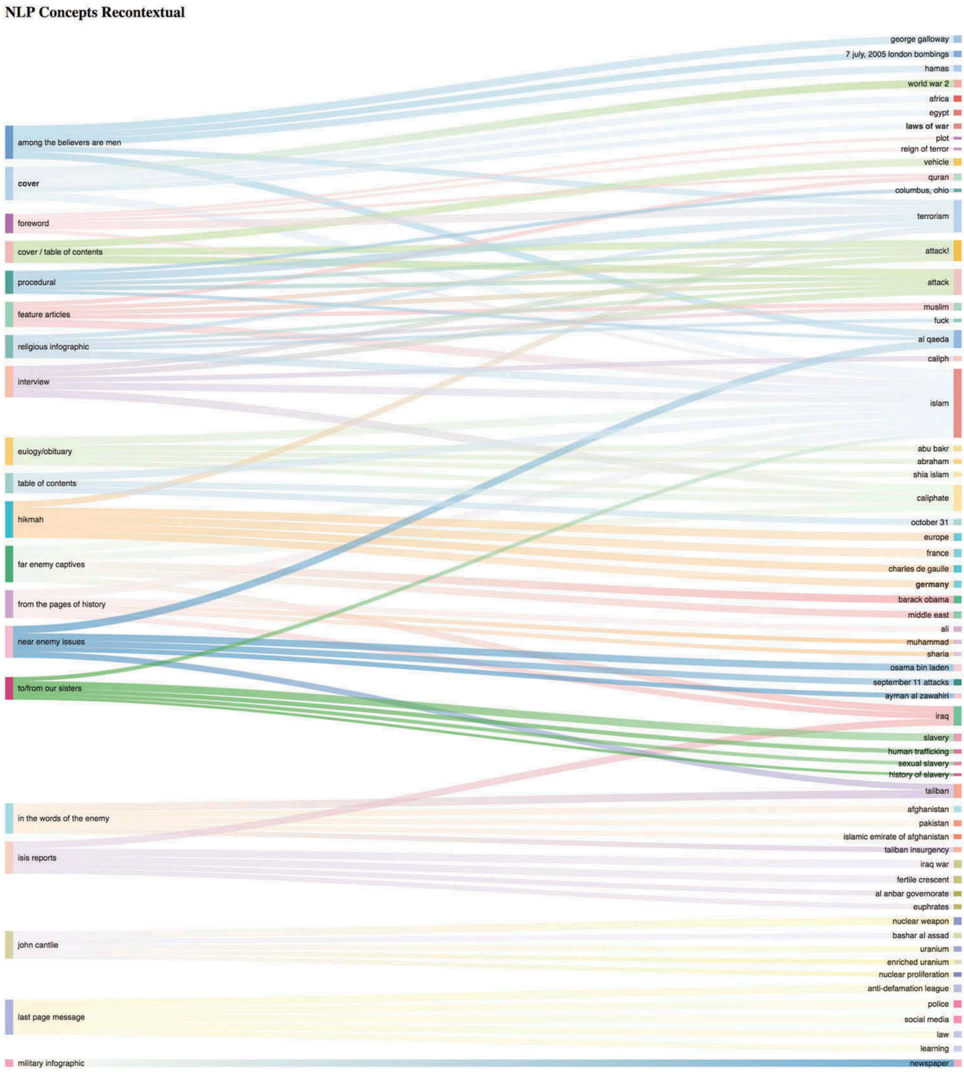

Figure 6. Screenshot of Sankey diagram showing results of analysis for concepts for texts containing recontextualised images from those texts (https://curtinic.github.io/mma-violent-extremism/sankey_nlp_concepts_rec.html).

Entities

Figure 8(a) and 8(b) shows word clouds displaying the results of the automated analysis according to *entities*. Figure 8(a) shows *entities* from *Dabiq* and *Rumiyah* and Figure 8(b) shows *entities* from the recontextualisations.

Entities are closely related to the content of a text and there is quite a bit of crossover depending on how closely related the article containing the recontextualised image is to the original article. For instance, where the recontextualisation directly references the original or is a comment on the content of the original such as with *Far Enemy Captives*, the *Obituary/Eulogy*, and *ISIS Reports*. The most visible difference between the word

(a) Word cloud of relative proportion of *concepts* in *Dabiq* and *Rumiyah*

(b) Word cloud showing relative proportion of *concepts* in the recontextualisations

Figure 7. Word clouds showing relative proportions of concepts in *Dabiq* and *Rumiyah* (7a) and in the recontextualisations (7b).

(a) Word cloud showing the results of *entities* for *Dabiq* and *Rumiyah* texts

(b) Word cloud showing the results of *entities* for texts containing recontextualised images from those texts

Figure 8. Word clouds displaying the results of analysis for entities: 8(a) Word cloud showing the results for *Dabiq* and *Rumiyah* texts and 8(b) word cloud showing the results for texts containing recontextualised images from those texts.

clouds is the foregrounding of ISIS in the recontextualisations. ISIS very rarely (if ever) refer to themselves as ISIS.

Other differences worth exploring include the use of Arabic words. In the recontextualisations, the Arabic words used are almost always either the names of people or of places. On the other hand, while there are the names of people and places in Arabic in the original texts, there are also far more other Arabic words. While *Dabiq* and *Rumiyah* are published in English, they are liberally and strategically sprinkled with Arabic words. These words are integral to the meaning of the text and often realise negative and positive sentiment. These words often refer to ISIS's enemies, for example *murtaddin* (apostates), *Rafidi* (Shi'a Muslims), and *mushrikin* (polytheists/idol worshippers). The English word "crusaders" is also used to refer to ISIS's Western enemies and their allies. Other Arabic words such as *bay'ah* (pledge of allegiance), *mujahidin* (fighter engaged in *jihad*), and

Rasulullulah (Messenger of Allah) are words which support and reinforce ISIS's world view and values.

Entities also frequently misclassifies, especially when it is classifying Arabic words. For example, in the analysis of "Last Page Message," it classifies the Arabic word *Dajjal* as a company when the word refers to the Islamic equivalent of the Christian Antichrist; in the text of "Among the Believers are Men" (Example 6 in Table 1), *Abu Qudamah* is classified both as a person and as a geographical feature (*Abu Qudamah* is a person's name). Classification is also sometimes close but not quite right. For example, *crusaders* in Figure 2(c) is classified as an *organization* when it is a term used by ISIS to refer to the "far enemy."[28] Other examples are the classification of towns and provinces as geographical features. Training in specific sets of Arabic words could perhaps remedy this issue. There is also quite a lot of overlap between *entities* and *keywords*, with many *entities* appearing as *keywords*, but not necessarily the other way around.

There is also some mix between *concept* and *entities*. For example, organisations such as Al Qaeda and Taliban are classified as *concepts* while an organisation of the same type, Islamic State, is classified as an *entity*.

Keywords

The word clouds in Figure 9 visualise the results of the analysis of *keywords*. The results of the automated analysis of the *Dabiq* and *Rumiyah* texts and texts containing recontextualised images according to *keywords* are shown in Figure 9(a) and 9(b). Lists of *keywords* are often extensive so only the top five in each example are used. For purposes of comparison, two word clouds derived from the aggregated text of each set of examples are shown in Figure 9(c) and 9(d).

The word clouds for just the top five keywords in each article in both sets of data are similar to the word clouds for the complete text of the articles in both sets of data, indicating that these keywords are a useful pointer. Keywords in both sets of articles do show quite a bit of crossover, largely related to the topic of the article. Where the content of the text with the recontextualised image is related to the original article, there is more crossover in keywords than when the connection is less direct. There are, however, some interesting differences. These differences can be observed in the word clouds. The word cloud for the *Dabiq* and *Rumiyah* texts, for example, highlights words related to Islam and words related to ISIS's "enemies." It also highlights a number of Arabic words. In contrast, the word cloud for the recontextualisations foregrounds the acronym ISIS and tends to focus on words with a political rather than a religious flavour. These patterns are also evident in the two word clouds derived from the analysis of *keywords* in the aggregated text of the original articles and the recontextualisations.

These differences can also be seen in the *keywords* analysis of the individual texts. For instance, ISIS do not refer to themselves as ISIS. They mostly refer to themselves as Islamic State or the Caliphate (*Khilafah* if Arabic is used). They are occasionally referred to as Islamic State in the recontextualisations but more frequently as ISIS, usually to show that the following word is associated with ISIS (e.g., ISIS propaganda, ISIS soldiers, ISIS fighters). ISIS never use this type of structure to refer to themselves. There are some other differences in the patterns of keywords. ISIS frequently refer to Allah and the Prophet Muhammad. These words do not appear in the recontextualisations. The words *propaganda, terrorist*, and *terror attacks* are found frequently in the recontextualistaions and either not at all or very infrequently in the original articles. Arabic

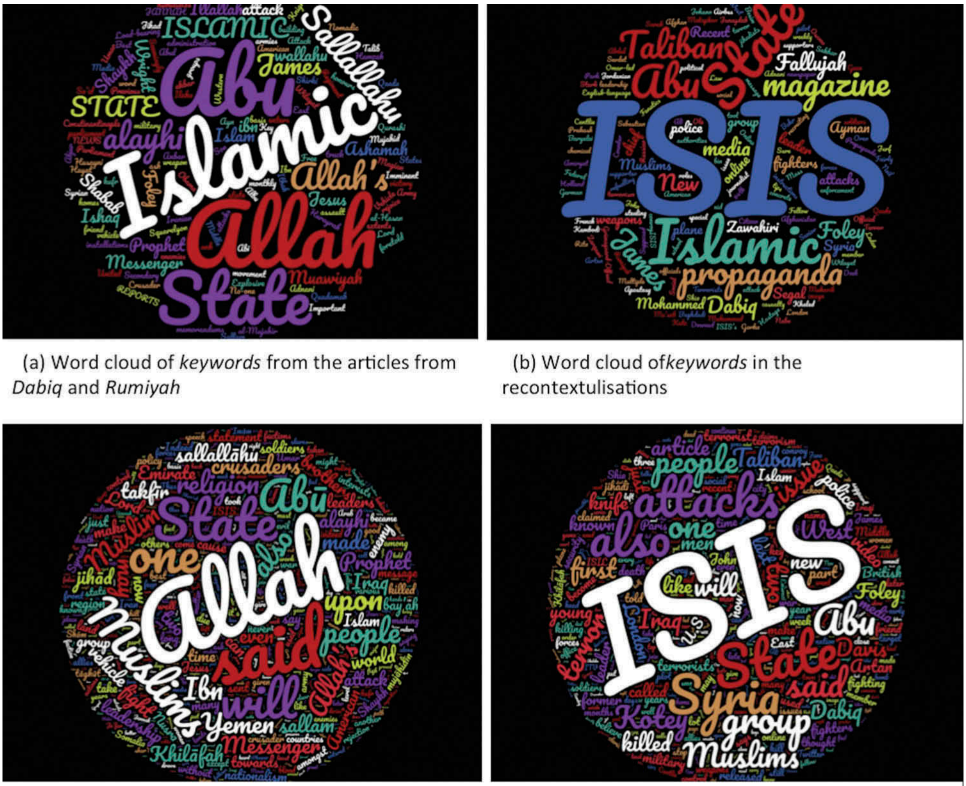

(a) Word cloud of *keywords* from the articles from *Dabiq* and *Rumiyah*

(b) Word cloud of*keywords* in the recontextulisations

(c) Word cloud of *keywords* for the aggregated text of the articles from *Dabiq* and *Rumiyah*

(d) Word cloud of *keywords* for the aggregated text of the recontextualisations

Figure 9. Word clouds for the keywords from the articles from *Dabiq* and *Rumiyah* and the keywords in the recontextulisations.

words are frequently used in the original articles. These are either words with religious and/or military associations, words relating to Islamic scripture, words of praise for ISIS and words denigrating their "enemies," names of people, and place names. Arabic words for individual people and place names are found frequently in the recontextualisations, but the other Arabic words are not found at all. This is a point developed further in the final summary.

Emotion

Figure 10 shows the results of the automated analysis of the *Dabiq* and *Rumiyah* texts and texts containing recontextualised images from those texts analysed according to *emotion*. *Emotions* are coded between 0 and 1. For example, a value of .01 for joy and a value of .80 for anger would indicate a very angry and almost joyless text.

This analysis provided some unexpected insights. For example, in 14 of the 20 recontextualisations, the value for fear was higher than in the corresponding article in *Dabiq* or *Rumiyah*. In three cases the values were equal and in three cases the original articles had higher values for *fear* than the recontextualisations. However, it is difficult to draw inferences from the *emotion* analysis without knowing the direction and source of the

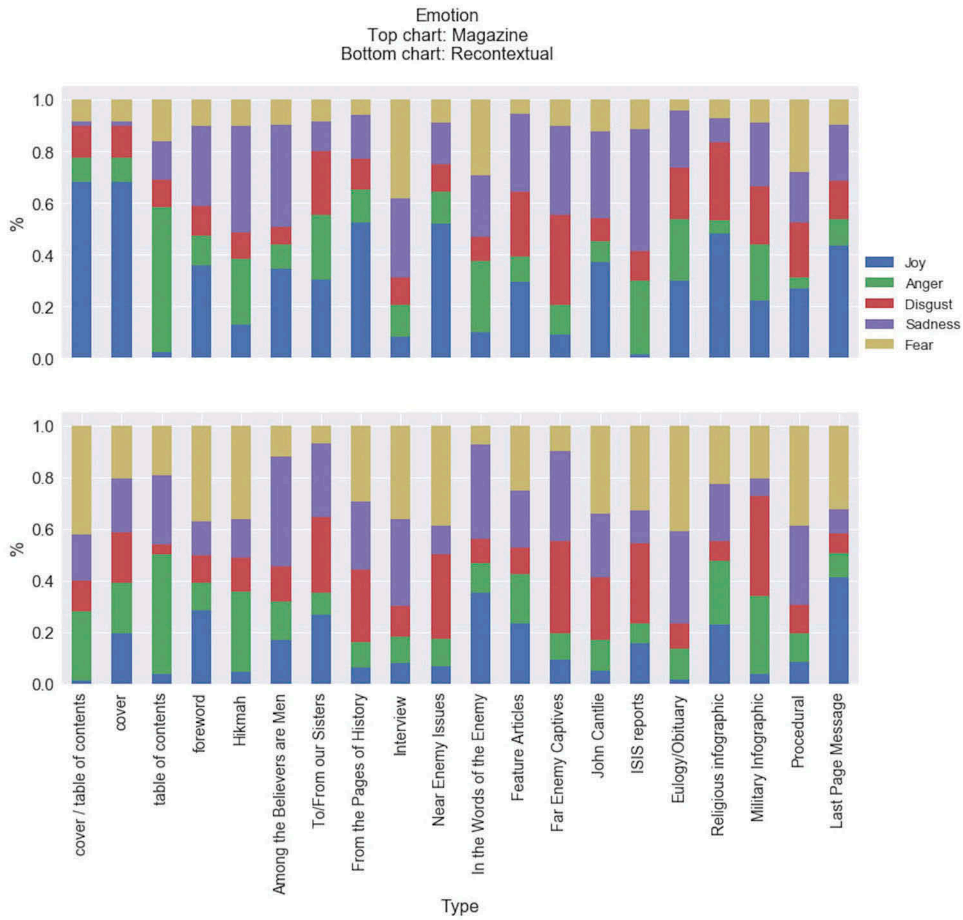

Figure 10. Visualisation of emotion analysis of *Dabiq* and *Rumiyah* texts and texts containing recontextualised images from those texts.

emotion. Information on things like who is angry at whom and who is scared of whom would greatly aid in drawing inferences from the results.

Sentiment

Figure 11 shows the results of the automated analysis of the *Dabiq* and *Rumiyah* texts and texts containing recontextualised images from those texts analysed according to *sentiment*. *Sentiment* is coded as a positive or negative value with a maximum range of between 1 and −1.

Of the 40 texts analysed, all but four showed negative *sentiment*. One, the cover of Issue 7 of *Dabiq*, had a 0 (neutral) rating. The articles which showed positive *sentiment* were all from *Rumiyah*. The two most positive ratings were for a procedural article on how to plan and conduct a vehicular terrorist attack in a Western country and then for a religious infographic which describes *Jannah*, the Islamic "paradise" reserved for the righteous in the afterlife. The other two positives were for an obituary for a "martyred" ISIS leader and for the cover and table of contents for the previous issue of *Rumiyah*, which contained a pointer to an article

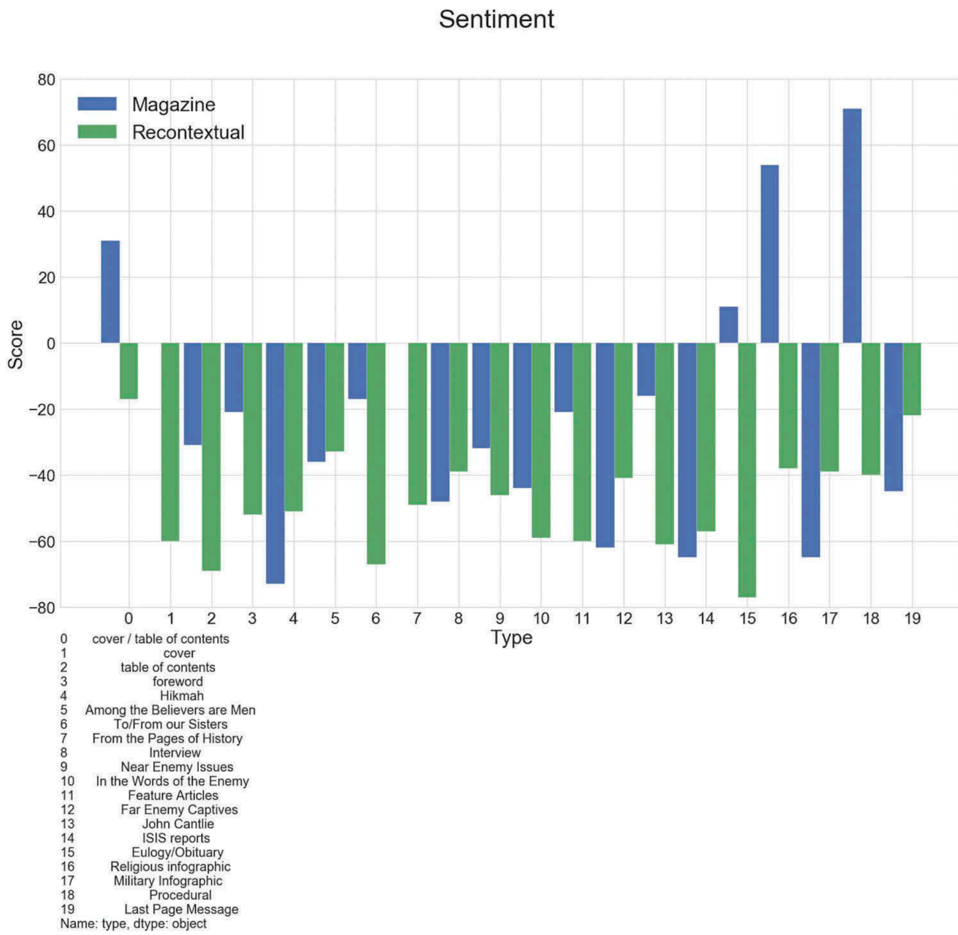

Figure 11. Sentiment results for *Dabiq* and *Rumiyah* texts and texts containing recontextualised images from those texts.

about the life and deeds of the same "martyr." All other articles had negative ratings for *sentiment*, with 13 of the recontextualisations having a greater negative value than their corresponding articles in *Dabiq* and *Rumiyah*. All of the recontextualisations showed negative sentiment towards ISIS, while the *Dabiq* and *Rumiyah* articles were all positive towards ISIS (especially ISIS heroes and Islam) and were all negative towards all of ISIS's many "enemies." In the aggregated/average *sentiment* analysis, these appeared to somewhat balance each other out, which possibly led to the generally smaller negative ratings for the *Dabiq* and *Rumiyah* articles. The *sentiment* analysis in the demonstration models as it stands, does not appear to provide many insights on its own. Like *emotion*, the analysis of *sentiment* would be strengthened if it could include the direction, source, and target of the emotion and sentiment.

Summary of observations and conclusions

When interpreted through a systemic functional lens, the demo focuses on the experiential component of ideational meaning. *Categories* and *concepts* relate to the register variable field, with *categories* the most general descriptor. *Categories* identifies generally what the text is about without necessarily identifying specific wording in the text. *Concepts* is also field-related but more specific. Similarly, *entities* and *keywords* are also experientially oriented but are more delicate and identify people and things in text. These appear to be based on identification, occurrence, and recurrence of individual words.

There is some overlap between *categories* and *concept*, between *concept* and *entities*, and between *entities* and *keywords*. Taking *keywords* and *entities* as the most delicate level of analysis, there is increasing abstraction as the analysis moves to *concepts* to *categories*, which appear to be derived from frequency and aggregation of semantic relatedness of *keywords* and *entities* in a text. All are closely related to field.

Interpersonal meaning is addressed through *sentiment* and *emotion*. When these are applied to a whole text, they do not reveal much about individual texts, but when these analyses are applied to groups of texts and the results compared, they do appear to offer some promise. *Sentiment* and *emotion* could possibly both also be used to infer the text's ideological position, although the direction and target of the sentiment and emotion would need to be known in order to do this. In their aggregated output these two variables do not, for example, reveal the direction of the *sentiment* or *emotion*, only overall "averages." Targeted *sentiment* and *emotion* could prove quite useful if they could be automated rather than just applied to individual segments of text.

Textual meaning is not addressed. For example, the demo software does not distinguish between a headline and words in the body of a text. It also does not identify points of information prominence at different scales in a text. As the software appears to rank all information as being of equal importance, this is a shortcoming that would need to be addressed in future developments, especially since lead images frequently have a headline superimposed over them, giving additional textual prominence to both the image and the headline.

Where the IBM Watson demonstration NLU models are strongest is in ideational analysis at the highest (e.g., *categories*) and lowest (e.g., *keywords*) ranks in the hierarchy. The automated analysis is heavily ideationally focussed. In systemic functional terms, it is strongly associated with the register variable field at all levels in its hierarchy. It identifies *categories* well. This information, if combined with metadata associated with URLs, could prove useful in helping to distinguish between material which supports terrorism, material which is either against or comments on violent extremism, and material which is unrelated. A grammatical approach that is integrated with information about context, such as that obtained from *categories* and also integrated with information that can be obtained from below, such as from *keywords*, would add a middle level to the analysis. Such an approach could also enhance the analysis of interpersonal features such as *emotion* and *sentiment*. This integration is essential for interpreting nuances of meaning, as "each utterance is dependent on the context, while also changing that context for subsequent utterances." A grammatical analysis using a selection of key features of a systemic functional approach would satisfy those conditions. A simplified set of selections would be necessary, at least in the short to medium term since, as Bateman et al. point out, "hundreds of thousands of words (or a smaller, highly specific sample) may be needed in order to develop frequency profiles for lexical alternatives."[29] That is, the more delicate the analysis, the larger the sample needed. Such approaches are in the process of being developed but are not yet fully operational.[30]

When interpreted through an SF lens, some useful and potentially useful insights can be drawn from the IBM Watson demonstration NLU models. The software as it stands does identify some tendencies and patterns in both sets of text, which can at least help to distinguish between terrorist/ISIS/Islamic extremist related and non-terrorist/ISIS/Islamic extremist related texts. That is, there appears to be enough commonality between material produced by ISIS and material produced in response to ISIS material to distinguish between Islamic extremist-related discourse (for and against) and discourse which is unrelated.

The automated analysis also produces some insights which, with some refinement, could aid in distinguishing between material which supports and promotes violent extremism and material which either opposes or comments or reports on that material. A first step towards such refinement could be to map the results of the IBM Watson analyses on to the stratal framework (see section "Theoretical Framework"). This would highlight how the analyses performed with IBM Watson fit within a systemic functional framework and also identify gaps. An example of such a mapping is shown in Table 2. As the table shows, ideational meaning is better accounted for in the IBM Watson analyses than is interpersonal meaning. Textual meaning is not covered at all by the IBM Watson analyses.

The IBM Watson analyses on their own, if applied without an intervening interpretive framework, are therefore of limited value. Even in conjunction with an interpretive framework, there is not one single feature or even one or two features which could be used by themselves to make distinctions between material which supports and promotes violent extremism and material which either opposes or comments or reports on that material. Any answer would appear to lie in combinations of features ranging from higher level features of the discourse to the lexical level. In systemic functional terms, these are features which relate to register, discourse semantics, and grammar. It is for these reasons that an interpretive approach which uses an integrated theory of language and context is necessary. By using such an approach, patterns found in analyses at different strata can be compared and calibrated.

One such possibility which would enhance the analysis from words through to register is to train any software used on propaganda texts, especially so that it recognises a relatively small set of Arabic words frequently, and in this context, exclusively used by ISIS to denigrate "enemies" and to praise themselves, their affiliates, and supporters. The frequent use of Arabic words relating to Islamic scripture is also a feature of the *Dabiq* and *Rumiyah* texts not found in the recontextualisations. This is one feature which distinguishes ISIS produced texts from other terrorism/extremism-related texts. Such a set could be easily compiled from the text of *Dabiq* and *Rumiyah* and used for training purposes. This is also important because most of the errors made in the automated analysis stem from misclassification of Arabic words.

Such a list would not only operate at the word level. Results could also feed up in the analytical hierarchy to aid in making analyses like *categories* and *concepts* more

Table 2. Mapping the results of the IBM Watson analyses on to a stratal framework.

Metafunction Strata	Ideational	Interpersonal	Textual
Register	Categories Concepts		
Discourse Semantics	Keywords organised taxonomically (Lexical relations)		
Lexicogrammar	Entities Keywords Semantic Roles	Emotion Sentiment	

sophisticated. For example, a *category* that is quite common across both sets of texts is *religion and spirituality/Islam*. If this could be extended to one or two "nodes" so that it was something like *religion and spirituality/Islam/Islamic extremism/jihad* based on choices from a religious and jihadist set of Arabic terms, this would provide a strong basis for identifying texts which were about Islamic extremism and terrorism. It would also exclude texts from Islamic sites which had nothing to do with extremism or terrorism. That is, the combination of Islam and jihad/terrorism would be a strong indicator of the field and ideology of a text. If information from the same lexical set were added to show *sentiment*, this would be a strong indicator of the ideological position of the source text. In addition, if analyses of *sentiment* and *emotion* could be extended to include noun and verb combinations in the grammar and/or adjective and noun combinations (see[31] for a systemic functional description of these types of relationships), this information would specify the source and target of the sentiment and the force of the sentiment or emotion. The introduction of some features of interpersonal systems which concern the expression of attitudes in text would also assist in interpretation.

At the level of lexis, analysis of *keywords* could also be useful. The analyses of *concept* and *entities*, while potentially useful, are perhaps less insightful than *categories* and *keywords*. Both of these seem to be derived from some kind of lexical analysis and show considerable overlap with other analyses. For example, *concept* overlaps with *categories, categories* overlaps with *entities*, and *entities* overlaps with *keywords*.

What is missing from the ideationally oriented analyses is attention to grammatical and discourse semantic features. For example, the discourse semantic feature lexical relations could be used to build a small number of key taxonomically related lexical sets (in English and Arabic). Such sets would be based on relationships among words rather than frequency. This would shift the lexical analysis up to the level of discourse semantics where relationships among lexical items are accounted for. In addition, while the automated analyses identify nouns (and sometimes nominal groups) well, they do not identify verbs very often and only occasionally identify noun and verb combinations infrequently. Information about things like who is doing what to whom and how someone feels about someone or something else would be likely to provide useful information for helping to distinguish pro-violent extremist texts from other texts.

Although the automated analysis is heavily experientially focused, some attention is paid to interpersonal meaning through analyses of *emotion* and *sentiment*. These analyses seem to be based on lexis. The analysis of *emotion* did provide some insights. For instance, the distribution and weighting of the emotions of fear and joy was interesting. The aggregated *sentiment* analysis is less useful as it stands. It seems to be based on values averaged out over a whole text. This leads to similarities between the *Dabiq* and *Rumiyah* texts and the recontextualisations which could be misleading since the analysis does not take into account the direction of the *sentiment*. If analysis of the direction of the *sentiment* could be automated, or at least informed by a more sophisticated grammatical analysis, this would provide more useful information. Similarly, the direction of *emotion* is as important as its intensity.

The demo NLU models do not really consider the textual metafunction at all. Textual prominence, in terms of both position and size of type and in the composition of features such as headline and image, play an important communicative role. For example, a religious/jihad inspired headline superimposed over an iconic image plays a large role in foregrounding what is being presented as important information in a text. The applicability of any automated analysis would be greatly enhanced if it could account for at least some textual features.

There is not enough data in this study to draw firm conclusions, but there are some things that show promise for further investigation. The Sankey diagram for *categories* appears to indicate an association between the *categories* religion and spirituality/Islam and society/unrest and war and certain article types. This association is found in both the *Dabiq* and *Rumiyah* texts and in the recontextualisations. The original texts associated with these *categories* tend to be those with both a religious and a jihadist flavour such as Last Page Messages and Forewords. These texts also almost always contain iconic ISIS images which are textually prominent and a strong interpersonal loading. Examination of patterns of distribution of ISIS images identified in Tan et al.[32] showed that the images which had the widest and most consistent pattern of redistribution over time were iconic ISIS images. In this trial, images were only used to identify text. The meanings arising from images and text combined requires further development. One possibility is that the combination of iconic images in association with the *categories* of religion and spirituality/Islam and society/unrest and war, keywords drawn from lexical sets associated with Islamic scripture and *jihad* and URL metadata might be a potential starting point for a larger scale analysis.

Lastly, the IBM Watson NLU tool is only one of many automated text analysis tools that we evaluated in this research mainly through its ease of use. There are many free and open source NLP/NLU tools including spaCy[33] and the Natural Language ToolKit (NLTK)[34] that may generate better or poorer results in certain NLP tasks when compared to the IBM Watson NLU tool in the context of analysing violent extremist text.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

The research is undertaken in the project “A Context-Based Mixed Methods Approach to Extremist Communications” with Lead Investigators Professor Kay O’Halloran (School of Education and Curtin Institute for Computation) and Dr Lucy Resnyansky (Defence, Science and Technology (DST) Group, Department of Defence, Australia), in collaboration with Dr Peter Wignell and Dr Sabine Tan (School of Education, Curtin University) and Dr Kevin Chai and Dr Rebecca Lange (Curtin Institute of Computation, Curtin University). The collaborative project is funded by the Defence, Science and Technology (DST) Group, South Australia.

Notes on contributors

Dr Peter Wignell is a Senior Research Fellow and a member of the Multimodal Analysis Group in the School of Education at Curtin University, Western Australia. He has an academic background in Systemic Functional Linguistics, discourse analysis and multimodal analysis. His research has been both theoretical and applied. For example, his work on the role of language in the construction of specialised knowledge systems has informed literacy theory and pedagogy. Recent research has led to publication in a diverse range of applications such as brand semiotics, performance studies and translation. Peter's current research focus is on the application of a Systemic Functional Multimodal analysis approach to violent extremist discourse and political discourse.

Dr Kevin Chai is currently working as the lead data scientist for Curtin's Institute for Computation (CIC) to provide data, computing, analytics and machine learning / artificial intelligence expertise to collaborate in projects with researchers and industry partners. Kevin completed his PhD in computer science at Curtin University, worked as Research Fellow at the University of New South Wales developing machine learning and natural language processing models to analyse large medical datasets and as a Data Scientist at Curtin University on developing predictive models. Kevin has research interests and experience in machine learning / AI and its application to solving important problems in domains such as health, medicine, education, science, engineering, the humanities and many other diverse areas of research.

Dr Sabine Tan is a Senior Research Fellow and a member of the Multimodal Analysis Group in the School of Education at Curtin University, Western Australia. She has a background in critical multimodal discourse analysis, social semiotics, and visual communication. She has applied multi-disciplinary perspectives for the analysis of institutional discourses involving traditional and new media. She has worked on interdisciplinary projects involving the development of interactive software for the multimodal analysis of images, videos and 360 videos for research and educational purposes. Sabine's current work involves developing multimodal approaches to big data analytics in the fields of online terrorism, political discourse, and 360 video.

Professor Kay O'Halloran is Professor and leader of the Multimodal Analysis Group in the School of Education at Curtin University, Western Australia. She has a background in mathematics, multimodal discourse analysis, and linguistics. She has extensive experience in establishing and leading interdisciplinary research teams to develop and make widely available new digital tools and techniques for analysing images, videos, and 360 videos. Kay is currently developing mixed methods approaches that combine multimodal analysis, data mining and visualisation for big data analytics in areas which include online extremism, political rhetoric and 360 video.

Dr Rebecca Lange is a Computational Specialist with the Curtin Institute for Computation. Rebecca completed her PhD in astronomy at the International Centre for Radio Astronomy Research at the University of Western Australia. She has worked on the development, testing and in situ implementation of instruments and software for the Imaging & Sensing for Archaeology, Art History & Conservation group at Nottingham Trent University in the UK. Throughout her studies and research work she has gained extensive programming as well as data analytics and visualisation experience in various programming languages. Rebecca is particularly interested in applying her technical skills to support and train researchers in the digital humanities.

Notes

1. Attila Kovács, "The New Jihadists and the Visual Turn from Al-Qaida to ISIL/ISIS/Daish, *BiztPol Affairs2*, no. 3 (2015): 67.
2. Stevan Rudinac, Iva Gornishka, and Marcel Worring, "Multimodal Classification of Violent Online Political Extremism Content with Graph Convolutional Networks," Thematic Workshops 17, October 23–27, 2017. Mountain View, California, USA (2017).
3. James P. Farwell, "The Media Strategy of ISIS," Survival: *Global Politics and Strategy* 56, no. 6 (2014): 49–55; Kovacs, "The New Jihadists" (see note 1); Peter Wignell, Sabine Tan, and Kay L. O'Halloran, "Violent Extremism and Iconisation: Commanding Good and Forbidding Evil?" *Critical Discourse Studies* 14, no. 1 (2017): 1–22.; Ahmed Al-Rawi, "Video Games, Terrorism, and ISIS's Jihad 3.0, " *Terrorism and Political Violence* 30, no. 4 (2018): 740–60, DOI:10.1080/09546553.2016.1207633; Carol K. Winkler, Kareem El Damanhoury, Aaron Dicker, and Anthony F. Lemieux, "The Medium Is Terrorism: Transformation of the About to Die Trope in Dabiq," *Terrorism and Political Violence* (2016): DOI:10.1080/09546553.2016.1211526; Kareem El Damanhoury, Carol Winkler, Wojciech Kaczkowski, and

Aaron Dicker, "Examining the Military–Media Nexus in ISIS's Provincial Photography Campaign," *Dynamics of Asymmetric Conflict* (2018): 1–20: DOI is: DOI:10.1080/17467586.2018.1432869
4. Peter Wignell, Kay L. O'Halloran, Sabine Tan, Rebecca Lange, and Kevin Chai, "Image and Text Relations in ISIS Materials and the New Relations Established through Recontextualisation in Online Media," *Discourse and Communication* (2018): vol. 12, 5: pp. 535-559; Sabine Tan, Kay L. O'Halloran, Peter Wignell, Kevin Chai, and Rebecca Lange, "A Multimodal Mixed Methods Approach for Examining Recontextualisation Patterns of Violent Extremist Images in Online Media, " *Discourse, Context and Media* 21 (2018): 18–35; Kay L. O'Halloran, Sabine Tan, Peter Wignell, and Rebecca Lange, "Multimodal Recontextualisations of Images in Violent Extremist Discourse," in *Advancing Multimodal and Critical Discourse Studies: Interdisciplinary Research Inspired by Theo van Leeuwen's Social Semiotics*, eds. Sumin Zhao, Emilia Djonov, Anders Björkvall, and Morten Boeriis (New York and London: Routledge, 2018), 181–202.
5. Kay L. O'Halloran, Sabine Tan, Peter Wignell, John Bateman, Duc-Son Pham, Michele Grossman, and Andrew Vande Moere, "Interpreting Text and Image Relations in Violent Extremist Discourse: A Mixed Methods Approach for Big Data Analytics," *Terrorism and Political Violence* (2016). DOI:10.1080/09546553.2016.1233871; Jan Christoffer Andersen and Sveinung Sandberg, "Islamic State Propaganda: Between Social Movement Framing and Subcultural Provocation," *Terrorism and Political Violence* (2018). DOI:10.1080/09546553.2018.1484356; Julian Droogan and Shane Peattie, "Reading Jihad: Mapping the Shifting Themes of Inspire Magazine," *Terrorism and Political Violence* 30, no. 4 (2018): 684–717, DOI:10.1080/09546553.2016.1211527; Logan Macnair and Richard Frank, "Changes and Stabilities in the Language of Islamic State Magazines: A Sentiment Analysis," *Dynamics of Asymmetric Conflict* 1, no. 2 (2018): 109–20. DOI:10.1080/17467586.2018.1470660; Douglas C. Derrick, Karyn Sporer, Sam Church, and Gina Scott Ligon, "Ideological Rationality and Violence: An Exploratory Study of ISIL's Cyber Profile," *Dynamics of Asymmetric Conflict* 9, no. 1–3 (2016): 57–81. DOI:10.1080/17467586.2016.1267866.
6. Jennifer Boutz, Hannah Benninger, and Alia Lancaster, "Exploiting the Prophet's Authority: How Islamic State Propaganda Uses Hadith Quotation to Assert Legitimacy," *Studies in Conflict & Terrorism* (2018). DOI:10.1080/1057610X.2018.1431363; Haroro J. Ingram, "An Analysis of Inspire and Dabiq: Lessons from AQAP and Islamic State's Propaganda War," *Studies in Conflict & Terrorism* 40, no. 5 (2017): 357–75. DOI:10.1080/1057610X.2016.1212551; Peter Wignell, Sabine Tan, and Kay L. O'Halloran, "Under the Shade of AK47s: A Multimodal Approach to Violent Extremist Recruitment Strategies for Foreign Fighters", *Critical Studies on Terrorism* 10, no. 3 (2017): 429–52.
7. Peter Wignell, Sabine Tan, Kay L. O'Halloran, and Rebecca Lange, "A Mixed Methods Empirical Examination of Changes in Emphasis and Style in the Extremist Magazines Dabiq and Rumiyah," *Perspectives on Terrorism*, Vol 11, no 2(2017): 2–20; Wignell, Tan, and O'Halloran, "Violent Extremism and Iconisation" (see note 3); Wignell, Tan, O'Halloran, and Lange, "A Mixed Methods Empirical Examination" (see note 23).
8. Droogan and Peattie, "Reading Jihad" (see note 5); Macnair and Frank, "Changes and Stabilities in the Language of Islamic State Magazines" (see note 5).
9. David B. Skillicorn, "Empirical Assessment of Al Qaeda, ISIS, and Taliban Propaganda" (Intelligence and Security Informatics (ISI), 2015 IEEE International Conference, 2015), 61–66: DOI:10.1109/*ISI*.2015.7165940; Matteo Vergani and Ana-Maria Bliuc, "The Evolution of the ISIS' Language: A Quantitative Analysis of the Language of the First Year of Dabiq Magazine," *Sicurezza, Terrorismo e Società* 1, no. 2 (2015): 7–20; Carey Jewitt, Jeff Bezemer, and Kay L. O'Halloran, eds., Introducing Multimodality (London: Routledge, 2015); Carey Jewitt, ed., *The Routledge Handbook of Multimodal Analysis*, 2nd ed. (London: Routledge, 2014).
10. Kay L. O'Halloran, ed., *Multimodal Discourse Analysis* (London: Continuum, 2004).

11. Gunther Kress and Theo Van Leeuwen, *Multimodal Discourse: The Modes and Media of Contemporary Communication Discourse* (London: Arnold, 2001).
12. Theo Van Leeuwen, *Introducing Social Semiotics* (London: Routledge, 2005).
13. Kay L. O'Halloran, "Systemic Functional-Multimodal Discourse Analysis (SF-MDA): Constructing Ideational Meaning Using Language and Visual Imagery," *Visual Communication* 7, no. 4 (2008): 443-75; Michael A. K. Halliday, *Explorations in the Functions of Language* (London: Edward Arnold, 1978).
14. Michael A. K. Halliday, *An Introduction to Functional Grammar*, 2nd ed. (London: Arnold, 1994); Michael A. K. Halliday and C. M. I. M. Matthiessen, *Halliday's Introduction to Functional Grammar*, 4th ed. (Oxford: Routledge, 2014);
15. James R. Martin: Evolving systemic functional linguistics: beyond the clause. *Functional Linguistics* 1:3, (2014) doi.org/10.1186/2196-419X-1-3
16. James R. Martin, *English Text: System and Structure* (Amsterdam: John Benjamins, 1992).
17. Ibid; James R. Martin and David Rose, *Working with Discourse: Meaning beyond the Clause* (London, New York: Continuum, 2003); Ibid.
18. Kay L. O'Halloran, Sabine Tan, and Peter Wignell, "SFL and Multimodal Discourse Analysis," in *The Cambridge Handbook of Systemic Functional Linguistics*,eds.Geoff Thompson, Wendy L. Bowcher, Lise Fontaine, and David Schönthal (Cambridge: Cambridge University Press, in press); Michael O'Toole, *The Language of Displayed Art*, 2nd ed. (London: Routledge, 2011).
19. John Bateman, Daniel McDonald, Tuomo Hiippala, Daniel Couto-Vale, and Eugeniu Costetchi, "Systemic Functional Linguistics and Computation: New Directions, New Challenges," in *The Cambridge Handbook of Systemic Functional Linguistics*, eds. Geoff Thompson, Wendy L. Bowcher, Lise Fontaine, and David Schönthal (Cambridge: Cambridge University Press, in press).
20. Wignell, Tan, O'Halloran, and Lange, "A Mixed Methods Empirical Examination" (see note 7).
21. Walaa Medhat, Ahmed Hassan, and Hoda Korashy, "Sentiment Analysis Algorithms and Applications: A Survey, " *Ain Shams Engineering Journal5*, no. 4 (2014): 1093–113.
22. Vidisha M. Pradhan, Jay Vala, and Prem Balani, "A Survey on Sentiment Analysis Algorithms for Opinion Mining," *International Journal of Computer Applications* 133, no. 9 (2016): 7–11.
23. Edward Chao-Chun Kao, Chun-Chieh Liu, Ting-Hao Yang, Chang-Tai Hsieh, and Von-Wun Soo, "Towards Text-Based Emotion Detection: A Survey and Possible Improvements" (Information Management and Engineering, 2009. ICIME '09. *International Conference on, IEEE, 2009*), 70–74: DOI: 10.1109/ICIME.2009.113; Kazi Saidul Hasan and Vincent Ng, "Automatic Keyphrase Extraction: A Survey of the State of the Art" (*Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers)), vol. 1, 2014), 1262–73.
24. Ralph Grishman and Beth Sundheim, "Message Understanding Conference-6: A brief History" (Copenhagen: *COLING 1996 Volume 1: The 16th International Conference on Computational Linguistics*, vol. 1, 1996); Erik F. Tjong Kim Sang and Fien De Meulder, "Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition" (Antwerp: *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL* 2003-Volume 4, Association for Computational Linguistics, 2003), 142–47.
25. Kjersti Aas and Line Eikvil, "Text Categorisation: A Survey" (Norwegian Computing Center, Oslo, 1999).
26. Fabrizio Sebastiani, " A Tutorial on Automated Text Categorisation" (*Proceedings of ASAI-99, 1st Argentinian Symposium on Artificial Intelligence*, Buenos Aires, AR, 1999), 7–35; Enrique Alfonseca and Suresh Manandhar, "An Unsupervised Method for General Named Entity Recognition and Automated Concept Discovery" (*Proceedings of the 1st International Conference on General WordNet*, Mysore, India, 2002), 34–43.
27. Martha Palmer, Daniel Gildea, and Paul Kingsbury, "The Proposition Bank: An Annotated Corpus of Semantic Roles, " *Computational Linguistics* 31, no. 1 (2005): 71–106; Lluís Màrquez, Xavier Carreras, Kenneth C. Litkowski, and Suzanne Stevenson, "Semantic Role Labeling: An Introduction to the Special Issue," MIT Press (2008): 145–59.

28. Jason Burke, "Al Qaeda: Think Again," Foreign Policy No. 142 (Washington, DC: *Washington Post Newsweek*, 2004, May–June).
29. Irm Haleem, *The Essence of Islamist Extremism: Recognition through Violence, Freedom through Death* (London: Routledge, 2012).
30. Bateman, McDonald, Hiippala, Couto-Vale, and Costetchi, "Systemic Functional Linguistics And Computation" (see note 23).
31. Halliday, *An Introduction to Functional Grammar* (see note 14).
32. Tan, O'Halloran, Wignell, Chai, and Lange, "A Multimodal Mixed Methods Approach" (see note 10).
33. Matthew Honnibal and Ines Montani, "spaCy 2: Natural Language Understanding with Bloom Embeddings, Convolutional Neural Networks and Incremental Parsing." (2017).
34. Steven Bird, Ewan Klein, and Edward Loper, "Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit" (Sebastopol, Ca: O'Reilly Media, Inc., 2009).