

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/126962/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Macdonald, Stuart, Grinnell, Daniel, Kinzel, Anina and Lorenzo-Dus, Nuria 2019. Daesh, Twitter and the social media ecosystem: a study of outlinks contained in tweets mentioning Rumiya. *RUSI Journal* 164 (4) , pp. 60-72.
10.1080/03071847.2019.1644775

Publishers page: <https://www.tandfonline.com/doi/full/10.1080/03071...>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Daesh, Twitter and the Social Media Ecosystem: A Study of Outlinks Contained in Tweets Mentioning *Rumiyah*

Stuart Macdonald, Daniel Grinnell, Anina Kinzel and Nuria Lorenzo-Dus

The presence of Daesh (also known as the Islamic State of Iraq and Syria, ISIS) on Twitter has greatly diminished over the past five years as Daesh's propaganda dissemination strategy has evolved. Yet some Daesh supporters have persevered in their use of Twitter, using throwaway accounts to share outlinks to pro-Daesh materials on other platforms. This article analyses 892 outlinks found in 11,520 tweets that contained the word Rumiyah (Daesh's online magazine). It evaluates Twitter's response to attempts to use its platform to signpost users to Rumiyah in the context of the wider social media ecosystem and highlights the role played by botnet activity in efforts to disseminate the magazine and the impact of traditional news media coverage.

Five years have now passed since Daesh (also known as the Islamic State of Iraq and Syria, ISIS) enjoyed its 'Golden Age' on Twitter.¹ The *ISIS Twitter Census* conducted by Berger and Morgan in 2015 found that during October and November 2014 there were no fewer than 46,000 overt Daesh supporter accounts on Twitter – and possibly as many as 90,000. The average number of followers of these accounts was 1004, and each account posted an average of 7.3 tweets per day over its lifetime.² As well as proselytisation, recruitment and firming up the resolve of followers, Daesh utilised the new capabilities offered by social media to employ the platform for psychological warfare purposes.³

Since then Daesh's presence on Twitter has been reduced significantly. Towards the end of 2014 Twitter began an aggressive campaign of suspensions. Berger and Morgan found that, by February 2015, Daesh supporters on Twitter were having to devote far more time to rebuilding their networks.⁴ A follow-up study conducted by Berger and Perez also found that suspension activity had a significant disruptive effect.⁵ Individual users who repeatedly created new accounts after being suspended 'suffered devastating reductions in their follower counts' and declines in networks persisted even when suspension pressure eased, 'suggesting that suspensions diminish activity in ways that extend beyond the simple removal of accounts'.⁶

In response, IS supporters resorted to the use of a variety of countermeasures. These included locking their accounts so that they were no longer publicly accessible, using an innocuous image or the default egg as the avatar image and selecting a random combination of letters and numbers as the user handle or screen name.⁷ However, 'A conscious, supportive and influential virtual community is almost impossible to maintain in the face of the loss of access to such group or

¹ Maura Conway et al., *Disrupting Daesh: Measuring Takedown of Online Terrorist Material and its Impacts* (VOX-Pol Network of Excellence, 2017), p. 28.

² J M Berger and Jonathon Morgan, 'The ISIS Twitter Census: Defining and Describing the Population of ISIS Supporters on Twitter', Analysis Paper No. 20, March 2015, Brookings Institution, Washington, DC.

³ Jytte Klausen, 'Tweeting the *Jihad*: Social Media Networks of Western Foreign Fighters in Syria and Iraq', *Studies in Conflict and Terrorism* (Vol. 38, No. 1, 2015), pp. 1–22.

⁴ Berger and Morgan, 'The ISIS Twitter Census', p. 38.

⁵ J M Berger and Heather Perez, 'The Islamic State's Diminishing Returns on Twitter: How Suspensions are Limiting the Social Networks of English-speaking ISIS Supporters', George Washington University Program on Extremism, Washington, DC, February 2016.

⁶ *Ibid.*, p. 4.

⁷ Other countermeasures include reverse shoutouts, guides on how to hack and assume control of other user accounts and account banks. See, Berger and Perez, 'The Islamic State's Diminishing Returns on Twitter'.

ideological symbols and the resultant breakdown in commitment'.⁸ Unsurprisingly, therefore, Daesh's community-building activities largely moved to other platforms, in particular Telegram.⁹ Importantly, in their study of English-speaking Daesh supporters on Telegram, Clifford and Powell found that, whilst Telegram's suite of features is used by Daesh supporters to interact and communicate, to distribute joinlinks to other groups and channels and to provide instructional materials, by far the most common function is the distribution of core IS media and, in particular, other pro-IS materials (regardless of their origin).¹⁰ As well as 'using Telegram's file-sharing features to disseminate content internally, IS sympathizers on Telegram use external file-sharing sites to ensure IS content remains on the internet and resilient to takedowns'.¹¹ A single piece of pro-Daesh material may be distributed using dozens of unique URLs on different file-sharing sites so that, even if content is removed from one site, stable access exists to others. File-sharing platforms are thus utilised as 'black boxes' to 'enable the rapid redistribution of content even under conditions of drastic policing and filtering'.¹² The result is a 'fragmentation' of Daesh propaganda that makes these materials 'less trackable by authorities' and results in a 'relatively closed and stable digital propaganda ecosystem'.¹³

Whilst Daesh's dissemination strategy has evolved, some Daesh supporters have persisted in their use of Twitter. This continued use largely focuses on the sharing of URLs made available on Telegram.¹⁴ As Conway et al conclude in their report *Disrupting Daesh*, Daesh's Twitter activity 'has largely been reduced to tactical use of throwaway accounts for distributing links to pro-IS content on other platforms, rather than as a space for public IS support and influencing activity'.¹⁵ Macdonald et al.'s examination of Daesh's attempts to use Twitter to disseminate its online magazines *Dabiq* and *Rumiyah* found that the throwaway disseminator accounts established by Daesh sympathisers were mostly recently established (often less than one day old at the time of suspension), had very few followers (sometimes none at all) and received few retweets.¹⁶ Some sought to compensate for this lack of visibility by repeat posting.

Given that this past research shows that Daesh throwaway disseminator accounts gain little traction on Twitter, it is interesting that Daesh and its supporters still persevere with this approach. To better

⁸ Maura Conway et al., *Disrupting Daesh: Measuring Takedown of Online Terrorist Material and its Impacts* (VOX-Pol Network of Excellence, 2017), p. 30.

⁹ Nico Prucha, 'IS and the Jihadist Information Highway – Projecting Influence and Religious Identity via Telegram', *Perspectives on Terrorism* (Vol. 10, No. 6, 2016), pp. 48–58.

¹⁰ Bennett Clifford and Helen Powell, 'Encrypted Extremism: Inside the English-Speaking Islamic State Ecosystem on Telegram', George Washington University Program on Extremism, Washington, DC, June 2019.

¹¹ *Ibid.*, p. 24.

¹² Teodor E Mitew and Ahman Shehabat, 'Black-boxing the Black Flag: Anonymous Sharing Platforms and ISIS Content Distribution Tactics', *Perspectives on Terrorism* (Vol. 12, No. 1, 2018), pp. 81–99, 97.

¹³ Laurence Bindner and Raphael Gluck, 'Trends in Islamic State's Online Propaganda: Shorter Longevity, Wider Dissemination of Content', ICCT Perspective, International Centre for Counter-Terrorism, The Hague, December 2018. See also Samantha Weirman and Audrey Alexander, 'Hyperlinked Sympathizers: URLs and the Islamic State', *Studies in Conflict and Terrorism* (April 2018), DOI: 10.1080/1057610X.2018.1457204.

¹⁴ Clifford and Powell, 'Encrypted Extremism'.

¹⁵ Conway et al., 'Disrupting Daesh: Measuring Takedown of Online Terrorist Material and its Impacts', p. 30.

¹⁶ Daniel Grinnell et al., 'Who Disseminates Rumiyah? Examining the Relative Influence of Sympathiser and Non-Sympathiser Twitter Users', paper presented at the 2nd European Counter Terrorism Centre Advisory Group conference, Europol Headquarters, The Hague, 17–18 April 2018, <<https://www.europol.europa.eu/publications-documents/who-disseminates-rumiyah-examining-relative-influence-of-sympathiser-and-non-sympathiser-twitter-users>>, accessed 12 July 2019; Daniel Grinnell, Stuart Macdonald and David Mair, 'The Response of, and on, Twitter to the Release of Dabiq Issue 15', paper presented at the 1st European Counter Terrorism Centre conference on online terrorist propaganda, Europol Headquarters, The Hague, 10–11 April 2017, <<https://www.europol.europa.eu/publications-documents/response-of-and-twitter-to-release-of-dabiq-issue-15>>, accessed 12 July 2019.

understand this phenomenon, how Daesh uses the platform and its place within the wider social media ecosystem, this article examines a different aspect of Daesh’s propaganda dissemination strategy. Specifically, it conducts an analysis of the outlinks contained in *Rumiyah*-mentioning tweets. After briefly outlining the methodology, the article first examines the hostnames found in the outlinks. In keeping with other studies, it highlights the prevalence of smaller platforms and file-sharing sites. Second, it examines the types of content to which the outlinks led. As well as PDFs of *Rumiyah*, the article shows that a significant proportion of outlinks led to reports and information about *Rumiyah* and explains that the effect of some traditional news media coverage is to amplify the message contained in *Rumiyah*. Third, the analysis focuses on outlinks to PDFs of *Rumiyah* and shows that Twitter was largely successful in frustrating efforts to use its platform to disseminate the issues of the magazine in the dataset. The conclusion offers some practical recommendations.

Methodology

Data Collection

Data was collected between 1 November 2016 and 31 October 2017 using the ‘Sentinel’ research tool.¹⁷ All collected Twitter posts: mentioned the term ‘Rumiyah’; were posted within 21 days of the release of a new issue; and were posted from an account that used the English language interface (US or UK). The last of these reflected the authors’ decision to focus specifically on users that posted about the English language version of *Rumiyah* (which is also published in multiple other languages). The authors also collected the publicly available user data of these posters, the details of each post (including when it was posted and whether it was a retweet), the onward distribution counts of these posts, and the account status (at the end of the data collection period).¹⁸ From these data, the authors extracted all posts that: contained original content (that is to say, were not retweets); and contained an outlink.

As Table 1 shows, 11 issues of *Rumiyah* were published during the data collection period (issues 3 to 13). As a result of collection drop outs, the data collection for issues 6 and 8 was incomplete.¹⁹ These issues were therefore excluded from the study. The dataset thus encompasses a total of nine issues.

Table 1: Issues for which data was collected

Table 1: Issues for which data were collected			
Issue	Date and time of first tweet collected	Date and time of last tweet collected	Notes
3	11/11/2016 17:18	02/12/2016 17:18	
4	07/12/2016 19:33	28/12/2016 19:33	
5	06/01/2017 16:57	27/01/2017 16:57	
6	04/02/2017 18:53	25/02/2017 18:53	Incomplete collection so not included in the dataset

¹⁷ Alun Preece et al., ‘Sentinel: A Codesigned Platform for Semantic Enrichment of Social Media Streams’, IEEE Transactions on Computational Social Systems (Vol. 5, No. 1, March 2018), pp. 118–31. The Sentinel platform supports semantic enrichment of streamed social media data for the purposes of situational understanding. It is founded upon a knowledge-based approach, in which input streams (channels) are characterized by spatial and terminological parameters, collected media is pre-processed to identify significant terms (signals), and data are tagged (framed) in relation to an ontology.

¹⁸ For the purposes of this study, Sentinel functioned only as a repository of structured data supplied by the Twitter Streaming API.

¹⁹ These were short-term drop outs in research data capture infrastructure that resulted in missing information in the collection.

7	07/03/2017 16:04	28/03/2017 16:04	
8	05/04/2017 15:23	26/04/2017 15:23	Incomplete collection so not included in the dataset
9	04/05/2017 14:40	25/05/2017 14:40	
10	08/06/2017 21:32	29/06/2017 21:32	Timing moderately uncertain due to similarly timed presence of a 'fake issue'
11	13/07/2017 15:00	03/08/2017 15:00	Timing highly uncertain due to hashtag flooding and a 'fake issue'
12	06/08/2017 14:47	27/08/2017 14:47	
13	09/09/2017 18:58	30/09/2017 18:58	

Source: Author's research.

An Overview of the Dataset

There was a total of 11,520 posts that contained both original content and an outlook. These 11,520 posts contained a total of 892 distinct links and were posted by a total of 1,493 distinct user accounts.

The data inclusion criteria did not include any requirement that the user posting the outlook was a Daesh sympathiser. It is important to note, therefore, that a previous analysis examined the profiles of the user accounts in the dataset.²⁰ This found that roughly two-thirds (67.1%) had been suspended by the end of the data collection period. The users that were not suspended included personal accounts, news organisations, intelligence analysts/practitioners and researchers. The tone of these users' posts was generally either factual or critical.²¹ By contrast, almost all of the suspended accounts (95.8%) were overtly sympathetic to Daesh, and there were hardly any users that were overtly sympathetic Daesh that were not suspended (0.3%).²² In the analysis that follows users who were suspended are accordingly regarded as Daesh sympathisers.

Data Analysis

The authors analysed data in three stages. In the first stage, they examined the hostnames found in the 892 distinct outlooks. The 20 most common hostnames were identified using three different measures: the number of posts containing the hostname; the number of distinct links containing the hostname; and by the number of distinct users outlooking to it. They then followed each of the 892 outlooks to see what type of content it linked to. When the link no longer worked, they instead looked at the wording of the URL and of the post in which it was contained. Completing this task enabled the authors to identify which of the outlooks took users to a PDF of *Rumiyah* (or had been intended to do so). The third stage of the analysis focused on these links by examining their hostnames, how many times they had been shared and the account status of those that shared them.

²⁰ It should also be noted that, since the previous study focused on *users* who posted outlooks, the dataset was limited to just the first *Rumiyah*-mentioning tweet posted by each user following the release of each issue. By contrast, since the current study focuses on *posts* containing outlooks, the dataset was not limited in this way and so contains all tweets that met the inclusion criteria detailed in the text.

²¹ Grinnell et al, 'Who Disseminates *Rumiyah*?' (see Table 6).

²² Of the 1,517 user accounts examined in the authors' previous study, 1,018 had been suspended by the end of the data collection period. Of these 1,018, 975 were overtly sympathetic to Daesh. See Table 7 and footnote 3 in Grinnell et al., 'Who disseminates *Rumiyah*?'

Outlinking to Where?

The outlinks in the 11,520 posts in the dataset contained a total of 244 different hostnames. Table 2 shows the 20 most common hostnames, respectively ordered by the number of posts and the number of distinct links containing the hostname.

Table 2: Top 20 Hostnames, by Number of Posts and Distinct Links

Table 2: Top 20 hostnames, by number of posts and distinct links					
<i>By number of posts</i>			<i>By number of distinct links</i>		
Hostname	Site type	Total	Hostname	Site type	Total
ref.gl	URL shortener	3733	ref.gl	URL shortener	84
drive.google.com	File sharing site	1634	memri.org	Website of a US-based not-for-profit organisation	69
cloud.mail.ru	File sharing site	878	archive.org	File sharing site	58
justpaste.it	File sharing site	535	drive.google.com	File sharing site	54
cldup.com	File sharing site	475	justpaste.it	File sharing site	39
archive.org	File sharing site	464	facebook.com	Social network	29
yadi.sk	File sharing site	417	4shared.com	File sharing site	19
1drv.ms	File sharing site	336	cldup.com	File sharing site	17
jpst.it	File sharing site	329	siteintelgroup.com	Website of a terrorism research/analysis group (paid subscription service)	15
dropbox.com	File sharing site	188	cloud.mail.ru	File sharing site	14
pc.cd	Apparently a URL shortener	160	mediafire.com	File sharing site	14
uploaded.net	File sharing site	125	1drv.ms	File sharing site	13
memri.org	Website of a US-based not-for-profit organisation	112	pietervanostaeyen.com	Privately run website (password protected)	12
uptobox.com	File sharing site	107	dropbox.com	File sharing site	12
4shared.com	File sharing site	104	yadi.sk	File sharing site	11
Weather.com	Weather site	104	almlf.com	File sharing site	11
turbobit.net	File sharing site	90	clarionproject.org	Website of a US-based not-for-profit organisation	9
cloudup.com	File sharing site	87	cloudup.com	File sharing site	9
load.to	File sharing site	80	dailymail.co.uk	Website of a UK-based newspaper	8
filefactory.com	File sharing site	70	sdb.esisc.org	Terrorism database (paid subscription service)	7

Source: Authors' research.

The prevalence of file sharing sites and smaller platforms in Table 2 is consistent with previous studies.²³ Twelve hostnames appear on both lists, ten of which are file sharing sites (drive.google.com, cloud.mail.ru, justpaste.it, cldup.com, archive.org, yadi.sk, 1drv.ms, dropbox.com, memri.org, 4shared.com, cloudup.com).²⁴ Five of these hostnames appear in the top ten of each list. In addition, of the sixteen other hostnames that only appear in one of the lists, eight are file sharing sites (jpst.it, uploaded.net, uptobox.com, turbobit.net, load.to, filefactory.com, mediafire.com, almlf.com). YouTube appears on neither list. Facebook appears on the 'distinct links' list only.²⁵

The importance of file-sharing sites is clearer still in Table 3, which lists the top 20 hostnames in order of the number of distinct users outlinking to them, as well as the suspension status of these users at the end of the data collection period.

Table 3: Top 20 Hostnames, by Number of Distinct Users Outlinking to it

Hostname	Site Type	Suspension Status		Total
		Extant	Suspended	
drive.google.com	File sharing site	3	561	564
cloud.mail.ru	File sharing site	1	351	352
cldup.com	File sharing site	3	291	294
archive.org	File sharing site	6	215	221
1drv.ms	File sharing site	0	180	180
dropbox.com	File sharing site	0	133	133
yadi.sk	File sharing site	1	124	125
pc.cd	Apparently a URL shortener	0	102	102
4shared.com	File sharing site	0	70	70
cloudup.com	File sharing site	1	60	61
justpaste.it	File sharing site	4	40	44
mediafire.com	File sharing site	0	36	36
counterjihadreport.com	Privately run website (password protected)	30	0	30
icct.nl	Website of a Netherlands-based think tank	29	1	30
memri.org	Website of a US-based not-for-profit organisation	29	0	29
up.top4top.net	File sharing site	1	26	27
clarionproject.org	Website of a US-based not-for-profit organisation	20	0	20

²³ See Conway et al., 'Disrupting Daesh'; see also and Clifford and Powell, 'Encrypted Extremism'.

²⁴ The other two are ref.gl (which is discussed further in the article) and memri.org.

²⁵ The data in this table differs from the table of outlinks contained in the authors' previous study, 'Who Disseminates *Rumiyah*?'. This is because the previous study focused on *users* who posted outlinks (so the dataset consisted of the first *Rumiyah*-mentioning tweet posted by 1,392 distinct users), whereas this article focuses on *posts* containing outlinks (so the dataset consists of all 11,520 tweets that contained an outlink).

terrortrendsbulletin.com	Personal blog	19	0	19
heavy.com	News and information website	15	4	19
express.co.uk	Website of a UK-based newspaper	17	1	18 ²⁶

Eleven of the top twelve hostnames in this Table were file sharing sites. The suspension rate for the users that posted the links to these hostnames was 99.1%.²⁷ As noted above, the vast majority of the user accounts in the dataset that were suspended were overtly sympathetic to Daesh. Table 3 thus indicates which file-sharing sites are most commonly used by Daesh sympathisers seeking to signpost other users to copies of *Rumiyah*. Again, the findings are consistent with previous studies.²⁸

Also noteworthy is the fact that ref.gl – which was top of both lists in Table 2 – does not feature in Table 3 at all. Known as the ‘Reffy Botnet’, ref.gl is a URL shortener that is commonly used in networks of ill-intent. In April 2018, all accounts using the ref.gl shortener were suspended by Twitter.²⁹ In the dataset, ref.gl was used in a total of 84 distinct links. These links appeared in a total of 3,733 posts (32.4% of the posts in the dataset). Just nine user accounts were responsible for these posts (and one of these accounts only posted a single tweet).³⁰ Table 4 shows the date of the first and last post captured that contained the ref.gl hostname in the three-week data collection period following the release of each relevant issue.

Table 4: Distribution of Tweets Posting an Outlink Containing the ref.gl Hostname

Issue	10	11	12	13
Date and time of first post following issue’s release	19 June 2017, 19:35	13 July 2017, 15:08	6 August 2017, 15:15	9 September 2017, 20:11
Date and time of last post collected following issue’s release	29 June 2017, 20:45	3 August 2017, 14:15	27 August 2017, 13:07	30 September 2017, 18:19
Total number of tweets	679	1086	877	1091

For issues 11, 12 and 13 the first post containing a ref.gl outlink appeared soon after the release of the new issue on Twitter – in the case of issue 11, just eight minutes after (see Table 1). This, coupled with the sheer volume of tweets (over 50 a day on average) and small number of user accounts posting them, are indicative of botnet activity.

Interestingly, the eight user accounts that were responsible for all but one of the posts using the ref.gl shortener were also responsible for the 104 posts outlinking to Weather.com (see Table 2), as

²⁶ Facebook also had a total of 18 distinct users (18 extant, 0 suspended) and so was joint 20th.

²⁷ Based on a total for the eleven hostnames of 19 extant users and 2,061 suspended users.

²⁸ See Conway et al., ‘Disrupting Daesh’, see also Clifford and Powell, ‘Encrypted Extremism’.

²⁹ Mike Farb, ‘The Reffy Botnet’, Medium, 29 April 2018, <<https://medium.com/@unhackthevote/the-refly-botnet-f8a7dc817e9a>>, accessed 18 June 2019.

³⁰ The nine accounts were @MiddleBeast, @ExtremePropa, @ExtremistWatch_, @Islamoinform, @JihadiInfo, @TabsTerror, @terror_history1, @terrorwatch1 and @VicPower87. The last of these was responsible for just one of the 3,733 tweets. All nine accounts have since been suspended.

well as a further 69 posts outlinking to AccuWeather.com and 59 posts outlinking to Climate-Data.org. This apparently indiscriminate use of the word ‘Rumiyah’ is also suggestive of botnet activity.

The hostname pc.cd also appeared to be connected to botnet activity. This hostname appeared in a total of 160 posts. These were all posted in the space of just over 25 hours,³¹ by a total of 102 distinct user accounts (whose user names were randomised collections of letters). All of these user accounts were subsequently suspended.

Outlinking to What?

Table 5 breaks down the 892 distinct outlinks, by the type of content each link led to. For those outlinks that no longer worked (for example, because the destination page had been removed or the hostname had been suspended), the type of content was determined by examining the text of both the URL and the post. Given the inherent limitations of relying on the wording of the URL and post, the authors drew a distinction between categorisations in which they had a high degree of confidence and those in which they had only a moderate degree of confidence. As Table 5 shows, even after completing this process, there remained a total of 29 outlinks that were so unclear that they could not be categorised.

Table 5: Types of Content Behind the Outlinks

Type of content	Number of distinct links (high confidence)	Number of distinct links (moderate confidence)	Total
Report/summary/information about <i>Rumiyah</i>	232	90	322 (36.1%)
PDF of <i>Rumiyah</i> (but no longer available)	79	228	307 (34.4%)
PDF of <i>Rumiyah</i> behind subscription or password protection	56	3	59 (6.6%)
Other content (not terrorism-related)	39	11	50 (5.6%)
Other content (terrorism-related)	27	4	31 (3.5%)
Academic analysis/writing	14	16	30 (3.3%)
Not possible to tell	-	-	29 (3.3%)
Picture of real <i>Rumiyah</i> issue	16	5	21 (2.4%)
Picture of fake <i>Rumiyah</i> issue	13	5	18 (2.0%)
PDF of <i>Rumiyah</i> (currently available)	15	0	15 (1.7%)
Suspicious “Download” button	5	0	5 (0.6%)
Fake issue of <i>Rumiyah</i>	5	0	5 (0.6%)
Total	501	362	892

A total of 381 of the outlinks led to a PDF of *Rumiyah*. Of these, 307 were no longer available, 59 were behind a subscription requirement or password protection and just 15 led directly to the PDF. These links are discussed further below.

³¹ From 20:06 on 9 September 2017 to 21:26 on 10 September 2017.

The next most common type of content was report/summary/information about *Rumiyah*. This category consisted largely of news items either specifically about a new issue of the magazine or in which *Rumiyah* featured prominently. In the entire dataset, there were just three posts that were retweeted more than 50 times during the data collection period.³² All three outlinked to the same news item in the UK's *Daily Mail* newspaper, titled 'ISIS calls on Islamists to carry out knife attacks in areas such as alleys, forests and quiet neighbourhoods and told to aim for "a reasonable kill count" in latest magazine'.³³

The *Daily Mail* was one of two UK newspapers that, in the aftermath of the recent Christchurch attack, posted the video of the attack on its website. This has raised concerns about the traditional news media, especially as it has been reported that the attack video only went viral after this happened.³⁴ The dataset raises similar questions. For example, one outlink led to an item published by the UK's *Metro* newspaper in September 2016. It describes how issue 1 of *Rumiyah* states that killing disbelievers is a form of worship to Allah, 'even the blood of the kafir street vendor selling flowers to those passing by'.³⁵ Below this is a photo of a market trader at a flower stall, which the magazine's producers had apparently taken from the trader's website. The magazine offers no indication of the trader's name or the location of his flower stall. By contrast, the item in *Metro* names the trader and states the area in which his stall is located. A simple Google search reveals that several other UK newspapers ran a similar story, including a local newspaper that named the street in which the flower stall can be found. The effect of this news coverage was thus to amplify – and, importantly, to sharpen – the message contained in *Rumiyah*.

There were 17 links in the dataset that led to a picture of *Rumiyah*. These included pictures of the front cover, the table of contents and excerpts from the foreword or a specific article.

For all nine issues of *Rumiyah* in the dataset, the release of the official issue was preceded by fake versions.³⁶ The fake version of issue 5, for example, included articles on similar themes to earlier issues, whilst fake issue 7 contained the exact titles of articles that appeared in the official copy of the same issue, revealing that an insider was involved in the disinformation campaign.³⁷ The

³² There were only 22 posts that were retweeted ten times or more. Four of these outlinked to password-protected copies of *Rumiyah*. One outlinked to a PDF of *Rumiyah* that is no longer available. The others outlinked to news items (eight posts), academic analyses (six posts) and other terrorism-related content (three posts).

³³ Hannah Al-Othman, 'ISIS Calls on Islamists to Carry Out Knife Attacks in Areas such as Alleys, Forests and Quiet Neighbourhoods and Told to Aim for "a Reasonable Kill Count" in Latest Magazine', *Daily Mail*, 5 October 2016. The three tweets received 220, 54 and 52 retweets respectively.

³⁴ Tech Against Terrorism, 'Analysis: New Zealand Attack and the Terrorist Use of the Internet', undated, <<https://www.techagainstterrorism.org/2019/03/26/analysis-new-zealand-attack-and-the-terrorist-use-of-the-internet/>>, accessed 18 June 2019.

³⁵ Ashitha Nagesh, 'Isis Magazine Urges Followers to Kill a Random Florist in Cheshire', *Metro*, 7 September 2016.

³⁶ The tweets in the dataset outlinked to fake versions of six issues (4, 5, 9, 11, 12 and 13). It is important to add, however, that, for each issue in the study, the data collection commenced once the official publication had been released. Since the fake issues tended to be released before the official one, it is possible that there were tweets outlinking to fake issues that were not part of the dataset. For example, as the main text explains, there was a fake version of issue 7, but none of the tweets in the dataset outlinked to it. As Table 1 above shows, there was also a fake version of issue 10, but again none of the tweets in the dataset outlinked to it. And it has been reported that there was a fake version of issue 3 (Antonis Samouris, 'Jihadist Strategies on the Internet: How the Decline in IS Official Propaganda Raises the Visibility of Pro-IS User-Generated Content', in Andreas Gofas (ed.), *Terrorism and European Security Governance* (Florence: European University Institute, 2018)). In short, there were fake versions of all nine issues in the study.

³⁷ Samouris, 'Jihadist Strategies on the Internet'.

objective of these fake issues was apparently to confuse supporters and/or to collect information on them through embedded malware.³⁸ In response, Daesh sympathisers were encouraged to validate the authenticity of any new content by checking for its simultaneous appearance on the Telegram channels of core disseminators. This system of validation meant that the fake issues were not picked up by the community of Daesh sympathisers.³⁹ Indeed, in the dataset there were only five links that led to a PDF of a fake issue of the magazine.⁴⁰ Five other links led to a suspicious download button that promised a copy of the magazine – the authors did not click on these, for security reasons. There were also 16 links to pictures of fake issues of *Rumiyah*. These were largely pictures of a front cover. Some of the fake front cover pictures were accompanied by a download button, which again the authors did not click for security reasons.

The remaining three categories were: academic analysis/writing (30 links); other content (terrorism-related) (31 links); and, other content (not terrorism-related) (50 links). The second of these categories covered content that, whilst unrelated to *Rumiyah*, did focus more generally on the activities of Daesh or another terrorist actor or group. The other content (not terrorism-related) was an eclectic mix. As well as the links to weather forecasts mentioned above, it also included such things as advertisements for a social media analytics company and a baby name website (offering the meaning of the name ‘Rumiyah’).

Finally, it should be noted that there were instances of an outlink leading to a news item or blog (and so was counted in the report/summary/information category), where the item/blog in turn contained a link to an issue of *Rumiyah*. There were four instances of this. Of these, two of the sub-links led to a freely available copy of *Rumiyah*, one led to a copy that was behind password protection and the final one no longer worked.

Outlinking to Full Copy PDFs of *Rumiyah*

The results presented so far have shown the variety of content that the *Rumiyah*-mentioning tweets in the dataset outlinked to, and indicated which file-sharing sites are most commonly used by Daesh sympathisers seeking to signpost other users to copies of *Rumiyah*.

The 381 outlinks to a PDF of *Rumiyah* contained a total of 48 different hostnames. These are listed in Table 6. Table 6 also shows the number of distinct outlinks containing each hostname, as well as the total number of posts containing these outlinks and the number of reposts these posts received. (Hostnames for which there were just one or two distinct links appear in a single category, ‘Other hostnames’).

Table 6: Outlinks to a PDF of *Rumiyah*, by Hostname and Number of Posts

Hostname	Number of distinct links to a PDF	Number of posts containing a link to the PDF	Number of reposts containing a link to the PDF	Total number of posts and reposts
drive.google.com	53	1631	85	1716

³⁸ *Ibid.*

³⁹ *Ibid.*

⁴⁰ These were identified as being fake by comparing them to the issues of *Rumiyah* stored in repositories maintained by subject experts.

archive.org	53	442	171	613
ref.gl	32	1450	81	1531
justpaste.it	29	82	138	220
cldup.com	26	562	70	632
4shared.com	18	103	15	118
cloud.mail.ru	13	877	37	914
siteintelgroup.com	13	27	70	97
mediafire.com	13	44	11	55
1drv.ms	12	335	10	345
dropbox.com	12	188	13	201
pietervanostaeyen.com	12	17	25	42
memri.org	12	15	2	17
yadi.sk	11	417	4	421
almlf.com	10	15	7	22
counterjihadreport.com	8	46	31	77
up.top4top.net	5	27	10	37
trackingterrorism.org	5	5	6	11
facebook.com	3	3	0	3
<i>Other hostnames</i>	<i>41</i>	<i>522</i>	<i>14</i>	<i>536</i>
Grand Total	381	6808	800	7608

Source: Authors' research

Two points in particular emerge from Table 6. The first is that 5,714 (83.9%) of the posts outlinked to just seven of the hostnames: drive.google.com; archive.org; ref.gl; cldup.com; cloud.mail.ru; 1drv.ms; and, yadi.sk. If ref.gl – which was discussed above – is excluded, a total of 4,264 (62.6%) of the posts outlinked to the other six hostnames. Table 3 above showed that a total of 1,736 distinct users posted links to these six hostnames. Of these, 1,722 (99.2%) had been suspended by the end of the data collection period. There was thus a concerted effort by Daesh sympathisers to use these six file-sharing sites as black boxes for the distribution of *Rumiyah*.⁴¹

The second point is the ratio of posts to reposts. Between them, the 6,808 posts containing an outlink to *Rumiyah* received a total of just 800 reposts during the data collection period (that is to say, 8.51 tweets per retweet). For the six file-sharing sites listed in the previous paragraph, the ratio is even higher (11.31 tweets per retweet). Together, the high suspension rate and low number of reposts for users outlinking to these six sites indicate that Twitter was successful in frustrating efforts to use its platform to disseminate new issues of *Rumiyah*.

Table 7 focuses exclusively on the fifteen outlinks to openly available PDFs of *Rumiyah*.⁴²

Table 7: Outlinks to Currently Available PDFs of *Rumiyah*

--

⁴¹ Mitew and Shehabat, 'Black-boxing the Black Flag'.

⁴² The researchers shared the fifteen URLs with the relevant authorities. Some have subsequently been removed.

Hostname ⁴³	Number of posts containing the link	Number of times the posts have been reposted	Number of users that posted the links	Status of user accounts
cloud.mail.ru	154	0	154	All suspended
qb5cc3pam3y2ad0tm1zxuhho-wpengine.netdna-ssl.com	14	3	6	Five extant, one suspended
adobe.ly	4	0	2	Both extant
drive.google.com	2	1	2	One extant, one suspended
azelin.files.wordpress.com	2	1	1	Extant
azelin.files.wordpress.com	2	0	1	Extant
jihadology.net	1	0	1	Extant
cloud.mail.ru	1	0	1	Suspended
cloudup.com	1	0	1	Suspended
pietervanostaeyen.com	1	0	1	Extant
clarionproject.org	1	2	1	Extant
magentacloud.de	1	0	1	Suspended
magentacloud.de	1	0	1	Suspended
reddit.com	1	0	1	Suspended
reddit.com	1	0	1	Suspended
Total	187	7	175	

Source: Author's research.

Whilst the fifteen outlinks appeared in a total of 187 posts, 154 of these posts contained the same URL. These 154 tweets were all posted by different users. The names of all these users were randomised collections of numbers and letters, and all accounts had been suspended by the end of the data collection period – though curiously the outlink remained functional. The other fourteen outlinks appeared in a combined total of 33 posts. These tweets were posted by a total of 19 distinct users.⁴⁴ They received just seven retweets.

Furthermore, the fact that twelve of the user accounts remain extant is not indicative of a failure on Twitter's part to enforce its terms of service (the *Twitter Rules*). The rules prohibit promoting and recruiting for a violent extremist group. In accordance with this, and as the authors' previous studies found, the accounts of non-Daesh sympathisers who post outlinks to the group's magazine (for example, for research purposes or general interest) are not suspended.⁴⁵

Following from this, it is noteworthy that four of the URLs in Table 7 outlink to repositories maintained by researchers (and a fifth outlinked to a repository maintained by a NGO, the Clarion

⁴³ For ethical reasons, the article lists only the hostname and not the complete URL. For this reason, some hostnames appear more than once.

⁴⁴ The figures in the relevant column in Table 7 add up to 21. The reason for this apparent disparity is that there were two users that shared more than one of the outlinks. (The same user shared both of the outlinks to reddit.com, and another user shared two distinct links to Jihadology).

⁴⁵ Daniel Grinnell et al., 'Who Disseminates Rumiyaq?'; Daniel Grinnell, Stuart Macdonald and David Mair, 'The Response of, and on, Twitter to the Release of Dabiq Issue 15'.

Project). Three of these URLs outlinked to the website Jihadology.⁴⁶ This site has received much scrutiny in recent months – in 2018, reports indicated that the UK government had urged WordPress.com to place the site’s contents behind password protection or close it altogether.⁴⁷ In April 2019, it was announced that, with funding from the Global Internet Forum to Counter Terrorism (GIFCT; a consortium of technology companies, founded by Facebook, Google, Twitter and Microsoft, that collaborate to tackle terrorist exploitation of their services), the Tech Against Terrorism initiative had developed a new interface for Jihadology, to ensure that particularly sensitive content is only accessible to users with registered academic/research, governmental, journalistic or humanitarian email addresses.⁴⁸ In terms of this specific study, however, it seems clear that Daesh sympathisers did not seek to use Twitter to signpost users to copies of *Rumiyah* on Jihadology. Not only were there only five posts containing outlinks to Jihadology in the entire dataset of 11,520 tweets – with these five posts receiving a total of just one retweet during our data collection period – none of these five tweets was posted by an Daesh sympathiser.⁴⁹

The other URL outlinked to the site pietervanostaeyen.com. This site is also hosted by WordPress.com, but is password protected. Users are required to register for an account. In spite of this, the outlink we collected took us directly to a PDF of an issue of *Rumiyah* without requiring us to enter a password.⁵⁰ Moreover, the owner has in any event publicly stated that he approves every request he receives for access to the website, explaining that he lacks the capacity to vet those who request access.⁵¹ This illustrates the wider point that smaller platforms may not be able to regulate their platforms adequately in the absence of additional support.

Conclusion and Recommendations

This article examined a total of 892 distinct outlinks found in tweets mentioning *Rumiyah* that were posted to Twitter between November 2016 and September 2017. It has shown how Daesh sympathisers sought to disseminate issues of the magazine by posting links to a number of different file-sharing sites, including several smaller platforms. This approach is in keeping with Daesh’s wider dissemination strategy, designed to ensure stable access to its propaganda outputs, although interestingly there was no evidence of Daesh seeking to signpost Twitter users to copies of *Rumiyah* that are freely available from repositories maintained by researchers or NGOs. The article also highlighted the role of botnet activity in efforts to disseminate *Rumiyah*, in particular the Reffy Botnet.

Twitter’s response to Daesh’s attempts to use the platform as a gateway to *Rumiyah* appeared effective. Each of the instances of botnet activity that uncovered resulted in suspension. There was a very high suspension rate (99.2%) for users who posted outlinks to a PDF of *Rumiyah* hosted on one of the six file-sharing sites apparently used by Daesh and the posts containing these outlinks received relatively few reposts.

⁴⁶ The two links using the hostname azelin.files.wordpress.com, plus the one using jihadology.net.

⁴⁷ David Bond, ‘How Extremist Videos are Hitting UK Relations with US Tech Groups’, Financial Times, 3 December 2018.

⁴⁸ Tech Against Terrorism, ‘Launching an Updated Version of Jihadology to Limit Terrorist Exploitation of the Site’, press release, 10 April 2019, <<https://www.techagainstterrorism.org/2019/04/10/press-release-10th-april-2019-launching-an-updated-version-of-jihadology-to-limit-terrorist-exploitation-of-the-site/>>, accessed 12 July 2019.

⁴⁹ The three users that posted these five tweets were two academic researchers and an individual tweeting in a personal capacity.

⁵⁰ The authors tested this for every other issue of *Rumiyah* stored on the website and found that this problem only existed for this individual issue.

⁵¹ Bob Garfield, ‘Archiving Terrorist Propaganda’, WNYC Studios, 22 March 2019, <<https://www.wnycstudios.org/story/archiving-terroristpropaganda-jihadology>>, accessed 18 June 2019.

The article has also shown that other Twitter users, not sympathetic to Daesh, also post links to reports about *Rumiyah*, academic analyses of it and PDFs of the magazine itself. Since the *Twitter Rules* focus on the promotion of, and recruitment for, violent extremist groups, these non-sympathisers were not suspended from the platform. Yet at the same time it is important to note that traditional news media coverage has had the effect of amplifying the message contained in *Rumiyah* – and posts outlinking to such coverage received the most retweets of all the posts in our dataset.

In conclusion, we offer the following three recommendations based on our findings. First, larger social media companies have automated means that employ behavioural cues to block content (such as abnormal posting volume or using trending hashtags to gain attention).⁵² This is valuable in the present context, given that botnet activity played a significant role in efforts to disseminate *Rumiyah*. By contrast, many smaller companies rely exclusively on humans to use content-based cues to identify and remove terrorist content. Where possible, GIFCT members should develop shared automated systems that use behavioural cues to block terrorist content.

Second, there is a pressing need to expand membership of the GIFCT. At present the network has fourteen members, a small number in comparison to the 244 different hostnames contained in the dataset. Many smaller technology companies lack the capacity needed to meet the standards imposed by the GIFCT eligibility criteria.⁵³ Some lack the willingness to abide by these criteria. Here policymakers have an important role to play, providing the support required by the former and offering appropriate incentives to the latter.

Finally, the role of the traditional news media in the dissemination of terrorist propaganda must be addressed. Dialogue is needed between the GIFCT and the traditional news media regarding the use of social media to share news coverage that has the effect of amplifying the terrorist message. As the UK's Counter Terrorism Policing Lead, Assistant Commissioner Neil Basu, stated in an open letter following the attacks in Christchurch, New Zealand, 'it's time to have a sensible conversation about how to report terrorism in a way that doesn't help terrorists'.⁵⁴

This article is an extended version of the Research Paper A Study of Outlinks Contained in Tweets Mentioning Rumiyah, which was published by the Global Research Network on Terrorism and Technology

Acknowledgements: The authors would like to thank the anonymous reviewers for their detailed comments on an earlier draft. They also gratefully acknowledge the funding for this research provided by the Global Research Network on Terrorism and Technology and by Swansea University's CHERISH-Digital Economy Centre (EPSRC grant number EP/M022722/1).

⁵² Isabelle van der Vegt, Paul Gill, Stuart Macdonald and Bennett Kleinberg, 'Shedding Light on Terrorist and Extremist Content Removal', Global Research Network on Terrorism and Technology No. 3, RUSI, July 2019.

⁵³ These can be found at <https://gifct.org/members/>.

⁵⁴ Neil Basu, 'Open Letter From Neil Basu To Deny Terrorists A Voice', 20 March 2019, <<https://www.counterterrorism.police.uk/acso-neil-basu-issues-open-letter-about-reporting-of-terrorism/>>, accessed 9 July 2019.