

Who could be behind QAnon? Authorship attribution with supervised machine-learning

FLORIAN CAFIERO* and JEAN-BAPTISTE CAMPS*

An anonymous series of posts, signed by the letter Q on various social media platforms, started a movement known as QAnon, which led some of its most radical supporters to violent and/or illegal actions. Even if some leads are evoked by the press, the authorship of these posts raises to these days many questions. This white paper tries to tackle this issue. For these texts to be attributed, serious challenges have to be addressed. The “Q drops” constitute a genre *per se*, with very unique formal constraints: numerous interrogative forms, telegraphic style with a scarce use of functors etc. They also have possibly been written by different authors. Finally, reference texts by the plausible authors are not readily available, and sometimes hard to find. We thus collected important samples (>12.000 words) for each candidate authors. We then trained Support-Vector Machines to recognize the style of these authors, with 3-grams of characters as input. We conclude that one candidate is likely to have written most of the Q drops, while other authors might have made localized contributions, especially at the beginning of the period. These potential authors are not high-ranked personality from the U.S. administration, but rather social media activists.

Additional Key Words and Phrases: **Authorship attribution; Computational forensics; Conspiracy theories; Domestic threats and terrorism; Machine learning.**

1 INTRODUCTION

The QAnon movement revolves around the posts of one or more individuals signing their message under the name “Q” on social media platform. These messages first appeared in October 2017 on 4chan, and were later on posted on 8chan.

1.1 The QAnon movement: specificity and social impact

The violence of the QAnon believers is not unprecedented nor unmatched on social networks. Discussions around QAnon on Voat even show less toxicity than elsewhere in average on the rest of this admittedly very specific platform [Papasavva et al. 2021]. The singularity of this online group mostly resides in how much its theories have spread, and on its important consequences in real life. Considered a potential domestic terrorist threat by the FBI since 2018 [Winter 2019], the most radical QAnon supporters were implicated in a variety of criminal incidents and violent events [Garry et al. 2021]. Recent research has shown that, beyond political orientation, the main common trait between people arrested for breaking into the United States Capitol in Washington on January 6, 2021 [Kaplan 2021] was that they believed in the QAnon narrative. Documented in the press [Gilbert 2021], the high impact of the QAnon theories on the social life of its believers of their beliefs has been compared to the “conversion to a cult or high-pressure religious group” [Kaplan 2021]. QAnon also affected people far beyond the U.S. borders, spreading in particular through social media and instant messaging software. A large-scale study on QAnon on Telegram for instance showed that messages in German even outnumbered at some point the posts written in English, and that posts in German and Portuguese used an even more “toxic” language [Hoseini et al. 2021].

*Both authors contributed equally to this research.

1.2 Who is Q?

Theories regarding the author behind the posts signed “Q” fit into two main categories. Believers in the authenticity of the source argue that a single source or a collective of people from U.S. intelligence agencies would have authored the various posts. Persons cited as a plausible author or co-author of Q range from General Michael Flynn to Donald J. Trump or his entourage [Huback 2021].

A second group of theory states that one person would have posted as Q, but without having any specific access to exclusive and reliable sources. NBC reporters claimed that they had traced the success of QAnon to three people: ‘Pamphlet Anon’, a.k.a. Coleman R., ‘BaruchtheScribe’, a.k.a. Paul F., and Tracy D. [Zadrozny and Collins 2018]. The second group of theories holds that the original Q posted continuously for a while, but then something changed. This second category includes Brennan’s idea that Ron and John W. are are paying someone to carry on as Q, or are even acting as Q themselves [Huback 2021].

The third group of theories holds that Q is a collective, with a small number of people sharing access to the account. This third category “includes the notion that Q is a new kind of open-source military-intelligence agency” [LaFrance 2020].

Previous research suggested that there could be more than one author to the Q drops. Stylometric analysis, based on factor analysis on character 3-grams, suggested that there were probably two authors who wrote these texts, one after the other [Orphanalytics 2020]. Analyses of the distribution of the number of character and words in each Qdrops, as well as the use of special characters suggested that two hands could have written the Q drops [Aliapoulios et al. 2021], posts written under one of the 10 tripcodes Q used exhibiting different properties from the rest of them. Examinations of the pictures posted by Q however show that they have mostly been posted from the same Time Zone in Asia, and that original pictures were taken from the same camera all along the period [Fox (pseudo) 2021]. This is interpreted as a sign that they were would one unique author to the Q drops .

2 MATERIAL AND METHOD

2.1 Corpus constitution

Because of their contested contents, many accounts have been deleted by the social media platforms. Other accounts were simply deleted by the user themselves. This raised many challenges for the corpus constitution, which had to rely on data collection realized before the deletion, or on difficulty searchable web archives. Moreover, all sources do not have the same time-span, and were created at different dates, discussing potentially different news, a source of heterogeneity in training material that creates challenge in the attribution procedure.

We list here the sources we used for each candidate:

Roger S. We collected Roger S.’s posts on Gab (<https://gab.com/RogerJStoneJr>), from June 14th to August 2nd, 2021.

Michael F. Michael F. wrote a series of 10 articles for the *Western Journal*, from June 29, 2020 to July 31, 2021. (<https://www.westernjournal.com/author/mflynn/>). He also wrote a letter to ask for support to Roger S.’s wife, published on FrankReport (<https://frankreport.com/2021/06/11/banned-by-twitter-gen-michael-flynn-is-published-on-frank-report-concerning-roger-stone/>).

His article for Fox News about ISIS has also been analyzed. (<https://web.archive.org/web/20161215042531/http://www.fox-news.com/opinion/2016/11/02/gen-michael-flynn-after-mosul-is-liberated-isis-could-attack-us-next.html>).

Finally, we collected posts on his Twitter account from October 19th, 2016 to September 18th, 2017.

Paul F. We collected Paul F.'s personal writing on his website (<https://paulfurber.net/>).

We retrieved archives from his twitter on account of a threadreader (<https://threadreaderapp.com/thread/115854052-3008905216.html>).

We also captured archives from the CBTS boards on 8chan, where he wrote as "The Board Owner".

A few posts he left on Discord were transcribed from pictures found online.

Finally, Paul F. wrote a book: *Q: Inside The Greatest Intelligence Drop In History*, included in our large corpus.

Jim W. Archives of Jim W.'s Twitter account were found on archive.today (<https://archive.is/https://twitter.com/xerxes-watkins>), which preserves 9 screenshots (22 Dec 2014, 6 Feb 2016, 15 Mar 2016, 30 Mar 2016, 4 Apr 2016, 5 May 2016, 28 May 2016, 25 Mar 2017, 8 Apr 2017) taken prior to the account suspension. A large amount of these tweets were only citing article titles from his own media *The Goldwater*. As these titles were not necessarily (and probably not) written by him, we chose to exclude them.

A small text was written by Jim W. on 5ch about a service problem on 8ch that he blames on a government attack. (<https://fox.5ch.net/test/read.cgi/poverty/1418027836/826>)

Finally, posts from his Parler account were added to the corpus.

Ron W. A sample of 3130 tweets by Ron W. have been collected through the Twitter API.

We also collected his posts on Telegram from November 30th, 2021 to December, 20th, 2021.

Coleman R. Under the pseudonym *PamphletAnon*, Coleman R. wrote a vast number of posts on 8chan, especially on the board "the Storm" (<https://8ch.net/thestorm/catalog.html>) of which he was the owner. We collected his posts on the Wayback Machine, which seems to provide a complete archive of the board.

A few messages he posted on Discord (Q Central, 2017) have been archived by DDOS (https://ddosecrets.com/wiki/Distributed_Denial_of_Secrets), at this address <https://whispers.ddosecrets.com/discord/user/376607495470448643>.

347 messages by PamphletAnon on Discord (Patriots' Soapbox, 2018) are also available at: <https://discordleaks.unicorn-riot.ninja/discord/user/41936>.

We also found a small text on Reddit, where Coleman R. announces a future talk on Infowars with Rob Dew.

On September 11, 2020, Pamphlet Anon wrote a text on "Patriots' Soapbox", the media he curates with Christina U. (<https://patriotsssoapbox.com/opinion/memories-of-9-11-surreal-and-terrifying/>).

Courtney T. Collaborating with Ron and Jim W., Courtney T. publicly announced that she knew the truth about the Q drops, and that it would be highly disappointing to their public [Huback 2021]. We collected her posts on Twitter, under her account IWillRedPillYou, archived here (http://web.archive.org/web/20180113162029if_/https://twitter.com/I-WillRedPillYou).

Tracy D. Under the pseudonym Tracy Beanz, Tracy D. published a large number of tweets, on an account now suspended, but of which archives.today retains 92 captures (<https://archive.vn/IDRyR>), from December 2016 to January 2021.

She also published a long post explaining herself about quarrels over Q related publications on Steemit (<https://steemit.com/drama/@tracybeanz/she-stood-in-the-storm>)

Christina U. Christina U. wrote a number of articles on PatriotsSoapbox (<https://patriotsssoapbox.com/>), of which the 5 most recent on July 7h, 2021 were collected.

We also found an online conversation she had on Muckrock with Homeland Security on the 7th and 22nd of July 2021. (<https://www.muckrock.com/foi/united-states-of-america-10/patriots-soapbox-department-of-homeland-security-115196/#file-956999>)

Archives of her GAB profile were also analyzed.

Donald T. Former president of the United States tweets were collected during the month of december 2020. We removed tweets suspected to have been written by his staff by the site <https://factba.se/>.

Melania T. We collected a sample of tweets by former FLOTUS thanks to the Twitter API, from January 21st, 2017 to January 19th, 2021.

Eric T. We collected a sample of 3000 tweets by the son of former POTUS thanks to the Twitter API from September 10th 2016, to June 24th 2021.

Dan S. We collected a sample of tweets by the former POTUS' deputy chief of staff thanks to the Twitter API, from October 25th, 2018 to January 20th, 2021.

2.2 Dealing with quotations and copy/paste

Quotes of authors outside of the corpus have been excluded as much as possible by close reading: in particular, quotes from Q, *Wikipedia*, *the Stanford Encyclopedia of Philosophy*, Abraham Lincoln, the Intelligence Resource Program (irp-fas), Steve Scully's biography etc. All these quotes have been removed.

Direct quotations (with or without quotation marks) and copy/paste between the writings of the different candidates can also occur. A good deal of them quote Donald, Eric or Melania T. – Q does it too. There is also a certain number of quotations from Q by the others (such as Paul F. for instance). This could lead to small biases in the constitution of idiolectal profiles. To avoid this, we then proceeded to systematically detect citation between the candidates themselves. Direct pairwise comparison being computationally too costly for a corpus of this size, we used a Locality-Sensitive Hashing (LSH) algorithm. To that end, we used the open source TextReuse package [Mullen 2020]. The corpus was tokenised into sentences, and broken word bi-grams (with skip of 1, that is, allowing for any one word to be inserted between the two words of the bigrams) were counted. For all pairs of sentences, a Jaccard similarity score was computed. Be A and B two samples considered as sets of bi-grams, the Jaccard similarity is computed as:

$$J(A, B) = \frac{A \cap B}{A \cup B}$$

All pairs of sentences with a Jaccard similarity score superior or equal to 0.5 (i.e., at least half of their bi-grams in common) were examined by a human expert, and quotations removed.

Even for $J = 1$, we were sometimes confronted to false positives. Dan S. and Melania T. both use once the sentence “we are all in this together”, without directly citing each other. We thus left this passage in both their texts. Rarely used, the sentence “the American people are not stupid” nevertheless appears in different texts. It was kept in the texts studied, as other simple sentences (“thank you for your service” etc.)

Other situations were trickier to address. For instance, Dan S. uses once the sentence: “the best is yet to come”. It is used five times by Q, himself quoting former President Donald Trump. This sentence could be used by anyone without directly quoting Q or Donald Trump. Yet, as its use by Dan S. starts with “As the President says...”, we considered it a direct quotation and proceeded to deletion from Dan S.'s text. Yet, we did not delete it from Q's own writing, as it is never used as an explicit quotation: the sentence could be used in another context, the person(s) writing the Qdrops

with this sentence could try to impersonate Donald Trump, etc. In any of these cases, it would be legitimate to leave the information. Same thing goes for expression such as “the world is watching” or “make America great again”, used by Donald Trump. but also by Q and some of the potential candidates here.

2.3 Definition of two subcorpus: dealing with generic difference and an imbalanced dataset

The difficulties in data collection for a variety of individual and profiles, as well as the ubiquity of deleted content, not recoverable to us, forces us to adopt a dual approach, and to build two corpora:

large corpus in which we include the larger number of candidates, whatever the number of samples available to us and the genre of said samples.

controlled corpus in which we removed authors for which only too cross-genre and/or too few samples are available, and do not include training material that is too different from the rest (in particular, books).

In the end:

- (1) **large corpus** contains everything described in section 2.1;
- (2) **controlled corpus** is the same, minus
 - interviews transcripts (Michael F., Paul F.);
 - a book by Paul F.;
 - the small amount of available data for Courtney T. and Roger S.

In both cases, due to the limitations in data collection and available material, the quantity of training material is imbalanced between authors, a potential problem in machine learning. To counter this effect, we used class weights during training, where errors for a given class are penalised not always by one, but by a specific weight inversely proportional to class size, where the weight for class i is computed as

$$w_i = \frac{N}{C(n_i + 1)}$$

where N is the total number of samples, and C the total number of unique classes and n_i the number of samples in class i (we used the sklearn ‘balance’ implementation [Pedregosa et al. 2011]).

2.4 The genre of “Q drops”: a methodological challenge

The study of the Q drops raises a number of specific challenges. First of all, the Q drops constitute per se a kind of a *genre*. It follows specific rules that most people would not use in another context: they do not look like a regular blog, media or social media post, nor do they belong to any specific literary genre etc. This forces us to consider our attribution problem as a cross-topic attribution problem.

This specific genre has consequences on many linguistic properties of the Q drops. The structural brevity of the sentences for instance **prevents us from taking the sentence lengths** as a clue of who wrote what. The overwhelming proportion of interrogative forms, especially in the first Q drops, **makes it difficult to reason on morpho-syntactic sequences**, as they are often extremely and artificially stereotyped. Part-of-Speech n-grams such as “Interrogative pronoun - conjugated verb - common noun” would for instance emerge as a signature of the Q drops, and could derail our analysis, by pointing to any of the suspects using interrogative forms the most in other contexts. Finally, the elliptic style of the Q drops, written almost as if they were a telegram, distorts the use of function words, less frequent than expected in the Q drops corpus. **Approaches relying only on function words could be made less robust by this distortion.**

We thus chose to work on character trigrams, the most flexible and reliable feature we could use in the very specific context of this study, and a widely acknowledge feature in stylometry, in particular for its supposed capacity to capture grammatical morphemes [Kestemont 2014; Sapkota et al. 2015] while bearing in mind potential greater sensibility to thematic attractions in comparison to function words.

2.5 Detecting style changes: rolling stylometry

Collaborative writing is not necessarily easy to handle. The scenario in which authors simply took turns, and divided the work between themselves is already complicated to address. But when the collaboration is more complex, especially when the various authors contribute together to the same passages, the style of the original authors can be hard to recognize. The collaboration then results in a new style, that does not match the style of one of the authors [Kestemont et al. 2015].

The principle of rolling stylometry [Eder 2016] is simple: rather than attributing a whole text, we arbitrarily decompose it in a series of overlapping smaller parts: from the 1st word to the 1001th word, from the 2nd word to the 1002nd word etc. Then, we attribute each of these parts to a certain author. We only have to define the length of these parts, and by how much they overlap.

Rolling stylometry has been successfully implemented in a wide variety of settings. With Burrows' delta, it has for instance been used to assess Ford's claims about his implications in collaborations with Joseph Conrad [Rybicki et al. 2014], to determine the beginning of Vostaert's intervention on Dutch Arthurian novel *Roman van Walewein* [van Dalen-Oskam and Van Zundert 2007], or to understand Lovecraft's and Eddy's implication in *The Loved Dead* [Gladwin et al. 2017]. Using support-vector machines, rolling stylometry more recently helped to confirm Fletcher and Shakespeare's collaboration for *Henry VIII* [Plecháč 2020] or Molière and Corneille's collaboration for *Psyché* [Cafiero and Camps 2021].

2.5.1 Support Vector Machine. We choose to train linear Support Vector Classifiers (SVC) to identify the style of each potential candidate. The family of Support Vector Machines algorithm has been widely and successfully used for authorship attribution in a variety of settings and languages, and for very diverse sources, ranging from e-mails or blogs to Shakespeare plays [De Vel et al. 2001; Diederich et al. 2003; Marukat et al. 2014; Mikros 2012; Ouamour and Sayoud 2012; Plecháč 2020]. At the PAN competition, a reference for digital text forensics and stylometry, it also served as a baseline for the "cross-domain authorship attribution" tasks the last time they were proposed in 2018 [Stamatatos et al. 2018] and 2019 [Kestemont et al. 2019]. The Q drops being a sort of "domain" in their own, our own task can be considered as a cross-domain authorship attribution task.

To determine the choice of features and the size of the training samples, we are constrained by two antagonistic goals: the shorter the samples, the more detailed and precise the results that we will get in terms of attribution, yet the longer, the more statistically reliable. Particularly, authorship attribution has proven to require relatively high amounts of data, with a floor for reliable authorship attribution between 1000 and 3000 words, depending on genre and language [Eder 2015, 2017]. The question of sample length can also be linked to the difficulty of the attribution task; cross-domain attribution with multiple candidates presents a challenge in this regard.

On the other hand, the features we retained, character 3-grams, could increase robustness, as they are known to reduce sparsity and perform well in attribution studies [Kestemont 2014; Sapkota et al. 2015]. While punctuation can strongly reflect authorial signature [Sapkota et al. 2015], we removed it because of variety of platforms

from where the data were recovered could cause inconsistencies in the use of signs that can be encoded in different fashions, e.g., apostrophes.

For these reasons, we retain a setup that is a compromise between reliability and finer grain analysis:

Sample length 1000 words;

features character 3-grams (all, including punctuation).

To evaluate our setups, we opt for a leave-one-out cross evaluation on the training corpus (Table 1). The confusion matrix gives more information on the nature of the small number of classification errors (Table 2). As can be expected, performance is slightly lower for authors for which training material is very limited (ColemanR, CourtneyT, RogerS). For the others, it is above 95%, if we except a few confusions between Michael F. and Roger S. (on the large corpus only). These can be explained by the limited size of Roger S. training data, thematic attractions, the fact that he talks about Flynn more or less directly, but also probably by a generational (age) bias.

Table 1. Results of the leave-one-out cross-evaluation for the large corpus (left) and the controlled corpus (right).

	precision	recall	f1-score	support		precision	recall	f1-score	support
ChristinaU	1.00	0.93	0.97	15	ChristinaU	1.00	1.00	1.00	15
ColemanR	0.89	1.00	0.94	8	ColemanR	1.00	0.88	0.93	8
CourtneyT	1.00	0.67	0.80	6					
DanS	1.00	1.00	1.00	9	DanS	1.00	1.00	1.00	9
DonaldT	1.00	1.00	1.00	9	DonaldT	1.00	1.00	1.00	9
EricT	1.00	1.00	1.00	26	EricT	1.00	1.00	1.00	26
JimW	1.00	1.00	1.00	24	JimW	1.00	0.92	0.96	12
MelaniaT	1.00	1.00	1.00	25	MelaniaT	1.00	1.00	1.00	25
MichaelF	1.00	0.84	0.91	19	MichaelF	1.00	0.94	0.97	17
PaulF	0.96	1.00	0.98	74	PaulF	1.00	0.96	0.98	24
RogerS	0.75	1.00	0.86	6					
RonW	0.98	1.00	0.99	43	RonW	0.93	1.00	0.97	43
TracyD	1.00	0.94	0.97	18	TracyD	0.95	1.00	0.97	18
<i>accuracy</i>			0.98	282	<i>accuracy</i>			0.98	206
<i>macro avg</i>	0.97	0.95	0.96	282	<i>macro avg</i>	0.99	0.97	0.98	206
<i>weighted avg</i>	0.98	0.98	0.97	282	<i>weighted avg</i>	0.98	0.98	0.98	206

Table 2. Confusion matrix for the leave-one-out evaluation on the larger corpus

<i>Predicted</i>	ChristinaU	ColemanR	CourtneyT	DanS	DonaldT	EricT	JimW	MelaniaT	MichaelF	PaulF	RogerS	RonW	TracyD
Expected													
ChristinaU	14	0	0	0	0	0	0	0	0	1	0	0	0
ColemanR	0	8	0	0	0	0	0	0	0	0	0	0	0
CourtneyT	0	0	4	0	0	0	0	0	0	1	0	1	0
DanS	0	0	0	9	0	0	0	0	0	0	0	0	0
DonaldT	0	0	0	0	9	0	0	0	0	0	0	0	0
EricT	0	0	0	0	0	26	0	0	0	0	0	0	0
JimW	0	0	0	0	0	0	24	0	0	0	0	0	0
MelaniaT	0	0	0	0	0	0	0	25	0	0	0	0	0
MichaelF	0	1	0	0	0	0	0	0	16	0	2	0	0
PaulF	0	0	0	0	0	0	0	0	0	74	0	0	0
RogerS	0	0	0	0	0	0	0	0	0	0	6	0	0
RonW	0	0	0	0	0	0	0	0	0	0	0	43	0
TracyD	0	0	0	0	0	0	0	0	0	1	0	0	17

We then apply our models to all successive overlapping slices of Q drops, arranged in chronological order, with window size of length 1000 words and step 200 words. We then plot the resulting decision functions for each classifier. The higher the value, the more likely the attribution of a sample to a given author

All analyses are implemented in Python, inside the *SuperStyl* package [Camps et al. 2021], and use internally the SVM and pipeline facilities provided by *Sklearn* [Pedregosa et al. 2011]. Plots are created using R and Python (*matplotlib*).

3 ETHICAL STATEMENT

Even if written independently, this study tried to abide as much as possible to the principles of the “*Pratiquer une recherche intègre et responsable*” guide by the *Centre National de la Recherche Scientifique*’s ethics board [Comité d’éthique du CNRS 2017]. This article does not reveal the identity of individuals that were not broadly known beforehand. In this case, all candidate authors were already either public figures, or individuals whose identity had been stated in major media outlets (NBC, HBO etc.). It only uses information that was in conscience made publicly available by the candidate authors and was accessible through standard internet searches at the time of data collection.

For ethical reasons and to respect the privacy policy of the platforms studied here, we do not freely release any content studied here. To respect the data sharing and open data principles, we however detail our data collection method, which should be sufficient to ensure immediate reproducibility. Some contents could not be available anymore when an attempt at reproducing our computations is performed. In that case, these missing materials could be delivered to research teams on request.

We choose to designate the candidates we study only by their first name and initial, not to impact internet searches on their names.

Finally, this paper does not assert in any way that other persons outside of the persons studied here could not have written the Q drops.

4 RESULTS

Results show that, for most of the slices, the highest decision function is by far by Ron W., even if it does not go above 0 (Fig. 1). The most significant deviation from this concerns the first period of the QDrops, before the switch to 8chan. In this period, the larger corpus analysis gives Paul F. as, by far, the top candidate, before a period where Paul F. and Ron W. signals are competing, until finally Ron W. signals takes over, after a second break that closely matches a tweet described by Paul F. himself as the last authentic Qdrop, that goes

*There will be no further posts on this board under this ID.
This will verify the trip is safeguarded and in our control.
This will verify this board is compromised.
God bless each and every one of you.
Fight, fight, fight!
Q*

The dominance of Paul F. in the first period is not seen at all on the smaller corpus analysis.

More secondarily, there are very localised spikes of Christina U. and Michael F. signals, especially in the more recent period of the QDrops. The rest of the candidates lag far behind.

Results obtained on the two rolling analyses, and their eventual difference, have to be contextualised by investigating the features who received the strongest coefficients in the different SVM classifiers (fig. 2). For some candidates, like Ron W., the features seem mostly idiolectal, like the 3-grams ‘nyb’, ‘ybo’ (in ‘anybody’) or the relative avoidance of ‘_th’ and ‘his’ and remain stable in between both analyses. This is also the case, for instance, for Donald T. whose most distinctive feature is ‘fak’, part of his very idiolectal ‘FAKE’, while other are more content related (‘mpg’ is even due to the regularity with which he mentioned ‘BrianKempGA’ in the training material), a consequence of the choice of characters 3-grams as features.

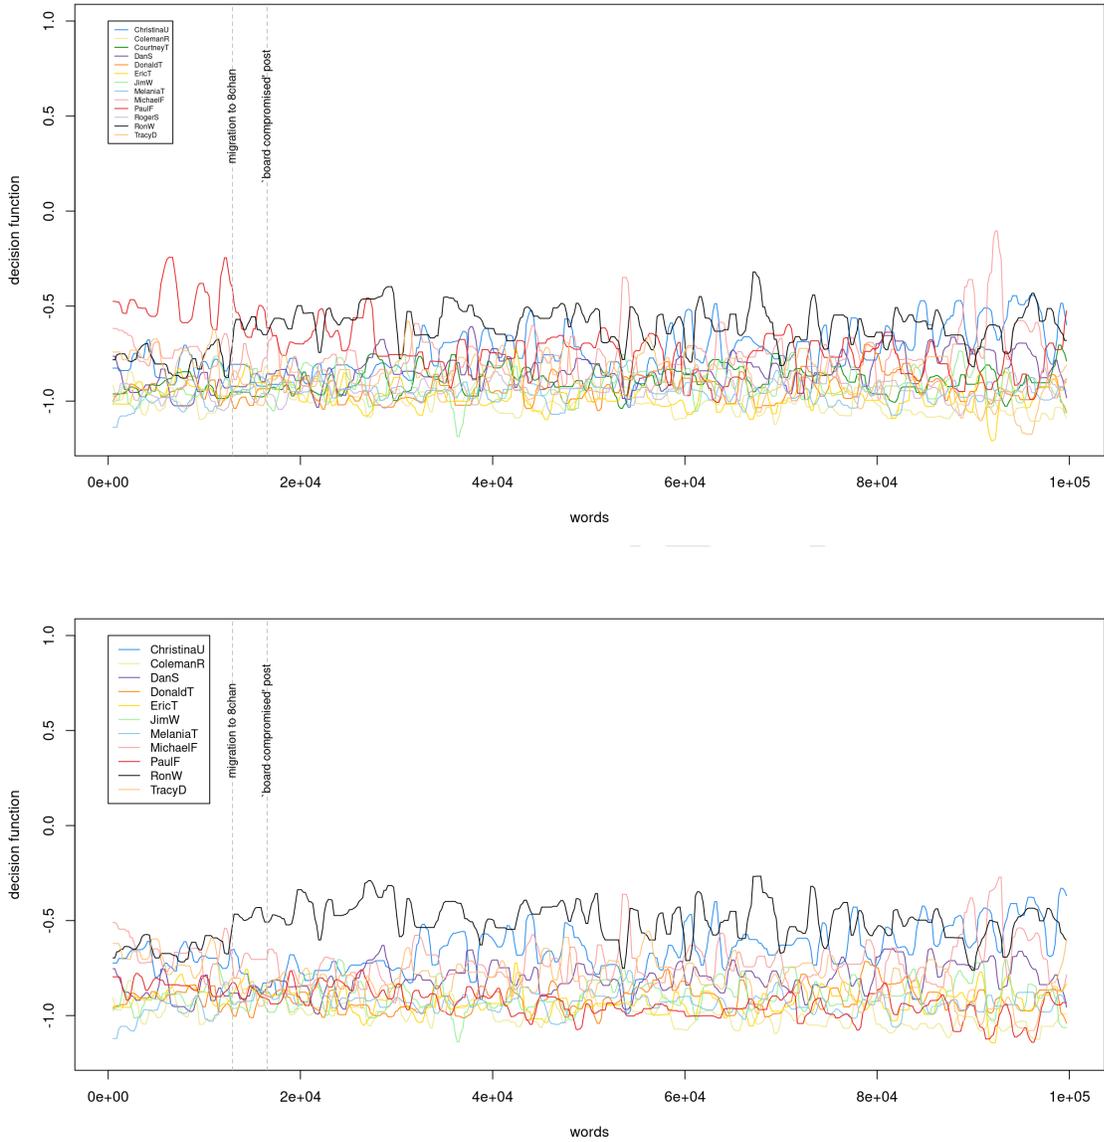


Fig. 1. Decision function of each classifier for each successive overlapping windows of Q drops, arranged in chronological order for the large corpus (top) and the controlled corpus (bottom)

For authors like Christina U., the features are very content and news-related, like the 3-grams extracted from 'Israel(i)', 'blm', 'psy' (psychologists, psychiatrists, ...), etc.

In the case of Michael F., the features seem very dependent on the small quantity of the available training material, and the grandiloquent and religious nature of the few material available, with features such as ‘god’ (‘God’), ‘hty’ (‘almighty’), ‘lib’ (‘liberty’).

Finally and more importantly, these features, in their variation between analyses, give very good insight in the different results concerning Paul F. In the small corpus, due to the exclusion of his book, the most distinctive features for him are all cursory words and racist insults (‘_fu’, ‘fuc’, ‘uck’, ‘shi’, ‘hit’, ‘_ni’, ‘nig’, ‘igg’, ‘gge’, etc.); on the larger corpus, on the other hand, with the book included, they seem revealing of more neutral idiolectal (and grammatical) features, with pronouns, auxiliaries, determiners (‘he_’, ‘had’, ‘was’, ‘the’, etc.). These elements point to the larger corpus analysis being more reliable in what concerns Paul F. (especially in a cross-genre setup) than the smaller corpus analysis.

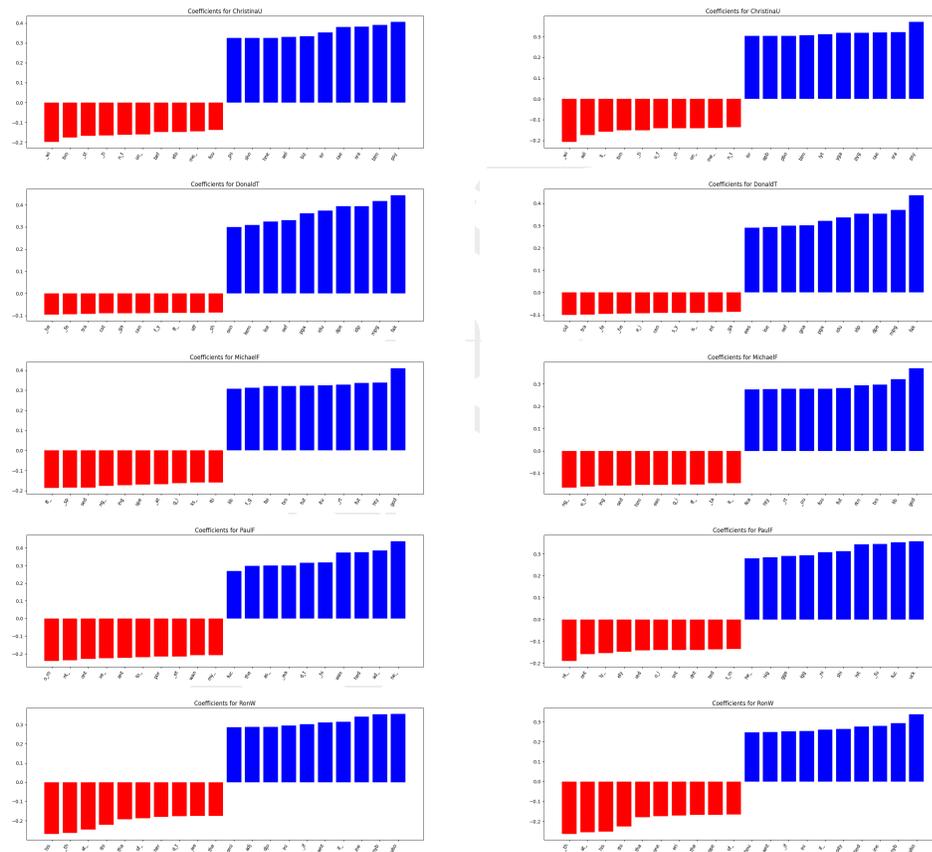


Fig. 2. 10 largest coefficient (negative and positive) for the Liner SVC classifiers of Christina U., Donald T., Michael F., Paul F. and Ron W., trained on the larger (left) and smaller (right) corpora

5 DISCUSSION

If the author of the Qdrops is among our candidates, the results here seem to demonstrate the major role of Ron W. in the writing of the Qdrops, at least since the switch to 8chan. The peak of Paul F. in one of the two analyses, for the

period before 8chan, could very well be revealing of a real participation, even a leading role for all the period before 8chan, with perhaps afterwards a brief period of collaboration (or competition), in between the migration to 8chan and what Paul F. himself describes as the last authentic Qdrop.

Localised peaks of Christina U. or Michael F. on the other hand, while they might very hypothetically be revealing of more occasional collaborations, should probably not be over-interpreted. Given the nature of the coefficients used by the model for them (fig. 2), they seem more likely to be caused by ‘topic similarities’ due to the news and topics dealt with in the Qdrops (keeping in mind that the training material of the different authors have different time scopes). This could result in the choice of a similar lexicon, and even in quotations or paraphrases. To say it otherwise, confusions between Michael F. or Christina U. on one hand, and other candidates (in training) or Q, seem due to attractions in terms of language register and generic peculiarities: samples that use a more elevated and grandiloquent type of patriotic address are brought somewhat closer to Michael F. samples, while those including heavier news-related content might be drawn towards Christina U.

This paper of course has limitations. The very nature of the Q drops, a genre in itself, and the brevity of these texts, make difficult to render a finer grain picture than the one we present here. It is thus plausible, even if we cannot demonstrate it, that other punctual interventions could have occurred. The training data we collected is important, and sufficient to obtain excellent performances. Yet, more training data could of course help increase the precision and reliability of the analyses. A media outlet for instance collected the Facebook posts written by one of our candidate author, for which more support could have been helpful [Zadrozny and Collins 2018]. More generally, other individuals not listed here could of course have participated in the writing.

REFERENCES

- Max Aliapoulos, Antonis Papasavva, Cameron Ballard, Emiliano De Cristofaro, Gianluca Stringhini, Savvas Zannettou, and Jeremy Blackburn. 2021. The gospel according to Q: Understanding the QAnon conspiracy from the perspective of canonical information. *arXiv preprint arXiv:2101.08750* (2021).
- Florian Cafiero and Jean-Baptiste Camps. 2021. ‘Psyché’ as a Rosetta Stone? Assessing Collaborative Authorship in the French 17th Century Theatre. *Proceedings http://ceur-ws.org ISSN 1613* (2021), 0073.
- Jean-Baptiste Camps, Thibault Clérice, and Florian Cafiero. 2021. *Supervised Stylometry: SuperStyl*. École des chartes, Paris. <https://github.com/SupervisedStylometry/SuperStyl/>
- Comité d’éthique du CNRS. 2017. *Pratiquer une recherche intégrée et responsable: Guide*. CNRS, Paris. <https://comite-ethique.cnrs.fr/wp-content/uploads/2019/10/GUIDE-2017-FR.pdf>
- Olivier De Vel, Alison Anderson, Malcolm Corney, and George Mohay. 2001. Multi-topic e-mail authorship attribution forensics. In *Proceedings ACM Conference on Computer Security-Workshop on Data Mining for Security Applications*.
- Joachim Diederich, Jörg Kindermann, Edda Leopold, and Gerhard Paass. 2003. Authorship attribution with support vector machines. *Applied intelligence* 19, 1 (2003), 109–123.
- Maciej Eder. 2015. Does size matter? Authorship attribution, small samples, big problem. *Literary and Linguistic Computing* 30, 2 (June 2015), 167–182. <https://doi.org/10.1093/llc/fqt066>
- Maciej Eder. 2016. Rolling stylometry. *Digital Scholarship in the Humanities* 31, 3 (2016), 457–469.
- Maciej Eder. 2017. Short Samples in Authorship Attribution: A New Approach.. In *DH. ADHO*, Montréal. <https://dh2017.adho.org/abstracts/341/341.pdf>
- Gavin Fox (pseudo). 2021. Where in the World is Q? Clues from Image Metadata. <https://catdogmix.com/2021/05/11/where-in-the-world-is-q-clues-from-image-metadata/>
- Amanda Garry, Samantha Walther, Rukaya Rukaya, and Ayan Mohammed. 2021. QAnon Conspiracy Theory: Examining its Evolution and Mechanisms of Radicalization. *Journal for Deradicalization* 26 (2021), 152–216.
- David Gilbert. 2021. How QAnon is Tearing Families Apart. *VICE* 26 (2021). <https://www.vice.com/en/article/dy8ayx/how-qanon-is-tearing-families-apart>
- Alexander AG Gladwin, Matthew J Lavin, and Daniel M Look. 2017. Stylometry and collaborative authorship: Eddy, Lovecraft, and ‘The Loved Dead’. *Digital Scholarship in the Humanities* 32, 1 (2017), 123–140.
- Mohamad Hoseini, Philippe Melo, Fabricio Benevenuto, Anja Feldmann, and Savvas Zannettou. 2021. On the Globalization of the QAnon Conspiracy Theory Through Telegram. *arXiv preprint arXiv:2105.13020* (2021). <http://arxiv.org/abs/2105.13020>
- Cullen Huback. 2021. Q: Into the storm. 6 episodes HBO Documentary.

- Jeffrey Kaplan. 2021. A Conspiracy of Dunces: Good Americans vs. A Cabal of Satanic Pedophiles? *Terrorism and Political Violence* 33, 5 (2021), 917–921. <https://doi.org/10.1080/09546553.2021.1932342>
- Mike Kestemont. 2014. Function words in authorship attribution. From black magic to theory?. In *Proceedings of the 3rd Workshop on Computational Linguistics for Literature (CLFL)*. Association for Computational Linguistics, Stroudsburg PA, 59–66.
- Mike Kestemont, Sara Moens, and Jeroen Deploige. 2015. Collaborative authorship in the twelfth century: A stylometric study of Hildegard of Bingen and Guibert of Gembloux. *Digital Scholarship in the Humanities* 30, 2 (2015), 199–224.
- Mike Kestemont, Efstathios Stamatatos, Enrique Manjavacas, Walter Daelemans, Martin Potthast, and Benno Stein. 2019. Overview of the Cross-domain Authorship Attribution Task at PAN 2019.. In *CLEF (Working Notes)*.
- Adrienne LaFrance. 2020. The Prophecies of Q. *The Atlantic* (May 2020). <https://www.theatlantic.com/magazine/archive/2020/06/qanon-nothing-can-stop-what-is-coming/610567/> Section: Ideas.
- Rangsipan Marukatat, Robroo Somkiadcharoen, Ratthanan Nalintasnai, and Tappasarn Aramboonpong. 2014. Authorship attribution analysis of thai online messages. In *2014 International Conference on Information Science & Applications (ICISA)*. IEEE, 1–4.
- George K Mikros. 2012. Authorship attribution and gender identification in Greek blogs. *Methods and Applications of Quantitative Linguistics* 21 (2012), 21–32.
- Lincoln Mullen. 2020. *textreuse: Detect Text Reuse and Document Similarity*. <https://CRAN.R-project.org/package=textreuse> R package version 0.1.5.
- Orphanalytics. 2020. *Style analysis by machine learning reveals that two authors likely shared the writing of QAnon’s messages at two different periods in time*. Technical Report. Vevey. <https://www.orphanalytics.com/en/news/whitepaper202012>
- Siham Ouamour and Halim Sayoud. 2012. Authorship attribution of ancient texts written by ten arabic travelers using a smo-svm classifier. In *2012 International Conference on Communications and Information Technology (ICCIT)*. IEEE, 44–47.
- Antonis Papasavva, Jeremy Blackburn, Gianluca Stringhini, Savvas Zannettou, and Emiliano De Cristofaro. 2021. ‘Is it a Qoincidence?’: An Exploratory Study of QAnon on Voat. In *Proceedings of the Web Conference 2021*. Association for Computing Machinery, New York, 460–471. <https://doi.org/10.1145/3442381.3450036>
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- Petr Plecháč. 2020. Relative contributions of Shakespeare and Fletcher in Henry VIII: An analysis based on most frequent words and most frequent rhythmic patterns. *Digital Scholarship in the Humanities* fqaa032 (2020). <https://doi.org/10.1093/llc/fqaa032>
- Jan Rybicki, David Hoover, and Mike Kestemont. 2014. Collaborative authorship: Conrad, Ford and rolling delta. *Literary and Linguistic Computing* 29, 3 (2014), 422–431. Publisher: Oxford University Press.
- Upendra Sapkota, Steven Bethard, Manuel Montes, and Thamar Solorio. 2015. Not all character n-grams are created equal: A study in authorship attribution. In *Proceedings of the 2015 conference of the North American chapter of the Association for Computational Linguistics: Human language technologies*. Association for Computational Linguistics, Stroudsburg PA, 93–102.
- Efstathios Stamatatos, Francisco Rangel, Michael Tschuggnall, Benno Stein, Mike Kestemont, Paolo Rosso, and Martin Potthast. 2018. Overview of PAN 2018. In *International conference of the cross-language evaluation forum for european languages*. Springer, 267–285.
- Karina van Dalen-Oskam and Joris Van Zundert. 2007. Delta for middle Dutch author and copyist distinction in Walewein. *Literary and Linguistic Computing* 22, 3 (2007), 345–362.
- Jana Winter. 2019. FBI document warns conspiracy theories are a new domestic terrorism threat. *Yahoo News* 1 (2019). <https://news.yahoo.com/fbi-documents-conspiracy-theories-terrorism-160000507.html>
- Brandy Zadrozny and Ben Collins. 2018. Who is behind the Qanon conspiracy? We’ve traced it to three people. *NBC News* (Aug. 2018). <https://www.nbcnews.com/tech/tech-news/how-three-conspiracy-theorists-took-q-sparked-qanon-n900531>