

## Impostors tending towards the wild Purposes of authorship analysis and a specific impostors method in the Karvounakis terrorism case

Isobelle Clarke <sup>1</sup> and Tim Grant <sup>2</sup>

<sup>1</sup>Linguistics and English Language, Lancaster University, i.clarke@lancaster.ac.uk

<sup>2</sup>Aston Institute for Forensic Linguistics, t.d.grant@aston.ac.uk

### Abstract

*Authorship analysis can be used at different stages of a criminal case, including the initial investigation (i.e. to narrow down the search for a suspect), to support a search warrant application, and to provide evidence for a criminal prosecution. Each stage has its own complexities and pressures, requiring some flexibility in the type of method applied. Indeed, it is not unusual in such cases for the conditions to be far from ideal, adding to the complexities of the particular stage of the case. For instance, cases often involve texts that are very short and comparison data that is not within the same genre as the questioned text. In some cases, there are strict security constraints, meaning that some tools and software are prohibited and that the forensic linguist must work alone. HMA v Nikolaos Karvounakis is a domestic terrorism case, which required authorship analyses at all three stages. Each stage had a particular purpose and specific security constraints and complexities, which influenced the kinds of authorship analyses employed. The final stage led to the development of the Specific Impostors method, which is variation of the General Impostors method. This article describes the case history, the methods employed and the outcome of their application to the data and concludes with reflections on the case and the methodology.*

**Keywords:** *Specific Impostor's Method, n-grams, comparative authorship analysis, domestic terrorism.*

### Resumo

*A análise de autoria pode ser utilizada em diferentes fases de um processo penal, incluindo a investigação inicial (ou seja, para restringir a procura de um suspeito), para apoiar um pedido de mandado de busca e para fornecer provas para uma ação penal.*

*Cada fase tem as suas próprias complexidades e pressões, exigindo alguma flexibilidade no tipo de método aplicado. De facto, não é raro que, nestes casos, as condições estejam longe de ser ideais, o que aumenta a complexidade da fase específica do processo. Por exemplo, trata-se frequentemente de textos muito curtos e de dados de comparação que não pertencem ao mesmo género do texto questionado. Nalguns casos, existem restrições rigorosas em matéria de segurança, o que significa que algumas ferramentas e programas informáticos são proibidos e que o linguista forense tem de trabalhar sozinho. HMA v Nikolaos Karvounakis é um caso de terrorismo doméstico, que exigiu análises de autoria nas três fases. Cada fase tinha um objetivo particular e restrições e complexidades de segurança específicas, que influenciaram os tipos de análises de autoria utilizados. A fase final levou ao desenvolvimento do método dos Impostores Específicos, que é uma variação do método dos Impostores Gerais. Este artigo descreve a história do caso, os métodos empregues e o resultado da sua aplicação aos dados e conclui com reflexões sobre o caso e a metodologia.*

**Palavras-chave:** *Método do Impostor Específico, n-gramas, análise comparativa de autoria, terrorismo doméstico.*

## 1. Introduction: Case history

A little before 8pm in the evening of 11th January 2018, a suspicious cardboard box was discovered by a park ranger on a sheltered seat in Princess Street Gardens, a public park in Central Edinburgh in Scotland. A controlled explosion was carried out and initially it was believed that the device was relatively harmless, but later inspection by forensic experts showed that it was a viable bomb and had the capacity to cause serious harm.

There were no clear leads to the identity or the motives of the person who had planted the bomb, until a month later when a journalist received an anonymous email which linked to Maldición Eco-Extremista, a blog from the Mexican eco-anarchist group called *Individualistas Tendiendo a lo Salvaje* (hereafter the ITS). In English their name translates to *Individualists Tending to the Wild*, and their central belief is that the devastation of world ecosystems was linked to scientific and technological progress to an extent that justifies direct and sometimes terrorist actions. On 6th February 2018 the blog published two new posts, both of which contained the same content but one was in English and the other in Spanish. Specifically, the posts contained a clear claim to responsibility for the bomb and a picture of the interior of the bomb (Figure 1).

The blog posts were sent to TG with a generalised request for commentary and a particular focus on whether the author could be identified from the writings. Initial analysis was constrained – as this analysis was part of a live terrorism investigation the text and the focus of the inquiry on the text was classified information – this meant that TG had to work alone and was largely prevented from using online analysis tools. However, through a manual analysis based on back translation of short phrases and idioms using Google translate, TG established that the English language text seemed to be the original and the Spanish language text was based on the English version. Further to this TG concluded that it was unlikely that the English text was written by a native



band had a website which contained lyrics and other texts supposedly written by Karvounakis. TG was asked to carry out a comparative authorship analysis to give evidence to ask a court to grant a warrant to search Karvounakis' properties, and to seize his computers and other devices. TG was able to carry out an analysis which showed some limited evidence that linked Karvounakis' writings from the *Face of Another* website with the Maldición Eco-Extremista blogs. This along with other evidence from Police Scotland became grounds for a successful application for a search warrant.

Once the warrant had been served this produced a whole lot more text from Karvounakis' computers, including blog posts known to be from the ITS website including the *Paroxysm of Chaos* 'zines' described below ('zine' is an abbreviated term for 'electronic magazine'). Unfortunately, however, the seized materials did not include a copy of, nor was there a link to the original claim of responsibility for planting the bomb from Karvounakis' devices. This meant that the investigation required one final forensic authorship analysis to be carried out. This time the comparison would be between, on the one hand, the set of blog posts seized from Karvounakis' computers, which the police believed they could reasonably show that he had written, and on the other hand, the original 'claim of responsibility' text. The purpose would be to determine whether or not these texts had a common author and therefore provide evidence that could contribute to the criminal case against Karvounakis.

By this time TG had been working with the police team of investigators for more than three years. He was clearly part of the investigation (albeit on the edges of it) and psychologically was not acting as a wholly independent expert. To avoid bias he contacted a second forensic linguist (IC) and provided her with the documents in the case with no story about what had happened, and he did so in a predetermined order following protocols set out in Grant (2022). In what follows we present IC's authorship analysis in which she identifies consistent and distinctive features of Karvounakis' known language. She also shows that these features occur in the 'claim of responsibility' text and this provides some evidence that Karvounakis is a likely author of that claim. Reinforcing the importance of these features in the analysis, IC goes on to show that they hardly occur in the other ITS writers' blogs, and nor do they occur in a comparison with other Greek speakers writing in English or in more general corpora of English.

Through this case history, we can see that the authorship analyses carried out were commissioned with different purposes and constraints and needed to be responsive to these. The next section of this article examines in more detail the variety of purposes across different types of forensic authorship analysis and considers how the design of the analysis can and should (or sometimes should not) bend to these purposes. Following this we discuss the variety of methods applied at each stage and how they serve the specifics of each of these different purposes and of the specifics of the case. Finally, in our conclusion, we reflect on how the required flexibility of method in the light of these specifics can be accommodated even where we also prioritise the push to validation of methods and acting ethically as forensic linguistic experts in investigative and evidential work.

## 2. Varied Purposes of Authorship Analysis

This case involved progressive authorship analyses, first with an investigative purpose, next with the purpose of providing evidence to support an application for a search warrant, and finally with the purpose of providing evidence which would form part of a criminal prosecution. In each of these three stages, the function of the analysis affected its design and outcome.

### 2.1. Stage 1 – Authorship analysis for investigation

Investigative forensic linguistics can involve analysis where the intended purpose is not to directly provide evidence for a court, but rather to help identify a suspect (see Grant & MacLeod, 2020, Chapter 6 for a discussion). Such an analysis is sometimes referred to as simply *investigative analysis* or *intelligence analysis* (as contrasted with evidential analysis). Depending on the Criminal Practice Directions (or equivalent) that apply in different jurisdictions, such analyses may have to be disclosed to defence teams (either as a duty of the prosecution as occurs for UK jurisdictions, or as a response to a subpoena in United States jurisdictions), or a decision might be made by prosecution teams, that an investigative analysis should actually be used as evidence in court. Thus, at the point the expert is approached to work on a case the final value or nature of their contribution will be unknown even when the primary purpose of commissioning the analysis is investigative. This context, alongside the knowledge that disclosure may be required in the future, can constrain the techniques used and the nature of the conclusions that can be expressed.

As described above, the ITS investigation was a live terrorism inquiry and as such came with security constraints. Electronic files of the posts and of the email to the journalist were sent to TG by police courier on encrypted drives, and the investigation was ‘indoctrinated’ – a term used by UK policing to indicate that materials and any progress in the investigation were classified. TG was required to not seek assistance in the investigation from any colleagues, and to not use online tools to analyse the texts. He also does not read or speak Spanish. These constraints clearly hampered the nature of the analysis that could be undertaken. Colleagues at the Aston Institute of Forensic Linguistics not only include Spanish speakers (including a Mexican Spanish speaker), but also these colleagues carry out research projects directed at the detection of different dialects of Spanish influence on English language texts (e.g. Batel, Abrams, & Pezik, 2022). TG was able to review the (then unpublished) papers associated with these projects and manually examine the English language texts for artefacts of Spanish influence. Although there were clear indications in the syntax and word choices that the writer was a non-native writer of English, no specifically Spanish interlingual features were found. He was unable to determine what other language(s) might be influencing the written English. These conclusions, fed back to the investigators by way of a brief report, indicated the uncertainty associated with the conclusions. This report was included in the ‘unused materials’ disclosed by the prosecution and could therefore have been used by a defence team at trial. The risk to the prosecution would have occurred if the defendant had turned out to be a Spanish speaker, in which case the linguistics report could have been of value to the defence.

The limited conclusions of the report were however presumably of little value to the ongoing investigation in identifying a suspect as nothing occurred for a further two years.

## 2.2. Stage 2 – Authorship analysis to a standard of reasonable suspicion

The purpose of evidential analysis is not always to provide evidence beyond a reasonable doubt to a criminal court. Authorship analysis has been used in a variety of civil contexts, in employment and immigration hearings and in other quasi-judicial contexts. One such context is providing evidence to be taken before a judge in an application for a search warrant. In Scotland (as in England and Wales), the legal standard that must be reached to grant a warrant is “Reasonable Suspicion”. This more or less equates to the United States standard of “Probable Cause” and is clearly a long way less certain than criminal standards of proof such as “Being Sure” or “Beyond a Reasonable Doubt”. This lower burden of proof may raise questions for the analyst in terms of method. There is a difference between, on the one hand, carrying out an analysis with a well-attested method, and the results from the analysis reaching a well-understood lower level of certainty, and on the other hand, carrying out an analysis which is less well-validated, or which has unavoidable weaknesses such that there is less confidence in the outcome. Consideration of the possible problems with the nature of the analysis were an issue with this stage of the case. In comparative authorship analyses there is now increasing evidence that cross-genre authorship analysis is difficult to do well. In recent computational authorship analysis competitions, such as PAN (<https://pan.webis.de>), the best computational systems that can typically score over 93% accuracy in within-genre verification tasks, drop to barely better than chance in cross-genre attribution (Stamatatos et al., 2022).

In carrying out the analysis to provide evidence to support a warrant application, the same security constraints were applied as in the Stage 1 analysis. This is because, at this stage of the case, the suspect was unaware of the investigation, and because the line of investigation could have led to the elimination of the particular individual from the inquiry. The analysis at this stage was clearly a cross-genre analysis. On the one hand there were the original ideological posts (and the claim of responsibility) from the Maldición Eco-Extremista blog, and on the other hand there were song lyrics from the *Face of Another* website and other brief texts describing the band’s music and ideologies. TG used an analysis based on a word n-gram approach. N-gram measures are frequently used in authorship studies and simply take each sequence of units (characters or words) that arise from a text. So, for instance, the phrase ‘The cat sat on the mat’ has the following 3-grams: “the\_cat\_sat”, “cat\_sat\_on”, “sat\_on\_the”, “on\_the\_mat”. N-grams can also be character sequences (e.g. “t\_h\_e, h\_e\_c, e\_c\_a, etc), but for this analysis, and due to the need for interpretability in court, we only considered word n-grams. Research into n-grams generally suggests that n-grams can be found that are characteristic of specific authors (whether or not these carry interpretable linguistic information). For example, Grieve et al. (2019) point to character n-grams as being the most useful features for discriminating authors in their problem, and yet only provide word n-grams for reasons of interpretability. Importantly, it has been repeatedly shown that collections of n-grams can be used to discriminate between authors (e.g. Grieve, 2007; Grieve et al., 2019; Stamatatos, 2009). Thus, one author may fall into a habitual use of

certain n-grams, whereas a second author never uses those particular strings of words. The purpose of such an authorship analysis is to determine from a text whether strings of n-grams can be derived that demonstrate a consistent within-author pattern of use, which is also distinctive in comparison to other potential authors of the text.

The analysis applied at this stage was computer assisted – TG used a simple R script (R Core Team, 2023) to list the words, bi-grams (two-word strings) and tri-grams (three-word strings) which were present in each set of texts. This was essentially a computer-assisted stylistic analysis which showed some limited evidence that some rare words, rare two-word strings, and rare three-word strings occurred in both the texts known to have been authored by Karvounakis, from the *Face of Another* website and the anonymous texts from the ITS website. TG was also able to show some distinctiveness for these n-grams in that they were either unique or very low frequency in the rest of the texts contained in the blog posts.

In his report, TG made clear that this was a cross-genre analysis and the difficulties with that, and the limitations of the method meant that levels of certainty were low, but also that the finding of individual rare words, bi-grams and tri-grams across the different set of texts did, in his opinion, carry some evidential weight. This evidence, along with other evidence that Police Scotland put before the court<sup>1</sup>, gave rise to the judgement that there was sufficient weight of evidence to amount to reasonable suspicion and the warrant was granted.

### **2.3. Stage 3 – Authorship analysis to a criminal standard**

The granting of the warrant literally opened the door to a wealth of further evidence. Karvounakis' premises were searched, and his devices seized and searched. These searches revealed a number of documents that had appeared on the *Maldición Eco-Extremista* blog and provided strong evidence that Karvounakis wrote under the pen name Archegonos. The newly available data and the full analysis is described below. Two principal issues arose as the analysis switched to providing evidence for criminal prosecution.

The first change was, of course, in the standard of proof that the new analysis would have to answer to. Whilst this does not apply to any individual item of evidence or fact in the case – the objective of a prosecution in Scotland is to convince a jury “looking at the evidence as a whole” that they are “satisfied of the guilt of the accused beyond reasonable doubt” (Judicial Institute for Scotland, 2023), it does require a high degree of certainty about the validity of the analysis and its conclusions. How this was achieved is described below.

Further to the increase in the burden of proof, there also arose the issue of potential bias in the analysis and how to mitigate this. As noted above TG had at this stage of the investigation been working on the case for over three years and considered himself more part of the investigation than an independent expert. As in England and Wales, the duties of an expert witness in Scotland are clear, they must understand that their “principal duty is to the court, not to the prosecution or defence. Evidence must come

---

<sup>1</sup>This comprised evidence from a forensic meteorologist who opined that a cloud formation in a picture from the *Face of Another* website matched a picture of a cloud posted on the *Maldición Eco-Extremista* blog site.

from the independent thinking of the expert witness. It must not be influenced by those who instructed them. Expert evidence is objective and unbiased” (Crown Office and Procurator Fiscal, 2024). Recognising that his ‘independent thinking’ might well be compromised and cognisant of the research base on cognitive biases, and specifically confirmation bias (e.g. Edmond et al., 2017; Found, 2015; Stoel, Berger, Kerkhoff, Mattijssen, & Dror, 2015), TG decided to recruit an independent forensic linguist (IC) to produce the evidential analysis which might be used at trial. By this time, Karvounakis had been interviewed and although the police were still concerned with investigative security, permission was obtained to proceed.

TG provided IC with the documents in the case with no story about what had happened, and (following Grant, 2022; Grant & Grieve, 2022 he did so in a predetermined order. First, he provided only the ITS blog posts that Karvounakis was now known to have written. On the basis of these texts, he asked IC to find language features which were common between texts that he had written and that were distinctive from other writers on the ITS blog site. Like the Stage 2 analysis, IC decided to focus on word n-grams in order to systematically account for a larger pool of features. He then informed IC that the writer’s English might be influenced by Greek, Karvounakis’ native tongue, and encouraged IC to check whether the identified distinctive language features of Karvounakis could be attributed to their language background. Finally, and only when this analysis was complete did TG pass to IC the claim of responsibility blog post. These measures were an attempt to create the conditions for as rigorous analysis as possible and one in which bias was mitigated wherever possible.

Throughout the case history, we thus see the tasks, and the pressures associated with the tasks, moving and changing: in Stage 1, the questions were wide open, but the methods maximally constrained by the security concerns. In Stage 2, those security issues were still a concern, and the authorship analysis was challenged by the small quantity of short texts involved and the cross-genre nature of the comparison. This was possibly mitigated by the lower burden that was at issue for the judge in granting the warrant. For Stage 3, the security issues were considerably mitigated and there was a wealth of within-genre material available to create good comparisons. The challenge was, however, to mitigate bias and provide the strongest possible evidence for criminal prosecution. The method applied to provide this evidence is a variation of the General Impostors method. Thus, before we present our method, we describe the General Impostors approach to provide some context.

### **3. The (General) Impostors Method**

The Impostors Method (IM) was originally proposed by Koppel and Winter (2014) to determine whether two texts Q and K are written by the same author, where the Q document is a document of questioned authorship and the K document is a document of known authorship. The IM assesses whether Q is more similar to K than to each one in a set of impostor texts. The comparison is based on a large number of features which are sampled randomly to create 100 separate feature sets. The overall similarity is then expressed as the percentage of those feature sets where Q is more similar to K than to any of the impostors.

Koppel and Winter (2014) suggest that the selection of impostor texts and how many impostors is key in this process. They argue that too few impostors or unconvincing impostors could lead to too many false positives (in Koppel and Winter's terminology<sup>2</sup>), while too many impostors or impostors that are in the same genre as the disputed text could lead to many false negatives. They thus proposed seeking "an optimal combination of impostor quality, impostor quantity, and score threshold" (Koppel & Winter, 2014, p. 182). They suggest generating three groups of impostors for the disputed text:

1. Fixed set – a fixed set of impostor texts that have no relation to the disputed text or the suspected author.
2. On-the-fly – a set of plausible impostors, which are generated by selecting small sets of random medium-frequency words from the known author's text and the disputed text and incorporating them into a search engine query. The top results of each query are aggregated.
3. Shared genre – a set of impostors that are in the same genre as both the disputed and the known author text.

The use of the impostor method is akin to a police lineup where a witness is asked whether a previously seen suspect is present. In visual parades (Darling, Valentine, & Memon, 2008) or indeed in voice lineups (Pautz et al., 2023) the selection of impostors (or 'foils') has been shown to be a crucial aspect of design. In language analysis this often overlooked by advocates of IM. For example, a face lineup where the previously seen suspect was described as having darker skin, and then at the identity parade only a single contender in the lineup was dark-skinned, would clearly create unfair identification. In the linguistic case therefore, the impostors in the IM need to be selected or constructed to be close enough to the Q text and the K text to create a fair competition for closest match. For Koppel and Winter's (2014) three groups of impostors this does not seem to be well accounted for. In the forensic context, it is important that the task is not too easy, as this might be challenged by an opposing side (e.g. in cross-examination). However, the conceptualisation of closeness or distance in style between texts is under-explored, in terms of the full range of linguistic variation and how this is created (see e.g. Grant & MacLeod, 2020). We address this through developing our specific impostors method (below).

With regard to features used in the IM, previous studies have used *inter alia*: the most frequent words (MFW), word and/or character n-grams, function words and part-of-speech tags (e.g. Andrijanić & Bąkowski, 2021; Khonji & Iraqi, 2014). The technique then randomly picks a percentage of the full feature set and then calculates whether Q, based on this set of features, is closer to the suspect's known writings or to the individual impostors using a distance measure. This is repeated for 100 different feature sets.

The IM has been hugely successful with variations of the technique winning first places in PAN 2013 and 2014 (Potha & Stamatatos, 2017). One of the winners was Seidman (2013) who sought to generalise the IM to deal with cases where the task is to assess whether a document is an outlier in a set of known documents (i.e. where there are multiple documents by the candidate author). They tested various methods and

<sup>2</sup>Koppel and Winter (2014) are confusing in their use of the terminology of false negatives and false positives in this context. Too similar impostors could lead to false positives, whilst too dissimilar impostors offers an easy comparison meaning that the known texts of a suspect will be deemed most similar, which could overemphasise the similarity between the Q and K.

found that the best performing approach involved running the IM on all pairs consisting of the questioned document and a single known document and then aggregating the results.

Attempting to address the issue raised above, Potha and Stamatatos (2017) sought to improve Seidman's (2013) General Impostors Method (GIM) by incorporating impostors that were close to the suspect. They acknowledge that in the traditional police lineup, individuals are not selected randomly, but rather they are selected to have characteristics similar to their suspect. They achieve this by submitting queries into Bing search engine using salient words from the set of known documents and download the first results. In addition, they noted how the GIM disregarded cases where at least one impostor text is found more similar to the disputed text in comparison to the candidate author. As a result, they propose ranking the similarities in decreasing order and considering the position of the candidate author's texts. These modifications were found to be competitive and sometimes better than other GIM variations.

As noted, we used a variation of the GIM for Stage 3. We call this the Specific Impostors Method, which we present below.

#### 4. Developing a Specific Impostors Method

This section describes the analysis conducted in Stage 3. This analysis was guided by the twin competing hypotheses:

- (H1) Karvounakis is the author of the disputed text.
- (H2) Someone other than Karvounakis is the author of the disputed text.

Initially, IC was provided with two zines that were posted to the ITS blog: Paroxysm of Chaos I (PoC1) and Paroxysm of Chaos II (PoC2). These zines were comprised of texts written by Karvounakis (under the pen name Archegonos) as well as writings from other individuals. The texts can be characterised as anarchist texts often expressing anti-civilisation attitudes. IC was tasked with identifying a consistent within-author pattern across the documents written by the suspect, and to identify if this within-author pattern is distinctive from the writings of the other individuals. Then, IC was provided with the disputed text for comparison.

##### 4.1. Data Description

###### 4.1.1. Comparison texts

###### 4.1.2. Paroxysm of Chaos I

PoC1 is a zine comprised of 17 sections (see Table 1). The zine was compiled under the pen name Archegonos and includes sections written by Archegonos as well as writings from other individuals (some of whom are known to be alive and some of whom are known to be dead). The first chapter is written by 'Archegonos' – our suspect's pseudonym - and is called 'Prelude to the abyss', which serves as the introduction to the publication. In a side note after the introduction it says: "[e]verything included in this zine is something that I see either as theoretical or practical gain and definitely not representative of mine". Files were created based on the presence of an authorial sign-

off. The majority of sections have a sign-off, except for sections 16 and 17. Section 17 is comprised of multiple poems from two different authors and is signed off after each poem within the section (3 poems were from Archegonos and 1 was from Bruno Filippi).

Section 16, however, is unique in its composition of numerous claims of responsibility for illegal acts, yet the sign-offs are often missing or unclear. For instance, in one part of section 16, it describes an incident where ten trucks owned by a prison officer were set on fire in France, but this is not claimed under any specific pseudonym. For some of these subsections, the content makes it seem as if they were written by the main editor of the zine, Archegonos, but this could not be depended upon for the analysis. Because of the missing sign-offs these were grouped into one file for this section and the authors were deemed to be 'unknown' (see Table 1).

One section is signed off as a "Collaboration of anarchists", whilst another section is signed off by 'Fallon and Amelie'. These explicit acknowledgements that some sections are multi-authored, gives at least some weight to the assumption that other signed sections are written by individual authors.

**Table 1. The Paroxysm of Chaos I**

	Title of section	No. word tokens	Signed Author(s)	Filename
1	Prelude to the abyss.	1056 words	Archegonos	Archegonos_1
2	Total liberation as an egoist and iconoclastic consideration.	8364 words	Archegonos	Archegonos_2
3	On radical moralism and wildness.	1309 words	Anonymous	Anonymous_1
4	Against the language of militancy.	400 words	Wolfi	Wolfi-
5	Concerning "good" and "evil"	578 words	Landstreicher	Landstreicher_1
6	Sacrilegious laughter.	962 words	Émile Armand	Emile-Armand_1
7	Into the abyss-chaos.	240 words	Erinne Vivani	Erinne-Vivani_1
8	Veganism from a nihilist and anti-civilization perspective.	1244 words	L	L_1
9	Veganism from a nihilist and anti-civilization perspective.	1244 words	Archegonos	Archegonos_3
9	A critique, not a program: For a non primitivist anti-civilization critique.	2359 words	Wolfi	Wolfi-
10	Mexico : Open letter of Amelie Pelletier and Fallon Poisson.	482 words	Landstreicher	Landstreicher_2
10	Mexico : Open letter of Amelie Pelletier and Fallon Poisson.	482 words	Fallon and Amelie	Fallon-Amelie_1

11	Mexico : Anarchist comrade Carlos Lopez 'Chivo's reply to International Revolutionary Left "Buenaventura Durruti".	897 words	Carlos Lopez "Chivo"	Carlos-Lopez_1
12	De Profundis Clamavi	938 words	RadioAzione [Croatia]	RadioAzione_1
13	War on the state: The subject of desire.	1746 words	Saul Newman	Saul-Newman_1
14	Tearing down the prison of civilization.	2773 words	Collaboration of anarchists – Memento Mori (Translated from Greek to English by Archegonos)	C-o-A_1
15	Some thoughts on the matter of responsibility claims.	299 words	Archegonos	Archegonos_4
16	Responsibility claims – Moments of war.	6353 words	Unknown (multiple authors)	Unknown_1
17	Individualist poetry.	801 words	Archegonos	Archegonos_5
		199 words	Archegonos	Archegonos_6
		255 words	Archegonos	Archegonos_7
		580 words	Bruno Filippi	Bruno-Filippi_1

#### 4.1.3. Paroxysm of Chaos II

PoC2 is the sequel to PoC1 and is made up of 13 sections, which were similarly compiled by Archegonos with sections written by Archegonos or other authors (see Table 2). Like PoC1, texts were created based on an authorial sign-off (as opposed to sections). Most texts had a clear sign-off. Exceptions to this were sections 8 and 10. Similar to PoC1, there was a section compiling numerous claims of responsibility (section 10). These were grouped together in a file with the author(s) specified as unknown (although we suspected parts of this to be written by Archegonos). Section 8 did not have a sign-off. After section 8, there was a note from the publisher (possibly Archegonos), specifying that this text was published "because it makes some points I agree with". We can therefore assume that it was likely not written by Archegonos. Section 12 is the introduction of a group/blog/project called 'Antisocial Evolution'. No other texts or files can be found by this group to connect it to a specific author. Although it seemed that this text in section 12 was likely written by Archegonos, because it was not specifically signed off by this pen name, we did not include it in the known writing samples of Archegonos.

Table 2. The Paroxysm of Chaos II

	Title of section	No. of word tokens	Signed Author(s)	Filename
1	Prologue – Desecration of universality and the theatrical significance of civilisation’s performances	2720 words	Archegonos	Archegonos_8 <sup>3</sup>
2	Chaotic iconoclasm and incineration of idols (Anticivilisation of paroxysm part 1)	1746 words	Archegonos	
3	The coronation of moralism upon the throne of the ghost of Nature (Anticivilisation paroxysm part 2)	6961 words	Archegonos	
4	Breaking the windows that sell the products of civilisation (Anticivilisation paroxysm part 3)	675 words	Archegonos	
5	My misanthropism (Anticivilisation paroxysm part 4)	1507 words	Archegonos	
6	Wildness in the city	2265 words	Chris Kortright	Chris-Kortright_1
7	A Life	1616 words	Renzo Novatore	Renzo-Novatore_1
8	Insurrection or revolution?	1792 words	Unknown	Unknown_2
9	The unique one meets the overhuman II	1008 words	John F Welsh	John-F-Welsh_1
10	Moments of war	3668 words	Unknown (multiple authors)	Unknown_3
11	Poetry of the void	590 words	Archegonos	Archegonos_9
12	Introduction to the project Antisocial Evolution	2782 words	Antisocial Evolution	Unknown_4
13	On being lyrical	1072 words	E. M. Cioran	E-M-Cioran_1

<sup>3</sup> The Archegonos sign-off appeared only at the end of part 4 and not for each section. As a result, these were combined into one file.

Because the known writing samples of Archegonos were in English, the English version of the disputed text was used for comparison. Archegonos’ known writing samples amounted to 25,173 words across nine sections of the two Paroxysm of Chaos texts (see Table 2). The only other author to have more than one text was Wolfi Landstreicher, with a total of two texts comprising 2759 words. All other authors only had one text with text lengths ranging from 240 to 6353 words.

#### 4.1.4. The Disputed Text – the claim of responsibility

The claim of responsibility was posted to the ITS blog on 6th February 2018 in English and in Spanish. The English version is 1,864-word tokens. Each version of the text in-

cludes the same 5 images, including an image of the bomb. The text claims responsibility not only for the bomb in Princess Street Gardens, but also the arson of two buses. The text includes long stretches of evil laughter, as can be seen in Figure 1 – reflective of the laughter of a fictional villain. Even though the text was posted to the ITS blog, there is also an overt attribution of the attack to ITS, with the author claiming themselves to be part of ITS. The Spanish version was a close translation with some evidence of automated translation of English idioms into Spanish, which may have been subjected to correction by a Spanish speaker, but it was not analysed further as part of the case.

#### 4.1.5. Methodology

To investigate H1 and H2 above, IC firstly needed to uncover consistent within-author language features for the suspect's known writing samples, and secondly, needed to assess whether this pattern is distinctive. As described above, IC elected to focus on word n-grams due to validated use in prior research and analyses and because computational n-gram analyses provide a way to take into account all structural features. When considering the criterion of distinctiveness, it is important to take into account the potential authors' linguistic backgrounds as well as the genre and topics of the texts under analysis. A set of n-grams that may be thought to pick out an individual could in fact be common within a certain online community, or common within a certain genre or subject matter. In this case, the suspect is known to be a Greek national, so we needed to check potentially discriminating n-grams against the writings in English of other Greek nationals. Furthermore, as noted previously, based on the compiler Archegonos' introduction, where he says that he selected the PoC1 and PoC2 documents for their "theoretical or practical gain" and/or because they "make[s] some points that I agree with", it was clear that the suspect's writing in the *Paroxysm of Chaos* texts is communicating with, and from, an online community of like-minded writers. It is thus important to establish whether the potentially discriminating n-grams are distinctive in this community, as opposed to the wider community of writers of English.

Consequently, we decided to use the texts from the other authors in the PoC1 and PoC2 as our specific comparative set. This approach had advantages and disadvantages. By using these other authors to create our impostors or foils, we automatically had access to a set of texts from a socially close speech community. We had no information as to whether any of these writers were also speakers of Greek (and so tackled this separately, below). This also limited the amount of text for the impostor group with samples from the other authors ranging from 240 to 6353 words.

The n-grams of size 1 to 5 from each text in PoC1 and PoC2 were extracted. Each list of n-grams of size  $n$  from each of the texts were collated to form a master list of n-grams of size  $n$ . For instance, the master bigram list included all 34,434 bigrams from 29 texts (from the PoC1 (20 texts) and PoC2 (9 texts)). Each text was then automatically analysed using R for the frequency of each of these bigrams and the results were recorded in a data matrix, such as Table 3, which presents an extract of this matrix.

**Table 3. Extract from the bigram data matrix recording the occurrence of each bigram in the corpus in each of the texts.**

	Anonymous_1.txt	Archegonos_1.txt	Archegonos_2.txt	Archegonos_3.txt
<b>i_can</b>	2	1	1	1
<b>can_only</b>	4	0	0	0
<b>only_be</b>	2	0	1	0

Each text was then compared with every other text in terms of which n-grams were shared (types) as well as the overall number of shared n-grams (tokens). This involved summing the counts of the n-grams for each comparison set. So, taking Table 3 as an example, Anonymous\_1 and Archegonos\_2 share the bigrams “i\_can” and “only\_be”. There are 3 occurrences of each n-gram across these texts. Thus, the total frequency of shared n-grams from Table 3 between these two texts is 6. Because the lengths of the texts differ substantially (see Tables 1 and 2), the number of shared n-grams were normalised by dividing that number by the sum of the number of word tokens in each text in the pair of texts that were being compared. So, for instance, Anonymous\_1 and Archegonos\_2 share 249 bigrams in total. Anonymous\_1 has 1309-word tokens, whilst Archegonos\_2 has a total of 8364-word tokens, meaning that they share 25.74 bigrams per thousand words.

This process was used to establish a consistent within-author pattern for Archegonos. Specifically, where n-grams were used across two or more of Archegonos’ texts, even if only in a single instance in each text, we considered this sufficient to indicate a consistent aspect of style across his known texts. For example, in Table 3 “I can” is consistently used across Archegonos 1, 2 and 3, even though in each text it only occurs once. Once IC had established the consistently-used n-grams, the process was repeated to include the n-grams from the disputed text. Each compared pair of texts were then ranked according to their relative frequency for the particular n-gram length. So, the text pair that had the highest relative frequency of shared n-grams (of length  $n$ ) was assigned 1, whilst the text that had the lowest relative frequency of shared n-grams was assigned 465. This was repeated for each length of n-gram. The median of the ranks for all n-gram lengths for each compared pair of texts was retrieved. (The median is the central number of the rankings when all the rankings for a comparison pair are arranged from smallest to largest.) The lowest median indicated the most similar text pair with respect to all n-gram lengths. Median was used for ranking in case the data was skewed by outliers.

In addition to retrieving an overall rate of sharedness between authors, within authors and with the disputed text, we needed to assess distinctiveness of the shared n-grams to answer our hypotheses. For H1, we searched for 1-5-word n-grams which were (1) repeated across Archegonos sections in the PoC texts; (2) exclusive to Archegonos as a writer in the PoC texts, i.e. which occurred only in the Archegonos text and not in the text written by the other writers; and (3) in the disputed text.

With respect to H2, we searched for 1-5-word n-grams which were (1) shared between the disputed and other writers in the PoC 1 and PoC2, but (2) not shared with any of Archegonos’ texts. Apart from Wolfi Landstrieher, no other authors had mul-

tiple texts. As a result, we could not check for consistency with respect to our other authors.

Finally, because Karvounakis was a Greek national, the shared and distinctive n-grams between Archegonos and the disputed text were examined for their usage in the Greek section of the International Corpus of Learner English to determine whether any of them were indicative of generic Greek-English idioms<sup>3</sup>.

#### 4.1.6. Results

This analysis revealed that the texts that are most similar to each other across all n-gram lengths across PoC1 and PoC2 are Archegonos\_2 and Archegonos\_8, indicating within-author consistency. This is shown in Table 4 along with the top ten most similar ranked file pairings. The next most similar pairs of text are Archegonos\_8 and Archegonos\_2 linked with the Unknown texts. As discussed above, we had always believed that some of these texts had been authored by Archegonos, but because the sign-offs were missing or unclear we did not want to include them as known to have been written by him. Given that these Unknown author texts when paired with known Archegonos texts share numerous n-grams, we concluded that it is consistent with the proposition that at least parts of these texts were indeed authored by Archegonos. Unknown\_1 linked with Unknown\_3 are the next in the ranking of most similar text pairs, sharing a high proportion of n-grams. This is perhaps not surprising given that they may well have the same author(s), and also are in a similar register in their announcement of and claims of responsibility to attacks. There remains the possibility that none of Archegonos nor the impostors wrote the Unknown author texts, and this remains a weakness of the IM, which might be mitigated by the non-linguistic evidence in the case.

**Table 4. Top 10 most similar text pairs by ranked n-gram analysis.**

Paired texts	Unigram rank	Bigram rank	Trigram rank	Fourgram rank	Fivegram rank	Median rank
Archegonos_2 + Archegonos_8	1	1	1	1	1	1
Archegonos_8 + Unknown_4	2	2	2	2	6	2
Archegonos_2 + Unknown_1	3	3	4	20	36	4
Archegonos_8 + Unknown_1	4	4	5	6	12	5
Unknown_1 + Unknown_3	11	9	7	5	7	7
Archegonos_2 + Unknown_4	5	7	8	13	19	8
Archegonos_2 + Archegonos_3	8	11	3	7	35	8
Archegonos_2 + Unknown_3	7	8	6	19	42	8
Archegonos_2 + CoA_1	9	5	9	11	17	9

When introducing the disputed claim of responsibility into this analysis, the most similar texts to the disputed text were written by Archegonos, specifically Archegonos\_8 and Archegonos\_2. These pairs of texts rank the 16th and the 30th most similar across all lengths of n-gram with all files included in the analysis (including those attributed

<sup>4</sup>Thanks go to one of our reviewers who writes: “‘Nichilistically’ is very interesting, because it is obviously not a Spanish spelling. Given the post-hoc knowledge that the suspect was Greek I searched and the word is spelled in Greek with the ‘ch’ symbol also used for Christ. So a focus on that spelling might have led to a suggestion that the author was Greek”.

to sign-offs other than Archegonos). The next most similar text to the disputed text was Unknown\_4, which was ranked the 35th most similar pair overall. The highest ranked text by similarity that was not written by Archegonos was by Wolfi Landsteicher\_2, which is ranked 45th in the overall table.

Given the strong within-author consistency between Archegonos\_2 and Archegonos\_8 and the strong similarity of these texts to the disputed text, it suggests that the disputed claim of responsibility is likely written by Archegonos. However, producing such a metric in a forensic context is problematic, especially in the context of a jury trial as it is not the task of the expert to tell the jury who wrote the text, but rather the linguist should serve as a “linguistic tour guide”, showing elements of similarity and difference and allowing the jury to make up their own mind (Grant, 2022; Solan, 1998). As a result, we looked more specifically at the types of shared n-grams to test our hypotheses.

To examine H1, n-grams that were exclusive to Archegonos and the disputed text were explored. This revealed 40 uni-grams, 141 bi-grams, 97 tri-grams, 19 four-grams, and 3 five-grams shared only between Archegonos’ texts and the disputed text. Yet, it was important that we only considered those that were used across two or more of his known texts, which in this small number of texts, we considered were part of his consistent style. This left us with 31 uni-grams, 55 bi-grams, 37 tri-grams, and 2 four-grams. Examples included “devours”, “nihil and”, “the human animal”, and “I am a nihilist”. An examination of all these n-grams in large general corpora of English, including the British National Corpus 2014 (spoken (Love, Dembry, Hardie, Brezina, & McEnery, 2017) and written (Brezina, Hawtin, & McEnery, 2021) on LancsBox’ X (Brezina & Platt, 2024)) and the English Web 2021 corpus on SketchEngine Kilgarrieff et al. (2014) show them to be rare, as can be seen in Table 5.

**Table 5. The frequencies and relative frequencies per million words (PMW) in the BNC2014 and enTenTen21 of a sample of the n-grams exclusive to Archegonos and the disputed text.**

	<b>enTenTen21</b>	<b>BNC2014</b>
devours	14526 (0.24 PMW)	11 (0.11 PMW)
nihil and	150 (<0.01 PMW)	1 (0.01 PMW)
the human animal	4352 (0.07 PMW)	8 (0.08 PMW)
I am a nihilist	43 (<0.01 PMW)	0 (0 PMW)

None of the five-grams occurred more than twice in known Archegonos’ texts. Rather, they only occurred once in Archegonos’ texts and once in the disputed. It is therefore not possible to argue that these are consistent features of Archegonos’ known style, but alongside the shared shorter n-grams (bi-grams, tri-grams, etc.) this does seem to provide some additional weight of evidence to link Archegonos with the claim of responsibility. The five-grams included: “unite on the basis of”, “for the human being to”, and “wouldn’t exist without the”. An exploration of these five-grams in the same large general corpora of English also show them to be exceedingly rare and therefore provide some measure of distinctiveness (see Table 6).

Because Karvounakis was a Greek national, each exclusive n-gram, which also appeared in the disputed text, was also examined to see if it was characteristic of the Greek section of the International Corpus of Learner English. Our purpose was to determine whether any of these n-grams were in fact indicative of generic Greek-English idioms, rather than an indication of individual style. However, none the exclusive n-grams were found in the ICLE corpus, so we concluded that none were in fact a function of Greek-English language contact.

**Table 6. The frequencies and relative frequencies per million words (PMW) in the BNC2014 and enTenTen21 of the five-grams in the disputed text shared with Archegonos' texts from PoC1 and PoC2.**

<b>five-grams</b>	<b>Freq. in enTenTen21</b>	<b>Freq. in BNC2014</b>
would n't exist without the	594 (0.01 PMW)	0 (0 PMW)
unite on the basis of	170 (0.01 PMW)	0 (0 PMW)
for the human being to	220 (0.01 PMW)	0 (0 PMW)

Additionally, it was important to consider H2 – that someone other than Archegonos wrote the disputed text – by finding n-grams that were exclusive to the disputed text and one of the other authors from PoC1 and PoC2 (but not Archegonos). This revealed in total across the other 19 sections (not written by Archegonos in PoC1 and PoC2) two five-grams, fourteen four-grams, 67 tri-grams, 150 bi-grams, and 80 uni-grams that were shared with the disputed text. This overlap is clearly important, but exploration of the n-grams across the BNC2014 and enTenTen21 deemed them to be more generic than those exclusive to Archegonos' text and the disputed text. For example, Table 7 presents some tri-grams, four-grams and five-grams that were shared by non-Archegonos' texts and with the disputed and their frequencies and relative frequencies in the BNC2014 and the English Web Corpus 2021. These frequencies show them to be more common than those exclusive to Archegonos and the disputed text and therefore carry little evidential weight. In addition, none of the other authors shared as many exclusive n-grams with the disputed text as Archegonos.

**Table 7. The frequencies and relative frequencies per million words (PMW) in the BNC2014 and enTenTen21 of n-grams shared across the disputed text and other authors from PoC1 and PoC2.**

	<b>Freq. in enTenTen21</b>	<b>Freq in BNC2014</b>
<b>who want to</b> (shared with Erinne Vivani and Unknown 4)	1148768 (18.65 PMW)	722 (7.06 PMW)
<b>and what is</b> (shared with Chris Kortright)	343176 (5.57 PMW)	460 (4.5 PMW)
<b>do not have</b> (shared with Chris Kortright)	8930217 (145 PMW)	1481 (14.48 PMW)
<b>there are people who</b> (shared with E. M Cioran)	66534 (1.08 PMW)	99 (0.97 PMW)
<b>those who want to</b> (shared with Erinne Vivani)	153082 (2.49 PMW)	127 (1.24 PMW)
<b>I do n't give a</b> (shared with Wolfi Landstreicher)	26751 (0.43 PMW)	157 (1.53 PMW)
<b>not to fall into the</b> (shared with Unknown 4)	2386 (0.04 PMW)	4 (0.04 PMW)

## 5. Reflections

In general, working on this case required drawing on a diverse set of methods and tools for forensic authorship analysis. The choice of tool was as much determined by the context at the stage of the case as by the structure of the authorship problem. Thus, in the Stage 1 analysis, which attempted to profile the writer, because of the security demands, a hand analysis performed by a single linguist was all that was possible. The outcome of this analysis, however, was not intended to be evidential, but simply to inform the investigation. For Stage 2, security was still a restriction. The analysis was now a cross-genre comparison intended to be taken before a court, but the evidential bar stayed relatively low. Stage 3 was for a criminal standard of proof but was free of the previous security constraints. The comparative analysis was also able to be carried out on within-genre texts drawn from the same speech community. In addition, steps were taken to avoid bias by recruiting a fresh linguist uncontaminated by previous work on the case.

These different contexts and demands influenced the complexity and the degree to which the approach could draw on well-attested methods from the literature. Even in the third stage, the specifics of this analysis required and allowed for variation from the published and tested General Impostors Method. A lot has been written on the necessity for validation of methods in all forms of forensic analysis (Morrison & Enzinger, 2016; Stoel et al., 2015), including forensic authorship analysis (e.g. Grant, 2022; Solan, 2012), and this is clearly crucial. However, as the progress of this case shows, adaptability of method to context is also a key component of at least forensic authorship analysis and this needs more thought and discussion in the literature.

### 5.1. Reflections on impostors methods

The Stage 3 analysis is an authorship verification task, one of the harder tasks of comparative authorship analysis. In contrast with classification tasks, in authorship verifi-

cation tasks, there is potentially an open set of other authors who might have written the disputed text, beyond the single suspect. One of the few published methods for authorship verification is the IM (Koppel & Winter, 2014). This effectively transforms the problem into a quasi-closed set problem by creating sets of random, plausible, and genre-matched impostor texts to the disputed text. Potha and Stamatatos (2017) argue that, like the police line-up, the impostor texts need to be matched to the suspect's known writings, rather than the disputed. We position the methodology we used for Stage 3 as being similar to Potha and Stamatatos' (2017) version of the GIM, in the sense that our impostor texts were socially similar to the known writings of the suspect. We did not follow Potha and Stamatatos' method of searching for internet texts with salient features. Yet arguably, our impostor texts are even more specific, having been chosen by our suspect to be included in the zines, which he compiled. They are anarchist texts, crossing a similar range of different genres, including poems, claims of responsibility, letters, introductions, and opinion pieces. Additionally, our approach is like Potha and Stamatatos' through our ranking of the pairs of compared texts. This does not entirely counter the argument that the quasi-close set can never be a true closed set. A badly collected or constructed comparison set of impostor documents can statistically emphasise similarity between Q and K documents. Although we tried to avoid this possibility by using the specific impostors provided by the case, this is a risk, or cross-examination issue, that cannot be entirely avoided.

In summary our approach is distinct from other impostor methods in a number of ways: First, we compare Archegonos' known writings to our impostor texts. This allows us to observe our within-author consistency, especially across genres, as well as the influence of genre on our impostor documents. Second, we have not artificially created the set of impostors. The convenience sample of potential impostors was provided naturally by the nature of the documents in the case. Whereas in the GIM, these are created through search engine hits of significant words (either in the disputed text or known writing samples), all of our impostors were found as authors in the PoC zines. This use of specific impostors allows us to assess the distinctiveness of both the disputed text and of Archegonos' known writings from writers drawn from the same very specific community and writing in similar genres. Just as in a line-up, choosing close foils makes the comparison harder and thus the identification of the suspect carries greater evidential weight. The third difference from the GIM is that previous advocates have used statistical distance metrics. As we felt that such approaches would add a layer of complexity in presentation to a jury, we instead used the ranking method described above based on the relative frequencies of shared n-grams. This preserves the usefulness of the IM whilst recognising the importance of the courtroom context for forensic work.

## 6. Outcomes of the case

On the basis of the Stage 3 analysis, TG and IC wrote up an expert witness report in which they were able to assert that:

- the probability of seeing these [language features] in the disputed text is high, if Karvounakis is the author; and that,
- the probability of seeing these [language features] in the disputed text is low if someone other than Karvounakis is the author.

This kind of conclusion did not say that Karvounakis wrote the text, but rather provides an opinion for the jury, which then allows them to draw their own conclusion. In UK jurisdictions, this is now considered best practice for the type of conclusion in an expert witness report. It should be noted that to our knowledge there was no linguist appointed as expert witness for the defence.

Karvounakis was presented with the full range of evidence against him, and this included our linguistic evidence. On the basis of this, he pled guilty to a contravention of Section 57 of the Terrorism Act 2000. Karvounakis maintained that he did not intend the device to detonate and that he had manufactured the device in such a way that it could not have ignited. The judge in the case, Lord Braid (2022), said: “the scientific evidence is that even if what you say is true, there remained a potential for explosion, due to the potential introduction of a hot filament wire in close proximity to a low explosive substance, for example if the fuse was not sufficiently well separated from the filament.” Lord Braid goes on to recognise the implications of our analysis saying to Karvounakis that “your present position is not consistent with the claim of responsibility which you posted, or at least caused to be posted, on the internet shortly after the device had been found.”

On the 16th February 2022, Karvounakis was sentenced to 8 years and 4 months in prison at the High Court in Edinburgh.

## Acknowledgements

We would like to thank our reviewers for their extremely helpful comments and feedback.

## References

- Andrijanić, I., & Bąkowski, J. (2021). Stylometry Then and Now: Authorship Verification of Vivekacūḍāmaṇi and Vedāntic Stotras. *Challenges of Interdisciplinary and Multidisciplinary Approach: New Horizons in Oriental Studies (Prace Orientalystyczne/Oriental Studies)*, 44, 30–46.
- Batel, A., Abrams, M., & Pezik, P. (2022). Native Dialect Influence Detection (NDID) differentiating between Mexican and Peninsular L1 Spanish in L2 English. *Language and Law=Linguagem e Direito*, 9(1), 120–145. Retrieved 2025-06-27, from <https://ojs.letras.up.pt/index.php/LLLD/article/view/12829/11685> doi: 10.21747/21833745/lanlaw/9\_1a6
- Brezina, V., Hawtin, A., & McEnery, T. (2021). The Written British National Corpus 2014 – design and comparability. *Text & Talk*, 41(5-6), 595–615. Retrieved 2025-06-27, from <https://www.degruyter.com/document/doi/10.1515/text-2020-0052/html> doi: 10.1515/text-2020-0052
- Brezina, V., & Platt, W. (2024). *#LancsBox X*. Lancaster University. Retrieved from <http://lancsbox.lancs.ac.uk>
- Crown Office and Procurator Fiscal. (2024). *Expert Witness Guide*. Retrieved from <https://www.copfs.gov.uk/for-professionals/expert-witness-guide>
- Darling, S., Valentine, T., & Memon, A. (2008). Selection of lineup foils in operational contexts. *Applied Cognitive Psychology*, 22(2), 159–169. Retrieved 2025-06-27, from <https://onlinelibrary.wiley.com/doi/10.1002/acp.1366> doi: 10.1002/acp.1366

- Edmond, G., Towler, A., Grouns, B., Ribeiro, G., Found, B., White, D., ... Martire, K. (2017). Thinking forensics: Cognitive science for forensic practitioners. *Science & Justice*, 57(2), 144–154. Retrieved 2025-06-27, from <https://linkinghub.elsevier.com/retrieve/pii/S1355030616301289> doi: 10.1016/j.scijus.2016.11.005
- Found, B. (2015). Deciphering the human condition: the rise of cognitive forensics. *Australian Journal of Forensic Sciences*, 47(4), 386–401. Retrieved 2025-06-27, from <http://www.tandfonline.com/doi/full/10.1080/00450618.2014.965204> doi: 10.1080/00450618.2014.965204
- Grant, T. (2022). *The Idea of Progress in Forensic Authorship Analysis* (1st ed.). Cambridge University Press. Retrieved 2025-06-27, from <https://www.cambridge.org/core/product/identifier/9781108974714/type/element> doi: 10.1017/9781108974714
- Grant, T., & Grieve, J. (2022). The Starbuck Case: Methods for Addressing Confirmation Bias in Forensic Authorship Analysis. In R. Perkins, I. Picornell, & M. Coulthard (Eds.), *Methodologies and Challenges in Forensic Linguistic Casework* (1st ed., pp. 13–28). Wiley. Retrieved 2025-06-27, from <https://onlinelibrary.wiley.com/doi/10.1002/9781394266661.ch2> doi: 10.1002/9781394266661.ch2
- Grant, T., & MacLeod, N. (2020). *Language and Online Identities: The Undercover Policing of Internet Sexual Crime* (1st ed.). Cambridge University Press. Retrieved 2025-06-27, from <https://www.cambridge.org/core/product/identifier/9781108766425/type/book> doi: 10.1017/9781108766425
- Grieve, J. (2007). Quantitative Authorship Attribution: An Evaluation of Techniques. *Literary and Linguistic Computing*, 22(3), 251–270. Retrieved 2025-06-27, from <https://academic.oup.com/dsh/article-lookup/doi/10.1093/llc/fqm020> doi: 10.1093/llc/fqm020
- Grieve, J., Clarke, I., Chiang, E., Gideon, H., Heini, A., Nini, A., & Waibel, E. (2019). Attributing the Bixby Letter using n-gram tracing. *Digital Scholarship in the Humanities*, 34(3), 493–512. Retrieved 2025-06-27, from <https://academic.oup.com/dsh/article/34/3/493/5145759> doi: 10.1093/llc/fqy042
- Judicial Institute for Scotland. (2023). *e-Jury Manual*. Retrieved from [https://judiciary.scot/docs/librariesprovider3/judiciarydocuments/judicial-institute-publications/export\\_e-jury\\_manual\\_2023-08-15\\_0911-\(1\)-\(1\).pdf?sfvrsn=ff85e46f\\_1](https://judiciary.scot/docs/librariesprovider3/judiciarydocuments/judicial-institute-publications/export_e-jury_manual_2023-08-15_0911-(1)-(1).pdf?sfvrsn=ff85e46f_1)
- Khonji, M., & Iraqi, Y. (2014). A slightly-modified gi-based author-verifier with lots of features (asgalf). In *CLEF 2014 Labs and Workshops, Notebook Papers*. CLEF and CEUR-WS.org.
- Kilgarriff, A., Baisa, V., Bušta, J., Jakubíček, M., Kovář, V., Michelfeit, J., ... Suchomel, V. (2014, July). The Sketch Engine. *Lexicography*, 1(1), 7–36. Retrieved 2025-06-27, from <https://utppublishing.com/doi/10.1007/s40607-014-0009-9> doi: 10.1007/s40607-014-0009-9
- Koppel, M., & Winter, Y. (2014). Determining if two documents are written by the same author. *Journal of the Association for Information Science and Technology*, 65(1), 178–187. Retrieved 2025-06-27, from <https://asistdl.onlinelibrary.wiley.com/doi/10.1002/asi.22954> doi: 10.1002/asi.22954
- Lord Braid. (2022). *Sentencing Statements*. Retrieved from <https://judiciary.scot/home/sentences-judgments/sentences-and-opinions/2022/02/16/hma-v-nikolaos-karvounakis>
- Love, R., Dembry, C., Hardie, A., Brezina, V., & McEnery, T. (2017). The Spoken BNC2014:

- Designing and building a spoken corpus of everyday conversations. *International Journal of Corpus Linguistics*, 319–344. Retrieved 2025-06-27, from <http://www.jbe-platform.com/content/journals/10.1075/ijcl.22.3.02lov> doi: 10.1075/ijcl.22.3.02lov
- Morrison, G. S., & Enzinger, E. A. (2016, April). New Paradigm for the Evaluation of Forensic Evidence. National Institute for Science and Technology. Retrieved from [https://www.nist.gov/system/files/documents/2020/04/03/04\\_morrisonenzinger\\_nist\\_workshop\\_2016\\_04\\_30a\\_optimized.pdf](https://www.nist.gov/system/files/documents/2020/04/03/04_morrisonenzinger_nist_workshop_2016_04_30a_optimized.pdf)
- Pautz, N., McDougall, K., Mueller-Johnson, K., Nolan, F., Paver, A., & Smith, H. M. J. (2023). Identifying unfamiliar voices: Examining the system variables of sample duration and parade size. *Quarterly Journal of Experimental Psychology*, 76(12), 2804–2822. Retrieved 2025-06-27, from <https://journals.sagepub.com/doi/10.1177/17470218231155738> doi: 10.1177/17470218231155738
- Potha, N., & Stamatatos, E. (2017). An Improved Impostors Method for Authorship Verification. In G. J. Jones et al. (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction* (Vol. 10456, pp. 138–144). Cham: Springer International Publishing. Retrieved 2025-06-27, from [https://link.springer.com/10.1007/978-3-319-65813-1\\_14](https://link.springer.com/10.1007/978-3-319-65813-1_14) (Series Title: Lecture Notes in Computer Science) doi: 10.1007/978-3-319-65813-1\_14
- R Core Team. (2023). *\_r: A Language and Environment for Statistical Computing\_*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org>
- Seidman, S. (2013). Authorship Verification Using the Impostors Method. In P. Forner, R. Navigli, & D. Tufis (Eds.), *CLEF 2013 Evaluation Labs and Workshop – Working Notes Papers*.
- Solan, L. M. (1998). Linguistic experts as semantic tour guides. *International Journal of Speech, Language and the Law*, 5(2), 87–106. Retrieved 2024-06-14, from <https://journal.equinoxpub.com/IJSLL/article/view/6582> doi: 10.1558/sll.1998.5.2.87
- Solan, L. M. (2012). Intuition versus Algorithm: The Case of Forensic Authorship Attribution. *Brooklyn Journal of Law and Policy*, 21(551).
- Stamatatos, E. (2009). A survey of modern authorship attribution methods. *Journal of the American Society for Information Science and Technology*, 60(3), 538–556. Retrieved 2025-06-27, from <https://onlinelibrary.wiley.com/doi/10.1002/asi.21001> doi: 10.1002/asi.21001
- Stamatatos, E., Kestemont, M., Potthast, M., Kredens, K., Pezik, P., Heini, A., ... Potthast, M. (2022). Overview of the Authorship Verification Task at PAN 2022. In *CEUR workshop proceedings* (Vol. 3180, pp. 2301–2313). CEUR-WS. org. (Series Title: Lecture Notes in Computer Science) doi: 10.1007/978-3-031-13643-6\_24
- Stoel, R. D., Berger, C. E., Kerkhoff, W., Mattijssen, E. J. A. T., & Dror, E. I. (2015). Minimizing contextual bias in forensic casework. In M. J. Hickman & K. Strom (Eds.), *Forensic Science and the Administration of Justice: Critical Issues and Directions* (Online-Ausg ed., pp. 67–86). Los Angeles: SAGE.