PILOT STUDY

Screening the Handwriting of Different Individuals Using CEDAR-FOX

Jessica Owen[1]

Abstract: *The use of computer software to assist in screening a database of anonymous writings for potential authorship links is investigated. Large databases of writings, such as the New Zealand Police Document Examination Section's Anonymous Letter Database (ALD) can be unwieldy and time consuming to search manually. CEDAR-FOX is software which can perform a similar role with documents scanned into a computer using some manual processing. The results generated by CEDAR-FOX, the amount of user intervention required, and the time taken to process documents in CEDAR-FOX were examined. On the whole, CEDAR-FOX performed well as a filtering tool, highly ranking documents completed by a common author. However, a number of poor performance results were also obtained. Further research into the operating processes of CEDAR-FOX may explain these observations. It is concluded that CEDAR-FOX cannot yet replace the ALD, although it can operate as a complementary tool for identifying possible authors of handwriting from a large pool.*

Reference: Jessica Owen (2014). Screening the Handwriting of Different Individuals Using CEDAR-FOX. J. Forensic Document Examination, Vol. 24, pp. 53-66.

Keywords: Handwriting identification, CEDAR-FOX, Automated handwriting systems, Anonymous writings

1. Introduction

1.1 Background

Offences such as robbery, blackmail, arson, breaching restraining orders, threatening to kill and stalking often involve written communications where the author hides under the cover of anonymity. A common feature of such behaviour is repeat offending which can extend over significant periods of time. Even when threats are not acted upon, the receipt of such communications can be distressing to the recipients and occupy significant amounts of police time.

At the New Zealand Police Document Examination Section (NZPDES), documents of this type are sorted and screened manually in an attempt to identify potential authors and link cases. The NZPDES Anonymous Letter Database (ALD) contains photocopies of handwritten letters and envelopes sent to public individuals and organisations. An examiner undertakes a preliminary screening between the handwriting in question and the handwritten passages in the ALD to determine whether there are enough similarities to warrant further comparison.

The ALD is constantly growing as new letters are added when they are submitted to the section. Accordingly, it is becoming more time-consuming to maintain. An investigation was therefore undertaken to determine whether using a computer-based system would be a more effective way of managing such a database. A small number of computer systems that could be capable of replacing the ALD exist worldwide (Kroon-Van der Kooij, 1996; Maguire & Moran, 1996; Schomaker, 2008). The Dutch police

1. Senior Document Examiner, New Zealand Police Document Examination Section, Wellington, NEW ZEALAND
Email:jessica.owen@police.govt.nz

created their own system, SCRIPT, as did the German Bundeskriminalamt, who created the Forensic Information System for Handwriting (FISH). However, at the time of writing, neither of these databases is publicly available. A third programme, CEDAR-FOX (Srihari, Srinivasan & Desai, 2007) is freely available and appears capable of being used as envisaged by the NZPDES (Srihari & Leedham, 2003).

The CEDAR-FOX software uses algorithms to perform handwriting identification and verification tasks and expresses its findings in the form of Log Likelihood Ratios (LLRs) (Srihari et al., 2007). A positive LLR suggests that the author of a questioned sample of handwriting completed the handwriting it is compared against (referred to in CEDAR-FOX as the 'known' handwriting). A negative LLR suggests that the questioned and known handwritings were completed by different authors. Handwriting identification results are presented in a ranked list, ordered from the LLR with the greatest magnitude to the least, positive LLR to negative.

Handwriting verification studies using CEDAR-FOX outline the myriad features that can affect the LLRs produced by the programme (Kabra, Srinivasan, Huang & Srihari, 2007; Srihari, Huang & Srinivasan, 2008). These features include the quantity and content of handwriting compared, which can affect both the polarity and magnitude of the LLR produced.

Literature on the subject of handwriting identification suggests that CEDAR-FOX is effective at identifying one author, or screening for similar authors, from many in a database (Srihari, 2010; Verduijn, Van Den Heuvel & Stoel, 2011). These findings suggest that CEDAR-FOX will be a useful tool in the NZPDES.

1.2 Focus of research

CEDAR-FOX is a user-interactive system (Srihari et al., 2007). It is not fully automated, meaning that the examiner can amend the automatic processing conducted by the software. Investigation into the effect of varying levels of user intervention on the comparison results will be explored in this study.

An assessment of the time required to input data into CEDAR-FOX will be undertaken. This will include determining how much information needs to be entered (and subsequently amended by the examiner). Also, the amount of writing needed to produce the most accurate results will be investigated.

The accuracy of the results can be measured in two ways – the list of results ranked in order or the LLR magnitude. Kabra et al. (2007) note that the magnitude alone of the LLR cannot be used to determine the strength of evidence, which suggests that the ranked order of results will be more significant than the LLR magnitude in determining the accuracy of CEDAR-FOX.

The handwriting of different authors with different content will be compared to test the accuracy of CEDAR-FOX in a simulation of the type of writings encountered in casework.

Taking into account the accuracy of the results obtained and the amount of user intervention required, conclusions will be drawn regarding whether the CEDAR-FOX software can effectively streamline the matching of handwriting samples contained in the NZPDES ALD.

2. Method

2.1 Controls

A single document examiner conducted the experiment. This provided a level of consistency regarding variables such as the amount of editing (for example, deciding whether a hyphenated word was considered as one or two separate words), the time taken to edit and/or process each document and so on.

It was known which questioned documents had been completed by a common author. This meant that the validity of the results generated by CEDAR-FOX could be evaluated.

2.2 Sample collection

Naturally-completed handwriting was collected from the course of business documents of 12 different authors. The samples were selected based on the quantity and quality of handwriting available. Features such as writing style, size and case were not considered when selecting samples for comparison. Each sample needed to contain an extended passage (around five or more lines) of naturally and speedily completed text, with minimal background interference (such as lined paper, or other background noise). The selection criteria mean that the content of each sample is unique. Three samples were collected from each of the twelve different authors, providing a total of 36 documents.

2.3 Method

Pre-processing of files

Each sample was scanned in Full Colour, at 300dpi as a .tif file using a Konica Minolta Bizhub C280 Multifunction Device. The .tif file was edited in Adobe Photoshop CS5 (Version 12.0.4 x32) to remove background noise that may interfere or be confused with the handwritten text, such as pre-printed ruled lines or postage cancellation stamps. No editing of the content of the handwriting was undertaken. Each image was then converted to greyscale and saved as a .png file.

Each .png file was then subjected to the following processing in CEDAR-FOX (Version 1.3, March 3, 2008, CEDAR TECH):

PL1: Processing Level One – PROC

The images were batch-processed using CEDAR-FOX. The batch-processing function automatically processes the image using basic processing (PROC) and saves each image as a .fox file.

PL2: Processing Level Two – CW, Truth

The .fox files produced in PL1 were duplicated and processed further with user intervention. Word segmentations that had been automatically defined by CEDAR-FOX during PL1 were manually corrected as needed (Appendix A). This process is referred to as 'correct word segmentation' (CW). The actual text of each word, referred to as Truth, was then manually entered.

The manually-processed images were Saved As new .fox files.

PL3: Processing Level Three – ROI – PROC

A Region of Interest (ROI) of around 130 characters was selected from each whole page document and basic processing was performed. New .fox files were then saved for each ROI file.

PL4: Processing Level Four – ROI – CW, Truth

The ROIs were processed as per PL2 above.

PL5: Processing Level Five – ROI – Check ID

The .fox files produced after manual processing in PL4 were subjected to a further level of manual editing. CEDAR-FOX automatically identifies some characters. These were manually checked, and where necessary, corrected or deleted. Each file was then saved as a new .fox file. Appendix B demonstrates examples of this process.

Result generation

At the end of each processing level, the *Identify 1:n* function was performed. This function compares a questioned document with a user-selected quantity of known documents. Each of the 36 samples was used as the questioned document and was compared to a known database containing all 36 samples. This meant that each document was compared to itself, to two other documents by the same author, and to the remaining 33 passages completed by different authors. The comparison of each document to itself established a baseline LLR.

3. Results

CEDAR-FOX generates results in list form, with the documents ranked in order of the greatest to least LLR value. In every comparison run, the document compared to itself was the first-ranking document in the results list. For each comparison, the lowest-ranking document of the three documents by a common author was used to measure some of the results.

Accordingly, in figures and tables relating to ranking order or lowest-ranking documents, a smaller value reflects a better ranking than a higher value.

Thus, having the three documents completed by the same author appearing in the first three ranking positions was termed an Ideal ranking result.

3.1 Levels of processing

The lowest-ranking position of each document is demonstrated in Figure 1[i].The ranking order positions were largely variable for each author. The best ranking results were produced using PL4, where the majority of the questioned authors had all three documents appearing in the top ten ranking positions. The notable exceptions were authors DW and MC.

Under more detailed inspection, the unusual results for author MC could be attributed to one of the three documents by that author skewing the ranking order. When this document was not considered, the other two documents by the author MC consistently returned with themselves in the first and second (i.e.

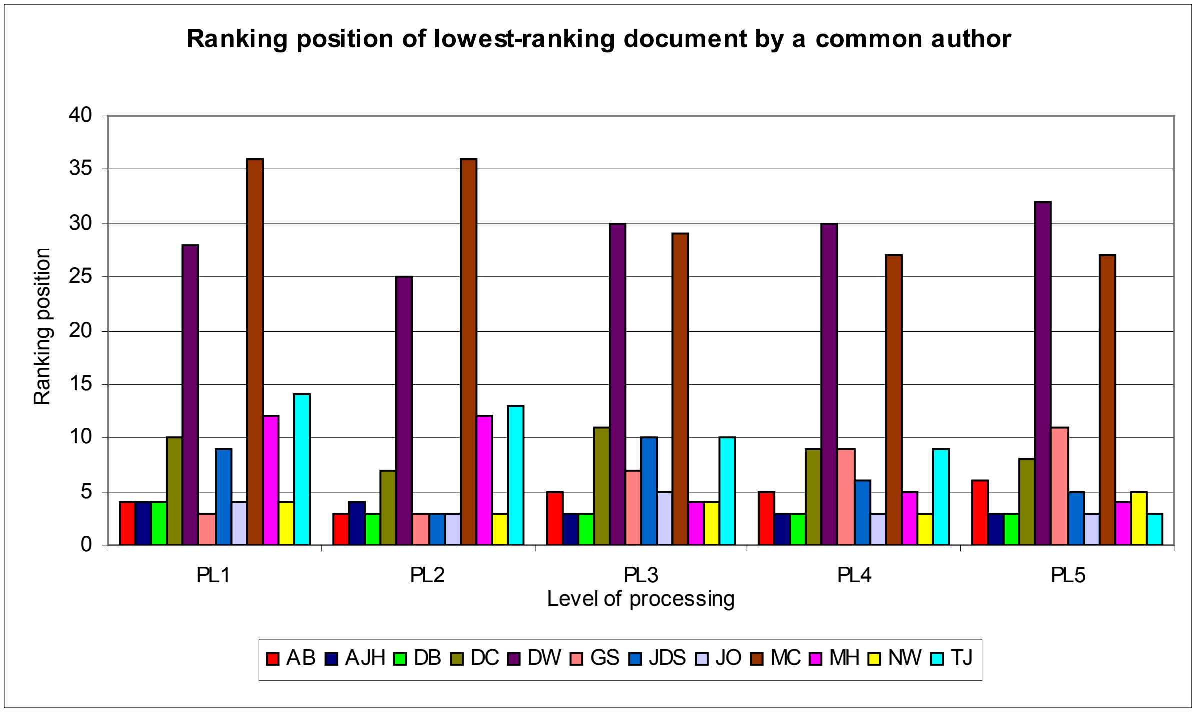

Figure 1 Ranking position of lowest-ranking document by a common author.

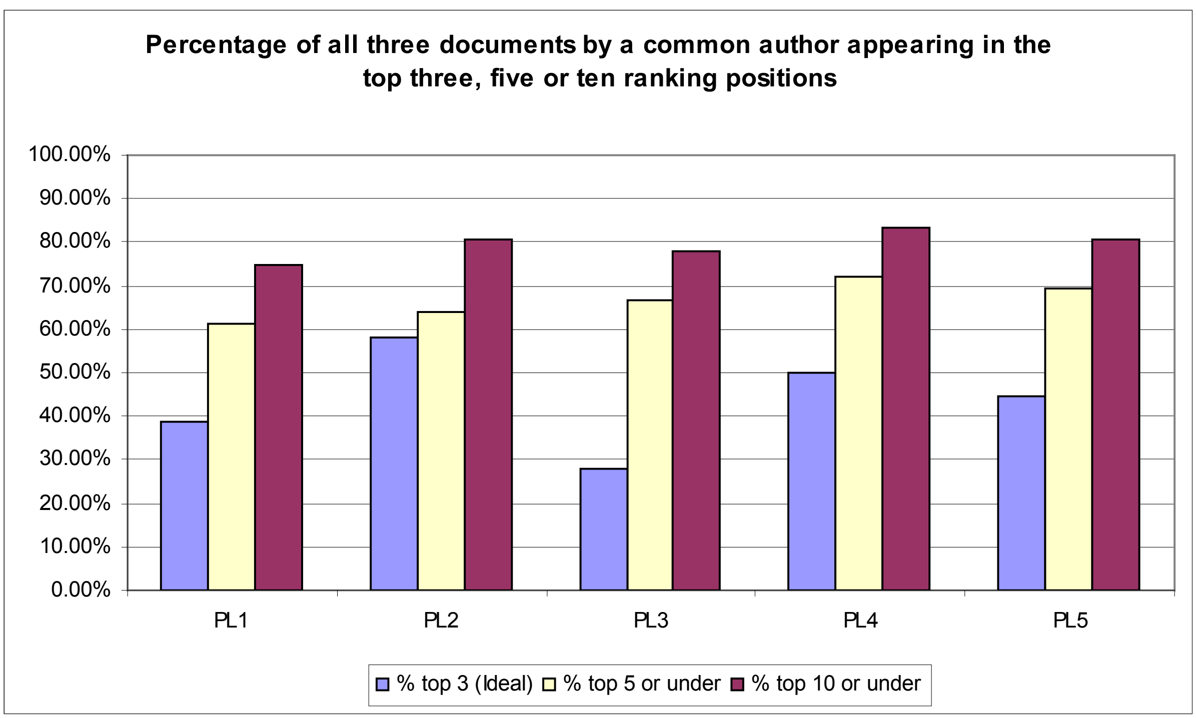

Figure 2 Percentage of all three documents by a common author appearing in the top three, five or ten ranking positions (PLs 1-5).

Ideal) ranking positions. That is, only one document ranked poorly when it was questioned against other documents by the same author.

This was not the case for the author DW, where there was no clear pattern to the poorly ranking results. When each of the three documents by DW were questioned, the second document consistently ranked within the top three positions at the whole-page processing levels (PL1 and 2), while whichever was the third document ranked poorly. However, at PLs 3-5, the second document ranked anywhere from 2nd to 18th place. PL2 was the level of processing that produced the best ranking order results for the third document by that author.

Figure 2 shows that different levels of processing performed better than others, depending on whether performance was measured by the three documents written by the author of the questioned document returning rankings in the top three, five or ten positions. For example, PL2 produced the greatest number of Ideal results (58%), while PL4 produced the most results in the top five or ten ranking positions (72% and 83%, respectively). In over 60% of the comparisons run, the three documents completed by the same author were returned within the top five ranking positions regardless of the level of processing.

3.2 Checking automatically-identified characters

The effect of amending automatically-identified characters varied. For 18 of the documents, there was no change in rankings between PL4 and PL5. Of the remaining 18 comparisons run, nine produced lower rankings and nine produced higher rankings.

3.3 Time taken

Table 1 shows the average time taken to process each document. There was a wide range of processing times due to the variable nature of the course of business specimens used in the experiment.

On average, it took around twelve minutes to scan, edit and process a whole-page document up to Processing Level Two. To process a ROI-sized document of around 130 characters, the time taken was estimated to be around seven minutes. This estimate is based on the editing in Photoshop taking three minutes for the ROI, rather than the four and a half minutes on average for a whole-page document.

3.4 Content and quantity of text

There was little correlation between the lowest-ranking document by a common author and the number of characters in the document. For example, a document with just 94 characters produced Ideal ranking results, while a document containing 294 characters ranked the third document by the common author in 36th (last) place. A document with 449 characters produced rankings at 31st and 34th positions for PL1 and 2, respectively, yet documents containing more than 450 characters consistently produced rankings within the top five results.

Activity	Approximate time
Scanning	1 minute
Editing in Photoshop (removing pre-printed lines, interfering backgrounds etc)	4.5 minutes
Correcting word segmentation, word truthing (whole-page)	6 minutes
Correcting word segmentation, word truthing (ROI)	2.5 minutes
Correcting automatically-identified characters (ROI)	2.5 minutes
Running Identify 1:n processing for each level of processing	under 1 minute

Table 1 Time taken to process each document

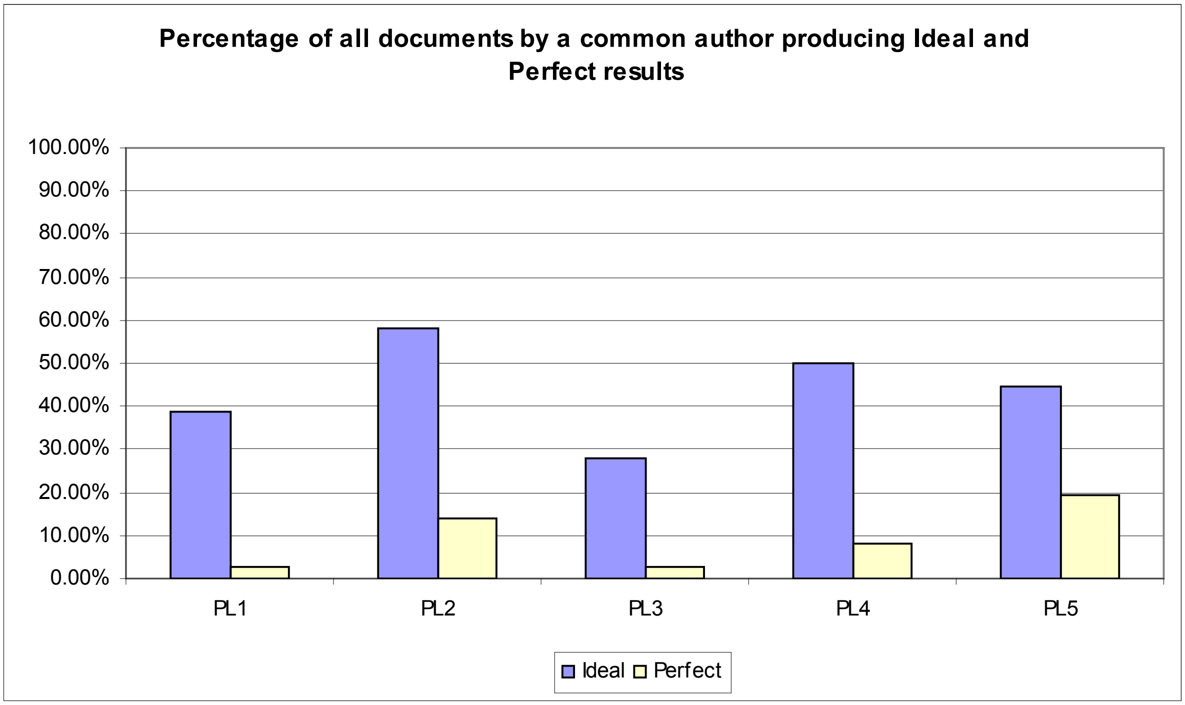

Figure 3 Percentage of all documents by a common author producing Ideal and Perfect results.

3.5 LLR magnitude

In general, the greater the number of characters in a document, the greater the LLR produced when that document was compared to itself. Accordingly, LLRs of a greater magnitude were produced for whole-page documents than Regions of Interest. PL3 returned the lowest LLR magnitudes of all the processing levels. This level of processing was the most basic, with no user input and it consisted of the smaller amount of text. Documents that had had word segmentation corrected and had been truthed (PLs 2 and 4) had greater LLRs than their basic processed (PROC) equivalents (PLs 1 and 3).

3.6 False positive/negative LLRs

A small number of false positive results (less than 1% of all the results generated) occurred only in the ROI-sized samples at PLs3-5. Such a small quantity of falsely positive LLR results was not considered problematic as it is preferable in this type of database searching to have incorrect results included in a result list than correct ones excluded.

The majority of true positive LLRs were for only the first result, when the questioned document was compared to itself. That is, even if the next two results were documents completed by the same author, and were correctly ranked second and third, they falsely had negative LLRs. This meant a large number of false negative LLRs were produced.

Accordingly, a comparison between the Ideal and Perfect results in both experiments was undertaken. A Perfect result was defined as all three documents by the same author ranking in the top three positions with a positive LLR. An Ideal result had the same ranking order as the Perfect result, but did not consider the polarity of the LLR.

A marked difference between the quantities of Ideal versus Perfect results was noted, due to the large number of falsely negative LLRs produced (Figure 3). Only 17 (9.44%) out of the 180 comparisons run (twelve authors x five processing levels x three documents per author) returned Perfect results. When only the ranking order position was considered, meaning the polarity of the LLR was ignored, this number rose to 79 (44%) out of 180. That is, 62 of the comparisons run which returned all three documents by the same author in the top three positions had one or more false negative LLRs. PL2 produced the greatest number of Ideal rankings, followed by PL4.

4. Discussion and Conclusions

4.1Efficacy of CEDAR-FOX

The present study showed that CEDAR-FOX tended to correctly rank documents by a common author in higher positions than those by different authors, even if the LLRs produced were falsely negative. Accordingly, the LLR magnitude was less descriptive than the ranking order as a measure of CEDAR-FOX's efficacy. These results supported the hypothesis based on the findings of Kabra et al. (2007) that the order of the results list was more significant than the LLR attached to each result.

4.2 Amount of user intervention needed

The investigation into the effect of user intervention on the results generated by CEDAR-FOX showed that two levels of processing consistently produced the best ranking results. PL2 produced the greatest quantity of Ideal ranking results and PL4 produced the most results within the top five or ten rankings.

PL2 and PL4 required a moderate amount of user intervention, namely correcting the word segmentation determined by CEDAR-FOX and applying the word truth. Both these levels of processing assessed a greater number of features within each document and produced LLRs with higher magnitudes than PL1 and PL3.

This finding is consistent with research on CEDAR-FOX which explains that the more information there is to process, the greater the LLRs will be (Kabra et al., 2007). Srihari et al. (2008) also noted that documents with the same content tended to produce greater LLR magnitudes than those with different contents. This finding was replicated in a preliminary study conducted by this author, who found that CEDAR-FOX produced more accurate ranking order results when similar content was present in the questioned documents than when the content varied from document to document (Owen, 2013).

4.3 Checking automatically-identified characters

After correcting word segmentation and truthing, the next level of processing requiring more user intervention was PL5, which involved checking and correcting, if necessary, characters that had been automatically-identified by CEDAR-FOX. This added around three minutes more processing time per document and produced highly variable results. In only one quarter of the results was there any improvement upon the results generated under the previous level of processing (PL4).

The checking of automatically-identified characters was trialled only on ROI-sized documents. Accordingly, the estimated time of 3 minutes per document would be considerably longer for whole-page samples.

The general lack of improvement in the ranking results led to the conclusion that for the database project envisioned checking the automatically-identified characters was not worth the extra time involved.

4.4 Time taken

Based on the average timings, processing a single document in CEDAR-FOX would be quicker than manually searching the ALD. However, if the questioned material consists of a greater quantity of text or multiple documents, the processing time in CEDAR-FOX will be longer than the time spent manually comparing them to the ALD. In addition, an initial back-capture of the existing documents in the ALD would need to be undertaken to create the known database of samples for the questioned documents to be compared against. This would take over 40 hours even if ROI sized samples were used.

The timings in the current research are estimates as they do not take into account the variations seen in actual casework. Some documents may contain much less questioned writing, which would result in reduced processing times, while others may require more pre-editing to reduce background noise, resulting in an increase in the time spent processing the document.

Additionally, it is unknown what effect background noise may have on the processing functions of CEDAR-FOX. This opens an avenue of potential future research to assess how background noise affects the ability of CEDAR-FOX to process documents, how much enhancement of text is needed and what other editing is required prior to documents being entered into CEDAR-FOX.

4.5 Quantity of text

There was little difference in the ranking orders produced for the ROI-sized samples which contained around 130 characters compared to the whole-page-sized documents which comprised around 320 characters on average. More research is required to determine possible reasons for these results, and whether there is an optimal quantity of text.

Due to the nature of anonymous letters, the amount of text available for processing is sometimes limited. For example, address blocks on envelopes contain a small, finite quantity of text. Accordingly, in these situations the optimum quantity of text needed to process a document in CEDAR-FOX becomes moot.

4.6 Anomalous results

In the case of documents completed by the author MC, CEDAR-FOX consistently ranked two documents (mc12 and mc33) more highly than the third by that author (mc5). This document (mc5) sometimes ranked as poorly as 36th (last) place in the ranked list of results.

It was noted that there were more pictorial similarities between the handwriting on mc12 and mc33 than between either of those two documents and mc5 (see Appendix C). Additionally, mc5 differed from the other two documents in other largely pictorial features such as writing size, slope and word spacing. It is therefore possible that CEDAR-FOX was placing a higher value on pictorial features than a human examiner would. This suggestion is reinforced by Verduijn et al (2011) noting that CEDAR-FOX looks only at pictorial similarity, and not at the "underlying movement execution" of the handwriting (p180). They note that the majority of the processing in CEDAR-FOX occurs behind the scenes. That is, the statistics are produced in a 'black box' and it can therefore be difficult to determine exactly which features are being assessed, and with what weighting, to produce the LLR (Verduijn et al., 2011).

Zhang, Srihari & Lee (2003) produced a ranked list of the discriminative strength of each character assessed by CEDAR-FOX. Some characters, such as 8, 0 and o, had low identification accuracies. However, document examiners considering constructional differences in these (and other) characters may regard them as more descriptive than other characters that CEDAR-FOX considered to be more discriminative. This suggests that CEDAR-FOX and document examiners are looking at and weighting their comparisons in different ways.

In addition, CEDAR-FOX examines only alpha-numeric characters (0-9, a-z and A-Z). This means that features such as #, $ and & which may be assessed by a human examiner are not examined by CEDAR-FOX. Accordingly, document examiners are assessing more and different features than CEDAR-FOX.

Despite the pictorial differences between mc5 and the other two documents completed by MC, there were many other handwriting features, such as direction of stroke, size relationships and individual letter constructions in common across all three documents. Accordingly, based on the similarities noted, a document examiner would likely conclude all three documents by MC were completed by a common author.

4.7 Limitations

The findings of this paper have possibly been limited by the small number of different authors (12 individuals) who produced the samples used in the experiments. This contrasts with the majority of the existing research into CEDAR-FOX which compared writings from around 1000 different individuals.

It is possible that the comparisons in CEDAR-FOX were limited by the content of the documents. Zhang et al. (2003) found CEDAR-FOX was able to identify 97.83% of writers using extracted Micro features from the characters 0-9, A-Z and a-z (62 characters in total). That is, all of these letterforms were present in all of the documents being compared. As the number of different characters compared decreased, so too did the percentage of authors identified. The documents used in the current study did not contain every letter of the alphabet in upper and lower case and every numeral 0-9, so the comparisons in CEDAR-FOX may have been limited by the content.

Further exploration of these results could be undertaken to determine whether there is any correlation between the ranking order results and the quantity of different characters available in each document.

The nature of anonymous letters is such that the content of the questioned documents varies and cannot be dictated. It would be useful to understand the effect the content of the documents being compared has on the results produced by CEDAR-FOX. This could aid in selecting Regions of Interest to focus on when faced with long anonymous letters.

4.8 Future research

One of the greatest limitations in this study has been a lack of information as to the processing CEDAR-FOX is conducting behind the scenes. Without knowing what weighting is placed on certain features, the optimum quantity of text required for future assessments cannot be accurately determined.

Additionally, determination of the features CEDAR-FOX is assessing, and how much weight is accorded to each feature may explain the causes of the poor ranking order results for some of the authors.

The effect of background noise on the ability of CEDAR-FOX to process documents could be examined. A focus on the quantity of editing in Photoshop needed for CEDAR-FOX to accurately process a document may provide a more accurate estimate of the time taken to input documents into the programme.

Srihari and Shi (2004) outline other features of CEDAR-FOX that may be useful in the collation and management of an anonymous letter database. Features such as searching the word truth would be useful in identifying questioned documents based on their content. This may highlight documents that have not been identified by a handwriting identification comparison in CEDAR-FOX, but warrant further comparison to one another. Accordingly, more research into these functions and how they may be applied to an anonymous letter database would be beneficial.

Testing could be undertaken with other similar document programmes, such as FISH and SCRIPT. A comparison between the results generated, amount of user input needed and the time involved in using these systems and CEDAR-FOX may provide context to the results obtained in the current paper.

4.9 Practical application of findings

The NZPDES ALD contains many samples determined to have been written by the same author which have already been grouped together. As the example of the documents completed by the author MC showed, even though one document by that author ranked poorly in the comparisons, the other two documents by that author still ranked highly. This suggests that while CEDAR-FOX may miss linking one document by a common author, there is a strong possibility that it will rank another or others by that same author highly. Accordingly, each of the documents in the existing ALD would need to be imported, in case one of them turns out to rank poorly.

Having to import as many documents as possible, rather than a representative sample, may impact any gains in efficiency created by using CEDAR-FOX instead of the ALD.

5. Conclusion

While only a preliminary study, the findings of the current paper suggest CEDAR-FOX will be a useful tool for database type searches. However, the efficiency and accuracy of CEDAR-FOX at present is limited by some as-yet unexplained anomalies. Accordingly, more research is required before the practicality of using CEDAR-FOX to manage an anonymous letter database can be fully assessed.

References

Kabra, S., Srinivasan H., Huang C. & Srihari S. N. (2007). On computing strength of evidence for writer verification. Published in *ICDAR 2007 Ninth International Conference on Document Analysis and Recognition* (pp 844-848). 23-26 September 2007 Curitiba, Brazil. doi:10.1109/ICDAR.2007.4377034

Kroon-Van der Kooij L. N. (1996). *The NIFO-TNO system "SCRIPT"*. Presented at the Fifty-Fourth Annual Meeting of the American Society of Questioned Document Examiners, 24-28 August 1996, Washington DC, and at the 14th Meeting of the Internal [sic] Association of Forensic Sciences, 26-30 August 1996, Tokyo, Japan.

Maguire K. B. and Moran T. L. (1996). *Identification of written text writings by the forensic information system for handwriting*. Prepared for the Fifty-Fourth Annual

Meeting of the American Society of Questioned Document Examiners. August 1996, Washington DC.

Owen J. J. L. (2013). *Screening a database of anonymous writings for authorship links using CEDAR-FOX computer software.* Research paper produced to meet requirements for promotion to Senior Document Examiner. September 2013. Wellington, New Zealand.

Schomaker L. (2007) Advances in writer identification and verification. Published in *ICDAR 2007 Ninth International Conference on Document Analysis and Recognition* (pp 1268-1273). 23-26 September 2007, Curitiba, Brazil. doi:10.1109/ICDAR.2007.4377119

Srihari S. N. (2010). Computational methods for handwritten questioned document examination. *National Criminal Justice Research Report*. Retrieved from https://www.ncjrs.gov/pdffiles1/nij/grants/232745.pdf

Srihari S. N., Huang C. & Srinivasan H. (2008). On the discriminability of the handwriting of twins. *Journal of Forensic Sciences 53* (2), 430-446.

Srihari S. N. & Leedham G. (2003). A survey of computer methods in forensic document examination. In Teulings H. L and Van Gemmert A. W. A. (Eds.), *Proceedings of the 11th Conference of the International Graphonomics Society (IGS2003)* (pp 278-281). 2-5 November 2003, Scottsdale, AZ.

Srihari S. N., Shi Z. (2004) Forensic handwritten document retrieval system. In *Proceedings of the First International Workshop on Document Image Analysis for Libraries (DIAL 04)* (pp 188-194). 23-24 January 2004, Palo Alto CA. doi:10.1109/DIAL.2004.1263248

Srihari S. N., Srinivasan H. & Desai K. (2007). Questioned document examination using CEDAR-FOX. *Journal of Forensic Document Examination 18*, 1-19.

Verduijn R. J., Van Den Heuvel C. E. & Stoel R. D. (2011). Exploratory investigation on the performance of the CEDAR-FOX system for forensic handwriting verification and identification. In *15th Conference of the International Graphonomics Society*, (pp177-180). 12-15 June 2011, Cancun, Mexico.

Zhang B., Srihari S. N. & Lee S. (2003). Individuality of handwritten characters. In *Proceedings of the 7th International Conference on Document Analysis and Recognition* (pp1086-1090). August 3-6, 2003, Edinburgh, Scotland.

(end note)

i As this graph can be difficult to interpret, the reader is invited to view the results for the author TJ with the following commentary: The x-axis displays the different levels of processing used and the y-axis shows the ranking order position of the lowest-ranking document completed by a common author. Accordingly, the results for the author TJ show improved results as the level of processing progressed. At PL1 the third document completed by TJ was ranked in 14th place. That is, the two other documents completed by TJ and 12 documents completed by other authors ranked more highly than the third document completed by TJ. At PL2 the worst-ranking document was ranked in 13th place; it ranked 10th at PL3 and 9th at PL4. At PL5 the Perfect result was achieved, with the third-ranked document completed by that author appearing in the third ranking position. That is, all three documents completed by TJ ranked in the top three places

Appendices

Appendix A
Correcting word segmentation (CW) after basic processing (PROC)

After basic processing (PL1), CEDAR-FOX automatically segments words by colour.

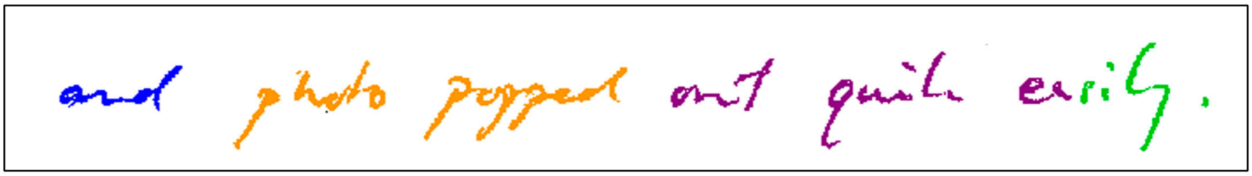

The programme may over-segment words, incorrectly considering separate words to be a single word.

In the above example, the words *photo popped* have incorrectly been considered to be a single word (shown in yellow), as has the text *out quite ea* (shown in purple).

CEDAR-FOX may under-segment a word, incorrectly considering a single word to be two separate words.

In the above example, the word *easily* is not a single colour, showing that CEDAR-FOX considers it to be two words: *ea* and *sily.*

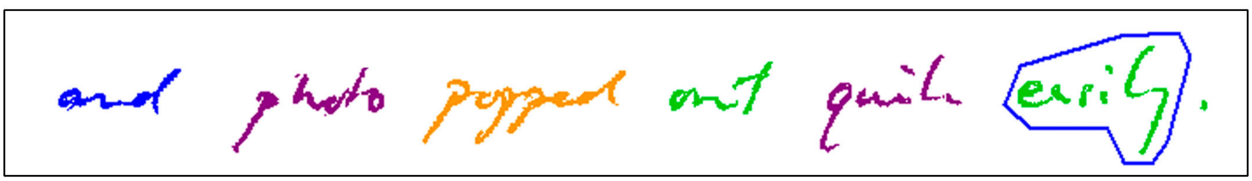

Incorrectly segmented words require user intervention to correct them. The user draws around the word as it should display, and selects the *Correct Word Segmentation* (CW) function. Each word then appears as a different colour.

Appendix B
Automatically-identified characters

CEDAR-FOX does not automatically identify every character. Only the characters boxed in red have been automatically identified.

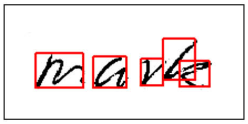

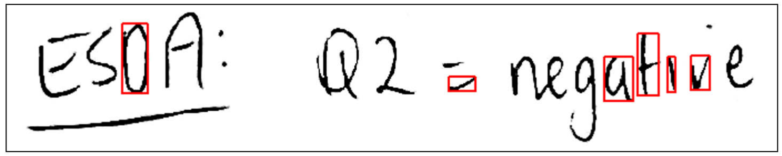

Right-clicking each character displays the value CEDAR-FOX has ascribed the character. The image below is for the *m* in *mark*.
CEDAR-FOX provides a figure of confidence for each character. If the character is changed (corrected) by the user, the confidence ratio reaches 1 (i.e. certainty).

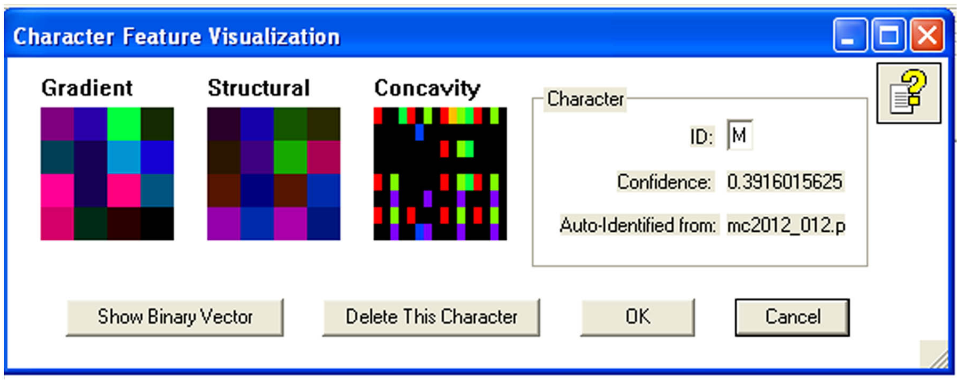

CEDAR-FOX often misidentified the automatically-identified characters. Sometimes this means mistaking an upper case character for a lower case one (or vice versa).
This is demonstrated above, where CEDAR-FOX has determined the lower case *m* to be an upper case *M*, with a confidence ratio of 0.39.

The images below depict the characters identified by CEDAR-FOX (in red) and the ground truth (in black).

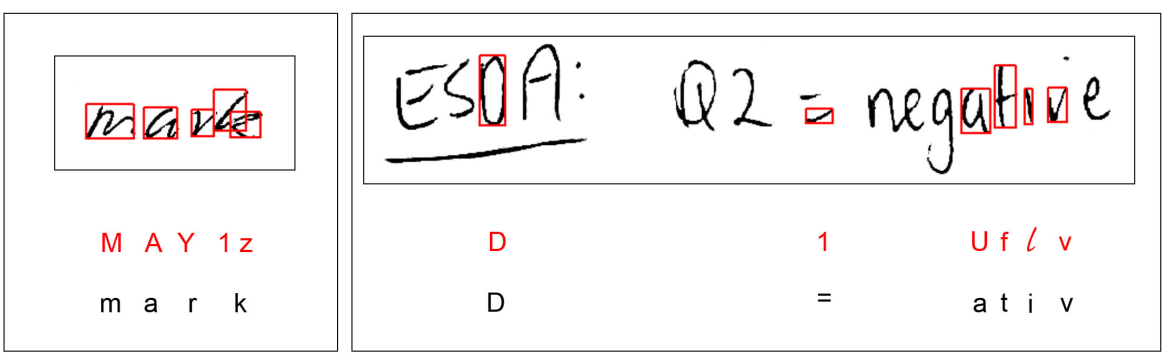

In these examples, CEDAR-FOX has misidentified the *k* in *mark*, perceiving it as being composed of a *1* and a *z*.
It has misidentified characters in the word *negative*, and has nominated part of the equals sign = (a character that CEDAR-FOX does not consider in its comparisons) as the numeral *1*.

Appendix C
Documents completed by the author MC

Q1 = + indentations from rear of envelope on front of card.
indentations from front of envelope on rear of card.
Nil indentations inside card.
Q2 = − Nil indentations on front & rear of envelope visible by sidelight.

ESDA: Q2 = negative −nil indentations located on front or rear.

Q1 = cover = + for indentations. Dry erase lift shows indentations sourced from g/ue text on envelope.

mc5

Counterfeit $100 (NZ). 2012/12

- Printed Offset lithography
- No serial #
- same characteristics as "no serial # series."
  - see photocopy.
- transparent windows not simulated → black/grey printed.
- no watermark.
- on paper.
- white mark near oval transparent window.
-
- entered in to CF database & No serial # series" spreadsheet.
- retain for CF library

mc12

Notes 2012/33

- Q Notebook contains handwritten entries on front and rear pages only.
- Specimen notebook KA1-9 has evidence of multiple pages having been removed.
  - handwritten text only on rear page of notebook. (KA9 rear).
- 4 x fingerprint forms ~~re~~ me for [redacted] received 5/3/12 − do not assist with h/w exam.
  → No fingerprint forms available for [redacted] (Livescan only).
- Indentations located on page ~~KA2~~ me KA2 of specimen notebook = similar to Q writing.
  * cell phone # ~~on~~ in indents for "TIM" is the same number recorded for [redacted] in NIA.

mc33

Names of involved parties have been pixellated for privacy.

The document mc5 is pictorially different from the other two documents in gross features such as handwriting style and slope.

Appendix D
Common handwriting features in documents completed by the author MC

All three documents completed by MC share common subtle features, such as directions of stroke, size relationships within words, letter connections and individual letter constructions.

Extracts from mc5	Extracts from mc12 and mc33	
		Common letter constructions and directions of stroke.
		Common letter connections and letter spacing.
		Common size relationships within words.