

Stylometry with R: A Package for Computational Text Analysis

by Maciej Eder, Jan Rybicki and Mike Kestemont

Abstract This software paper describes ‘Stylometry with R’ (**stylo**), a flexible R package for the high-level analysis of writing style in stylometry. Stylometry (computational stylistics) is concerned with the quantitative study of writing style, e.g. authorship verification, an application which has considerable potential in forensic contexts, as well as historical research. In this paper we introduce the possibilities of **stylo** for computational text analysis, via a number of dummy case studies from English and French literature. We demonstrate how the package is particularly useful in the exploratory statistical analysis of texts, e.g. with respect to authorial writing style. Because **stylo** provides an attractive graphical user interface for high-level exploratory analyses, it is especially suited for an audience of novices, without programming skills (e.g. from the Digital Humanities). More experienced users can benefit from our implementation of a series of standard pipelines for text processing, as well as a number of similarity metrics.

Introduction

Authorship is a topic which continues to attract considerable attention with the larger public. This claim is well illustrated by a number of high-profile case studies have recently made headlines across the popular media, such as the attribution of a pseudonymously published work to acclaimed *Harry Potter* novelist, J. K. Rowling (Juola, 2013), or the debate surrounding the publication of Harper Lee’s original version of *To Kill a Mocking Bird* and the dominant role which her editor might have played therein (Gamerman, 2015). The authorship of texts clearly matters to readers across the globe (Love, 2002) and therefore it does not come as a surprise that computational authorship attribution increasingly attracts attention in science, because of its valuable real-world applications, for instance, related to forensics topics such as plagiarism detection, unmasking the author of harassment messages or even determining the provenance of bomb letters in counter-terrorism research. Interestingly, the methods of stylometry are also actively applied in the Humanities, where multiple historic authorship problems in literary studies still seek a definitive solution – the notorious Shakespeare-Marlowe controversy is perhaps the best example in this respect.

Authorship attribution plays a prominent role in the nascent field of stylometry, or the computational analysis of writing style (Juola, 2006; Stamatatos et al., 2000; Stamatatos, 2009; Koppel et al., 2009; Van Halteren et al., 2005). While this field has important historical precursors (Holmes, 1994, 1998), recent decades have witnessed a clear increase in the scientific attention for this problem. Because of its emergent nature, replicability and benchmarking still pose significant challenges in the field (Stamatatos, 2009). Publicly available benchmark data sets are hard to come across, mainly because of copyright and privacy issues, and there are only a few stable, cross-platform software packages out there which are widely used in the community. Fortunately, a number of recent initiatives lead the way in this respect, such as the recent authorship tracks in the PAN competition (<http://pan.webis.de>), where e.g. relevant data sets are efficiently interchanged.

In this paper we introduce ‘Stylometry with R’ (**stylo**), a flexible R package for the high-level stylistic analysis of text collections. This package explicitly seeks to further contribute to the recent development in the field towards a more advanced level of replicability and benchmarking in the field. Stylometry is a multidisciplinary research endeavour, attracting contributions from divergent scientific domains, which include researchers from Computer Science – with a fairly technical background – as well as experts from the Humanities – who might lack the computational skills which would allow them easy access to the state-of-the-art methods in the field (Schreibman et al., 2004). Importantly, this package has the potential to help bridge the methodological gap lurking between these two communities of practice: on the one hand, **stylo**’s API allows to set up a complete processing pipeline using traditional R scripting; on the other hand, **stylo** also offers a rich Graphical User Interface which allows non-technical, even novice practitioners to interface with state-of-the-art methods without the need for any programming experience.

Overview of stylometry

Stylometry deals with the relationship between the writing style in texts and meta-data about those texts (such as date, genre, gender, authorship). Researchers in ‘stylochronometry’, for instance, are interested in inferring the date of composition of texts on the basis of stylistic aspects (Stamou, 2008;

Juola, 2007). Authorship studies are currently the most popular application of stylometry. From the point of view of literary studies, stylometry is typically concerned with a number of recent techniques from computational text analysis that are sometimes termed ‘distant reading’, ‘not reading’ or ‘macroanalysis’ (Jockers, 2013). Instead of the traditional practice of ‘close reading’ in literary analysis, stylometry does not set out from a single direct reading; instead, it attempts to explore large text collections using computational techniques (and often visualisation). Thus, stylometry tries to expand the scope of inquiry in the humanities by scaling up research resources to large text collections in order to find relationships and patterns of similarity and difference invisible to the eye of the human reader.

Usually, stylometric analyses involve a complex, multi-stage pipeline of (i) preprocessing, (ii) feature extraction, (iii) statistical analysis, and finally, (iv) presentation of results, e.g. via visualisation. To this end, researchers presently have to resort to an ad hoc combination of proprietary, language-dependent tools that cannot easily be ported across different platforms. Such solutions are difficult to maintain and exchange across (groups of) individual researchers, preventing straightforward replication of research results and reuse of existing code. **stylo**, the package presented, offers a rich, user-friendly suite of functionality that is ideally suited for fast exploratory analysis of textual corpora as well as classification tasks such as are needed in authorship attribution. The package offers an implementation of the main methods currently dominant in the field. Its main advantage therefore lies in the integration of typical (e.g. preprocessing) procedures from stylometry and statistical functionality by other, external libraries. Written in the R language, the source code and binaries for the package are freely available from the Comprehensive R Archive Network, guaranteeing a straightforward installation process across different platforms (both Unix- and Windows-based operating systems). The code is easily adaptable and extensible: the developers therefore continue to welcome user contributions, feedback and feature requests. Our code is open source and GPL-licensed: it is being actively developed on GitHub.¹

In the rest of this paper, we will first illustrate the functionality of the package for unsupervised multivariate analysis through the high-level function `stylo()`. Secondly, we will discuss a number of Graphical User Interfaces which we provide for quick exploration of corpora, in particular by novice users or students in an educational setting, as well as for scholars in the Humanities without programming experience. Next, we move on to the function `classify()`, implementing a number of supervised classification procedures from the field of Machine Learning. Finally, we concisely discuss the `oppose()`, `rolling.delta()` and `rolling.classify()` functionality which allow, respectively, to inspect differences in word usage between two subsets of a corpus, and to study the evolution of the writing style in a text.

Overview of the package

Downloading, installing and loading **stylo** is straightforward. The package is available at CRAN and at GitHub repository. The main advantages and innovative features of **stylo** include:

Feature extraction

Crucial in stylometry is the extraction of quantifiable features related to the writing style of texts (Sebastiani, 2002). A wide range of features have been proposed in the literature, considerably varying in complexity (Stamatatos, 2009). ‘Stylometry with R’ focuses on features that can be automatically extracted from texts, i.e. without having to resort to language-dependent preprocessing tools. The features that the package allows to extract are n -grams on token- and character level (Houvardas and Stamatatos, 2006; Kjell, 1994). Apart from the fact that this makes the package considerably language-independent, such shallow features have been shown to work well for a variety of tasks in stylometry (Daelemans, 2013; Kestemont, 2014). Moreover, users need not annotate their text materials using domain-specific tools before analyzing them with ‘Stylometry with R’. Apart from the standard usage, however, the package does allow the users load their own annotated corpora, provided that this is preceded by some text pre-processing tasks. An example of such a non-standard procedure will be shown below. Thus, **stylo** does not aim to supplant existing, more targeted tools and packages from Natural Language Processing (Feinerer et al., 2008) but it can easily accommodate the output of such tools as a part of its processing pipeline.

¹<https://github.com/computationalstylistics/stylo>

Metrics

A unique feature of **stylo** is that it offers reference implementations for a number of established distance metrics from multivariate statistical analysis, which are popular in stylometry, but uncommon outside the field. Burrows's Delta is the best example here (Burrows, 2002); it is an intuitive distance metric which has attracted a good share of attention in the community, also from a theoretical point of view (Hoover, 2004a,b; Argamon, 2011).

Graphical user interface

The high-level functions of the package provide a number of Graphical User Interfaces (GUIs) which can be used to intuitively set up a number of established experimental workflows with a few clicks (e.g. unsupervised visualization of texts based on word frequencies). These interfaces can be easily invoked from the command line in R and provide an attractive overview of the various experimental parameters available, allowing users to quickly explore the main stylistic structure of corpora. This feature is especially useful in an educational setting, allowing (e.g. undergraduate) students from different fields, typically without any programming experience, to engage in stylometric experimentation. The said high-level functions keep the analytic procedure from corpus pre-processing to final results presentation manageable from within a single GUI. More flexibility, however, can be achieved when the workflow is split into particular steps, each controlled by a dedicated lower-level function from the package, as will be showcased below.

Example workflow

An experiment in stylometry usually involves a workflow whereby, subsequently, (i) textual data is acquired, (ii) the texts are preprocessed, (iii) stylistic features are extracted (iv) a statistical analysis is performed, and finally, (v) the results are outputted (e.g. visualised). We will now illustrate how such a workflow can be performed using the package.

Corpus preparation

One of the most important features of **stylo** is that it allows loading textual data either from R objects, or directly from corpus files stored in a dedicated folder. Metadata of the input texts are expected to be included in the file names. The file name convention assumes that any string of characters followed by an underscore becomes a class identifier (case sensitive). In final scatterplots and dendrograms, colors of the samples are assigned according to this convention; common file extensions are dropped. E.g. to make the samples colored according to authorial classes, files might be named as follows:

ABronte_Agnes.txt	ABronte_Tenant.txt	Austen_Pride.txt
Austen_Sense.txt	Austen_Emma.txt	CBronte_Professor.txt
CBronte_Jane.txt	CBronte_Villette.txt	EBronte_Wuthering.txt

All examples below can be reproduced by the user on a data sets which can be downloaded from the authors' project website.² For the sake of convenience, however, we will use the datasets that come with the package itself:

```
data(novels)
data(galbraith)
data(lee)
```

Our first example uses nine prose novels by Jane Austen and the Brontë sisters, provided by the dataset `novels`.

Preprocessing

stylo offers a rich set of options to load texts in various formats from a file system (preferably encoded in UTF-8 Unicode, but it also supports other encodings, e.g. under Windows). Apart from raw text, **stylo** allows to load texts encoded according to the guidelines of the Text Encoding Initiative, which is relatively prominent in the community of text analysis researchers.³ To load all the files saved in a directory (e.g. `'corpus_files'`), users can use the following command:

²<https://sites.google.com/site/computationalstylistics/corpora>

³<http://www.tei-c.org/index.xml>

```
raw.corpus <- load.corpus(files = list.files(), corpus.dir = "corpus_files",
  encoding = "UTF-8")
```

If the texts are annotated in e.g. XML, an additional pre-processing procedure might be needed:

```
corpus.no.markup <- delete.markup(raw.corpus, markup.type = "xml")
```

Since the dataset that we will use has no annotation, the markup deletion can be omitted. We start the procedure with making the data visible for the user:

```
data(novels)
summary(novels)
```

To preprocess the data, **stylo** offers a number of tokenizers that support a representative set of European languages, including English, Latin, German, French, Spanish, Dutch, Polish, Hungarian, as well as basic support for non-Latin alphabets such as Korean, Chinese, Japanese, Hebrew, Arabic, Coptic and Greek. Tokenization refers to the process of dividing a string of input texts into countable units, such as word tokens. To tokenize the English texts, e.g. splitting items as ‘don’t’ into ‘do’ and ‘n’t’ and lowercasing all words, the next command is available:

```
tokenized.corpus <- txt.to.words.ext(novels, language = "English.all",
  preserve.case = FALSE)
```

The famous first sentence of Jane Austen’s *Pride and Prejudice*, for instance, looks like this in its tokenized version (the 8th to the 30th element of the corresponding vector):

```
tokenized.corpus$Austen_Pride[8:30]
```

[1] "it"	"is"	"a"	"truth"	"universally"
[6] "acknowledged"	"that"	"a"	"single"	"man"
[11] "in"	"possession"	"of"	"a"	"good"
[16] "fortune"	"must"	"be"	"in"	"want"
[21] "of"	"a"	"wife"		

To see basic statistics of the tokenized corpus (number of texts/samples, number of tokens in particular texts, etc.), one might type:

```
summary(tokenized.corpus)
```

For complex scripts, such as Hebrew, custom splitting rules could easily be applied:

```
tokenized.corpus.custom.split <- txt.to.words(tokenized.corpus,
  splitting.rule = "[^A-Za-z\u005C\u005D\u005E\u005F\u0060\u0061\u0062\u0063\u0064\u0065\u0066\u0067\u0068\u0069\u006A\u006B\u006C\u006D\u006E\u006F\u0070\u0071\u0072\u0073\u0074\u0075\u0076\u0077\u0078\u0079\u007A\u007B\u007C\u007D\u007E\u007F\u0080\u0081\u0082\u0083\u0084\u0085\u0086\u0087\u0088\u0089\u008A\u008B\u008C\u008D\u008E\u008F\u0090\u0091\u0092\u0093\u0094\u0095\u0096\u0097\u0098\u0099\u009A\u009B\u009C\u009D\u009E\u009F\u00A0\u00A1\u00A2\u00A3\u00A4\u00A5\u00A6\u00A7\u00A8\u00A9\u00AA\u00AB\u00AC\u00AD\u00AE\u00AF\u00B0\u00B1\u00B2\u00B3\u00B4\u00B5\u00B6\u00B7\u00B8\u00B9\u00BA\u00BB\u00BC\u00BD\u00BE\u00BF\u00C0\u00C1\u00C2\u00C3\u00C4\u00C5\u00C6\u00C7\u00C8\u00C9\u00CA\u00CB\u00CC\u00CD\u00CE\u00CF\u00D0\u00D1\u00D2\u00D3\u00D4\u00D5\u00D6\u00D7\u00D8\u00D9\u00DA\u00DB\u00DC\u00DD\u00DE\u00DF\u00E0\u00E1\u00E2\u00E3\u00E4\u00E5\u00E6\u00E7\u00E8\u00E9\u00EA\u00EB\u00EC\u00ED\u00EE\u00EF\u00F0\u00F1\u00F2\u00F3\u00F4\u00F5\u00F6\u00F7\u00F8\u00F9\u00FA\u00FB\u00FC\u00FD\u00FE\u00FF\u0100\u0101\u0102\u0103\u0104\u0105\u0106\u0107\u0108\u0109\u010A\u010B\u010C\u010D\u010E\u010F\u0110\u0111\u0112\u0113\u0114\u0115\u0116\u0117\u0118\u0119\u011A\u011B\u011C\u011D\u011E\u011F\u0120\u0121\u0122\u0123\u0124\u0125\u0126\u0127\u0128\u0129\u012A\u012B\u012C\u012D\u012E\u012F\u0130\u0131\u0132\u0133\u0134\u0135\u0136\u0137\u0138\u0139\u013A\u013B\u013C\u013D\u013E\u013F\u0140\u0141\u0142\u0143\u0144\u0145\u0146\u0147\u0148\u0149\u014A\u014B\u014C\u014D\u014E\u014F\u0150\u0151\u0152\u0153\u0154\u0155\u0156\u0157\u0158\u0159\u015A\u015B\u015C\u015D\u015E\u015F\u0160\u0161\u0162\u0163\u0164\u0165\u0166\u0167\u0168\u0169\u016A\u016B\u016C\u016D\u016E\u016F\u0170\u0171\u0172\u0173\u0174\u0175\u0176\u0177\u0178\u0179\u017A\u017B\u017C\u017D\u017E\u017F\u0180\u0181\u0182\u0183\u0184\u0185\u0186\u0187\u0188\u0189\u018A\u018B\u018C\u018D\u018E\u018F\u0190\u0191\u0192\u0193\u0194\u0195\u0196\u0197\u0198\u0199\u019A\u019B\u019C\u019D\u019E\u019F\u01A0\u01A1\u01A2\u01A3\u01A4\u01A5\u01A6\u01A7\u01A8\u01A9\u01AA\u01AB\u01AC\u01AD\u01AE\u01AF\u01B0\u01B1\u01B2\u01B3\u01B4\u01B5\u01B6\u01B7\u01B8\u01B9\u01BA\u01BB\u01BC\u01BD\u01BE\u01BF\u01C0\u01C1\u01C2\u01C3\u01C4\u01C5\u01C6\u01C7\u01C8\u01C9\u01CA\u01CB\u01CC\u01CD\u01CE\u01CF\u01D0\u01D1\u01D2\u01D3\u01D4\u01D5\u01D6\u01D7\u01D8\u01D9\u01DA\u01DB\u01DC\u01DD\u01DE\u01DF\u01E0\u01E1\u01E2\u01E3\u01E4\u01E5\u01E6\u01E7\u01E8\u01E9\u01EA\u01EB\u01EC\u01ED\u01EE\u01EF\u01F0\u01F1\u01F2\u01F3\u01F4\u01F5\u01F6\u01F7\u01F8\u01F9\u01FA\u01FB\u01FC\u01FD\u01FE\u01FF\u0200\u0201\u0202\u0203\u0204\u0205\u0206\u0207\u0208\u0209\u020A\u020B\u020C\u020D\u020E\u020F\u0210\u0211\u0212\u0213\u0214\u0215\u0216\u0217\u0218\u0219\u021A\u021B\u021C\u021D\u021E\u021F\u0220\u0221\u0222\u0223\u0224\u0225\u0226\u0227\u0228\u0229\u022A\u022B\u022C\u022D\u022E\u022F\u0230\u0231\u0232\u0233\u0234\u0235\u0236\u0237\u0238\u0239\u023A\u023B\u023C\u023D\u023E\u023F\u0240\u0241\u0242\u0243\u0244\u0245\u0246\u0247\u0248\u0249\u024A\u024B\u024C\u024D\u024E\u024F\u0250\u0251\u0252\u0253\u0254\u0255\u0256\u0257\u0258\u0259\u025A\u025B\u025C\u025D\u025E\u025F\u0260\u0261\u0262\u0263\u0264\u0265\u0266\u0267\u0268\u0269\u026A\u026B\u026C\u026D\u026E\u026F\u0270\u0271\u0272\u0273\u0274\u0275\u0276\u0277\u0278\u0279\u027A\u027B\u027C\u027D\u027E\u027F\u0280\u0281\u0282\u0283\u0284\u0285\u0286\u0287\u0288\u0289\u028A\u028B\u028C\u028D\u028E\u028F\u0290\u0291\u0292\u0293\u0294\u0295\u0296\u0297\u0298\u0299\u029A\u029B\u029C\u029D\u029E\u029F\u02A0\u02A1\u02A2\u02A3\u02A4\u02A5\u02A6\u02A7\u02A8\u02A9\u02AA\u02AB\u02AC\u02AD\u02AE\u02AF\u02B0\u02B1\u02B2\u02B3\u02B4\u02B5\u02B6\u02B7\u02B8\u02B9\u02BA\u02BB\u02BC\u02BD\u02BE\u02BF\u02C0\u02C1\u02C2\u02C3\u02C4\u02C5\u02C6\u02C7\u02C8\u02C9\u02CA\u02CB\u02CC\u02CD\u02CE\u02CF\u02D0\u02D1\u02D2\u02D3\u02D4\u02D5\u02D6\u02D7\u02D8\u02D9\u02DA\u02DB\u02DC\u02DD\u02DE\u02DF\u02E0\u02E1\u02E2\u02E3\u02E4\u02E5\u02E6\u02E7\u02E8\u02E9\u02EA\u02EB\u02EC\u02ED\u02EE\u02EF\u02F0\u02F1\u02F2\u02F3\u02F4\u02F5\u02F6\u02F7\u02F8\u02F9\u02FA\u02FB\u02FC\u02FD\u02FE\u02FF\u0300\u0301\u0302\u0303\u0304\u0305\u0306\u0307\u0308\u0309\u030A\u030B\u030C\u030D\u030E\u030F\u0310\u0311\u0312\u0313\u0314\u0315\u0316\u0317\u0318\u0319\u031A\u031B\u031C\u031D\u031E\u031F\u0320\u0321\u0322\u0323\u0324\u0325\u0326\u0327\u0328\u0329\u032A\u032B\u032C\u032D\u032E\u032F\u0330\u0331\u0332\u0333\u0334\u0335\u0336\u0337\u0338\u0339\u033A\u033B\u033C\u033D\u033E\u033F\u0340\u0341\u0342\u0343\u0344\u0345\u0346\u0347\u0348\u0349\u034A\u034B\u034C\u034D\u034E\u034F\u0350\u0351\u0352\u0353\u0354\u0355\u0356\u0357\u0358\u0359\u035A\u035B\u035C\u035D\u035E\u035F\u0360\u0361\u0362\u0363\u0364\u0365\u0366\u0367\u0368\u0369\u036A\u036B\u036C\u036D\u036E\u036F\u0370\u0371\u0372\u0373\u0374\u0375\u0376\u0377\u0378\u0379\u037A\u037B\u037C\u037D\u037E\u037F\u0380\u0381\u0382\u0383\u0384\u0385\u0386\u0387\u0388\u0389\u038A\u038B\u038C\u038D\u038E\u038F\u0390\u0391\u0392\u0393\u0394\u0395\u0396\u0397\u0398\u0399\u039A\u039B\u039C\u039D\u039E\u039F\u03A0\u03A1\u03A2\u03A3\u03A4\u03A5\u03A6\u03A7\u03A8\u03A9\u03AA\u03AB\u03AC\u03AD\u03AE\u03AF\u03B0\u03B1\u03B2\u03B3\u03B4\u03B5\u03B6\u03B7\u03B8\u03B9\u03BA\u03BB\u03BC\u03BD\u03BE\u03BF\u03C0\u03C1\u03C2\u03C3\u03C4\u03C5\u03C6\u03C7\u03C8\u03C9\u03D0\u03D1\u03D2\u03D3\u03D4\u03D5\u03D6\u03D7\u03D8\u03D9\u03DA\u03DB\u03DC\u03DD\u03DE\u03DF\u03E0\u03E1\u03E2\u03E3\u03E4\u03E5\u03E6\u03E7\u03E8\u03E9\u03EA\u03EB\u03EC\u03ED\u03EE\u03EF\u03F0\u03F1\u03F2\u03F3\u03F4\u03F5\u03F6\u03F7\u03F8\u03F9\u03FA\u03FB\u03FC\u03FD\u03FE\u03FF\u0400\u0401\u0402\u0403\u0404\u0405\u0406\u0407\u0408\u0409\u040A\u040B\u040C\u040D\u040E\u040F\u0410\u0411\u0412\u0413\u0414\u0415\u0416\u0417\u0418\u0419\u041A\u041B\u041C\u041D\u041E\u041F\u0420\u0421\u0422\u0423\u0424\u0425\u0426\u0427\u0428\u0429\u042A\u042B\u042C\u042D\u042E\u042F\u0430\u0431\u0432\u0433\u0434\u0435\u0436\u0437\u0438\u0439\u043A\u043B\u043C\u043D\u043E\u043F\u0440\u0441\u0442\u0443\u0444\u0445\u0446\u0447\u0448\u0449\u044A\u044B\u044C\u044D\u044E\u044F\u0450\u0451\u0452\u0453\u0454\u0455\u0456\u0457\u0458\u0459\u045A\u045B\u045C\u045D\u045E\u045F\u0460\u0461\u0462\u0463\u0464\u0465\u0466\u0467\u0468\u0469\u046A\u046B\u046C\u046D\u046E\u046F\u0470\u0471\u0472\u0473\u0474\u0475\u0476\u0477\u0478\u0479\u047A\u047B\u047C\u047D\u047E\u047F\u0480\u0481\u0482\u0483\u0484\u0485\u0486\u0487\u0488\u0489\u048A\u048B\u048C\u048D\u048E\u048F\u0490\u0491\u0492\u0493\u0494\u0495\u0496\u0497\u0498\u0499\u049A\u049B\u049C\u049D\u049E\u049F\u04A0\u04A1\u04A2\u04A3\u04A4\u04A5\u04A6\u04A7\u04A8\u04A9\u04AA\u04AB\u04AC\u04AD\u04AE\u04AF\u04B0\u04B1\u04B2\u04B3\u04B4\u04B5\u04B6\u04B7\u04B8\u04B9\u04BA\u04BB\u04BC\u04BD\u04BE\u04BF\u04C0\u04C1\u04C2\u04C3\u04C4\u04C5\u04C6\u04C7\u04C8\u04C9\u04CA\u04CB\u04CC\u04CD\u04CE\u04CF\u04D0\u04D1\u04D2\u04D3\u04D4\u04D5\u04D6\u04D7\u04D8\u04D9\u04DA\u04DB\u04DC\u04DD\u04DE\u04DF\u04E0\u04E1\u04E2\u04E3\u04E4\u04E5\u04E6\u04E7\u04E8\u04E9\u04EA\u04EB\u04EC\u04ED\u04EE\u04EF\u04F0\u04F1\u04F2\u04F3\u04F4\u04F5\u04F6\u04F7\u04F8\u04F9\u04FA\u04FB\u04FC\u04FD\u04FE\u04FF\u0500\u0501\u0502\u0503\u0504\u0505\u0506\u0507\u0508\u0509\u050A\u050B\u050C\u050D\u050E\u050F\u0510\u0511\u0512\u0513\u0514\u0515\u0516\u0517\u0518\u0519\u051A\u051B\u051C\u051D\u051E\u051F\u0520\u0521\u0522\u0523\u0524\u0525\u0526\u0527\u0528\u0529\u052A\u052B\u052C\u052D\u052E\u052F\u0530\u0531\u0532\u0533\u0534\u0535\u0536\u0537\u0538\u0539\u053A\u053B\u053C\u053D\u053E\u053F\u0540\u0541\u0542\u0543\u0544\u0545\u0546\u0547\u0548\u0549\u054A\u054B\u054C\u054D\u054E\u054F\u0550\u0551\u0552\u0553\u0554\u0555\u0556\u0557\u0558\u0559\u055A\u055B\u055C\u055D\u055E\u055F\u0560\u0561\u0562\u0563\u0564\u0565\u0566\u0567\u0568\u0569\u056A\u056B\u056C\u056D\u056E\u056F\u0570\u0571\u0572\u0573\u0574\u0575\u0576\u0577\u0578\u0579\u057A\u057B\u057C\u057D\u057E\u057F\u0580\u0581\u0582\u0583\u0584\u0585\u0586\u0587\u0588\u0589\u058A\u058B\u058C\u058D\u058E\u058F\u0590\u0591\u0592\u0593\u0594\u0595\u0596\u0597\u0598\u0599\u059A\u059B\u059C\u059D\u059E\u059F\u05A0\u05A1\u05A2\u05A3\u05A4\u05A5\u05A6\u05A7\u05A8\u05A9\u05AA\u05AB\u05AC\u05AD\u05AE\u05AF\u05B0\u05B1\u05B2\u05B3\u05B4\u05B5\u05B6\u05B7\u05B8\u05B9\u05BA\u05BB\u05BC\u05BD\u05BE\u05BF\u05C0\u05C1\u05C2\u05C3\u05C4\u05C5\u05C6\u05C7\u05C8\u05C9\u05CA\u05CB\u05CC\u05CD\u05CE\u05CF\u05D0\u05D1\u05D2\u05D3\u05D4\u05D5\u05D6\u05D7\u05D8\u05D9\u05DA\u05DB\u05DC\u05DD\u05DE\u05DF\u05E0\u05E1\u05E2\u05E3\u05E4\u05E5\u05E6\u05E7\u05E8\u05E9\u05EA\u05EB\u05EC\u05ED\u05EE\u05EF\u05F0\u05F1\u05F2\u05F3\u05F4\u05F5\u05F6\u05F7\u05F8\u05F9\u05FA\u05FB\u05FC\u05FD\u05FE\u05FF\u0600\u0601\u0602\u0603\u0604\u0605\u0606\u0607\u0608\u0609\u060A\u060B\u060C\u060D\u060E\u060F\u0610\u0611\u0612\u0613\u0614\u0615\u0616\u0617\u0618\u0619\u061A\u061B\u061C\u061D\u061E\u061F\u0620\u0621\u0622\u0623\u0624\u0625\u0626\u0627\u0628\u0629\u062A\u062B\u062C\u062D\u062E\u062F\u0630\u0631\u0632\u0633\u0634\u0635\u0636\u0637\u0638\u0639\u063A\u063B\u063C\u063D\u063E\u063F\u0640\u0641\u0642\u0643\u0644\u0645\u0646\u0647\u0648\u0649\u064A\u064B\u064C\u064D\u064E\u064F\u0650\u0651\u0652\u0653\u0654\u0655\u0656\u0657\u0658\u0659\u065A\u065B\u065C\u065D\u065E\u065F\u0660\u0661\u0662\u0663\u0664\u0665\u0666\u0667\u0668\u0669\u066A\u066B\u066C\u066D\u066E\u066F\u0670\u0671\u0672\u0673\u0674\u0675\u0676\u0677\u0678\u0679\u067A\u067B\u067C\u067D\u067E\u067F\u0680\u0681\u0682\u0683\u0684\u0685\u0686\u0687\u0688\u0689\u068A\u068B\u068C\u068D\u068E\u068F\u0690\u0691\u0692\u0693\u0694\u0695\u0696\u0697\u0698\u0699\u069A\u069B\u069C\u069D\u069E\u069F\u06A0\u06A1\u06A2\u06A3\u06A4\u06A5\u06A6\u06A7\u06A8\u06A9\u06AA\u06AB\u06AC\u06AD\u06AE\u06AF\u06B0\u06B1\u06B2\u06B3\u06B4\u06B5\u06B6\u06B7\u06B8\u06B9\u06BA\u06BB\u06BC\u06BD\u06BE\u06BF\u06C0\u06C1\u06C2\u06C3\u06C4\u06C5\u06C6\u06C7\u06C8\u06C9\u06CA\u06CB\u06CC\u06CD\u06CE\u06CF\u06D0\u06D1\u06D2\u06D3\u06D4\u06D5\u06D6\u06D7\u06D8\u06D9\u06DA\u06DB\u06DC\u06DD\u06DE\u06DF\u06E0\u06E1\u06E2\u06E3\u06E4\u06E5\u06E6\u06E7\u06E8\u06E9\u06EA\u06EB\u06EC\u06ED\u06EE\u06EF\u06F0\u06F1\u06F2\u06F3\u06F4\u06F5\u06F6\u06F7\u06F8\u06F9\u06FA\u06FB\u06FC\u06FD\u06FE\u06FF\u0700\u0701\u0702\u0703\u0704\u0705\u0706\u0707\u0708\u0709\u070A\u070B\u070C\u070D\u070E\u070F\u0710\u0711\u0712\u0713\u0714\u0715\u0716\u0717\u0718\u0719\u071A\u071B\u071C\u071D\u071E\u071F\u0720\u0721\u0722\u0723\u0724\u0725\u0726\u0727\u0728\u0729\u072A\u072B\u072C\u072D\u072E\u072F\u0730\u0731\u0732\u0733\u0734\u0735\u0736\u0737\u0738\u0739\u073A\u073B\u073C\u073D\u073E\u073F\u0740\u0741\u0742\u0743\u0744\u0745\u0746\u0747\u0748\u0749\u074A\u074B\u074C\u074D\u074E\u074F\u0750\u0751\u0752\u0753\u0754\u0755\u0756\u0757\u0758\u0759\u075A\u075B\u075C\u075D\u075E\u075F\u0760\u0761\u0762\u0763\u0764\u0765\u0766\u0767\u0768\u0769\u076A\u076B\u076C\u076D\u076E\u076F\u0770\u0771\u0772\u0773\u0774\u0775\u0776\u0777\u0778\u0779\u077A\u077B\u077C\u077D\u077E\u077F\u0780\u0781\u0782\u0783\u0784\u0785\u0786\u0787\u0788\u0789\u078A\u078B\u078C\u078D\u078E\u078F\u0790\u0791\u0792\u0793\u0794\u0795\u0796\u0797\u0798\u0799\u079A\u079B\u079C\u079D\u079E\u079F\u07A0\u07A1\u07A2\u07A3\u07A4\u07A5\u07A6\u07A7\u07A8\u07A9\u07AA\u07AB\u07AC\u07AD\u07AE\u07AF\u07B0\u07B1\u07B2\u07B3\u07B4\u07B5\u07B6\u07B7\u07B8\u07B9\u07BA\u07BB\u07BC\u07BD\u07BE\u07BF\u07C0\u07C1\u07C2\u07C3\u07C4\u07C5\u07C6\u07C7\u07C8\u07C9\u07CA\u07CB\u07CC\u07CD\u07CE\u07CF\u07D0\u07D1\u07D2\u07D3\u07D4\u07D5\u07D6\u07D7\u07D8\u07D9\u07DA\u07DB\u07DC\u07DD\u07DE\u07DF\u07E0\u07E1\u07E2\u07E3\u07E4\u07E5\u07E6\u07E7\u07E8\u07E9\u07EA\u07EB\u07EC\u07ED\u07EE\u07EF\u07F0\u07F1\u07F2\u07F3\u07F4\u07F5\u07F6\u07F7\u07F8\u07F9\u07FA\u07FB\u07FC\u07FD\u07FE\u07FF\u0800\u0801\u0802\u0803\u0804\u0805\u0806\u0807\u0808\u0809\u080A\u080B\u080C\u080D\u080E\u080F\u0810\u0811\u0812\u0813\u0814\u0815\u0816\u0817\u0818\u0819\u081A\u081B\u081C\u081D\u081E\u081F\u0820\u0821\u0822\u0823\u0824\u0825\u0826\u0827\u0828\u0829\u082A\u082B\u082C\u082D\u082E\u082F\u0830\u0831\u0832\u0833\u0834\u0835\u0836\u0837\u0838\u0839\u083A\u083B\u083C\u083D\u083E\u083F\u0840\u0841\u0842\u0843\u0844\u0845\u0846\u0847\u0848\u0849\u084A\u084B\u084C\u084D\u084E\u084F\u0850\u0851\u0852\u0853\u0854\u0855\u0856\u0857\u0858\u0859\u085A\u085B\u085C\u085D\u085E\u085F\u0860\u0861\u0862\u0863\u0864\u0865\u0866\u0867\u0868\u0869\u086A\u086B\u086C\u086D\u086E\u086F\u0870\u0871\u0872\u0873\u0874\u0875\u0876\u0877\u0878\u0879\u087A\u087B\u087C\u087D\u087E\u087F\u0880\u0881\u0882\u0883\u0884\u0885\u0886\u0887\u0888\u0889\u088A\u088B\u088C\u088D\u088E\u088F\u0890\u0891\u0892\u0893\u0894\u0895\u0896\u0897\u0898\u0899\u089A\u089B\u089C\u089D\u089E\u089F\u08A0\u08A1\u08A2\u08A3\u08A4\u08A5\u08A6\u08A7\u08A8\u08A9\u08AA\u08AB\u08AC\u08AD\u08AE\u08AF\u08B0\u08B1\u08B2\u08B3\u08B4\u08B5\u08B6\u08B7\u08B8\u08B9\u08BA\u08BB\u08BC\u08BD\u08BE\u08BF\u08C0\u08C1\u08C2\u08C3\u08C4\u08C5\u08C6\u08C7\u08C8\u08C9\u08CA\u08CB\u08CC\u08CD\u08CE\u08CF\u08D0\u08D1\u08D2\u08D3\u08D4\u08D5\u08D6\u08D7\u08D8\u08D9\u08DA\u08DB\u08DC\u08DD\u08DE\u08DF\u08E0\u08E1\u08E2\u08E3\u08E4\u08E5\u08E6\u08E7\u08E8\u08E9\u08EA\u08EB\u08EC\u08ED\u08EE\u08EF\u08F0\u08F1\u08F2\u08F3\u08F4\u08F5\u08F6\u08F7\u08F8\u08F9\u08FA\u08FB\u08FC\u08FD\u08FE\u08FF\u0900\u0901\u0902\u0903\u0904\u0905\u0906\u0907\u0908\u0909\u090A\u090B\u090C\u090D\u090E\u090F\u0910\u0911\u0912\u0913\u0914\u0915\u0916\u0917\u0918\u0919\u091A\u091B\u091C\u091D\u091E\u091F\u0920\u0921\u0922\u0923\u0924\u0925\u0926\u0927\u0928\u0929\u092A\u092B\u092C\u092D\u092E\u092F\u0930\u0931\u0932\u0933\u0934\u0935\u0936\u0937\u0938\u0939\u093A\u093B\u093C\u093D\u093E\u093F\u0940\u0941\u0942\u0943\u0944\u0945\u0946\u0947\u0948\u0949\u094A\u094B\u094C\u094D\u094E\u094F\u0950\u0951\u0952\u0953\u0954\u0955\u0956\u0957\u0958\u0959\u095A\u095B\u095C\u095D\u095E\u095F\u09
```

Features

After these preprocessing steps, users will want to extract gaugeable features from the corpus. In a vast majority of approaches, stylometrists rely on high-frequency items. Such features are typically extracted in the level of (groups of) words or characters, called n -grams (Kjell, 1994). Both word-token and character n -grams are common textual features in present-day authorship studies. **Stylo** allows users to specify the size of the n -grams which they want to use. For third order character trigrams ($n = 3$), for instance, an appropriate function of **stylo** will select partially overlapping series of character groups of length 3 from a string of words (e.g. 'tri', 'rig', 'igr', 'gra', 'ram', 'ams'). Whereas token level features have a longer tradition in the field, character n -grams have been fairly recently borrowed from the field of language identification in Computer Science (Stamatatos, 2009; Eder, 2011). Both n -grams at the level of characters and words have been listed among the most effective stylistic features in survey studies in the field. For $n = 1$, such text representations model texts under the so-called 'bag-of-words' assumption that the order and position of items in a text is negligible stylistic information. To convert single words into third order character chains, or tri-grams:

```
corpus.char.3.grams <- txt.to.features(corpus.no.pronouns, ngram.size = 3,
  features = "c")
```

Sampling

Users can study texts in their entirety, but also draw consecutive samples from texts in order to effectively assess the internal stylistic coherence of works. The sampling settings will affect how the relative frequencies are calculated and allow users to normalize text length in the data set. Users can specify a sampling size (expressed in current units, e.g. words) to divide texts into consecutive slices. The samples can partially overlap and they can be also be extracted randomly. As with all functions, the available options are well-documented:

```
help(make.samples)
```

To split the current corpus into non-overlapping samples of 20,000 words each, one might type:

```
sliced.corpus <- make.samples(tokenized.corpus, sampling = "normal.sampling",
  sample.size = 20000)
```

Counting frequent features

A crucial point of the dataset preparation is building a frequency table. In stylometry, analyses are typically restricted to a feature space containing the n most frequent items. It is relatively easy to extract e.g. the 3,000 most frequent features from the corpus using the following function:

```
frequent.features <- make.frequency.list(sliced.corpus, head = 3000)
```

After the relevant features have been harvested, users have to extract a vector for each text or sample, containing the relative frequencies of these features, and combine them into a frequency table for the corpus. Using an appropriate function from **stylo**, these vectors are combined in a feature frequency table which can be fed into a statistical analysis (external tables of frequencies can be loaded as well):

```
freqs <- make.table.of.frequencies(sliced.corpus, features = frequent.features)
```

Feature selection and sampling settings might interact: an attractive unique feature of **stylo** is that it allows users to specify different 'culling' settings. Via culling, users can specify the percentage of samples in which a feature should be present in the corpus in order to be included in the analysis. Words that do not occur in at least the specified proportion of the samples in the corpus will be ignored. For an 80% culling rate, for instance:

```
culled.freqs <- perform.culling(freqs, culling.level = 80)
```

Analysis

Stylo offers a seamless wrapper for a variety of established statistical routines available from R's core library or contributed by third-party developers; these include t-Distributed Stochastic Neighbor Embedding (van der Maaten and Hinton, 2008), Principal Components Analysis, Hierarchical Clustering and Bootstrap Consensus Trees (a method which will be discussed below). An experiment can be initiated with a pre-existing frequency table with the following command:

```
stylo(frequencies = culled.freqs, gui = FALSE)
```

When the input documents are loaded directly from text files, the default features are most frequent words (MFWs), i.e. 1-grams of frequent word forms turned into lowercase. Also, by default, a standard cluster analysis of the 100 most frequent features will be performed. To perform e.g. a Principal Components Analysis (with correlation matrix) of the 200 most frequent words, and visualize the samples position in the space defined by the first two principal components, users can issue the following commands:

```
stylo(corpus.dir = "directory_containing_the_files", mfw.min = 200, mfw.max = 200,
      analysis.type = "PCR", sampling = "normal.sampling", sample.size = 10000,
      gui = FALSE)
```

In Fig. 1, we give an example of how Principal Components Analysis (the first two dimensions) can be used to visualize texts in different ways, e.g. with and without feature loadings. Because researchers are often interested in inspecting the loadings of features in the first two components resulting from such an analysis, **stylo** provides a rich variety of flavours in PCA visualisations. For an experiment in the domain of authorship studies, for instance, researchers will typically find it useful to plot all texts/samples from the same author in the same color. The coloring of the items in plots can be easily controlled via the titles of the texts analyzed across the different R methods that are used for visualisation – a commodity which is normally rather painful to implement across different packages in R. Apart from exploratory, unsupervised analyses, **stylo** offers a number of classification routines that will be discussed below.

The examples shown in Fig. 1 were produced using the following functions:

```
stylo(frequencies = culled.freqs, analysis.type = "PCR",
      custom.graph.title = "Austen vs. the Bronte sisters",
      pca.visual.flavour = "technical",
      write.png.file = TRUE, gui = FALSE)

stylo(frequencies = culled.freqs, analysis.type = "PCR",
      custom.graph.title = "Austen vs. the Bronte sisters",
      write.png.file = TRUE, gui = FALSE)

stylo(frequencies = culled.freqs, analysis.type = "PCR",
      custom.graph.title = "Austen vs. the Bronte sisters",
      pca.visual.flavour = "symbols", colors.on.graphs = "black",
      write.png.file = TRUE, gui = FALSE)

stylo(frequencies = culled.freqs, analysis.type = "PCR",
      custom.graph.title = "Austen vs. the Bronte sisters",
      pca.visual.flavour = "loadings",
      write.png.file = TRUE, gui = FALSE)
```

Return value

Stylo makes it easy to further process the objects returned by an analysis. To cater for the needs of less technical users, the results returned by an analysis are saved by default to a number of standard files and outputted on screen. Advanced users can easily use the returned objects in subsequent processing:

```
stylo.results = stylo() # optional arguments might be passed

print(stylo.results)
summary(stylo.results)
```

The list of features created, for instance, can be easily accessed (and manipulated) subsequently, and the same applies to tables of frequencies or other results:

```
stylo.results$features
stylo.results$table.with.all.freqs
stylo.results$distance.table
stylo.results$pca.coordinates
```

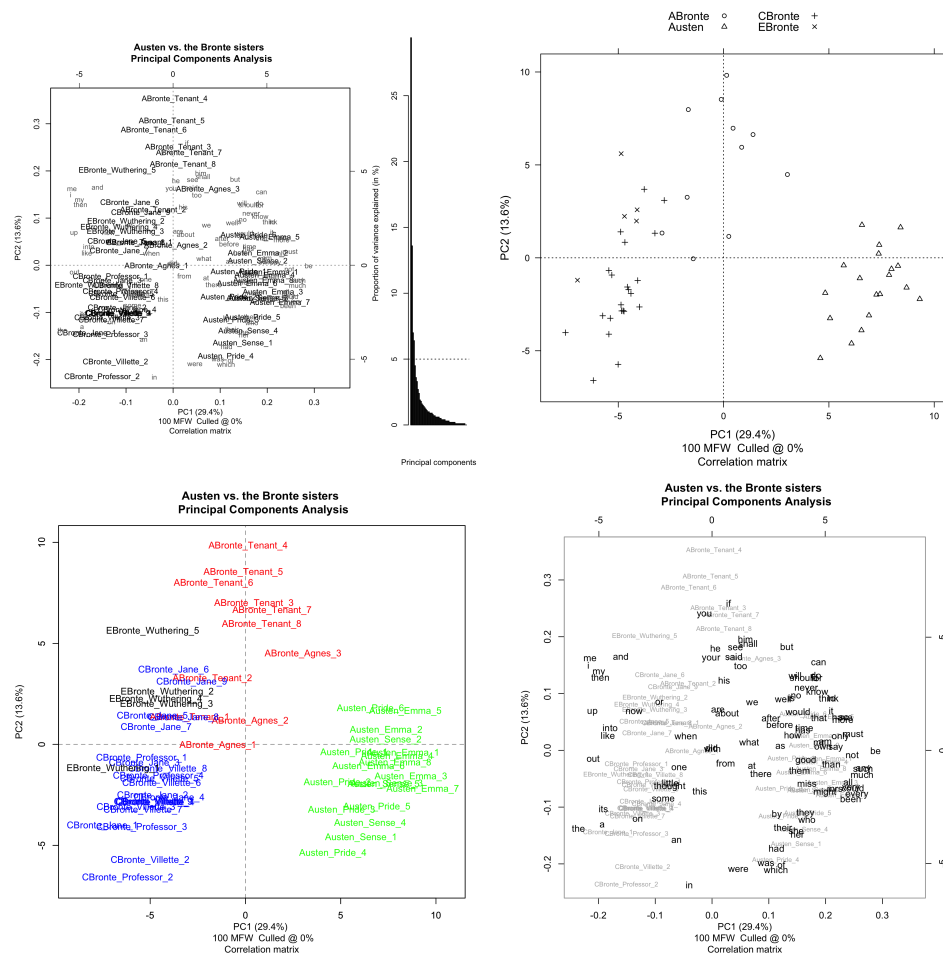


Figure 1: Illustration of different visualisation options for the first two dimensions outputted by a Principal Components Analysis (applied to 9 novels by 4 authors from our dummy corpus). Four different visualisation flavours are presented: ‘Technical’ (Fig. 1a), ‘Classic’ (Fig. 1b), ‘Symbols’ (Fig. 1c) and ‘Loadings’ (Fig. 1d). Users whose file names follow **stylo**’s naming conventions can easily exploit different coloring options.

GUI mode

Apart from the various functions to perform actual stylometric tasks, the package comes with a series of GUIs that can be used to set up typical experimental workflows in a quick and intuitive fashion. This unique feature renders **stylo** especially useful in educational settings involving students and scholars without programming experience. The cross-platform Graphical User Interface (automatically installed along with the rest of the package) has been written for Tcl/Tk and can be easily invoked from the command line. Four GUIs are currently available, which all come with extensive tooltips to help users navigate the different options. In this section, we will illustrate the use of these GUIs via an unsupervised stylometric experiment involving Bootstrap Consensus Trees.

The currently most widely used GUI component of ‘Stylometry with R’ is the eponymous GUI for **stylo()**, which is useful for the unsupervised stylistic exploration of textual corpora. It can be easily invoked using a single intuitive command (without the need to specify additional arguments):

```
stylo()
```

The various tabs of the **stylo** GUI (see Figure 2) present in a clear fashion the various parameters which can be specified before running the analysis by clicking the OK button. Users can freely switch between tabs and revisit them before running an experiment. Moreover, **stylo** will remember the experimental settings last used, and automatically default to these when users re-launch the GUI (which is useful for authors running a series of consecutive experiments with only small changes in parameters).

To illustrate the GUI mode, we will now concisely discuss a sample experiment involving Bootstrap Consensus Trees (BCT, selectable under the STATISTICS tab in the GUI). In stylometry, BCT exploits the

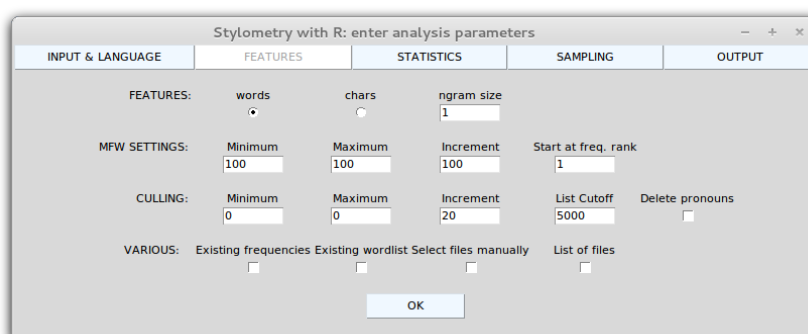


Figure 2: The Graphical User Interface for **stylo** (FEATURES tab). The high-level functions of **stylo** provide Graphical User Interfaces (GUIs) to intuitively set up experimental workflows. This feature is especially useful in an educational setting, allowing students without programming experience to engage in stylometric exploration and experimentation.

idea that the results become stable when one divides the list of MFW in non-identical, yet potentially overlapping frequency bands and analyzes these independently from each other (Eder, 2012). BCT were originally borrowed by Eder from the field of Language Evolution and Genetics; since a number of successful applications of the technique have been reported in the literature (Rybicki and Heydel, 2013; van Dalen-Oskam, 2014; Stover et al., in press). If the user specifies that different frequency bands should be used on the FEATURES tab, the bootstrap procedure will run different (virtual) cluster analyses and aggregate the results into a single (unrooted) consensus tree. This visualisation will only consider nodes for which there exists a sufficiently large consensus among the individual cluster analyses. The user in the corresponding text field (e.g. 0.5, which comes down to a majority vote for the cluster nodes). As such, users can assess the similarities between texts across different frequency bands.

Under the FEATURES tab, users can define the minutes of the MFW division and sampling procedure, using the increment, the minimum and maximum parameters. For minimum = 100, maximum = 3000, and increment = 50, **stylo** will run subsequent analyses for the following frequency bands: 100 MFW, 50–150 MFW, 100–200 MFW, ..., 2900–2950 MFW, 2950–3000 MFW. This is an attractive feature because it enables the assessment of similarities between texts across different bands in the frequency spectrum. A parallel logic underpins the CULLING text fields, where experiments will be carried out iteratively for different culling rates.

We illustrate the working of the BCT procedure in **stylo** using the recently covered case study on *Go Set a Watchman*, the second novel by Harper Lee, written before *To Kill a Mockingbird*. The novel itself attracted a reasonable attention worldwide, also because of its alleged authorship issues. Suspicion resurfaced about the strange fact that one of the greatest bestsellers in American history was its author's only completed work; Lee's childhood friendship with Truman Capote (portrayed as Dill in *To Kill A Mockingbird*) and their later association on the occasion of *In Cold Blood* fueled more speculations on the two Southern writers' possible, or even just plausible, collaboration; finally, the role of Tay Hohoff, Lee's editor on her bestseller, was discussed.

The stylometric study on this novel, featured in *Wall Street Journal* (Gamerman, 2015), revealed that the truth proved to be at once much less sensational than most of the rumors. Very strong stylometric evidence shows clearly that Harper Lee is the author of both *To Kill A Mockingbird* and *Go Set A Watchman*. In our replication of the experiment, the following code was used to produce the plots:

```
data(lee)

stylo(frequencies = lee, analysis.type = "CA",
      write.png.file = TRUE, custom.graph.title = "Harper Lee",
      gui = FALSE)

stylo(frequencies = lee, analysis.type = "CA",
      mfw.min = 1500, mfw.max = 1500, custom.graph.title = "Harper Lee",
      write.png.file = TRUE, gui = FALSE)

stylo(frequencies = lee, analysis.type = "BCT",
      mfw.min = 100, mfw.max = 3000, custom.graph.title = "Harper Lee",
      write.png.file = TRUE, gui = FALSE)
```

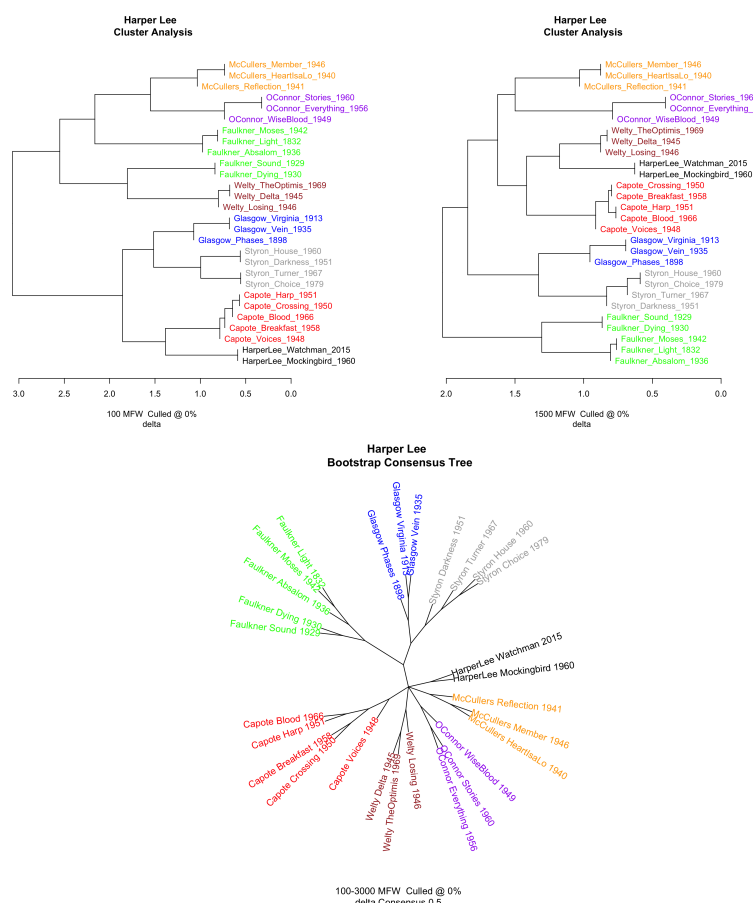


Figure 3: Analysis of the corpus of 28 novels by Harper Lee, Truman Capote as well as a number of comparable control authors writing in the American South. A frequency table of this corpus is provided by the package **stylo**, so that all are experiments can be replicated. In all plots, Lee's writing style is clearly very consistent, even if for some input parameters Lee's novels are close to Capote's. Figure panel 3a-3b: Traditional dendrograms outputted by cluster analyses with Burrows's Classic Delta Metric for 100 MFV and 1,500 MFV respectively (default settings; entire novels). Figure panel 3c: Bootstrap consensus tree for 100 MFV to 3,000 MFV (with an incremental step size of 50 words). Unrooted tree which combines clade information from analyses such as the ones presented in Fig. 1a-1b. The tree collapses nodes which were observed in at least 50% of the underlying trees (majority vote).

Classify

Apart from the already-discussed explanatory multivariate tests and the associated visualisations, stylometry has borrowed a number of advanced classification methods from the domain of Machine Learning. Some of them have simply been transferred to stylometry (e.g. Support Vector Machines or Naïve Bayes Classifier); others have been tailored to the needs of humanities researchers. The best example in this respect is Delta, a so-called 'lazy' learner developed by Burrows (Burrows, 2002). The **stylo** package offers an interface to a selection of established classifiers: including Burrows's original Delta and other distance-based classifiers, Nearest Shrunken Centroids, Support Vector Machines and Naïve Bayes Classifier. These are available through a single function:

```
classify() # optional arguments might be passed
```

If any non-standard text preprocessing procedures are involved, the above function can be fed with the result of a multi-stage custom pipeline. Combining the function `classify()` with spreadsheet tables of frequencies is also possible.

In a typical classification experiment, the analysis is divided in two stages. In the first stage, representative text samples for each target category (e.g. authorial group) are collected in a training corpus. The remaining samples form the test corpus. The first set, being a collection of texts, e.g. written by known authors ('candidates'), serves as a sub-corpus for fine-tuning the hyperparameters of a classifier and model architecture selection. The second set is a pool that consists of test texts of known

authorship and anonymous texts of disputed authorial provenance. The classifier's performance can be measured by applying a standard evaluation metric to the classifier's output on the test set (e.g. the number of correct attributions to authors in the training set). In **stylo**, users can divide their data over two subdirectories (or input custom-created R objects using the low-level functions discussed above); one directory should contain the training samples, the other the test samples. Other options can be specified via the parameters that run parallel to those of the `stylo()` function, such as the desired feature type or culling rate. Function-specific parameters for `classify()` include the number of cross-validation folds or the type of classifier (e.g. Support Vector Machine).

We illustrate the performance of classification methods in **stylo** using the well-known case study of the pseudonymous author Galbraith/Rowling, which recently attracted a good deal of press attention. In July 2013, the *Sunday Times* (UK) revealed that J. K. Rowling, the successful author behind the bestselling series of *Harry Potter* novels, had published a new detective novel (*The Cuckoo's Calling*) under the pseudonym of 'Robert Galbraith'. (The paper had received an anonymous tip with respect to this pen name over Twitter). For covering this case study, the *Sunday Times* has collaborated with Patrick Juola, an authority in the field of authorship attribution, and Peter Millican (Juola, 2013). They reported in a blog post on the Language Log that their stylometric analysis showed the writing style (e.g. on the level of function words) found in *The Cuckoo's Calling* to be broadly consistent with Rowling's writing in other works. Below, we report on a dummy attribution experiment which illustrates a supervised procedure.

In this experiment we will confront Galbraith's *The Cuckoo's Calling* with 25 other phantasy novels and thrillers by 4 famous novelists: H. Coben (e.g. *Tell No One*), C. S. Lewis (e.g. *The Chronicles of Narnia*), J. R. R. Tolkien (e.g. the *Lord of the Rings* trilogy) and J. K. Rowling (e.g. the *Harry Potter* series). Our replication experiments indeed confirm that Galbraith's writing style is more consistent with that of Rowling than that of any other author included. Instead of loading particular text files, we will use a computed table of frequencies provided by the package; the table has to be split into two tables (training set and test set). As an illustration, we specify the training set manually (with two training texts per class):

```
# specify a table with frequencies:
data(galbraith)
freqs <- galbraith

# specify class labels:
training.texts <- c("coben_breaker", "coben_dropshot", "lewis_battle",
                  "lewis_caspian", "rowling_casual", "rowling_chamber",
                  "tolkien_lord1", "tolkien_lord2")

# select the training samples:
training.set <- freqs[(rownames(freqs) %in% training.texts),]

# select remaining rows as test samples:
test.set <- freqs[!(rownames(freqs) %in% training.texts),]
```

To perform Delta on the Rowling corpus (50 MFWs, no sampling), we type:

```
classify(training.frequencies = training.set, test.frequencies = test.set,
         mfw.min = 50, mfw.max = 50, classification.method = "delta",
         gui = FALSE)
```

The results are automatically outputted to a log file 'final_results.txt':

```
galbraith_cuckoos --> rowling rowling coben

50 MFW, culled @ 0%, 17 of 17 (100%)

General attributive success: 17 of 17 (100%)

MFWs from 50 to 50 @ increment 100
Culling from 0 to 0 @ increment 20
Pronouns deleted: FALSE; standard classification
```

The overall performance of the classifier for our dummy corpus is optimal, since 100% of the test samples were correctly attributed to the correct authors. The experiment adds support to the identification of the author of *The Cuckoo's Calling* as Rowling. To combat model overfitting, cross-validation on the training data can be applied. It has been shown that for linguistic datasets a standard

10-fold cross validation might overestimate the performance of models, especially if languages other than English are assessed (Eder and Rybicki, 2013). To neutralize class imbalance, **stylo** therefore provides stratified cross-validation protocols for stylometric experiments. To perform a classification with a 'plain vanilla' 20-fold CV, using Nearest Shrunken Centroids classification and a series of tests for 50, 100, 150, 200, ..., 500 MFWs, one might type:

```
results <- classify(training.frequencies = training.set,
  test.frequencies = test.set,
  mfw.min = 50, mfw.max = 500, mfw.incr = 50,
  classification.method = "nsc", cv.folds = 20, gui = FALSE)
```

To inspect the classification accuracy for particular cross-validation folds, the user can type:

```
results$cross.validation.summary
```

Average scores of the cross-validation outcome (note that the overall performance is now slightly worse, ca. 95%) can be accessed via:

```
colMeans(results$cross.validation.summary)
```

Miscellaneous other functions

Apart from the above discussed functions, the package offers miscellaneous other, less established functions to stylometrically analyze documents. With the `oppose()` function, users can contrast two sets of documents and extract the most characteristic features in both sets of texts. The most discriminative features can be visualized and fed into other components of the package as part of a pipeline. Several metrics are implemented that can select features which display a statistically significant difference in distributions between both sets. Craig's Zeta, for instance, is an extension of the Zeta metric originally proposed by Burrows (Burrows, 2007), which remains a popular choice in the stylometric community to select discriminative stylometric features in binary classification settings (Craig and Kinney, 2009). An example of another more widely used metric for feature selection in corpus linguistics is the Mann-Whitney ranks test (Kilgariff, 2001). As a dummy example, we can confront the above mentioned texts; be it the novels by Jane Austen and Anne Brontë:

```
data(novels)

corpus.all <- txt.to.words.ext(novels, language = "English.all",
  preserve.case = TRUE)

corpus.austen <- corpus.all[grepl("Austen", names(corpus.all))]
corpus.abronte <- corpus.all[grepl("ABronte", names(corpus.all))]

zeta.results <- oppose(primary.corpus = corpus.austen,
  secondary.corpus = corpus.abronte, gui = FALSE)
```

As can be seen in the results (first 20 most discriminating words), Jane Austen is an enthusiast user of terms related to socio-cultural phenomena (e.g. *situation*, *opinion*, *party*, *engaged*, ...), whereas Anne Brontë's vocabulary can be characterized by a variety of auxiliary verbs with contractions, as well as religious and light-related vocabulary (e.g. *bright*, *dark*).

```
zeta.results$words.preferred[1:20]
```

```
[1] "Her"      "farther"  "behaviour" "opinion"  "party"
[6] "point"    "perfectly" "afterwards" "Colonel"  "directly"
[11] "spirits"  "situation" "settled"    "hardly"   "Jane"
[16] "Emma"    "equal"    "family"    "engaged"  "They"
```

```
zeta.results$words.avoided[1:20]
```

```
[1] "don't"    "I^m"      "I^ll"      "beside"    "Arthur"
[6] "can't"    "I^ve"     "it^s"      "won't"     "Huntingdon"
[11] "presence" "Helen"    "face"      "bright"    "God"
[16] "mamma"    "further"  "heaven"    "dark"      "feet"
```

Of course, the above results of this simple feature selection tool can be fed into one of the package's classification routines:




Figure 4: The Rolling Stylometry visualization. The medieval French allegoric story *Roman de la Rose* assessed using Rolling SVM and 100 MFWs; window size: 5,000 words, sample overlap: 4,500 words. Sections attributed to Guillaume de Lorris are marked red, those attributed to Jean de Meun are green. The level of certainty of the classification is indicated by the thickness of the bottom stripe. The commonly-accepted division into two authorial parts is marked with a vertical dashed line ‘b’.

```
combined.features <- c(zeta.results$words.preferred[1:20],
                      zeta.results$words.avoided[1:20])
stylo(parsed.corpus = corpus.all, features = combined.features, gui = FALSE)
```

Other functionality worth mentioning are `rolling.delta()` and `rolling.classify()`. These functions implement a procedure meant to progressively analyze the development of a style in a text, using e.g. one of the stylometric distance metrics discussed (Rybicki et al., 2014; Eder, in press). In many works, specific parts of the text are conjectured to have been plagiarized or contributed by other authors: `rolling.delta()` and `rolling.classify()` offer an easy way to visualize local stylistic idiosyncrasies in texts. In Fig. 4 we have plotted a `rolling.classify()` analysis of the well-known French allegorical romance *Roman de la Rose* from the Middle Ages. It has been written by two authors: Guillaume de Lorris is the author of the opening 4,058 lines (ca. 50,000 words), and the second part by Jean de Meun consists of 17,724 lines (ca. 218,000 words). This knowledge is supported by the text itself, since Jean de Meun explicitly points out the takeover point (it is marked with a dashed vertical line ‘b’ in Fig. 4). In this example, the aim is to verify whether two authorial styles can indeed be discerned in the text, that is, before and after the authorial takeover. First a Support Vector Machine classifier is trained on four 5,000-word samples: two extracted from the beginning of the text and two near the middle of the text (yet well beyond the hypothesized takeover: they are marked with the dashed line ‘a’ and ‘c–d’, respectively). Next, we apply a windowing procedure and we extract consecutive and partially overlapping samples from the entire text. Finally, the trained classifier is applied to each of these ‘windows.’ In Fig. 4 we plot the respective classification scores for both authors in each sample: in this case, these scores represent the probability, estimated by a Support Vector Machine, that a particular sample should be attributed to one of the two authors involved. Although the result is not flawless, a clear shift in authorial style can be discerned around the position of the takeover, as indicated verbatimly in the text by one of the authors.

The dataset to replicate the test can be downloaded from this page: https://sites.google.com/site/computationalstylistics/corpora/Roman_de_la_Rose.zip. The following code should be typed to perform the classification:

```
# unzipping the dataset
unzip("Roman_de_la_Rose.zip")

# changing working directory
setwd("Roman_de_la_Rose")

rolling.classify(write.png.file = TRUE, classification.method = "svm", mfw = 100,
                 training.set.sampling = "normal.sampling", slice.size = 5000,
                 slice.overlap = 4500)
```

Conclusion

‘Stylometry with R’ targets two distinct groups of users: experienced coders and beginners. Novice users have found it useful to work with the intuitive Graphical User Interface (GUI), which makes it easy to set and explore different parameters without programming experience. We wish to emphasize, however, that **stylo** is useful beyond these high-level functions and GUIs: it also offers experienced users a general framework that can be used to design custom processing pipelines in R, e.g. in other text-oriented research efforts. The current version of **stylo** (version number 0.6.3) is available from GitHub under a GPL 3.0 open-source licence; binary installation files are available from CRAN. **stylo** has been used in a number of innovative studies in the field of computational stylistics (Kestemont et al., 2013; van Dalen-Oskam, 2014; Lauer and Jannidis, 2014; Anand et al., 2014; Oakes and Pichler, 2013; Boot, 2013), and we encourage the future application of **stylo** to challenging new problems and languages in stylometry.

Acknowledgments

We would like to thank the users of **stylo** for the valuable feedback and feature requests which we have received over the past years. MK was partially founded for this research as a postdoctoral fellow by The Research Foundation of Flanders (FWO). ME was partially supported by Poland’s National Science Centre (grant number 2014/12/W/ST5/00592).

Bibliography

- S. Anand, A. K. Dawn, and S. K. Saha. A statistical analysis approach to author identification using latent semantic analysis. Notebook for PAN at CLEF 2014. In *CLEF2014 Working Notes*, pages 1143–1147, Sheffield, UK, 2014. CLEF. [p13]
- S. Argamon. Interpreting Burrows’s Delta: Geometric and probabilistic foundations. *Literary and Linguistic Computing*, 23(2):131–147, 2011. [p3]
- P. Boot. Online boekdiscussie van een afstand gelezen. *TNTL. Journal of Dutch Linguistics and Literature*, 129(4):215–232, 2013. [p13]
- J. Burrows. ‘Delta’: A measure of stylistic difference and a guide to likely authorship. *Literary and Linguistic Computing*, 17(3):267–287, 2002. [p3, 9]
- J. Burrows. All the way through: Testing for authorship in different frequency strata. *Literary and Linguistic Computing*, 22(1):27–48, 2007. [p11]
- H. Craig and A. Kinney, editors. *Shakespeare, Computers, and the Mystery of Authorship*. Cambridge University Press, 2009. [p11]
- W. Daelemans. Explanation in computational stylometry. In *Proceedings of the 14th International Conference on Computational Linguistics and Intelligent Text Processing – Volume 2, CICLing’13*, pages 451–462, Berlin, Heidelberg, 2013. Springer-Verlag. [p2]
- M. Eder. Style-markers in authorship attribution: A cross-language study of the authorial fingerprint. *Studies in Polish Linguistics*, 6:99–114, 2011. URL <http://www.wuj.pl/page,art,artid,1923.html>. [p5]
- M. Eder. Computational stylistics and biblical translation: How reliable can a dendrogram be? In T. Piotrowski and Ł. Grabowski, editors, *The Translator and the Computer*, pages 155–170. WSF Press, Wrocław, 2012. [p8]
- M. Eder. Rolling stylometry. *Digital Scholarship in the Humanities*, in press. doi: 10.1093/llc/fqv010. [p12]
- M. Eder and J. Rybicki. Do birds of a feather really flock together, or how to choose training samples for authorship attribution. *Literary and Linguistic Computing*, 28(2):229–236, 2013. [p11]
- I. Feinerer, K. Hornik, and D. Meyer. Text mining infrastructure in R. *Journal of Statistical Software*, 25(5):1–54, 3 2008. URL <http://www.jstatsoft.org/v25/i05>. [p2]
- E. Gamerman. Data miners dig into ‘Watchman’. *Wall Street Journal*, page D5, July 2015. [p1, 8]
- D. Holmes. Authorship attribution. *Computers and the Humanities*, 28(2):87–106, 1994. [p1]

- D. Holmes. The evolution of stylometry in Humanities scholarship. *Literary and Linguistic Computing*, 13(3):111–117, 1998. [p1]
- D. Hoover. Testing Burrows’s Delta. *Literary and Linguistic Computing*, 19(4):453–475, 2004a. [p3, 4]
- D. Hoover. Delta prime. *Literary and Linguistic Computing*, 19(4):477–495, 2004b. [p3, 4]
- J. Houvardas and E. Stamatatos. N-gram feature selection for authorship identification. In J. Euzenat and J. Domingue, editors, *Proceedings of Artificial Intelligence: Methodologies, Systems, and Applications*, pages 77–86. Springer Verlag, 2006. [p2]
- M. Jockers. *Macroanalysis: Digital Methods and Literary History*. Topics in the Digital Humanities. University of Illinois Press, 2013. [p2]
- P. Juola. Authorship attribution. *Foundations and Trends in Information Retrieval*, 1(3):233–334, 2006. [p1]
- P. Juola. Becoming Jack London. *Journal of Quantitative Linguistics*, 14(2):145–147, 2007. [p2]
- P. Juola. Rowling and “Galbraith”: An authorial analysis, 2013. URL <http://languagelog.ldc.upenn.edu/n11/?p=5315>. [p1, 10]
- M. Kestemont. Function words in authorship attribution: From black magic to theory? In *Proceedings of the 3rd Workshop on Computational Linguistics for Literature (CLFL)*, pages 59–66, Gothenburg, Sweden, April 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/W14-0908>. [p2]
- M. Kestemont, S. Moens, and J. Deploige. Collaborative authorship in the twelfth century: A stylometric study of Hildegard of Bingen and Guibert of Gembloux. *Literary and Linguistic Computing*, 28: 1–15, 2013. doi: doi:10.1093/lilc/fqt063. [p13]
- A. Kilgariff. Comparing corpora. *International Journal of Corpus Linguistics*, 6(1):97–133, 2001. [p11]
- B. Kjell. Discrimination of authorship using visualization. *Information Processing and Management*, 30(1):141–50, 1994. [p2, 5]
- M. Koppel, J. Schler, and S. Argamon. Computational methods in authorship attribution. *Journal of the American Society for Information Science and Technology*, 60(1):9–26, 2009. [p1]
- G. Lauer and F. Jannidis. Burrows’s Delta and its use in German literary history. In M. Erlin and L. Tatlock, editors, *Distant Readings. Topologies of German Culture in the Long Nineteenth Century*, pages 29–54. Rochester, New York, 2014. [p13]
- H. Love. *Attributing Authorship. An Introduction*. Cambridge University Press, Cambridge, 2002. [p1]
- M. Oakes and A. Pichler. Computational stylometry of Wittgenstein’s ‘Diktat für Schlick’. *Bergen Language and Linguistics Studies*, 3(1):221–240, 2013. [p13]
- J. Pennebaker. *The Secret Life of Pronouns: What our Words Say about Us*. Bloomsbury Press, New York, 2011. [p4]
- J. Rybicki and M. Heydel. The stylistics and stylometry of collaborative translation: Woolf’s ‘Night and Day’ in Polish. *Literary and Linguistic Computing*, 28(4):708–717, 2013. [p8]
- J. Rybicki, D. Hoover, and M. Kestemont. Collaborative authorship: Conrad, Ford and rolling Delta. *Literary and Linguistic Computing*, 29(3):422–431, 2014. [p12]
- S. Schreibman, R. Siemens, and J. Unsworth, editors. *A Companion to Digital Humanities*. Blackwell, 2004. URL <http://www.digitalhumanities.org/companion/>. [p1]
- F. Sebastiani. Machine Learning in automated text categorization. *ACM Computing Surveys*, 34(1):1–47, 2002. [p2]
- E. Stamatatos. A survey of modern authorship attribution methods. *Journal of the American Society for Information Science and Technology*, 60(3):538–556, 2009. [p1, 2, 5]
- E. Stamatatos, N. Fakotakis, and G. Kokkinakis. Automatic text categorization in terms of genre and author. *Computational Linguistics*, 26(4):471–495, 2000. [p1]
- C. Stamou. Stylochronometry: Stylistic development, sequence of composition, and relative dating. *Literary and Linguistic Computing*, 23(2):181–199, 2008. [p1]

- J. Stover, Y. Winter, M. Koppel, and M. Kestemont. Computational authorship verification method attributes new work to major 2nd century African author. *Journal of the American Society for Information Science and Technology*, in press. doi: 10.1002/asi.23460. [p8]
- K. van Dalen-Oskam. Epistolary voices: The case of Elisabeth Wolff and Agatha Deken. *Literary and Linguistic Computing*, 29(3):443–451, 2014. [p8, 13]
- L. van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9: 2579–2605, 2008. [p5]
- H. Van Halteren, H. Baayen, F. Tweedie, M. Haverkort, and A. Neijt. New machine learning methods demonstrate the existence of a human stylome. *Journal of Quantitative Linguistics*, 12(1):65–77, 2005. [p1]

Maciej Eder
Institute of Polish Language
Polish Academy of Sciences
al. Mickiewicza 31, 31-120 Kraków
Poland
maciejeder@gmail.com

Jan Rybicki
Institute of English Studies
Jagiellonian University
al. Mickiewicza 9A, 31-120 Kraków
Poland
jkrybicki@gmail.com

Mike Kestemont
Department of Literature
University of Antwerp
Prinsstraat 13, B-2000 Antwerp
Belgium
mike.kestemont@uantwerp.be