THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Computational Linguistics and Natural Language Processing

Citation for published version:

Luz, S 2022, Computational Linguistics and Natural Language Processing. in F Zanettin & C Rundle (eds), *The Routledge Handbook of Translation and Methodology*. Routledge, New York, pp. 373-391.
<<https://www.taylorfrancis.com/chapters/edit/10.4324/9781315158945-27/computational-linguistics-natural-language-processing-saturnino-luz>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

The Routledge Handbook of Translation and Methodology

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Computational linguistics and Natural Language Processing

Saturnino Luz

Introduction and definitions

Broadly defined, the term *computational linguistics* refers to the use of computational methods and tools in the study of linguistic phenomena. A distinction is sometimes made between computational linguistics and *natural language processing*. The former is usually regarded as the study of linguistic ability as a computational process, and the latter as an “engineering” pursuit directed towards the application of algorithmic methods to practical problems such as automatic categorisation of text, parsing, prediction of the part-of-speech (categories) of words, automatic translation, text summarisation and other tasks, all of which involve processing of natural languages (as opposed to formal or programming languages).¹ As the boundaries between computational linguistics and natural language processing are not always clearly defined, we will not concern ourselves with this distinction here. Therefore, our use of the term computational linguistics in this chapter will encompass the methods and tools these fields share in common, with a focus on methodology and its relation to translation practice and enquiry.

As computing systems have become increasingly important in the practice of translation, their use in scholarly studies of translation and its practice have followed a similar path. A typical example of this trend is the consolidation of corpus-based translation studies as an active area of research (Baker 1993). As with all data intensive, corpus-based linguistic studies, the applications of corpora in the area of translation studies owe their existence to computational tools and resources, mirroring the evolution of translation practice (Karamanis et al. 2011).

While computational linguistics is usually associated with quantitative or formal approaches, its methods have also been used to support qualitative analysis. In fact, it can be argued that the application of computational linguistics to corpus analysis, and corpus-based translation studies in particular, is generally based on an interplay of qualitative and quantitative techniques (McEnery and Hardie 2011). The combined use of concordancing (a qualitative method) and word frequency analysis (a quantitative method) provides a well known example of such interplay. However, more intricate methods and combinations exist and continue to appear as the field of computational linguistics evolves. In the following sections, we review the main developments in the field, the tools they produced, and the uses these tools have found in translation and methodology.

Before we proceed, however, a few basic definitions are in order. In addition to the central concept of computational linguistics, we will often refer to related concepts such a corpus, metadata, corpus

linguistics, corpus based translation studies, statistical methods, machine learning, and text visualisation. A *corpus*, in this context, is simply a collection of texts often accompanied by data that describe and categorise the texts in the collection, commonly referred to as *metadata*, and sometimes complemented by resources of a non-textual nature, such as images, recorded speech and video. We will limit ourselves here to corpora comprising text and metadata only. *Corpus linguistics*, sometimes referred to as *data-intensive linguistics* (Church and Mercer 1993), is the use of corpora, in digital form, aided by computing technology, for the study of linguistic phenomena. *Corpus based translation studies* employs and extends corpus linguistics methods in the study of translated text (Laviosa 2004) with the aim of uncovering the particular “nature of translated text as a mediated communicative event” (Baker 1993). The data-intensive nature of corpus linguistics implies a preponderance of *statistical methods*, which encompass both descriptive (e.g. word frequency, co-occurrence counts) and inferential statistics (e.g. inference of collocation patterns from co-occurrence probabilities). Traditional statistical methods are increasingly being complemented by *machine learning* methods. Machine learning is a sub-discipline of artificial intelligence which aims to automate the processes of creating structured representations from raw data (e.g. representing text as feature vectors) and performing inference on these representations (e.g. classifying text into high-level categories, labelling a sequence of words with the corresponding sequence of grammatical categories, parsing sentences into tree structures). Finally, the analytic framework of corpus linguistics, and corpus-based translation studies, includes descriptive tools that enhance the ability of the translation scholar to inspect large volumes of text data interactively, with the help of a visual computer interface. These tools will be referred to as *text visualisation* tools, and include the familiar concordance list, as well as a number of new visual presentations that have been developed more recently.

Methods, Tools and Users

Computational linguistic methods can be categorised as *symbolic methods* and *statistical learning methods*. The former consists of approaches that originated from the formal logic and symbolic artificial intelligence (AI) tradition. The latter have their roots in probability theory, statistics and connectionist (neural network) approaches. While handcrafted symbolic parsers and machine translation systems have all but disappeared, symbolic approaches still play a role in an area that is relevant to translation, namely, semantics, where large scale ontologies and lexical databases such as Wordnet (Miller 1995) have been employed in corpus linguistics (Budanitsky and Hirst 2006) as well as in knowledge based machine translation systems (Costa-Jussa and Fonollosa 2015).

However, it is statistical methods that currently dominate the contribution of computational linguistics to translation methodology, studies, and practice. The above mentioned techniques of concordancing and frequency list comparison can be regarded as statistical methods. Relative frequencies are clearly statistical in nature, and text concordances are a form of descriptive statistics

combined with a text visualisation device, as can be seen clearly in the abstract representations proposed by Luz and Sheehan (2014), and by Wattenberg and Viégas (2008), for instance. In addition to these basic methods, other statistical methods have long been used in corpus linguistics (Sinclair 1991) and translation studies.

More recently, machine learning methods have become closely linked to computational linguistics, and their influence is also evident in corpus-based translation studies. Automatic classification methods (Emms and Luz 2007; Sebastiani 2002) have been used, for instance, in the identification of linguistic and stylistic patterns that might be characteristic of translated text, or “translationese” (Baroni and Bernardini 2005). Machine learning has also extended the repertoire of distributional semantics methods beyond simple comparison of statistical and information theoretic measures, providing this area with sophisticated modelling tools which allow the corpus analyst to automatically identify textual patterns not easily spotted by visual comparison alone. These tools include latent semantic analysis (Deerwester et al. 1990), latent Dirichlet allocation for topic modelling (Blei 2012; Blei et al. 2003), and word embedding models such as word2vec (Mikolov et al. 2013) and t-distributed stochastic neighbour embedding, t-SNE (Van der Maaten and Hinton 2008). Machine translation is also a field in which there have been significant improvements in recent years (Wu et al. 2016) due to advances in the area of deep learning (LeCun et al. 2015).

Underlying most of these methods, is an infrastructure for pre-processing the corpora before they are analysed. This involves collecting, storing, indexing and providing access to language resources. The word-wide web has been used both as a source of such data (Baroni et al. 2009) and as a means of providing access to corpora (Luz 2011).

We will now describe these different facets of computational linguistics in translation methodology and study, from basic infrastructure issues and tools to advanced techniques and methods.

Building and managing a corpus

Gathering, storing, indexing and managing access to a collection of electronic texts are usually the first steps in any type of corpus based study. A number of tools have been developed over the years to support these tasks, sometimes individually, as in the case of UNIX directory structures and command line tools for search and filtering of text (Schmitt et al. 2007), sometimes as part of larger general-purpose packages such as the indexing tools provided by the Xapian project², and sometimes as specialised corpus linguistics tools. While reviewing the first two kinds of tools falls outside the scope of this chapter, and an exhaustive review of specialised corpus linguistics tools would be impractical, we describe three tools used for corpus management and analysis which illustrate the main characteristics of the various existing tools and their differences.

Concerning their differences, existing tools can be grouped into broad categories on the basis of their mode of distribution, the way they store and provide access to corpora, and their licensing

terms. Corpus software can be distributed through conventional stand-alone software packages, through web-services and on-line interfaces, or by adopting a mixed web-based mode of distribution, with specialised software “clients” (the software component running on the end-user’s computer) interacting with one or more “servers” (centralised software running on a remote computer).

One of the most popular tool sets distributed in the conventional way is WordSmith Tools (Scott 2021). WordSmith Tools is a robust and stable set of tools which has been used in a number of papers in lexicography and translation studies. However, it only runs on Windows platforms, which limits its accessibility. Other similar concordance software such as MonoConc and ParaConc (Barlow 1999) are also limited in this respect. Web-based applications mitigate this limitation by allowing access from different platforms and operating systems. Possibly the best-known example of web-based corpus tools is the Sketch Engine (Kilgarriff et al. 2014). This tool, which runs entirely on web browsers, has had a significant impact on corpus and lexicography research. Finally, an example of a mixed distribution mode is the modnlp/tec tool (Luz 2011), which is distributed using Java Web Start technology and runs as a stand-alone client on the user’s computer, typically interacting with a remote server. The modnlp software suite runs on most computer operating systems, and has been used extensively in translation studies, having been initially developed for the translational English corpus, TEC (Baker 1999), and having since been used in a number of other applications, most recently the Genealogies of Knowledge project (Jones 2019; Luz and Sheehan 2020).³

In terms of the way they provide access to corpora, WordSmith Tools only allow the user to index and access data which is stored on his/her computer. They are therefore typically used with smaller corpora which are not subject to copyright restrictions. Sketch Engine is used mainly for access to large, pre-indexed, web-stored corpora, but also allows users to load and index their own corpus, as well as define new web-based corpora. The modnlp/tec tool is capable of both accessing corpora stored on the user’s computer, as well as remote web-based corpora.⁴

These tools also differ with respect to the licenses under which they are distributed. Both WordSmith tools and the Sketch Engine are commercial proprietary software.⁵ WordSmith is distributed under a proprietary license as executable binary code. No access to source code is provided. Sketch Engine is fully web based, and operates on a subscription basis. As in the case of WordSmith Tools, no access to source code is provided with the Sketch Engine software, although a version called NoSketch Engine has been released recently as open source which provides limited functionality. The modnlp suite on the other hand is free software (Stallman 2002) and is distributed under the terms of the GNU general public license, a free/libre and open source software (FLOSS) license.

Common functions provided by these and most other corpus management tools include tokenisation, indexing, and data presentation. Wordsmith tools and modnlp have built-in indexing capabilities, and the latter can either access a pre-built index via the network, or create a local index, using different tokenisers (e.g. the tokenisers created by the Apache Foundation’s Lucene project, the Stanford

tokeniser for Arabic, etc.) as needed. Sketch Engine only provides access to indexing via its web interface; the index is otherwise inaccessible to the user. While both Wordsmith Tools and Sketch Engine have basic features for storing metadata, modnlp provides comprehensive support for combining text search and metadata constraints. Metadata files in modnlp are encoded in XML (Luz 2011) and allow the user to select sub-corpora according to common metadata attributes, such as author, source language, genre etc.. Finally, data presentation minimally covers concordance, collocation and frequency lists. Presentation techniques are described below.

Data presentation and visualisation

The concordance list is perhaps the oldest and most basic form of data presentation for corpus analysis. It consists of arranging passages of a text or collection of texts in alphabetical order according to user-defined keywords. Combined with a “keyword-in-context” (KWIC) indexing method, which aligns different occurrences of the keyword specified by the user at the centre of each concordance line, and with interactive capabilities for sorting and rearranging the list, concordancing is a powerful text presentation method.

Visual generalisations of concordances have been proposed, which seek to address one of the main limitations of this technique, namely its poor use of space, since a concordance could comprise several thousand lines of text. Wattenberg and Viégas (2008) proposed a form of visual encoding of concordances called Word Tree, which displays alternatively the left or the right context of a concordance as a tree where words are vertices linked in textual order and scaled in size according to their frequencies. Although they provide a better visual summary than concordance lists, Word Trees limit the display to half of the text (the keyword plus the left or right context) of its underlying concordances. It is therefore impossible for the user to read the full sentences in which the keyword appears. For certain corpus linguistics tasks, such as detection of phrases that span left and right contexts (as in the expression ‘run the whole *gamut* of ...’, for instance), frequency information for words occurring on each context is usually more useful to the analyst than the linear structure of a single context (Sinclair, 2003). A mosaic style visualisation has been proposed which overcomes this limitation (Luz and Sheehan, 2014). This “*concordance mosaic*” uses a tabular structure which preserves the relative position of each word and scales the rectangles they occupy proportionally to word occurrence probabilities or collocation statistics. Interactive restructuring of a concordance is enabled through the interface. This restructuring combined with color highlighting of the concordance lines creates a powerful technique for investigating significant collocation patterns. The space filling design allows for perceptually efficient comparison of word collocation statistics such as mutual information, z-score and log-log score, while preserving some of the context available on a concordance list. Figure 24.24.1 shows two examples of concordance mosaics for the word ‘regime’, showing the effect of using a collocation metric to re-scale a full concordance summary.

Figure 24.1: Raw frequency distribution view (top) and collocation strength view (bottom) for the word “regime” in the Genealogies of Knowledge modern corpus. The collocation strength view is scaled according to mutual information statistics and reveals patterns not readily evident in the frequency distribution

Frequency lists enable quick comparisons between words and terms for a particular corpus, collocation or sub-corpus. In translation studies, it is common for scholars to compare several lists from different sub-corpora. In studying translation styles one might wish, for example, to compare sub-corpora corresponding to translations by a given translator, from a given source language, etc. As the frequency distributions for different sub-corpora will typically vary, comparisons of words in the same rank order in different lists sometimes may be misleading. For this reason, modnlp/tec introduced a tool which enables visual comparisons of frequency lists using slope lines to connect words on the same relative positions of logarithmic scaled axes (Sheehan et al. 2018) which corrects any distortions due to differing text sizes.

A less commonly encountered feature of corpus tools is the visualisation of concordances in parallel corpora. While this feature is of great relevance to translation and translation studies, visualising source and target translated texts side by side disrupts the visual scanning of concordance patterns.

Perhaps for this reason, this form of display is not commonly used outside specialised localisation tools (Lewis et al. 2009). However, Sketch Engine and modnlp have add-on modules for basic parallel concordancing, while Wordsmith Tools handles alignment and parallel concordancing through utility programs. Paraconc is one of the few corpus tools designed specifically for multilingual parallel corpora.

Basic computational linguistics methods

Some of the functions of the tools described above, and others like them, draw mostly on the notion of *collocation* (Sinclair 1991). Collocations are words that occur together often enough to be noticed as textual units, being often referred to in the context of translation as *terms* (though in the broader context of corpus research not all collocations would be regarded as terms). In this sense, terminology extraction is a task related to translation which employs computational linguistic methods. In addition to descriptive methods, such as concordancing, quantitative analytic methods are employed to determine how often is often enough as a pattern of word co-occurrence. The basic framework is one of hypothesis testing, where one wishes to assess whether a given co-occurrence pattern forms a collocation or is a merely “accidental” co-occurrence (Manning and Schütze 1999). Accidental co-occurrence has been equated with statistical independence, in a probabilistic framework. Thus, the null hypothesis in a test of statistical significance is formulated as the statement that $P(x,y)=P(x)P(y)$ where $P(x,y)$ is the (joint) probability that words x and y appear together in a given context, while $P(x)$ is the (prior) probability of occurrence of word x . Different statistical tests have been used in the corpus analysis literature. The most common approach is to employ the t-test, where the t statistic is computed to determine the probability that the difference between the expected co-occurrence distribution (modelled as a Bernoulli distribution) and the empirically observed distribution occurred by chance. The null hypothesis is rejected if this probability is smaller than a set value, usually 0.05. Other tests have been proposed to account for the fact that word probabilities are generally not normally distributed (an implicit assumption of the t-test). Among the non-parametric tests used, the most common is Pearson’s chi-squared test. This test is of particular interest in translation as it has also been employed to identify translation pairs in parallel texts. To identify word associations, for instance, a contingency table containing the frequencies of the four possible co-occurrences of a candidate pair (say, the English word ‘house’ and the French word ‘chambre’) is created, and the chi-squared test or similar can be applied to determine the strength of their association (Gale and Church 1991).

Other statistical tests and information theoretic measures of word association used in translation and corpus analysis include the z-score, likelihood ratios, information gain, and mutual information. The last two can be regarded as information theoretic measures, as they are motivated by properties of distributions. The mutual information score has been widely employed.⁶ It essentially computes the

following function of the word frequencies: $PMI(x, y) = \log \frac{P(x, y)}{P(x)P(y)} = \log \frac{P(xy)}{P(x)}$. PMI thus gives a measure of association; the larger its value, the greater the dependency between x and y (and the likelihood that they will form collocations) in the text. This metric is, in fact, the basis of the re-scaling of words on the second (lower) mosaic shown in Figure 24.24.1. The words are scaled according to a metric $M(w, k) = \frac{\text{count}(w, k)}{N} \times \frac{\sum \text{count}(x)}{\text{count}(w)}$ which is the ratio of the relative frequency of word w in the context of word k to the relative frequency of k in corpus. It can be easily shown that this scaling factor is closely related to mutual information, that is, $M(w, k) \propto 2^{PMI(k, w)}$.

Machine learning

In the last few decades, work in computational linguistics has been increasingly dominated by machine learning approaches. Machine learning is an umbrella term for a number of methods which aim to detect patterns and regularities in data by automatic algorithmic means. According to Mitchell (1997: 2), a learner is a system whose performance ‘with respect to some class of tasks T and performance measure P ’ improves as the system is exposed to experience E . In the case of computational linguistics, the tasks range from classification of texts, words or segments, to labelling and mapping sequences to sequences. Performance measures include accuracy, precision, recall, and F scores (among other measures), and the experience (data) corresponds to the corpus itself, or part of it, presented to the learning algorithm as a formal object, which we will refer to as its *representation*.

The learning task is usually conceptualised as a function approximation task: given a true function f , the system must learn an approximation g of that function. The exact nature of the function to be learnt depends on the task. In a clustering task, for instance, the function to be approximated could map individual text documents to sets of documents. A classification task might map texts to sets of labels, a sequence-to-sequence mapping task might map sequences of words to sequences of parts of speech, or sequences of words in a source language to sequences of words in a target language, and so on.

Of the categorisation tasks, the most relevant to translation practice and studies are word category disambiguation (part-of-speech tagging) and word sense disambiguation (Jurafsky and Martin 2008). In category disambiguation, a word is assigned a grammatical category (a part-of-speech tag), for instance, the word “word” may be assigned the category of noun in most contexts or the category of verb in a few contexts. Corpus analysis tools such as those described above often allow the user to specify a category for the keyword in order to restrict the search. Such categories are typically assigned to the indexed text automatically, through machine learning algorithms. In sense disambiguation, words or terms are assigned semantic labels. Although search by word senses is less common in corpus tools, it is often useful in machine translation, and computer-assisted translation.

Both tasks can be conceptualised, alternatively, as word-to-category or sequence-to-sequence mappings. Machine learning algorithms that have been employed to learn word-to-category mappings of parts of speech or sense tags, include maximum entropy (Ratnaparkhi 1996) and decision trees (Màrquez and Rodríguez 1998) among others. Sequence-to-sequence models include hidden Markov models (Cutting et al. 1992), conditional random fields (Lafferty et al. 2001), and more recently, “deep learning” models have been proposed for multilingual tagging (Plank et al. 2016).

Machine translation systems often make use of the results of categorisation tasks, and most modern systems incorporate additional machine learning methods. This is true of both statistical machine translation (Lopez 2008) and of hybrid approaches which incorporate corpus and rule-based components (Costa-Jussa and Fonollosa 2015). Deep learning methods have also had a significant impact on machine translation performance (Sutskever et al. 2014) and continue to be an active field of research (Neubig and Watanabe 2016).

A distinction is sometimes made between supervised and unsupervised learning, meaning learning tasks where the “experience” (in the sense of Mitchell’s definition) is guided by human feedback (often in the form of annotation of the training data), or solely by features of the data themselves. The tasks described above might be regarded as supervised tasks. Typical unsupervised tasks in computational linguistics include the clustering of terms or documents into homogeneous sets. Unsupervised tasks are commonly used for the extraction of input representations for supervised tasks (Emms and Luz 2007), for example, in the creation of distributed representation of word semantics as in the case of word embeddings and latent semantic indexing, or directly in inference tasks such as word category and sense disambiguation. In the case of distributed word representations, it is common to represent words as vectors, so that a collection of documents can be represented as a co-occurrence matrix where each row corresponds to a numerical vector representing a word, as shown in Table 24.1. This is a $v \times v$ matrix where v is the size of the vocabulary and each entry (m,n) corresponds to the number of times word m occurs in the same document as word n . In Table 24.1, for example, the word ‘usair’ co-occurs with ‘voting’ twice. Based on this representation, a simple clustering of these terms (using, for instance, the unsupervised k-means algorithm) would create separate clusters for {“stake”}, {“usair”, “merger”, “twa”}, {“acquire”, “acquisition”}, and {“voting”, “buyout”, “ownership”, ...} which intuitively appear to share similar semantics.

	usair	voting	buyout	stake	santa	merger	ownership	...	manufactures	...	attractive	undisclosed	aquisition	twa	interested
usair	20	2	0	1	0	4	1	...	0	...	0	0	2	14	1
voting	2	10	0	2	0	1	0		0		0	0	0	2	0
buyout	0	0	8	1	0	2	0		0		1	0	1	0	0
stake	1	2	1	62	0	0	0		0		0	0	0	1	0
santa	0	0	0	0	7	3	0		0		0	2	0	0	1
merger	4	1	2	0	3	48	0		0		1	0	5	4	3
ownership	1	0	0	0	0	0	6		0		0	0	0	1	0
manufactures	...	0	0	0	0	0	0		4		0	0	0	0	0
attractive	...	0	0	1	0	0	1		0		4	0	0	0	0
undisclosed		0	0	0	0	2	0		0		0	24	4	0	0
aquisition		2	0	1	0	0	5		0		0	4	60	2	0
twa		14	2	0	1	0	4		0		0	0	2	20	1
interested		1	0	0	0	1	3		0		0	0	0	1	6

Table 24.1: Sample co-occurrence matrix for a subset of REUTERS-21578 (Emms and Luz, 2007)

The centroid of a cluster (i.e. the average values for each entry in the word vectors in the cluster) could then be taken to be word representations to be used in text classification and other analytical tasks.

Distributional semantics

The example above illustrates the basics of the *distributional semantics* approach, which has its roots in corpus linguistics and early machine translation research (Weaver 1955), and has been influential in the application of computational linguistics to translation.

Distributional semantics posits that the semantics of a word is determined by its *context*, that is, the words that occur in its vicinity. An early computational realisation of this idea is the *vector space model*, commonly used in information retrieval (Salton et al. 1975; Baeza-Yates and Ribeiro-Neto 1999) where a textual entity such as a word, a document or a text segment is represented as a vector in a common algebraic model. As mentioned before, each row (or, equivalently, each column) in the co-occurrence matrix of Table 24.1, for instance, can be regarded as a vector representation of a word. Thus, the word ‘usair’ may be represented by vector $v=(20,2,0,1,...)$, where each value

indicates the number of times ‘usair’ occurs in the same context as other words in the lexicon. This type of representation is also known as *bag of words* in text categorisation (Sebastiani, 2002).

The vector space model has served as a basis for a number of feature extraction methods. Feature extraction is the process of reducing the dimensionality of the vector space by projecting it onto a small number of dimensions. Each projection onto a lower dimensional space is regarded, in this framework, as a semantic representation of a word. Different methods exist which compute such projections. The most commonly used methods are principal component analysis (PCA) and latent semantic indexing (LSI). Both methods convert the original data (in our case vector representations of words, typically consisting of hundreds or thousands of co-occurrence counts) onto much more compact vector representations by applying linear algebraic operations which seek to maximise the variance of the new, compact representation, and thus preserve as much of the information contained in the original representation as possible. Manning and Schütze (1999: ch. 15) describe the use of LSI to project a high dimensional vector space of words onto a reduced space of independent dimensions by applying an operation known as singular value decomposition to the original document matrix. The intuitive appeal of these methods can be illustrated by a simple projection of the word vectors of Table 24.1 onto the vector space's two dimensions (defined by the first two principal components), as given by singular value decomposition (Figure 24.24.2). While there is considerable sparsity, it can be seen that the model succeeds in placing a number of semantically related words in close proximity on the 2-dimensional plane (e.g. 'usair' and 'twa' are both airlines, the vectors for 'voting', 'ownership', and 'acquisition' are in close proximity, and so on).

Figure 24.2: Projection of a set of REUTERS-21578 words (Table 24.1) onto two dimensions, by principal component analysis.

A recent development in this field is the combination of the vector space model with neural networks for language modelling (Turian et al. 2010). A *language model* is a probabilistic model which encodes the probability of occurrence of a word in a context, most often for predicting the next word given a sequence of words. In neural network implementations of such models, the network's inputs are word vectors and the learning task is to determine the optimal parameters for the model (the word sequence probabilities). The input vectors are mapped onto a lower dimensional space through a network layer called the *embedding layer*. In addition to providing compact and effective representations for machine learning, these embeddings have interesting semantic properties from a distributional semantic perspective. In an often-cited example, Mikolov et al. (2013) show how the vector representations for 'king', 'queen', 'man' and 'woman', learnt from a corpus, can be combined through algebraic operations so that 'king - man + woman' results in a vector that is very near the vector for 'queen' in the lower dimensional space yielded by the embedding. Embedding techniques have since been decoupled from their original application to language modelling and become an area of research in itself, with applications in several areas of computational linguistics, including translation (Almeida and Xexéo 2019). Similar techniques have also been applied to produce cross-lingual representations –that is, word embeddings comprising a common semantic representation for words in different languages– from aligned words (Luong et al. 2015) or sentences (Hermann and Blunsom 2013).

Contexts of application

Contexts of application for computational linguistic methods include, besides stand-alone machine translation systems, the use of corpora for translator training and education (Zanettin et al. 2014), the use of comparable and parallel texts, concordancing and machine translation tools in translation practice (Doherty et al. 2012), the use of corpora and text visualisation tools in the analysis of the linguistic behaviour of translators (Luz 2011; Baker 1999), and the use of machine learning and text categorisation methods for the characterisation of translationese (Baroni and Bernardini 2005; Ilisei et al. 2010), as mentioned above.

Machine translation has been an important area of application for many of the techniques and tools described above. While rule-based machine translation systems have benefited from these tools and techniques, it is in the area of statistical machine translation that they have had the greatest impact. As mentioned, methods such as word category and sense disambiguation are often incorporated into these systems to improve performance (Carpuat and Wu 2007), and conversely new developments in neural machine translation have recently been used in disambiguation tasks (Gonzales et al. 2017). As pointed out earlier, new computational linguistics methods such as cross-lingual embeddings,

which can be induced automatically from parallel (and sometimes from comparable) corpora have also contributed semantically compelling representations of terms across language pairs.

Machine translation is being increasingly used for “gisting” (understanding the essential message of a text) both in informal multi-lingual communication and in commercial settings (Koponen and Salmi 2015). Machine translation has also been incorporated into the work of professional translators and post-editors, often with unexpected implications for their established work practices (Karamanis et al. 2011; Moorkens and O’Brien 2017) and their collaboration with colleagues in the workplace (Doherty et al. 2012).

Corpus management tools have also been extensively used in commercial translation settings, more commonly so than machine translation. The use of concordances as a means of selecting a suitable translation among a set of alternatives is a feature of most “localisation” systems (Lewis et al. 2009; Moorkens and O’Brien 2017). Corpora and corpus software also play an important role in translator education. Tools such as concordance browsers, statistical analysis packages, and visualisation software form an important part of practical translator training. See articles in the collection edited by Zanettin et al. (2014) for different perspectives and approaches in this area. Technologies such as machine translation are also being incorporated into the curriculum (Doherty and Kenny 2014).

Finally, it is indisputable that the methodological foundations of corpus-based translation studies are firmly based on the capabilities afforded by computational processing of large volumes of text. Work in this discipline was strongly influenced by the work of John Sinclair (Sinclair, 2003, 1991), and the analyses based on the translational English corpus through the use of the TEC corpus tools (Luz 2011) greatly contributed to the methodology of some of the field’s seminal work (Baker 1999; 2004). Methods of computational linguistics that rely strongly on machine learning are used less often in translation studies, but the work of Baroni and Bernardini (2005) and Ilisei et al. (2010), in which the authors employed machine learning models to identify characteristic features of translated texts, are interesting exceptions.

Critical issues and topics

As computational linguistics and machine learning methods make further inroads into translation and translation studies, potential issues arise which need to be examined from a methodological perspective. These issues include: the need to support the translator and the translation scholar with tools and user interfaces that enable usable and effective access to large volumes of text and facilitate the selection and comparison of sub-corpora, the interpretability of the statistical and connectionist models employed, and issues concerning the generation and validation of hypotheses and conclusions reached through the use of computational linguistics methods.

Clearly, good user interface support for users of computational linguistics tools is crucial in applications such as translator training and computer-assisted translation. With the advent of corpus-

based translation studies, it has also become important to provide the research community with usable tools which will also allow researchers to document and share their work. Tools such as the Sketch Engine (Kilgariff et al. 2014), the BYU corpora (Davies and Fuchs 2015; Davies 2010), the CQPweb corpus analysis system (Hardie 2012), and the modnlp/tec (Luz 2011; Luz and Sheehan 2020) software suite can be seen as efforts towards achieving these goals. However, the use of analytic tools and corpora is still hampered by software access and licensing constraints. With the exception of modnlp/tec, all of the above mentioned tools are exclusively web-based. While using the web as a platform certainly facilitates access, the absence of a stand-alone, offline tool limits the more experienced user's flexibility and their ability to explore language resources available to them privately, which they might not wish or have the legal right to share. Another issue concerns licensing terms which might limit access to software source code, or prevent it entirely. The ability to inspect and modify source code has been increasingly regarded as a crucial aspect of reproducibility in data-intensive research (Hutson 2018). If corpus-based studies are to develop a robust methodology for the use of computational linguistics tools and methods, the issue of sharing source code as well as data needs to be addressed.

A related issue is the availability and stability of corpus analysis software. Users of purely web-based tools are entirely dependent on the software provider for their analytic work. If the tool's underlying algorithm changes, or the tool is withdrawn from public access, the corpus scholar potentially faces the situation of having their analyses invalidated (in the case of algorithm changes) or uncorroborated (in the case of access withdrawal). While several web-based text visualisation tools have appeared recently⁷, these tools are often prototypes built to demonstrate new visualisation approaches. Ideally, standardised FLOSS platforms will be built in the future which will allow corpus and translation researchers to document and share their analyses, perhaps along the lines of what has been done for "vernacular visualisation" (Viégas and Wattenberg 2008), where users are encouraged to publish their analyses along with the text visualisations.

As machine learning models and algorithms start to become part of the translator's and translation scholar's toolbox, the issue arises of ensuring that these models and algorithms are well understood and properly used. This is not a trivial issue, as is confirmed by the concerns in other traditionally data-intensive research communities regarding the misuses and misinterpretations of simple statistical tests. As machine learning models tend to be a lot more complex, and in many cases more opaque than the models used in traditional statistical testing, their potential misuse should be a cause for concern. For instance, while the interpretation of a co-occurrence matrix is straightforward, an embedding vector which might involve non-linear transformations of the original data would not be as easy to interpret. Therefore, somewhat in opposition to the need we just discussed of providing usable interfaces, it seems that ensuring the proper use of computational linguistics methods based on machine learning will require a level of proficiency in these methods which cannot be achieved simply through better user interfaces. This is, however, an active area of research, and new

techniques of visualisation and model explanation are being proposed which might alleviate this problem in future. A related issue concerns the incorporation of new technology into translator education. Doherty and Kenny (2014) argue for ways of incorporating statistical machine translation into the translation training curriculum in a way that promotes a good understanding of the underlying technology and therefore empowers the translator.

From education to scholarly studies, it is clear that the arrival of new language technologies has disrupted established practices, so that emerging fields such as corpus-based translation studies now find themselves at a stage where further progress in methodology can only be made through the joint efforts of researchers from several disciplines, including translation scholars, linguists, statisticians and computer scientists. This state of affairs was anticipated to some extent (Baker, 2000), but as the field evolves the need for interdisciplinary collaboration is becoming increasingly evident.

Recommendations for practice

Although there are several implementations of computational linguistics tools, many of which have been released as FLOSS software, which could, in principle, be used in translation and translation studies projects, integrating such tools into usable and coherent tool sets for use by translator scholars can be challenging. Most of these tools require the user to have at least basic programming skills, and in many cases considerably more advanced skills. Access to suitable corpora and other data resources is also often an issue. Fortunately, however, there have been developments and efforts towards standardisation and methodological consolidation (Zanettin 2014), which may ease the burden of translators and translation scholars wishing to use computational linguistic tools.

Based on the developments, tools and methodological challenges discussed in this chapter so far, we can now summarise some recommendations for practice regarding the use of computational linguistics in translation work and research:

- In undertaking corpus-based work, consider the accessibility and potential restrictions on sharing of linguistic material. Acquisition and use of such resources is often a demanding and time-consuming first step in both research and commercial translation.
- Choose usable computing tools which give you the necessary flexibility with regards to use and management of your own corpora, in addition to online corpora, and which allow the user to progress naturally from simpler to more complex modes of analysis.
- Learn the basic underlying theory of the computational linguistics methods used in your analysis or work, in order to ensure that these methods are adequately used. This is important in any data analysis task, but specially important in the analysis of translation, where data sparsity and complex, ‘black box’ models often combine to produce invalid results. Learning basic programming skills may also save the corpus and translation scholar a lot of time and frustration.

- Consider the reproducibility of research results. Whenever possible, prefer FLOSS software which allows other researchers to closely inspect your methods and replicate your results. Prefer open, well-documented standards for text encoding, storage, and access.
- Monitor the literature for new developments in text visualisation technology. This is a fast-moving field which is gaining increasing importance with the widespread use of neural networks and other machine learning models in computational linguistics. Text visualisation tools can act as an effective complement to computational linguistics methods. For the translation researcher, this is important as visualisation tool may help the user gain an accurate overview of the data and formulate research hypotheses. For the practitioner, these tools have the potential to improve user interface, and the user's overall experience of the translation process.

Further reading

Manning, Christopher D. and Hinrich Schütze (1999) *Foundations of Statistical Natural Language Processing*. Cambridge, Massachusetts: The MIT Press.

This book provides a good overview of the basic methods in computational linguistics, including the basic statistical concepts that underpin the main methods.

Jurafsky, Daniel and James H. Martin (2008) *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition* (2 ed.). New Jersey: Prentice Hall.

This is a widely used natural language processing textbook. It describes a large number of techniques used in natural language processing, including knowledge based methods.

Mitchell, Tom M. (1997) *Machine Learning*. New York: McGraw-Hill.

This is a very readable introduction to machine learning, covering the machine learning foundations of many computational linguistics and natural language processing methods.

Lopez, Adam (2008) "Statistical machine translation." *ACM Computing Survey* 40(3), 1–49.

This survey article presents a comprehensive overview of statistical machine translation technology. It can be read together with the paper by Cost Jussa and Fonollosa (2015), which presents the latest trends in systems that combine rule-based and statistical methods in machine translation.

Almeida, Felipe and Geraldo Xexéo (2019) "Word embeddings: A survey." *arXiv preprint arXiv:1901.09069*.

The latest developments in vector representations of text such as word embeddings are concisely reviewed in this paper, complementing the extensive survey of cross-lingual models is presented by (Ruder et al., 2017).

Cao, Nan and Weiwei Cui (2016) *Introduction to text visualization*. New York: Springer.

This book provides an introduction to methods for visualisation of text and language models. Work in this field overlaps with the field of human-computer interaction, where studies of translator work and collaboration in translation, and the relation between computational linguistics and qualitative methods have been published (Karamanis et al. 2011; O'Brien 2012).

References

Almeida, Felipe and Geraldo Xexéo (2019) "Word embeddings: A survey." *arXiv preprint arXiv:1901.09069*.

Baeza-Yates, Ricardo and Berthier Ribeiro-Neto (1999) *Modern Information Retrieval*. New York: ACM Press.

Baker, Mona (1999) "The role of corpora in investigating the linguistic behaviour of professional translators." *International Journal of Corpus Linguistics* 4(2), 281–298.

Baker, Mona (2000) "Towards a methodology for investigating the style of a literary translator." *Target. International Journal of Translation Studies* 12(2), 241–266.

Baker, Mona (2004) "A corpus-based view of similarity and difference in translation." *International Journal of Corpus Linguistics* 9, 167–193(27).

Baker, Mona et al. (1993) "Corpus linguistics and translation studies: Implications and applications." *Text and technology: In honour of John Sinclair* Amsterdam: John Benjamins.

Barlow, Michael (1999) "Monoconc 1.5 and paraconc." *International journal of corpus linguistics* 4(1), 173–184.

Baroni, Marco and Silvia Bernardini (2005) "A new approach to the study of translationese: Machine-learning the difference between original and translated text." *Literary and Linguistic Computing* 21(3), 259–274.

Baroni, Marco, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta (2009) "The WaCky wide web: a collection of very large linguistically processed web-crawled corpora." *Language resources and evaluation* 43(3), 209–226.

Blei, David M (2012) "Probabilistic topic models." *Communications of the ACM* 55(4), 77–84.

Blei, David M, Andrew Y Ng, and Michael I Jordan (2003) "Latent Dirichlet allocation." *Journal of machine Learning research* 3, 993–1022.

Budanitsky, Alexander and Graeme Hirst (2006) "Evaluating wordnet-based measures of lexical semantic relatedness." *Computational Linguistics* 32(1), 13–47.

Cao, Nan and Weiwei Cui (2016) *Introduction to text visualization*. New York: Springer.

Carpuat, Marine and Dekai Wu (2007), “Improving statistical machine translation using word sense disambiguation.” *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, New York: ACL Press.

Church, Kenneth W and Robert L Mercer (1993) “Introduction to the special issue on computational linguistics using large corpora.” *Computational linguistics* 19(1), 1–24.

Costa-Jussa, Marta R and José AR Fonollosa (2015) “Latest trends in hybrid machine translation and its applications.” *Computer Speech & Language* 32(1), 3–10.

Cutting, Doug, Julian Kupiec, Jan Pedersen, and Penelope Sibun (1992), “A practical part-of-speech tagger.”. in *Third Conference on Applied Natural Language Processing* New York. ACL Press.

Davies, Mark (2010) “More than a peephole: Using large and diverse online corpora.” *International Journal of Corpus Linguistics* 15(3), 412–418.

Davies, Mark and Robert Fuchs (2015) “Expanding horizons in the study of world englishes with the 1.9 billion word global web-based english corpus (GloWbE).” *English World-Wide* 36(1), 1–28.

Deerwester, Scott, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman (1990) “Indexing by latent semantic analysis.” *Journal of the American society for information science* 41(6), 391.

Doherty, Gavin, Nikiforos Karamanis, and Saturnino Luz (2012, August) “Collaboration in translation: The impact of increased reach on cross-organisational work.” *Computer Supported Cooperative Work (CSCW)* 21(6), 525–554.

Doherty, Stephen and Dorothy Kenny (2014) “The design and evaluation of a statistical machine translation syllabus for translation students.” *The Interpreter and Translator Trainer* 8(2), 295–315.

Emms, Martin and Saturnino Luz (2007) *Machine Learning for Natural Language Processing*. ESSLLI’07, Dublin: European Summer School of Logic, Language and Information, course reader.

Gale, William A and Kenneth W Church (1991), “Identifying word correspondences in parallel texts.” *Speech and Natural Language: Proceedings of a Workshop Held at Pacific Grove, California, February 19-22*, New York: ACL Press.

Glozzo, A. and C. Strapparava (2009) *Semantic Domains in Computational Linguistics*. New York: Springer.

Gonzales, Annette Rios, Laura Mascarell, and Rico Sennrich (2017), “Improving word sense disambiguation in neural machine translation with sense embeddings.” *Proceedings of the Second Conference on Machine Translation* pp. 11–19. New York: ACL Press.

Hardie, Andrew (2012) “CQPweb – combining power, flexibility and usability in a corpus analysis tool.” *International journal of corpus linguistics* 17(3), 380–409.

Hermann, Karl Moritz and Phil Blunsom (2013) “Multilingual distributed representations without word alignment.” *arXiv preprint arXiv:1312.6173*.

Hutson, Matthew (2018). “Artificial intelligence faces reproducibility crisis.” *Science* 359(6377), 725–726.

Ilisei, Iustina, Diana Inkpen, Gloria Corpas Pastor, and Ruslan Mitkov (2010), “Identification of translationese: A machine learning approach.”. *International Conference on Intelligent Text Processing and Computational Linguistics* pp. 503–511. New York: Springer.

Jones, Henry (2019) “Searching for statesmanship: A corpus-based analysis of a translated political discourse.” *Polis: The Journal for Ancient Greek and Roman Political Thought* 36(2), 216–241.

Jurafsky, Daniel and James H. Martin (2008) *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition* (2 ed.). New Jersey: Prentice Hall.

Karamanis, Nikiforos, Saturnino Luz, and Gavin Doherty (2011) “Translation practice in the workplace: contextual analysis and implications for machine translation.” *Machine Translation* 25(1), 35–52.

Kilgariff, Adam, Vít Baisa, Jan Bušta, Miloš Jakubíček, Vojtěch Kovář, Jan Michelfeit, Pavel Rychlý, and Vít Suchomel (2014) “The Sketch Engine: ten years on.” *Lexicography* 1(1), 7–36.

Koponen, Maarit and Leena Salmi (2015) “On the correctness of machine translation: A machine translation post-editing task.” *The Journal of Specialised Translation* 23, 118–136.

Lafferty, John D., Andrew McCallum, and Fernando C. N. Pereira (2001), “Conditional random fields: Probabilistic models for segmenting and labeling sequence data.”. in *Proceedings of the 18th International Conference on Machine Learning* pp. 282–289. San Francisco: Morgan Kaufmann.

Laviosa, Sara (2004) “Corpus-based translation studies: Where does it come from? where is it going?” *Language Matters* 35(1), 6–27.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton (2015) “Deep learning.” *Nature* 521(7553), 436.

Lewis, David, Stephen Curran, Gavin Doherty, Kevin Feeney, Nikiforos Karamanis, Saturnino Luz, and John McAuley (2009) “Supporting flexibility and awareness in localisation workflows.” *The International Journal of Localisation* 8(1), 29–38.

Lopez, Adam (2008) “Statistical machine translation.” *ACM Computing Survey* 40(3), 1–49.

- Luong, Thang, Hieu Pham, and Christopher D Manning (2015), “Bilingual word representations with monolingual quality in mind.”. in *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing* pp. 151–159. New York: ACL Press.
- Luz, Saturnino (2011) “Web-based corpus software.” in Alet Kruger, Kim Wallmach, and Jeremy Munday (Eds.), *Corpus-based Translation Studies – Research and Applications* Chapter 5, pp. 124–149. London: Continuum.
- Luz, Saturnino and Shane Sheehan (2020) “Methods and visualization tools for the analysis of medical, political and scientific concepts in genealogies of knowledge.” *Palgrave Communications* 6(1).
- Manning, Christopher D. and Hinrich Schütze (1999) *Foundations of Statistical Natural Language Processing*. Cambridge, Massachusetts: The MIT Press.
- Màrquez, Lluís and Horacio Rodríguez (1998), “Part-of-speech tagging using decision trees.” *European Conference on Machine Learning* New York pp. 25–36. New York: Springer.
- McEnery, Tony and Andrew Hardie (2011) *Corpus Linguistics: Method, Theory and Practice*. Cambridge Textbooks in Linguistics. Cambridge: Cambridge University Press.
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean (2013), “Distributed representations of words and phrases and their compositionality.” in *Advances in neural information processing systems* pp. 3111–3119. Cambridge, MA: The MIT Press.
- Miller, George A (1995) “Wordnet: a lexical database for english.” *Communications of the ACM* 38(11), 39–41.
- Mitchell, Tom M. (1997) *Machine Learning*. New York: McGraw-Hill.
- Moorkens, Joss and Sharon O’Brien (2017) “Assessing user interface needs of post-editors of machine translation.” in *Human issues in translation technology* pp. 127–148. New York: Routledge.
- Neubig, Graham and Taro Watanabe (2016, March) “Optimization for statistical machine translation: A survey.” *Computational Linguistics* 42(1), 1–54.
- O’Brien, Sharon (2012) “Translation as human-computer interaction.” *Translation Spaces* 1(1), 101–122.
- Plank, Barbara, Anders Søgaard, and Yoav Goldberg (2016) “Multilingual part-of-speech tagging with bidirectional long short-term memory models and auxiliary loss.” *arXiv preprint arXiv:1604.05529*.
- Ratnaparkhi, Adwait (1996), “A maximum entropy model for part-of-speech tagging.” *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. New York: ACL Press.

- Ruder, Sebastian, Ivan Vulić, and Anders Søgaard (2017) “A survey of cross-lingual word embedding models.” *arXiv preprint arXiv:1706.04902*.
- Salton, Gerard, Anita Wong, and Chung-Shu Yang (1975) “A vector space model for automatic indexing.” *Communications of the ACM* 18(11), 613–620.
- Schmitt, Lothar M, Kiel Christianson, and Renu Gupta (2007) “Linguistic computing with UNIX tools.” in *Natural language processing and text mining* pp. 221–258. New York: Springer.
- Scott, Mike (2021). Wordsmith tools version 7, stroud: Lexical analysis software. (accessed 15 May 2021).
- Sebastiani, Fabrizio (2002) “Machine learning in automated text categorization.” *ACM Computing Surveys* 34(1), 1–47.
- Sheehan, Shane, Masood Masoodian, and Saturnino Luz (2018), “COMFRE: A visualization for comparing word frequencies in linguistic tasks.” *Proceedings of Advanced Visual Interfaces AVI’18* pp. 36–40. New York: ACM Press.
- Sinclair, John (1991) *Corpus, Concordance, Collocation*. Oxford: Oxford University Press.
- Sinclair, John (2003) *Reading concordances: an introduction*. London: Pearson/Longman.
- Stallman, Richard M. (2002) *Free software, free society: Selected essays of Richard M. Stallman*. Cambridge, MA: Free Software Foundation.
- Sutskever, Ilya, Oriol Vinyals, and Quoc V Le (2014), “Sequence to sequence learning with neural networks.”. in *Advances in neural information processing systems* pp. 3104–3112. Cambridge, MA: The MIT Press.
- TextTexture (2021). <https://texttexture.com/> (accessed 20 May 2021).
- Turian, Joseph, Lev Ratinov, and Yoshua Bengio (2010), “Word representations: a simple and general method for semi-supervised learning.”. in *Proceedings of the 48th annual meeting of the association for computational linguistics* pp. 384–394. New York: ACL Press
- Van der Maaten, Laurens and Geoffrey Hinton (2008) “Visualizing high-dimensional data using t-SNE.” *Journal of Machine Learning Research* 9(2579-2605), 85.
- Viégas, Fernanda and Martin Wattenberg (2008) “Tag clouds and the case for vernacular visualization.” *Interactions* 15(4), 49–52.
- Voyant Tools (2021). <https://voyant-tools.org/> (accessed 20 May 2021).
- Wattenberg, Martin and Fernanda Viégas (2008) “The word tree, an interactive visual concordance.” *IEEE Transactions on Visualization and Computer Graphics* 14(6), 1221–1228.
- Weaver, Warren (1955) “Translation.” *Machine translation of languages* 14, 15–23.

Wu, Yonghui, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. (2016) “Google’s neural machine translation system: Bridging the gap between human and machine translation.” *arXiv preprint arXiv:1609.08144*.

Zanettin, Federico (2014) *Translation-driven corpora: Corpus resources for descriptive and applied translation studies*. New York: Routledge.

Zanettin, Federico, Silvia Bernardini, and Dominic Stewart (2014) *Corpora in translator education*. New York: Routledge.

¹Gliozzo and Strapparava (2009:1) attribute the distinction between computational linguistics and natural language processing to the late Martin Kay.

²<https://xapian.org/>

³<http://www.genealogiesofknowledge.net/>

⁴The original intention was to provide the research community with on-line access to material subject to copyright restrictions in a way that did not violate copyright.

⁵Although they can currently be used for free by researchers within the EU.

⁶The version we will describe here is also known as “pointwise mutual information.” as it deals with *values* of the random variables that describe word occurrences, rather than the random variables themselves.

⁷See for instance, Voyant tools (2021) which generates interactive word clouds and summary statistics from text uploaded by the user or harvested from the web, and TextTexture (2021) which renders texts as visually appealing graphs.