Semantic Excel: An Introduction to a User-Friendly Online Software Application for Statistical Analyses of Text Data

Sikström^{1*} S., Kjell¹ O.N.E. & Kjell¹ K.

¹Department of Psychology, Lund University, Sweden

*Correspondence to: Sverker.Sikstrom@psy.lu.se

Acknowledges. We would like to thank Igor Marchetti for suggesting improvements on earlier drafts of this article.

Abstract

Semantic Excel (www.semanticexcel.com) is an online software application with a simple, yet powerful interface enabling users to perform statistical analyses on texts. The purpose of this software is to facilitate statistical testing based on words, rather than numbers. The software comes with semantic representations, or an ordered set of numbers describing the semantic similarity between words/texts that are generated from Latent Semantic Analysis. These semantic representations are based on large datasets from Google N-grams for a dozen of the most commonly used languages in the world. This *small-by-big data* approach enables users to conduct analyses of small data that is enhanced by semantic knowledge from big data. First, we describe the theoretical foundation of these representations. Then we show the practical steps involved in carrying out statistical calculation using these semantic representations in Semantic Excel. This includes calculation of semantic similarity scores (i.e., computing a score describing the semantic similarity between two words/texts), semantic t-tests (i.e., statistically test whether two sets of words/texts differ in meaning), semantic-numeric correlations (i.e., statistically examine the relationship between words/texts and a numeric variable) and semantic predictions (i.e., using statistically trained models to predict numerical values from words/texts).

Keywords: Natural language processing; Statistical semantics; Latent Semantic Analysis; Machine learning; Computer software; Automated text analyses.

Behavioral science is dominated by the use of rating scales as a measure of outcome. At the same time, the natural way for people to communicate their mental states is by means of words. We believe that a significant reason for this use of rating scales is a lack of easily accessible methods and software for quantifying and statistically measure texts. This paper is an attempt to address this problem by introducing Semantic Excel (www.semanticexcel.com), which is an online statistical software application enabling automated text analyses in a user-friendly interface. Thus, the purpose is to facilitate statistical analyses of words by mapping them to numbers.

First, we provide a theoretical background, focusing on semantic spaces and their semantic representations. This involves natural language processing by employing data-driven approaches that rely on the patterns of human word usage to represent the meaning (i.e., semantics) of words (Landauer & Dumais, 1997). Then, we describe how these semantic representations are used in various ways in Semantic Excel, including computing semantic similarity scores, semantic t-tests, semantic-numeric correlation and semantic predictions. Thus, this paper primarily focuses on the practical usages of the Semantic Excel software, for a description of potential applications of these methods in psychology, see Kjell, Kjell and Sikström (under review), which shows how statistical semantics and related methods have objective, systematic and quantitative potentials as a quantitative research tool.

The aim of Semantic Excel is to provide an easily accessible tool for statistically testing words as a variable of interest in behavioral science. This software makes natural language processing, and more specifically statistical semantics, available as a tool suitable for research in psychology and related fields. Statistical semantics with a focus on Latent Semantic Analysis (LSA; Landauer & Dumais, 1997) may for instance be carried out using the S-Space package in JAVA (Jurgens & Stevens, 2010), the DISSECT package in Python (Dinu, 2013), the *lsa* (Wild, 2015) and *LSAfun* (Günther, Dudschig, & Kaup, 2015) packages in R or the LSA website (<http://lsa.colorado.edu/>; Dennis, 2007). However, these packages require coding and current programs mostly focus on the *creation* of semantic spaces

(however, see *LSAfun*) and/or studying the nature of semantic spaces. For example, Dennis (2007) states that the LSA website is created to enable users “to investigate the properties of LSA” (p. 66). Semantic Excel, on the other hand, focuses on simplifying the mapping of semantic representations to text data and apply these in various statistical analyses. It is programmed in MATLAB, but is easily accessible as it runs online within a point-and-click environment. The software has been developed on the basis of two major objectives: First, to provide a wide range of statistical tools for analyzing different features of language, alone or in relation to numerical variables, and, second, to make natural language processing available in a highly accessible and user-friendly format. The rest of this paper is organized into two parts. In the first part, we describe the theoretical foundation of the statistical methods and in the second part we describe the practical steps involved in carrying out analyses in Semantic Excel.

The Semantic Space and its Semantic Representations

The semantic space and its semantic representations form the backbone of Semantic Excel, as they are used for the majority of analyses. These semantic representations are preprogrammed and thus directly available for analysis in a number of languages. A *semantic space* refers to a table, as exemplified in Figure 1, where the rows comprise words described on several *semantic dimensions* (i.e., in the columns), and where the numbers in the cells represent how every word is related to one another. Each row in this semantic space is referred to as a *semantic representation*, which in other words is an ordered vector of numbers describing how a given word relates to every other word in the semantic space.

Creating a high-quality semantic space (i.e., mapping numbers to words) requires a vast amount of text data, which is impractical and rarely possible to collect in research based on experiments or surveys (as opposed to corpus). However, it is worth noting that it is possible to use a pre-existing semantic space, where its semantic representations are applied to the smaller data set. In other words, big data is used for enabling analyses of small data, which we refer to as *small-by-big data analyses*

(Kjell, Kjell & Sikström, under review). Hence, Semantic Excel provides “general” semantic spaces based on large text corpora, such as the Google N-gram database (see <https://books.google.com/ngrams>), which is based on scanned books. These semantic spaces may be seen as domain-general rather than domain-specific. A domain-general space refers to when domain-unspecific text data was used for creating the semantic space, whereas a domain-specific space refers to when text related to the research data is used for creating the semantic space. If, for example, diaries are examined, then the domain-specific semantic space has been constructed using a vast amount of text from other diaries. Employing a domain-general space has several advantages. Perhaps most importantly, it allows the usage of high quality semantic representations on extremely small datasets (i.e. they may consist of single word responses generated by just a few participants). It allows direct comparisons across studies. It is easy to use, since the semantic representations can be re-used. On the other hand, employing a domain-specific space has the advantage of comprising semantic representations that may be more precise for the domain-specific context. However, for this to be practically beneficial, it requires the presence of vast amounts of domain-specific text data, which in many cases are not available, too time-consuming or too costly to collect.

Figure 1 illustrates the three major steps involved in constructing the semantic spaces used in Semantic Excel, which are based on a method similar to LSA as described by Landauer and Dumais (1997). The first step involves gathering text data in the language for which you want to create a space. Most of the semantic spaces in Semantic Excel are based on *5-gram* word contexts (i.e., sequences of five words in a row taken from a text corpus; see the left side of Figure 1).

Figure 1.

Conceptual illustration of the creation of a semantic space and its semantic representations.

The second step involves creating a word by word frequency co-occurrence table, where the most common words in all 5-gram contexts are listed in the rows (typically, the 120,000 most common words) and the columns (typically, the 10,000 most common words). The algorithm then counts how many times words co-occur within the 5-gram word contexts, so that, for example, each time *apple* and *pie* co-occur, this is indicated by raising the number in the cells indicated by *apple* and *pie* by one. When all co-occurrences have been counted, the matrix is normalized by computing the natural logarithm plus one, which increases and decreases the importance of infrequent and frequent words, respectively (e.g., see Nakov, [2001] and Lintean, Moldovan, Rus & McNamara [2010] for different weight functions for transforming the frequency co-occurrence matrix).¹

The third step involves producing and selecting the number of semantic dimensions. The semantic dimensions are produced by a type of factor analysis referred to as Singular Value Decomposition (SVD; Golub & Kahan, 1965). This reduces the 10,000 columns to a smaller number of informative semantic dimensions (c.f. factors; for a more detailed mathematical description of SVD using LSA, see for example the appendix in the article of Landauer & Dumais [1997]). The number of dimensions to subsequently retain for the semantic space is typically selected by means of some external criteria

¹ Landauer and Dumais (1997) also transform the raw matrix using an information-theoretic measure called *entropy*.

(See Table 1). This is, for example, done by means of a multiple-choice synonym test, where the number of dimensions performing the best on this test is used for the semantic space. With regard to the semantic spaces in Semantic Excel, this involved computing a semantic space error index (SSEI) by selecting synonym word pairs from a thesaurus and then examine their level of proximity to other words in the semantic space. The quality of semantic representations may then be measured as the rank order of how many words are positioned closer to one of the synonym words compared to the two synonym words. We define the SSEI as the average rank order divided by the total number of words in the space, where the highest quality is $SSEI = 0$ (i.e., all synonyms are ranked first) and $\frac{1}{2}$ for a random rank ordering. Usually, the SSEI is carried out using a complete thesaurus with thousands of word pairs.

In order to test the number of first semantic dimensions to keep, we tested the number sequence of the power of 2 (i.e., 1, 2, 4, 8, 16... 512, 1024 dimensions). Typically, a semantic space comprises 200 to 800 dimensions. The semantic space used in the forthcoming examples in this article is referred to as *English 1* and comprises 512 dimensions and 120,000 words. Table 1 displays information regarding semantic spaces available in Semantic Excel.

Table 1.

Information Regarding Semantic Spaces Available in Semantic Excel

Language Space	Word Corpus				Output		
	Size (nr. of words)	Contexts; Origin	Rows (words)	Columns (contexts)	Criteria for Dim. selection	Dim.	Performance/ Validation test
English 1	8.70 e+11	5-gram; Google ¹	120,701	10,000	SSEI	512	SSEI = .0003
English 2	8.70 e+11	5-gram; Google ¹	120,599	50,000	SSEI	512	SSEI = .0001
Swedish 1	2.99 e+13	5-gram; Google ²	120,448	10,000	SSEI	256	SSEI = .0164
Swedish 2	2.99 e+13	5-gram; Google ²	110,006	50,000	SSEI	256	SSEI = .0112
Spanish	3.23 e+11	5-gram; Google ¹	120,406	10,000	Default	300	Not yet done
Dutch	4.58e+12	5-gram; Google ²	120,386	10,000	Default	300	Not yet done
Romanian	1.21 e+13	5-gram; Google ²	120,386	10,000	Default	300	Not yet done
Italian	5.62 e+10	5-gram; Google ¹	107,235	10,000	Default	300	Not yet done
German	1.90 e+11	5-gram; Google ¹	120,408	10,000	Default	300	Not yet done
French	4.97 e+11	5-gram; Google ¹	120,406	10,000	Default	300	Not yet done
Chinese	-	5-gram; Google ¹	38,714	10,000	Default	300	Not yet done
Czech	3.73 e+13	5-gram; Google ²	120,384	10,000	Default	300	Not yet done
Finnish	7.22 e+11	5-gram; Google ²	120,384	10,000	Default	300	Not yet done
Hebrew	-	5-gram; Google ²	94,033	10,000	Default	300	Not yet done
Polish	3.80 e+13	5-gram; Google ²	120,352	10,000	Default	300	Not yet done
Portuguese	2.49 e+13	5-gram; Google ²	120,384	10,000	Default	300	Not yet done
Russian	4.06e+11	5-gram; Google ²	120,334	10,000	Default	300	Not yet done
Persian	-	-	104,352	10,000	Default	300	Not yet done
Norwegian	1.94 e+11	-	120,384	-	Default	300	Not yet done
Danish	2.18e+10	-	120,384	-	Default	300	Not yet done

Notes. The Logarithmic frequency plus 1 was used to transform the matrix for all languages. Singular value decomposition to generate 2000 dimensions were used for all languages. Dim = Dimensions; SSEI = Semantic space error index.

¹ The 5-grams come from Google N-gram Database, version July 1, 2012; see <https://books.google.com/ngrams>.

² Google N-gram CD with 10 European languages (<https://catalog.ldc.upenn.edu/LDC2009T25>)

Applying and adding semantic representations.

Semantic representations (which are created from big data, meaning a very large corpus) may be applied to a word or a very small data set of words (i.e., the small-by-big data analyses technique; see Figure 2). Collected text responses containing more than one word, such as sentences or paragraphs, may be captured in a semantic representation by adding the values of each dimension of the semantic representations of each word in the text. This is achieved by summing the values of the first dimension for all semantic representations and then summing all values of the second dimensions and so on for

all dimensions. Finally, the length of this aggregated semantic representation is normalized to one (i.e., the length is changed to one while its direction remains the same), as this simplifies the use of the semantic representations in computations, such as calculating semantic similarity.

Figure 2.

The small-by-big data analyses approach: Applying semantic representations from a semantic space to experimental data and adding them together. Notes. *The normalization of the semantic representation to 1 is not illustrated.

Semantic Similarity

Semantic representations can be seen as coordinates in a high-dimensional space, and the closer two points are situated in the semantic space, the more similar they are in meaning. This semantic similarity can be quantified numerically by computing the cosine of the angle between the two points. The semantic similarity score ranges from -1 to 1, but in practice rarely goes much below 0 (Landauer, McNamara, Dennis, & Kintsch, 2007). Importantly, high positive values indicate a high level of semantic similarity, whereas values close to 0 indicate that the meanings of the words are unrelated (i.e., orthogonal vectors). In our English semantic spaces, the density distribution of two randomly selected words exhibits a positive skew, where a single word typically has few words with similarity scores above .5 and none below -.5.

Figure 3a-b demonstrates how the level of semantic similarity between *apple* and *banana* is higher than *apple* and *pain*, and how the sentences “*In the afternoon I want apple pie*” has a higher level of semantic similarity to “*She likes to eat fruits for a snack*” compared to “*The car needs to be insured immediately.*” Hence, semantic similarity may be used for determining the degree two words/texts are similar, which may for instance be used in semantic t-tests.

Figure 3 a-b.

Conceptual illustration of semantic similarity as measured by cosine (i.e., the values in blue and brown).

Semantic T-Tests

Based on semantic representations and the semantic similarities between these, the difference in meaning between two sets of words/texts may be tested statistically. This may be particularly useful when examining whether there is a significant difference in the text responses between an intervention group and a control group. As another example, we might want to examine whether there is a statistically significant difference between participants’ answers to two different semantic questions (i.e., questions that enables individuals to answer with open-ended word responses); for example, one question concerning their harmony in life and another question concerning their satisfaction with life, as exemplified below.

The semantic t-test is carried out in three steps. In the first step, we add up all semantic representations for the harmony and satisfaction responses separately and then subtract one semantic representation from the other creating a *semantic comparison representation*. In step two, we compute the semantic similarities between the semantic comparison representation and each semantic representation of the

two sets of responses (see illustration in Figure 4). We refer to the resulting values as *semantic comparison scores*. To avoid biasing the results, a leave-n-out procedure is employed in step one and two: 10% of the responses are left out when creating the semantic comparison representation in step one, where these 10% are used in step two for computing semantic comparison scores. Step one and two are then repeated until semantic similarities to the semantic comparison representation have been computed for all responses. In the last step, a t-test is conducted comparing the semantic comparison scores between the two sets of responses. Hence, the semantic t-test tests the null hypothesis asserting that there is no mean difference between the semantic comparison scores. The results provide both a *p*-value and an effect size indicating whether or not the two sets of texts differ significantly in meaning.

Figure 4.

Conceptual illustration of how semantic representations of responses related to harmony in life and satisfaction with life are measured against the semantic comparison representation in order to create semantic comparison scores that may be tested in a t-test. (Illustrated values are made up.)

Semantic Training: Semantic-Numeric Correlation and Semantic Prediction

The semantic representations can also be used for analyzing the relationship between words/text and a numerical (or categorical) variable. This is primarily achieved by using multiple linear regression ($y = \beta_0 + \beta_1 * x_1 \dots \beta_m * x_m + \varepsilon$) and correlation. In the regression model, y is the numerical value to be predicted, x_1 through x_m are the semantic dimensions from the semantic representations,

β_0 is a constant and β_1 through β_m are the coefficients defining the relationship between the numerical values and the words. Lastly, ε is the error term.

To evaluate and avoid overfitting, the regression model is employed to a leave-n-out cross-validation method. This is a set of techniques used for evaluating the accuracy of a model whilst taking into account how well it generalizes to other data sets (for a review of cross-validation methods, see Browne, 2000). As demonstrated in Figure 5, cross-validation is carried out in several steps, where the data is split into training sets and testing sets. The parameters in the regression model (i.e., β_{0-m} and ε) are first estimated using the training set data, after which they are applied in the testing set data in order to predict the numerical values. In this way, the testing is performed on data independent from the training. The predicted values (i.e., $\hat{y}$) may ultimately be correlated with the actual scores (y) in order to assess the strength of the relationship between the text and the numerical values. Figure 5 shows leave-10%-out cross-validation zooming in on the last two steps. At each step, 90% of the data is used as a training set and 10% is used as a testing set, and this procedure is repeated until all rows have predicted values that may be correlated with the actual scores.

To further prevent overfitting in the regression model, a limited number of semantic dimensions are used. The selection process of deciding which semantic dimensions are chosen is optimized by selecting the number of dimensions offering the best prediction of the numerical values. By default, this is achieved using a leave-10%-out cross-validation procedure (however, an alternative is to use lasso regression; see Tibshirani, 1996). In practice, the model tries out different numbers of dimensions, always keeping the first dimension(s), which carries the most important information; and then successively adding more dimensions.² The model then includes the number of dimensions

² The procedure we use for adding dimensions is by adding 1 and multiplying it with 1.3 rounded to the nearest integer (e.g., 1, 3, 5, 8, 12, 17, where the next number in the series is given by $(17+1) * 1.3$ and rounded to 23) until all dimensions in the space are utilized.

providing the highest correlation between predicted ($\hat{y}$) and empirical values (y). We refer to the final predicted scores (i.e., $\hat{y}$) as semantic trained scales. Semantic training can be used for *semantic-numeric correlations* (i.e., analyzing the relationship between a text and a numeric variable within a study) or *semantic predictions* (i.e., applying trained models from other data sets) as further described next.

Figure 5. Conceptual illustration of training and testing using leave-10%-out cross-validation.

Semantic-numeric correlation

Employing cross-validation, semantic training may be used on data within a study (*within semantic trained scales*) for obtaining predicted values that are correlated with the actual/empirical values. This correlation is an estimation of the strength of the relationship between texts (e.g., word responses) and numerical values (e.g., rating scale scores). We refer to this analysis as a *semantic-numeric correlation* between a text variable and a numerical variable.

Semantic predictions

The flexibility of training models enables researchers to apply semantic trained models from other data sets on a new data set in order to predict semantic-psychological features (*independently semantic trained scales*). Thus, in Semantic Excel it is possible to create your own library of models that may be used on other data sets as well as making them publically available for other users. For example, Semantic Excel includes semantic trained models of the Affective Norms for English Words (ANEW; Bradley & Lang, 1999). The ANEW comprises a large number of words that have been rated by individuals in relation to three dimensions, including valence, arousal and dominance, which means that three separate models have been trained in Semantic Excel (i.e., the semantic representations of each word have been trained to predict their specific ratings of valence, arousal and dominance separately). When using one of the English spaces, these may now be applied to any new data; for example, enabling a semantic prediction of a text's ANEW-based valence.

Analyzing in Semantic Excel

Getting started

To get started, set-up a user account at www.semanticexcel.com and sign in. You may familiarize yourself with the website at the **Website Tour** in the upper right corner. In the main menu, under **Publications**, you find examples of a broad range of articles using Semantic Excel, and under **User Community**, you may read, ask and answer questions regarding Semantic Excel.

Import and export data

Under **Home**, you may import your data by selecting **Create** or **Import Sheet** (see Figure 6a). Selecting **Create** opens up a pop-up window asking you to name your new document, select the language (i.e., semantic space) corresponding to the language of your data and press **Save changes**. You may now find the newly created document in the list of documents in the **Name** column. Clicking

on it takes you to the document, where you may start inserting your data by hand or by using your computer's copy and paste function.

Alternatively, selecting **Import sheet** opens up a pop-up window (Figure 6b) that allows you to browse for your data file. The required formats are either a comma-separated values (CSV) file or a Microsoft Excel (XLSX) file. After choosing a data file, select the semantic space you want to use and click **Submit**. The imported document is now found in your list under **Name** and you may open it by clicking on it.

Scrolling around within a document reveals that the rows are numbered and the columns are referred to by letters. By default, one may access a fixed number of cells, but it is possible to apply for more cells in the upper right corner: **Apply for more cells**. The default cell limitation is set to avoid crashing the servers, as a large number of big analyses require considerable computer power. In the upper right corner, the document name may be changed by clicking on its current name (Figure 6c). Finally, under **File**, in the left corner, the sheet may be exported to CSV or XLSX formats. This gives you flexibility in going back and forth between different statistical programs, where non-text statistical analyses are preferably performed in programs such as R, MATLAB, SPSS or Excel.

Figure 6a-c. Display of the main page and menu (a), the **Import sheet** function (b) and an overview of the demo data (c).

The Demo Data

The demo document contains data that is used in the following examples (Figure 6c shows responses from three participants). The data is a subset from a study by (Kjell, Kjell, Garcia & Sikström, 2018), where the authors develop and validate semantic questions enabling respondents to report their subjective states of harmony in life, satisfaction with life, depression and worry using open-ended responses. The demo data includes participant number (column A), responses to semantic questions concerning harmony in life and satisfaction with life using descriptive words (column B and C, respectively) or text (column D and E, respectively) as well as composite scores for two numerical rating scales, the Harmony in life scale (column F; Kjell, Daukantaitė, Hefferon, & Sikström, 2016) and the Satisfaction with life scale (column G; Diener, Emmons, Larsen, & Griffin, 1985).

Word norms

Word norms are a list of words used for measuring the semantic similarity to an empirically measured construct (Kjell et al., 2018). Word norms may for example be created by asking an independent set of participants to describe their view of a target construct, such as harmony in life. This word norm is typically used for computing the semantic similarity to the responses to the semantic question concerning harmony in life. The higher the semantic similarity between a participant's answer concerning harmony in life to the word norm describing harmony in life, the higher the level of harmony in life we assume they enjoy (see Figure 7). Semantic Excel has a feature facilitating saving word norms, which may be used for your analyses as well as for sharing publically so that other users may use them.

Figure 7.

Conceptual illustration of semantic similarity as measured by the cosine between word responses and a word norm.

Saving and publically sharing word norms.

In this example, we save a harmony in life word norm that will be used in a later analysis. To save a word norm, go to **Scales** in the main menu and select **My norms**. Clicking **Create** opens up a pop-up window as shown in Figure 8. Select to which language space you want to save the word norm, but note that if you save it to “English 1” it will not be available in another space, such as the “English 2” space. Give the norm a suitable name (e.g., *Harmony in life 2018*) and insert the text that makes

up the word norm in the “Text norm” box. Provide your email, and making it public is optional but increases the chances of the word norm becoming useful for others. Make a careful description of the norm, including details of the methods used to generate the norm. For example, we insert the following information to describe the norm:

Word norm collected by asking 120 participants (age range: 18–51 years; mean age 29.43 (SD 7.89) years; Female = 40.8%, Male = 59.2%) to answer the question: “Please write 10 words that best describe your view of harmony in life”. Reference Kjell, Kjell, Garcia and Sikström, 2018.

Clicking **Save changes** saves the word norm and displays it in your personal list of norms. This word norm will be used in an example with regard to computing semantic similarity, but let us first familiarize the reader with some of the existing functions.

Create New Norm

English 1

Name of norm

Harmony in life 2018

Text norm

friends warmth shelter family ocean rain laughter music yoga peace quiet contentedness beauty love happiness acceptance goodwill care appreciation balance peace meaning purpose unity friendship tranquility consistent cooperation union peace quiet meaning life death reentry happiness destiny elder peace togetherness cleanliness peace cooperation caring knowledge wellness balance happiness trust love peace honour charity hope friendship respect loyalty caring leadership peace contentment happiness satisfied fulfilled married family healthy entertained in balance respect courtesy truth generosity trust love patience clarity

Description

Participants reported being from: US=95.8%; India=1.7%; Other=2.5%. Reference Kjell, Kjell, Garcia and Sikström, 2018.

Email

oscar_kjell@psy.lu.se

Make the norm public ☒

Close Save changes

Figure 8.
Displays of settings that save and publically share a word norm.

Becoming familiar with functions

The analytic tools are made available by first positioning the cursor in an empty cell where the result output will be presented and then clicking **Functions** in the upper left corner. This opens up a pop-up window giving you access to the functions. The functions are displayed in the left menu and are categorized under four headings: **Measure** (including options for measuring various aspects of words/text), **Analyze** (including options for analyzing text), **Visualize** (including options for visually highlighting various aspects of your data, such as significant words or word categories) and **Numerical functions** (including options for carrying out basic functions that are only based on numbers). The right side of the pop-up window enables you to insert the information necessary for the analysis. First, there is a **Formula box** (which is also shown above the data sheet) that shows the settings you have applied to the analysis dynamically. This is followed by a presentation of the most important settings, for example requiring you to insert the cells containing the data you want to analyze. At the bottom, you may find **Advanced options** with a large number of different parameters.

Computing semantic similarity

The following example shows how to compute semantic similarity scores between participants' word responses to the harmony in life question and the harmony in life word norm. **Semantic similarity** is the first alternative under **Measure** (Figure 9). Selecting the **Multiple text** alternative enables you to compute several pairs of words or texts simultaneously. First, in the **Text 1 (start cell) - (last cell)**, insert B2 to B101, and in the box titled **Store calculated values in**, insert J2 to J101. As we want the responses in column B to be compared with a word norm, we ignore the **Text2 (start cell) - (last cell)** option and instead select the *Harmony in life 2018* word norm in the **Select norm** drop-down menu. Clicking **OK** computes the semantic similarity between each of the semantic responses in the cells of column B with the harmony in life word norm and saves each similarity score in J2 to J101.

In addition, the stored values may be correlated with another numerical column by clicking **Correlation** in **Functions**. Inserting J2 to J101 and F2 to F101 correlates the semantic similarity

scale and the HILS-scores, revealing that these scales exhibit a significant correlation ($r = .45$; $p < .001$).

Advanced settings.

It is possible to perform more advanced analyses by setting parameters under **Advanced options**. For example, Kjell et al., (2018) found that it is possible to achieve results with a higher degree of validity by controlling for artifacts related to word frequency.³ This option may be found by clicking **Advanced options** and selecting the **Correct for frequency artifacts during creation of a semantic representation** under the heading of **Context generations**. To turn on this setting, change **0** (which means that this specific parameter is turned off) to **1** (which means this parameter is turned on) and click **OK**. However, for this tutorial keep it off (i.e., 0). To also apply this setting to the word norm, set the setting **Allow parameter settings to change norm representation** (also found under **Context generations**) to **1**.

³ To control for artefacts that relate to frequently occurring words, Kjell et al., (in revision) apply a normalization when summing the semantic representation of words. This is achieved by computing, based on Google N-gram, a frequency-weighted average of every semantic representation in the space. Hence, the weighting in this representation is thus proportional to the words' frequency in Google N-gram. This representation is then used in the normalization procedure by subtracting it from all the words that are to be summed; and then adding it to the final semantic representation.

Figure 9.
Displays of settings that compute semantic similarity scores.

Computing properties of texts

Semantic Excel allows to measure many different features of texts (e.g., see Table 2). For example, functions exist for computing word frequencies, word length, semantic coherence, semantic associates, etc. To, for example, measure the average number of letters within the harmony in life responses, go to **Properties of texts** under **Measure** and select **Multiple Text** (see Figure 10). Insert B2 and B101 in the **Text start cell - last cell** boxes, and insert K2 and K101 in the **Store calculated values in** boxes. Select **Functions** in the **Choose category of properties** drop-down menu; then select **_wordlength** in the **Choose an instance of the category** drop-down menu. Press **OK**; the results are provided in column K. The value in each cell represents the average number of letters comprising each word.

Table 2.

Description of some of the functions available for calculating properties of texts.

Function name	Brief description
_nwordsfound	Calculates how many words that are recognised by the space (i.e., that has a semantic representation). For example, the input “mother daughter son asdfe”, gives the output 3. The first three words are recognised whereas “asdfe” is not.
_frequency	Shows how frequent a word is within the text corpus that was utilised to create the space. For example “mother” gives the result .00017 (i.e., 0.017%).
_wordlength	Computes the average number of letters of the words. For example, the input “mother daughter son”, gives the output 5.67 (i.e., $(6+8+3)/3$ letters).
_coherence	Measures the degree of semantic coherence among words within a cell. The semantic similarity is calculated for a sliding window over the text. Higher scores represent higher coherence. For example, the input “mother daughter son” gives the output .49; whereas the input “mother moon down” indeed gives a lower coherence score of 0.03.
_associates	Generates the 10 words that have the highest semantic simliarty score to the semantic representation of a word/text. For example, “mother”, generates the following output: father 0.68 husband 0.59 wife 0.58 parents 0.44 woman 0.42 parent 0.40 friend 0.40 family 0.36 children 0.33 patient 0.31.

Figure 10.

Display settings for computing the number of adjectives in a text using the **Properties of texts** function.

Computing Semantic T-Tests

This example examines whether the responses to the harmony in life versus satisfaction with life semantic questions differ significantly in meaning by using a semantic t-test. Position the cursor in the cell where you want the result output to be presented and then click **Functions**. Under **Analyse**, select **Semantic t-test** and specify which columns of text data you want to compare (see Figure 11a). In relation to the current example, compare B2 to B101 with C2 to C101 by inserting them in the **Text 1** and **Text 2** boxes respectively. Without any further settings, this computes a Student's t-test. However, since we are here comparing responses *within* participants, we select a **Paired semantic t-test** under **Advanced options**. Finally, click **OK**.

To see the results, right-click on the cell where the cursor is positioned and select **Show value** (see Figure 11b). Descriptive statistics are presented first, followed by inferential results. The results reveal that the responses concerning harmony in life (Mean = 0.06; SD = .14) and satisfaction with

life (Mean = -0.08; SD = 0.13) differ significantly in meaning ($t(198) = 7.67, p < .001$) with a moderate effect size (Cohen's $d = .77$).

As a matter of fact, the result output also presents the semantic similarity scores between the semantic responses and the semantic comparison representations for the responses concerning harmony in life (set 1) and satisfaction with life (set 2). These are presented as **x1** and **x2**, respectively, with a descriptive text (e.g., for word set 1: “%Values on the semantic scale for set 1”). These numbers may be copied to another statistical program in case you want to use them for further analyses.

Functions

Measure

Semantic similarity

Properties of texts

Predict

Analyse

Semantic t-test

Semantic dimension test

Keyword test

Cluster

Train

Visualise

Plot

+

Numerical functions

Sum

Average

Standard deviation

T-test

Correlation

semanticTest(Set1Start;Set1End;Set2Start;Set2End)

Test whether two sets of texts have statistically different semantic representations.

=semanticTest(B2;B101;C2;C101)

Texts 1: start cell/last cell

B2 - B101

Texts 2: start cell/last cell

C2 - C101

[Hide Advanced Options](#)

☒ Paired semantic t-test

Select texts with the following criteria (e.g. $x > 2.3$, where x is a numerical value in cells):

Close

OK

Figure 11a-b.

Display settings for a semantic t-test (a) and its result output (b).

Computing Semantic-Numeric Correlations

Semantic-numeric correlations may be used for analyzing whether trained predictions based on semantic representations correlate with a numerical outcome variable. Here, this is illustrated by examining the relationship between the word responses to the semantic question on harmony in life and the numerical composite score of the Harmony in life scale (Kjell et al., 2016). After positioning the cursor in a cell, go to **Functions** and select **Train** under the heading **Analyze**. Select the word responses to the harmony in life question by inserting B2 to B101 in the first set of boxes labeled *Train on text data in start cell/last cell* and select corresponding rating scale scores by inserting F2 to F101 in the second set of boxes labelled *Train to predict numerical values in start cell/last cell* (Figure 12a). The prediction model that is produced as a result may be saved by giving it a name (e.g., *Harmony words trained to HILS demo*) and description (e.g., *Training harmony semantic question to Harmony in life scale scores using the demo (N=100) data; see Kjell, Kjell, Garcia & Sikström, 2018*). This enables you to apply the model to other word-based data and thus produce HILS-score estimates. Ticking the box for **Public** makes this model available for other Semantic Excel users.

Clicking **OK** produces the results (and saves the prediction). Right-clicking on the active cell and selecting **Show value** displays the result output (Figure 12b). The results show that the semantic representations of the word responses predict the corresponding rating scale scores with a Pearson correlation of $r = .54$ ($p < .001$). These results may now be added to the model description text by going to **Scales**, selecting **My predictions** and then clicking the paper icon in the **Update description** column. Clicking the green check mark changes whether or not the trained model should be available to other users. Note that a public model is only available for those using the same semantic space. The final prediction description now reads: “*Training harmony semantic questions to Harmony in life scale scores using the demo (N=100) data, $r = .54375$, $p = 3.6036e-09$; see Kjell, Kjell, Garcia & Sikström, 2018*” (Figure 12c).

Changing type of regression in advanced options.

The predicted variable in the previous example was a numerical rating scale, which is seen as an interval scale. However, if the to-be-predicted variable is a category, such as different conditions or gender, one should use (multinomial) logistic regression. The type of regression model is changed under **Advanced options**. In the **Set parameters here** drop-down menu, under the heading **Train**, select **Type of predictor during training**. This presents another drop-down menu with the following alternatives: regression, logistic, ridge and lasso.

Functions

Measure

Semantic similarity
Properties of texts
Predict

Analyse

Semantic t-test
Semantic dimension test
Keyword test
Cluster

Train

Visualise

Plot

+

Numerical functions

Sum
Average
Standard deviation
T-test
Correlation

train(ItemStart;ItemEnd;AssignedStart;AssignedEnd)

Train from text (and numerical data) to predict numerical data.

=train(B2;B101;F2;F101;'predicti

Train on text data in start cell/last cell

B2 - B101

Train to predict numerical values in start cell/last cell

F2 - F101

Prediction name

Harmony words trained to HILS c ☒ Public

Description

Training harmony semantic questions to Harmony in life scale scores using the demo (N=100) data; see Kjell, Kjell, Garcia & Sikström, in revision 2018

Close

OK

a

Value

%'Information';

modelName='_harmonywordstrainedtohilsdemo'; %Identifier of the trained model

algorithm='regression';

date='20-Sep-2018 08:59:01';

predictors='_semantic';

%'Main descriptive statistics';

n=98; %Number of data points

nWordsFound=960; %Total number of words in the texts that also are present in the semantic representation

nWords=973; %Total number of words in the texts

%'Main inferential statistics';

p=3.6036e-09; %The probability of obtaining results as extreme or more extreme given that the null hypothesis is true

r=0.54375; %Pearson correlation coefficient between predicted and empirical values

rVersusN='r=0.62329 p=0.0043568 N(words)=8.6316 r=0.59138 p=0.0060269 N(words)=10 r=0.6957 p=0.00094108 N(won

%'Information about regression';

c=20.7183; %constant term in the regression

Nremoved=2; %Number of removed or missing data points

ndim=28.1429; %Mean number of dimensions used in the prediction

Close

b

Update Description

Name

_Harmony words trained to HILS

Description

Training harmony semantic questions to Harmony in life scale scores using the demo (N=100) data, r= .54375, p= 3.6036e-09; see Kjell, Kjell, Garcia & Sikström, 2018.

Close

OK

c

Figure 12a-c.

Display settings for computing a semantic-numeric correlation (a), its result output (b) and settings to publically save (a) and update (c) a trained model.

Computing Semantic Predictions

Previously constructed models that have been saved may be used as *semantic predictions* in order to estimate a value from a text. Applying a previously trained model to a text is achieved by going to **Functions** and selecting **Predict** under the heading **Measures** (Figure 13). First select **Multiple texts** to allow predicting several values at once. Specify which text the model should be applied to by filling in the *Text start cell - last cell* boxes. Once again, select the harmony in life responses in B2 to B101. Specify where the estimated values will be saved by completing the *Store calculated values in* boxes and select the currently empty cells L2 to L101. Finally, select the variable to predict (i.e., a previously trained regression model) from the drop-down menu. For this example, select the *Valence ANEW 1999* model (Bradley & Lang, 1999). Click **OK** to run the prediction.

Kjell et al., (2018) found that the semantic trained Valence ANEW scales from semantic responses correlate strongly with corresponding rating scales. This may now be tested using the **Correlation** function under **Numerical functions**. Inserting cells L2 to L101 and F2 to F101 reveals a correlation of $r = .65$ ($p = .001$).

Figure 13.
Displays settings for predicting ANEW valence.

Closing Remarks

In this text, we have described the underlying theories of statistical semantics analyses pertaining to semantic spaces, semantic similarity scores, semantic t-tests, semantic-numeric correlations, semantic predictions as well as the visualization of words in plots. We then presented instructions for how to import data and carry out these analyses using Semantic Excel. Although some standard frequentist analyses may be carried out in Semantic Excel (e.g., t-test and correlation), the focus is on analyses related to statistical semantics. Therefore, Semantic Excel also enables easy ways of exporting data so that further analyses may be carried out in other statistical programs. Importantly, the statistical analyses using Semantic Excel (and its offline predecessor Semantic) have been used in several different contexts for many different purposes. For example, it has been used for: conceptualizing psychological constructs by investigating how individuals construe their pursuit of different types of well-being (Kjell et al., 2016); developing and validating a method for measuring and describing psychological constructs using words as the response format rather than closed-ended numerical scales (Kjell et al., 2018); investigating the relationship between written positive and negative

narratives and self-reported personality traits (Garcia, et al., 2015); and studying individuals' recall of autobiographical memories (Karlsson, Sikström & Willander, 2013). These articles are just a few examples of different research contexts where statistical semantics and Semantic Excel have been used as a research tool.

References

- Bradley, M. M., & Lang, P. J. (1999). Affective norms for English words (ANEW): Instruction manual and affective ratings.
- Browne, M. W. (2000). Cross-validation methods. *Journal of mathematical psychology*, 44(1), 108-132.
- Dennis, S. (2007). How to use the LSA web site. *Handbook of latent semantic analysis*, 57-70.
- Diener, E., Emmons, R. A., Larsen, R. J., & Griffin, S. (1985). The satisfaction with life scale. *Journal of Personality Assessment*, 49(1), 71-75.
- Dinu, G. (2013). DISSECT-DIStributional SEmantics Composition Toolkit.
- Fridolin Wild (2015). lsa: Latent Semantic Analysis. R package version 0.73.1. <https://CRAN.R-project.org/package=lsa>
- Garcia, D., Anckarsäter, H., Kjell, O. N., Archer, T., Rosenberg, P., Cloninger, C. R., & Sikström, S. (2015). Agentic, communal, and spiritual traits are related to the semantic representation of written narratives of positive and negative life events. *Psychology of Well-Being*, 5(1), 1.
- Golub, G., & Kahan, W. (1965). Calculating the singular values and pseudo-inverse of a matrix. *Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis*, 2(2), 205-224.
- Günther, F., Dudschig, C., & Kaup, B. (2015). LSAfun-An R package for computations based on Latent Semantic Analysis. *Behavior research methods*, 47(4), 930-944.
- Jurgens, D., & Stevens, K. (2010). *The S-Space package: an open source package for word space models*. Paper presented at the Proceedings of the ACL 2010 System Demonstrations, Uppsala, Sweden.
- Karlsson, K., Sikström, S., & Willander, J. (2013). The Semantic Representation of Event Information Depends on the Cue Modality: An Instance of Meaning-Based Retrieval. *PLoS ONE*, 8(10), 1-8. doi:10.1371/journal.pone.007337

Kjell, O. N. E., Kjell, K., Garcia, D., & Sikström, S. (2018). Semantic measures: Using natural language processing to measure, differentiate, and describe psychological constructs. *Psychological Methods*, doi:10.1037/met0000191

Kjell, Kjell & Sikström (under review). Statistical Semantics: Measuring and Describing Psychological Phenomena through Natural Language.

Lintean, M. C., Moldovan, C., Rus, V., & McNamara, D. S. (2010). *The Role of Local and Global Weighting in Assessing the Semantic Similarity of Texts Using Latent Semantic Analysis*. Paper presented at the FLAIRS Conference.