Semantic Norm Extrapolation is a Missing Data Problem

Bryor Sneffjella & Idan Blank

University of California, Los Angeles

Department of Psychology

Author Note

7581 Franz Hall, Los Angeles California, 90095-1563

bsneffjella@psych.ucla.edu

Kind thanks are due to MacData, the audiences of the 2019 Society for Personality and Social and Psychology Language Preconference and 2018 Psychonomics conferences, as well as Craig Enders for support and critiques of the early stages of the project. We thank Alexander Robitzsch for help with multiple imputation software. We thank the Institute for Digital Research and Education at UCLA for high performance computing support.

Abstract

For close to 70 years psychologists have studied word meaning using a simple method: participants rate words on some theoretically motivated property (e.g. pleasantness, familiarity) using a Likert scale as the measurement instrument. Such semantic judgments serve as a means of interrogating the underlying structure of lexical semantic constructs, to select stimuli for experiments, or as covariates in models predicting brain or behaviour. Recently, there has been a surge of interest in using computational distributional semantic word representations and supervised learning to predict semantic judgments on Likert scales for words lacking empirical measurements. We call this task semantic norm extrapolation. A significant body of work has developed showing methods for semantic norm extrapolation are often highly accurate. The impressive performance of models for this task may give the appearance that non-empirical, machine learning derived estimates of semantic norms are interchangeable with empirical measurements of semantic norms. Herein, we argue that this is not the case, and that *all* extant methods for semantic norm extrapolation are more problematic than the literature suggests. Naive use of extrapolated semantic norms should be expected to yield biased and anti-conservative analyses. We make this argument using a mixture of 1) the principles of analysis of partially observed data, 2) simulations, and 3) a real-data example. Achieving sound inference when using semantic norm extrapolation requires a conceptual and methodological shift from treating semantic norm extrapolation as a prediction problem to treating it as a missing data problem. This shift in perspective also lays bare problems in default analytical procedures of semantic norms and megastudy data, and surprisingly suggests that semantic norm extrapolation — when done using recommended procedures for analysis of partially observed data — should be default methodological practice.

Keywords: semantic norms, missing data, valence, arousal, concreteness, lexical semantics

Semantic Norm Extrapolation is a Missing Data Problem

Introduction

The science of word meaning has suffered from data sparsity since its inception. By a conservative definition of “word,” an average English speaker knows more than 40,000 words (Brysbaert, Stevens, Mandera, & Keuleers, 2016), but only a handful of studies of word meaning include this many items. The seminal work of Osgood, Suci, and Tannenbaum (1957) entrenched a simple and still widely used method to study meaning: survey participants about a theoretically motivated dimension of meaning (e.g. a word’s positive or negative connotation, how animate it is), using Likert scales as the measurement instrument. Item means gleaned from these instruments — a word’s semantic norm — are commonly used to select stimuli, as covariates in analysis of brain or behaviour, or to automatically code natural language corpora for semantic features. Recent advances in computational linguistics have revealed that semantic norms can be conveniently and accurately reproduced through a computational synthesis of natural language data and statistical methods. These statistically estimated semantic norms are already being published for researchers to use (cf. van Paridon and Thompson, 2019), and used in inference (cf. Dingemanse and Thompson, 2020). It may seem that high correlations with empirical semantic norms make it straightforward to use extrapolated norms as covariates in regression analyses. However, we feel the apparent accuracy of these methods masks subtle methodological problems when extrapolated norms are substituted for empirical ones. Herein, we describe why this emerging practice is statistically problematic, and offer solutions for how to put the use of extrapolated semantic norms as covariates on firmer statistical ground. As well, we argue that while semantic norm extrapolation is currently done problematically, it is also actually *underused*. Default statistical practice in many analyses using semantic norms risks biases that a properly constructed extrapolation may correct. Fundamentally, the problem with semantic norm extrapolation to date is that it has been conceptualized as a prediction problem, when a

more optimal approach is to treat norm extrapolation as a missing data problem.

What do we mean by a missing data problem? In practical statistical analysis, it is often the case that some observations are only partially observed. For example, if a participant skips a question on a survey, a blink occurs during a fixation on an important stimulus in an eye-tracking study, or an experimental device fails halfway through an experiment (for reviews of missing data problems and methodology in psychology, see Graham, 2009). We argue that missing data problems emerge often in single-word processing studies when experimental stimuli are selected based on their availability in existing lexical resources, or when secondary analyses of data from existing single word processing experiments are conducted. For example, when two datasets of semantic norms are combined, but only shared words included in both the datasets are analyzed. Words that are known to exist but that lack easily accessible measurements of their features are routinely deselected from analysis, and the unknown features of the deselected words become missing data. When data is partially observed, the analysis risks either a loss of statistical power or bias, depending on why the data are missing and the response to the missing data by the data analyst. Since the seminal work of Rubin (1976), statisticians have pursued how sound inferences can be drawn in the face of partially observed data. We adapt techniques from the missing data literature to perform our semantic norm extrapolation.

In contemporary times words excluded from any analysis almost always have at least some partially observed data: features of words vectors from distributional semantic models (DSM) (Landauer & Dumais, 1997; Mikolov, Grave, Bojanowski, Puhersch, & Joulin, 2018; Mikolov, Sutskever, Chen, Corrado, & Dean, 2013; Pennington, Socher, & Manning, 2014). Rather than discarding words lacking measurements from analysis, semantic norm extrapolation leverages DSMs to “fill in” missing empirical semantic norms with extrapolated semantic norms, so words with partially observed empirical data can still be included in the analysis. In essence, the semantic norm extrapolation methods that have

already been proposed are actually possible solutions to the missing data induced by inconsistent coverage of the lexicon between different experiments and datasets of word features. However, the semantic norm extrapolation literature has not framed the issue this way nor taken inspiration from the extensive statistical literature on missing data. In the missing data literature “filling in” missing values is known as *imputation*. We explicitly evaluate proposed methods for norm extrapolation as an imputation technique for missing data, using evaluation procedures from the missing data literature. Unfortunately, the methods for semantic norm extrapolation that have been used to date, despite their often high predictive accuracy, are also expected to be a sub-optimal solution to missing data when recast as an imputation technique. Yet, we also argue that even these problematic methods can sometimes be better than discarding words without complete data from the analysis. We find that a superior method to existing ones is to extrapolate missing semantic norms using procedures recommended by statisticians when data is partially observed, using a technique known as multiple imputation (Rubin, 2004). We present simulations and a real data example to support these views. To begin our foray into issues of missing data, what data is missing?

Many lexical semantic constructs and their measurement scales have been developed. Taking studies publicly sharing semantic norms in supplementary materials or online repositories as representative, prominent threads in the literature include a word’s emotional connotations, as captured by dimensions of valence, arousal, and dominance (Bradley & Lang, 1999; Cso, Guasch, Ferr, & Hinojosa, 2019; Dodds, Harris, Kloumann, Bliss, & Danforth, 2011; Eilola & Havelka, 2010; Gilet, Grhn, Studer, & Labouvie-Vief, 2012; Ho, Mak, Yeung, Duan, Tang, Yeung, & Ching, 2015; Imbir, 2015; Mohammad, 2018a; Moors, De Houwer, Hermans, Wanmaker, Van Schie, Van Harmelen, De Schryver, De Winne, & Brysbaert, 2013; Mukherjee & Heise, 2017; Redondo, Fraga, Padrn, & Comesaa, 2007; Schauenburg, Ambrasat, Schrder, von Scheve, & Conrad, 2015; Schmidtke, Schrder, Jacobs, & Conrad, 2014; Scott, Keitel, Becirspahic, Yao, & Sereno,

2018; Sianipar, van Groenestijn, & Dijkstra, 2016; Söderholm, Häyry, Laine, & Karrasch, 2013; Stadthagen-Gonzalez, Imbault, Sánchez, & Brysbaert, 2017; Verheyen, De Deyne, Linsen, & Storms, 2019; Warriner, Kuperman, & Brysbaert, 2013; Yao, Wu, Zhang, & Wang, 2017; Yee, 2017). Alternatively, words have been normed for the degree to which they connote discrete emotional categories, such as disgust, sadness, or anger (Binder, Conant, Humphries, Fernandino, Simons, Aguilar, & Desai, 2016; Briesemeister, Kuchinke, & Jacobs, 2011; Hinojosa, Martínez-García, Villalba-García, Fernández-Folgueiras, Sánchez-Carmona, Pozo, & Montoro, 2016; Mohammad, 2018b; Stadthagen-González, Ferré, Pérez-Sánchez, Imbault, & Hinojosa, 2018; Stevenson, Mikels, & James, 2007; Syssau, Yakhoulfi, Giudicelli, Monnier, & Anders, 2020). An alternative account to valence and arousal puts words' danger and usefulness at the centre of their emotional connotations (Heard, Madan, Protzner, & Pexman, 2019; Wurm, 2007). Word tabooeness has been normed (Janschewitz, 2008; Reilly, 2019a; Roest, Visser, & Zeelenberg, 2018). Concreteness, the degree to which a word's referent is accessible to the senses (e.g. *pitbull* versus *democracy*) has been studied extensively (Altarriba, Bauer, & Benvenuto, 1999; Benjafield & Muckenheim, 1989; Bonin, Méot, & Bugaiska, 2018; Brysbaert, Warriner, & Kuperman, 2014; Campos, 1990; Coto et al., 2019; Gilhooly & Logie, 1980; Scott et al., 2018; Soares, Costa, Machado, Comesaña, & Oliveira, 2017). As well, numerous theories distinguish between types of abstract and concrete words or posit specific roles for sensorimotor and perceptual systems in the mental lexicon. Examples include imageability (Benjafield et al., 1989; Gilhooly et al., 1980; Grandy, Lindenberger, & Schmiedek, 2020; Schock, Cortese, & Khanna, 2012; Scott et al., 2018; Soares et al., 2017; Stration, Jacobus, & Brinley, 1975; Walker, 1970), context availability (Altarriba et al., 1999), body object interaction ratings (M. A. Alonso, Díez, Díez-Alamo, & Fernandez, 2018; Bennett, Burnett, Siakaluk, & Pexman, 2011; Pexman, Muraki, Sidhu, Siakaluk, & Yap, 2019; Tillotson, Siakaluk, & Pexman, 2008), motor content norms (Abella & González-Nosti, 2019), relative embodiment of verbs (Sidhu, Kwan, Pexman, & Siakaluk, 2014), graspability

ratings (Amsel, Urbach, & Kutas, 2012; Heard et al., 2019), manipulability (Medler, Arnoldussen, Binder, & Seidenberg, 2005; Moreno-Martínez, Montoro, & Rodríguez-Rojo, 2014), sensory experience ratings (Bonin, Méot, Ferrand, & Bugajska, 2015; Díez-Álamo, Díez, Wojcik, Alonso, & Fernandez, 2018; Hinojosa, Rincón-Pérez, Romero-Ferreiro, Martínez-García, Villalba-García, Montoro, & Pozo, 2016; Juhasz & Yap, 2013), iconicity ratings (Perry, Perlman, & Lupyan, 2015; Winter, Perlman, Perry, & Lupyan, 2017), mode of acquisition ratings (Della Rosa, Catricalà, Vigliocco, & Cappa, 2010), sensory modality specific concreteness ratings (Lynott & Connell, 2009; Lynott, Connell, Brysbaert, Brand, & Carney, 2019; Miklashevsky, 2018; Morucci, Bottini, & Crepaldi, 2019; Speed & Majid, 2017; Vergallito, Petilli, & Marelli, 2019), and ratings of which bodily effectors are used to interact with a word’s referent (Lynott et al., 2019). Age of acquisition norms measure the age at which participants believe they learned a word (M. A. Alonso, Fernandez, & Díez, 2015; M. Á. Alonso, Díez, & Fernandez, 2016; Birchenough, Davies, & Connelly, 2017; Hinojosa, Rincón-Pérez, et al., 2016; Kuperman, Stadthagen-Gonzalez, & Brysbaert, 2012; Montefinese, Vinson, Vigliocco, & Ambrosini, 2019). Pain semantic norms were gathered by Borelli, Crepaldi, Porro, and Cacciari (2018), and Binder et al. (2016). Binder et al. (2016) proposed 65 relevant dimensions of meaning according to a review of lexical semantics in cognitive psychology, and collected norms for a set of more than 500 words on all these dimensions, including cognition, colorfulness, and nearness. Clark and Paivio (2004) and Scott et al. (2018) measured binary gender connotations. Engelthaler and Hills (2018) measured word humourousness. All told, we identified approximately 118 unique semantic scales/constructs¹ in the literature (see below and Table 1). Clearly the use of Likert scales to measure meaning has and continues to have broad appeal and impact for the study of meaning in cognitive science.

¹ We use both the terms *scale* and *construct* interchangeably because it is difficult to disentangle how much differences and overlaps between measurements in this literature can be attributed to scales capturing the same or different latent factors or due to scale biases introduced from different norming instructions.

While these diverse constructs have been proposed and studied, the number of words actually normed on a given semantic norm scale is often only a small proportion of an average English speaker’s 40,000 word lexicon. Figure 1 plots year of publication against number of normed words and number of normed constructs for all available open-source English Likert-scale semantic norm datasets we could locate, (see Table 1 for a list of datasets). Although an increasing number of words and semantic constructs are normed, few datasets contain more than a few thousand normed words. With increasing use of crowdsourcing (mass online collection of semantic norms) after 2010, some datasets in the tens of thousands emerge. Yet, these large datasets cover only a handful of the most studied semantic constructs such as valence and concreteness. Figure 2 shows the number of words normed for each of 118 unique Likert-scale measured constructs with English open-source norms we could locate in the lexical semantics literature. The majority of constructs have only hundreds of words with norms. For most words of English most semantic constructs do not have a semantic norm, this data sparsity only worsens as one considers other Indo-European languages, and resources become all but nonexistent for languages from other families. This poor coverage of the lexicon — actually most lexicons — for most semantic constructs is one strong motivation for semantic norm extrapolation. Outside of a handful of intensively studied constructs in English, most lexical semantic constructs lack measurements for most words.

Previous literature has acknowledged the poor coverage of the lexicon, and the recent remedy other than semantic norm extrapolation has been to gather semantic norms for most known words through online crowdsourcing platforms (cf. Brysbaert et al., 2014; Kuperman et al., 2012; Warriner et al., 2013). These massive online data collection efforts provide a more representative sample, usually using only a cutoff value of corpus word frequency to select words for norming. While crowdsourcing allows for norm collection at scale, it is also costly: paying a mere one cent per semantic judgment to each participant, assuming at least 30 ratings are necessary to get the central tendency for each word, and

that 40,000 words is an ideal sample, yields a cost of \$8,000 for semantically norming a single construct. The presence of more than a hundred proposed semantic constructs, multiple measurement formats beyond Likert scales, many languages, the presence of systematic individual differences, and that these norms are not stable over time, multiplies the cost. Paying participants a cent for each semantic judgment, a “back of the envelope” calculation yields a cost of \$180,000,000 assuming the objective is to norm 40,000 words for 100 constructs, 30 judgments per word are used to get its central tendency, in 3 response formats, assuming 5 relevant individual differences exist, and for ten languages. Depending on how quickly meaning changes, we would be paying this cost repeatedly over time. This is probably still a conservative estimate, because semantic norming is rarely done for lexical units not conventionally defined as “words” in the literature, such as proper nouns, emojis, multi-word expressions, or for less dominant senses of words with multiple senses. Even paying a fraction of this number by conducting true random sampling from the lexicons of the world’s languages cost can be prohibitive. Some semantic norms will probably always be missing. For these practical reasons, and not scientific reasons, a computational replacement for empirical measurements of semantic norms is quite desirable.

Concerns about data sparsity and cost of semantic norms become more acute in the current research environment in single word processing. It has become common for databases of behavioural responses to thousands to tens of thousands of word stimuli to be published for other researchers to analyze (Balota, Yap, Hutchison, Cortese, Kessler, Loftis, Neely, Nelson, Simpson, & Treiman, 2007; Cortese, Khanna, & Hacker, 2010; Cortese, McCarty, & Schock, 2015; Ferrand, New, Brysbaert, Keuleers, Bonin, Méot, Augustinova, & Pallier, 2010; Keuleers, Lacey, Rastle, & Brysbaert, 2012; Mandera, Keuleers, & Brysbaert, 2019; Pexman, Heard, Lloyd, & Yap, 2017; Sze, Liow, & Yap, 2014; Tse, Yap, Chan, Sze, Shaoul, & Lin, 2017; Tucker, Brenner, Danielson, Kelley, Nenadić, & Sims, 2019; Yap, Liow, Jalil, & Faizal, 2010). Megastudies improve the replicability, power, and generalizability of psycholinguistic research. But, these solutions to the problems of

yesterday create the problems of today. Frequent re-use of megastudy data leads to more cases of missing data. This occurs when, for example, a researcher has developed a new measurement of a word property (e.g. a new semantic construct and its measurement scale), wishes to use megastudy data to determine if this measurement is relevant to word processing, but that measurement is only available for a subset of words in the megastudy. Indeed, it is routine that when re-analyzing megastudy data, only a subset of items in the megastudy will have complete measurements (cf. large proportions of megastudy items discarded in Brysbaert and Biemiller, 2017; Kuperman, Estes, Brysbaert, and Warriner, 2014; Mandera et al., 2019; Pexman et al., 2019; Snefjella and Kuperman, 2016). Missing semantic norms lead to only portions of megastudy items being included in the analysis, causing at minimum a reduction in statistical power and potentially returning researchers to the problematic, small-sample analyses that have damaged scientific credibility (Ioannidis, 2005).

Above and beyond the sheer quantity of missing norms and partially observed data routinely discarded because of them, the complete data may also be systematically different from the missing data. Ideally, the words chosen to be semantically normed would be a true random sample from the lexicon, to ensure results from studies of meaning generalize to other words and have external validity about linguistic processing and behaviour. Yet, in many cases the normed words are explicitly nonrandomly selected. Shifting to a missing data framework reveals that if the megastudy items discarded due to missing data are not a random sample from the full data, results will also suffer from bias (Rubin, 1976). We quantitatively confirm in Study 1 that in general words with existing empirical semantic norms are systematically different than those words missing them. This means that whenever an experiment is constructed or an analysis is conducted using only the word stimuli with complete existing data, the analysis may reach biased conclusions that fail to generalize to the entire lexicon. An extrapolation that incorporates procedures for analyzing cases where the complete and missing data are systematically different can

potentially restore unbiased inference (see below) without the need to collect empirical norming data. This is another reason why we argue extrapolation should be done with methods recommended to address missing data problems, why we believe extrapolation is underused, and why extrapolation should be default practice.

If semantic norm extrapolation is needed to solve data sparsity because most semantic norms are missing, to make unbiased generalizations about word processing in general, and to not bankrupt an esoteric corner of psychological science, how has it been done? How can you measure meaning with nothing but computation and statistics? There are four components: a training sample of semantically normed words, a supervised learning method, a corpus of unstructured natural language data, and a distributional semantic model (e.g. Landauer et al., 1997; Mikolov et al., 2018; Mikolov et al., 2013; Pennington et al., 2014). Distributional semantic models operate through the distributional hypothesis (Firth, 1957; Harris, 1954), which states that we can take similarity of words' linguistic contexts as an implicit measurement of their similarity in meaning. Quantitatively speaking, training a distributional semantic model yields a word-by-feature matrix, placing each word in an n -dimensional space, with n usually set to 300. Words with more similar linguistic contexts, and therefore meanings, are embedded in more similar locations within this space. The word-by-feature matrix becomes training data for the supervised model, which is optimized to predict the training set of semantic norms. The most common supervised learning method for the task is linear regression, which amounts to finding a hyperplane through the 300 dimensional space that minimizes squared error. For words lacking empirical norms but attested in the training corpus, a semantic norm can be estimated based on that word's position in the semantic space relative to the learned hyperplane.

What supervised learning methods have been proposed for semantic norm extrapolation? Methods include linear regression (Gultchin, Patterson, Baym, Swinger, & Kalai, 2019; Hollis, Westbury, & Lefsrud, 2017; Mandera, Keuleers, & Brysbaert, 2015;

Martínez-Huertas, Jorge-Botana, Luzón, & Olmos, 2020; Thompson & Lupyan, 2018; Turton, Vinson, & Smith, 2020; Utsumi, 2018; Westbury, Shaoul, Hollis, Smithson, Briesemeister, Hofmann, & Jacobs, 2013), ridge regression (Paetzold & Specia, 2016; Recchia & Louwerse, 2015; Richie, Zou, & Bhatia, 2019; Turton et al., 2020; van Paridon et al., 2019), locally-weighted regression (Wang, Yu, Lai, & Zhang, 2016), random forests (Mandera et al., 2015; Turton et al., 2020), k-nearest neighbours (Bestgen, 2002; Bestgen & Vincze, 2012; Mandera et al., 2015; Martínez-Huertas et al., 2020; Turton et al., 2020; Van Rensbergen, De Deyne, & Storms, 2016), exemplar based methods (Köper & Im Walde, 2016; Paetzold et al., 2016), feed-forward neural networks (Buechel, Rücker, & Hahn, 2020; Köper & im Walde, 2017; Martínez-Huertas et al., 2020; Turton et al., 2020; Utsumi, 2018), support vector machines with a linear kernel (Li, Lu, Long, & Gui, 2017), support vector machines with a polynomial kernel (Paetzold et al., 2016), and support vector machines with a radial basis function kernel (Ljubešić, Fišer, & Peti-Stantić, 2018). As well, numerous semi-supervised or unsupervised methods for the task have been proposed in the computational linguistics literature (see Hamilton, Clark, Leskovec, and Jurafsky, 2016 and references therein). The predictive accuracy of many of these norm extrapolation methods is impressive, by the standards of psychology. For constructs with the largest available semantic norms, R^2 on held-out validation samples generally ranges from .4 to .8, depending on the reliability of the semantic norm scale, supervised learning model used, and choice of distributional semantic model. Buechel et al. (2020) report the formidable r of .94, .76, and .88 for a validation set of English valence, arousal, and dominance norms respectively. This is approaching the theoretical maximum correlation possible: the reliability of these semantic norm scales.

While the results of these studies are impressive, we can immediately call attention to some issues. Most extrapolation papers to date extrapolate the largest available sets of norms: valence, arousal, dominance, and concreteness. Accurately extrapolating semantic norms that there are already an abundance of is less valuable than extrapolating rarer

norms. The effectiveness of the methods cited above for other semantic norm scales is not known. Smaller samples would be more prone to overfitting, and overfitting is more likely with the complex nonlinear methods used to achieve highest accuracies on the largest sets of norms. Accurate extrapolation of rarer norms is more desirable, but less often done. The issue of extrapolation in small samples is compounded when these extrapolation methods are admitted to be imputation techniques. Failure to account for the “known unknowns” of imputation model error and parameter uncertainty, which would be larger in small samples, is a critical deficiency of the extrapolation methods cited above (see below for extensive discussion). Based on the fact that most semantic norm scales contain only hundreds of words and the typical dimension of word vectors is 300, more appropriate imputation/extrapolation methods may be ones suitable for “high p, low n” statistical problems and not necessarily flexible nonparametric methods that benefit from large samples.

A second issue with all these methods is how they are evaluated. How does a researcher decide which extrapolation method is the best to use? The literature on semantic norm extrapolation has proceeded by exploring what supervised learning algorithms and what type of corpus and machine learning derived word representations yield the highest predictive accuracy. Most often, predictive accuracy is measured by cross validation. The “best” method achieves the highest R^2 on samples of semantic norms withheld during training. However, cross validation is not immune to sampling concerns and model dependence. For cross validated error to be an unbiased estimator of generalization error, the observed data used to form training/testing partitions must be a random sample from the population (Murphy, 1983). We show in Study 1 that words with available semantic norms are a highly non-random sample from the lexicon. Therefore cross validated error should be expected to be an optimistic estimate of the generalization error of existing semantic norm extrapolation methods. More problems with cross validation appear when we switch to a missing data perspective. In the missing data

literature imputation methods are usually not evaluated by their generalization error. They are evaluated by their ability to provide unbiased parameter estimates with nominal confidence interval coverage in a downstream analysis (Rubin, 2004; Van Buuren, 2018). As true values of parameters are rarely known, most often imputation techniques are evaluated through simulations. No methods for semantic norm extrapolation to date have been evaluated this way (though see Hollis et al., 2017; Mandera et al., 2015 for discussion of how extrapolated norms compare to empirical ones). Good cross validation accuracy is not the recommended objective for solving a missing data problem.

Why doesn't evaluating extrapolation methods by their predictive accuracy do what psycholinguists want? Because psycholinguists are not merely interested in predicting semantic norms. The extrapolation process is just "one step on the road" to inference about brain or behaviour. Predicting that, for illustration, the word *truck* has an animacy norm of 2.1 but *river* has an animacy of 3.5 isn't where our scientific labour ends. A typical analysis aims to relate semantic norms to behaviour, brain responses, or probe the underlying structure of different constructs' norms. The end goal is actually to put theories about how the mind and language works to strong hypothesis tests. Ideally, researchers would have a truly random and representative sample of empirical semantic norms. Lacking this, researchers have a complex black box, distributional semantic models, which bear systematic relations to empirical semantic norms, and psycholinguists wish to stand in for empirical semantic norms in the analysis of brain or behaviour. But, such extrapolations are not empirical measurements. Which is to say that when extrapolated semantic norms are used, analyses should provide maximal confidence that the use of extrapolations isn't biasing the results. However, a good cross-validated accuracy does not entail unbiased parameter estimates in downstream analysis (Rubin, 2004). If an analysis is based only on empirical data, it should provide a more confident result than an analysis based on partially observed data. However, achieving a good cross validated accuracy does nothing to adjust confidence intervals of parameters to reflect a distinction between

estimates from empirical data and estimates from imputed data. Psycholinguists don't care what the predicted animacy of *truck* is; Psycholinguists care about falsifying hypotheses about the construct of animacy using strong statistical tests. Semantic norm extrapolation is not a prediction problem. Semantic norm extrapolation is a missing data problem.

Accepting that in analyses using extrapolated semantic norms the objective should be unbiased estimates of scientific parameters of interest and for confidence intervals for these parameters to reflect uncertainty introduced by using extrapolations, can the extrapolation methods given above be expected to provide them? Unfortunately the missing data literature strongly suggests such methods will be biased. They must be anti-conservative. Within missing data analyses there are terms for what the proposed methods for semantic norm extrapolation summarized above do — they are called single imputation², regression mean imputation³, or naive imputation, and these methods are dispreferred (Allison, 2001; Brown, 1994; Enders, 2010; Horton & Kleinman, 2007; R. J. Little & Rubin, 1989; Van Buuren, 2018) due to inherent overconfidence and expected biases in downstream analyses. The reasons for both bias and overconfidence of these methods is that they discard known sources of uncertainty about the extrapolation/imputation model's predictions: the estimated error term and the uncertainty introduced from estimating parameters in finite sample data. When an analyst neglects uncertainty in their imputation model for semantic norms and imputes missing data with a conditional mean, they are assuming an unreasonably favourable and highly implausible scenario. They are assuming that if they actually sampled empirical semantic norms for the words missing them, the empirical measurements would fall precisely on the most likely predicted value of the model

² Imputation refers to 'filling in' missing data with plausible values. Single imputation refers to any method which corrects for missing data by replacing a missing value with a single value, as opposed to multiple imputation, which produces $m > 1$ values for the missing data.

³ Regression mean imputation refers to a single imputation where missing values are filled in with predicted values of the conditional mean of a regression model

for every observation, with zero variation.

An alternative to regression mean imputation is to impute the missing values using a model’s conditional mean plus a random draw from the estimated error: an imputation approach known as stochastic regression imputation. This imputation method can be unbiased, but is still anti-conservative. Any single imputation method that fills in missing data with a single value and with the analysis treating the full data matrix as if it had been complete, empirically measured data all along must be overconfident. This amounts to testing hypotheses only in one plausible realization of the missing data, when there are in fact many.

Accepting that semantic norm extrapolation is a missing data problem, and that single regression mean imputation — therefore all methods proposed in the semantic norm extrapolation literature so far — is expected to be both biased and overconfident, what can be done? One solution to these issues, and to in principle achieving sound inference when using extrapolated semantic norms, is to use a multiple imputation technique (R. J. Little & Rubin, 2019; Rubin, 2004) for the extrapolation. This amounts to making assumptions about the joint distribution of the complete data, repeating the analysis under different plausible realizations of the missing data, and diminishing confidence in rejections of null hypotheses based on how strongly the extra variance introduced by missing data affects parameter estimates. Within a multiple imputation framework, unbiased inference and proper coverage of confidence intervals can be achieved by filling in missing data with plausible values, drawn from the posterior predictive distribution of the missing data. Rather than “filling in” the data once to provide one complete dataset to analyze, instead the missing data is filled in $m > 1$ times, yeilding $m > 1$ datasets, where the empirical data is same, but the imputed values vary. Applying multiple imputation to semantic norm extrapolation would yield $m > 1$ values for each word without an available semantic norm, where variability in these values captures imputation model uncertainty about that missing semantic norm. In a downstream statistical analysis, rather than performing a single

analysis, instead $m > 1$ analyses are conducted, with estimates of the parameters of interest (their means and standard errors) pooled using a set of simple equations known as Rubin’s Rules (Rubin, 2004). These equations serve to inflate the variance of parameter estimates, depending on how much uncertainty variability in the missing values introduces to those estimates. As well, multiple imputation can restore full-sample inference, as no partially observed rows of the data matrix are discarded due to missingness. In effect, a multiple imputation analysis answers the question: given my assumptions about the missing data, the observed data I have, and the “known unknowns” in my model of the missing data, does the evidence still support rejection of the null hypothesis?

To explain why we feel multiple imputation is the optimal solution for semantic norm extrapolation, we must first introduce some missing data terminology. The missing data literature defines three types of missing data mechanisms. Missing data mechanisms describe the possible patterns of missing data: missing completely at random (MCAR), missing at random (MAR), and missing not at random (MNAR).

When the data is missing completely at random (MCAR) the probability of the data being missing is not conditional on either the observed data, or any unobserved data. An example would be if someone maliciously deleted some entries for a variable on the basis of flips of a fair coin. If the missing data is missing at random (MAR), the probability of the data being missing depends on the observed data, but not on any unobserved data. This would be the case, for example, if we had part of speech information but were more likely to gather semantic norms for nouns: missingness would be conditional on part of speech. If the missing data is missing not at random (MNAR), the probability of the data being missing depends on unobserved data. This would be the case, for example, if we set out to gather tabooess norms, but tended to exclude the most extremely taboo words from our stimuli. Under these sampling conditions estimates of mean tabooess would not converge to the true population mean tabooess. We would also lack data to correct this bias. The appropriate choice of remedy for missing data depends on which of these mechanisms

accurately describes the missing data.

The simplest response to missing data is to simply discard any rows of the data matrix with missing values and perform the analysis over the remaining data: a complete case analysis. This is the default in most statistical software, and as noted above, in many use cases of semantic norms, particularly when megastudy data is analyzed. If the data is MCAR then a complete case analysis will yield unbiased inference at the expense of statistical power. In the case of semantic norm extrapolation it seems that a MCAR missing data mechanism is unlikely. There are many candidate variables that missingness could be conditional on, such as corpus word frequency and word length. In Study 1, we show that missingness for semantic norms in general is conditional on word frequency and other lexical variables, making an MCAR assumption and complete case analysis dubious default practice.

If the missing data is MAR, where the probability of the data being missing depends on the observed data, then a complete case analysis will be biased (R. J. Little, 1992; Rubin, 1976; Schafer & Graham, 2002). “Missing at random” is a misnomer. It refers to that missingness is random after we have taken into account its dependence on observed variables. MAR indicates the completely observed data and partially observed data are not random samples of the population. However, under MAR the observed data contains information on the missingness mechanism and therefore steps can be taken by the analyst that correct for the missing data and yield unbiased inference. For semantic norm extrapolation, we demonstrate missingness in semantic norms is likely to be conditional on many lexical variables. We feel the default assumption of semantic norm extrapolation should be that the missing data is MAR, missing conditional on observed lexical variables as well as features of word vectors from DSMs.

If the missing data is missing not at random then the probability of the data being missing depends on unobserved data. A MNAR mechanism is also impossible to quantitatively verify, since if the analyst had the data to do so, the mechanism would be

MAR. In the face of an MNAR mechanism, sound inference rests on correct specification of the missingness mechanism by the analyst (Mustillo & Kwon, 2015; Rubin, 2004).

Why is multiple imputation our preferred solution? Multiple imputation can provide sound inference in the face of MCAR, MAR, and even MNAR mechanisms under certain conditions. Under MCAR, multiple imputation will find no systematic relations in the data and will impute the missing data with a random draw from the scale of the observed data. Multiple imputation under MCAR still allows the complete data to be analyzed and can increase statistical power relative to a complete case analysis. Under an assumption of MNAR, extensions of multiple imputation exist allowing for imposition of priors on the missing data. Under MAR, assuming the model of the posterior distribution of the missing data is correct, the number of imputations m is large enough, the model used by the downstream analysis is correct, and the parameter(s) of interest are normally distributed across the imputed datasets, multiple imputation can restore unbiased inference. For these reasons, statisticians often consider multiple imputation a more optimal default behaviour when data is missing than complete case analysis (Enders, 2017; van Ginkel, Linting, Rippe, & van der Voort, 2019; Wilkinson, 1999).

A further challenge in semantic norm extrapolation is that the multiple imputation method used should be able to handle multivariate missing data. How can we perform multiple imputation if multiple variables have missing data? Multivariate imputation by chained equations (MICE) is one approach (Buuren & Groothuis-Oudshoorn, 2010). MICE yields a Gibbs sampler for the missing data, under the assumption that a multivariate distribution for the missing data exists. This method uses a “divide and conquer” approach, breaking down the multivariate imputation problem into a sequence of univariate imputations. First, missing values are filled in with random draws from the observed data for each variable. For each variable with missing data, the observed data is used to supervise training of a predictive model. This predictive model uses any of the fully observed or partially observed variables as predictors, as determined by the researcher or

an automated variable selection method (e.g. LASSO, random forests). The missing data in that variable are then filled in with random draws from a predictive distribution conditioned on the observed data (in the simplest case a draw from normal distribution with a mean equal the estimated conditional mean and standard deviation equal to the estimated standard deviation of the model’s residuals). This process is repeated for all variables with missing data, allowing imputations of the missing values in one variable to serve as predictors for others. This process is iterated, with imputations from previous iterations serving as predictors for later ones, until any trends in the imputations have stabilized. Repeating this process $m > 1$ times yields $m > 1$ imputed datasets, to be used in downstream analyses with results pooled using Rubin’s rules. We use MICE to perform our semantic norm extrapolation when multivariate missing data is present.

The shift to a missing data framework also clarifies how the model filling in the missing data much be structured, in order for that model to be able to provide unbiased coefficient estimates and nominal coverage of confidence intervals. Imputation models can be “proper” or “improper” depending on whether they take uncertainty of the imputation model parameters into account (Rubin, 2004). Multiple imputation was developed in a Bayesian context, and for Rubin’s rules to be a valid way of pooling estimates across imputations, imputation models must take into account that parameter estimates (e.g. regression coefficients, residual standard deviations) are estimated from sample data (Rubin, 1976). This means that the marginal imputation models in MICE must either 1) be fully specified Bayesian models, or 2) do bootstrapping of the observed data at each iteration of MICE to approximate a fully Bayesian model. Again, no bootstrapping has been done in any extrapolation models we could locate, nor have Bayesian methods been used, making all such methods improper by the definitions of Rubin (2004). All semantic norm extrapolation methods fail to quantify uncertainty in estimates introduced from estimating model parameters in finite samples. To embed existing methods into a multiple imputation framework, we opted to use bootstrapping for computational convenience.

Two more important pieces of missing data terminology are imputation models and analysis models. An advantage — but also pitfall — of multiple imputation is that it separates the filling in of missing values, imputation, from the later analysis of the data. The advantage is that, similar to a megastudy, multiply imputed data can potentially be published for other researchers to analyze. The pitfall is that unlike with empirical measurements the analyst is constrained in the assumptions their analysis can make by the assumptions made by the imputer. If, for example, the analyst wishes to test for nonlinear relation between variables but the imputer assumed a linear relation, the analyst’s and imputer’s models are deemed uncongenial (for a formal definition of uncongeniality, see Meng, 1994) and the analyst’s model will be biased. Thus, the substantive question of interest informs how imputation and analysis should proceed. Simply using published extrapolated semantic norms, without considering potential sources of uncongeniality between the imputer’s and analyst’s models, is another reason why current approaches to semantic norm extrapolation may lead to be biased results, even in the face of a good cross-validated accuracy.

Critically, avoiding uncongeniality requires that the dependent variable for the analyst’s model should be included as a predictor in the imputer’s model. While this may seem at first to lead to overfitting, it does not. The point of multiple imputation is to assume that the data contains all relevant information (i.e. assume a MAR mechanism) about the missing data, and to discover the relations between variables in the observed data and preserve these relations (as well as known sources of uncertainty about them) into the missing values. Failing to include the dependent variable as a predictor of the missing data is expected to downwardly bias the very relations the analyst wishes to discover (van Ginkel et al., 2019). It is worth noting that we could not find a single example of semantic norm extrapolation that included the dependent variable of a downstream analysis as a predictor in the extrapolation/imputation model. Almost universally, only features of word vectors are used as predictors. Therefore, analyses using currently existing

sets of extrapolated norms implicitly assume something rather controversial: that all relevant information about the variables of interest with missing data, the mechanism by which the data are missing, and their relation to the dependent variable in the analysis, is contained in the word vectors. This is a strong assumption. A safer approach is to include all analysis model variables in the imputation, including the dependent variable, to ensure congeniality. Thus poor selection of variables to include in the imputation model represents another source of bias that is obscured when we fail to treat norm extrapolation as a missing data problem.

Similar congeniality concerns arise when we consider many analyses involve multiple semantic constructs. To insure that we accurately induce dependencies between extrapolated semantic norms for different constructs, all types of semantic norms in the analysis model should be combined in the imputation model. This would allow as much information as possible from partially observed sets of norms to be preserved into the extrapolated norms. With the exception of Buechel et al. (2020) who used feed-forward neural networks and multi-task learning, existing semantic norm extrapolation methods have proceeded through independent models for each kind of semantic norm extrapolated. More than being potentially uncongenial, independent models for each kind of semantic norm is also inefficient, in that information contained in other kinds of semantic norms or relevant lexical variables — but not in word vectors — may be lost. Nor do independent models for individual semantic norms correct for missing data in any non-semantic variables. This may lead to a MNAR rather than MAR missing data mechanism and therefore more bias. Multiple imputation imputes missing data in all variables while attempting to exploit all sources of information, in all variables both completely and partially observed, to maximize efficiency in the estimates of the missing data.

Another important piece of missing data terminology is an auxiliary variable. This refers to a variable that is used in the imputation model, but is not used in the analysis model. When semantic norm extrapolation is shifted to a missing data framework and uses

multiple imputation, this is what the features of the word vectors function as. Features of the word vectors themselves are unlikely to be of substantive interest and included in the analysis model. As the word vectors are auxiliary variables, features of the word vectors are not even necessary to perform multiple imputation at all. The imputation model could proceed using only variables found in the analysis model. Auxiliary variables are included to increase the efficiency of the estimates of the missing data, and to make the assumption of MAR more plausible (Hardt, Herke, & Leonhart, 2012; Yoo, 2009). Yet, in certain cases, such as when correlations between auxiliary variables and the variable with missing data are low, inclusion of auxiliary variables can lead to less precision and more bias (Hardt et al., 2012). Surprisingly, it could be the case in terms of unbiased parameter estimation that using word vectors could actually be worse than not using them. The choice to use word vectors or not depends on whether the imputer believes the word vectors contain information absent from the variables in the analysis model about the missing data. Thus, in our simulations and real data analyses we perform semantic norm extrapolation both with word vectors included as auxiliary variables, and with no auxiliary information at all (i.e. no features of word vectors, only analysis model variables) in a multiple imputation. Again, no semantic norm extrapolation method has been evaluated this way, and it has never been considered if features of words vectors need to be used at all, or may in fact be detrimental, in the restoration of sound inference in downstream analyses.

A further advantage of pivoting to a multiple imputation framework is a suite of statistical tests which augment familiar statistics such as tests for significance for regression coefficients and model comparisons testing relative goodness of fit (Enders, 2010). These statistics adjust confidence by taking into account the variance between imputations induced by the missing data, and can be used to test if a predictor appears to contribute significantly to model fit even if that predictor is only partially observed. As well, these statistics provide information beyond their regular, full-data counterparts on how strongly the missing data has affected the estimate of the parameter. Again, simply including

outputs of a regression model as covariates, then proceeding with statistical tests as one would with an empirical measurement, will necessarily yield anti-conservative analyses compared to methods that take into account between-imputation variance. As well, simply performing the usual statistical tests will fail to indicate how strongly a result “leans” on the imputed or complete data.

To summarize, most semantic norms are actually missing, we have good reason to suspect that missing semantic norms can be described with a MAR or MNAR missing data mechanism, and that for practical reasons there probably always will be some words with missing norms. When data is MAR or MNAR ignoring this missing data leads to biased analyses. While approaches for extrapolating semantic norms exist, if we use them to impute missing data (which is being already done, but not being called imputation), we would expect analyses to be 1) anti-conservative and biased due to neglect of imputation model error and parameter uncertainty, 2) biased due to uncongeniality between imputation and analysis models, and 3) inefficient, due to discarding information from partially observed or important variables beyond features of word vectors and by failing to correct for missingness in non-semantic variables. Favourable evaluation by cross validated accuracy does not necessarily fix any of these issues. However, existing approaches for semantic norm extrapolation can be modified to incorporate them into a multiple imputation, which can provide unbiased and appropriately confident analyses under a properly specified imputation model, analysis model, and if the true missing data mechanism is MAR.

Having summarized how the principles of missing data analysis apply to semantic norm extrapolation, we now lay out evidence for our argument that semantic norm extrapolation is a missing data problem in 3 sections. First, we explore the most likely missing data mechanism for semantic norms in general and determine variables that missingness in semantic norms will probably be conditional on. Second, we explore machine learning methods suitable to high-dimensional cases, we adapt these methods for

use in MICE, and evaluate these methods against existing ones using simulations. Third, we apply proper multiple imputation techniques to a real empirical question involving semantic norms, and compare results between a single regression mean imputation (i.e. the most common current method used for semantic norm extrapolation) and a proper and congenial multiple imputation.

Study 1: Missingness Patterns in Semantic Norms

Appropriate remedy for missing data depends on the mechanism by which that data is missing. While the precise missing data mechanism will depend on particular research questions and data, we can get a broad view of whether an MCAR or MAR assumption in analyses involving semantic norms is more plausible. Ideally, the words chosen for semantic norming would be a random sample from the lexicon, licensing the commonplace use of complete case analysis and an implicit default assumption of MCAR in single word processing studies. Yet, the papers accompanying open-source datasets of semantic norms name many ways words are nonrandomly selected for norming. What factors influence the probability that a word is selected to be semantically normed?

A non-exhaustive survey reveals many nonrandom reasons that words are selected. Common reasons are inclusion in earlier datasets (Bradley et al., 1999; Engelthaler et al., 2018; Lynott et al., 2019; Moreno-Martínez et al., 2014; Pexman et al., 2019; Stadthagen-Gonzalez et al., 2017; Vergallito et al., 2019; Warriner et al., 2013), corpus word frequency (Brysbaert et al., 2014; Engelthaler et al., 2018; Moors et al., 2013; Warriner et al., 2013), concreteness (Amsel et al., 2012), simpler morphology (Moors et al., 2013), number of syllables (Juhasz et al., 2013; Tillotson et al., 2008), semantic category membership (Ferré, Guasch, Moldovan, & Sánchez-Casas, 2012), part of speech (Lynott et al., 2009, 2013) or arbitrarily by the experimenter (Janschewitz, 2008; Roest et al., 2018). From these examples, it seems that a word is more likely be normed if it occurs in corpora with higher frequency, is shorter or simpler, and is more concrete.

To quantitatively verify and explore this apparent nonrandom selection of words for norming, we combine an expert word list, open-source datasets of semantic norms, and quantitative measurements of word frequency and other lexical characteristics, and present scatterplots of these variables against the rate of a word’s missingness among many semantic constructs with open-source datasets of norms. The expert word list stands in as the best representation for the total population that ideally would have been sampled randomly from, and the word lists from each semantic norm dataset represent what researchers actually did sample. If the missing semantic norms are really MCAR we expect no lexical characteristics of the words to predict their rate of missingness. A rejection of the null hypothesis of an overall MCAR missing data mechanism would suggest that MAR is the safer default assumption in analyses involving missing semantic norms. A basic choice in this analysis is whether to explore missingness in each type of semantic norm separately, or somehow pool together norms to look at broad patterns in missingness. Rather than explore missingness mechanisms for each type of semantic norm separately, we opted to perform an overall test, combining multiple sources of semantic norms and predicting the rate by which words are missing semantic norms from many sources. In this case, because we don’t have a specific analysis model in mind, we are more interested in the overall picture of what data is missing, and whether the evidence supports the default assumption and practice if MCAR and complete case analysis, or MAR and multiple imputation.

Materials

Table 1

Open-Source Semantic Norm datasets used in Study 1

Semantic Construct/Scale	Sources
Age of Acquisition (Integer Ratings)	Hollis and Westbury (2018), Kuperman et al. (2012)
animacy	Heard et al. (2019)

Table 1

Open-Source Semantic Norm datasets used in Study 1 (continued)

Semantic Construct/Scale	Sources
Age of Acquisition (Likert Scale Ratings)	Clark et al. (2004), Coltheart (1981), Gilhooly et al. (1980), Juhasz, Lai, and Woodcock (2015), Schock, Cortese, Khanna, and Toppi (2012), Scott et al. (2018), Stration et al. (1975)
arousal	Binder et al. (2016), Eilola et al. (2010), Hollis et al. (2018), Janschewitz (2008), Mohammad (2018a), Scott et al. (2018), Warriner et al. (2013)
association with babies	Perry et al. (2015)
attention	Binder et al. (2016)
audition	Binder et al. (2016)
auditory strength	Amsel et al. (2012), Binder et al. (2016), Lynott et al. (2009, 2013), Lynott et al. (2019), Medler et al. (2005), Reilly (2019b)
away	Binder et al. (2016)
benefit	Binder et al. (2016)
biomotion	Binder et al. (2016)
body	Binder et al. (2016)
body object interaction	Bennett et al. (2011), Pexman et al. (2019), Tillotson et al. (2008)
brightness	Binder et al. (2016)
categor	Toglia (2009)
caused	Binder et al. (2016)
cognition	Binder et al. (2016), Reilly (2019b)
color	Amsel et al. (2012), Binder et al. (2016), Medler et al. (2005), Reilly (2019b)
communication	Binder et al. (2016)
complexity	Binder et al. (2016)
concreteness	Altarriba et al. (1999), Benjafield et al. (1989), Brysbaert et al. (2014), Clark et al. (2004), Coltheart (1981), Eilola et al. (2010), Gilhooly et al. (1980), Hollis et al. (2018), Scott et al. (2018), Toglia (2009)
consequential	Binder et al. (2016)
context availability	Altarriba et al. (1999), Clark et al. (2004)

Table 1

Open-Source Semantic Norm datasets used in Study 1 (continued)

Semantic Construct/Scale	Sources
danger	Heard et al. (2019), Wurm (2007)
darkness	Binder et al. (2016)
discrete anger	Binder et al. (2016), Mohammad (2018b), Stevenson et al. (2007)
discrete disgust	Binder et al. (2016), Stevenson et al. (2007)
discrete fear	Binder et al. (2016), Mohammad (2018b), Stevenson et al. (2007)
discrete happy	Binder et al. (2016), Mohammad (2018b), Stevenson et al. (2007)
discrete sad	Binder et al. (2016), Mohammad (2018b), Stevenson et al. (2007)
discrete surprise	Binder et al. (2016)
dominance	Mohammad (2018a), Scott et al. (2018), Warriner et al. (2013)
drive	Binder et al. (2016)
duration	Binder et al. (2016)
ease of definition	Clark et al. (2004)
embodiment	Sidhu et al. (2014)
emotion	Reilly (2019b)
emotionality	Clark et al. (2004)
emotionality goodness	Clark et al. (2004)
emotionality pleasantness	Clark et al. (2004)
face	Binder et al. (2016)
fast	Binder et al. (2016)
foot leg	Lynott et al. (2019)
gender	Clark et al. (2004), Scott et al. (2018)
goodness	Benjafield et al. (1989)
graspability	Amsel et al. (2012), Heard et al. (2019)
gustatory strength	Binder et al. (2016), Lynott et al. (2009, 2013), Lynott et al. (2019)
hand arm	Lynott et al. (2019)

Table 1

Open-Source Semantic Norm datasets used in Study 1 (continued)

Semantic Construct/Scale	Sources
haptic strength	Binder et al. (2016), Lynott et al. (2009, 2013), Lynott et al. (2019)
harm	Binder et al. (2016)
head	Binder et al. (2016), Lynott et al. (2019)
high	Binder et al. (2016)
human	Binder et al. (2016)
iconicity	Winter et al. (2017)
imageability	Altarriba et al. (1999), Benjafield et al. (1989), Bennett et al. (2011), Clark et al. (2004), Coltheart (1981), Gilhooly et al. (1980), Janschewitz (2008), Juhasz et al. (2015), Schock, Cortese, and Khanna (2012), Scott et al. (2018), Stration et al. (1975), Toglia (2009), Walker (1970)
interoception	Lynott et al. (2019)
landmark	Binder et al. (2016)
large	Binder et al. (2016)
left meaning dominance	Juhasz et al. (2015)
long	Binder et al. (2016)
loud	Binder et al. (2016)
low	Binder et al. (2016)
lowerlimb	Binder et al. (2016)
manipulability	Medler et al. (2005)
humour mean	Engelthaler et al. (2018)
meaningfulness	Clark et al. (2004), Coltheart (1981), Stration et al. (1975), Toglia (2009)
moral	Reilly (2019b)
motion	Amsel et al. (2012), Binder et al. (2016), Medler et al. (2005)
motionself	Reilly (2019b)
mouth	Lynott et al. (2019)
music	Binder et al. (2016)
near	Binder et al. (2016)
needs	Binder et al. (2016)

Table 1

Open-Source Semantic Norm datasets used in Study 1 (continued)

Semantic Construct/Scale	Sources
negative valence	Binder et al. (2016)
number	Binder et al. (2016)
number of actions	Heard et al. (2019)
number of attributes	Toglia (2009)
number of meanings	Clark et al. (2004)
offensiveness	Eilola et al. (2010), Janschewitz (2008)
olfactory strength	Amsel et al. (2012), Binder et al. (2016), Lynott et al. (2009, 2013), Lynott et al. (2019)
pain	Amsel et al. (2012), Binder et al. (2016)
pantomime	Heard et al. (2019)
path	Binder et al. (2016)
pattern	Binder et al. (2016)
personal use	Janschewitz (2008)
polarity	Reilly (2019b)
positive valence	Binder et al. (2016)
practice	Binder et al. (2016)
pronouncability	Clark et al. (2004)
quantity	Reilly (2019b)
scene	Binder et al. (2016)
self	Binder et al. (2016)
sensory experience	Juhasz et al. (2015), Juhasz et al. (2013)
shape	Binder et al. (2016)
short	Binder et al. (2016)
size	Heard et al. (2019), Scott et al. (2018)
slow	Binder et al. (2016)
small	Binder et al. (2016)
social	Binder et al. (2016), Reilly (2019b)
space	Reilly (2019b)
speech	Binder et al. (2016)

Table 1

Open-Source Semantic Norm datasets used in Study 1 (continued)

Semantic Construct/Scale	Sources
subjective familiarity	Amsel et al. (2012), Benjafield et al. (1989), Clark et al. (2004), Coltheart (1981), Eilola et al. (2010), Gilhooly et al. (1980), Janschewitz (2008), Juhasz et al. (2015), Scott et al. (2018), Stration et al. (1975), Toglia (2009)
tabooness	Janschewitz (2008), Reilly (2019a)
tactile	Reilly (2019b)
taste pleasantness	Amsel et al. (2012)
tastesmell	Reilly (2019b)
temperature	Binder et al. (2016)
texture	Binder et al. (2016)
time	Binder et al. (2016), Reilly (2019b)
torso	Lynott et al. (2019)
toward	Binder et al. (2016)
transparency	Juhasz et al. (2015)
upperlimb	Binder et al. (2016)
usefulness	Heard et al. (2019), Wurm (2007)
valence	Clark et al. (2004), Dodds et al. (2011), Eilola et al. (2010), Hollis et al. (2018), Janschewitz (2008), Medler et al. (2005), Mohammad (2018a), Scott et al. (2018), Toglia (2009), Warriner et al. (2013)
visual strength	Binder et al. (2016), Lynott et al. (2009, 2013), Lynott et al. (2019)
visualform	Reilly (2019b)
weight	Binder et al. (2016)

In order to confirm and explore the trends in missing data, we opted to combine an expert constructed word list and all available open-source English semantic norm datasets we could find. A full list of included datasets is given in Table 1. As many datasets measure the same semantic constructs, we opted to merge datasets measuring the same construct into a single word list per construct. We identified 46 datasets with approximately 118 unique lexical semantic constructs or scales. Some datasets contain related constructs

gathered with different instructions, e.g. both Binder et al. (2016) and Reilly (2019b) normed words on scales related to thought or cognition, but used different instructions for participants. Others contain similar but not synonymous concepts e.g. graspability, manipulability and body object interaction norms. We opted to merge constructs into a single word list if they set out to measure substantially similar constructs, and norms of words shared between the datasets had a linear relation and $r > .8$. Merged constructs are grouped together by row in table 1. Merging these sources of data allows for quantitative prediction of a word’s probability of being normed on a construct, assuming the expert word list reflects the population of words we ideally would have semantic norms for.

As our expert word list we took the dataset of Brysbaert, Mandera, McCormick, and Keuleers (2019). This contains some 62,000 English word forms compiled from expert databases, subtitle corpora, and open-source dictionaries. Furthermore Brysbaert et al. (2019) used online crowdsourcing to ask 220,000 people if they know these words. This provides a measurement of word prevalence, the proportion of individuals who know the word. Word prevalence is intended to overcome some of the limitations of corpus-derived word frequency estimates. It has greater discrimination than corpus-frequency among very low frequency words, whereas corpus-derived frequency estimates discriminate more among high frequency words. This dataset also provides subtitle word frequencies from the SUBTLEX-US (Brysbaert & New, 2009) corpus. As well, the English Crowdsourcing Project (Mandera et al., 2019), a recent online lexical decision megastudy, gathered response times for this massive sample of words. Means of response times z-scored within participants allow verification of whether a word’s lexical decision response time is related to its being missing from more semantic constructs and scales.

A limitation of this analysis is that we, ideally, would have complete measurements for semantic properties of words, to determine if missingness is conditional on lexical semantics. Yet, such measurements are of course unavailable for the complete data, as determining how to provide sound estimates is exactly the exercise of this paper. The sole

semantic variable we consider has the largest available norms and the least missing data, making a complete case analysis more defensible: concreteness from (Brysbaert et al., 2014).

However, as distributional semantic models contain latent dimensions corresponding to many psycholinguistically relevant word properties (Hollis & Westbury, 2016), including their semantics, we can test if a word’s probability of being semantically normed has a relation to features of its word vector. Therefore, we chose as our distributional semantic models two pretrained, open-source sets of word vectors recently found to be highly effective across many natural language processing tasks and that we found had excellent predictive performance in the early stages of this project. Our first source is fastText vectors (Mikolov et al., 2018). We use the subword variant trained on the 600 billion word common crawl corpus. As word forms in these vectors are cased, but most semantic norm datasets are uncased, we merged words based on their lowercase forms in the fastText word vectors. Although this set of embeddings contains 2 million word types, it did not contain vectors for all words in Brysbaert et al. (2019). We therefore included an additional set of conceptnet numberbatch vectors (Speer, Chin, & Havasi, 2017). For words without fastText vectors, we used conceptnet numberbatch vectors aligned to fastText vectors through a procrustes rotation (Schönemann, 1966). Of the 62,000 words in Brysbaert et al. (2019), 60629 had available vectors when both sources of vectors are combined.

As a dependent variable in these analyses, we calculated a variable representing the proportion of the 118 constructs which was missing a semantic norm for each word in Brysbaert et al. (2019). Proportions ranged from 0.08 to 1, indicating that no word has been semantically normed on all proposed semantic constructs, and some words had never been normed at all. Figure 3 presents a histogram of missingness proportions of words. We then use a semi-parametric binomial regression to predict the proportion of missingness for each word, using a separate model for each lexical variable, and plot results. While the pattern of missing data may vary by individual semantic constructs, this approach provides

an overall “audit” of which factors seem to be related to selection of words for semantic norming. Additionally, we begin by visualizing words by their proportion of missingness, to explore possible trends in missingness.

Methods. We begin by visualizing words by the proportion of constructs for which they are missing, and qualitatively explore what factors missingness may be conditional on. To assess missingness quantitatively, for each candidate lexical variable on which missingness may be conditional, we train a binomial generalized additive model (GAM) predicting proportion of constructs that are missing a semantic norm measurement. We used the *bam()* function of the *mgcv* R package (Wood, 2017) with $k = 20$, except for word length which used a $k = 10$, and thin plate spline basis function. Effective degrees of freedom (EDF) values for the splines indicate whether that variable has a linear or nonlinear relation to missingness on the log-odds scale, with $\text{EDF} > 3$ indicating a significant departure from linearity (Wood, 2017). We then provide scatterplots of proportion of missingness by each lexical variable superimposed with this spline on the probability scale. We take deviance explained by the GAM as the relative indicator of how influential each lexical variable is in the prediction of missingness.

To test for possible effects of the raw features of the word vectors on missingness, we fit three binomial GAMs, one containing only smooths of the lexical predictors, one containing cubic splines of each of the 300 features of the word vectors, and one containing both. Each spline included an additional penalty on the null space (via the *select = T* argument to *bam()*), allowing for shrinkage of the spline towards zero, in order to perform variable selection. As well, this allows meaningful likelihood ratio tests of GAM models fit with the fREML method (Wood, N., Pya, & Säfken, 2016). If the word vectors contribute in the prediction of missingness beyond the lexical variables, some of their features should be selected (i.e. not shrunk to zero). Additionally, likelihood ratio tests of the fREML scores of these models provide quantified confirmation that the additional deviance explained by the addition of word vectors or lexical variables justifies additional model

complexity. This allows us to assess whether word vectors contain information above and beyond the lexical variables, in the prediction of missingness.

Results

Before exploring the effects of specific lexical features on missingness, we visualize and qualitatively assess groups of words clustered by the proportion of datasets where they are missing. Figure 4 presents words within ranges of proportion of missingness. Where too many words fall within that range of missingness, we present a random sample of words within that range. Words with the lowest probability of being missing are short, concrete nouns (e.g. *car* and *tree*). As the probability of being missing increases, words generally become longer, more complex, and less concrete (e.g. *blasphemousness* and *inferential*). As well, we observe more jargon (e.g. *diode*, *polycystic*). Morphologically derived forms are often missing, even if their stems are common known words (e.g. *romancer*, *dreaded*, *sifter*, *unwanted*). Many commonly known words are nonetheless normed on less than 30% of scales, such as *ten*, *sweatpants*, and *rig*. Although some abstract words (*hope*, *peace*, *knowledge*, and *deceit*) are among the few words that are normed on most constructs, abstract words appear to be more often missing norms. All these trends lend plausibility that missingness will be conditional on word frequency and prevalence, length, and concreteness. In the next sections, we confirm this quantitatively variable by variable using binomial generalized additive models.

Word Frequency. Figure 5 presents a scatterplot of word frequency and missingness. Word frequency has a moderate (deviance explained = 45.5%), non-linear on the log-odds scale (effective degrees of freedom = 18.02) relation with proportion of missingness. The proportion of missingness is lowest at mid to high ranges of word frequency, but increases among both very high and very low frequency words. The trend of greater missingness among high frequency words likely represents the absence of function words from semantic norming. Words attested less frequently in SUBTLEX-US are less

likely to be semantically normed, particularly those at the lowest possible frequency values (i.e. a single occurrence in SUBTLEX US).

Word Prevalence. Figure 6 presents a scatterplot of word prevalence and missingness. Word prevalence has a large (deviance explained = 62.72%), non-linear on the log-odds scale (effective degrees of freedom = 16.3) relation with proportion of missingness. The proportion of missingness is lowest at the highest levels of word prevalence (i.e. among words 100% of people say they know), but increases as fewer people report knowing the word.

Word Length. Figure 7 presents a scatterplot of word length and missingness. Word length has a smaller (compared to word frequency and prevalence) effect on missingness (deviance explained = 7.06%). The effect is non-linear on the log odds scale (effective degrees of freedom = 8.95). The proportion of missingness is lowest among words 3 to 4 characters long, but becomes higher for very short or for very long words.

Concreteness. Figure 8 presents a scatterplot of concreteness and missingness. Concreteness has a small (compared to word frequency and prevalence) effect on missingness (deviance explained = 10.13%). The effect is non-linear on the log odds scale (effective degrees of freedom = 17.8). Words are missing in more datasets if they are more abstract, and more likely to be included if they are more concrete.

English Crowdsourcing Project Response Times. Figure 9 presents a scatterplot of English Crowdsourcing Project Z-scored response times and missingness. Response time has a large effect on missingness (deviance explained = 58.02%). This is the second largest effect next to word prevalence. The effect is non-linear on the log-odds-scale (effective degrees of freedom = 17.39). Words are more likely to be missing as their response times become larger. An increase in response time among low response time words leads to a large increase in overall missingness, but levels off as response times increase.

Word Vectors. Can raw features of word vectors predict a word’s missingness from semantic norm datasets? And do these tap into different components of the

missingness mechanism than the lexical features such as word length, word prevalence, word frequency, and z-scored RT? The GAM model containing only features of the word vectors explained 57.09% of the deviance in proportions of missingness. This is smaller than the deviance explained by the model with only lexical features (71.45%). As well, the model containing the word vectors was more complex ($df = 5.8449868 \times 10^4$ versus $df = 6.049318 \times 10^4$) and had a worse fREML score ($fREML = 1.8592937 \times 10^5$ versus $fREML = 2.570489 \times 10^5$). A model containing only lexical features is a simpler and more predictive model than one containing features of the word vectors, yet the word vectors do seem to predict missingness. A model containing both features of the word vectors and the lexical predictors was preferred over both the simpler models. The χ^2 test on the fREML scores indicates that model with both types of features is significantly better than model with only features of word vectors ($\chi^2(18.00)=16666.42, p<2e-16$), and the model with only lexical features ($\chi^2(300.00)=3953.362, p<2e-16$). The model containing both features explained 74.73% of the deviance in the proportion of datasets missing the words. This suggests features of the word vectors are predictive of missingness above and beyond the lexical features alone. However, the black box nature of the word vectors prevents easy interpretation of what these possibly semantic factors underlying missingness are.

Discussion

For this study, we have explored missingness in semantic norms by seeing if lexical characteristics and/or the raw features of word vectors predict the proportion of semantic norm constructs missing a word. Problematically for complete cases analyses of megastudy response times involving semantic norms, z-scored response times from the English Crowdsourcing Project are a strong predictor of missingness. Results also indicate that corpus word frequency and word prevalence, a variable introduced by Brysbaert et al. (2019) that measures the proportion of participants who indicate they know a word, are strong predictors of missingness. Word concreteness also appears to be related to

missingness, with abstract words being more missing. Less frequent and less known words are less likely to be normed. Very short and very long words are less likely to be normed. The relations of these variables to missingness was also complex: e.g. both very short and very long, and very low, and very high, frequency words were more often missing. Taken together, it seems very dubious to assume words selected for semantic norming are a random sample from the lexicon, and we should assume that semantic norms are MAR by default. The individual missing data mechanism will likely vary depending on the precise selection criteria for words for individual semantic norm datasets, but the variables we have identified here are likely candidates for missingness to be conditional on. This should also make us pessimistic about the validity of complete case analyses with semantic norms, as keeping only complete data vectors in a data matrix after merging existing semantic norms can be expected to systematically drop words that have longer response times, are more abstract, are very short or are five characters or longer, are infrequent or very highly frequent, and are less known from the sample. This supports our argument that multiple imputation, and not complete case analysis, should be default behaviour for analyses involving semantic norms and missing data. Additionally, we note that raw features of word vectors seem to predict missingness well, but less well than the lexical features. Though it is difficult to know what additional information the word vectors contain about missingness, the best performing model of missingness contained both smooths of the lexical features and smooths of raw features of the word vectors.

Study 2: Proper versus Improper Imputation

Having established that, by default, we should expect semantic norms to be missing conditional on observed lexical variables, we now seek statistical remedies for this missing data problem. Complete case analysis is expected to be biased in this context, where the data is MAR. In this study, we demonstrate the problematic nature of both semantic norm extrapolation techniques that do not use multiple imputation and complete case analyses.

Instead of evaluating extrapolation methods by their predictive accuracy, we evaluate whether methods provide unbiased inferences and optimal converge of confidence intervals in downstream analyses. We perform two simulations. To compare methods, we compare results obtained from 1) analyses of the complete data 2) complete case analyses, 3) improper single regression mean imputation, as per the most common approaches in semantic norm extrapolation 4) multiple imputation, but without any simulated word vectors, and 4) multiple imputation with simulated word vectors included as auxiliary variables in the imputation. All multiple imputations are followed by pooling of parameter estimates and adjustment of standard errors through Rubin’s Rules. An ideal semantic norm extrapolation method would yield unbiased parameter estimates, as well as confidence intervals containing those parameters with close to optimal coverage rates. We also include root-mean-squared error of parameter estimates, which collapses bias and variance of imputation methods into a single measure.

Methods

We perform two simulations. In Simulation 1 we recreate a “simple” example, with missingness in just one predictor variable conditional on the dependent variable. This accords with the likely missingness mechanism identified in Study 1, where words with longer response times more often had missing semantic norms. In Simulation 2, we consider a case where there is missingness in multiple variables, and the analysis model involves squared terms and an interaction. This allows us to explore what effect uncongeniality between imputation and analysis models have on downstream inferences.

A basic concern for imputation in this context is the presence of a high dimension set of auxiliary variables, i.e. word vectors (usually of dimension 300) versus the typically small sample sizes for semantic norms, which often exist for mere hundreds of words. For all simulations, we include a high dimensional set of auxiliary variables, to mimic the semantic norm extrapolation task. Thus, we require imputation methods suitable for

high-dimensional cases to avoid overfitting, or to allow the marginal univariate regression models in MICE to be estimable. Thankfully, we can use closely related methods to those already used for improper single regression mean imputation of semantic norms, and modify them slightly to make them suitable for use in MICE.

The most common method for semantic norm extrapolation is linear regression, with a possible ridge penalty on the coefficients (Gultchin et al., 2019; Hollis et al., 2017; Mander et al., 2015; Martínez-Huertas et al., 2020; Paetzold et al., 2016; Recchia et al., 2015; Richie et al., 2019; Thompson et al., 2018; Turton et al., 2020; Utsumi, 2018; van Paridon et al., 2019; Westbury et al., 2013). Thus, we chose a form of linear regression as our stand-in in the simulations for existing methods of extrapolating semantic norms. As well, we used a modified version of the same regression method in proper multiple imputations. We chose LASSO regression (Tibshirani, 1996), but use this method in either in 1) improper single regression mean imputation, where missing values are filled in once with predicted conditional mean values, or 2) in a proper multiple imputation, possibly including multiple imputation by chained equations in cases where multiple variables have missing values. Embedding LASSO regression into a proper multiple imputation requires modifications. One, rather than imputing the missing values with the conditional mean, we instead impute values using the conditional mean plus a random draw from a normal distribution with standard deviation equal to the standard deviation of the model’s residuals. As well, for a multiple imputation to be proper in the sense of Rubin (1996), we must incorporate uncertainty in all model parameters (intercepts, regression coefficients, residual standard deviation, the LASSO λ hyperparameter), requiring either that we use a fully Bayesian version of LASSO regression, or we perform bootstrap resampling of the observed data at each iteration of imputation of MICE for each variable. The use of bootstrapped LASSO in multiple imputation has been explored by Deng, Chang, Ido, and Long (2016) and the fully Bayesian LASSO in Zhao and Long (2016). In the presence of high dimensional auxiliary variables, LASSO outperformed multiple imputation with

unregularized linear regression in simulations in these studies. As well, Deng et al. (2016) found that rather than use the LASSO regularized coefficients themselves in the imputation, less bias occurred in simulations when LASSO was used to select variables, but unpenalized coefficient estimates were used to impute the missing values. We opted for the bootstrapped version of LASSO for computational efficiency, and include both LASSO regularized coefficient estimates, and estimates after LASSO was used to selected variables, in Simulation 1.

A further concern is that many common analyses involving megastudy data incorporate squared terms or interactions. As noted in the introduction, if an imputation model fails to take into account nonlinearities or interactions in the analysis model, it is likely to bias estimates of these effects. Recently, Bartlett, Seaman, White, Carpenter, and Alzheimer’s Disease Neuroimaging Initiative (2015) proposed a multiple imputation method, substantive model compatible fully conditional specification (SMCFCS), which is a parametric method which can provide minimally biased estimates of squared terms and interactions in a multiple imputation. Each univariate imputation performs sampling compatible with a joint distribution of the complete data implied by an explicitly stated analysis model either through rejection sampling or a Metropolis-Hastings algorithm (see Carpenter and Kenward, 2012 for a comparison). For imputations involving squared terms and interactions, we use SMCFCS. However, as default SMCFCS performs no variable selection or regularization, we modify standard SMCFCS with a variable selection step. At each iteration of MICE, for each variable with missing data, prior to performing SMCFCS imputation, we perform a bootstrap of the observed data, do LASSO regression and then include only LASSO selected variables and variables in the substantive analysis model in the SMCFCS step. This is similar to the method used by Deng et al. (2016), except SMCFCS is used after the LASSO, rather than a simple linear regression. When LASSO is used in MICE we used 5-fold cross validated mean-squared error to select the optimal λ hyperparameter at each iteration and for each imputed variable. We also use 5 cross

validation folds to select λ in the single regression mean imputation.

For all multiple imputations, we use R package MICE (van Buuren & Groothuis-Oudshoorn, 2011). MICE provides an extensible and highly customizable interface for multiple imputation, allowing for easy use of self-coded imputation functions. We use the R package glmnet implementation of LASSO (Friedman, Hastie, & Tibshirani, 2010). We use the function *mice.impute.smcfcfs* of R package miceadds (Robitzsch & Grund, 2020) when performing SMCFCFS, or a modification of this function to include a LASSO variable selection step. We also wrote our own imputation wrapper functions for LASSO multiple imputation with bootstrapping, and include all these R functions in the supplementary materials.

Simulation 1

For this simulation, we attempt to create a simple and plausible example, mimicking a case where we are analyzing response times, but the probability of a semantic norm being missing is higher for words with longer response times.

Let Z be our set of auxiliary variables. Z consists of n observations drawn from a multivariate normal distribution. The variables $z_1, \dots, z_{300}$, represent a high dimensional set of auxiliary variables, i.e. word vectors. The variable x represents the unobservable dimension of meaning that a semantic norm scale is meant to measure, and y stands in for response time. All variables have a standard deviation of 1 and a mean of 0. The variables x and y have a correlation of 0.45, and six variables $z_1, \dots, z_6$ have correlations of (.1,.2,.3,-.3,-.2,-.1) with x , and all other correlations are zero.

As semantic norms are indirect measurements of a latent variable through a Likert scale, we replace x with a proxy variable a as follows. For each x_i we take 30 draws from a normal distribution with a mean of x_i and with a standard deviation of 2. This random deviation from x_i is meant to represent participant-level random variation during semantic norm collection. We also attempt to simulate the forced-choice nature of Likert scales. We

round the 30 draws to the nearest of $(-4, -3, -2, -1, 0, 1, 2, 3, 4)$, and average them to reach a . This technically bounds a at $(-4, 4)$. The addition of this measurement error leads to an attenuation of the regression coefficient of a on y compared to x on y (0.47 compared to 0.5). It is this attenuated effect we are trying to recover through the multiple imputation. While ideally we would want to estimate the true effect of x , the issue of measurement error in the use of semantic norms has not generally been considered to date, and we discard it for the present exercise (See general discussion).

For the missingness mechanism in this simulation, we use a MAR mechanism, with missingness in a and conditioned on y . This mimics the pattern found in Study 1, where missingness in semantic norms was highly predicted by their lexical decision response times, with words with longer response times having more missing semantic norms. We therefore set probability of missingness in a to the logit transformation of y . As the amount of missingness in semantic norms tends to be large, we set a missingness proportion of 50% of a . We set the sample size per simulation n to 500 and run the simulation 1000 times. In each simulation, for the multiple imputation methods, we set the number of imputed datasets, m , to 50. Estimates from all multiple imputations are pooled using Rubin's Rules. For each iteration of the simulation, we use the following imputation and analysis methods.

1. an analysis of the effect of a on y with no missing data
2. a complete case analysis of the effect of a on y
3. a single regression mean imputation using LASSO regression of a with $z_1, \dots, z_{300}$ as predictors (this is our proxy for existing approaches to semantic norm extrapolation)
4. a multiple imputation using a bootstrapped LASSO regression, with $z_1, \dots, z_{300}$ and y , as predictors of a in the multiple imputation
5. a multiple imputation using a bootstrapped LASSO regression to select variables, with selected variables from $z_1, \dots, z_{300}$ and y , used as predictors in unregularized linear regression of a in the multiple imputation
6. a multiple imputation using a bootstrapped linear regression of a with only y as

predictor

Simulation 2. In this simulation, we explore a more complex analysis model and an imputation with MICE, involving squared terms and interactions. Z is again a high dimensional set of auxiliary variables, i.e. word vectors and consists of n observations drawn from a multivariate normal distribution with mean vector zero and standard deviation vector of 1. Variables in the analysis model are x_1 , x_2 , x_3 , and y , which represent a set of 3 kinds of semantic norms and response times. x_1 , x_2 , x_3 , are related to y by the following equation: $y = \beta_0 + \beta_1 * x_1 + \beta_2 * x_2 + \beta_3 * x_3 + \beta_4 * x_1 * x_2 + \beta_5 * x_3^2 + \epsilon$. $\beta_1 \dots \beta_5$ are set to 0.5, and the standard deviation of the error in y is 1. The three semantic norm variables x_1 , x_2 , and x_3 each have correlations of (.1,.2,.3,-.3,-.2,-.1) with six different of $z_1, \dots, z_{300}$, and correlations of .3 with each other. Unlike Simulation 1, we did not include measurement error in x_1 , x_2 and x_3 for simplicity. We set the probability of missingness in x_1 , x_2 and x_3 , to $\text{logit}(y)$, and eliminate 50% of each of those variables in each iteration of the simulation.

Because the true data generating mechanism in this simulation involves a squared term and an interaction, we compare several methods for performing multiple imputation with such terms. Three parametric alternatives for such imputation include just-another-variable (JAV) imputation, passive imputation (see *Derived Variables* in Van Buuren, 2018), and substantive model compatible fully conditional specification (SMCFCS, Bartlett et al., 2015). JAV imputation involves creating squared terms and interactions with observed data before imputation, and then using standard MICE with a separate univariate imputation model for these terms. Passive imputation computes a new product or squared terms of predictors within each iteration of the MICE algorithm, and uses these variables as predictors in the imputation. SMCFCS uses an explicit analysis model to perform sampling compatible with the joint distribution implied by the analysis model via a metropolis-Hastings algorithm. For each of these methods of imputing squared terms or interactions, we include both an imputation with auxiliary variables and a LASSO

variable selection step, or a imputation with no auxiliary variables and no LASSO variable selection. In all cases, we use the correctly specified regression equation giving the effect of $x_1, x_2, x_3, x_1 * x_2$ and x_3^2 on y . Thus the complete set of imputation/analysis methods compared in this simulation are (see Table 3 for a summary of methods)

1. a properly specified analysis of the effect of $x_1, x_2, x_3, x_1 * x_2$ and x_3^2 on y with no missing data
2. a complete case analysis of the effect of $x_1, x_2, x_3, x_1 * x_2$ and x_3^2 on y
3. a multiple imputation by chained equations using a bootstrapped LASSO regression of x_1, x_2 and x_3 with variables selected by a LASSO regression from $z_1, \dots, z_{300}, y, x_1, x_2$ and x_3 as predictors. Variables selected by the LASSO are then used in an SMCFCFS imputation with a properly specified analysis model.
4. a multiple imputation by chained equations using bootstrapped linear regression of x_1, x_2 and x_3 with $z_1, \dots, z_{300}, y$, and x_1, x_2 and x_3 as predictors
5. a multiple imputation by chained equations using default SMCFCFS with x_1, x_2, x_3 and y and no auxiliary variables.
6. a multiple imputation by chained equations using default SMCFCFS with x_1, x_2, x_3 and y and, and with the dimensions of $z_1, \dots, z_{300}$ with true nonzero correlation with x_1, x_2, x_3 as auxiliary variables. We include this method as a comparison case between regular SMCFCFS with auxiliary variables and our modification of SMCFCFS to include LASSO variable selection of the auxiliary variables. In real data analysis we would likely never know which of $z_1, \dots, z_{300}$ had true relations to x_1, x_2 , and x_3 .
7. a multiple imputation by chained equations using bootstrapped linear regression of x_1, x_2 and x_3 with x_1, x_2 and x_3 as predictors, with passive imputation of $x_1 * x_2$ and x_3^2
8. a multiple imputation by chained equations using a bootstrapped LASSO regression of x_1, x_2 and x_3 , with passive imputation of $x_1 * x_2$ and x_3^2 , and with variables selected by a LASSO regression from $z_1, \dots, z_{300}, y, x_1, x_2, x_3, x_1 * x_2$ and x_3^2 as predictors. Variables selected by the LASSO are then used in a normal linear regression imputation

with bootstrapping.

9. a multiple imputation by chained equations using bootstrapped linear regression of x_1 , x_2 and x_3 with x_1 , x_2 and x_3 as predictors, with JAV imputation of $x_1 * x_2$ and x_3^2

10. a multiple imputation by chained equations using a bootstrapped LASSO regression of x_1 , x_2 and x_3 , with JAV imputation of $x_1 * x_2$ and x_3^2 , and with variables selected by a LASSO regression from $z_1, \dots, z_{300}$, y , x_1 , x_2 , x_3 , $x_1 * x_2$ and x_3^2 as predictors. Variables selected by the LASSO are then used in a normal linear regression imputation with bootstrapping.

11. a single regression mean imputation using independent LASSO regressions of x_1 , x_2 and x_3 with $z_1, \dots, z_{300}$ as predictors (again, this is our stand in for current norm extrapolation methods)

Estimates from all multiple imputations are pooled using Rubin's Rules. These comparisons allow us to see how using a single regression mean imputation does when the imputation and analysis models are uncongenial, i.e. when the analyst is testing for relations that the imputation model cannot support. As well, comparing different specifications of SMCFCFS, passive imputation, and JAV imputation allows us to see whether a LASSO variable selection step before performing SMCFCFS, passive, or JAV imputation leads to bias in a complex analysis model. JAV and passive imputation are expected to perform worse than SMCFCFS, and conducting each variant with a LASSO variable selection step allows us to verify the effectiveness of the novel LASSO + SMCFCFS imputation we introduce.

Results

Simulation 1. Figure 10 plots the relationship between the simulated semantic norms a , the latent x and simulated response times y , from a single iteration of Simulation 1, as well as results of linear regression of a on y , x on y , and the results of a complete case analysis where only non-missing a is regressed on y . As expected for a MAR missing data

mechanism, the complete case analysis is biased. The effect of x is slightly attenuated by measurement error in a .

Table 2 displays the results of Simulation 1. With these parameters, any kind of semantic norm extrapolation yields bias and sub-optimal coverage rates. The best method is in fact multiple imputation with no auxiliary variables included at all, both in terms of bias (0), and coverage of confidence intervals (0.95). Multiple imputation with LASSO regularization performs better than an improper imputation or complete case analysis, but is biased (0.06). The approach meant to simulate extrapolating semantic norms with an improper single LASSO regression mean imputation is badly biased (-0.12), but not as badly biased as the complete case analysis (0.15). The variant of multiple imputation where LASSO is used to select variables, but imputations are drawn from an unregularized linear regression, performs better than LASSO, in terms of bias (0.06 versus 0.04) and coverage rates. (0.89 versus 0.95). Thus, the default of analyzing only complete data gives a poor estimate of the effect of a on y . The best approach to restoring inference, in this simple example, is in fact to use multiple imputation and pooling with a simple bootstrapped linear model as the imputation method and not include auxiliary information at all. This is true even though some of $z_1, \dots, z_{300}$ are correlated with x and a . Using LASSO variable selection followed by imputation by an unregularized linear regression performs second best, but with some bias.

Simulation 2. Tables 4-7 present mean bias, root mean squared error, 95% confidence interval coverage rates, and 95% confidence interval widths for methods used in Simulation 2. As expected, the complete case analysis is biased, with a lower intercept (β_0 bias = -0.11), and generally underestimating the slopes of all coefficients (biases β_1 = -0.11, β_2 = -0.11, β_3 = -0.06, β_4 = -0.07, β_5 = -0.1). MICE imputation using bootstrapped normal models for each univariate imputation is severely biased. This reflects that the imputation is uncongenial with the analysis model, which contains a squared term and an interaction. The bias of this method is worse than a complete case analysis. An improper single

regression mean imputation, our comparison case to how semantic norms are most often extrapolated, is also biased, producing an intercept and estimates of β_1 , β_2 , and β_3 that are too high (β_0 bias = 0.83, β_1 bias = 0.26, β_2 bias = 0.19, β_3 bias = 0.09). While the interaction term in this case is only slightly biased (β_0 bias = 0.04), the quadratic effect is underestimated (β_4 bias = -0.3). Passive imputation performs poorly for most coefficients, in some cases worse than the complete case analysis. JAV imputation with or without LASSO performs mostly as poorly as a complete case analysis. The best overall methods are all variants of SMCFCFS, and our modification of SMCFCFS to include LASSO variable selection performs close to, or better than, SMCFCFS with no auxiliary variables. All SMCFCFS variants are less biased than a complete case analysis.

In terms of root mean squared error of parameter estimates, which captures both bias and variance in estimating correct parameters in the downstream analysis (Table 5). LASSO + SMCFCFS and SMCFCFS with no auxiliary variables perform well, comparably or better than SMCFCFS with correctly specified auxiliary variables. It is worth noting that SMCFCFS with correct auxiliary variables is an intentionally unrealistic comparison case included for illustration only, as we rarely have ground truth information on which dimensions of word vectors have true relations to semantic norms. We include this variant to evaluate our modification of SMCFCFS to include a LASSO variable selection step.

In terms of coverage rates of 95% confidence intervals, the complete case analysis has sub-optimal coverage for all coefficients. Improper single regression mean imputation also has sub-optimal coverage for all coefficients. As well, the uncongenial MICE multiple imputation has generally poor coverage. LASSO + SMCFCFS has slightly overwide coverage of β_1 , β_2 , β_3 , and β_5 . SMCFCFS with no auxiliary variables performs well, but has slightly narrow intervals. JAV with LASSO variable selection of auxiliary variables also performs fairly well, having the closest to optimal coverage of β_3 , β_4 , and β_5 . Taken together, SMCFCFS variants have the overall best confidence interval coverage.

Discussion

We have presented two simulations plausibly recreating a common use case for semantic norms. The simulations are intended to recreate a situation in which semantic norms predict response times, but semantic norms are more likely to be missing when response time is high. As expected for a MAR missing data mechanism, in both simulations, a complete case analysis is biased. As well, using an improper single regression mean imputation tends to produce biased estimates of relations of substantive interest and confidence intervals with poor coverage.

Surprisingly, even though some dimensions of the simulated word vectors have small to moderate correlation with the missing semantic norms, there is minimal gain in efficiency and some bias introduced by the inclusion of simulated word vectors in the imputation. In both simulations, imputation methods with no auxiliary information had consistently good performance, including the least bias and the best confidence interval coverage rates in Simulation 1.

As well, in Simulation 2, a complex analysis model with a squared term and an interaction clearly demonstrates the negative consequences of uncongeniality. Both an uncongenial multiple imputation and an uncongenial single regression mean imputation produced severely biased estimates of at least one of squared or interaction terms. SMCFCFS, a recent method for performing imputation that ensures congeniality with analyses models with squared terms and interactions, performed best. Our augmented version of this imputation technique, which precedes SMCFCFS imputation with a LASSO variable selection step, performed comparably to default SMCFCFS, but was able to accommodate a high-dimensional set of auxiliary variables with a low sample size. Also, LASSO + SMCFCFS worked better than JAV or passive imputation of the interaction and squared terms, replicating previous research on SMCFCFS imputation. With slight modification, regularized regression methods such as LASSO used for semantic norm extrapolation can be adapted for inclusion in the MICE algorithm, and can provide

estimates with some bias, but still less bias than a complete case analysis.

Taken together, the two simulations show that both improper single mean regression imputation and complete case analysis tend to produce biased results, under the assumptions of the simulations. Multiple imputation, when properly specified, is a superior approach that leads to less bias and confidence intervals with more optimal coverage. This supports of view that 1) current semantic norm extrapolation methods, which are single regression mean imputations, are less effective than cross validation paints them as, and 2) complete case analysis is poor default methodological practice. We now turn to a real data example to further support this view.

Study 3: Real Data Example

In Study 2 we have presented evidence from simulations showing how the most common semantic norm extrapolation method (regularized linear regression in a single regression mean imputation) performs poorly when evaluated as a solution to a missing data problem. A complete case analysis also performed poorly. Both suffered from bias and poor confidence interval coverage in downstream analyses compared to some form of properly specified multiple imputation. In this section, we conduct an original analysis comparing results when using proper multiple imputation methods, versus using improper regression mean imputation, and show these methods produce qualitatively and quantitatively different conclusions about a substantive question involving semantic norms.

The hypothesis we explore concerns the semantic constructs of danger and usefulness (Kryuchkova, Tucker, Wurm, & Baayen, 2012; Witherell, Wurm, Seaman, Brugnone, & Fulford, 2012; Wurm, 2007, 2015; Wurm & Vakoch, 2000) and their possible interactive effect on lexical decision response times. Danger and usefulness are an alternative account of automatic evaluative effects (cf. Pratto and John, 1991) in lexical decision to valence (stimulus pleasantness) and arousal (stimulus excitingness). While valence and arousal correlate positively with usefulness and danger norms, respectively, effects of valence and

arousal in word recognition are predicted to superseded or rendered non-significant once danger and usefulness are taken into account. When a word's connotation is both high in danger and usefulness (e.g. *knife*) it leads to activation of both approach and avoidance responses which slows response times. Thus, danger and usefulness semantic norms are expected to be 1) predictive of lexical decision response times above and beyond the effects of valence and arousal 2) show an interaction, where high-danger, high-usefulness words yield slower reaction times.

Yet, in the studies on danger and usefulness cited above, only some hundreds of words have been normed and used in experiments. In contrast, studies showing effects of valence and arousal have involved many thousands of words (Kuperman et al., 2014). As well, the largest study of such effects by Kuperman et al. (2014) show that in the English Lexicon Project, valence and arousal both show interactive effects with word frequency, where low-frequency words show a more pronounced effect of both variables. Higher valence speeds reaction times, higher arousal slows them, and both more so when a word is less frequent. No studies of danger and usefulness on response times have evaluated their effects once these valence by frequency and arousal by frequency interactions are taken into account.

As shown in Study 1, a complete case analysis of megastudy data testing for semantic effects via currently published semantic norms is likely to exclude words with longer response times, longer words, and lower frequency words. Thus, effects of valence, arousal, danger, and usefulness could have also been biased in existing studies of these constructs. Taken together, an ideal study would pit an equally large sample size of danger and usefulness norms against valence and arousal norms across the range of word frequency. Lacking danger and usefulness norms for many common words, as well as valence and arousal norms for many rarer words, some kind of semantic norm extrapolation is the only option (besides collecting empirical norms for the four constructs) to potentially restore sound inference and test the hypothesized effects of valence, arousal, danger, and usefulness

among all megastudy items. As well, all critical hypothesized effects involve interactions, either valence and arousal with word frequency, or danger and usefulness with each other. An uncongenial single-regression mean imputation is likely to bias tests for these interactions. While we do not have access to ground truth, we do have prior literature and theory on the constructs of interest. In this section, we show differences in substantive conclusions reached depending on how imputation of the missing data is performed.

Materials

As our dependent variables in these analyses, we use log-transformed item-mean lexical decision response times from the English Lexicon Project (ELP) (Balota et al., 2007), British Lexicon Project (BLP) (Keuleers et al., 2012), Massive Auditory Lexical Decision Database (MALD) (Tucker et al., 2019), and the English Crowdsourcing Project (ECP) (Mandera et al., 2019). For the semantic variables of primary interest we use Warriner et al. (2013) and Dodds et al. (2011) as our sources of valence and arousal norms, and danger and usefulness norms from Wurm (2007) and Heard et al. (2019). Mandera et al. (2019) present a strong baseline model of lexical decision response times. We use this model to select control predictors: subtitle word frequency (Brysbaert, New, & Keuleers, 2012) and a quadratic effect of word frequency, word prevalence (Brysbaert et al., 2019), orthographic lexical density and phonological lexical density from Balota et al. (2007), word length and a quadratic effect of word length, number of syllables, number of morphemes from Balota et al. (2007) and Tucker et al. (2019), concreteness (Brysbaert et al., 2014), age of acquisition (Kuperman et al., 2012) and a quadratic effect of age of acquisition. These predictors are used for the visual lexical decision megastudies. For MALD, we substitute stimulus duration for word length in characters. For simplicity of interpretation, and to match norms collected on Likert scales with a different number of points, we scaled all semantic norms except for age of acquisition to a 0-1 scale by subtracting one and dividing the norms for each dataset by the range of its Likert scale.

We limited our analysis to words with vectors available in Mikolov et al. (2018) and Speer et al. (2017), once again aligning conceptnet word vectors to fastText ones using a Procrustes rotation. Two non-trivial choices must be made. The first choice concerns whether to also impute word vectors for words included in megastudies but without vectors available in Mikolov et al. (2018) and Speer et al. (2017). While this would provide a truly complete analysis of the megastudies, this would require an additional 300 imputation models at each iteration of MICE, which was unfeasible computationally. Second, we must choose whether or not to impute lexical decision response times for words not shared between the megastudies, i.e. should we impute a response time for the ELP of an item included in the ECP but not in the ELP. Indeed, technically we could impute a response time for any string with a word vector from Mikolov et al. (2018) and Speer et al. (2017), though this would include many strings not generally considered “words” for lexical decision experiments, such as numbers and emojis. For this exercise, we chose not to impute any response times, and only analyze empirical response times for each megastudy. We note that there is no reason in principle why a dependent variable should not also be imputed. The choice depends on what population generalizations are sought for. For present purposes we are therefore commenting on what happens when we fully analyze virtually all megastudy items, and we leave the more difficult question of defining the ideal set of words we would like to generalize about to future work.

For all megastudies, data was only partially observed. Number of observed data for each variable in each megastudy are given in Table 8. A striking feature of the data is that even for valence, arousal, and concreteness — the semantic norms with the largest datasets and with significant theoretical interest — between 75% to 35% of the lexical decision megastudy items lack a semantic norm. Empirical danger and usefulness norms exist for only 1-2% of the data. Once all missing data is taken into account, a mere 539 to 629 rows of the data matrix for each megastudy are complete.

Methods

We compare four analyses in this section. A complete case analysis, an analysis using improper single regression mean imputation, a multiple imputation analysis using no auxiliary information from word vectors, and a multiple imputation using features of the word vectors as auxiliary variables. We also include word duration from MALD as an auxiliary variable in multiple imputations. For all methods, the final analysis model is the same, except that 1) in the MALD analysis model, we substitute word length in characters for stimulus duration, 2) in the multiple imputation analyses we apply the analysis model to all imputed datasets and pool parameter estimates using Rubin’s Rules.

For the improper single regression mean imputation, we again use a series of independent LASSO regressions for each variable with missing data, and fill in missing values in all variables with the conditional mean of their respective model, and no pooling is performed. Each imputation model contains only features of the word vectors as predictors. As number of syllables and number of morphemes are count variables, we use a LASSO regularized Poisson regression for these variables.

For the multiple imputation models, we use SMCFCFS on LASSO selected variables when using features of the word vectors, or default SMCFCFS in models containing no features of the word vectors. SMCFCFS is employed because the full analysis model contains several interactions and squared terms. For the marginal imputation models of count variables number of syllables and number of morphemes, we subtracted one from the variable, and used an unregularized hurdle Poisson regression from package R package *countimp* (Kleinke & Reinecke, 2013). We used this imputation method for the two count variables due to convergence issues in we encountered in MICE when a LASSO regularized Poisson regression was used.⁴ All marginal models in MICE use the same set of predictors

⁴ Modelling separate processes for the shortest, versus longer words, appeared to solve the convergence issues. It could be that nonlinear or other methods may be more effective for imputing these variables. As they are not of interest here, we leave this for future work.

(minus the dependent variable for each marginal model), save for arousal, which included the square of valence as an additional predictor. We use 20 iterations of mice, and set m to 500, due to the high amount of missing data. Convergence was assessed using visual checks on variable means and standard deviations from 500 MICE chains. By 5 iterations it appeared trends in chains had largely stabilized, and we observed no visual trends in iterations 5-20. Plots of MICE chains are included in the supplementary materials. When LASSO is used in MICE we used 5-fold cross validated mean-squared error to select the optimal λ hyperparameter at each iteration and for each imputed variable. We also use 5 cross validation folds to select λ in the single regression mean imputation.

The full analysis model contains the following variables: word frequency, word frequency squared, word length/word duration, OLD, PLD, word prevalence, age of acquisition, age of acquisition squared, concreteness, number of morphemes, number of syllables, valence, arousal, danger, usefulness, an interaction of valence with word frequency, an interaction of arousal with word frequency, and an interaction of danger and usefulness. When reporting coefficient estimates from the multiple imputation analyses, we include an additional statistic λ , or the proportion of total variation attributable to missing data (Rubin, 2004). This statistic ranges from 0 to 1, with 0 indicating the coefficient estimate is completely unaffected by missing data (i.e. all data is complete and observed). λ values > 0.5 indicate the imputation model has more influence on the estimate than the complete data, and λ of 1 indicate a complete lack of information from complete data.

The complexity of this model, along with expected high correlations between many control variables, leads to some concerns about collinearity. To address this, we also perform a hierarchical regression, taking a strong baseline model, models with additional effects of key variables with this baseline model nested within them, and perform tests of model goodness of fit. We do this for both the single regression mean and multiple imputation analyses. We use the R *anova()* function for comparisons of the single regression mean imputations. For the multiple imputations, we use the *D1()* function of

package MICE. The D_1 statistic is an adaptation of the F test performed by the *anova()* function, i.e. is a test of nested models' goodness of fit, which takes into account between imputation variance in the tested variable(s), and adjusts the confidence of the test statistic accordingly (see Enders, 2010, p. 333-339). As well, this test provides an additional statistic, the relative inflation of variance, which is an estimate of the proportion of increased sample variance attributable to the missing data. Importantly, these tests are only concerned with model goodness of fit, not the interpretation of regression coefficients and their signs, and thus collinearity is not a concern.

For these comparisons, we compare five models: 1) a baseline model without any main effects or interactions of valence, arousal, danger, or usefulness, 2) the baseline model plus main effects of valence, arousal, and their interactions with word frequency, 3) the baseline model, plus main effects of danger and usefulness, 4) the baseline model, plus effects of valence, arousal, their interactions with word frequency, and danger and usefulness, and 5) the same model as 4, but with a danger and usefulness interaction. We call these models Baseline, Model A, Model B, Model C, Model D, respectively. An important theoretical question concerns whether danger and usefulness and valence and arousal have independent or overlapping effects. By comparing Model C to both Model A and Model B, we can see if both valence and arousal, and danger and usefulness, contribute uniquely to goodness of fit. As well, comparing Model D to Model C tests for the interaction of danger and usefulness above and beyond an effect of valence and arousal, and the main effects of danger and usefulness.

Results

First, we explore estimated effects of valence, arousal, danger and usefulness in the full analysis model. We follow this by model comparisons tests of goodness of fit, when these effects and included above the baseline model.

Estimated Coefficients of Semantic Variables. Figure 11 presents estimated coefficients of valence, arousal, danger, usefulness, the interaction of valence and arousal with word frequency, and a danger usefulness interaction, for the ELP, BLP, ECP, and MALD. Supplementary Tables S1-S16 contain the full list of estimated coefficients and inferential statistics for the full analysis models. Across all megastudies, confidence intervals from multiple imputation are wider than those from the single regression mean imputation. This reflects the inflation of standard errors from taking between-imputation variance into account via Rubin’s Rules. Yet, confidence intervals from the multiple imputation analyses are also much narrower than those from the complete case analysis. This indicates that, provided the imputation model is correct, taking into account the partially observed data has recovered information compared to the complete case analysis. In what follows, we discuss effects of individual variables and interactions.

Valence. The complete case analysis shows no significant main effect of valence in any megastudy ($\beta_{ELP} = 0.002$, $t(609) = 0.02$, $p = 0.9863$, $\beta_{BLP} = 0.113$, $t(519) = 1.16$, $p = 0.2468$, $\beta_{ECP} = 0.006$, $t(609) = 0.1$, $p = 0.9225$, $\beta_{MALD} = 0.085$, $t(609) = 0.39$, $p = 0.6982$). Single regression mean imputation shows a significant main effect of valence in the visual lexical decision megastudies ($\beta_{ELP} = -0.046$, $t(43310) = -2.73$, $p = 0.0063$, $\beta_{BLP} = -0.142$, $t(31607) = -7.86$, $p < 0.001$, $\beta_{ECP} = -0.045$, $t(60608) = -4.11$, $p < 0.001$) where increasing valence speeds response times, but no effect in MALD ($\beta = 0$, $t(26648) = 0.01$, $p = 0.9888$). This differs from SMCFCFS + LASSO and SMCFCFS with no word vectors as auxiliary variables which show the main effect of valence as only marginally reliable (SMCFCFS + LASSO $\beta = -0.072$, $t(428) = -1.88$, $p = 0.0613$, $\lambda = 0.96$, SMCFCFS with no auxiliary word vectors $\beta = -0.147$, $t(248) = -1.96$, $p = 0.0506$, $\lambda = 0.99$). Both multiple imputation methods show a reliable main effect in the ELP (SMCFCFS + LASSO $\beta = -0.076$, $t(631) = -2.55$, $p = 0.011$, $\lambda = 0.85$, SMCFCFS with no auxiliary word vectors ($\beta = -0.203$, $t(430) = -3.64$, $p < 0.001$, $\lambda = 0.95$) and BLP (SMCFCFS + LASSO ($\beta = -0.273$, $t(494) = -7.18$, $p < 0.001$, $\lambda = 0.91$), SMCFCFS with no auxiliary word vectors $\beta = -0.643$,

$t(394) = -14.03, p < 0.001, \lambda = 0.96$), though the multiple imputation methods estimate a larger main effect of valence compared to the single regression mean imputation. For all multiple imputation analyses, λ values are very high, indicating the imputation model has much more influence on the estimated coefficients than complete data.

Arousal. For arousal, the complete case analysis shows no significant main effect in any megastudy ($\beta_{ELP} = 0.171, t(609) = 1.27, p = 0.2035, \beta_{BLP} = 0.035, t(519) = 0.33, p = 0.7443, \beta_{ECP} = -0.03, t(609) = -0.45, p = 0.6496, \beta_{MALD} = 0.078, t(609) = 0.34, p = 0.7361$). Single regression mean imputation shows a significant main effect of arousal in the visual lexical decision megastudies ($\beta_{ELP} = 0.444, t(43310) = 17, p < 0.001, \beta_{BLP} = 0.274, t(31607) = 9.23, p < 0.001, \beta_{ECP} = 0.174, t(60608) = 9.89, p < 0.001$) where increasing arousal slows response times, but no effect in MALD ($\beta = 0.048, t(26648) = 1.03, p = 0.303$). SMCFCS + LASSO only shows a significant main effect in the ELP ($\beta = 0.223, t(501) = 5.32, p < 0.001, \lambda = 0.92$). SMCFCS with no auxiliary features from word vectors only shows a significant main effect in the BLP ($\beta = -0.189, t(351) = -3.93, p < 0.001, \lambda = 0.97$). Again, λ values are high for all megastudies for both multiple imputation methods ranging from, 0.76 to 0.99.

Usefulness. For usefulness, the complete case analysis shows a significant main effect in the visual lexical decision megastudies, and a marginal one in MALD ($\beta_{ELP} = -0.092, t(609) = -2.02, p = 0.0435, \beta_{BLP} = -0.121, t(519) = -3.29, p = 0.0011, \beta_{ECP} = -0.051, t(609) = -2.29, p = 0.0222, \beta_{MALD} = -0.146, t(609) = -1.87, p = 0.062$). Single regression mean imputation shows a significant main effect in the ELP ($\beta = -0.064, t(43310) = -2.57, p = 0.01$), BLP ($\beta = -0.205, t(31607) = -7.84, p < 0.001$), and MALD ($\beta = -0.179, t(26648) = -4.45, p < 0.001$). SMCFCS + LASSO shows a significant main effect in all megastudies ($\beta_{ELP} = -0.097, t(365) = -2.61, p = 0.0094, \lambda = 0.97, \beta_{BLP} = -0.183, t(270) = -4.01, p < 0.001, \lambda = 0.98, \beta_{ECP} = -0.207, t(259) = -4.56, p < 0.001, \lambda = 0.99, \beta_{MALD} = -0.141, t(430) = -3.15, p = 0.0018, \lambda = 0.93$). SMCFCS with no auxiliary features from word vectors also shows significant main effects of usefulness for all

megastudies ($\beta_{ELP} = -0.123$, $t(293) = -2.62$, $p = 0.0092$, $\lambda = 0.98$, $\beta_{BLP} = -0.197$, $t(202) = -3.82$, $p < 0.001$, $\lambda = 0.99$, $\beta_{ECP} = -0.26$, $t(182) = -5.23$, $p < 0.001$, $\lambda = 1$, $\beta_{MALD} = -0.184$, $t(303) = -2.74$, $p = 0.0065$, $\lambda = 0.97$). Again, λ values are very high for the multiple imputations ranging from .97 to .99.

Danger. For danger, the complete case analysis shows no significant main effect in any megastudy ($\beta_{ELP} = 0.018$, $t(609) = 0.2$, $p = 0.8429$, $\beta_{BLP} = -0.096$, $t(519) = -1.28$, $p = 0.2022$, $\beta_{ECP} = -0.075$, $t(609) = -1.74$, $p = 0.0821$, $\beta_{MALD} = -0.109$, $t(609) = -0.72$, $p = 0.4726$). Single regression mean imputation shows a significant main effect of danger in the BLP ($\beta = -0.225$, $t(31607) = -4.23$, $p < 0.001$) and the ECP ($\beta = 0.079$, $t(60608) = 2.02$, $p = 0.0439$). SMCFCFS + LASSO only shows a significant effect in the ECP ($\beta = -0.159$, $t(196) = -2.06$, $p = 0.0403$, $\lambda = 0.99$). SMCFCFS with no auxiliary variables show significant effects in the BLP ($\beta = -0.197$, $t(182) = -2.44$, $p = 0.0158$, $\lambda = 0.99$) and the ECP ($\beta = -0.169$, $t(120) = -2.41$, $p = 0.0175$, $\lambda = 1$). Again, λ values are very high for the multiple imputations ranging from .92 to .99.

Word Frequency Valence Interaction. A complete case analysis shows no significant interaction between valence and word frequency in any megastudy ($\beta_{ELP} = 0.006$, $t(609) = 0.21$, $p = 0.8326$, $\beta_{BLP} = -0.018$, $t(519) = -0.77$, $p = 0.4433$, $\beta_{ECP} = -0.003$, $t(609) = -0.2$, $p = 0.839$, $\beta_{MALD} = -0.021$, $t(609) = -0.4$, $p = 0.6869$). Single regression mean imputation yields a significant interaction in the ELP ($\beta = 0.027$, $t(31607) = 5.24$, $p < 0.001$) and BLP ($\beta = 0.027$, $t(31607) = 5.24$, $p < 0.001$). SMCFCFS + LASSO shows a significant interaction only in the BLP ($\beta = 0.068$, $t(586) = 7.06$, $p < 0.001$, $\lambda = 0.86$). SMCFCFS with no auxiliary features of word vectors shows a significant interaction in the ELP ($\beta = 0.154$, $t(558) = 18.45$, $p < 0.001$, $\lambda = 0.87$) and BLP ($\beta = 0.154$, $t(558) = 18.45$, $p < 0.001$, $\lambda = 0.87$). Figure 12 plots predicted effects of valence by word frequency according to each of the 4 analysis types. The functional shape of the interaction only matches prior literature (Kuperman et al., 2014), where among low frequency words valence more strongly speeds response times, when a form of SMCFCFS multiple imputation

is performed. Single regression mean imputation yields small interactions of inconsistent direction.

Word Frequency Arousal Interaction. A complete case analysis shows no significant interaction between arousal and word frequency in any megastudy ($\beta_{ELP} = -0.05$, $t(609) = -1.52$, $p = 0.1288$, $\beta_{BLP} = -0.022$, $t(519) = -0.84$, $p = 0.4014$, $\beta_{ECP} = 0.007$, $t(609) = 0.42$, $p = 0.6754$, $\beta_{MALD} = -0.023$, $t(609) = -0.41$, $p = 0.6851$). Single regression mean imputation shows a significant interaction in the visual lexical decision megastudies ($\beta_{ELP} = -0.05$, $t(609) = -1.52$, $p = 0.1288$, $\beta_{BLP} = -0.022$, $t(519) = -0.84$, $p = 0.4014$, $\beta_{ECP} = 0.007$, $t(609) = 0.42$, $p = 0.6754$). LASSO + SMCFCFS imputation shows a significant interaction in the ELP ($\beta = -0.07$, $t(562) = -5.97$, $p < 0.001$, $\lambda = 0.89$), ECP ($\beta = -0.038$, $t(447) = -2.6$, $p = 0.0096$, $\lambda = 0.96$), and MALD ($\beta = -0.033$, $t(841) = -2.39$, $p = 0.0173$, $\lambda = 0.72$). SMCFCFS with no auxiliary features of word vectors only shows a significant interaction in only the BLP ($\beta = 0.036$, $t(409) = 2.84$, $p = 0.0047$, $\lambda = 0.95$). Figure 13 plots predicted effects of arousal by word frequency according to each of the 4 analysis types. The functional shape of the interaction matches that of (Kuperman et al., 2014), among most analysis methods, though the interaction is reversed in the BLP when SMCFCFS imputation with no auxiliary word vectors is used.

Danger Usefulness Interaction. A complete case analysis shows no significant interaction between danger and usefulness in any megastudy ($\beta_{ELP} = 0.02$, $t(609) = 0.11$, $p = 0.9112$, $\beta_{BLP} = 0.291$, $t(519) = 1.93$, $p = 0.0544$, $\beta_{ECP} = 0.151$, $t(609) = 1.76$, $p = 0.0792$, $\beta_{MALD} = 0.226$, $t(609) = 0.74$, $p = 0.4574$). Single regression mean imputation only shows a significant interaction in the BLP ($\beta = 0.463$, $t(31607) = 3.59$, $p < 0.001$). LASSO + SMCFCFS shows a significant interaction in the BLP ($\beta = 0.393$, $t(305) = 2.42$, $p = 0.0161$, $\lambda = 0.98$) and ECP ($\beta = 0.459$, $t(308) = 3.34$, $p < 0.001$, $\lambda = 0.99$). SMCFCFS with no auxiliary features of word vectors also shows a significant interaction in the BLP ($\beta = 0.485$, $t(249) = 3.04$, $p = 0.0026$, $\lambda = 0.98$) and ECP ($\beta = 0.469$, $t(299) = 5.56$, $p < 0.001$, $\lambda = 0.99$). Figure 14 plots predicted effects of usefulness by danger according to each

of the 4 analysis types. The SMCFCs imputation methods show an interaction matching the predicted functional shape in the ELP, BLP, and MALD: usefulness slows response times when danger is high, and speeds response times when danger is low. Single regression mean imputation produces a smaller interaction save for the estimate in the BLP.

In terms of coefficient estimates, the complete case analysis shows almost no reliable effects outside of usefulness. Single regression mean imputation yields many qualitatively different effects to the multiple imputation methods. In particular, the valence by word frequency and danger by usefulness interactions do not reliably accord with their predicted functional shapes. The multiple imputation methods generally produce effects qualitatively similar to those from prior literature, though these are not uniform across the megastudies (e.g. no danger usefulness interaction in the ELP, no valence by word frequency interaction in MALD). The observed effects appear sensitive to the inclusion of word vectors as auxiliary variables in the multiple imputation, where SMCFCs with no word vector features shows a stronger word frequency valence interaction in the visual lexical decision megastudies. The multiple imputation methods are more conservative than the single imputation, and also reveal that the missing data and imputation model have very large influence on the variability in the data for all variables. This includes estimates of coefficients of valence and arousal, even though the data includes more than ten thousand empirical semantic norms. Thus, we would expect results to be very sensitive to changes to the imputation model, a fact that the single regression mean imputation obscures.

Model Comparisons. Due to possible collinearity, we also performed a series of model comparisons to determine if addition of semantic effects significantly improves model goodness of fit. As well, we assess how taking into account variance between imputations, as well as using SMCFCs in MICE as the imputation method, changes our confidence in conclusions and interpretation of the data. For the multiple imputations, we do this by performing model comparisons using the D_1 statistic, which inflates the model comparison test statistics based on the extent of between imputation variance. We compare results

obtained from this statistic to those from an anova done on the single regression mean imputed data.

Improper Single Regression Mean Imputation. Table 9 presents results of model comparisons for the LASSO improper single regression mean imputations. In the ELP, ECP and MALD, addition of valence and arousal frequency interactions, and main effects of danger and usefulness, significantly improve model goodness of fit. The only insignificant model comparison is between model D and model C, the addition of the danger usefulness interaction. In the BLP, all model comparisons are significant. However, in all megastudies, relative improvement in R^2 from the addition of semantic variables is very small, amounting to 1% to 0.1% additional variance explained.

SMCFCS With No Auxiliary Word Vector Features. Table 10 presents results of D_1 nested model comparisons of goodness of fit, when SMCFCS imputation is done with no features of word vectors included as auxiliary variables. For the ELP, addition of valence and arousal word frequency interactions, and main effects of danger and usefulness, improve model goodness of fit. The addition of a danger usefulness interaction is not supported. For the BLP, addition of all semantic effects significantly improves model goodness of fit. For the ECP, addition of valence and arousal and their interactions with word frequency does not improve model fit over baseline. An addition of main effects of danger and usefulness does significantly improve over the baseline model. A model with both valence, arousal, their interactions with word frequency, and danger and usefulness, is marginally preferred over a model with just effects of valence and arousal and their interactions, and preferred over a model with only main effects of danger and usefulness. The model with a danger usefulness interaction is preferred over the model with only main effects of danger and usefulness. For MALD, the model with valence and arousal and their interactions with word frequency is not preferred over the baseline model. A model with main effects of danger and usefulness is preferred over baseline. A model with all semantic effects is preferred over one with just valence and arousal, but not preferred over a model

with just danger and usefulness. The model containing a danger usefulness interaction is not reliably supported by the data.

SMCFCS + LASSO. Table 11 presents results of MICE multiple imputation using LASSO + SMCFCS imputation with features of word vectors as auxiliary variables. For the ELP, addition valence and arousal word frequency interactions, and main effects of danger and usefulness, improve model goodness of fit. The addition of a danger usefulness interaction is not supported. For the BLP, addition of all semantic effects significantly improve model goodness of fit. For the ECP, both models with the valence and arousal interactions and danger and usefulness main effects are preferred over baseline. However, a model with effects of valence, arousal, their interactions with word frequency, and danger and usefulness main effects, is not preferred over a model with just the effects of valence and arousal and their interactions with word frequency. A model with a danger and usefulness interaction is preferred over a model without one. For MALD, both models with the valence and arousal interactions and danger and usefulness main effects are preferred over baseline. However, a model with effects of valence, arousal, their interactions with word frequency, and danger and usefulness main effects, is not preferred over a model with just the effects of valence and arousal and their interactions with word frequency. Inclusion of the danger and usefulness interaction is not preferred over a model without this interaction.

Similar to the estimated coefficients, single regression mean imputation tends to not support the danger usefulness interaction, which is reliably supported in 2 of 4 megastudies by the multiple imputation methods. For both methods and all models, the multiple imputation indicates larger increases R^2 from the addition of semantic variables. For the BLP, the multiple imputation analyses indicate that valence, arousal, and their interaction with frequency, account for an additional 5% to 9% variance explained. Other megastudies show smaller improvements of goodness of fit, i.e. less than 1%, but still larger than the single regression mean imputation.

All told, the estimated model coefficients and models comparisons indicate that 1) the single regression mean imputation is less reliable, seems to shrink interactions towards zero, agrees with prior literature and theory less, and is less conservative than the multiple imputations, and 2) we see generally favourable results for a role of both valence, arousal, their interactions with word frequency, and danger and usefulness and their interaction, in predicting lexical decision response times. However, supported effects are not uniform over the megastudies, and are likely to be highly sensitive to changes in the imputation model due to a very large amount of variability in the parameter estimates across imputations attributable to the missing data. The model comparisons indicate that valence, arousal, their interactions with word frequency, and danger, usefulness, and their interactions with each other, are mostly accounting for unique variance in response times. Thus it appears all of valence, arousal, danger, and usefulness, may be important semantic effects in word recognition.

Discussion

In this section, we have applied complete case analysis, an improper single regression mean imputation, and two kinds of multiple imputation, to a substantive question about semantic effects in word recognition. We see generally favourable evidence for previously established interactive effects of valence and arousal with word frequency (Kuperman et al., 2014), but extend hypothesis tests of these effects to almost all items in the ELP, BLP, ECP, and MALD. As well, we see somewhat favourable evidence for an account of semantic effects in word recognition, positing that words' danger and usefulness effect lexical decision response times. Usefulness in particular seems to be an important variable for all analyses and megastudies. In two megastudies, the BLP and ECP, we see a theoretically predicted interaction between danger and usefulness leading to longer response times. In MALD, this interaction has the expected functional shape but is not reliable. Importantly, we only consistently see this interaction present and in the predicted direction in these

megastudies when we use a multiple imputation using SMCFCS to ensure congeniality between imputation and analysis models to perform semantic norm extrapolation. As well these effects survive model comparisons tests of significance even once uncertainty about the missing data is taken into account, and even though around 98% of danger and usefulness norms are missing. However, this result is highly dependent on the missing data, and therefore the imputation model being correctly specified. Thus, single regression mean imputation provides qualitatively different results compared to more principled multiple imputation methods, especially for interaction effects as expected. As expected, a single regression mean imputation is also less conservative than the multiple imputations.

General Discussion

Considerable recent evidence has been amassed that semantic judgments measured via Likert scales can be easily and accurately predicted from a combination of supervised machine learning and distributional semantic models. However, the full extent of how problematic extrapolated semantic norms are to use as covariates has gone unrecognized. Through a series of simulation and a real data example, we highlight why such approaches are problematic. For a psycholinguist the final goal of extrapolating semantic norms is putting scientific hypotheses to strong tests, even in the face of systematically missing data. Evaluating extrapolation methods by cross-validation does not directly optimize on this goal. Sound inference in the face of missing data requires careful consideration of the missing data mechanism, the question of scientific interest, the joint distribution of the data, and whether the imputation model is proper and compatible with the planned analysis. Existing methods for the task have not considered the problem from this perspective. For all methods existing in the literature, is it deductively true that they are overconfident, only testing hypotheses in one highly implausible realization of missing data, where the missing data conforms to exactly the imputation model's most likely predicted value with no error. The principles of analysis of partially observed data (Rubin, 1976,

2004) suggest single regression mean imputation will yield biased parameter estimates. Viewing the task through the lens of missing data analysis also reveals that extrapolation of missing semantic norms — when done with procedures for analysis of partially observed data — should be default methodological practice. The necessary conditions for a MCAR missing data mechanism to hold and for complete case analysis to provide unbiased estimates appear robustly violated. The words currently lacking semantic norm data are a nonrandom sample of the lexicon, and this nonrandom sampling is strongly related to a common dependent variable: lexical decision response times.

We first quantify that words most often selected for semantic norming are a highly nonrandom selection from the lexicon. Yet, the methodological default is to perform a complete case analysis and simply not analyze words for which semantic norms are missing. This represents an unacknowledged reason to be skeptical of studies on word meaning, though how problematic complete case analyses are depends on the missing data mechanism and quantity of missing data. Our results suggest psycholinguists should be extrapolating more, but current methods for extrapolation should be modified to use proper procedures for handling missing data, such as multiple imputation.

Through two simulations and one real data example, we have shown how existing extrapolation methods are problematic. Our Simulation 1 attempts to recreate a simple and plausible example, similar to a single word processing study reusing megastudy data, where a MAR missing data mechanism has eliminated some values in a predictor at high levels of the dependent variable. The simulation reveals the surprising result that using high-dimensional auxiliary variables (such as features of word vectors) in an imputation can actually bias results away from the truth. In this simple example a properly specified imputation model with no auxiliary variables achieves unbiased parameter estimation with nominal confidence interval coverage. An improper single regression mean imputation using LASSO regularization, our stand in for the most common semantic norm extrapolation method of linear regression, is biased and overconfident. While automated

variable selection via the LASSO in a proper multiple imputation provides close to unbiased inference, it is still biased. Importantly, doing the de-facto methodological default in reanalysis of megastudies of single word processing — a complete case analysis — is more badly biased than all these methods. This simulation demonstrates the need to perform a proper imputation, use an effective means of incorporating auxiliary variables (or perhaps not include them at all), and shows the problematic nature of complete case analysis when a MAR missing data mechanism is present.

Simulation 2 highlights the need to consider the congeniality of the analysis model with the imputation model. No previous creators of semantic norm extrapolation methods have treated the issue as a missing data problem, and considered whether their imputation models can possibly provide unbiased answers to specific scientific questions of interest. By simulating a true data generating mechanism and a complex analysis model involving a squared term and an interaction, Simulation 2 demonstrates that performing independent improper regression mean imputations, or an uncongenial multiple imputation, can severely bias estimates of quadratic terms and interactions, and are again overconfident. A properly specified and flexible imputation model is necessary to capture this nonlinear relation. A combination of LASSO variable selection followed by SMCFCFS imputation (Bartlett et al., 2015) extended SMCFCFS multiple imputation to a case of high dimensional auxiliary data with low sample size (as would often occur with extrapolation of small semantic norm datasets) and provides close to unbiased estimates of a squared term and an interaction in the analysis model.

Our real data example compares a strong baseline model of lexical decision response times against a model containing additional semantic effects of valence, arousal, danger, usefulness, and interactions with word frequency or between danger and usefulness. This example again shows the problematic nature of a single regression mean imputation which only uses features of word vectors and treats imputed missing data as empirical measurements. This imputation method is more confident and provides qualitatively

different estimates of some regression coefficients, particularly for interactions, compared to a multiple imputation using a combination of LASSO variable selection and SMCFCFS imputation. As well, the value of imputation is shown, given that a complete case analysis of existing, empirical data shows almost no reliable effects. Our example is the first test of effects of valence, arousal, and their interaction with word frequency, for — almost — all megastudy items for all 4 of the largest English lexical decision megastudies. Results are more consistent with prior literature (Kuperman et al., 2014; Wurm, 2007) in visual lexical decision megastudies only when a multiple imputation by SMCFCFS is used. As well, our results suggest that danger and usefulness are influential semantic effects in word recognition, above and beyond valence and arousal. However, results are not consistent across megastudies and can be influenced by the choice of including word vectors as auxiliary variables in the imputation. A danger and usefulness interaction is not supported in the ELP and a valence frequency interaction is not supported in MALD. These conclusions rest on us having correctly specified our imputation model of the missing data, due to the very high amount of missing data, but using proper procedures for missing data problems is transparent about this limitation. Accordingly, results must still be taken with a “grain of salt” compared to an analysis with empirical measurements. But, unlike all previous semantic norm extrapolations we have taken the necessary steps to ensure our imputations are compatible with the analysis we wished to undertake, we have taken uncertainty introduced by missing data into account, and the analysis is transparent about its high dependency on imputed, rather than empirical, data.

While we feel for principled reasons that multiple imputation should be the default approach to semantic norm extrapolation in psycholinguistics, what arguments against that view are available? What defenses can we mount to defend the status-quo of viewing norm extrapolation as a prediction problem? First, we should do justice to the valuable and hard work of the authors we are criticizing. Many semantic norm extrapolation methods were developed with machine learning or computational linguistic applications in

mind which might count as true prediction problems. The use case of semantic norms to analyze brain or behaviour, and concerns specifically about unbiased and appropriately confident scientific inference when using them as covariates, is just one of many possible applications. Indeed, we are indebted to these authors, as we have merely embedded the methods they discovered into multiple imputation techniques. We note that as extrapolation models approach their maximal accuracy, we would expect multiple imputation to make little difference in many cases. How problematic it is to only use word vectors as features in an imputation depends on precisely what information about the missing data is contained in the “black box” of the word vectors. As well, the non-parametric nature of the best-performing models provides some assurances against uncongeniality (though not entirely due to the omission of an analysis model dependent variable and unknowns about what information word vectors contain). Given the very high accuracy of some of the newest extrapolation methods (cf. Buechel et al., 2020) for some constructs, we would expect minimal differences between a naive analysis and a missing data analysis when these extrapolated norms are used in place of empirical ones to replace a small amount of missing data.

The degree to which a multiple imputation will differ from a “naive” single regression mean imputation analysis depends on many factors, and in some cases (e.g. a small percentage of missing data, the imputation model is extremely accurate), it will likely not produce qualitatively different results. We feel multiple imputation will yield divergent results more often among smaller datasets, with more missing data, with less reliable semantic norms, when the predictive ability of word vectors for the extrapolated semantic dimension is poorest, when more flexible models with higher variance are used in the imputation or analysis models, and when hypothesis tests concern interactions or non-linearities.

A critical question we might ask of extrapolation methods in general is: do psycholinguists really need semantic norms for rare words? If people don’t know or use

these words often, aren't they of little relevance to brain and behaviour? We think they are relevant. Just because fewer participants know a word doesn't mean that word's semantic features aren't highly relevant for the word processing of those who do know it. In lexical decision for example it is common to only analyze correct responses to attested words, which amounts to studying how meaning effects processing for the people who do know the word. Rareness is also relative to how much empirical data we have for different constructs. Danger and usefulness norms in our analysis, for example, were missing for around 98% of words, including high frequency and widely known words. Even highly studied constructs such as valence and arousal were missing for more than half of lexical decision megastudy items. As well, we should consider the word frequency Zipf distribution, and its long tail of infrequent words. We feel it is a mistake to view rare words from the tail of word frequency as not representative of the lexicon. While any individual rare word will make up tiny proportion of the probability mass of word frequency, collectively the long tail of rare words may contain a large proportion of the total word frequency probability mass. Put differently, speakers use a lot of rare words, but use each one rarely. Our own analyses here and by others, such as Kuperman et al. (2014), note that semantic effects can become more pronounced among infrequent words. Greater missingness of rare words may lead to downwardly biased estimates of semantic effects, and thus taking into account longer, less frequent, or less often known words may in fact be critical to uncovering the true role of meaning in word recognition and ensuring our experiments have external validity for how words are processed.

Other criticisms can be leveled at the design of our imputations. We note that ridge regression, not LASSO regression, is the most common regularized regression method used in semantic norm extrapolation studies. Thus there may be a concern that we are not giving fair comparison to proposed methods. LASSO and ridge regression differ in their assumptions about the coefficient vector in the regression: ridge regression assumes a normal, non-sparse distribution of coefficients, while LASSO assumes a sparse, Laplace

distribution of coefficients. We opted for LASSO regression for three reasons. First, we found in initial stages of this project that LASSO and ridge regression had all but identical performance in terms of cross-validated error for semantic norm extrapolation. Second, LASSO induced sparsity in coefficient vectors and variable selection was desirable in our modification of SMCFCFS, where selecting a small number of variables allowed SMCFCFS imputation to function in a high-dimensional case. Third, the LASSO has had favourable results in other high-dimensional multiple imputation studies (cf. Deng et al., 2016). We admit that the best regularized regression method is unknown, but feel that LASSO regression is a fair comparison for improper single regression mean imputation methods (i.e. linear and ridge regressions in Gultchin et al., 2019; Hollis et al., 2017; Mandera et al., 2015; Recchia et al., 2015; Richie et al., 2019; Thompson et al., 2018; van Paridon et al., 2019; Westbury et al., 2013). Which type of regularized regression has the more desirable bias in this context is difficult to determine, due to the black-box nature of word vectors, and we leave it for future work.

Similarly, we have chosen parametric multiple imputation methods, whereas nonparametric ones exist, (e.g. random forest imputation, Shah, Bartlett, Carpenter, Nicholas, and Hemingway, 2014). As well, we were unable to modify the most state of the art method (in terms of cross-validated error) (Buechel et al., 2020) for use in MICE, as quantifying parameter uncertainty and making probabilistic predictions with deep neural networks are active areas of research. As well, deep neural networks tend to have a large number of hyperparameters to optimize (learning rates, width and depth of hidden units, weight regularization per layer, activation functions, dropout rates) making their incorporation in MICE computationally burdensome. Our analyses were done on item-averaged lexical decision response times, rather than trial level data. Performing multiple imputation in multi-level analyses is an active area of research (c.f. Enders, Mistler, and Keller, 2016), and thus application of multi-level imputation methods to trial-level data is important future work. As well, we have glossed over the important issue

of measurement error in semantic norms. Extensions of multiple imputation offer simultaneous treatment of missing data and measurement error (Blackwell, Honaker, & King, 2012) and these extensions of multiple imputation should be considered in future work as well. Alternatives to multiple imputation exist, including inverse probability weighting (Seaman & White, 2013) and full information maximum likelihood (Enders & Bandalos, 2001). Regardless of what extrapolation methods are developed, we strongly feel these methods should be evaluated in the framework of missing data problems: based on bias and confidence interval coverage in downstream analyses and not by cross-validated accuracy.

As well, we chose to only extrapolate missing predictor variables, not to impute any missing word vectors, and not to impute any response times for words not shared between megastudies. This implies some model dependence to the results depending on selection of words for each megastudy. Future work will have to establish what the ideal selection of words should be for an imputation, and potentially impute response times and word vectors for these words. As well, the extent to which inclusion of word vectors as auxiliary variables leads to bias should be established. While we feel it is clear psycholinguists should be imputing more, the choice to include word vectors or not depends on what information about the missing data is believed to be present in the vectors as compared to variables in the analysis model. This insight also yields an important direction of future work: determining which choice of distributional semantic model yields the best performance when methods are evaluated by downstream parameter estimates rather than by cross-validation. We have also focused specifically on semantic norms (Likert-scale item mean averages) though other kinds of word features (cf. discrete feature norms, McRae, Cree, Seidenberg, and McNorgan, 2005) could also be extrapolated in a missing data framework.

As stated in the introduction both the cost of collecting empirical norms and the fact that word meaning is a “moving target,” mean we will likely never have truly

comprehensive data on word meaning. This is one reason why putting semantic norm extrapolation on firm statistical footing is important. In missing data literature, there are procedures for operating under conditions of *planned missingness*, in which some missing data is incorporated into experimental design, and the experiment is constructed to minimize biases introduced by this missing data (cf. Baraldi and Enders, 2010; T. D. Little and Rhemtulla, 2013). Future studies of word meaning should incorporate these procedures into the development of experiments and analyses when empirical data collection is unfeasible.

Conclusion

While it has become “easy” to statistically estimate responses from surveys about word meaning, putting these estimated responses on firm statistical ground is challenging. If researchers want to use extrapolated semantic norms as covariates, they should treat the extrapolation process as a missing data problem. Extrapolation methods should be evaluated based on their bias and confidence interval coverage in simulations of specific analyses and missing data mechanisms, rather than by cross-validation as has been done to this point. Surprisingly, scientific studies of word meaning should be using extrapolation more often before because words lacking semantic norms are a demonstrably non-random subset of the lexicon. The standard practice of complete case analysis of megastudy data and semantic norms is problematic and likely to yield biased results. Worse, the apparent missing data mechanism suggests that a consequence of complete case analysis may be systematic underestimation the role of meaning in word recognition, by failing to analyze words for which semantic effects are most pronounced. Multiple imputation should be default practice in semantic norm extrapolation and the analysis of megastudy data. Whether including word vectors in the imputation actually helps in this endeavor is less clear when the issue is reframed as a missing data problem. In our real data example, we replicate and extend effects of valence and arousal, and their interactions with word

frequency, to almost all items in the 4 largest English lexical decision megastudies. As well, we find evidence that usefulness may be an important predictor in lexical decision, and perhaps danger and its interaction with usefulness as well. Importantly, the hypothesized effects of these variables are put to the strong tests using recommended procedures for the analysis of partially observed data, and (mostly, with some wrinkles worthy of future work) survive them.

DRAFT

References

- Abella, R. A. S. M., & González-Nosti, M. (2019). Motor content norms for 4,565 verbs in spanish. *Behavior Research Methods*, 1–8.
- Allison, P. D. (2001). *Missing data* (Vol. 136). Sage publications.
- Alonso, M. A., Díez, E., Díez-Alamo, A. M., & Fernandez, A. (2018). Body–object interaction ratings for 750 spanish words. *Applied Psycholinguistics*, 39(6), 1239–1252.
- Alonso, M. A., Fernandez, A., & Díez, E. (2015). Subjective age-of-acquisition norms for 7,039 spanish words. *Behavior Research Methods*, 47(1), 268–274.
- Alonso, M. Á., Díez, E., & Fernandez, A. (2016). Subjective age-of-acquisition norms for 4,640 verbs in spanish. *Behavior Research Methods*, 48(4), 1337–1342.
- Altarriba, J., Bauer, L. M., & Benvenuto, C. (1999). Concreteness, context availability, and imageability ratings and word associations for abstract, concrete, and emotion words. *Behavior Research Methods, Instruments, & Computers*, 31(4), 578–602.
- Amsel, B. D., Urbach, T. P., & Kutas, M. (2012). Perceptual and motor attribute ratings for 559 object concepts. *Behavior Research Methods*, 44(4), 1028–1041.
- Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., Neely, J. H., Nelson, D. L., Simpson, G. B., & Treiman, R. (2007). The english lexicon project. *Behavior Research Methods*, 39(3), 445–459.
- Baraldi, A. N., & Enders, C. K. (2010). An introduction to modern missing data analyses. *Journal of School Psychology*, 48(1), 5–37.
- Bartlett, J. W., Seaman, S. R., White, I. R., Carpenter, J. R., & Alzheimer’s Disease Neuroimaging Initiative. (2015). Multiple imputation of covariates by fully conditional specification: Accommodating the substantive model. *Statistical Methods in Medical Research*, 24(4), 462–487.
- Benjafield, J., & Muckenheimer, R. (1989). Dates of entry and measures of imagery, concreteness, goodness, and familiarity for 1,046 words sampled from theoxford

- english dictionary. *Behavior Research Methods, Instruments, & Computers*, 21(1), 31–52.
- Bennett, S. D., Burnett, A. N., Siakaluk, P. D., & Pexman, P. M. (2011). Imageability and body–object interaction ratings for 599 multisyllabic nouns. *Behavior Research Methods*, 43(4), 1100–1109.
- Bestgen, Y. (2002). Détermination de la valence affective de termes dans de grands corpus de textes, In *Actes du colloque international sur la fouille de texte*.
- Bestgen, Y., & Vincze, N. (2012). Checking and bootstrapping lexical norms by means of word similarity indexes. *Behavior Research Methods*, 44(4), 998–1006.
- Binder, J. R., Conant, L. L., Humphries, C. J., Fernandino, L., Simons, S. B., Aguilar, M., & Desai, R. H. (2016). Toward a brain-based componential semantic representation. *Cognitive Neuropsychology*, 33(3-4), 130–174.
- Birchenough, J. M., Davies, R., & Connelly, V. (2017). Rated age-of-acquisition norms for over 3,200 german words. *Behavior Research Methods*, 49(2), 484–501.
- Blackwell, M., Honaker, J., & King, G. (2012). Multiple overimputation: A unified approach to measurement error and missing data. URL: <http://gking.harvard.edu/files/gking/files/measure.pdf>.
- Bonin, P., Méot, A., & Bugaïska, A. (2018). Concreteness norms for 1,659 French words: Relationships with other psycholinguistic variables and word recognition times. *Behavior Research Methods*, 50(6), 2366–2387.
- Bonin, P., Méot, A., Ferrand, L., & Bugaïska, A. (2015). Sensory experience ratings (sers) for 1,659 french words: Relationships with other psycholinguistic variables and visual word recognition. *Behavior Research Methods*, 47(3), 813–825.
- Borelli, E., Crepaldi, D., Porro, C. A., & Cacciari, C. (2018). The psycholinguistic and affective structure of words conveying pain. *PloS one*, 13(6), e0199658.

- Bradley, M. M., & Lang, P. J. (1999). *Affective norms for english words (ANEW): Instruction manual and affective ratings* (tech. rep.). Technical report C-1, the Center for Research in Psychophysiology.
- Briesemeister, B. B., Kuchinke, L., & Jacobs, A. M. (2011). Discrete emotion norms for nouns: Berlin affective word list (denn-bawl). *Behavior Research Methods*, 43(2), 441.
- Brown, R. L. (1994). Efficacy of the indirect approach for estimating structural equation models with missing data: A comparison of five methods. *Structural Equation Modeling: A Multidisciplinary Journal*, 1(4), 287–316.
- Brysbaert, M., & Biemiller, A. (2017). Test-based age-of-acquisition norms for 44 thousand english word meanings. *Behavior Research Methods*, 49(4), 1520–1523.
- Brysbaert, M., Mandera, P., McCormick, S. F., & Keuleers, E. (2019). Word prevalence norms for 62,000 english lemmas. *Behavior Research Methods*, 51(2), 467–479.
- Brysbaert, M., & New, B. (2009). Moving beyond kučera and francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for american english. *Behavior Research Methods*, 41(4), 977–990.
- Brysbaert, M., New, B., & Keuleers, E. (2012). Adding part-of-speech information to the subtex-us word frequencies. *Behavior Research Methods*, 44(4), 991–997.
- Brysbaert, M., Stevens, M., Mandera, P., & Keuleers, E. (2016). How many words do we know? practical estimates of vocabulary size dependent on word definition, the degree of language input and the participant’s age. *Frontiers in Psychology*, 7, 1116.
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known english word lemmas. *Behavior Research Methods*, 46(3), 904–911.
- Buechel, S., Rücker, S., & Hahn, U. (2020). Learning and evaluating emotion lexicons for 91 languages. *arXiv preprint arXiv:2005.05672*.

- Buuren, S. v., & Groothuis-Oudshoorn, K. (2010). Mice: Multivariate imputation by chained equations in r. *Journal of Statistical Software*, 1–68.
- Campos, A. (1990). Concreteness, imagery, emotionality, and interest values of words when meaning is controlled. *Perceptual and Motor Skills*, 71(2), 603–610.
- Carpenter, J., & Kenward, M. (2012). *Multiple imputation and its application*. John Wiley & Sons.
- Clark, J. M., & Paivio, A. (2004). Extensions of the Paivio, Yuille, and Madigan (1968) norms. *Behavior Research Methods, Instruments, & Computers*, 36(3), 371–383.
- Coltheart, M. (1981). The MRC psycholinguistic database. *The Quarterly Journal of Experimental Psychology Section A*, 33(4), 497–505.
- Cortese, M. J., Khanna, M. M., & Hacker, S. (2010). Recognition memory for 2,578 monosyllabic words. *Memory*, 18(6), 595–609.
- Cortese, M. J., McCarty, D. P., & Schock, J. (2015). A mega recognition memory study of 2897 disyllabic words. *The Quarterly Journal of Experimental Psychology*, 68(8), 1489–1501.
- Ćoso, B., Guasch, M., Ferré, P., & Hinojosa, J. A. (2019). Affective and concreteness norms for 3,022 croatian words. *Quarterly Journal of Experimental Psychology*, 1747021819834226.
- Della Rosa, P. A., Catricalà, E., Vigliocco, G., & Cappa, S. F. (2010). Beyond the abstract—concrete dichotomy: Mode of acquisition, concreteness, imageability, familiarity, age of acquisition, context availability, and abstractness norms for a set of 417 italian words. *Behavior Research Methods*, 42(4), 1042–1048.
- Deng, Y., Chang, C., Ido, M. S., & Long, Q. (2016). Multiple imputation for general missing data patterns in the presence of high-dimensional data. *Scientific Reports*, 6(1), 1–10.

- Dingemanse, M., & Thompson, B. (2020). Playful iconicity: Structural markedness underlies the relation between funniness and iconicity. *Language and Cognition*, 1–22.
- Díez-Álamo, A. M., Díez, E., Wojcik, D. Z., Alonso, M. A., & Fernandez, A. (2018). Sensory experience ratings for 5,500 spanish words. *Behavior Research Methods*, 1–11.
- Dodds, P. S., Harris, K. D., Kloumann, I. M., Bliss, C. A., & Danforth, C. M. (2011). Temporal patterns of happiness and information in a global social network: Hedonometrics and twitter. *PloS one*, 6(12), e26752.
- Eilola, T. M., & Havelka, J. (2010). Affective norms for 210 british english and finnish nouns. *Behavior Research Methods*, 42(1), 134–140.
- Enders, C. K. (2010). *Applied missing data analysis*. Guilford press.
- Enders, C. K. (2017). Multiple imputation as a flexible tool for missing data handling in clinical research. *Behaviour Research and Therapy*, 98, 4–18.
- Enders, C. K., & Bandalos, D. L. (2001). The relative performance of full information maximum likelihood estimation for missing data in structural equation models. *Structural Equation Modeling*, 8(3), 430–457.
- Enders, C. K., Mistler, S. A., & Keller, B. T. (2016). Multilevel multiple imputation: A review and evaluation of joint modeling and chained equations imputation. *Psychological Methods*, 21(2), 222–240.
- Engelthaler, T., & Hills, T. T. (2018). Humor norms for 4,997 english words. *Behavior Research Methods*, 50(3), 1116–1124.
- Ferrand, L., New, B., Brysbaert, M., Keuleers, E., Bonin, P., Méot, A., Augustinova, M., & Pallier, C. (2010). The french lexicon project: Lexical decision data for 38,840 french words and 38,840 pseudowords. *Behavior Research Methods*, 42(2), 488–496.

- Ferré, P., Guasch, M., Moldovan, C., & Sánchez-Casas, R. (2012). Affective norms for 380 spanish words belonging to three different semantic categories. *Behavior Research Methods*, 44(2), 395–403.
- Firth, J. R. (1957). A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis*.
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1), 1–22.
<http://www.jstatsoft.org/v33/i01/>
- Gilet, A.-L., Grün, D., Studer, J., & Labouvie-Vief, G. (2012). Valence, arousal, and imagery ratings for 835 french attributes by young, middle-aged, and older adults: The french emotional evaluation list (feel). *Revue européenne de psychologie appliquée/European Review of Applied Psychology*, 62(3), 173–181.
- Gillhooly, K. J., & Logie, R. H. (1980). Age-of-acquisition, imagery, concreteness, familiarity, and ambiguity measures for 1,944 words. *Behavior Research Methods & Instrumentation*, 12(4), 395–427.
- Graham, J. W. (2009). Missing data analysis: Making it work in the real world. *Annual Review of Psychology*, 60, 549–576.
- Grandy, T. H., Lindenberger, U., & Schmiedek, F. (2020). Vampires and nurses are rated differently by younger and older adults—age-comparative norms of imageability and emotionality for about 2500 german nouns. *Behavior Research Methods*, 1–10.
- Gultchin, L., Patterson, G., Baym, N., Swinger, N., & Kalai, A. T. (2019). Humor in word embeddings: Cockamamie gobbledegook for nincompoops. *arXiv preprint arXiv:1902.02783*.
- Hamilton, W. L., Clark, K., Leskovec, J., & Jurafsky, D. (2016). Inducing domain-specific sentiment lexicons from unlabeled corpora, In *Proceedings of the conference on empirical methods in natural language processing. conference on empirical methods in natural language processing*. NIH Public Access.

- Hardt, J., Herke, M., & Leonhart, R. (2012). Auxiliary variables in multiple imputation in regression with missing x: A warning against including too many in small sample research. *BMC Medical Research Methodology*, 12(1), 184.
- Harris, Z. S. (1954). Distributional structure. *Word*, 10(2-3), 146–162.
- Heard, A., Madan, C. R., Protzner, A. B., & Pexman, P. M. (2019). Getting a grip on sensorimotor effects in lexical–semantic processing. *Behavior Research Methods*, 51(1), 1–13.
- Hinojosa, J. A., Martínez-García, N., Villalba-García, C., Fernández-Folgueiras, U., Sánchez-Carmona, A., Pozo, M. A., & Montoro, P. R. (2016). Affective norms of 875 spanish words for five discrete emotional categories and two emotional dimensions. *Behavior Research Methods*, 48(1), 272–284.
- Hinojosa, J. A., Rincón-Pérez, I., Romero-Ferreiro, M. V., Martínez-García, N., Villalba-García, C., Montoro, P. R., & Pozo, M. A. (2016). The madrid affective database for spanish (mads): Ratings of dominance, familiarity, subjective age of acquisition and sensory experience. *PloS one*, 11(5), e0155866.
- Ho, S. M., Mak, C. W., Yeung, D., Duan, W., Tang, S., Yeung, J. C., & Ching, R. (2015). Emotional valence, arousal, and threat ratings of 160 chinese words among adolescents. *PloS one*, 10(7), e0132294.
- Hollis, G., & Westbury, C. (2016). The principals of meaning: Extracting semantic dimensions from co-occurrence models of semantics. *Psychonomic Bulletin & Review*, 23(6), 1744–1756.
- Hollis, G., & Westbury, C. (2018). When is best-worst best? a comparison of best-worst scaling, numeric estimation, and rating scales for collection of semantic norms. *Behavior Research Methods*, 50(1), 115–133.
- Hollis, G., Westbury, C., & Lefsrud, L. (2017). Extrapolating human judgments from skip-gram vector representations of word meaning. *The Quarterly Journal of Experimental Psychology*, 70(8), 1603–1619.

- Horton, N. J., & Kleinman, K. P. (2007). Much ado about nothing: A comparison of missing data methods and software to fit incomplete data regression models. *The American Statistician*, 61(1), 79–90.
- Imbir, K. K. (2015). Affective norms for 1,586 polish words (anpw): Duality-of-mind approach. *Behavior Research Methods*, 47(3), 860–870.
- Ioannidis, J. P. (2005). Why most published research findings are false. *PLoS med*, 2(8), e124.
- Janschewitz, K. (2008). Taboo, emotionally valenced, and emotionally neutral word norms. *Behavior Research Methods*, 40(4), 1065–1074.
- Juhasz, B. J., Lai, Y.-H., & Woodcock, M. L. (2015). A database of 629 english compound words: Ratings of familiarity, lexeme meaning dominance, semantic transparency, age of acquisition, imageability, and sensory experience. *Behavior Research Methods*, 47(4), 1004–1019.
- Juhasz, B. J., & Yap, M. J. (2013). Sensory experience ratings for over 5,000 mono-and disyllabic words. *Behavior Research Methods*, 45(1), 160–168.
- Keuleers, E., Lacey, P., Rastle, K., & Brysbaert, M. (2012). The british lexicon project: Lexical decision data for 28,730 monosyllabic and disyllabic english words. *Behavior Research Methods*, 44(1), 287–304.
- Kleinke, K., & Reinecke, J. (2013). Multiple imputation of incomplete zero-inflated count data. *Statistica Neerlandica*, 67(3), 311–336.
- Köper, M., & Im Walde, S. S. (2016). Automatically generated affective norms of abstractness, arousal, imageability and valence for 350 000 german lemmas, In *Proceedings of the tenth international conference on language resources and evaluation (lrec'16)*.
- Köper, M., & im Walde, S. S. (2017). Improving verb metaphor detection by propagating abstractness to words, phrases and individual senses, In *Proceedings of the 1st workshop on sense, concept and entity representations and their applications*.

- Kryuchkova, T., Tucker, B. V., Wurm, L. H., & Baayen, R. H. (2012). Danger and usefulness are detected early in auditory lexical processing: Evidence from electroencephalography. *Brain and Language*, 122(2), 81–91.
- Kuperman, V., Estes, Z., Brysbaert, M., & Warriner, A. B. (2014). Emotion and language: Valence and arousal affect word recognition. *Journal of Experimental Psychology: General*, 143(3), 1065.
- Kuperman, V., Stadthagen-Gonzalez, H., & Brysbaert, M. (2012). Age-of-acquisition ratings for 30,000 english words. *Behavior Research Methods*, 44(4), 978–990.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211–240.
- Li, M., Lu, Q., Long, Y., & Gui, L. (2017). Inferring affective meanings of words from word embedding. *IEEE Transactions on Affective Computing*, 8(4), 443–456.
- Little, R. J. (1992). Regression with missing x’s: A review. *Journal of the American Statistical Association*, 87(420), 1227–1237.
- Little, R. J., & Rubin, D. B. (1989). The analysis of social science data with missing values. *Sociological Methods & Research*, 18(2-3), 292–326.
- Little, R. J., & Rubin, D. B. (2019). *Statistical analysis with missing data* (Vol. 793). John Wiley & Sons.
- Little, T. D., & Rhemtulla, M. (2013). Planned missing data designs for developmental researchers. *Child Development Perspectives*, 7(4), 199–204.
- Ljubešić, N., Fišer, D., & Peti-Stantić, A. (2018). Predicting concreteness and imageability of words within and across languages via word embeddings. *arXiv preprint arXiv:1807.02903*.
- Lynott, D., & Connell, L. (2009). Modality exclusivity norms for 423 object properties. *Behavior Research Methods*, 41(2), 558–564.

- Lynott, D., & Connell, L. (2013). Modality exclusivity norms for 400 nouns: The relationship between perceptual experience and surface word form. *Behavior Research Methods*, 45(2), 516–526.
- Lynott, D., Connell, L., Brysbaert, M., Brand, J., & Carney, J. (2019). The lancaster sensorimotor norms: Multidimensional measures of perceptual and action strength for 40,000 english words.
- Mandera, P., Keuleers, E., & Brysbaert, M. (2015). How useful are corpus-based methods for extrapolating psycholinguistic variables? *The Quarterly Journal of Experimental Psychology*, 68(8), 1623–1642.
- Mandera, P., Keuleers, E., & Brysbaert, M. (2019). Recognition times for 62 thousand english words: Data from the english crowdsourcing project. *Behavior Research Methods*, 1–20.
- Martínez-Huertas, J. Á., Jorge-Botana, G., Luzón, J. M., & Olmos, R. (2020). Redundancy, isomorphism, and propagative mechanisms between emotional and amodal representations of words: A computational study. *Memory & Cognition*, 1–16.
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37(4), 547–559.
- Medler, D., Arnoldussen, A., Binder, J., & Seidenberg, M. (2005). The wisconsin perceptual attribute ratings database. Retrieved from [www. neuro. mcw. edu/ ratings](http://www.neuro.mcu.edu/ratings).
- Meng, X.-L. (1994). Multiple-imputation inferences with uncongenial sources of input. *Statistical Science*, 538–558.
- Miklashevsky, A. (2018). Perceptual experience norms for 506 russian nouns: Modality rating, spatial localization, manipulability, imageability and other variables. *Journal of psycholinguistic research*, 47(3), 641–661.

- Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., & Joulin, A. (2018). Advances in pre-training distributed word representations, In *Proceedings of the international conference on language resources and evaluation (lrec 2018)*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality, In *Advances in neural information processing systems*.
- Mohammad, S. M. (2018a). Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 english words, In *Proceedings of the annual conference of the association for computational linguistics (acl)*, Melbourne, Australia.
- Mohammad, S. M. (2018b). Word affect intensities, In *Proceedings of the 11th edition of the language resources and evaluation conference (lrec-2018)*, Miyazaki, Japan.
- Montefinese, M., Vinson, D., Vigliocco, G., & Ambrosini, E. (2019). Italian age of acquisition norms for a large set of words (itaoa). *Frontiers in Psychology*, 10, 278.
- Moors, A., De Houwer, J., Hermans, D., Wanmaker, S., Van Schie, K., Van Harmelen, A.-L., De Schryver, M., De Winne, J., & Brysbaert, M. (2013). Norms of valence, arousal, dominance, and age of acquisition for 4,300 dutch words. *Behavior Research Methods*, 45(1), 169–177.
- Moreno-Martínez, F. J., Montoro, P. R., & Rodríguez-Rojo, I. C. (2014). Spanish norms for age of acquisition, concept familiarity, lexical frequency, manipulability, typicality, and other variables for 820 words from 14 living/nonliving concepts. *Behavior Research Methods*, 46(4), 1088–1097.
- Morucci, P., Bottini, R., & Crepaldi, D. (2019). Augmented modality exclusivity norms for concrete and abstract italian property words.
- Mukherjee, S., & Heise, D. R. (2017). Affective meanings of 1,469 bengali concepts. *Behavior Research Methods*, 49(1), 184–197.
- Murphy, K. R. (1983). Fooling yourself with cross-validation: Single sample designs. *Personnel Psychology*, 36(1), 111–118.

- Mustillo, S., & Kwon, S. (2015). Auxiliary variables in multiple imputation when data are missing not at random. *The Journal of Mathematical Sociology*, 39(2), 73–91.
- Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The measurement of meaning*. University of Illinois press.
- Paetzold, G., & Specia, L. (2016). Inferring psycholinguistic properties of words, In *Proceedings of the 2016 conference of the north american chapter of the association for computational linguistics: Human language technologies*.
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation, In *Proceedings of the 2014 conference on empirical methods in natural language processing (emnlp)*.
- Perry, L. K., Perlman, M., & Lupyan, G. (2015). Iconicity in english and spanish and its relation to lexical category and age of acquisition. *PloS One*, 10(9), e0137147.
- Pexman, P. M., Heard, A., Lloyd, E., & Yap, M. J. (2017). The calgary semantic decision project: Concrete/abstract decision data for 10,000 english words. *Behavior Research Methods*, 49(2), 407–417.
- Pexman, P. M., Muraki, E., Sidhu, D. M., Siakaluk, P. D., & Yap, M. J. (2019). Quantifying sensorimotor experience: Body–object interaction ratings for more than 9,000 english words. *Behavior Research Methods*, 51(2), 453–466.
- Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, 61(3), 380–391.
- Recchia, G., & Louwerse, M. M. (2015). Reproducing affective norms with lexical co-occurrence statistics: Predicting valence, arousal, and dominance. *The Quarterly Journal of Experimental Psychology*, 68(8), 1584–1598.
- Redondo, J., Fraga, I., Padrón, I., & Comesaña, M. (2007). The spanish adaptation of anew (affective norms for english words). *Behavior Research Methods*, 39(3), 600–605.

- Reilly, J. (2019a). Psycholinguistic databases & corpora
[https://www.reilly-coglab.com/s/TabooWordCodingReillyLab_3222019.csv].
- Reilly, J. (2019b). Psycholinguistic databases & corpora
[https://www.reilly-coglab.com/s/JoshBrmMDS_MASTER_952014.xlsx].
- Richie, R., Zou, W., & Bhatia, S. (2019). Predicting high-level human judgment across diverse behavioral domains. *Collabra: Psychology*, 5(1).
- Robitzsch, A., & Grund, S. (2020). *Miceadds: Some additional multiple imputation functions, especially for 'mice'* [R package version 3.9-14]. R package version 3.9-14.
<https://CRAN.R-project.org/package=miceadds>
- Roest, S. A., Visser, T. A., & Zeelenberg, R. (2018). Dutch taboo norms. *Behavior Research Methods*, 50(2), 630–641.
- Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3), 581–592.
- Rubin, D. B. (1996). Multiple imputation after 18+ years. *Journal of the American statistical Association*, 91(434), 473–489.
- Rubin, D. B. (2004). *Multiple imputation for nonresponse in surveys* (Vol. 81). John Wiley & Sons.
- Schafer, J. L., & Graham, J. W. (2002). Missing data: Our view of the state of the art. *Psychological Methods*, 7(2), 147.
- Schauenburg, G., Ambrasat, J., Schröder, T., von Scheve, C., & Conrad, M. (2015). Emotional connotations of words related to authority and community. *Behavior Research Methods*, 47(3), 720–735.
- Schmidtke, D. S., Schröder, T., Jacobs, A. M., & Conrad, M. (2014). Angst: Affective norms for german sentiment terms, derived from the affective norms for english words. *Behavior Research Methods*, 46(4), 1108–1118.
- Schock, J., Cortese, M. J., & Khanna, M. M. (2012). Imageability estimates for 3,000 disyllabic words. *Behavior Research Methods*, 44(2), 374–379.

- Schock, J., Cortese, M. J., Khanna, M. M., & Toppi, S. (2012). Age of acquisition estimates for 3,000 disyllabic words. *Behavior Research Methods*, 44(4), 971–977.
- Schönemann, P. H. (1966). A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1), 1–10.
- Scott, G. G., Keitel, A., Becirspahic, M., Yao, B., & Sereno, S. C. (2018). The glasgow norms: Ratings of 5,500 words on nine scales. *Behavior Research Methods*, 1–13.
- Seaman, S. R., & White, I. R. (2013). Review of inverse probability weighting for dealing with missing data. *Statistical Methods in Medical Research*, 22(3), 278–295.
- Shah, A. D., Bartlett, J. W., Carpenter, J., Nicholas, O., & Hemingway, H. (2014). Comparison of random forest and parametric imputation models for imputing missing data using MICE: A caliber study. *American Journal of Epidemiology*, 179(6), 764–774.
- Sianipar, A., van Groenestijn, P., & Dijkstra, T. (2016). Affective meaning, concreteness, and subjective frequency norms for indonesian words. *Frontiers in Psychology*, 7, 1907.
- Sidhu, D. M., Kwan, R., Pexman, P. M., & Siakaluk, P. D. (2014). Effects of relative embodiment in lexical and semantic processing of verbs. *Acta Psychologica*, 149, 32–39.
- Snefjella, B., & Kuperman, V. (2016). It’s all in the delivery: Effects of context valence, arousal, and concreteness on visual word processing. *Cognition*, 156, 135–146.
- Soares, A. P., Costa, A. S., Machado, J., Comesaña, M., & Oliveira, H. M. (2017). The minho word pool: Norms for imageability, concreteness, and subjective frequency for 3,800 portuguese words. *Behavior Research Methods*, 49(3), 1065–1081.
- Söderholm, C., Häyry, E., Laine, M., & Karrasch, M. (2013). Valence and arousal ratings for 420 finnish nouns by age and gender. *PloS one*, 8(8), e72859.
- Speed, L. J., & Majid, A. (2017). Dutch modality exclusivity norms: Simulating perceptual modality in space. *Behavior Research Methods*, 49(6), 2204–2218.

- Speer, R., Chin, J., & Havasi, C. (2017). ConceptNet 5.5: An open multilingual graph of general knowledge.
<http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14972>
- Stadthagen-Gonzalez, H., Imbault, C., Sánchez, M. A. P., & Brysbaert, M. (2017). Norms of valence and arousal for 14,031 spanish words. *Behavior Research Methods*, 49(1), 111–123.
- Stadthagen-González, H., Ferré, P., Pérez-Sánchez, M. A., Imbault, C., & Hinojosa, J. A. (2018). Norms for 10,491 spanish words for five discrete emotions: Happiness, disgust, anger, fear, and sadness. *Behavior Research Methods*, 50(5), 1943–1952.
- Stevenson, R. A., Mikels, J. A., & James, T. W. (2007). Characterization of the affective norms for english words by discrete emotional categories. *Behavior Research Methods*, 39(4), 1020–1024.
- Stratton, R. P., Jacobus, K. A., & Brinley, B. (1975). Age-of-acquisition, imagery, familiarity and meaningfulness norms for 543 words. *Behavior Research Methods & Instrumentation*, 7(1), 1–6.
- Syssau, A., Yakhloufi, A., Giudicelli, E., Monnier, C., & Anders, R. (2020). Fancat: French affective norms for ten emotional categories. *Behavior Research Methods*, 1–19.
- Sze, W. P., Liow, S. J. R., & Yap, M. J. (2014). The chinese lexicon project: A repository of lexical decision behavioral responses for 2,500 chinese characters. *Behavior Research Methods*, 46(1), 263–273.
- Thompson, B., & Lupyan, G. (2018). Automatic estimation of lexical concreteness in 77 languages, In *The 40th annual conference of the cognitive science society (cogsci 2018)*. Cognitive Science Society.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.
- Tillotson, S. M., Siakaluk, P. D., & Pexman, P. M. (2008). Body—object interaction ratings for 1,618 monosyllabic nouns. *Behavior Research Methods*, 40(4), 1075–1078.

- Toglia, M. P. (2009). Withstanding the test of time: The 1978 semantic word norms. *Behavior Research Methods*, 41(2), 531–533.
- Tse, C.-S., Yap, M. J., Chan, Y.-L., Sze, W. P., Shaoul, C., & Lin, D. (2017). The chinese lexicon project: A megastudy of lexical decision performance for 25,000+ traditional chinese two-character compound words. *Behavior Research Methods*, 49(4), 1503–1519.
- Tucker, B. V., Brenner, D., Danielson, D. K., Kelley, M. C., Nenadić, F., & Sims, M. (2019). The massive auditory lexical decision (mald) database. *Behavior Research Methods*, 51(3), 1187–1204.
- Turton, J., Vinson, D., & Smith, R. (2020). Extrapolating binder style word embeddings to new words, In *Proceedings of the second workshop on linguistic and neurocognitive resources*.
- Utsumi, A. (2018). A neurobiologically motivated analysis of distributional semantic models. *arXiv preprint arXiv:1802.01830*.
- Van Buuren, S. (2018). *Flexible imputation of missing data*. CRC press.
- van Buuren, S., & Groothuis-Oudshoorn, K. (2011). mice: Multivariate imputation by chained equations in r. *Journal of Statistical Software*, 45(3), 1–67.
<https://www.jstatsoft.org/v45/i03/>
- Van Rensbergen, B., De Deyne, S., & Storms, G. (2016). Estimating affective word covariates using word association data. *Behavior Research Methods*, 48(4), 1644–1652.
- van Ginkel, J. R., Linting, M., Rippe, R. C., & van der Voort, A. (2019). Rebutting existing misconceptions about multiple imputation as a method for handling missing data. *Journal of Personality Assessment*, 1–12.
- van Paridon, J., & Thompson, B. (2019). Subs2vec: Word embeddings from subtitles in 55 languages.

- Vergallito, A., Petilli, M. A., & Marelli, M. (2019). Perceptual modality norms for 1,121 italian words: A comparison with concreteness and imageability scores and an analysis of their impact in word processing tasks. *Behavior Research Methods*, 1–18.
- Verheyen, S., De Deyne, S., Linsen, S., & Storms, G. (2019). Lexicosemantic, affective, and distributional norms for 1,000 dutch adjectives. *Behavior Research Methods*, 1–14.
- Walker, H. J. (1970). Imagery ratings for 338 nouns. *Behavior Research Methods & Instrumentation*, 2(4), 165–167.
- Wang, J., Yu, L.-C., Lai, K. R., & Zhang, X. (2016). Locally weighted linear regression for cross-lingual valence-arousal prediction of affective words. *Neurocomputing*, 194, 271–278.
- Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior Research Methods*, 45(4), 1191–1207.
- Westbury, C. F., Shaoul, C., Hollis, G., Smithson, L., Briesemeister, B. B., Hofmann, M. J., & Jacobs, A. M. (2013). Now you see it, now you don't: On emotion, context, and the algorithmic prediction of human imageability judgments. *Frontiers in Psychology*, 4, 991.
- Wilkinson, L. (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist*, 54(8), 594.
- Winter, B., Perlman, M., Perry, L. K., & Lupyan, G. (2017). Which words are most iconic? *Interaction Studies*, 18(3), 443–464.
- Witherell, D., Wurm, L. H., Seaman, S. R., Brugnone, N. A., & Fulford, E. T. (2012). Danger and usefulness effects as a function of concept ancientness. *The Mental Lexicon*, 7(2), 183–209.
- Wood, S. (2017). *Generalized additive models: An introduction with r* (2nd ed.). Chapman; Hall/CRC.

- Wood, S., N., Pya, & Säfken, B. (2016). Smoothing parameter and model selection for general smooth models (with discussion). *Journal of the American Statistical Association*, 111, 1548–1575.
- Wurm, L. H. (2007). Danger and usefulness: An alternative framework for understanding rapid evaluation effects in perception? *Psychonomic Bulletin & Review*, 14(6), 1218–1225.
- Wurm, L. H. (2015). Auditory and visual processing of novel stimuli are affected by subjective connotations of danger and usefulness. *The Mental Lexicon*, 10(1), 1–31.
- Wurm, L. H., & Vakoch, D. A. (2000). The adaptive value of lexical connotation in speech perception. *Cognition & Emotion*, 14(2), 177–191.
- Yao, Z., Wu, J., Zhang, Y., & Wang, Z. (2017). Norms of valence, arousal, concreteness, familiarity, imageability, and context availability for 1,100 chinese words. *Behavior Research Methods*, 49(4), 1374–1385.
- Yap, M. J., Liow, S. J. R., Jalil, S. B., & Faizal, S. S. B. (2010). The malay lexicon project: A database of lexical statistics for 9,592 words. *Behavior Research Methods*, 42(4), 992–1003.
- Yee, L. T. (2017). Valence, arousal, familiarity, concreteness, and imageability ratings for 292 two-character chinese nouns in cantonese speakers in hong kong. *PloS one*, 12(3), e0174569.
- Yoo, J. E. (2009). The effect of auxiliary variables and multiple imputation on parameter estimation in confirmatory factor analysis. *Educational and Psychological Measurement*, 69(6), 929–947.
- Zhao, Y., & Long, Q. (2016). Multiple imputation in the presence of high-dimensional data. *Statistical Methods in Medical Research*, 25(5), 2021–2035.

Table 2

Results of Simulation 1.

	Bias	RMSE	CI Coverage	CI Width
Full Data	0 (0.05)	0.05 (0.05)	0.95	0.18
Complete Case	0.15 (0.05)	0.16 (0.13)	0.20	0.22
Improper LASSO Regression	-0.12 (0.06)	0.13 (0.12)	0.54	0.24
Multiple Imputation LASSO	0.06 (0.08)	0.1 (0.12)	0.89	0.31
Multiple Imputation, Lasso Selected Variables	0.04 (0.07)	0.09 (0.11)	0.95	0.33
Multiple Imputation, No Auxiliary Variables	0 (0.07)	0.07 (0.08)	0.95	0.27

Table 3

Imputation methods compared in Simulation 2.

	Imputation Type	LASSO Variable Selection	Quadratic/Interaction Imputation Method	Auxiliary Variables Included
Full Data	none			
Complete Case	none			
LASSO + SMCFCFS	multiple	yes	SMCFCFS	yes
MICE, bootstrap normal	multiple	no	none	yes
SMCFCFS, no auxiliary variables	multiple	no	SMCFCFS	no
SMCFCFS, correct auxiliary variables	multiple	no	SMCFCFS	only true $ r > 0$
MICE + Passive Imputation, no auxiliary variables	multiple	yes	passive imputation	no
MICE LASSO, passive imputation	multiple	no	passive imputation	yes
MICE LASSO, JAV	multiple	yes	JAV	yes
MICE, JAV	multiple	no	JAV	no
LASSO Improper Single Regression Mean	single	yes	none	yes

Table 4

Mean bias of imputation methods compared in Simulation 2.

	β_0	β_1	β_2	β_3	β_4	β_5
Full Data	0.004	0	0	-0.002	0	-0.001
Complete Case	-0.871	-0.105	-0.112	-0.06	-0.074	-0.1
LASSO + SMCFCFS	0.006	0.01	-0.03	0.02	0.054	-0.034
MICE, bootstrap normal	0.509	-0.493	-0.192	-0.103	-0.193	-0.498
SMCFCFS, no auxiliary variables	-0.072	0.012	0.006	0.019	0.029	0.023
SMCFCFS, correct auxiliary variables	-0.066	0.007	0.003	0.023	0.035	0.025
MICE + Passive Imputation, no auxiliary variables	0.368	0.042	0.039	-0.067	-0.342	-0.212
MICE LASSO, passive imputation	0.414	0.111	0.005	0.017	-0.21	-0.207
MICE LASSO, JAV	0.078	0.125	0.025	0.069	-0.086	-0.006
MICE, JAV	-0.081	0.133	0.13	0.047	0.089	0.041
LASSO Improper Single Regression Mean	0.831	0.255	0.191	0.094	-0.298	0.043

Bold numbers indicate least biased coefficient estimates.

Table 5

Root mean squared error of imputation methods compared in Simulation 2

	β_0	β_1	β_2	β_3	β_4	β_5
Full Data	0.057	0.048	0.047	0.049	0.032	0.045
Complete Case	0.881	0.156	0.161	0.182	0.167	0.159
LASSO + SMCFCFS	0.088	0.063	0.078	0.081	0.079	0.08
MICE, bootstrap normal	0.515	0.493	0.2	0.116	0.197	0.498
SMCFCS, no auxiliary variables	0.118	0.078	0.078	0.089	0.066	0.086
SMCFCS, correct auxiliary variables	0.112	0.075	0.077	0.088	0.069	0.086
MICE + Passive Imputation, no auxiliary variables	0.375	0.099	0.101	0.13	0.345	0.216
MICE LASSO, passive imputation	0.426	0.137	0.091	0.093	0.22	0.211
MICE LASSO, JAV	0.136	0.151	0.098	0.12	0.167	0.079
MICE, JAV	0.152	0.164	0.164	0.13	0.147	0.107
LASSO Improper Single Regression Mean	0.834	0.271	0.216	0.153	0.31	0.123

Bold numbers indicate lowest (best) root mean squared error.

Table 6

95% Confidence interval coverage rates of imputation methods compared in Simulation 2.

	β_0	β_1	β_2	β_3	β_4	β_5
Full Data	0.94	0.95	0.96	0.94	0.95	0.95
Complete Case	0	0.76	0.73	0.86	0.82	0.76
LASSO + SMCFCFS	0.95	0.98	0.97	0.96	0.87	0.97
MICE, bootstrap normal	0	0	0.49	0.98	0.33	0
SMCFCS, no auxiliary variables	0.83	0.95	0.95	0.92	0.91	0.91
SMCFCS, correct auxiliary variables	0.84	0.96	0.95	0.91	0.89	0.91
MICE + Passive Imputation, no auxiliary variables	0.01	0.95	0.95	0.93	0.01	0.11
MICE LASSO, passive imputation	0.02	0.84	0.98	0.98	0.59	0.53
MICE LASSO, JAV	0.91	0.77	0.97	0.95	0.91	0.96
MICE, JAV	0.83	0.64	0.67	0.9	0.81	0.89
LASSO Improper Single Regression Mean	0	0.16	0.55	0.84	0.03	0.93

Bold numbers indicate 95% confidence intervals with closest to nominal coverage.

Table 7

Width of 95% confidence intervals of imputation methods compared in Simulation 2.

	β_0	β_1	β_2	β_3	β_4	β_5
Full Data	0.22	0.19	0.19	0.19	0.13	0.17
Complete Case	0.44	0.39	0.39	0.54	0.45	0.39
LASSO + SMCFCS	0.35	0.29	0.36	0.35	0.24	0.33
MICE, bootstrap normal	0.40	0.05	0.38	0.43	0.34	0.06
SMCFCS, no auxiliary variables	0.33	0.30	0.30	0.33	0.23	0.29
SMCFCS, correct auxiliary variables	0.33	0.30	0.30	0.33	0.22	0.29
MICE + Passive Imputation, no auxiliary variables	0.37	0.40	0.40	0.45	0.34	0.30
MICE LASSO, passive imputation	0.44	0.49	0.43	0.47	0.38	0.32
MICE LASSO, JAV	0.47	0.37	0.43	0.48	0.64	0.34
MICE, JAV	0.44	0.35	0.35	0.45	0.41	0.34
LASSO Improper Single Regression Mean	0.30	0.33	0.40	0.43	0.29	0.45

Table 8

Total number of megastudy response times, versus number of observed cases of the predictors. Percentage of RTs with observations for each variable are given in parentheses.

	ELP	BLP	ECP	MALD
Danger	778 (1.8%)	664 (2.1%)	779 (1.28%)	660 (2.47%)
Usefulness	778 (1.8%)	664 (2.1%)	779 (1.28%)	660 (2.47%)
Arousal	13352 (30.81%)	8335 (26.35%)	14269 (23.54%)	10847 (40.67%)
Valence	17209 (39.72%)	10806 (34.17%)	16104 (26.56%)	13720 (51.45%)
Number of Morphemes	23579 (54.42%)	13234 (41.84%)	17552 (28.95%)	24642 (92.4%)
Age of Acquisition	21364 (49.31%)	12213 (38.62%)	30764 (50.74%)	15062 (56.48%)
Duration	25312 (58.42%)	14010 (44.3%)	18936 (31.23%)	26669 (100%)
Concreteness	24079 (55.57%)	13397 (42.36%)	37205 (61.37%)	17490 (65.58%)
OLD	39158 (90.37%)	18968 (59.97%)	27386 (45.17%)	25309 (94.9%)
PLD	39158 (90.37%)	18968 (59.97%)	27386 (45.17%)	25309 (94.9%)
Word Prevalence	27640 (63.79%)	16254 (51.39%)	60628 (100%)	18937 (71.01%)
Subtitle Word Frequency	27640 (63.79%)	16254 (51.39%)	60628 (100%)	18937 (71.01%)
Number of Syllables	39528 (91.23%)	28223 (89.24%)	37605 (62.03%)	26669 (100%)
Length	43330 (100%)	31627 (100%)	60628 (100%)	26669 (100%)
Total Megastudy Response Times	43330	31627	60628	26669
Complete Cases	629	539	629	629

Table 9

Model comparisons using improper single regression mean imputation.

Comparison	Residual Degrees of Freedom	Degrees of Freedom	F	P Value	R^2
ELP					
Baseline	43317				0.579
Model A > Baseline	43313	4	142.42	< 0.001	0.584
Model B > Baseline	43315	2	71.88	< 0.001	0.580
Model C > Model A	43311	4	121.83	< 0.001	0.585
Model C > Model B	43311	2	31.10	< 0.001	0.585
Model D > Model C	43310	1	0.49	0.4825	0.585
BLP					
Baseline	31614				0.335
Model A > Baseline	31610	4	64.26	< 0.001	0.340
Model B > Baseline	31612	2	113.31	< 0.001	0.340
Model C > Model A	31608	4	50.61	< 0.001	0.344
Model C > Model B	31608	2	86.03	< 0.001	0.344
Model D > Model C	31607	1	12.88	< 0.001	0.344
ECP					
Baseline	60615				0.860
Model A > Baseline	60611	4	78.37	< 0.001	0.860
Model B > Baseline	60613	2	64.80	< 0.001	0.860
Model C > Model A	60609	4	48.10	< 0.001	0.860
Model C > Model B	60609	2	4.43	0.0119	0.860
Model D > Model C	60608	1	1.38	0.2408	0.860
MALD					
Baseline	26655				0.264
Model A > Baseline	26651	4	7.55	< 0.001	0.265
Model B > Baseline	26653	2	45.97	< 0.001	0.267
Model C > Model A	26649	4	4.65	< 0.001	0.267
Model C > Model B	26649	2	40.15	< 0.001	0.267
Model D > Model C	26648	1	0.03	0.8575	0.267

Table 10

Model comparisons using LASSO + SMCFCs multiple imputation, no auxiliary word vectors.

Comparison	Degrees of Freedom	Residual Degrees of Freedom	Relative Inflation of Variance	D_1	P Value	Pooled R^2	2.5 Quantile of R^2	Upper 97.5 Quantile of R^2
ELP								
Baseline						0.640	0.633	0.646
Model A > Baseline	4	2482	6.65	44.84	< 0.001	0.651	0.644	0.658
Model B > Baseline	2	997	70.35	10.63	< 0.001	0.653	0.641	0.664
Model C > Model A	4	2128	16.53	8.49	< 0.001	0.658	0.648	0.668
Model C > Model B	2	991	89.51	4.10	0.0169	0.658	0.648	0.668
Model D > Model C	1	509	53.25	0.07	0.7962	0.659	0.648	0.669
BLP								
Baseline						0.491	0.480	0.502
Model A > Baseline	4	2252	9.53	149.53	< 0.001	0.577	0.558	0.595
Model B > Baseline	2	981	84.53	17.91	< 0.001	0.539	0.508	0.569
Model C > Model A	4	2042	20.54	52.61	< 0.001	0.598	0.574	0.621
Model C > Model B	2	976	100.37	7.27	< 0.001	0.598	0.574	0.621
Model D > Model C	1	502	64.44	9.25	0.0025	0.607	0.578	0.634
ECP								
Baseline						0.861	0.859	0.864
Model A > Baseline	4	1979	67.84	1.06	0.3724	0.863	0.860	0.866
Model B > Baseline	2	974	349.25	14.45	< 0.001	0.884	0.866	0.899
Model C > Model A	4	1965	88.03	2.38	0.0501	0.886	0.867	0.902
Model C > Model B	2	971	402.49	13.23	< 0.001	0.886	0.867	0.902
Model D > Model C	1	504	82.84	30.93	< 0.001	0.891	0.870	0.909
MALD								
Baseline						0.288	0.278	0.298
Model A > Baseline	4	2976	3.37	1.67	0.1531	0.289	0.279	0.299
Model B > Baseline	2	1016	33.12	5.17	0.0058	0.299	0.282	0.316
Model C > Model A	4	2349	7.32	0.62	0.6518	0.300	0.283	0.317
Model C > Model B	2	1005	40.59	4.14	0.0162	0.300	0.283	0.317
Model D > Model C	1	527	25.66	2.80	0.0947	0.303	0.285	0.321

Table 11

Model comparisons using LASSO + SMCFCs multiple imputation, auxiliary word vectors.

Comparison	Degrees of Freedom	Residual Degrees of Freedom	Relative Inflation of Variance	D_1	P Value	Pooled R^2	2.5 Quantile of R^2	Upper 97.5 Quantile of R^2
ELP								
Baseline						0.617	0.611	0.623
Model A > Baseline	4	2769	4.58	44.73	< 0.001	0.626	0.619	0.633
Model B > Baseline	2	1032	33.39	11.29	< 0.001	0.625	0.616	0.633
Model C > Model A	4	2404	7.70	18.32	< 0.001	0.630	0.622	0.638
Model C > Model B	2	1024	37.98	5.39	0.0047	0.630	0.622	0.638
Model D > Model C	1	523	31.92	0.02	0.8916	0.631	0.623	0.638
BLP								
Baseline						0.397	0.386	0.408
Model A > Baseline	4	2157	12.72	20.14	< 0.001	0.418	0.405	0.432
Model B > Baseline	2	1006	43.49	13.58	< 0.001	0.421	0.401	0.440
Model C > Model A	4	2108	15.29	10.61	< 0.001	0.434	0.416	0.453
Model C > Model B	2	1003	46.70	8.36	< 0.001	0.434	0.416	0.453
Model D > Model C	1	512	42.12	5.86	0.0158	0.439	0.418	0.461
ECP								
Baseline						0.863	0.860	0.865
Model A > Baseline	4	2068	27.65	2.92	0.02	0.864	0.861	0.866
Model B > Baseline	2	988	168.96	7.83	< 0.001	0.869	0.863	0.876
Model C > Model A	4	2033	36.66	2.33	0.0536	0.870	0.863	0.877
Model C > Model B	2	986	184.68	7.21	< 0.001	0.870	0.863	0.877
Model D > Model C	1	506	77.58	11.19	< 0.001	0.873	0.863	0.881
MALD								
Baseline						0.281	0.271	0.291
Model A > Baseline	4	4536	1.52	3.45	0.0081	0.282	0.272	0.292
Model B > Baseline	2	1103	13.19	7.63	< 0.001	0.287	0.276	0.299
Model C > Model A	4	3128	2.99	1.41	0.2269	0.288	0.277	0.300
Model C > Model B	2	1078	15.89	6.05	0.0024	0.288	0.277	0.300
Model D > Model C	1	594	9.41	2.03	0.1544	0.289	0.277	0.301

Figure 1. Number of words and number of constructs semantically normed by year. Size of dots indicates number of semantic constructs/scales included in the dataset.

Figure 2. Number of words normed for unique semantic constructs, order by the rank of their size.

Figure 3. Histogram of proportions of semantic norm datasets missing word from Brysbaert, Mandera, McCormick, and Keuleers (2019).

Figure 4. Example words group by ranges of probability of being missing. Words in the upper part of the figure are a random sample from that range of proportion of scales/constructs missing the word. Shade is for readability only.

Figure 5. Proportion of constructs/scales missing word by word frequency (Zipf Score).

Black dashed line is a semiparametric penalized thin plate spline from a binomial generalized additive model. Smaller dashed lines are 95% confidence intervals for the smooth. Shaded areas are hexagonal bins of the 2d space, with darker areas indicating where the data is more concentrated.

Figure 6. Proportion of constructs/scales missing word by word prevalence. Black dashed line is a semiparametric penalized thin plate spline from a binomial generalized additive model. Smaller dashed lines are 95% confidence intervals for the smooth. Shaded areas are hexagonal bins of the 2d space, with darker areas indicating where the data is more concentrated.

Figure 7. Proportion of Datasets missing constructs/scales by word length. Black dashed line is a semiparametric penalized thin plate spline from a binomial generalized additive model. Smaller dashed lines are 95% confidence intervals for the smooth. Shaded areas are hexagonal bins of the 2d space, with darker areas indicating where the data is more concentrated.

Figure 8. Proportion of Scales/Constructs missing word by concreteness. Black dashed line is a semiparametric penalized thin plate spline from a binomial generalized additive model. Smaller dashed lines are 95% confidence intervals for the smooth. Shaded areas are hexagonal bins of the 2d space, with darker areas indicating where the data is more concentrated.

Figure 9. Proportion of scales/constructs missing word by z-scored response time. Black dashed line is a semiparametric penalized thin plate spline from a binomial generalized additive model. Smaller dashed lines are 95% confidence intervals for the smooth. Shaded areas are hexagonal bins of the 2d space, with darker areas indicating where the data is more concentrated.

Figure 10. Illustration of bias introduced by missing data from a single iteration of simulation 1. Circles represent observed response times with a missing semantic norm, crosses represent cases with observed response times and observed semantic norms. Size is proportional to probability of being missing. Lines are estimated regression coefficients estimated from true latent semantic norms (dotted line), estimated coefficient of semantic norms with measurement error (solid line), and estimated coefficient of semantic norm when using a complete case analysis.

Figure 11. Coefficient estimates of critical variables. Bars are 95% confidence intervals of the coefficients. The complete case analysis yields very large confidence intervals mostly overlapping zero. The multiple imputations yield wider confidence than the single regression mean imputation, reflecting the correction of the standard errors for between imputation variance.

Figure 12. Estimated interaction of valence and word Frequency. Darker lines indicate lower frequency bands. In the BLP, ELP, and MALD, valence decreases response time more for low frequency words, but only when LASSO + SMCFCS is used as the imputation method. Single regression mean imputation yields an inconsistent interaction.

Figure 13. Estimated interaction of arousal and word frequency. Darker lines indicate lower frequency bands. In all megastudies, an increase in arousal increases response time more for low frequency words. The effect for low frequency words is larger when LASSO + SMCFCS used for the imputation.

Figure 14. Estimated interaction of danger and usefulness. Darker lines indicate lower danger values. For LASSO + SMCFCS, the expected danger usefulness interaction occurs in the BLP, ECP, and MALD (though the interaction is not significant in MALD, $p > 0.05$), where among higher danger words an increase in usefulness slows response times, but an increase in usefulness speeds up response times for low danger words. The interaction is inconsistent and smaller when single regression mean imputation is used.

Appendix A
Supplementary Material

DRAFT

Appendix B

Plugin Imputation Functions for MICE

This is the plugin function used in MICE imputation with LASSO + SMCFCs. It requires R packages *glmnet* and *miceadds*.

```
mice.impute.glmnet.smcfc <-
  function(y,
           ry,
           x,
           wy = NULL,
           sm, ### formula for the substantive model to pass to smcfc
           dep_type = "norm",
           sm_type = "norm",
           fac_sd_proposal = 1,
           mh_iter = 20,
           alpha = 1,
           ### alpha = 1 for glmnet, LASSO
           nfold = 5,
           ### number of cross validation folds to find optimal lambda
           ...) {
  if (is.null(wy)) {
    wy <- !ry
  }

  ### do a bootstrap, build model matrices for glmnet
  n1 <- sum(ry)
  s <- sample(n1, n1, replace = TRUE)
  xobs <- as.matrix(x)[ry, , drop = F][s, ]
```

```

yobs <- y[ry][s]

keep_vars <-
  all.vars(sm) ### get a vector of all variables in the substantive model

mod <- cv.glmnet(
  x = xobs,
  y = yobs,
  nfolds = nfolds,
  alpha = alpha,
  penalty_factor = ifelse(colnames(x) %in% keep_vars, 0, 1) ### set penalization of
)

nzero <-
  unique(c(
    which(colnames(x) %in% keep_vars),
    predict(mod, type = "nonzero")$X1
  )) ### get which variables have nonzero coefficients

#### now do a run of smcfcs, but only include substantive model variables and lass
mice.impute.smcfcs(
  y = y,
  ry = ry,
  x = x[, nzero],
  wy = wy,
  dep_type = dep_type,
  sm = sm,

```

```

    sm_type = sm_type,

    fac_sd_proposal = fac_sd_proposal,

    mh_iter = mh_iter,

    ...
)

}

```

Table S1

Coefficient Estimates of Analysis Model, ELP , Smefcs, No Auxiliary Word Vectors

Term	β	SE_{β}	t	DF	p	RIV	FMI	λ
Intercept	6.868	0.034	201.21	541	< 0.001	8.87	0.90	0.90
Prevalence	-0.082	0.003	-24.05	577	< 0.001	7.14	0.88	0.88
Old	0.004	0.003	1.29	764	0.1985	3.44	0.78	0.78
Pld	0.014	0.002	6.88	840	< 0.001	2.86	0.74	0.74
Concreteness	-0.015	0.005	-3.22	825	0.0013	2.96	0.75	0.75
Age Of Acquisition	0.012	0.002	7.15	1241	< 0.001	1.56	0.61	0.61
Age Of Acquisition Squared	0.000	0.000	-3.13	1126	0.0018	1.78	0.64	0.64
Nummorphs	-0.010	0.002	-4.71	672	< 0.001	4.62	0.82	0.82
Numsylls	0.003	0.002	1.33	489	0.1844	12.73	0.93	0.93
Length	-0.015	0.002	-9.08	1242	< 0.001	1.56	0.61	0.61
Length Squared	0.002	0.000	30.38	3190	< 0.001	0.59	0.37	0.37
Subtitle Frequency	-0.033	0.007	-4.49	750	< 0.001	3.58	0.78	0.78
Subtitle Frequency Squared	-0.003	0.001	-3.76	1040	< 0.001	2.00	0.67	0.67
Valence	-0.203	0.056	-3.64	430	< 0.001	20.62	0.95	0.95
Arousal	0.013	0.039	0.34	472	0.7362	14.56	0.94	0.94
Danger	0.049	0.072	0.68	250	0.4996	87.40	0.99	0.99
Usefulness	-0.123	0.047	-2.62	293	0.0092	62.66	0.98	0.98
Subtitle Frequency Valence Interaction	0.040	0.012	3.40	536	< 0.001	9.12	0.90	0.90
Subtitle Frequency Arousal Interaction	-0.010	0.011	-0.84	520	0.399	10.17	0.91	0.91
Danger Usefulness Interaction	-0.037	0.144	-0.26	314	0.7963	53.25	0.98	0.98

Table S2

Coefficient Estimates of Analysis Model, ELP , Lasso + Smcfcs, Auxiliary Word Vectors

Term	β	SE_{β}	t	DF	p	RIV	FMI	λ
Intercept	6.590	0.022	293.18	740	< 0.001	3.71	0.79	0.79
Prevalence	-0.080	0.002	-32.54	754	< 0.001	3.56	0.78	0.78
Old	0.005	0.002	2.13	1114	0.0336	1.81	0.65	0.64
Pld	0.015	0.002	9.04	1331	< 0.001	1.43	0.59	0.59
Concreteness	-0.021	0.004	-5.21	1234	< 0.001	1.58	0.61	0.61
Age Of Acquisition	0.014	0.002	8.09	1201	< 0.001	1.64	0.62	0.62
Age Of Acquisition Squared	0.000	0.000	-4.88	1060	< 0.001	1.95	0.66	0.66
Nummorphs	-0.008	0.002	-4.44	767	< 0.001	3.43	0.78	0.77
Numsylls	0.002	0.001	1.24	732	0.2145	3.80	0.79	0.79
Length	0.007	0.002	4.50	2224	< 0.001	0.82	0.45	0.45
Length Squared	0.001	0.000	14.58	3223	< 0.001	0.59	0.37	0.37
Subtitle Frequency	0.013	0.006	2.31	1013	0.0212	2.09	0.68	0.68
Subtitle Frequency Squared	-0.004	0.001	-4.86	1042	< 0.001	2.00	0.67	0.67
Valence	-0.076	0.030	-2.55	631	0.011	5.50	0.85	0.85
Arousal	0.223	0.042	5.32	501	< 0.001	11.72	0.92	0.92
Danger	0.034	0.063	0.54	328	0.5881	47.79	0.98	0.98
Usefulness	-0.097	0.037	-2.61	365	0.0094	35.47	0.97	0.97
Subtitle Frequency Valence Interaction	0.007	0.008	0.84	751	0.4035	3.59	0.78	0.78
Subtitle Frequency Arousal Interaction	-0.070	0.012	-5.97	562	< 0.001	7.80	0.89	0.89
Danger Usefulness Interaction	-0.018	0.129	-0.14	378	0.8916	31.92	0.97	0.97

Table S3

Coefficient Estimates of Analysis Model, ELP , Complete Case Analysis

Term	β	SE_{β}	t	DF	p
Intercept	7.069	0.113	62.66	609	< 0.001
Prevalence	-0.108	0.017	-6.16	609	< 0.001
Old	0.006	0.013	0.48	609	0.6332
Pld	-0.004	0.010	-0.40	609	0.6919
Concreteness	-0.010	0.018	-0.56	609	0.5739
Age Of Acquisition	0.001	0.009	0.07	609	0.9448
Age Of Acquisition Squared	0.000	0.001	0.79	609	0.4306
Nummorphs	-0.015	0.014	-1.02	609	0.3078
Numsylls	0.012	0.009	1.32	609	0.186
Length	-0.009	0.011	-0.81	609	0.4155
Length Squared	0.002	0.001	2.18	609	0.0299
Subtitle Frequency	-0.134	0.039	-3.41	609	< 0.001
Subtitle Frequency Squared	0.012	0.004	2.64	609	0.0084
Valence	0.002	0.127	0.02	609	0.9863
Arousal	0.171	0.134	1.27	609	0.2035
Danger	0.018	0.089	0.20	609	0.8429
Usefulness	-0.092	0.046	-2.02	609	0.0435
Subtitle Frequency Valence Interaction	0.006	0.030	0.21	609	0.8326
Subtitle Frequency Arousal Interaction	-0.050	0.033	-1.52	609	0.1288
Danger Usefulness Interaction	0.020	0.178	0.11	609	0.9112

Table S4

Coefficient Estimates of Analysis Model, ELP , Lasso Improper Single Regression Mean

Term	β	SE_{β}	t	DF	p
Intercept	6.509	0.017	386.92	43310	< 0.001
Prevalence	-0.083	0.001	-58.32	43310	< 0.001
Old	-0.004	0.002	-2.62	43310	0.0088
Pld	0.013	0.001	11.25	43310	< 0.001
Concreteness	-0.025	0.003	-8.31	43310	< 0.001
Age Of Acquisition	0.013	0.001	8.70	43310	< 0.001
Age Of Acquisition Squared	0.000	0.000	-5.96	43310	< 0.001
Nummorphs	-0.021	0.001	-21.79	43310	< 0.001
Numsylls	0.005	0.001	6.35	43310	< 0.001
Length	0.010	0.001	8.61	43310	< 0.001
Length Squared	0.001	0.000	18.95	43310	< 0.001
Subtitle Frequency	0.056	0.004	12.56	43310	< 0.001
Subtitle Frequency Squared	-0.006	0.000	-11.82	43310	< 0.001
Valence	-0.046	0.017	-2.73	43310	0.0063
Arousal	0.444	0.026	17.00	43310	< 0.001
Danger	-0.085	0.048	-1.77	43310	0.0767
Usefulness	-0.064	0.025	-2.57	43310	0.01
Subtitle Frequency Valence Interaction	-0.011	0.005	-2.25	43310	0.0244
Subtitle Frequency Arousal Interaction	-0.125	0.007	-16.70	43310	< 0.001
Danger Usefulness Interaction	0.081	0.116	0.70	43310	0.4825

Table S5

Coefficient Estimates of Analysis Model, BLP , Smcfs, No Auxiliary Word Vectors

Term	β	SE_{β}	t	DF	p	RIV	FMI	λ
Intercept	6.880	0.033	205.80	509	< 0.001	9.25	0.90	0.90
Prevalence	-0.060	0.004	-15.18	436	< 0.001	15.43	0.94	0.94
Old	0.011	0.004	2.88	646	0.0042	4.71	0.83	0.82
Pld	0.011	0.003	3.53	663	< 0.001	4.40	0.82	0.81
Concreteness	-0.009	0.005	-1.69	662	0.0906	4.42	0.82	0.82
Age Of Acquisition	0.021	0.002	10.03	641	< 0.001	4.79	0.83	0.83
Age Of Acquisition Squared	-0.001	0.000	-8.10	561	< 0.001	6.88	0.87	0.87
Nummorphs	-0.028	0.003	-9.05	521	< 0.001	8.60	0.90	0.90
Numsylls	0.002	0.003	0.70	542	0.487	7.62	0.88	0.88
Length	0.010	0.003	3.94	1705	< 0.001	1.04	0.51	0.51
Length Squared	0.000	0.000	-1.42	2167	0.1558	0.81	0.45	0.45
Subtitle Frequency	-0.092	0.010	-9.36	478	< 0.001	11.37	0.92	0.92
Subtitle Frequency Squared	-0.004	0.001	-5.26	668	< 0.001	4.33	0.81	0.81
Valence	-0.643	0.046	-14.03	394	< 0.001	21.27	0.96	0.96
Arousal	-0.189	0.048	-3.93	351	< 0.001	29.60	0.97	0.97
Danger	-0.197	0.081	-2.44	182	0.0158	110.32	0.99	0.99
Usefulness	-0.197	0.052	-3.82	202	< 0.001	93.21	0.99	0.99
Subtitle Frequency Valence Interaction	0.154	0.008	18.45	558	< 0.001	6.99	0.88	0.87
Subtitle Frequency Arousal Interaction	0.036	0.013	2.84	409	0.0047	18.85	0.95	0.95
Danger Usefulness Interaction	0.485	0.160	3.04	249	0.0026	64.44	0.98	0.98

Table S6

Coefficient Estimates of Analysis Model, BLP , Lasso + Smcfcs, Auxiliary Word Vectors

Term	β	SE_{β}	t	DF	p	RIV	FMI	λ
Intercept	6.405	0.038	167.65	511	< 0.001	9.20	0.90	0.90
Prevalence	-0.043	0.003	-13.28	543	< 0.001	7.63	0.88	0.88
Old	0.025	0.004	6.48	795	< 0.001	3.01	0.75	0.75
Pld	0.009	0.003	2.92	761	0.0037	3.28	0.77	0.77
Concreteness	-0.003	0.005	-0.51	800	0.6086	2.97	0.75	0.75
Age Of Acquisition	0.032	0.002	16.55	798	< 0.001	2.99	0.75	0.75
Age Of Acquisition Squared	-0.001	0.000	-14.07	646	< 0.001	4.73	0.83	0.83
Nummorphs	-0.006	0.003	-1.94	573	0.0526	6.48	0.87	0.87
Numsylls	-0.005	0.002	-2.14	778	0.0331	3.14	0.76	0.76
Length	0.035	0.003	11.28	1719	< 0.001	1.03	0.51	0.51
Length Squared	-0.002	0.000	-10.10	2046	< 0.001	0.86	0.46	0.46
Subtitle Frequency	-0.031	0.009	-3.34	553	< 0.001	7.19	0.88	0.88
Subtitle Frequency Squared	-0.003	0.001	-3.55	779	< 0.001	3.13	0.76	0.76
Valence	-0.273	0.038	-7.18	494	< 0.001	10.26	0.91	0.91
Arousal	0.062	0.067	0.92	347	0.357	30.50	0.97	0.97
Danger	-0.130	0.076	-1.71	256	0.0893	61.02	0.98	0.98
Usefulness	-0.183	0.046	-4.01	270	< 0.001	54.85	0.98	0.98
Subtitle Frequency Valence Interaction	0.068	0.010	7.06	586	< 0.001	6.09	0.86	0.86
Subtitle Frequency Arousal Interaction	-0.032	0.018	-1.79	412	0.0744	18.53	0.95	0.95
Danger Usefulness Interaction	0.393	0.162	2.42	305	0.0161	42.12	0.98	0.98

Table S7

Coefficient Estimates of Analysis Model, BLP , Complete Case Analysis

Term	β	SE_{β}	t	DF	p
Intercept	6.927	0.092	75.03	519	< 0.001
Prevalence	-0.119	0.014	-8.80	519	< 0.001
Old	-0.004	0.011	-0.32	519	0.7517
Pld	0.004	0.009	0.50	519	0.6147
Concreteness	-0.017	0.014	-1.25	519	0.2137
Age Of Acquisition	0.005	0.007	0.82	519	0.4137
Age Of Acquisition Squared	0.000	0.000	0.27	519	0.7892
Nummorphs	-0.024	0.013	-1.80	519	0.0723
Numsylls	-0.007	0.008	-0.94	519	0.35
Length	-0.001	0.012	-0.07	519	0.9468
Length Squared	0.001	0.001	0.65	519	0.5143
Subtitle Frequency	-0.115	0.030	-3.79	519	< 0.001
Subtitle Frequency Squared	0.011	0.003	3.27	519	0.0012
Valence	0.113	0.098	1.16	519	0.2468
Arousal	0.035	0.107	0.33	519	0.7443
Danger	-0.096	0.075	-1.28	519	0.2022
Usefulness	-0.121	0.037	-3.29	519	0.0011
Subtitle Frequency Valence Interaction	-0.018	0.023	-0.77	519	0.4433
Subtitle Frequency Arousal Interaction	-0.022	0.026	-0.84	519	0.4014
Danger Usefulness Interaction	0.291	0.151	1.93	519	0.0544

Table S8

Coefficient Estimates of Analysis Model, BLP , Lasso Improper Single Regression Mean

Term	β	SE_{β}	t	DF	p
Intercept	6.296	0.019	327.28	31607	< 0.001
Prevalence	-0.039	0.001	-26.86	31607	< 0.001
Old	-0.008	0.002	-3.25	31607	0.0011
Pld	0.003	0.002	1.30	31607	0.1921
Concreteness	-0.011	0.003	-3.24	31607	0.0012
Age Of Acquisition	0.030	0.002	19.00	31607	< 0.001
Age Of Acquisition Squared	-0.001	0.000	-17.55	31607	< 0.001
Nummorphs	-0.023	0.001	-19.73	31607	< 0.001
Numsylls	-0.002	0.001	-1.14	31607	0.2545
Length	0.043	0.002	17.97	31607	< 0.001
Length Squared	-0.002	0.000	-9.76	31607	< 0.001
Subtitle Frequency	0.032	0.005	6.67	31607	< 0.001
Subtitle Frequency Squared	-0.006	0.001	-12.68	31607	< 0.001
Valence	-0.142	0.018	-7.86	31607	< 0.001
Arousal	0.274	0.030	9.23	31607	< 0.001
Danger	-0.225	0.053	-4.23	31607	< 0.001
Usefulness	-0.205	0.026	-7.84	31607	< 0.001
Subtitle Frequency Valence Interaction	0.027	0.005	5.24	31607	< 0.001
Subtitle Frequency Arousal Interaction	-0.086	0.008	-10.42	31607	< 0.001
Danger Usefulness Interaction	0.463	0.129	3.59	31607	< 0.001

Table S9

Coefficient Estimates of Analysis Model, ECP , Smcfc, No Auxiliary Word Vectors

Term	β	SE_{β}	t	DF	p	RIV	FMI	λ
Intercept	7.411	0.044	170.25	404	< 0.001	34.08	0.97	0.97
Prevalence	-0.172	0.006	-28.14	289	< 0.001	89.67	0.99	0.99
Old	-0.025	0.006	-4.04	390	< 0.001	38.59	0.97	0.97
Pld	0.017	0.004	4.00	399	< 0.001	35.85	0.97	0.97
Concreteness	-0.029	0.006	-5.07	542	< 0.001	10.10	0.91	0.91
Age Of Acquisition	0.005	0.003	1.57	460	0.1175	20.22	0.95	0.95
Age Of Acquisition Squared	0.000	0.000	-1.28	415	0.2021	30.94	0.97	0.97
Nummorphs	0.005	0.003	1.61	437	0.108	25.11	0.96	0.96
Numsylls	-0.001	0.004	-0.15	295	0.8812	85.51	0.99	0.99
Length	-0.004	0.002	-1.85	565	0.0647	8.58	0.90	0.90
Length Squared	0.002	0.000	14.81	542	< 0.001	10.10	0.91	0.91
Subtitle Frequency	-0.060	0.010	-5.76	470	< 0.001	18.49	0.95	0.95
Subtitle Frequency Squared	0.006	0.001	7.05	663	< 0.001	5.13	0.84	0.84
Valence	-0.147	0.075	-1.96	248	0.0506	123.31	0.99	0.99
Arousal	0.039	0.061	0.63	263	0.5279	110.32	0.99	0.99
Danger	-0.169	0.070	-2.41	120	0.0175	381.15	1.00	1.00
Usefulness	-0.260	0.050	-5.23	182	< 0.001	212.19	1.00	1.00
Subtitle Frequency Valence Interaction	0.031	0.016	1.98	383	0.0483	41.35	0.98	0.98
Subtitle Frequency Arousal Interaction	-0.020	0.019	-1.07	354	0.2846	53.15	0.98	0.98
Danger Usefulness Interaction	0.469	0.084	5.56	299	< 0.001	82.84	0.99	0.99

Table S10

Coefficient Estimates of Analysis Model, ECP , Lasso + Smcfcs, Auxiliary Word Vectors

Term	β	SE_{β}	t	DF	p	RIV	FMI	λ
Intercept	7.395	0.027	269.95	530	< 0.001	11.11	0.92	0.92
Prevalence	-0.182	0.004	-45.21	406	< 0.001	33.73	0.97	0.97
Old	-0.030	0.003	-8.72	546	< 0.001	9.89	0.91	0.91
Pld	0.027	0.002	10.66	551	< 0.001	9.56	0.91	0.91
Concreteness	-0.034	0.005	-7.15	663	< 0.001	5.16	0.84	0.84
Age Of Acquisition	-0.002	0.002	-1.22	703	0.2212	4.42	0.82	0.82
Age Of Acquisition Squared	0.000	0.000	2.11	622	0.0349	6.24	0.86	0.86
Nummorphs	0.006	0.002	2.91	502	0.0038	14.02	0.93	0.93
Numsylls	-0.001	0.002	-0.53	400	0.597	35.76	0.97	0.97
Length	-0.002	0.002	-1.12	757	0.2636	3.71	0.79	0.79
Length Squared	0.002	0.000	13.19	560	< 0.001	8.98	0.90	0.90
Subtitle Frequency	-0.057	0.007	-8.66	644	< 0.001	5.60	0.85	0.85
Subtitle Frequency Squared	0.007	0.001	10.74	930	< 0.001	2.48	0.71	0.71
Valence	-0.072	0.038	-1.88	428	0.0613	27.52	0.97	0.96
Arousal	0.100	0.051	1.95	362	0.0525	49.88	0.98	0.98
Danger	-0.159	0.077	-2.06	196	0.0403	187.68	0.99	0.99
Usefulness	-0.207	0.045	-4.56	259	< 0.001	114.07	0.99	0.99
Subtitle Frequency Valence Interaction	0.015	0.011	1.42	497	0.1569	14.57	0.94	0.94
Subtitle Frequency Arousal Interaction	-0.038	0.015	-2.60	447	0.0096	22.94	0.96	0.96
Danger Usefulness Interaction	0.459	0.137	3.34	308	< 0.001	77.58	0.99	0.99

Table S11

Coefficient Estimates of Analysis Model, ECP , Complete Case Analysis

Term	β	SE_{β}	t	DF	p
Intercept	7.295	0.055	133.64	609	< 0.001
Prevalence	-0.149	0.008	-17.59	609	< 0.001
Old	-0.004	0.006	-0.61	609	0.5417
Pld	-0.013	0.005	-2.69	609	0.0073
Concreteness	0.011	0.009	1.29	609	0.1986
Age Of Acquisition	0.006	0.004	1.49	609	0.1357
Age Of Acquisition Squared	0.000	0.000	-0.48	609	0.6305
Nummorphs	0.005	0.007	0.78	609	0.435
Numsylls	0.009	0.004	2.19	609	0.0286
Length	-0.004	0.006	-0.77	609	0.4392
Length Squared	0.001	0.000	2.84	609	0.0047
Subtitle Frequency	-0.088	0.019	-4.64	609	< 0.001
Subtitle Frequency Squared	0.008	0.002	3.68	609	< 0.001
Valence	0.006	0.062	0.10	609	0.9225
Arousal	-0.030	0.065	-0.45	609	0.6496
Danger	-0.075	0.043	-1.74	609	0.0821
Usefulness	-0.051	0.022	-2.29	609	0.0222
Subtitle Frequency Valence Interaction	-0.003	0.015	-0.20	609	0.839
Subtitle Frequency Arousal Interaction	0.007	0.016	0.42	609	0.6754
Danger Usefulness Interaction	0.151	0.086	1.76	609	0.0792

Table S12

Coefficient Estimates of Analysis Model, ECP , Lasso Improper Single Regression Mean

Term	β	SE_{β}	t	DF	p
Intercept	7.313	0.012	588.52	60608	< 0.001
Prevalence	-0.188	0.001	-286.23	60608	< 0.001
Old	-0.036	0.001	-24.96	60608	< 0.001
Pld	0.019	0.001	16.79	60608	< 0.001
Concreteness	-0.027	0.002	-12.47	60608	< 0.001
Age Of Acquisition	0.002	0.001	2.23	60608	0.0258
Age Of Acquisition Squared	0.000	0.000	0.21	60608	0.8365
Nummorphs	-0.009	0.001	-10.69	60608	< 0.001
Numsylls	-0.001	0.001	-1.38	60608	0.1665
Length	0.001	0.001	1.63	60608	0.1029
Length Squared	0.002	0.000	40.16	60608	< 0.001
Subtitle Frequency	-0.062	0.003	-18.81	60608	< 0.001
Subtitle Frequency Squared	0.009	0.000	23.92	60608	< 0.001
Valence	-0.045	0.011	-4.11	60608	< 0.001
Arousal	0.174	0.018	9.89	60608	< 0.001
Danger	0.079	0.039	2.02	60608	0.0439
Usefulness	0.004	0.019	0.23	60608	0.8177
Subtitle Frequency Valence Interaction	0.004	0.003	1.09	60608	0.2766
Subtitle Frequency Arousal Interaction	-0.060	0.005	-11.34	60608	< 0.001
Danger Usefulness Interaction	-0.113	0.096	-1.17	60608	0.2408

Table S13

Coefficient Estimates of Analysis Model, MALD , Smcfcs, No Auxiliary Word Vectors

Term	β	SE_{β}	t	DF	p	RIV	FMI	λ
Intercept	6.846	0.051	133.12	666	< 0.001	4.19	0.81	0.81
Prevalence	-0.082	0.006	-14.29	657	< 0.001	4.33	0.81	0.81
Old	-0.004	0.003	-1.06	1440	0.2915	1.23	0.55	0.55
Pld	-0.009	0.003	-2.89	1337	0.0039	1.35	0.57	0.57
Concreteness	-0.008	0.008	-1.09	986	0.2766	2.03	0.67	0.67
Age Of Acquisition	0.005	0.003	1.67	1446	0.0945	1.23	0.55	0.55
Age Of Acquisition Squared	0.000	0.000	0.13	1330	0.8945	1.35	0.58	0.58
Nummorphs	-0.008	0.003	-2.74	775	0.0063	3.05	0.75	0.75
Numsylls	-0.005	0.003	-1.53	748	0.1266	3.28	0.77	0.77
Subtitle Frequency	-0.004	0.012	-0.34	932	0.7306	2.22	0.69	0.69
Subtitle Frequency Squared	-0.001	0.001	1.33	1606	0.1824	1.09	0.52	0.52
Valence	-0.003	0.075	-0.04	481	0.9648	10.17	0.91	0.91
Arousal	0.060	0.055	1.09	542	0.2749	7.11	0.88	0.88
Danger	-0.126	0.090	-1.41	293	0.1607	39.05	0.98	0.98
Usefulness	-0.184	0.067	-2.74	303	0.0065	36.29	0.97	0.97
Duration	0.000	0.000	7.25	12925	< 0.001	0.15	0.13	0.13
Duration Squared	0.000	0.000	3.69	12632	< 0.001	0.15	0.13	0.13
Subtitle Frequency Valence Interaction	-0.004	0.016	-0.21	600	0.8308	5.39	0.84	0.84
Subtitle Frequency Arousal Interaction	-0.022	0.016	-1.45	596	0.1474	5.49	0.85	0.85
Danger Usefulness Interaction	0.295	0.176	1.67	350	0.095	25.66	0.96	0.96

Table S14

Coefficient Estimates of Analysis Model, MALD , Lasso + Smcfcs, Auxiliary Word Vectors

Term	β	SE_{β}	t	DF	p	RIV	FMI	λ
Intercept	6.848	0.034	201.88	1366	< 0.001	1.31	0.57	0.57
Prevalence	-0.087	0.004	-20.37	1063	< 0.001	1.82	0.65	0.64
Old	-0.004	0.003	-1.23	2724	0.2202	0.64	0.39	0.39
Pld	-0.010	0.003	-3.64	2730	< 0.001	0.64	0.39	0.39
Concreteness	-0.011	0.006	-1.82	2206	0.0685	0.78	0.44	0.44
Age Of Acquisition	0.005	0.003	2.00	2506	0.0458	0.69	0.41	0.41
Age Of Acquisition Squared	0.000	0.000	-0.32	2248	0.7455	0.77	0.43	0.43
Nummorphs	-0.009	0.003	-3.61	1014	< 0.001	1.94	0.66	0.66
Numsylls	-0.004	0.002	-1.63	1317	0.1025	1.37	0.58	0.58
Subtitle Frequency	-0.012	0.009	-1.32	2180	0.1869	0.79	0.44	0.44
Subtitle Frequency Squared	0.003	0.001	2.88	4542	0.004	0.41	0.29	0.29
Valence	-0.003	0.039	-0.09	1001	0.9311	1.98	0.67	0.66
Arousal	0.091	0.050	1.82	727	0.0693	3.46	0.78	0.78
Danger	-0.075	0.063	-1.19	442	0.2359	13.09	0.93	0.93
Usefulness	-0.141	0.045	-3.15	430	0.0018	14.28	0.93	0.93
Duration	0.000	0.000	7.56	16655	< 0.001	0.11	0.10	0.10
Duration Squared	0.000	0.000	3.56	16979	< 0.001	0.10	0.09	0.09
Subtitle Frequency Valence Interaction	-0.002	0.010	-0.16	1333	0.8759	1.35	0.57	0.57
Subtitle Frequency Arousal Interaction	-0.033	0.014	-2.39	841	0.0173	2.62	0.72	0.72
Danger Usefulness Interaction	0.184	0.129	1.43	492	0.1545	9.41	0.90	0.90

Table S15

Coefficient Estimates of Analysis Model, MALD , Complete Case Analysis

Term	β	SE_{β}	t	DF	p
Intercept	6.698	0.212	31.58	609	< 0.001
Prevalence	-0.020	0.030	-0.66	609	0.5105
Old	-0.014	0.018	-0.77	609	0.4439
Pld	-0.004	0.017	-0.26	609	0.7957
Concreteness	-0.033	0.031	-1.05	609	0.2944
Age Of Acquisition	0.007	0.015	0.48	609	0.6324
Age Of Acquisition Squared	0.000	0.001	-0.11	609	0.9096
Nummorphs	0.018	0.024	0.75	609	0.4552
Numsylls	-0.001	0.014	-0.09	609	0.9256
Subtitle Frequency	-0.075	0.067	-1.13	609	0.2601
Subtitle Frequency Squared	0.010	0.008	1.27	609	0.2059
Valence	0.085	0.218	0.39	609	0.6982
Arousal	0.078	0.231	0.34	609	0.7361
Danger	-0.109	0.151	-0.72	609	0.4726
Usefulness	-0.146	0.078	-1.87	609	0.062
Duration	0.001	0.000	1.78	609	0.0761
Duration Squared	0.000	0.000	-0.60	609	0.5482
Subtitle Frequency Valence Interaction	-0.021	0.052	-0.40	609	0.6869
Subtitle Frequency Arousal Interaction	-0.023	0.057	-0.41	609	0.6851
Danger Usefulness Interaction	0.226	0.304	0.74	609	0.4574

Table S16

Coefficient Estimates of Analysis Model, MALD , Lasso Improper Single Regression Mean

Term	β	SE_{β}	t	DF	p
Intercept	6.932	0.032	213.68	26648	< 0.001
Prevalence	-0.088	0.003	-29.80	26648	< 0.001
Old	-0.002	0.002	-0.99	26648	0.3234
Pld	-0.010	0.002	-4.57	26648	< 0.001
Concreteness	-0.014	0.005	-2.79	26648	0.0053
Age Of Acquisition	0.003	0.003	1.21	26648	0.2271
Age Of Acquisition Squared	0.000	0.000	0.48	26648	0.6284
Nummorphs	-0.009	0.001	-6.18	26648	< 0.001
Numsylls	-0.006	0.002	-3.75	26648	< 0.001
Subtitle Frequency	-0.025	0.009	-2.95	26648	0.0032
Subtitle Frequency Squared	0.004	0.001	4.50	26648	< 0.001
Valence	0.000	0.030	0.01	26648	0.9888
Arousal	0.048	0.047	1.03	26648	0.303
Danger	-0.009	0.079	-0.12	26648	0.9054
Usefulness	-0.179	0.040	-4.45	26648	< 0.001
Duration	0.000	0.000	7.56	26648	< 0.001
Duration Squared	0.000	0.000	4.14	26648	< 0.001
Subtitle Frequency Valence Interaction	-0.005	0.009	-0.55	26648	0.5798
Subtitle Frequency Arousal Interaction	-0.024	0.013	-1.87	26648	0.0621
Danger Usefulness Interaction	0.033	0.183	0.18	26648	0.8575

Appendix C

Convergence of Mice

ELP Mean and SD by iteration, Word Vectors

ELP Mean and SD by iteration, Word Vectors

BLP Mean and SD by iteration, Word Vectors

BLP Mean and SD by iteration, Word Vectors

MALD Mean and SD by iteration, Word Vectors

ELP Mean and SD by iteration, No Word Vectors

ELP Mean and SD by iteration, No Word Vectors

BLP Mean and SD by iteration, No Word Vectors

BLP Mean and SD by iteration, No Word Vectors

ECP Mean and SD by iteration, No Word Vectors

ECP Mean and SD by iteration, No Word Vectors

MALD Mean and SD by iteration, No Word Vectors

MALD Mean and SD by iteration, No Word Vectors