

Concept Class Analysis: A Method for Identifying Cultural Schemas in Texts

Marshall A. Taylor,^a Dustin S. Stoltz^b

a) New Mexico State University; b) Lehigh University

Abstract: Recent methodological work at the intersection of culture, cognition, and computational methods has drawn attention to how cultural schemas can be “recovered” from social survey data. Defining cultural schemas as slowly learned, implicit, and unevenly distributed relational memory structures, researchers show how schemas—or rather, the downstream consequences of people drawing upon them—can be operationalized and measured from domain-specific survey modules. Respondents can then be sorted into “classes” on the basis of the schema to which their survey response patterns best align. In this article, we extend this “schematic class analysis” method to text data. We introduce concept class analysis (CoCA): a hybrid model that combines word embeddings and correlational class analysis to group documents across a corpus by the similarity of schemas recovered from them. We introduce the CoCA model, illustrate its validity and utility using simulations, and conclude with considerations for future research and applications.

Keywords: concept class analysis; concept mover’s distance; correlational class analysis; schemas; culture; cognition

SOCIOLOGISTS often use the concept of “schema” to articulate the role of personal culture in social action (DiMaggio 1997; Blair-Loy 2001; Swidler 2001; Vaisey 2009; Cerulo 2010; Spillman 2016). Measuring cultural schemas—which are defined as implicit and unevenly distributed relational memory structures that are slowly learned through prior life experiences—has nonetheless proved tricky for traditional methods, like surveys, not unlike many other kinds of culture (Mohr et al. 2020:3). For this reason, Goldberg’s (2011) article was particularly influential: it presented a formal method for not only operationalizing and measuring schemas but doing so in an unlikely place—social surveys.

Goldberg’s method (2011), relational class analysis (RCA), takes respondents’ response vectors across an array of domain-specific ordinal survey items as input and—using a schematic similarity metric and graph-partitioning technique that we briefly summarize later—sorts respondents into classes on the basis of the schema to which their survey response patterns best align. Boutyline’s follow-up approach—correlational class analysis (CCA)—further exemplifies the utility of these “schematic class analysis” methods for quantitative cultural analysis (Boutyline 2017:385). Given the preponderance of interest in measuring cultural schemas (and finding what they do and what predicts them) and the plethora of types of data for cultural analysis, a logical next step is to develop methods for extracting schemas from sources other than surveys.

We propose such a method—namely, a method for identifying schemas in texts. We introduce concept class analysis (CoCA): a hybrid model that combines word

Citation: Taylor, Marshall A., and Dustin S. Stoltz. 2020. “Concept Class Analysis: A Method for Identifying Cultural Schemas in Texts.” *Sociological Science* 7: 544–569.


Received: July 31, 2020

Accepted: October 4, 2020

Published: November 9, 2020

Editor(s): Jesper Sørensen, Gabriel Rossman

DOI: 10.15195/v7.a23

Copyright: © 2020 The Author(s). This open-access article has been published under a Creative Commons Attribution License, which allows unrestricted use, distribution and reproduction, in any form, as long as the original author and source have been credited. 

embeddings—specifically, concept mover’s distance (Stoltz and Taylor 2019) and semantic directions (Bolukbasi et al. 2016; Kozlowski et al. 2019; Taylor and Stoltz 2020)—and CCA to find groups of documents across a corpus that share similar schematic traces.

In contrast to inductive methods for clustering documents, such as topic modeling, our proposed workflow incorporates a deductive approach in which the analyst can define the dimensions of interest; and, unlike traditional supervised classifiers, our workflow requires no hand-coding of documents. More importantly, unlike surveys, the set of meaningful dimensions that can be estimated from a corpus are not limited by the questions of a survey instrument. If theory and empirical research suggest other semantic directions might be salient for a particular question or corpus, all that is needed is a list of juxtaposed terms that approximate this dimension. This opens up a vast landscape of applications for CoCA, given the ever-growing storehouse of relatively inexpensive and easily accessed natural language corpora. Any corpus can be used with our proposed workflow.

In what follows, we first overview the theoretical foundations of the schematic class analysis framework and further extend these foundations to show how schemas can also be approximated using their “traces” in natural language. We then detail the CoCA model and report the results from a series of simulations that illustrate CoCA’s validity and utility as a schema-identifying tool for document corpora. We then conclude with a discussion for future research and applications using the method.

Culture, Cognition, and Text Analysis

Defining Cultural Schemas

Simply put, a schema is a “flexible memory structure, automatically acquired and updated from patterned activity, [and] composed of multimodal neural associations” (Wood et al. 2018:246). In other words, a schema is a nondeclarative cognitive network that people acquire slowly over the life course to mentally represent situations and their prototypical features as well as typical sensorimotor procedures for interacting with objects and situations (Squire 1992; Ghosh and Gilboa 2014; Lizardo 2017).¹ Schemas are central to human learning, thinking, and action, as they are activated (literally, as a weighted network of co-activating neurons) to recognize, fill, and update patterns in one’s everyday observations (Rumelhart 1980:39)—a sort of mental shortcut that is especially crucial when faced with “conditions of incomplete information” (DiMaggio 1997:269). Schemas, then, are central neurocognitive structures behind the “cognitive-miser” model of human thinking (Fiske and Taylor 2013).²

For instance, one may see a 23-wheeled, translucent, octagon-shaped object with a person inside of it moving down a freeway—a previously unobserved object for this person, to be sure. But the fact the object is on a *road*, has *wheels*, is *moving*, and has a *person inside* of it will likely satisfy enough of the prototypical features of their internalized “automobile” schema for it to be activated as the relevant sense-making mechanism; the object will almost instantaneously be understood as a road vehicle

of some kind. The schema comprises the co-activation of the neural representations of these prototypical features. Furthermore, the internalization of this “automobile” schema is due to the person’s repeated and multisensory exposure to cars, trucks, roads, and the like—*seeing* cars, *driving* cars, *hearing* wheels squeal, *smelling* gasoline, etc., all correlationally experienced.

A schema is then said to be *cultural* to the extent that it is a learned cognitive structure (Wood et al. 2018:246). Furthermore, for sociologists, the primary focus of investigation has been the extent to which people *share* similar schemas (Boutyline and Soter 2020:6).³ By exploring the distributions of patterns of schematic traces in social life, or what Sperber (1996:1–2) refers to as an “epidemiology of representation,” the analyst “‘dissolv[es]’ the individual, unlock[ing] the potential to theorize the ‘collective-social’ at a level below the person” (Lizardo et al. 2020:9).

The automobile schema in the previous illustration is an example of a cultural schema, as it is certainly learned through one’s experience of the social world. Importantly, though, although it may be widely shared (in terms of relative similarity between people), it is not universal, and it is not necessarily the most salient sense-making mechanism. If a person grew up and lived in a militarized zone, for instance, then the features of this object may instead prime a “tank” schema (or perhaps a more general “military” schema).

For our purposes, there are two important points to be gleaned from the discussion thus far. First, schemas are inherently *relational* in the sense that they are “implicit recognition procedures that emerge from intricate *associational* links” between elements represented in memory (Goldberg 2011:1401, emphasis added; Boutyline and Vaisey 2017; Hunzaker and Valentino 2019). Second, cultural schemas will be *shared, to the extent that people have similar life histories and therefore occupy similar social positions and life-course trajectories*. Cultural schemas—by their nature of being learned and unevenly shared across populations—can be conceptualized as a means of distinguishing between “thought communities,” or social groups with more or less distinct meaning-making and meaning-maintenance practices (Fleck 1979; Zerubavel 2009; Lee and Martin 2018). Methodologically, then, we can “reverse engineer” and identify cultural schemas by identifying the social groups that demonstrate similar patterns of activities, attitudes, or products.

Schematic Class Analysis with Survey Data

Cultural schemas have been measured in a number of ways—e.g., through experiments in cultural and social psychology and semistructured interviews and participant observation in cognitive anthropology (Strauss and Quinn 1997; Shore 1998; Morris and Mok 2011; Hunzaker 2014; Homan et al. 2017; Hunzaker and Valentino 2019; Miles 2019). In a now classic 2011 article, Goldberg put forth RCA as an innovative method for measuring cultural schemas in social survey data. Specifically, RCA quantifies the extent to which participants can be grouped by formal similarities in their responses and then uses group-specific inter-item correlations to interpret the schemas that define each group.

Goldberg (2011:1402–1405) gives the following example to show the intuition behind RCA. Consider Figure 1, which reports musical tastes for four (A, B, C, and

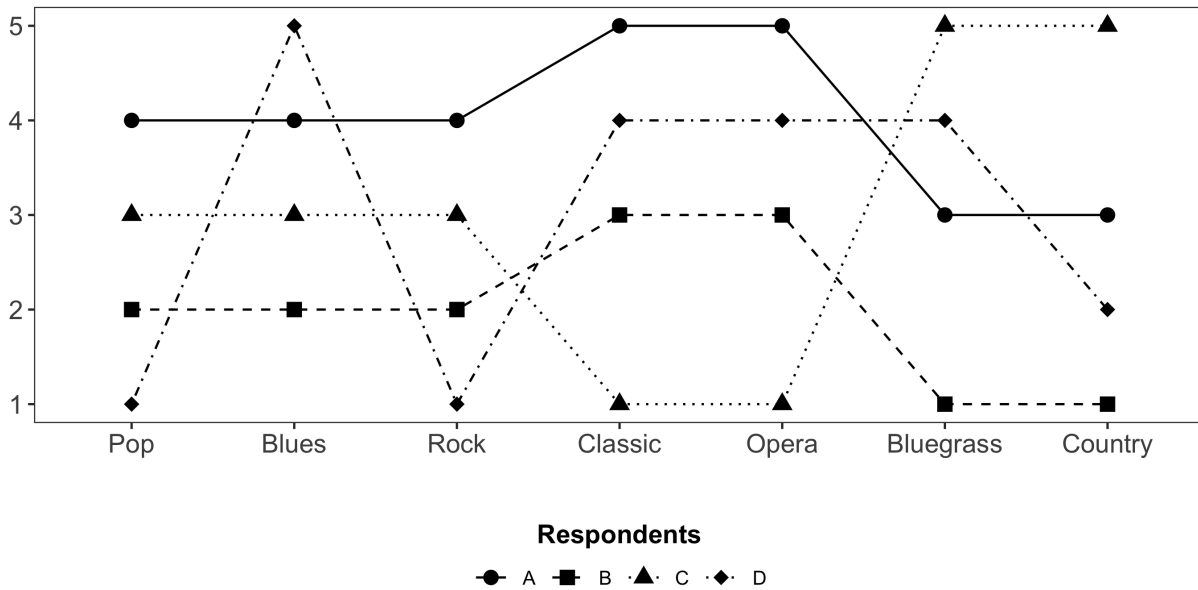


Figure 1: Hypothetical survey responses for four respondents. *Note:* Adapted from Goldberg (2011:1405) and Boutyline (2017:357, 361). The figure shows survey responses are musical tastes from four hypothetical respondents, ranging from “strongly dislike” (1) to “strongly like” (5).

D) hypothetical survey respondents. Each respondent rates the extent to which they “strongly dislike” (1) or “strongly like” (5) seven different music genres, listed along the x axis of panel A. Respondent A likes pop, blues, and rock, strongly likes classical and opera, and is fairly “in the middle” when it comes to bluegrass and country. Respondent B does not respond with any of the same ordinal categories across the genres: they dislike pop, blues, and rock, are neutral about classical and opera, and strongly dislike bluegrass and country.

We may default to using Euclidean distances as a measure of the extent to which these two respondents are similar in their musical tastes, but, as Goldberg notes (2011:1406), Euclidean distances only account for respondent differences specific to each variable (e.g., $A_{Pop} - B_{Pop}$) and rely on summed squares (i.e., $\sqrt{\sum_{g=1}^G (A_g - B_g)^2}$, where g is a genre in the total set of genres G). This means that Euclidean distances account for neither (a) interrespondent similarities in within-respondent response patterns across the set of survey items nor (b) the directionality of the similarities. Consequently, respondents A and B exhibit a (standardized) Euclidean distance of about 0.5.

If, however, we are interested in the extent to which respondents A and B are *schematically* similar, then we should instead identify similarities in *patterns of relative adjudication* (McDonnell et al. 2020)—that is, the extent to which respondents show similarity in the oppositions they implicitly assign between survey items. Though the respondents do not give the same answers for any of the taste items, they nonetheless “weight” their relative like or dislike for each genre in the same direction: respondents A and B both equally (dis)like pop, blues, and rock, like

classical and opera equal to one another (and more so than they like pop, blues, and rock), and they both dislike bluegrass and country more than they dislike any of the other genres. These respondents, then, are *schematically identical*: they use the same musical taste schema to evaluate musical genres.

Goldberg introduces a metric called “relationality” to quantify the extent to which two respondents are schematically similar. For respondents A and B in Figure 1, their relationality is equal to 1 (on a -1 to 1 scale)—indicating they are schematically identical (notice that A’s and B’s “trend line” across the genres is exactly the same; B is simply shifted down by two scale categories for each genre).

Importantly, two people can acquire similar schemas and yet express diametrically opposed valuations of individual objects within that same schematic understanding. For example, respondent C does not share any of the same answers across the genres as respondents A and B. Furthermore, their valuations go in opposite directions: they are neutral about pop, blues, and rock, strongly dislike classical and opera, and strongly like bluegrass and country. Although C’s valuations go in opposite directions as A and B, they nonetheless exhibit the *same pattern of evaluation*: e.g., classical, opera, bluegrass, and country are highly distinguishing genres in their meaning-making. Appropriately, C’s relationality with A and B is negative. Finally, respondent D shows no discernible similarities in how they organize their tastes; as such, they have a relationality close to zero for all three pairwise comparisons.

RCA then applies a graph-partitioning algorithm to the symmetric matrix of absolute pairwise interrespondent relationalities to identify the schematic classes—i.e., those groups of respondents who approximate the same cultural schema for organizing their beliefs, attitudes, or tastes within the domain of interest.

An important variant to RCA is CCA (Boutyline 2017). In short, CCA differs from RCA only in the measure of schematic similarity; the theoretical foundations, default graph-partitioning algorithm, and interpretation techniques are otherwise the same. Boutyline observed that two respondents—say, i and j —draw upon exactly the same schema for finding meaning in some domain if their associated response vectors across the set of K domain-specific survey items are perfect *linear transformations* of one another (Boutyline 2017:356). The absolute Pearson correlation is—like absolute relationality—a continuous measure in the $[0,1]$ interval of the extent to which i and j follow the same schema. Boutyline reported a series of simulations and found that the interrespondent absolute Pearson correlation was consistently more accurate than the absolute relationality as a schematic similarity metric and also more computationally efficient. For these reasons carefully laid out by Boutyline, we rely on CCA as our model for schematic class analysis.

Together, RCA and CCA have enjoyed wide empirical application. They have been used to study taste (Daenekindt 2019; Willekens and Daenekindt 2020), public opinion and politics (Baldassarri and Goldberg 2014; Wu 2014; Boutyline and Vaisey 2017; Daenekindt et al. 2017; Barbet 2020; Hu and Yin 2020), science and religion (DiMaggio et al. 2018), moral schemas about markets (DiMaggio and Goldberg 2018; McDonnell et al. 2020), organizational values (Rossoni et al. 2020), and literary schemas (Rawlings and Childress 2019). What all of these applications have in common, of course, is the reliance on attitudinal measures for identifying schemas.

How might RCA/CCA be applicable for researchers with nonsurvey data—say, texts?⁴

Can Cultural Schemas Be Identified in Natural Language?

Before outlining our method, we must first (briefly) lay out some theoretical and empirical support for the notion that schemas can be recovered from written natural language. Although the production of any piece of text is shaped by a collective process that is always embedded in social organization, and thus shaped by the exigencies of specific fields (e.g., Childress 2017), writing is also a fundamentally cognitive process. Furthermore, although the specific content is often the subject of explicit consideration and deliberate “framing” (DiMaggio et al. 2013; Wood et al. 2018), there are also intuitive aspects of writing that are often never more explicit than quickly judging whether a string of words “sounds right.” As such, along with practice-theoretic traditions in the study of the production of culture (Bourdieu 1993; Breiger 2000), we argue that there is a *relative homology* between the schemas authors draw upon to produce texts and the latent patterns learned from the texts themselves (Ignatow 2016).

This argument is supported by a growing body of cognitive, psychological, and linguistic research using distributed representations of words (specifically, word embeddings) to recover schematic conceptual information from the statistical properties of natural language corpora. In a review and comprehensive expansion of this line of research, Utsumi (2020) found that, with some variation, many distinct kinds of conceptual knowledge (concrete, abstract, spatial, temporal, perceptual, emotional, etc.) are accurately encoded in word embeddings (hereafter just *embeddings*). For example, one team (Fulda et al. 2017) found that embeddings trained on Wikipedia articles encoded prototypical object affordances, whereas another team found that embeddings encoded visual and evaluative features of objects, such as size and dangerousness (Grand et al. 2018) and can thus be used to accurately model human judgement about a variety of domains (Richie et al. 2019). Also working with word embeddings, Joseph and Morgan (2020) found that these distributional models reveal stereotypical biases (e.g., engineers are men) that largely mirror those found in survey data (see also Caliskan et al. 2017; Kozlowski et al. 2019).

But why do embeddings have this range of ecological validity? At core, this research is based on the “distributional hypothesis” in linguistics, which Harris (1954:156) defined as “difference of meaning correlates with difference of distribution.” Specifically (Arseniev-Koehler and Foster 2020; Caliskan and Lewis 2020), this is because distributed representations are a neurally plausible model of how human brains actually represent meaning—following decades of research on connectionism and artificial neural networks (Günther et al. 2019). Put another way, the key point of similarity between human minds and these computational models of language is that both store information as *distributions*:

In contrast to maintaining slots to represent all possible concepts in the world, storing a concept as a distributed representation means that many concepts can be represented by a limited number of neurons because

even a limited number of neurons may have many, many possible patterns of activation (Arseniev-Koehler and Foster 2020:7)

As a striking demonstration of this, a team was able to predict stories people were reading from functional magnetic resonance imaging data by using semantic representations of the stories obtained with embedding methods. Even more compelling, “this was possible even when the distributed representations were calculated using stories in a different language than the participant was reading” (Dehghani et al. 2017:6096).

Arseniev-Koehler and Foster (2020) further argue that one of the most common word embedding models, word2vec, is also an accurate model for the process by which human minds *acquire* or *build* these distributed representations (see also Caliskan and Lewis 2020). There are alternative procedures to produce word embeddings that have outputs comparable to word2vec, namely GloVe (Pennington et al. 2014). As our proposed method begins with embeddings, it is more or less agnostic to the underlying embedding model that produced them.⁵

Concept Class Analysis: The Building Blocks

Now that we have outlined the mechanics of RCA/CCA and shown how cognitive-schematic information can be recovered from written natural language, we now outline how CCA can be combined with embedding methods to group texts into schematic classes.

Concept Mover's Distance

Concept mover's distance (CMD) is a method for measuring the extent to which a document is similar to a location in a word embedding vector space (Stoltz and Taylor 2019). It begins by conceiving of documents as clouds of locations, drawing on prior work showing that similarity between two documents can be quantified by the “cost” of moving the clouds constituting one document to the locations of the clouds constituting another (Kusner et al. 2015; Atasu et al. 2017). CMD measures the cost of moving all words in a document to one or very few points in the embedding space. In its simplest form, this point is defined by a single word vector. Treating this as a transportation problem, in which term counts constitute amounts to be moved and word vectors define distances to be moved, this becomes a special case of earth mover's distance (EMD; i.e., the Wasserstein metric for comparing discrete probability distributions).

Imagine two documents represented as bags of words (i.e., a vector of unique term counts), doc_1 and doc_2 . These documents are normalized by dividing each by their respective total term count such that the sum of their vectors will equal 1. EMD finds how much of word p in doc_1 has to flow to word q in doc_2 , and the cost of moving word p in doc_1 to word q in doc_2 . Formally, this is defined as follows (Atasu et al. 2017:890; see also Rubner et al. 1998):

$$EMD(doc_1, doc_2) = \min \sum_{p,q} y[p, q] c[p, q] \quad (1)$$

where y and c are $n \times n$ matrices (where n is the total vocabulary in the corpus), $y[p, q]$ is how much of word p in doc_1 flows to q , and $c[p, q]$ is the cosine similarity between words p and q in an embeddings space (i.e., the cost of movement). The amount a word can flow from doc_1 is ultimately constrained by the relative frequency of its nearest neighbors in doc_2 , such that (Atasu et al. 2017:890):

$$\sum_q y[p, q] = doc_1[p], \sum_p y[p, q] = doc_2[q] \quad (2)$$

Let's say that the nearest neighbor of the word "dog" in doc_1 is "cat" in doc_2 , but 20 percent of the words in doc_1 are "dog," whereas only 10 percent of the words in doc_2 are "cat." The remaining 10 percent of doc_1 must be distributed to the next nearest neighbor in doc_2 , and on and on. The row vectors of the resulting flow matrix, y , between doc_1 and doc_2 must therefore sum to the relative frequency of word p in doc_1 (i.e., $\sum_q y_{pq} = doc_1(p)$) (see Kusner et al. 2015:3). Similarly, the columns vectors of y must also sum to the relative frequency of word q in doc_2 — $\sum_p y_{pq} = doc_2(q)$. The strength of this, however, is that the documents can be very different lengths, up to the limiting case where one document is a single word. CMD exploits this strength.

In the simplest case, CMD calculates the cost of moving to a single word denoting a concept of interest—for example, "love." The concept can be further specified either by including additional terms—e.g., "romantic love"—or by finding the average vector of various words denoting the concept of interest—e.g. "love," "romance," "devotion," "affection," and so on. CMD accomplishes this task by treating the word (or words) denoting a concept as a "pseudo-document," appending that pseudo-document to the document-term matrix (consisting of all 0s except for the columns that are for concept-denoting terms, which get 1s), then using a computationally efficient solver for the EMD problem⁶ to arrive at a vector with scores indicating how similar each real corpus document is to the pseudo-document—that is, to the concept of interest.

Finally, and central to our proposed method, we can also measure the cost of moving a document in a particular *direction* defined by juxtaposed terms—e.g., "love" as opposed to "hate." We now move to a discussion of this particular procedure.

Semantic Directions

To understand semantic directions, consider gender bias in texts (Bolukbasi et al. 2016). This is a major area of study using word embeddings because an analyst can easily define a direction toward *women* and away from *men*, for example. This is implicit in the king:man as queen:woman analogy task introduced by Mikolov et al. (2013). If we subtracted "man" from "woman," the result would be a vector that defined a direction pointing toward "woman" and away from "man"—it is not a word vector per se but a kind of *relation vector*. If we measured the cosine similarity of "queen" and "king" with this new relation vector, we would find that queen is close (defined by the cosine of the angle of the vectors) and king is far. But queen and king are obviously gendered labels; what about monarch or castle? We

might find they are somewhere between queen and king. If the cosine similarity is zero—i.e., the angle is perpendicular—this means they are equidistant to both man and woman. However, if they lean to one side or the other, we would say these terms have a bias toward man or woman (Bolukbasi et al. 2016; Caliskan et al. 2017).

To increase the accuracy of this direction, we can find several gendered word vectors corresponding to man on the one hand (e.g., men, gentlemen, boys) and woman on the other (e.g., women, ladies, girls), subtract the vectors for each pair, and then average the result (Kozlowski et al. 2019; Arseniev-Koehler and Foster 2020; Taylor and Stoltz 2020). Therefore, we define a semantic direction, \mathbf{d} , as the mean of a set of vector differences between a collection of juxtaposed word pairs:

$$\mathbf{d} = \frac{\sum_p^P (\mathbf{p}_1 - \mathbf{p}_2)}{P} \quad (3)$$

where p is a word pair in the total set of P relevant juxtaposed word pairs, \mathbf{p}_1 and \mathbf{p}_2 are the vector representations of the two words in juxtaposed pair p , and \mathbf{d} points toward \mathbf{p}_1 and away from \mathbf{p}_2 .⁷

We can then take a sample of terms and measure whether they are closer to one side or the other of this semantic direction to determine a range of gender bias. For example (Jones et al. 2020:14–15), we could determine the gender bias of “Science” by measuring the cosine similarity between our gender relation vector and the following terms: science, technology, physics, chemistry, Einstein, NASA, experiment, and astronomy. We can then average each similarity to arrive at a summarized measure of gender bias in “Science.” If the average cosine is zero, it neither leans toward women (a cosine of 1) nor toward men (a cosine of -1).

We can easily generalize this procedure to domains beyond gender. For example, Kozlowski et al. (2019) estimate directions for affluence and education, whereas Arseniev-Koehler and Foster (2020) estimate directions for health and morality, and Taylor and Stoltz (2020) get directions for life/death and political ideology (see Table 1 for examples of juxtaposing terms for a few semantic directions). To estimate a semantic direction, we simply need a list of juxtaposing terms—or what Arseniev-Koehler and Foster (2020) refer to as “anchor” words—for each dimension of interest within a given domain. This can be accomplished using precompiled dictionaries, prior theory and research, synonym/antonym tools like WordNet (Fellbaum 1998), and seeing which other terms in the word embeddings have high cosine similarity to our anchor terms.

Finally, once we derive a vector for our semantic direction using the above procedure, we can use CMD to measure engagement with a particular pole of a semantic direction. For example, in a corpus of U.S. political texts, one document may engage more with the concept of “conservative” relative to “liberal”—meaning they would have a larger and positive “political ideology” CMD score, assuming that the “liberal” word was subtracted from the “conservative” word across the political ideology word pairs (Taylor and Stoltz 2020). These semantic direction CMDs are the foundational input to CoCA.

Table 1: Examples of juxtaposing terms for semantic directions

HEALTH	
Healthy	healthy, safer, strength, sanitary, clean, healing, heal
Unhealthy	ill, riskier, weakness, filthy, dirty, harming, hurt
AGE	
Young	young, current, juvenile, immature, junior, novice, new, green, budding, fledgling, unaged
Old	old, ancient, elderly, mature, senior, veteran, old, seasoned, withering, expert, aged
POLITICS	
Liberal	liberal, progressive, democrat, untraditional
Conservative	conservative, unprogressive, republican, traditional
LIFE	
Life	life, survivor, birth, living, endure
Death	death, casualty, demise, dying, fatality
EDUCATION	
Educated	educated, learned, knowledgeable, trained, taught, literate, schooled
Uneducated	uneducated, unlearned, ignorant, untrained, untaught, illiterate, unschooled

Robustness Checks for Semantic Directions

As robustness checks, we should consider the following: sensitivity to juxtaposed terms, sensitivity to the method of averaging the vector offsets, and overall face validity. Regarding the first, we could compare each document's engagement with the averaged semantic direction to the component directions defined by the individual juxtaposed term pairs (see Taylor and Stoltz 2020:10) to ensure that a single set of term pairs is not driving the majority of the engagement. Put another way, are all term pairs approximating a similar semantic direction? Second, we could compare our proposed method of averaging (as defined in Equation 3) with potential alternatives (see, for examples, Arseniev-Koehler and Foster 2020:19; Taylor and Stoltz 2020, appendix). Finally, we could check the face validity of the direction by using it as a term-classifier (e.g., Arseniev-Koehler and Foster 2020:19). Using CMD, we could also do this at the level of document classification by checking whether key document-level covariates roughly track engagement as expected.

CCA and Modularity Maximization

Once the semantic direction CMDs have been computed—where each document has a standardized score indicating the extent to which they engage with a particular pole of a binary concept—the last step is to simply pass this document-by-CMD matrix to the CCA algorithm.




Figure 2: (A) Documents and (B) associated pairwise similarities across seven semantic direction CMDs concerning immigration. *Note:* Panel A shows semantic direction CMDs from a political document corpus from four hypothetical documents. Panel B reports each pairwise Euclidean distance (standardized by the maximum possible distance, 10.58301), Pearson correlation, and relationality.

Consider the hypothetical example in Figure 2, a CoCA variant of the RCA/CCA example presented in Figure 1. Instead of four respondents with answers across seven ordinal and equidistant survey items, we now have four documents explicitly discussing immigration (panel A). These documents have seven scores, each indicating the extent to which they engage with one pole of a direction pertaining to the topic of immigration. These scores are continuous.

Documents A and B follow the same cultural schema when they write on immigration issues: they both engage very little with the racial, skills, legality, employment, and education dimensions of immigration discourse (or otherwise engage with the binary oppositions that define those dimensions relatively equally), but display pronounced biases toward discussions of open borders and low economic capital when it comes to the dimensions of borders and socioeconomic status. Just like in the Figure 1 example, document C does not have any of the same CMD scores as A and B, but instead displays an inversion of this pattern. As such, we might say that documents A, B, and C carry “traces” of the same schema: a “closed/open borders and high/low SES” schema, perhaps, where the concepts of borders and socioeconomic capital are highly distinguishing meaning-making features. Document D, however, displays a substantively different relational pattern across the CMD scores, and therefore carries traces of a different schema.

Figure 2 panel B shows the relevance of the (now interdocument) Pearson correlation coefficient as a measure of schematic similarity. The Pearson correlation between documents A and B is precisely 1 (as it should be) because A and B are perfect linear transformations of one another with only a vertical shift (Boutyline 2017:356) necessary to turn one into the other ($A = 0.4 + B$; $B = -0.4 + A$). Similarly, the Pearson correlation between A and C and B and C are appropriately -1 —because, as before, A and C and B and C are perfect linear transformations of one another ($C = 0.1 - A$; $C = -0.3 - B$). The variance of document D, however, is not fully accounted for by its bivariate relationship with any of

the other documents ($D = 0.4548 + 0.2903A + \epsilon_1$, $D = 0.571 + 0.2903B + \epsilon_2$, and $D = 0.4839 - 0.2903C + \epsilon_3$, where $\epsilon_{1,2,3}$ are residual variance terms and where each $\epsilon_i \neq 0$). As expected, the Pearson correlations between D and the other three documents are close to zero.

The absolute values of the interdocument correlations are taken prior to subsetting the documents into schematic classes so that documents with inverted patterns are grouped together. Specifically, the symmetric document-by-document matrix, A , of absolute Pearson correlations is treated as a fully connected and weighted graph and partitioned using modularity maximization (Newman 2004; Goldberg 2011:1409–1410; Boutyline 2017:386). CoCA returns the set of schematic classes that maximizes graph A 's modularity score (after optionally decreasing the density of the graph by removing statistically nonsignificant correlations at some researcher-specified α -level).⁸ A higher graph modularity score, Q , indicates stronger absolute correlations between documents assigned to the same class than one would expect if the correlations were placed between documents at random.

As a class detection method, this modularity-based technique has the benefit of reporting the class solution that best maximizes Q —as opposed to, for example, agglomerative clustering, which requires researcher decisions on where to “cut” and define the classes. Lastly, modularity maximization also aligns with cognitive “theories of schematic transmission,” as Goldberg notes (2011:1410).⁹

Simulation Analysis

As a document's engagement with semantic directions is both continuous and less precise than are the ordinal scales found in survey responses, it is important to determine whether CCA is robust to variation on continuous measures. Therefore, we present a simulation analysis to illustrate how CoCA accurately identifies “ground truth” schemas from a range of artificial corpora with varying amounts of random noise injected into the schematic patterns.¹⁰ Next, we use one of the simulated class solutions to present some schema-specific interdirection correlation networks—common visualization strategies for interpreting the derived cultural schemas.

Partition Validity

Boutyline illustrated how CCA is robust to varying levels of noise (2017:365–366). We ran a series of simulations to test how accurate CCA continued to be as a schema-partitioning method when the input data are continuous (semantic direction CMDs) rather than ordinal (survey items with Likert and Likert-type response scales).

We ran a series of simulations to shed light on this question. First, we created three separate within-document, cross-direction patterns that varied within a $[-3, 3]$ interval: one that follows document A 's pattern in Figure 2, one that follows document D 's pattern in Figure 2, and another third pattern. The three patterns are shown in Figure 3, in blue. Then, noting that two documents i and j (with x and y being i 's and j 's score vectors across a set of semantic direction CMDs, respectively) have traces of the exact same schema if $y = \alpha + \beta x$ and where $\beta \neq 0$, we created

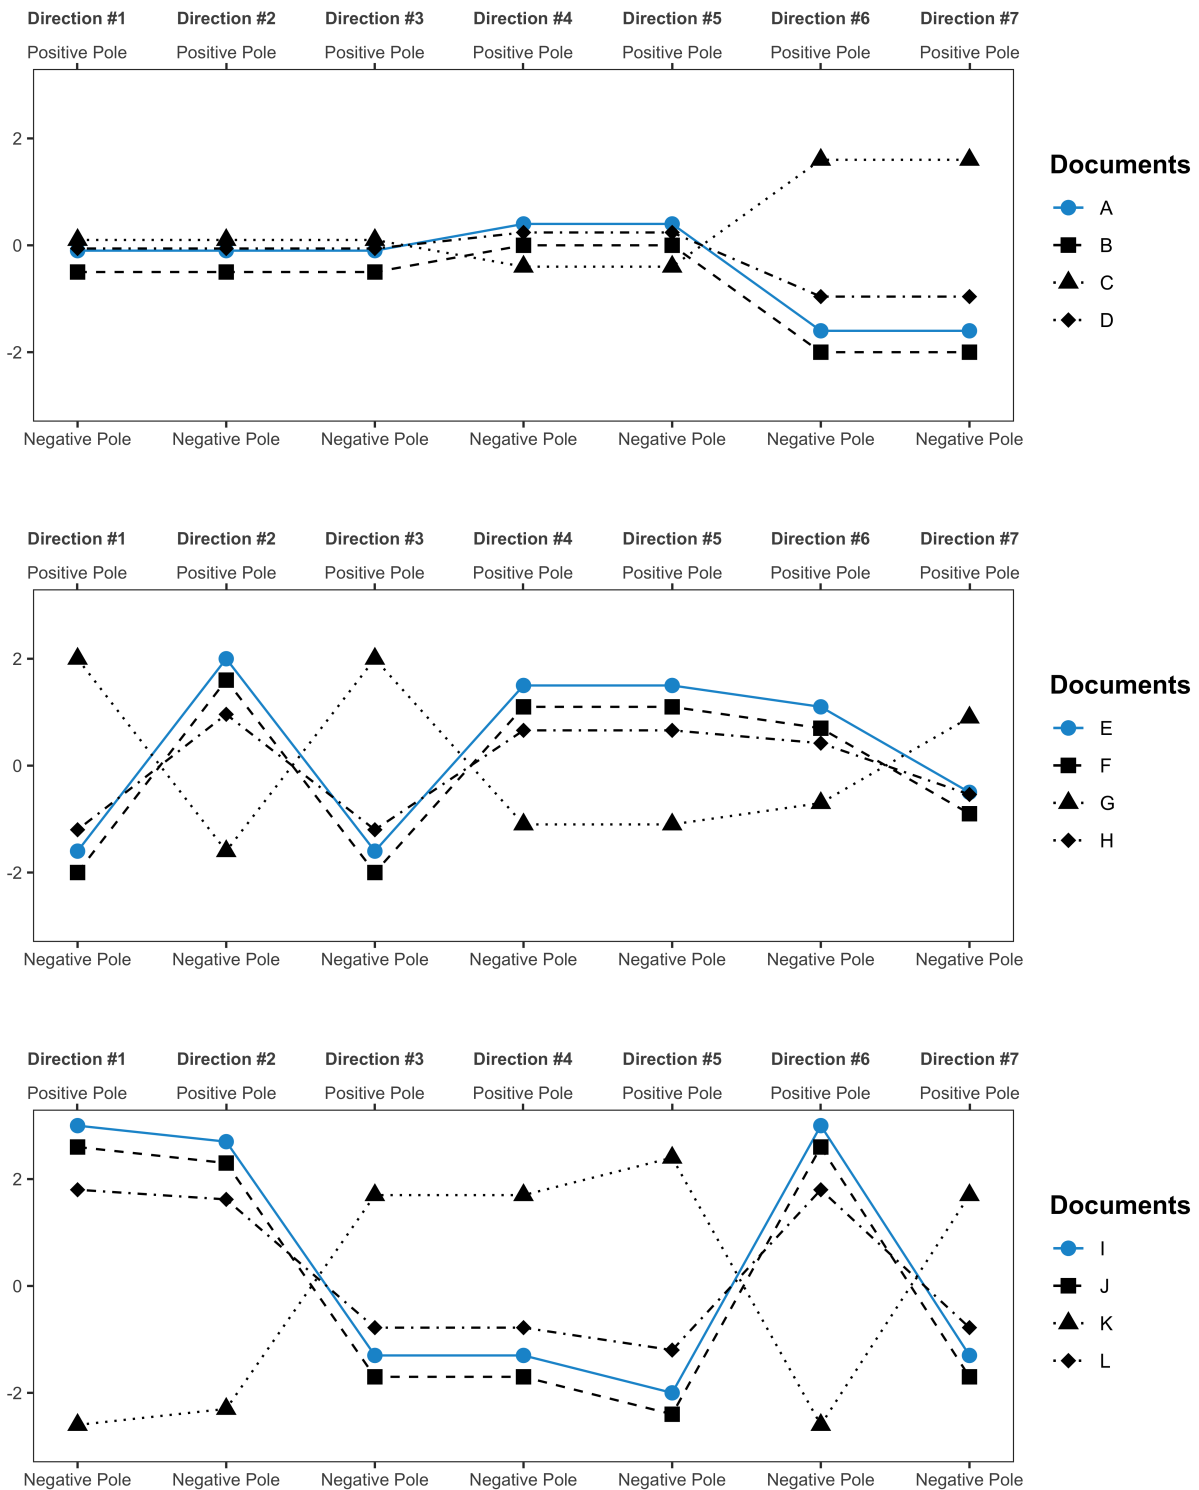


Figure 3: Three simulated schematic patterns. *Note:* The blue lines indicate the patterns from which each of the transformed patterns were derived.

Table 2: Pearson correlations between simulated documents

	A	B	C	D	E	F	G	H	I	J	K	L
A	1	1	-1	1								
B	1	1	-1	1								
C	-1	-1	1	-1								
D	1	1	-1	1								
E					1	1	-1	1				
F					1	1	-1	1				
G					-1	-1	1	-1				
H					1	1	-1	1				
I									1	1	-1	1
J									1	1	-1	1
K									-1	-1	1	-1
L									1	1	-1	1

Note: Only within-set correlations shown.

Table 3: Cramér's V summary statistics

Mean	SD	Min	25p	50p	75p	Max
0.81	0.15	0.47	0.66	0.84	0.95	1

three separate linear transformations of each of these patterns. One pattern shifted by $\alpha = -0.4$, one scaled by $\beta = 0.6$, and we then had an inversion of the original pattern ($y = -x$). Those transformed patterns are visualized with each original pattern in Figure 3. We then repeated each of these 12 patterns 10 times to create a simulated data set of 120 documents across seven semantic direction CMDs.

Each pattern correlates perfectly with every other pattern within the same set (see Table 2). This means that each set of documents— $\{A, B, C, D\}$, $\{E, F, G, H\}$, and $\{I, J, K, L\}$, respectively—follow their own schematic pattern. Indeed, a CCA of this $N = 120$ corpus accurately partitions the documents into three schematic classes, with 40 documents in each.

We then created 981 additional simulated data sets, each with random noise injected into the within-document, cross-direction patterns (i.e., the document row vectors). Each document across these data sets was perturbed with a random deviation drawn from a uniform distribution (see the “jitter” function, a base feature in R Core Team 2020). Let the “noise factor” be f . For the first of the 981 data sets, $f = 1$, corresponding to a relatively small amount of noise. Each subsequent data set was perturbed with an additional 0.05 to f , leading to the final data set (number 981) that had a noise factor of 50. CCA was run on each of the simulations, leading to 982 total schematic class solutions (the original, representing the true class partition, and 981 class assignments from the noise-added simulations).

How accurate was CCA at finding the correct class solution with varying levels of random noise? To assess this, we calculated Cramér's V between the true class solution and each of the other 981 simulated solutions. A Cramér's V of 1 indicates a perfect association: each category of the row variable (the true three-class solution) corresponds perfectly to one (and only one) category of the column variable (one of the simulated solutions). A Cramér's V of 0 indicates a purely random association, and values between 0 and 1 indicate approximations of these extremes. The 981 V statistics are summarized in Table 3. The mean V was 0.81 with a standard

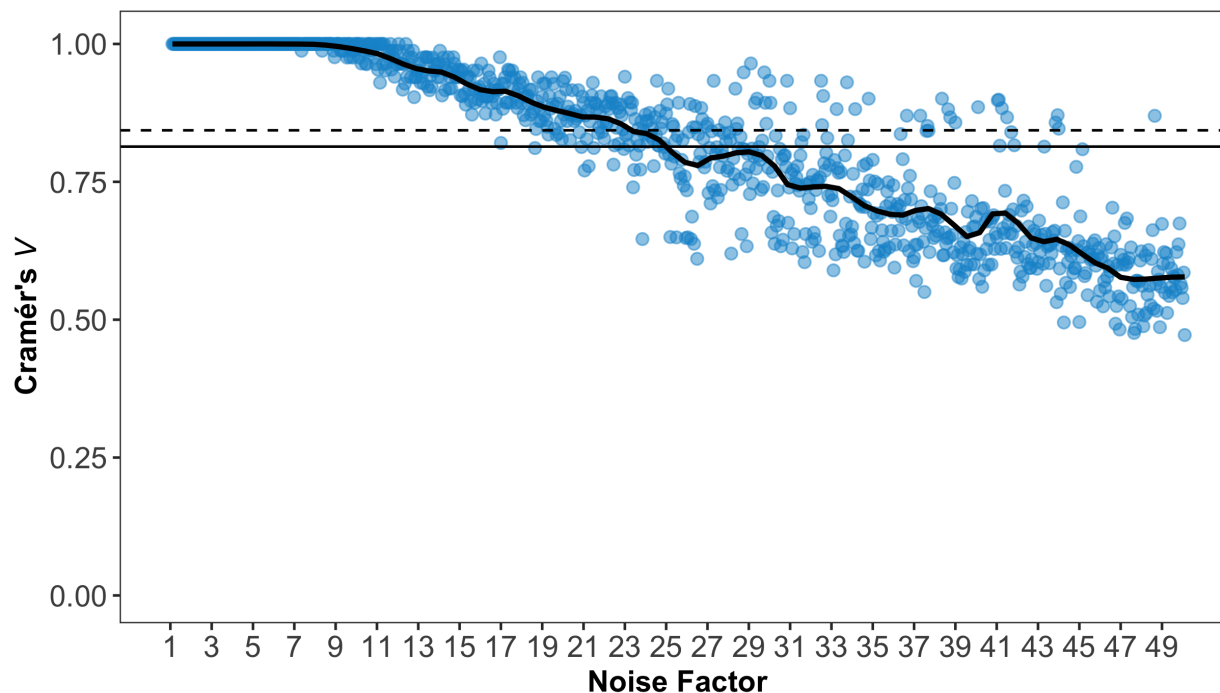


Figure 4: Accuracy. *Note:* The solid line is the mean Cramér's V ; the dashed line is the median.

deviation of 0.15; the median was 0.84. The interquartile range was 0.29, with a 25th percentile of 0.66 and a 75th percentile of 0.95. The maximum observed V was a perfect association of 1—which occurred 184 times (about 19 percent of the time).

These results suggest that CCA does an adequate job of finding the correct class solutions when the input data are continuous (i.e., semantic direction CMDs), even at high levels of noise. Indeed, as Figure 4 shows, most of the lower V statistics were from simulations with relatively high levels of noise. Altogether, these simulations suggest the CCA is a robust method for detecting schematic patterns across semantic direction CMDs.

Correlation Networks

In a simulation such as the one just described, the schematic patterns are known by design. Of course, in the vast majority of empirical use cases, this will not be the case. How, then, can researchers interpret the schema associated with each class after deriving them?

In RCA/CCA, the norm is to present class-specific inter-item correlation networks. The same visualization strategy can be used with CoCA, but with semantic direction CMDs rather than survey variables. The result, then, is interdirection correlation networks in which the edges indicate the extent that engagement with a binary concept is predictive of engagement with another binary concept.

Consider the initial simulation data set, consisting of 120 documents and with each document a perfect linear transformation of one of three schematic patterns

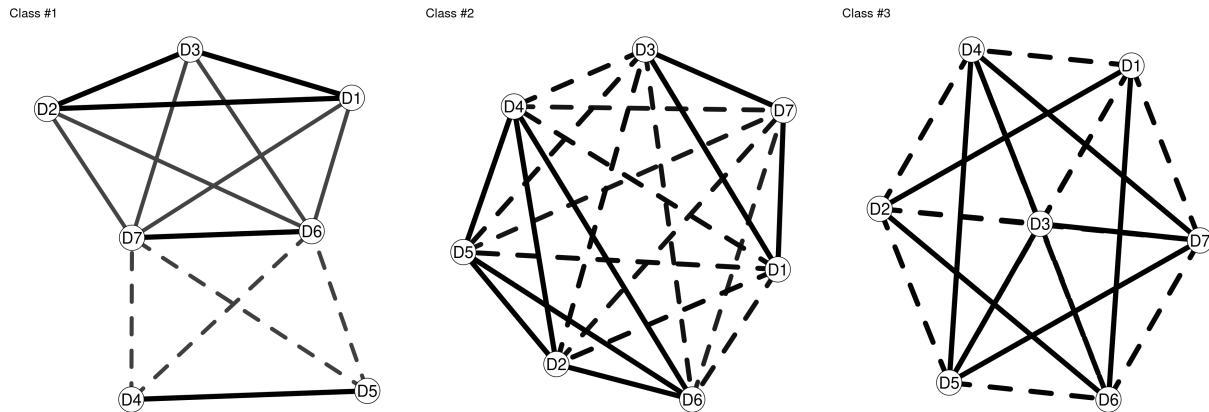


Figure 5: Graphs of three simulated schematic patterns. *Note:* The correlation networks plot all statistically significant Pearson correlation coefficients ($\alpha = 0.5$, two-tailed). Correlations where $|r|_{ij} = 0.15$ are not visualized. Positive correlations are shown with straight edges; negative correlations are shown as dashed edges. Darker and thicker edges represent stronger bivariate correlations. The network is spatially arranged using the Fruchterman-Reingold force-directed algorithm so that directions with stronger correlations are placed closer together and those with weaker correlations are “repulsed” from one another (with a repulsion factor equal to 1.86). All classes consist of 40 documents.

(see Figure 3). As stated above, CoCA accurately partitioned this set of documents into three schematic classes, each consisting of 40 documents. Correlation networks for each class are presented in Figure 5. The left panel corresponds to the top schematic class in Figure 3, the middle to the middle schema, and the right panel corresponds to the bottom schema.

Correlation networks such as these provide the information necessary to “re-construct” the schemas that constitute each class. For example, looking at the Class #1 network in Figure 5, we see that semantic directions #4 and #5 are strongly and positively correlated. This is in line with the schema as shown in the top panel of Figure 3, which shows that each document engages the same poles of each direction. Let’s say, for instance, direction #4 was an “Employment” direction opposing “employed” with “unemployed” and direction #5 was an “Education” direction opposing “educated” with “uneducated,” then, if a document in Class #1 engages with the concept of “employed” (in opposition to “unemployed”), that document also engages the concept of “educated” (as opposed to “uneducated”)—and vice versa. We also see negative correlations between direction #7 and direction #5 on the Class #1 network. This again is in line with the top panel of Figure 3, where engagement with direction #5 is inversely related to #7 for each document. If, for example, #5 is “Education” and #7 is “Class” (opposing “rich” with “poor”), then documents that engage highly with the “rich” pole of #7 will also engage highly with the “uneducated” pole of #5 (and vice versa).

Discussion and Conclusion

Summary

In this article, we proposed a workflow for recovering schemas in naturally occurring language. Defining cultural schemas as slowly learned, implicit, and unevenly distributed relational memory structures that allow people to make sense of the world, recent work has demonstrated ways of recovering schemas—or, rather, the downstream consequences of people drawing upon them—from domain-specific survey modules. We extended these methods to text data. Specifically, we draw on “schematic class analysis,” which sorts respondents into “classes” on the basis of the schema to which their survey response patterns best align, and introduce concept class analysis (CoCA): a hybrid model that combines word embeddings and correlational class analysis to group documents across a corpus by the similarity of schemas recovered from them.

Word embeddings, which use term co-occurrence statistics to determine semantic similarity between words, form the backbone of this workflow. We can then aggregate up from the meaning of words to the meaning of documents by conceiving of documents as clouds of locations defined by the words they contain.

Next, we proposed a deductive method for measuring a document’s engagement with specific concepts by measuring the cost of moving all the words contained in a document to a specific vector in the N -dimensional semantic space (Stoltz and Taylor 2019). Although vectors may be specific terms denoting a concept, we can extract “semantic directions,” which point toward one pole of a binary concept and away from the other pole (Taylor and Stoltz 2020). For example, a document may “point” toward “feminine” and away from “masculine.”

Finally, by measuring a document’s engagement with several key directions as continuous variables, we can then apply correlational class analysis to extract classes of similar engagement patterns (Goldberg 2011; Boutyline 2017). Using this workflow allows the text analyst to partition documents by recovering the “traces” of schematic information in the texts.

Further Research and Applications

As stated at the onset, one of the key strengths of our proposed workflow is that, unlike surveys, the set of meaningful dimensions that can be estimated are not limited by the questions of a survey instrument. Therefore, if theory and empirical research suggest other semantic directions might be salient for a particular question or corpus, the analyst only needs a list of juxtaposed terms that approximate this dimension.

It may be the case, however, that certain directions are more tightly organized as continuous bipolarities—such as gender in many contexts—whereas other semantic directions that could be estimated (such as movies juxtaposed to books or college versus professional sports) may not be so cleanly organized. Furthermore, although prior research, theory, and semantic dictionaries (e.g., WordNet) can aid in selecting juxtaposed term pairs, there is no agreed upon number of pairs required. Thus this greater flexibility for estimating variables along which documents may be

partitioned comes with the additional burden of considering construct validity for each distinct dimension.

Related, the fit between context and dimensions of interest directly bears on selecting an appropriate corpus. That is, like network analysis, text analysis runs into a boundary specification problem. Laumann, Marsden, and Prensky (1989:18–9) note that, for example, “it is obviously of great consequence if a key intervening actor or ‘bridging’ tie is omitted due to oversight or use of data that are merely convenient,” and that “misrepresentation of the process under study” is “precisely the outcome of errors in the definition of system boundaries in a network analysis.” Relational analyses using texts must consider this same issue. CoCA falls squarely within this discussion because the recovered schemas are interpreted from document classes partitioned on the basis of similarities in edge placement between documents in the corpus. In other words, like most dimension reduction techniques, the schemas are derived from patterns *internal* to the data themselves.

RCA and CCA analyses of surveys can often bypass this concern because they can be estimated using probability samples derived from representative sampling frames where each population member has a nonzero probability of selection. Most text corpora are not compiled with probability theory and representativeness in mind. Future methodological research should therefore think more critically about some of the statistical steps in CoCA—namely, the use of asymptotically derived p values and “statistical significance” to filter out correlations in both the interdocument graph and the interdimension correlation networks where the edges indicate whether engagement with a binary concept is predictive of engagement with another binary concept.

Finally, any corpus can be used with our proposed workflow (in fact, one does not even need the raw texts—the texts represented as a “bag-of-words” document-term matrix is sufficient). Furthermore, CoCA may even be particularly well suited for analyzing open-ended questions from surveys and interviews, especially those with well-defined domains and individual “authors” (i.e., respondents and interviewees, respectively). There has been text-analytic work on these types of data—e.g., analysts have applied topic modeling to open-ended responses from experiments and surveys (Roberts et al. 2014; Finch et al. 2018; Pietsch and Lessmann 2018) and unsupervised document clustering to interview transcripts (Janasik et al. 2009; Sherin 2013). We believe CoCA has much potential with these types of data: the single-author nature of each “document” and, importantly, their domain-specific content (where the participants are answering or ruminating on a specific, researcher-delivered question or prompt) means that CoCA is better able to isolate individual-level schematic traces. Additionally, open-ended responses from representative surveys may alleviate some of the statistical concerns raised in the preceding paragraph (assuming that there is also no nonresponse bias in the open-ended questions). We hope to see these types of CoCA applications in the near future.

Notes

- 1 See Wood et al. (2018) and Arseniev-Koehler and Foster (2020) for a discussion of some of the various types of sociologically relevant schemas and how they are different from the related concept of frames.
- 2 As Fiske and Taylor note (2013:15), “[t]he idea behind [the cognitive miser model] is that people are limited in their capacity to process information, so they take shortcuts whenever they can.... People adopt strategies that simplify complex problems; the strategies may not be correct processes or produce correct answers, but they are efficient. The capacity-limited thinker searches for rapid, adequate solutions rather than for slow, accurate solutions” (see also Taylor 1981; Spears et al. 1999).
- 3 It is an important, although minor point, to note that we consider schemas *cultural* so long as they are learned—the extent schemas may be shared is an important empirical question but makes them no more or less “cultural” per se (see White 1959; Weiss 1973).
- 4 After drafting the current article, we discovered one earlier article that applies RCA to texts, Miranda et al. (2015). The authors, along with three research assistants, hand-coded firms’ documents discussing their “initiatives” regarding social media. The documents were coded for the presence of six “principles.” Each unique initiative then received a score for each principle: the number of times a principle was present in a document divided by the total number of documents relating to the initiative. The authors then apply RCA to the pattern of scores to extract schematic classes.
- 5 The key difference is that GloVe begins with a matrix that represents the global term co-occurrence statistics for a given corpus; by contrast, word2vec iterates through the text, slowly extracting co-occurrence statistics (Levy and Goldberg 2014; Levy et al. 2015). The former is not neurally plausible because humans do not begin with schemas weighted by “global” probabilities of features in some domain. Rather, like word2vec, humans slowly accrue these probabilities from their ongoing lived experience.
- 6 EMD is computationally demanding. There are several efficient solvers that provide good enough approximations. CMD relies on the “relaxed word mover’s distance” (RWMD) algorithm, originally detailed by Kusner and colleagues (2015). Readers interested in the technical details relating RWMD to CMD are encouraged to read the original CMD article by Stoltz and Taylor (2019). Lastly, because the CMD algorithm relies heavily on the text2vec R package (Selivanov and Wang 2016), CMD uses a considerably faster approximation of RWMD written into the package known as linear complexity relaxed word mover’s distance (Atasu et al. 2017).
- 7 See appendix in Taylor and Stoltz (2020) for alternative ways to estimate a semantic direction.
- 8 RCA and CCA—and by extension, CoCA—use the leading eigenvector method for approximating the modularity score because full maximization can be quite computationally demanding (see Newman 2006).
- 9 As Boutyline (2017:386) points out, though, more research is needed to compare modularity maximization to other community detection techniques for schematic class analysis.
- 10 We focus on “partition validity” here and work from the assumption that the semantic direction CMDs have already been validated—see Stoltz and Taylor (2019) and Taylor and Stoltz (2020) for more on CMD validation.

References

- Arseniev-Koehler, Alina, and Jacob G. Foster. 2020. "Machine Learning as a Model for Cultural Learning: Teaching an Algorithm What It Means to Be Fat." Preprint, submitted March 24. <https://arxiv.org/abs/2003.12133>.
- Atasu, Kubilay, Thomas Parnell, Celestine Dünner, Manolis Sifalakis, Haralampos Pozidis, Vasileios Vasileiadis, Michail Vlachos, Cesar Berrospi, and Abdel Labbi. 2017. "Linear-Complexity Relaxed Word Mover's Distance with GPU Acceleration." Paper presented at the IEEE International Conference on Big Data, December 11–14, Boston, MA. <https://doi.org/10.1109/BigData.2017.8258005>.
- Baldassarri, Delia, and Amir Goldberg. 2014. "Neither Ideologues nor Agnostics: Alternative Voters' Belief System in an Age of Partisan Politics." *American Journal of Sociology* 120(1):45–95. <https://doi.org/10.1086/676042>.
- Barbet, Berta. 2020. "Nobody Seems to Be Fully Representing Me: Differential Inter-attitudinal Cohesion Systems and Their Effects on Satisfaction with the Political System." *Electoral Studies* 64:102–116. <https://doi.org/10.1016/j.electstud.2019.102116>.
- Blair-Loy, Mary. 2001. "Cultural Constructions of Family Schemas: The Case of Women Finance Executives." *Gender and Society* 15(5):687–709. <https://doi.org/10.1177/089124301015005004>.
- Bolukbasi, Tolga, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam T. Kalai. 2016. "Man Is to Computer Programmer as Woman Is to Homemaker? Debiasing Word Embeddings." Preprint, submitted July 21. <https://arxiv.org/abs/1607.06520>.
- Bourdieu, Pierre. 1993. *The Field of Cultural Production: Essays on Art and Literature*. New York, NY: Columbia University Press.
- Boutyline, Andrei. 2017. "Improving the Measurement of Shared Cultural Schemas with Correlational Class Analysis: Theory and Method." *Sociological Science* 4(15):353–93. <https://doi.org/10.15195/v4.a15>.
- Boutyline, Andrei, and Laura Soter. 2020. "Cultural Schemas: What They Are, How to Find Them, and What to Do Once You've Caught One." Preprint, submitted May 12. <https://doi.org/10.31235/osf.io/ksf3v>.
- Boutyline, Andrei, and Stephen Vaisey. 2017. "Belief Network Analysis: A Relational Approach to Understanding the Structure of Attitudes." *American Journal of Sociology* 122(5):1371–447. <https://doi.org/10.1086/691274>.
- Breiger, Ronald L. 2000. "A Tool Kit for Practice Theory." *Poetics* 27(2):91–115. [https://doi.org/10.1016/S0304-422X\(99\)00026-1](https://doi.org/10.1016/S0304-422X(99)00026-1).
- Caliskan, Aylin, Joanna J. Bryson, and Arvind Narayanan. 2017. "Semantics Derived Automatically from Language Corpora Contain Human-like Biases." *Science* 356(6334):183–6. <https://doi.org/10.1126/science.aal4230>.
- Caliskan, Aylin, and Molly Lewis. 2020. "Social Biases in Word Embeddings and Their Relation to Human Cognition." Preprint, submitted July 16. <https://doi.org/10.31234/osf.io/d84kg>.
- Cerulo, Karen A. 2010. "Mining the Intersections of Cognitive Sociology and Neuroscience." *Poetics* 38(2):115–32. <https://doi.org/10.1016/j.poetic.2009.11.005>.

- Childress, Clayton. 2017. *Under the Cover: The Creation, Production, and Reception of a Novel*. Princeton, NJ: Princeton University Press. <https://doi.org/10.23943/princeton/9780691160382.001.0001>.
- Daenekindt, Stijn. 2019. "Out of Tune. How People Understand Social Exclusion at Concerts." *Poetics* 74:101341. <https://doi.org/10.1016/j.poetic.2018.12.002>.
- Daenekindt, Stijn, Willem de Koster, and Jeroen van der Waal. 2017. "How People Organise Cultural Attitudes: Cultural Belief Systems and the Populist Radical Right." *West European Politics* 40(4):791–811. <https://doi.org/10.1080/01402382.2016.1271970>.
- Dehghani, Morteza, Reihane Boghrati, Kingson Man, Joe Hoover, Sarah I. Gimbel, Ashish Vaswani, Jason D. Zevin, Mary Helen Immordino-Yang, Andrew S. Gordon, Antonio Damasio, and Jonas T. Kaplan. 2017. "Decoding the Neural Representation of Story Meanings across Languages." *Human Brain Mapping* 38(12):6096–106. <https://doi.org/10.1002/hbm.23814>.
- DiMaggio, Paul. 1997. "Culture and Cognition." *Annual Review of Sociology* 23(2):263–87. <https://doi.org/10.1146/annurev.soc.23.1.263>.
- DiMaggio, Paul, and Amir Goldberg. 2018. "Searching for Homo Economicus: Variation in Americans' Construals of and Attitudes toward Markets." *European Journal of Sociology* 59(2):151–89. <https://doi.org/10.1017/S0003975617000558>.
- DiMaggio, Paul, Manish Nag, and David Blei. 2013. "Exploiting Affinities between Topic Modeling and the Sociological Perspective on Culture: Application to Newspaper Coverage of U.S. Government Arts Funding." *Poetics* 41(6):570–606. <https://doi.org/10.1016/j.poetic.2013.08.004>.
- DiMaggio, Paul, Ramina Sotoudeh, Amir Goldberg, and Hana Shepherd. 2018. "Culture out of Attitudes: Relationality, Population Heterogeneity and Attitudes toward Science and Religion in the US." *Poetics* 68:31–51. <https://doi.org/10.1016/j.poetic.2017.11.001>.
- Fellbaum, Christiane, ed. 1998. *WordNet: An Electronic Lexical Database*. Cambridge, MA: MIT Press. <https://doi.org/10.7551/mitpress/7287.001.0001>.
- Finch, W. Holmes, Maria E. Hernández Finch, Constance E. McIntosh, and Claire Braun. 2018. "The Use of Topic Modeling with Latent Dirichlet Analysis with Open-ended Survey Items." *Translational Issues in Psychological Science* 4(4):403–24. <https://doi.org/10.1037/tps0000173>.
- Fiske, Susan T., and Shelley E. Taylor. 2013. *Social Cognition: From Brains to Culture*. Los Angeles, CA: SAGE Publications.
- Fleck, Ludwik. 1979. *Genesis and Development of a Scientific Fact*. Chicago, IL: University of Chicago Press.
- Fulda, Nancy, Daniel Ricks, Ben Murdoch, and David Wingate. 2017. "What Can You Do with a Rock? Affordance Extraction via Word Embeddings." Preprint, submitted March 9. <https://arxiv.org/abs/1703.03429>.
- Ghosh, Vanessa E., and Asaf Gilboa. 2014. "What Is a Memory Schema? A Historical Perspective on Current Neuroscience Literature." *Neuropsychologia* 53:104–14. <https://doi.org/10.1016/j.neuropsychologia.2013.11.010>.

- Goldberg, Amir. 2011. "Mapping Shared Understandings Using Relational Class Analysis: The Case of the Cultural Omnivore Reexamined." *American Journal of Sociology* 116(5):1397–436. <https://doi.org/10.1086/657976>.
- Grand, Gabriel, Idan Asher Blank, Francisco Pereira, and Evelina Fedorenko. 2018. "Semantic Projection: Recovering Human Knowledge of Multiple, Distinct Object Features from Word Embeddings." Preprint, submitted February 5. <https://arxiv.org/abs/1802.01241>.
- Günther, Fritz, Luca Rinaldi, and Marco Marelli. 2019. "Vector-Space Models of Semantic Representation from a Cognitive Perspective: A Discussion of Common Misconceptions." *Perspectives on Psychological Science* 14(6):1006–33. <https://doi.org/10.1177/1745691619861372>.
- Harris, Zellig S. 1954. "Distributional Structure." *Word* 10(2–3):146–62.
- Homan, Patricia, Lauren Valentino, and Emi Weed. 2017. "Being and Becoming Poor: How Cultural Schemas Shape Beliefs about Poverty." *Social Forces* 95(3):1023–48. <https://doi.org/10.1080/00437956.1954.11659520>.
- Hu, Anning, and Chen Yin. 2020. "Schematic Categorization of Governments Moderates the Association between the Orientation toward Authority and Government Trust in China." *Poetics*, first published on July 24, 2020 as <https://doi.org/10.1016/j.poetic.2020.101478>.
- Hunzaker, M. B. Fallin. 2014. "Making Sense of Misfortune: Cultural Schemas, Victim Redefinition, and the Perpetuation of Stereotypes." *Social Psychology Quarterly* 77(2):166–84. <https://doi.org/10.1177/0190272514521219>.
- Hunzaker, M. B. Fallin, and Lauren Valentino. 2019. "Mapping Cultural Schemas: From Theory to Method." *American Sociological Review* 84(5):950–81. <https://doi.org/10.1177/0003122419875638>.
- Ignatow, Gabe. 2016. "Theoretical Foundations for Digital Text Analysis." *Journal for the Theory of Social Behaviour* 46(1):104–20. <https://doi.org/10.1111/jtsb.12086>.
- Janasik, Nina, Timo Honkela, and Henrik Bruun. 2009. "Text Mining in Qualitative Research: Application of an Unsupervised Learning Method." *Organizational Research Methods* 12(3):436–60. <https://doi.org/10.1177/1094428108317202>.
- Jones, Jason J., Mohammad Ruhul Amin, Jessica Kim, and Steven Skiena. 2020. "Stereotypical Gender Associations in Language Have Decreased over Time." *Sociological Science* 7:1–35. <https://doi.org/10.15195/v7.a1>.
- Joseph, Kenneth, and Jonathan H. Morgan. 2020. "When Do Word Embeddings Accurately Reflect Surveys on Our Beliefs about People?" Preprint, submitted April 25. <https://arxiv.org/abs/2004.12043>.
- Kozlowski, Austin C., Matt Taddy, and James A. Evans. 2019. "The Geometry of Culture: Analyzing the Meanings of Class through Word Embeddings." *American Sociological Review* 84(5):905–49. <https://doi.org/10.1177/0003122419877135>.
- Kusner, Matt, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. "From Word Embeddings to Document Distances." Paper presented at the 32nd International Conference on Machine Learning, July 7–9, Lille, France. <http://proceedings.mlr.press/v37/kusnerb15.html>.

- Laumann, Edward O., Peter V. Marsden, and David Prensky. 1989. "The Boundary Specification Problem in Network Analysis." In *Research Methods in Social Network Analysis*, edited by L. C. Freeman, D. R. White, and A. K. Romney, pp. 61–88. Fairfax, VA: George Mason University Press.
- Lee, Monica, and John Levi Martin. 2018. "Doorway to the Dharma of Duality." *Poetics* 68:18–30. <https://doi.org/10.1016/j.poetic.2018.01.001>.
- Levy, Omer, and Yoav Goldberg. 2014. "Neural Word Embedding as Implicit Matrix Factorization." Paper presented at the 27th International Conference on Neural Information Processing Systems, December 8–13, Quebec, Canada. <http://papers.nips.cc/paper/5477-neural-word-embedding-as-implicit-matrix-factorization.pdf>.
- Levy, Omer, Yoav Goldberg, and Ido Dagan. 2015. "Improving Distributional Similarity with Lessons Learned from Word Embeddings." *Transactions of the Association for Computational Linguistics* 3:211–25.
- Lizardo, Omar. 2017. "Improving Cultural Analysis: Considering Personal Culture in Its Declarative and Nondeclarative Modes." *American Sociological Review* 82(1):88–115. https://doi.org/10.1162/tac1_a_00134.
- Lizardo, Omar, Brandon Sepulvado, Dustin S. Stoltz, and Marshall A. Taylor. 2020. "What Can Cognitive Neuroscience Do for Cultural Sociology?" *American Journal of Cultural Sociology* 8(1):3–28. <https://doi.org/10.1057/s41290-019-00077-8>.
- McDonnell, Erin Metz, Dustin S. Stoltz, and Marshall A. Taylor. 2020. "Multiple Market Moralities: Identifying Distinct Patterns in How Consumers Evaluate the Fairness of Price Changes." *Socio-Economic Review*, first published on October 9, 2020 as <https://doi.org/10.1093/ser/mwaa034>.
- Mikolov, Tomáš, Wen-tau Yih, and Geoffrey Zweig. 2013. "Linguistic Regularities in Continuous Space Word Representations." Paper presented at the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, June 9–14, Atlanta, GA. <https://www.aclweb.org/anthology/N13-1090/>.
- Miles, Andrew. 2019. "An Assessment of Methods for Measuring Automatic Cognition." In *The Oxford Handbook of Cognitive Sociology*, edited by W. H. Brekhus and G. Ignatow. Oxford, England: Oxford University Press.
- Miranda, Shaila M., Inchan Kim, and Jama D. Summers. 2015. "Jamming with Social Media: How Cognitive Structuring of Organizing Vision Facets Affects IT Innovation Diffusion." *Mis Quarterly* 39(3):591–614. <https://doi.org/10.25300/MISQ/2015/39.3.04>.
- Mohr, John W., Christopher A. Bail, Margaret Frye, Jennifer C. Lena, Omar Lizardo, Terence E. McDonnell, Ann Mische, Iddo Tavory, and Frederick F. Wherry. 2020. *Measuring Culture*. New York, NY: Columbia University Press. <https://doi.org/10.7312/mohr18028>.
- Morris, Michael W., and Aurelia Mok. 2011. "Isolating Effects of Cultural Schemas: Cultural Priming Shifts Asian-Americans' Biases in Social Description and Memory." *Journal of Experimental Social Psychology* 47(1):117–26. <https://doi.org/10.1016/j.jesp.2010.08.019>.
- Newman, Mark EJ. 2004. "Finding and Evaluating Community Structure in Networks." *Physical Review E* 69:1–16. <https://doi.org/10.1103/PhysRevE.69.026113>.

- Newman, Mark E. J. 2006. "Finding Community Structure in Networks Using the Eigenvectors of Matrices." *Physical Review E* 74(3):036104. <https://doi.org/10.1103/PhysRevE.74.036104>.
- Pennington, Jeffrey, Richard Socher, and Christopher D. Manning. 2014. "GloVe: Global Vectors for Word Representation." Paper presented at the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), October 25–29, Doha, Qatar. <https://doi.org/10.3115/v1/D14-1162>.
- Pietsch, Andra-Selina, and Stefan Lessmann. 2018. "Topic Modeling for Analyzing Open-ended Survey Responses." *Journal of Business Analytics* 1(2):93–116. <https://doi.org/10.1080/2573234X.2019.1590131>.
- R Core Team. 2020. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rawlings, Craig M., and Clayton Childress. 2019. "Emergent Meanings: Reconciling Dispositional and Situational Accounts of Meaning-Making from Cultural Objects." *American Journal of Sociology* 124(6):1763–809. <https://doi.org/10.1086/703203>.
- Richie, Russell, Wanling Zou, and Sudeep Bhatia. 2019. "Distributional Semantic Representations Predict High-Level Human Judgment in Seven Diverse Behavioral Domains." Paper presented at the 41st Annual Meeting of the Cognitive Science Society, July 24–27, Montreal, Canada. <https://cogsci.mindmodeling.org/2019/papers/0454/index.html>.
- Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson, and David G. Rand. 2014. "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* 58(4):1064–82. <https://doi.org/10.1111/ajps.12103>.
- Rossoni, Luciano, Clayton Pereira Gonçalves, Mônica Pereira da Silva, and Alex Ferreira Gonçalves. 2020. "Mapping Organizational Culture Schemes Based on Correlational Class Analysis." Preprint, submitted May 6. <https://doi.org/10.31235/osf.io/sf2v4>.
- Rubner, Yossi, Carlo Tomasi, and Leonidas J. Guibas. 1998. "A Metric for Distributions with Applications to Image Databases." Paper presented at the Sixth International Conference on Computer Vision, January 4–7, Bombay, India. <https://doi.org/10.1109/ICCV.1998.710701>.
- Rumelhart, David E. 1980. "Schemata: The Building Blocks of Cognition." In *Theoretical Issues in Reading Comprehension*, pp. 33–58. Taylor and Francis Group.
- Selivanov, Dmitriy, and Qing Wang. 2016. *text2vec: Modern Text Mining Framework for R*. R package version 0.4.0, <https://CRAN.R-project.org/package=text2vec>.
- Sherin, Bruce. 2013. "A Computational Study of Commonsense Science: An Exploration in the Automated Analysis of Clinical Interview Data." *Journal of the Learning Sciences* 22(4):600–38. <https://doi.org/10.1080/10508406.2013.836654>.
- Shore, Bradd. 1998. *Culture in Mind: Cognition, Culture, and the Problem of Meaning*. Oxford, England: Oxford University Press.
- Spears, Russell, S. Alexander Haslam, and Ruurd Jansen. 1999. "The Effect of Cognitive Load on Social Categorization in the Category Confusion Paradigm." *European Journal of Social Psychology* 29(5–6):621–39. [https://doi.org/10.1002/\(SICI\)1099-0992\(199908/09\)29:5/6%3C621::AID-EJSP969%3E3.0.CO;2-W](https://doi.org/10.1002/(SICI)1099-0992(199908/09)29:5/6%3C621::AID-EJSP969%3E3.0.CO;2-W).

- Sperber, D. 1996. *Explaining Culture: A Naturalistic Approach*. Hoboken, NJ: Wiley-Blackwell.
- Spillman, Lyn. 2016. "Culture's Coherence: How the Trees Compose the Woods." *American Sociologist* 47(4):416–29. <https://doi.org/10.1007/s12108-016-9311-3>.
- Squire, Larry R. 1992. "Declarative and Nondeclarative Memory: Multiple Brain Systems Supporting Learning and Memory." *Journal of Cognitive Neuroscience* 4(3):232–43. <https://doi.org/10.1162/jocn.1992.4.3.232>.
- Stoltz, Dustin S., and Marshall A. Taylor. 2019. "Concept Mover's Distance: Measuring Concept Engagement via Word Embeddings in Texts." *Journal of Computational Social Science* 2(2):293–313. <https://doi.org/10.1007/s42001-019-00048-6>.
- Strauss, Claudia, and Naomi Quinn. 1997. *A Cognitive Theory of Cultural Meaning*. Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/CB09781139167000>.
- Swidler, Ann. 2001. *Talk of Love: How Culture Matters*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226230665.001.0001>.
- Taylor, Marshall A., and Dustin S. Stoltz. 2020. "Integrating Semantic Directions with Concept Mover's Distance to Measure Binary Concept Engagement." *Journal of Computational Social Science*, first published on July 14, 2020 as <https://doi.org/10.1007/s42001-020-00075-8>.
- Taylor, Shelley E. 1981. "The Interface of Cognitive and Social Psychology." In *Cognition, Social Behavior, and the Environment*, volume 1, pp. 189–211. Hillsdale, NJ: L. Erlbaum.
- Utsumi, Akira. 2020. "Exploring What Is Encoded in Distributional Word Vectors: A Neurobiologically Motivated Analysis." *Cognitive Science* 44(6):e12844. <https://doi.org/10.1111/cogs.12844>.
- Vaisey, Stephen. 2009. "Motivation and Justification: A Dual-Process Model of Culture in Action." *American Journal of Sociology* 114(6):1675–715. <https://doi.org/10.1086/597179>.
- Weiss, Gerald. 1973. "A Scientific Concept of Culture." *American Anthropologist* 75(5):1376–413. <https://doi.org/10.1525/aa.1973.75.5.02a00130>.
- White, Leslie A. 1959. "The Concept of Culture." *American Anthropologist* 61(2):227–51. <https://doi.org/10.1525/aa.1959.61.2.02a00040>.
- Willekens, Mart, and Stijn Daenekindt. 2020. "Cultural Logics and Modes of Consumption: Unraveling the Multiplicity of Symbolic Distinctions among Concert Audiences." *Musicae Scientiae*, first published on March 31, 2020 as <https://doi.org/10.1177/1029864920908305>.
- Wood, Michael Lee, Dustin S. Stoltz, Justin Van Ness, and Marshall A. Taylor. 2018. "Schemas and Frames." *Sociological Theory* 36(3):244–61. <https://doi.org/10.1177/0735275118794981>.
- Wu, Angela Xiao. 2014. "Ideological Polarization over a China-as-Superpower Mind-set: An Exploratory Charting of Belief Systems among Chinese Internet Users, 2008–2011." *International Journal of Communication* 8:2243–72.
- Zerubavel, Eviatar. 2009. *Social Mindscapes: An Invitation to Cognitive Sociology*. Cambridge, MA: Harvard University Press.

Acknowledgments: A replication repository for this article can be found at: <https://github.com/Marshall-Soc/CoCA>. We thank Jesper Sørensen, the deputy editor, and the consulting editor for their thoughtful comments on this article.

Marshall A. Taylor: Department of Sociology, New Mexico State University.
E-mail: mtaylor2@nmsu.edu.

Dustin S. Stoltz: Department of Sociology and Anthropology, Lehigh University.
E-mail: dss219@lehigh.edu.