Thinking through the past, present, and future: Language convergence-entropy is influenced by when you think of and how you feel

Constance M. Bainbridge
(cbainbridge@ucla.edu)
Department of Communication
University of California, Los Angeles

Zachary P. Rosen
Department of Communication
University of California, Los Angeles

Rick Dale
Department of Communication
University of California, Los Angeles

Abstract

Under construal-level theory, psychologically distant concepts such as a far away land are generally seen more abstractly compared to concepts viewed as closer in time, space, or identity. As we mentally travel through time, dynamics of language contained in streams-of-consciousness may provide a look into how we drift through topical space. Here, we investigated self-convergence and entropy using a BERT-based method to see how language drift over the course of typed streams-of-consciousness may be shaped by temporal framing. We applied this method to a dataset where undergraduate students during COVID-19 shared their thoughts imagining life before, during, and after the pandemic. We find that post-pandemic, future-directed thoughts showcase greater drift compared to past and present thoughts, suggesting greater exploration. Interestingly, past thoughts showed the least drift, suggesting there may be differences in concreteness depending on the direction in time you travel and the ability to have impact over temporally-tethered events.

Keywords: Thoughts, stream-of-consciousness, natural language processing, entropy, construal-level theory

Introduction

The human mind, empowered by active imagination, is able to travel through time. We recall past selves and generate future models of who we could become living whatever kind of life. One tool by which we do this is through language, which may be applied in a self-communicative manner, such as through journaling or thinking in a stream-of-consciousness style. Psychological distance, or how proximal or distal something is experienced in time, space, or perspective–such as similarity or relatability to another individual–is proposed to relate to the abstractness or concreteness of associated thoughts. For example, the here and now may be spoken of and perceived as more concrete compared to once-upon-a-time. This concept, "construal-level theory," suggests thinking about the future would likely explore more possibilities than when thinking about the present (Trope and Liberman, 2010).

Psychological distance in language can present in various forms, such as the perspective taken. For example, talking about the self from a third-person perspective can lead to the use of more abstract language, and can in turn more strongly influence self-conceptualization (Gainsburg and Kross, 2020). Some of the abstraction that results from distanced language use can lead to greater rationality in a dictator game (Gainsburg et al., 2022) or emotional regulation in reaction to a range of topics (Orvell et al., 2021). Under construal-level theory, this is perhaps achieved by detaching the individual from an overly rigid self-perception, allowing for more consideration of alternate possibilities.

Previous work has suggested that the concreteness of the past versus the future may be near mirror images of each other, with similar levels at identical temporal distances such as "last week" versus "next week" (Snefjella and Kuperman, 2015). Consistent with construal-level theory, the level of concreteness is higher as time becomes closer, regardless of the direction in time. While this work used concreteness ratings of individual words (Brysbaert, Warriner, and Kuperman, 2014), how language shifts across a text also may provide insights into how we explore possibilities differently as a result of psychological distancing.

Here, we examine stream-of-consciousness style texts written about life before, during, and after COVID-19 at a time when undergraduate students were engaged in online classes due to the pandemic. We use a convergence-entropy measure (Rosen and Dale, 2023) based on Bidrectional Encoder Representations from Transformers (BERT), a transformer model, to explore how predictable text is across each document, based on an initial segment (e.g., at segment 1, how predictable is the text in $1 + n$?). We find the greatest amount of entropy for thoughts imagining a post-pandemic future, compared to when writing about life during or before the pandemic. This suggests that topical space may be more exploratory for thinking about the future.

Methods

Analyzing lexico-semantic similarity

Given the subtlety of the phenomenon, a number of mechanisms to measure convergence have been proposed over the years (Srivastava et al., 2024). Convergence-entropy as first described in Rosen and Dale (2023) is a relatively context-agnostic measurement of how much information there is between sets of utterances, and has been shown to be sensitive to a number of social and cognitive factors. Intuitively, convergence-entropy replicates the conditions of the classic Shannon experiment, wherein an observer, having overheard some utterance from a particular speaker, is tasked to analyze "the linguistic output of a speaker and [give] a 'yes' or 'no' answer to the question of whether what they've said is conceptually similar between them and some other set of speakers," (Rosen and Dale, 2023). The resulting entropy measure is thus a good representation of how much variation

(a) A text is broken up into windows at time t=0 and several other windows of text at time $t + k$

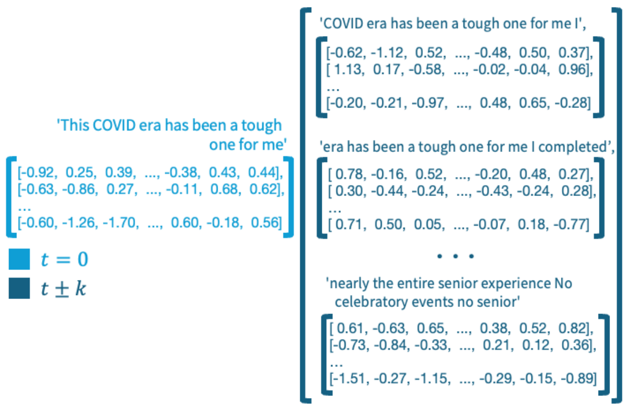

(b) Text for windows at time t=0 and all windows at time $t + k$ are converted to word vectors using BERT

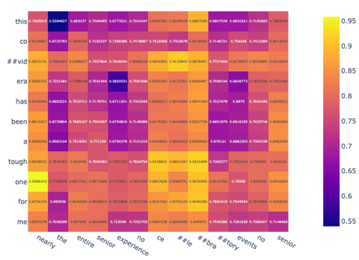

(c) For each comparison of window t=0 to each window $t + k$, the Cosine Error is calculated for the raw similarity between every token in t=0 and $t + k$

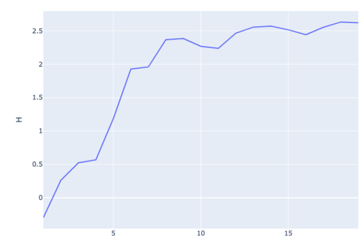

(d) Using the minimum Cosine Error values in each row (i.e. $\min_j$), we calculate the total entropy via $P_{N_{[0,\infty]}}(\mu = 0, \sigma)$.

Figure 1: Breakdown of the process used to calculate convergence-entropy

or "choice" was exerted by authors in the lexico-semantic content of their utterances at different points in time.

Methodologically, convergence-entropy leverages a transformer language model to generate word vectors for each token (in this case, an individual word) in an utterance, and then uses pairwise comparisons to calculate how much novel information exists in one utterance when compared to another. Some utterance x is converted to a set of word vectors–one vector per token i in x (E_{xi}). That process is repeated for every token j in an utterance y (E_{yj}). The total entropy for the utterance x when given an utterance y is then calculated as follows:

$$H[x|y] = -\sum_i P(E_{xi}|E_y) \log P(E_{xi}|E_y) \tag{1}$$

where the probability of the i^{th} token in the utterance x upon having read the utterance y is found based on the lowest Cosine Error (CoE) between E_{xi} and the vector for some token E_{yj}. A Half-Gaussian prior with $\mu = 0$ and arbitrary scale σ is used to convert CoE values to probabilities based on the logic that if the CoE value for the comparison between vectors is 0, then the two word vectors would be in the same location in vector space and thus they would effectively refer to the same, precise meaning.

$$P(E_{xi}|E_y) = P_{\mathcal{N}_{[0,\infty]}}\left(\min_j CoE(E_{xi}, E_{yj})\middle|\mu = 0, \sigma\right) \tag{2}$$

Practically, this value is only minimized in a transformer model when the contexts surrounding two tokens are sufficiently similar to one another, thus indicating that there is significant lexico-semantic similarity between the token i and at least one token j in the utterance y. In their paper, Rosen and Dale utilize the BERT transformer language model–among the first language transformer models that's been shown to be particularly sensitive to differences in context across uses of the same lexical items (Devlin et al., 2019). Like most transformer language models, BERT represents lexical items via the weighted sum of the word vectors collocated with a particular contextual item. The result is a representation per token that is quite sensitive to differences in context, and thus to subtle changes in lexico-semantic meaning (Devlin et al., 2019; Song et al., 2020; Utsumi, 2020).

Convergence-Entropy treats every lexical unit in an utterance as being a draw from a Bernoulli distribution with some probability p of observing the idea being referenced by a word as coming from some other distribution. In the current implementation, that means that we treat each word in some window at time $t = 0$ as being Bernoulli distributed with some probability p conditioned on each of the subsequent sliding windows, thus $P(w_i \in x \mid y)$. Each new window y occurs at time $(t + k) > t(x)$ is treated as its own unique distribution. Superficially this measurement resembles Kullback-Leibler divergence (Kullback & Leibler, 1951) in notation, but each lexical unit is actually independently conditioned on the distribution of lexical and conceptual units contained in the utterance y. And because of the emphasis on the similarity between semantic meaning of the various lexical units in the utterances x and y, it more directly tests whether the concepts underlying the two utterances match one another than Uniform Information Density (Jaeger, 2010), which tracks the predictability of the surface forms of words based on normative corpus statistics.

The data

The analyses use data collected around the COVID-19 pandemic, spanning around 2021-2022. Undergraduate students were asked to express their thoughts in typed streams-of-consciousness in an online study, instructed to let these thoughts flow spontaneously. First, the participants were prompted to describe life either before or during the pandemic. All participants then expressed their thoughts imagining life after the pandemic. Finally, whatever temporal frame remained (pre-pandemic or during pandemic) was given, resulting in three documents per participant, each constructed over ten minutes. These streams-of-consciousness are intended to be spontaneous, and are not reviewed by the participants after completion. Afterwards, participants received a series of questions, which included a rumination scale (Treynor, Gonzalez, and Nolen-Hoeksema 2003) and a social connectedness and assurance scale (Lee and Robbins, 1995). Previous exploratory findings from this dataset can be seen in Bainbridge and Dale (2023).

Analyzing temporal dynamics

In order to assess lexico-semantic differences in ideas across the duration of a writing prompt, we split each transcript into a series of segments. Each of these windows was then compared to 10-word windows forward in time in a sliding window fashion from 1 to 40 windows ahead. This represents a sliding window analysis, comparing the next 10 words (the context) to the original, target 10 words. We do not implement any padding in this analysis. We initially chose a 10-word window because it likely captures entire sentential patterns in this range, thus pairwise convergence measures capture comparisons of full idea units across time. Future investigation should test the effect of this parameter, especially because varying window sizes may subtly reflect levels of linguistic and semantic structure (Moscoso del Prado Martín et al., 2011).

We then calculate the convergence-entropy for the target from each window to assess the degree to which the lexico-semantic content drifts with a given response. If convergence-entropy rises rapidly across these 40 windows, it represents "semantic drift" in someone's stream-of-consciousness task. On the other hand, if the convergence-entropy rises more slowly, it may represent a kind of "semantic trapping" as the participant stays near to the original i^{th} context.

After generating these pairwise comparisons across $n = 89$ for 2,106,074 total comparisons across these windows. In this large dataset, we have variables relating to temporal framing, temporal disparity of windows, and the convergence-entropy measure.

Study 1: Temporal framing

Model

We deployed an Ordinary Least Squares regression model to test the degree to which the following variables predicted differences in convergence-entropy measurements: past (writing about pre-pandemic life), present (writing about life during the pandemic), future (about their expectations for what the future might hold), and the distance between the target and context window.

Results

The range of convergence-entropy for the future condition significantly differs from the past condition ($\beta = -0.223$; $t(2106058) = -3.048$, $p < .005$), with a slight trend compared to the present condition ($\beta = -0.142$; $t(2106058) = -1.941$, $p = .053$). The difference over time in convergence-entropy can be seen in Figure 1.

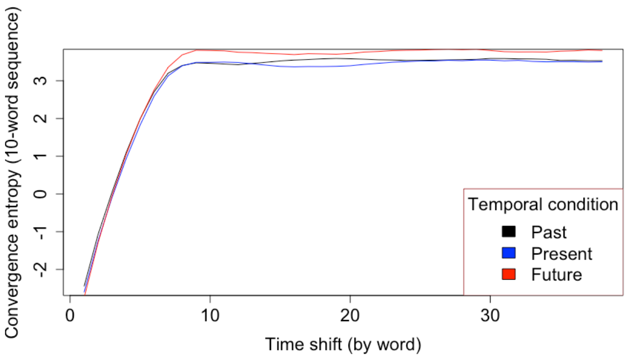

Figure 2: The texts written about life in a post-pandemic future showed the greatest change in convergence-entropy.

The increase in convergence-entropy is most pronounced within the first ten words for each condition, before mostly plateauing. The future condition led to a greater spike here, reaching greater entropy than the other two conditions reach at any point. This convergence-entropy highlights that future-oriented texts drift more as they progress, exploring more topical space.

In the original data, participants always started by either discussing life before or during the pandemic, followed by the future (post-pandemic). To test potential effects of ordering (i.e., does the time period you think about first bleed into what time you think of next?), we used the following equation for the subset of the data featuring future-oriented texts only:

*Convergence entropy ~ # of 10-word windows * ordering*

Here we find significant effects across the board: the placement of the sliding window ($\beta = 0.064$; $t(2106058) = 323.193$, $p < 0.001$), the ordering of conditions ($\beta = 0.022$; $t(2106058) = 6.775$, $p < .001$), and the interaction ($\beta = 1.31e\text{-}3$; $t(2106058) = 4.528$, $p < .001$) are all significant. When participants were primed by first writing about the present, they drifted more in their future-oriented texts compared to when first primed with the past.

Study 2: Social connectedness, assurance, and rumination

Model

We modeled the degree to which several psychometric measurements might affect lexico-semantic drift per each condition. To do so, we deployed a linear effects (LME) regression model to test the degree to which a number of psychometric measurements, per condition, predicted changes in convergence entropy. We assumed that each condition (past, present, future) was a categorical variable. Measurements for social assurance (social_A), social connectedness (social_C), rumination, and distance from the target were then used, within each condition, to predict changes in convergence-entropy. For reference, our LME equation is as follows:

Convergence entropy ~ [(condition)/(social_A + social_C + rumination + # of 10-word windows)] + ordering

We include random effects for each target window.

We then used a contrast-based F-test to assess whether the parameter estimates for social assurance, Social connectedness and rumination in each condition (i.e. past vs. present, past vs. future, or present vs. future) were statistically significantly different from one another in the final model.

We further tested whether there were statistically significant regularities in the parameter estimates for each psychometric measure across conditions (i.e. if their combined effect across all conditions was statistically significant from the null hypothesis–that their combined effect is equivalent to 0). To perform this test, we used a Wald F-test to assess the combination of psychometric parameter estimates across all conditions.

Results

We summarize a few results of the model here. First, across all conditions, rumination was predictive of larger convergence-entropy values (i.e. rumination contributed to greater lexico-semantic drift). Consistent with prior results, both condition past and present were predictive of lower convergence-entropy than the future condition–($\beta = -0.048$; $t(2106058) = -4.01$, $p < .001$) and ($\beta = -0.111$; $t(2106058) = -9.47$, $p < .001$) respectively. Social assurance in the present was predictive of increased convergence-entropy ($\beta = 9.3e\text{-}04$; $t(2106058) = 3.0$, $p = 0.003$). Social connectedness in both the past ($\beta = -0.002$; $t(2106058) = -7.74$, $p < .001$) and present ($\beta = 0.002$; $t(2106058) = 8.94$, $p < .001$) were predictive of differences in convergence-entropy, though in opposite directions. Rumination was predictive of increased convergence entropy in all conditions–past: ($\beta = 0.002$; $t(2106058) = 10.4$, $p < .001$), present: ($\beta = 5.35e\text{-}04$; $t(2106058) = 3.58$, $p < .001$), future: ($\beta = 5.74e\text{-}04$; $t(2106058) = 3.76$, $p < .001$). Distance between the target and the window for comparison was predictive of increased convergence entropy in all conditions, as expected based on the prior results–past: ($\beta = 0.061$; $t(2106058) = 453.0$, $p < .001$), present: ($\beta = 0.063$; $t(2106058) = 504.0$, $p < .001$), future: ($\beta = 0.064$; $t(2106058) = 478.0$, $p < .001$).

Comparisons for each psychometric value across conditions was considered significant only if the p-value for the comparison was less than .016 (Bonferroni correction–3 temporal conditions tested per psychometric value). Differences in estimates for social assurance in the past versus the present condition were statistically significant from one another ($F(1,2106058) = 29.142$, $p < .001$). As was the difference in estimates for social assurance for the present versus the future ($F(1,2106058) = 15.45$, $p < .001$). In all three comparisons for social connectedness, the differences between parameter estimates were statistically significant–past vs. present: ($F(1,2106058) = 468.384$, $p < .001$), past vs. future: ($F(1,2106058) = 96.741$, $p < .001$), present vs. future: ($F(1,2106058) = 142.476$, $p < .001$). Differences in parameter estimates for past versus present ($F(1,2106058) = 78.016$, $p < .001$) and past versus future ($F(1,2106058) = 72.758$, $p < .001$) were statistically significant. Despite seeming quite similar to one another, differences between parameter estimates for distance between windows in the past vs. present conditions ($F(1,2106058) = 55.754$, $p < .001$), past vs. future conditions ($F(1,2106058) = 219.773$, $p < .001$), and present vs. future conditions ($F(1,2106058) = 62.999$, $p < .001$) were all statistically significant.

For the aggregate psychometric values across conditions, significant trends only existed for rumination (harmonic mean = 7.07e-3; $F(1,2106058) = 44.501$, $p < .001$) and distance between the target and window (harmonic mean = .0627; $F(1,2106058) = 683279.259$, $p < .001$). To save space, our full model results can be found at https://bit.ly/COVIDRef-CogSci25.

Discussion

When temporally oriented around the global event of the COVID-19 pandemic, undergraduate students' typed streams-of-consciousness drift through topical space differently depending on what time period they think about. Thinking about a post-pandemic future brings greater drift or convergence-entropy as their thoughts flow compared to thinking about life before or during the pandemic. This drift is especially pronounced when thinking first about the present (during the pandemic), compared to the past (before the pandemic). This may suggest fluidity to construal-level theory–psychological distance may be modulated based on prior distances, or construals may become restricted by previous topical spaces.

Additionally, individual differences such as social connectedness, social assurance, and rumination tendencies may add complexity to the ways we drift through thoughts. Interestingly, while rumination is associated with recurrent, repetitive thoughts, we found an increase in convergence-entropy for higher levels of rumination. Social connectedness and assurance may predict some aspects of drift, though there is no overall trend in the aggregate across the temporal conditions. As the social connectedness and assurance scales may represent temporary states (and especially in this case, be specific to the time participants completed the task), we may instead expect trait-based individual differences to cause more distinctions in any or all conditions.

Part of the complexity inherent to those same individual differences lies in their interaction with the different timeframes participants described. While this is surprising, it is not unprecedented in other domains of research. Prior research studying the use of various temporal metaphors in English, for example, show that preferences in the use of the nascent TIME IS A MOVING OBJECT versus TIME IS A LOCATION metaphor depends on the psychological associations of the event being described (Margolies and Crawford, 2008)–people tend to prefer to use the TIME IS A MOVING OBJECT metaphor when events are negatively valenced, in part out of a hope that negative events will "pass" quickly (Lee and Ji, 2014). And for strongly negative events, like periods of grief, the selection of the TIME IS A MOVING OBJECT metaphor is correlated with longer predicted durations of said negative event (Ruscher, 2011). Such work, however, has not traditionally posited what psychological qualia would likely yield those preferential differences. While different lexicalization strategies are correlated with differences in other behaviors, a chicken-and-egg problem remains of what comes first–psychological qualia or selection of a particular lexico-semantic construct? The current study shows that individuals' psychological traits influence lexico-semantic drift throughout the course of a given writing task in a way that dovetails with other, extant research programs assessing peoples' views on temporally located events. And while the locus of analysis is quite different, our results hint at the causal relationship underlying prior observations: pre-existing psychological traits affect how we interact with various descriptions of events in the past, present, and future.

This prior research might also help us to unpack the unlikely positive relationship of rumination to convergence-entropy. While individuals may exhibit high rumination, that does not mean that they want the events that they ruminate on to linger. Additionally, in the present study a sliding window of 10 tokens was selected–however, rumination or other distinct features of an individual may present on different scales that may eventually be explored systematically in greater depth.

While this study takes an initial step towards understanding the dynamics of topic exploration over time when positioned in the past, present, or future, there remain details worth including in future designs using this methodology. First, greater control of the experience of time may shift how the future versus past appear: this sample was composed of undergraduate students located in the United States, who likely perceive the future as a longer time frame compared to the past. Extending across age groupings, and including measures of what exactly is the future, past, and present may reveal important differences. It is also worth noting that how far participants projected themselves into the past and future may have differed across the sample. Individual differences such as socio-economic status may also play a mediating or moderating role, as is suggested by life-history theory accounts (Griskevicius et al., 2011).

Future work should explore how convergence-entropy methods may operate differently when controlling for different temporal distances, regardless of whether projecting into the past or future. Similar to Snefjella and Kuperman (2015), how might yesterday or tomorrow differ from traveling to the scale of years, decades, or beyond? How the topic interacts with convergence-entropy may also be revealing - while the prompt for the present dataset frames specifically around a major event, other framings or an entire lack of one may encourage more or less entropy. Possibilities for outcome-based work, such as in mental health interventions and detection of mental wellbeing, may then be ripe for future work. For example, writing with more or less psychological distancing may benefit individuals differently based on their attachment styles (Wang, Lin, Huang, and Yeh, 2012), suggesting encouragement of specific temporal framings or levels of convergence-entropy may hold applied potential. Exploring spoken versus typed or written thoughts may also reveal differing benefits based on potentially distinct convergence-entropy patterns across these modalities.

References

Bainbridge, C. M., & Dale, R. (2023). Thinking about life in COVID-19: An exploratory study on the influence of temporal framing on streams-of-consciousness. *PLOS ONE*, 18(4), e0285200.

Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)* (pp. 4171-4186).

Gainsburg, I., & Kross, E. (2020). Distanced self-talk changes how people conceptualize the self. *Journal of Experimental Social Psychology*, 88, 103969.

Gainsburg, I., Sowden, W. J., Drake, B., Herold, W., & Kross, E. (2022). Distanced self-talk increases rational self-interest. *Scientific Reports*, 12(1), 1-8.

Griskevicius, V., Tybur, J. M., Delton, A. W., & Robertson, T. E. (2011). The influence of mortality and socioeconomic status on risk and delayed rewards: a life history theory approach. *Journal of Personality and Social Psychology*, 100(6), 1015.

Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, 61(1), 23–62. https://doi.org/10.1016/j.cogpsych.2010.02.002

Kullback, S., & Leibler, R. A. (1951). On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1), 79–86. https://doi.org/10.1214/aoms/1177729694

Lee, A., & Ji, L.-J. (2014). Moving away from a bad past and toward a good future: Feelings influence the metaphorical understanding of time. *Journal of Experimental Psychology: General*, 143(1), 21–26. https://doi.org/10.1037/a0032233

Lee, R. M., & Robbins, S. B. (1995). Measuring belongingness: The social connectedness and the social assurance scales. *Journal of Counseling Psychology*, 42(2), 232.

Margolies, S. O., & Crawford, L. E. (2008). Event valence and spatial metaphors of time. *Cognition and Emotion*, 22(7), 1401–1414. https://doi.org/10.1080/02699930701810335

Moscoso del Prado Martín, F. (2011). The universal "shape" of human languages: spectral analysis beyond speech. *Nature Proceedings*. https://doi.org/10.1038/npre.2011.6097.2

Orvell, A., Vickers, B. D., Drake, B., Verduyn, P., Ayduk, O., Moser, J., ... & Kross, E. (2021). Does distanced self-talk facilitate emotion regulation across a range of emotionally intense experiences?. *Clinical Psychological Science*, 9(1), 68-78.

Rosen, Z. P., & Dale, R. (2023). BERTs of a feather: Studying inter- and intra-group communication via information theory and language models. *Behavior Research Methods*. https://doi.org/10.3758/s13428-023-02267-2

Ruscher, J. B. (2011). Moving forward: The effect of spatiotemporal metaphors on perceptions about grief. *Social Psychology*, 42(3), 225–230. https://doi.org/10.1027/1864-9335/a000066

Snefjella, B., & Kuperman, V. (2015). Concreteness and psychological distance in natural language use. *Psychological Science*, 26(9), 1449-1460.

Song, Y., Wang, J., Liang, Z., Liu, Z., & Jiang, T. (2020). Utilizing BERT intermediate layers for aspect based sentiment analysis and natural language inference. *arXiv preprint arXiv:2002.04815*.

Srivastava, S., Wentzel, S. D., Catala, A., & Theune, M. (2024). Measuring and implementing lexical alignment: A systematic literature review. *Computer Speech & Language*, 101731. https://doi.org/10.1016/j.csl.2024. 101731arXiv:2002.04815 [cs]. Retrieved December 21, 2021, from http://arxiv.org/abs/2002.04815

Treynor, W., Gonzalez, R., & Nolen-Hoeksema, S. (2003). Rumination reconsidered: A psychometric analysis. *Cognitive Therapy and Research*, 27, 247-259.

Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440.

Utsumi, A. (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

Wang, Y.-L., Lin, Y.-C., Huang, C.-L., & Yeh, K.-H. (2012). Benefitting from a different perspective: The effect of a complementary matching of psychological distance and habitual perspective on emotion regulation. *Asian Journal of Social Psychology*, 15(3), 198–207.