Event construal and temporal distance in natural language

Sudeep Bhatia [a,*], Lukasz Walasek [b,*]

[a] Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA
[b] Department of Psychology and Behavioral Science Group at the Warwick Business School, University of Warwick, Coventry, United Kingdom

ABSTRACT

Construal level theory proposes that events that are temporally proximate are represented more concretely than events that are temporally distant. We tested this prediction using two large natural language text corpora. In study 1 we examined posts on Twitter that referenced the future, and found that tweets mentioning temporally proximate dates used more concrete words than those mentioning distant dates. In study 2 we obtained all New York Times articles that referenced U.S. presidential elections between 1987 and 2007. We found that the concreteness of the words in these articles increased with the temporal proximity to their corresponding election. Additionally the reduction in concreteness after the election was much greater than the increase in concreteness leading up to the election, though both changes in concreteness were well described by an exponential function. We replicated this finding with New York Times articles referencing US public holidays. Overall, our results provide strong support for the predictions of construal level theory, and additionally illustrate how large natural language datasets can be used to inform psychological theory.

© 2016 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

With the ongoing digitization of information on the internet, it is now possible to access a large variety of natural language datasets. For a psychologist, these data offer an unprecedented gateway to study the formation of beliefs and attitudes, the dynamics of interpersonal relationships, and the preferences and behaviors of decision makers (Boyd & Crawford, 2012). Although still scarcely used in psychological science, these data have a potential to provide a unique and statistically powerful approach to evaluating models of human cognition, including models that are not specifically about language use.

In the following paper, we demonstrate how large corpora of natural language can be used to study the relationship between temporal distance and mental representations of events in the world. More specifically, we test the core assumption of the Construal Level Theory (CLT, Trope & Liberman, 2010) that the mental representation of a given object or event becomes less concrete and more abstract with increasing temporal distance. In two studies, we show that this hypothesis can be tested using online communication on Twitter and a large database of New York Times articles, by examining the concreteness of words used in these tweets and articles, and the distance between when they were written and when the events they pertain to occurred.

CLT has been used to describe the influence of psychological distance on the representations of physical objects (Liberman, Sagristano, & Trope, 2002), choice alternatives (Borovoi, Liberman, & Trope, 2010), events (Wakslak, Trope, Liberman, & Alony, 2006), consumer goods (Trope, Liberman, & Wakslak, 2007), actions (Liberman & Trope, 1998), and individuals (Rim, Uleman, & Trope, 2009). Psychological distance includes dimensions such as time, space, social distance and hypotheticality, although in the following paper we focus exclusively on the effect of temporal distance (Trope & Liberman, 2010). To illustrate, a person may describe a date as a "candlelight dinner" if the event occurred recently, but as a more abstract "romantic evening" if it occurred in the distant past. Higher level construals are not just impoverished versions of more proximal representations. Rather they involve different types of information, including information about the meaning of the object or action (Semin & Fiedler, 1988). According to CLT, the effect of psychological distance on our thought is functional as it allows us to plan for the future, learn from our mistakes or to communicate more efficiently.

The predictions of CLT find support in existing empirical studies (Soderberg, Callahan, Kochersberger, Amit, & Ledgerwood, 2015). The positive association between temporal distance and construal abstractness has been demonstrated in categorization tasks (Liberman et al., 2002), event descriptions (Liberman & Trope,

* Corresponding author.
  E-mail addresses: bhatiasu@sas.upenn.edu (S. Bhatia), L.Walasek@warwick.ac.uk (L. Walasek).

2008) and even at the level of visual perception (Förster, Friedman, & Liberman, 2004). Often simply imagining an event occurring in the distant future leads people to describe it using words with a high level of abstraction (Semin & Fiedler, 1988). Crucially, CLT has been shown to predict actual behavior, behavioral intentions, negotiation style, self-control, risk perception, and temporal discounting (see Trope et al., 2007 for a review). CLT thus attempts to provide a unifying framework for understanding the effect of distance on perceptual processes, social interactions, moral reasoning, consumer behavior, and even decision making under risk and uncertainty.

Most existing work on CLT uses laboratory experiments, in which psychological distance is induced or manipulated using hypothetical frames and primes (but see Magee, Milliken, & Lurie, 2010). For example, participants in these studies are often found to represent objects and events less concretely after being instructed to think of these items as distant (e.g., in spatial or temporal terms), once the information is presented as distant due to a specific reference to time or space (e.g., occurring in a distant future or happening far away), or when construal is induced by the question framing (e.g., thinking *why* rather than *how* an event occurred; Trope et al., 2007). While the exact methods of framing and priming vary greatly between the studies, many of these techniques require some form of artificial manipulation of psychological distance with reference to a given object or event. However, CLT also predicts that real-world psychological distance to an event should influence its representation in everyday thought and discourse. This is not easily tested in the lab.

In this paper we hope to test the predictions of CLT by analyzing the level of concreteness and abstraction of language in real-world settings. In study 1, we collected and analyzed millions of time-indexed posts on Twitter. Twitter is an excellent source of data – in 2015, it averaged 236 million active users (http://www.statista.com/statistics), who posted close to 500 million messages (tweets) per day. These tweets contain up to 140 characters and are shared among each user's social network (Reips & Garaizar, 2011). In this study we obtained a large number of tweets that referenced dates in the future, and were able to use these tweets to determine the concreteness of the language used to describe events at these dates. This allowed us to observe how psychological distance influences everyday discourse, and put the key assumptions of the CLT to a real-world test.

In study 2, we analyzed word concreteness in news articles using the New York Times (NYT) Annotated Corpus (Sandhaus, 2008). This corpus contains over 1.8 million NYT articles written between 1987 and 2007. Importantly for our purposes, these articles are tagged with keywords describing the topics of the articles. In this study we obtained all NYT articles written before and after the 1988, 1992, 1996, 2000, and 2004 US Presidential elections, which were tagged as pertaining to these elections. We subsequently tested how the concreteness of the words used in the articles varied as a function of temporal distance to the election they reference. We also performed this analysis with NYT articles referencing three popular public holidays. Unlike study 1 and prior work (such as Snefjella & Kuperman, 2015), study 2 allowed us to examine the influence of temporal distance in the past and in the future, while controlling for the exact time when specific events occurred. Recent findings show that future events are perceived as more proximal than equally distant past events (Caruso, Van Boven, Chin, & Ward, 2013). For example, events occurring 1 year in the future are rated as psychologically closer than events that occurred 1 year in the past. Our NYT dataset allows us to test whether this asymmetry is reflected in the abstractness with which various events are described.

2. Study 1: temporal distance and tweet concreteness

2.1. Methods

We first examined the effect of temporal distance on object representation by studying the relationship between the concreteness of the words used in tweets about objects and events at various points in the future. In particular, we obtained tweets that used the phrases "next week", "next month" or "next year", as well as tweets whose text mentioned the years "2015", "2016", "2017", and "2018". This was done using Twitter's data streaming feature, which allows researchers to download tweets as they are created. We filtered the Twitter data stream using the terms 'next week', 'next month', 'next year', '2015', '2016', '2017', and '2018' to obtain only the tweets that mentioned one of these phrases. We streamed Twitter over the course of one week in August 2014. The length of our data collection was set to one week as this provided enough time to obtain a very large number of tweets. Our collection was pruned to exclude retweets (that is, tweets that were copied and reposted).

We formalized word concreteness using a database of 40,000 English word ratings obtained by Brysbaert, Warriner, and Kuperman (2014). Brysbaert et al. collected these from over four thousand participants who were asked to rate different words on a 5-point scale based on how abstract or concrete the meanings of the words were to them. Using this database we scored each tweet on the average concreteness of its component words. The score for each tweet ranged from 1 for highly abstract to 5 for highly concrete. Tweets composed entirely of words absent from the Brysbaert et al. database were excluded from our dataset.

2.2. Results

Our final dataset for examining temporal distance consisted of 1,746,788 tweets that mentioned one of the three phrases or one of the four years of interest to us, and included at least one of the 40,000 words from the Brysbaert et al. database. The distribution of concreteness ratings for these tweets is displayed in the top panel of Fig. 1.

Construal level theory predicts that tweets referencing objects that are far away into the future should use relatively less concrete words than tweets referencing objects that are nearby in the future. Thus the average concreteness of tweets that mention "next week" should be higher than those that mention "next month", which in turn should be higher than those that mention "next year". Likewise the average concreteness of tweets that reference the years 2015–2018 should decrease with the year that they reference.

These predictions were supported by our data, which is summarized in Figs. 2 and 3, respectively. Particularly, tweets that mentioned "next week" had an average concreteness of 2.69 ($SD = 0.47$), tweets that mentioned "next month" had an average concreteness of 2.64 ($SD = 0.43$), and tweets that mentioned "next year" had an average concreteness of 2.48 ($SD = 0.39$). The differences between these three groups were statistically significant ($t = 36.65$, $p < 0.001$ for the difference between week and month, $t = 139.70$, $p < 0.001$ for the difference between month and year, and $t = 246.30$, $p < 0.001$ for the difference between week and year).

We obtained similar results for tweets that referenced the years 2015–2018. Particularly, tweets that referenced 2015 had an average concreteness of 2.79 ($SD = 0.58$), tweets that referenced 2016 had an average concreteness of 2.72 ($SD = 0.51$), tweets that referenced 2017 had an average concreteness of 2.70 ($SD = 0.46$), and tweets that referenced 2018 had an average concreteness of 2.61 ($SD = 0.40$). Overall, the concreteness of a tweet was lower if it referenced a year that was further into the future. Formally, we tested this using a linear regression of the effect of year on tweet
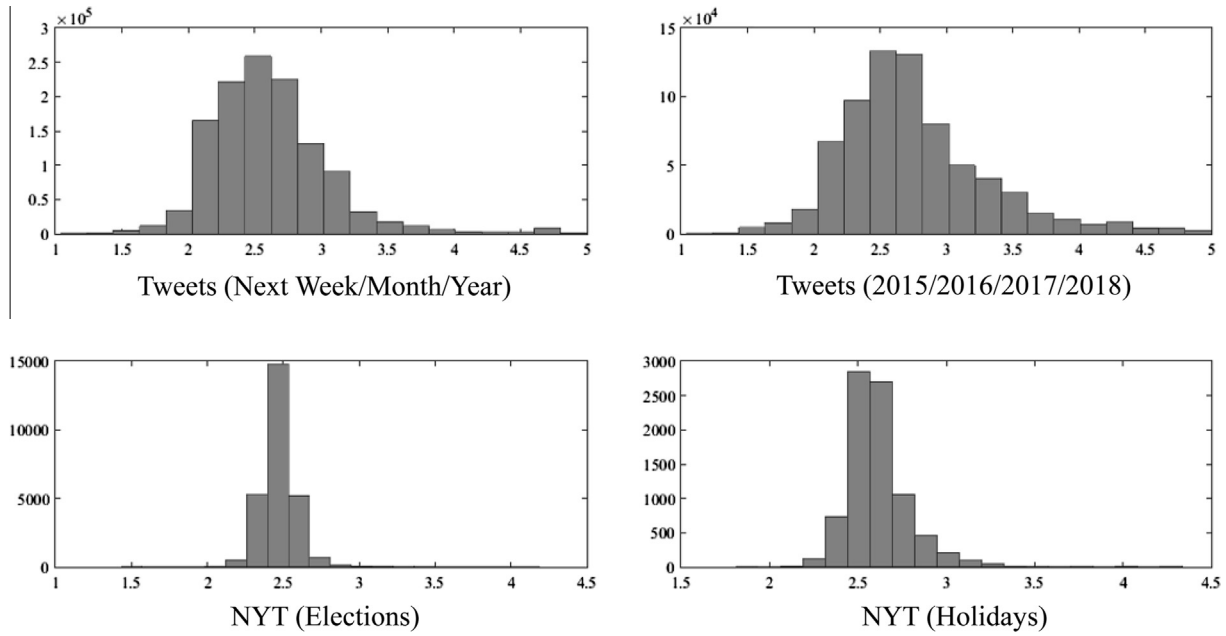
Fig. 1. Frequency distribution of concreteness for tweets in study 1 (top panel) and NYT articles in study 2 (bottom panel). The tweet distributions have increased variability due to the fewer number of words in tweets compared to news articles.
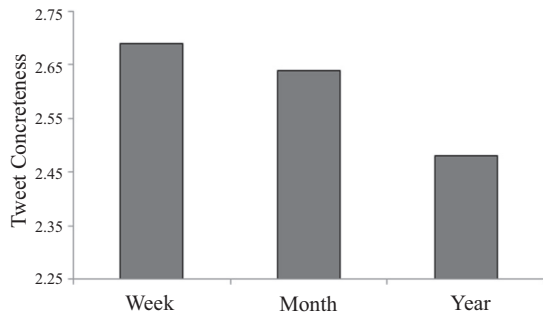

Fig. 2. Average tweet concreteness for tweets that mention "next week", "next month" and "next year" respectively. Note that standard error bars are omitted from the graph as they are too small to be visible ($SE < 0.01$ for all groups).
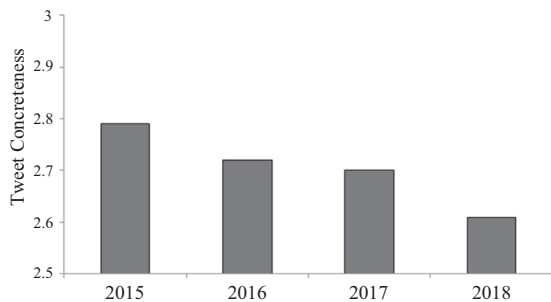

Fig. 3. Average tweet concreteness for tweets that mention "2015", "2016", "2017" and "2018" respectively. Note that standard error bars are omitted from the graph as they are too small to be visible ($SE < 0.01$ for all groups).

concreteness, which revealed a highly significant negative effect of year ($\beta = -0.06$, $t = -87.55$, $p < 0.001$). A more stringent statistical test which examined differences between pairs of years individually similarly revealed significant differences between all pairs of years ($p < 0.001$ for all pairs).

2.3. Discussion

The results of study 1 are consistent with the predictions of the CLT: As the temporal distance increases, people's communication uses higher-feature, more abstract and less concrete words. They are also consistent with closely related recent findings regarding construal and temporal distance in online forums (Snefjella & Kuperman, 2015; study 2). That said, study 1 is limited in a number of important ways. First, it is impossible to establish what kind of events people post about when they refer to future dates. Thus one cannot rule out the possibility that the types of events referenced in tweets about the distant future are different to those referenced in tweets about the proximate future, and that the results of our study are purely due to different mental representations for different events rather than different mental representations for similar events at different points in time. This is also a limitation of Snefjella and Kuperman's (2015) study on construal and temporal distance.

Secondly, our results are mute about the evolution of concreteness over time, with respect to a certain fixed event. Indeed, this is also a feature of Snefjella and Kuperman's (2015) study, as well as most existing laboratory research on CLT. This work typically does not track mental representations over long stretches of time, and thus cannot establish how representations for a particular event that happens at a particular point of time change over the course of many weeks, months, or even years. Likewise, existing evidence cannot demonstrate how these changes vary as a function of whether the event has occurred in the past or is yet to occur in the future. We address these concerns in study 2.

3. Study 2: concreteness and temporal distance in New York Times articles

3.1. Methods

In this study we used the New York Times Annotated Corpus (Sandhaus, 2008) to examine changes in the construal of event representations over time. This corpus contains over 1.8 million articles written and published by the New York Times between

January 1, 1987 and June 19, 2007. Additionally, the majority of these articles are manually tagged by library scientists or tagged using computer algorithms (with tags verified by the NYT online production staff). These tags are drawn from a vocabulary of people, organizations, locations and topic descriptors, and provide a simple but rigorous way to determine the topics that the articles refer to.

The NYT Annotated Corpus is especially valuable for our purposes due to its broad topical and temporal scope. Using the same methods as in study 1, it is possible to obtain a concreteness rating for each article in the corpus. Combining these data with the information about the time of occurrence for any given event, we can directly assess the temporal dynamics of the concreteness ratings. Particularly, we can examine how the concreteness changes over time, paying special attention to the temporal distance between the date the article was written and the time of the event it refers to. This allows us to study the effect of temporal distance on construal keeping the event in consideration constant, and thus controlling for event-specific confounds that could be affecting observed changes in concreteness in our first study. Additionally, for events that are anticipated (i.e., are covered in the articles preceding the events), we can examine articles' concreteness both before and after the events happen, providing a richer understanding of the relationship between time and event construal.

In this study we examined two different sets of events. Our first set of events involved presidential elections in the United States. Conveniently for our purposes, the NYT corpus has tags for each separate presidential election, so that we were able to recover all articles pertaining to a presidential election in a given year. In our analysis, we considered the presidential elections of the years 1988, 1992, 1996, 2000, and 2004, which were the elections within the time period covered by our corpus, and we obtained articles referencing these individual elections written before, during, and after each of the elections. Note that some articles referencing a particular election were actually written around the time of the subsequent election, and this made it difficult for us to rigorously determine which of the elections the articles were primarily about. To avoid issues with classifying these ambiguous articles, we limited our analysis to consider only articles written within two years of the election their tag referenced. Thus, for example, an article that was tagged as referencing the 1996 election was used in our analysis only if the article was published between November 1994 and November 1998, and thus was closer to the 1996 election than any other election.

Our second set of events involved popular holidays in the United States. We considered three popular holidays: July 4 (Independence Day), December 25 (Christmas), and January 1 (New Year's Day). Each of these holidays have tags in the NYT corpus, and we were able to obtain all articles that referred to these three holidays between 1987 and 2007. Again for all three holidays we had articles that were written before and after the dates of the actual holidays. We categorized an article as being about a particular holiday in a particular year, if it was closer to the holiday in that year compared to the previous or subsequent year. Thus an article about Christmas written in February of 1990 would be categorized as referring to Christmas in 1989 rather than Christmas in 1990, and subsequently being two months after rather than ten months before the event it refers to. For robustness, we also used a 3-month window to classify each holiday.

3.2. Results

3.2.1. Overview of data

There were a total of 26,873 NYT articles written about the presidential elections in our corpus, out of which 25,461 articles were published within two years of the elections they referenced.

These 25,461 articles were used in our analysis. The average distance between the dates of the articles and the dates of the elections they referenced was −153.02 days, (median = −114, SD = 198.31), implying that the bulk of the articles about an election were written before the election. However, the day with the most number of articles pertaining to the election was one day after the election. The average concreteness of the articles about the elections was 2.47 (SD = 0.12).

Likewise there were a total of 8,316 NYT articles written about the three holidays we consider. The average distance between the dates of the articles and the dates of the holidays was −10.20 (median = −5, SD = 25.16), and the day with the most number of articles pertaining to the holiday was the day of the holiday itself. The average concreteness of the articles about these dates was 2.61 (SD = 0.17). The distribution of concreteness ratings for both sets of articles is displayed in the bottom panel of Fig. 1.

3.2.2. Concreteness and temporal distance

Recall that for each of the elections in our corpus and for each of the different holidays in our corpus, we have both the dates on which articles about these events were written, as well as the concreteness of these articles (based on the same methods as in study 1; Brysbaert et al., 2014). We can thus examine whether increasing temporal distance from the date of the election or holiday does indeed decrease the concreteness of the articles written about the election or holiday, that is, whether event representations become more abstract with distance, controlling for the event in consideration.

We first examined presidential elections, pooling our data for the five election years. We used a linear regression model, with the concreteness of the articles as the dependent variable, and the natural logarithm of the absolute distance in days between the date of the article and the date the election mentioned in the article as the predictor. Our use of a logarithm to transform absolute distance was motivated by the fact that distance appeared to be roughly exponentially distributed. We also included a variable capturing whether or not the article was written before the election or after the election. Finally, we allowed for indicator variables corresponding to the year of the election, to allow for heterogeneity in the concreteness of the articles across different elections.

Our linear regression showed a strongly negative effect of log-distance on article concreteness (β = −0.01, SE = 0.001, t = −9.06, p < 0.001). This regression also indicated that articles written before the election were typically more concrete than articles written after the election (β = 0.04, SE = 0.002, t = 16.82, p < 0.001). We repeated the above analysis using absolute distance instead of log-absolute distance, and once again found a significant negative effect of distance and a significant positive effect of articles being written before vs. after the election (p < 0.001 for both) indicating that the above results are not a product of how we transformed distance in our analysis.

We performed a similar analysis on news articles written about Independence Day, Christmas, and New Year's. Again, allowing for indicator variables corresponding to the holiday in question to allow for heterogeneity in the concreteness of the articles across different holidays, we found a significant negative effect of log-distance (β = −0.01, SE = 0.002, t = −3.41, p < 0.001), and a significant positive effect of an article having been written before the holiday (β = 0.04, SE = 0.005, t = 10.27, p < 0.001) on its concreteness. These effects were invariant to whether we considered absolute distance and not the log of absolute distance (p < 0.001 for both effects).

We computed rolling averages for our two datasets, in order to better visualize the effect of temporal distance on concreteness. Thus for our elections dataset, we first computed the average concreteness of articles written at days at different distances from

the elections (e.g. average concreteness of articles written on the day of the elections (distance = 0 days), one day after the elections (distance = 1 day), one day before the elections (distance = −1 day), and so on). Then for each of the days before and after the events, we computed the average concreteness of articles written in a 90 day window around that day, giving us a three month rolling average for concreteness as a function of distance. Rolling averages are useful to smooth otherwise noisy data, and are commonly used in time series datasets like ours. We use a similar method for the dataset of three public holidays, except for the fact that our windows for computing the rolling averages were only 21 days long, giving us a three week rolling average. The choice of the smaller window was motivated by the fact that articles for holidays in this dataset are written at most a half a year before or after the holiday, compared to articles in the elections dataset, which are written up to two years before or after the election day.

Figs. 4 and 5 plot the rolling averages for concreteness as a function of distance from the elections and from the holidays respectively. The 0-days on the *x*-axes correspond to the date that the elections were held, and the dates on which the holidays occurred. As expected, there is a negative relationship between distance and concreteness, so that articles written many days before or many days after the events are much less concrete than articles written around the time of the events, which is consistent with the predictions of the CLT. Secondly, as suggested by our regressions, articles written before the events are more concrete than articles written after the event. This is particularly pronounced for the elections dataset. Although this is not directly predicted by the CLT, it is not inconsistent with its main tenets.

Note that there is a mild increase in concreteness of the NYT articles written for the holidays roughly two months after the holidays. It is likely that this is due to the fact that the articles published in this time period refer to the future holiday rather than the previous holiday. As mentioned above, we had categorized articles as pertaining to a holiday in a year if they were published within a six-month window of that holiday (that is, if they were closer in distance to that holiday in that year compared to the holiday in the previous or subsequent year). If NYT articles about holidays are more likely to mention future holidays relative to past holidays, then this method could generate occasional incorrect categorizations. Due to the size of our dataset, it is difficult to manually avoid this. However, we have replicated all our analysis with
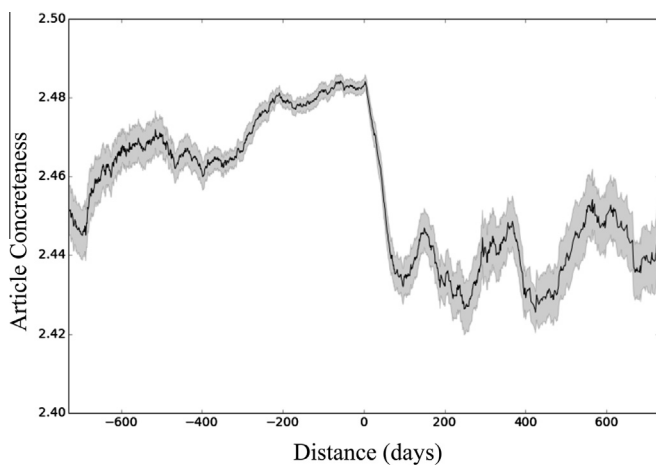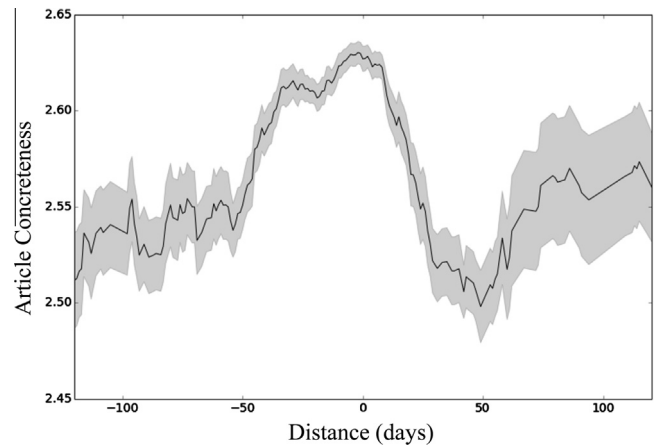

Fig. 5. A plot of the rolling average of article concreteness for various holidays, as a function of the distance between the date the articles were published and the dates of the holidays they correspond to. The rolling average uses a window of 21 days. Additionally the gray shaded areas correspond to ±one standard error involved in computing these rolling averages at each day.

3-month rather than 6-month time windows for categorizing articles, and all the results discussed in this paper hold (indeed, these results are slightly strengthened).

3.2.3. Model fitting

The above analysis shows that temporal distance has a negative effect on concreteness and that articles written before an event are more concrete than articles written after the event. In the remainder of this section, we wish to examine the functional forms which characterize this relationship. We use three common functional forms: the polynomial function (Eq. (1)), the power function (Eq. (2)), and the exponential function (Eq. (3)). These functions describe the relationship between the concreteness, *c*, of an article written about an event, and the absolute time difference, *t*, between the article and the event:

$$\text{Polynomial function}: \ c = \alpha + \beta \cdot t + \gamma \cdot t^2 \tag{1}$$
$$\text{Power function}: \ c = \alpha + \beta \cdot t^\gamma \tag{2}$$
$$\text{Exponential function}: \ \alpha + \beta \cdot e^{\gamma \cdot t} \tag{3}$$

All our fits involve three free parameters, α, β, and γ. In all three functions, α can be interpreted as an intercept term that shifts concreteness a number of units above zero, and β and γ can be seen as parameters that determine whether time has a positive or negative effect on concreteness. The specific interpretation of these parameters varies across functions, but in the case of the exponential function, γ corresponds to a discount rate. Exponential functions are commonly used to model time discounting in choice, a behavioral phenomenon that construal level theory has attempted to predict (Trope & Liberman, 2003). For this reason, this interpretation of γ will be useful in our analysis.

We performed model fits on rolling average data displayed in Figs. 4 and 5. In order to fit the three functional forms, we employed non-linear least squares and fitted the models separately for articles written before the events and articles written after the events. Overall, we found that the exponential model outperforms polynomial and power models for both datasets, both in terms of model R^2 and mean-squared error, when describing the concreteness of articles written after the events in question. In contrast, both the exponential and polynomial models fit equally well when describing the concreteness of articles written before the events in question. A summary of our model fits is provided in Table 1.

Fig. 4. A plot of the rolling average of article concreteness for various elections, as a function of the distance between the date the articles were published and the dates of the elections they reference. The rolling average uses a window of 90 days. Additionally the gray shaded areas correspond to ±one standard error involved in computing these rolling averages at each day.

Finally, let us examine the parameters of the exponential model, which is the best fitting overall model. As can be seen in Table 1, we find higher values for γ for articles written after the events compared to articles written before the events. Additionally γ is positive for all fits. These differences indicate that the exponential decrease in concreteness is larger after compared to before the events, corresponding to our earlier finding that concreteness is typically higher for articles written before the events.

3.3. Discussion

In study 2 we have examined how the concreteness of the words used in New York Times articles varies with the temporal distance between the dates these articles are published and the dates of the events they reference. Consistent with the results of study 1 and predictions of CLT we have found a negative relationship between concreteness and temporal distance, i.e. language used to describe the most temporally proximate events is the most concrete. Unlike study 1, however, we were able to demonstrate this relationship by controlling for the specific event being referenced. We interpret our results as evidence that the observed differences in concreteness of the articles can be confidently attributed to changes in temporal distance, rather than changes in the type of events discussed with reference to the proximal and distant future.

Extending the results of study 1 and previous work (Snefjella & Kuperman, 2015), we also found that articles published prior to events are more concrete than articles published after the events. Although this is not a prediction of CLT, it is consistent with its main tenets. According to CLT, the relationship between construal and distance is functional (Trope & Liberman, 2010) allowing us to better plan for the future or learn from our mistakes. It is not unlikely that the differences in concreteness before vs. after events are a product of these functional properties. Similar conclusions were drawn by Caruso et al. (2013), who observed that future events are perceived as more proximate than past events, even when the objective distance between an individual and these events is the same (e.g., 1 year). The authors argued that such asymmetry is functional, and that this "bias toward the future" may serve as an adaptive mechanisms for dealing with, and preparing for, future threats and rewards. It is important to note that the asymmetry in abstractness was not observed by Snefjella and Kuperman (2015). Since the authors did not control for the objective temporal distance separating the event from the time when it was referenced, it is possible that the lack of asymmetry is due to different topics/events being the target of online communication about the past and future.

Based on the findings of our work and those presented by Caruso et al. (2013) it could be argued that abstractness follows the subjective (i.e. perceived), rather than objective distance. In other words, if future appears to be more psychologically proximate, then increased language concreteness may represent just that. While our studies were not designed to dissociate between the two possibilities, neither is inconsistent with the main tenets of the CLT. Nonetheless, our findings make an important contribution toward a better understanding of the role of psychological distance.

Finally, we have also attempted to fit various mathematical functions to the data in study 2 in an attempt to describe the relationship between concreteness and temporal distance. By doing so we have found that the exponential function is better at describing changes in concreteness over time compared to similarly parametrized polynomial and power functions. A recent meta analysis by Soderberg et al. (2015) finds that the effect of distance on construal are nonlinear, and our model fits provide further evidence in support of this. The exponential model is also frequently used to model intertemporal choice (e.g. Frederick, Loewenstein, & O'donoghue, 2002), particularly, the preference for proximate rewards over distant rewards. Event construal has previously been used to explain some of the findings in this domain (see Trope & Liberman, 2003), and our model fits suggest that the drop in concreteness with event distance may be quantitatively related to the drop in reward desirability with reward distance.

4. General discussion

In this paper, we used large corpora of natural language data to test key premises of the construal level theory (CLT) (Trope et al., 2007). Particularly, using Twitter posts (study 1) and New York Times articles (study 2), we showed that temporally distant events are represented less concretely than similar, proximate events. The results of study 1 were established by comparing the concreteness of words used in tweets referencing proximate dates compared to distant dates. This is very similar to the methodology in a recent study (Snefjella & Kuperman, 2015; study 2) which uses references to dates on an online discussion forum to examine the relationship between construal and temporal distance. Although both sets of studies provide valuable converging evidence, they also admit an important confound: It is possible that the distant events referenced in the data are different to the proximate events in the data. Thus changes in word concreteness with temporal distance may be caused not by changes in construal with increasing distance, but rather changes in what is discussed with increasing distance.

Our study 2 addresses this limitation by examining changes in word concreteness keeping the event being referenced fixed. We achieve this by using a large database of NYT articles that are tagged with the events they reference, making it possible to exclusively examine articles published at different points of time, but referencing a single event. By examining this type of data, study 2 is unique in testing the role of temporal distance in event construal with rigorous controls for the event in consideration. It is also unique in being able to look at changes in construal for an event over many months and years before *and* after the event occurs. Indeed, we found that the overall level of concreteness is higher before an event compared to after the event, and that the changes in concreteness over time are best described using an

Table 1
Model fits for exponential, polynomial, and power function on our data. This table also displays best fitting model parameters for the exponential model (which is the best fitting model).

	Exponential		Polynomial		Power		Exponential parameters		
	R^2	RMSE	R^2	RMSE	R^2	RMSE	α	β	γ
Elections before	0.785	0.004	0.785	0.004	0.606	0.006	2.306	0.179	0.0002
Holidays before	0.900	0.012	0.905	0.012	0.755	0.019	2.480	0.163	0.012
Elections after	0.641	0.007	0.345	0.009	0.435	0.009	2.439	0.057	0.030
Holidays after	0.621	0.022	0.335	0.030	0.430	0.027	2.542	0.114	0.082
Average	0.737	0.011	0.592	0.014	0.556	0.015			

exponential function (matching Soderberg et al.'s (2015) recent finding that the effect of distance on construal is nonlinear). As discussed above, these findings highlight the functional role of construal in planning for the future and learning from the past, and potentially in discounting future rewards or past experiences.

According to CLT, psychological distance determines our internal representation of the world. At the level of cognitive processing, an object's abstraction can influence the storage and retrieval of information (Paivio, 1990; Schwanenflugel & Shoben, 1982), as well as its affective connotation (Kousta, Vigliocco, Vinson, Andrews, & Del Campo, 2011). Abstraction also affects social perception and categorization, with abstractly represented individuals being more likely to be stereotyped (Milkman, Akinola, & Chugh, 2012). Finally, a large body of work has shown that the abstraction or concreteness of an object influences judgments and decisions involving that object (e.g., Henderson, Wakslak, Fujita, & Rohrbach, 2011; Liberman & Trope, 1998). Since construal level manifests itself so clearly in language used on Twitter and in New York Times articles, it can be expected that these processes have serious consequences in the real-world setting. That is, people's predictions and evaluation of discussed topics will be affected by the psychological distance in the manner predicted by the CLT. For example, consider the process of opinion formation about potential presidential candidates based on news articles written about them. According to CLT and consistent with our data, representations of the candidates long before the election will be characterized by high-level construal; their abilities and characteristics will be described in abstract, simple and decontextualized way (Trope & Liberman, 2003). This, in turn, will influence people's evaluative judgments of the candidates.

Naturally, there are limitations to what online communication can reveal about the individual psychological processes. Only partial information about the individuals can be extracted from Twitter, although recent work has shown that it is possible to accurately establish users' personality profile (Qiu, Lin, Ramsay, & Yang, 2012), or basic demographics (Cavazos-Rehg, Krauss, Grucza, & Bierut, 2014). As these methods continue to develop, care needs to be taken when making inferences about cognitive processes from big data of online communication. Also, it needs to be noted that our analysis used a computer algorithm to determine tweet abstractness. Some researchers used crowdsourcing platforms to code tweets according to their content (e.g. Chorley, Colombo, Allen, & Whitaker, 2015). Although not necessary for the purposes of this paper, such methods may be more suitable for a more in-depth analysis of the cognitive–affective components of attitudes and behaviors.

There are also limitations with our use of New York Times articles. First, it needs to be noted that articles may not be most representative of natural discourse. Second, it is plausible that events other than presidential elections and public holidays are characterized by a different relationship between temporal distance and construal level. Future research should determine whether our findings generalize to other types of natural language corpora and other types of events or objects.

A related issue involves the fact that differences in concreteness over time look relatively minor on the five-point scale on which concreteness is being measured. This is due to the fact that tweets and news articles have considerable variability in terms of their content, and can span a very large number of topics even if they refer to the same event (e.g. election or holiday). This variability does not pertain to the effect of distance on construal, but nonetheless serves to bring averages for tweets or news articles at different points in time closer together. Future work should attempt to more closely restrict the topic of the text being analyzed, to cleanly observe the size of the effect of temporal distance on the concreteness of words.

Finally, it is important to note that our analysis of tweet and article concreteness uses very simple approach: it merely averages the concreteness of words used in the text. Although this is a suitable preliminary method for assigning concreteness ratings to texts composed of multiple words, it is clear that more rigorous results could be achieved by parsing the tweets and articles into their various syntactic components, and then analyzing the concreteness of these components separately. This would not only ensure that the effects of various logical connectives, such as negations, are adequately controlled for, but would also shed considerable light on the linguistic features of distance-based changes to event construal. For example, it is possible that temporal distance only affects the predicates of a particular sentence, and not the arguments that these predicates apply to. Future work should examine this hypothesis, as well as others, as to better understand the relationships between construal, temporal distance, and expressions in natural language.

Acknowledgments

References

Borovoi, L., Liberman, N., & Trope, Y. (2010). The effects of attractive but unattainable alternatives on the attractiveness of near and distant future menus. *Judgment and Decision Making, 5*(2), 102–109.

Boyd, D., & Crawford, K. (2012). Critical questions for big data. *Information, Communication & Society, 15*(5), 662–679.

Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods, 46*(3), 904–911.

Caruso, E. M., Van Boven, L., Chin, M., & Ward, A. (2013). The temporal Doppler effect: When the future feels closer than the past. *Psychological Science, 24*, 530–536.

Cavazos-Rehg, P., Krauss, M., Grucza, R., & Bierut, L. (2014). Characterizing the followers and tweets of a marijuana-focused Twitter handle. *Journal of Medical Internet Research, 16*(6).

Chorley, M. J., Colombo, G. B., Allen, S. M., & Whitaker, R. M. (2015). Human content filtering in Twitter: The influence of metadata. *International Journal of Human–Computer Studies, 74*, 32–40.

Förster, J., Friedman, R. S., & Liberman, N. (2004). Temporal construal effects on abstract and concrete thinking: Consequences for insight and creative cognition. *Journal of Personality and Social Psychology, 87*(2), 177–189.

Frederick, S., Loewenstein, G., & O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature, 40*(2), 351–401.

Henderson, M. D., Wakslak, C. J., Fujita, K., & Rohrbach, J. (2011). Construal level theory and spatial distance. *Social Psychology, 42*(3), 165–173.

Kousta, S.-T., Vigliocco, G., Vinson, D. P., Andrews, M., & Del Campo, E. (2011). The representation of abstract words: Why emotion matters. *Journal of Experimental Psychology: General, 140*(1), 14–34.

Liberman, N., Sagristano, M. D., & Trope, Y. (2002). The effect of temporal distance on level of mental construal. *Journal of Experimental Social Psychology, 38*(6), 523–534.

Liberman, N., & Trope, Y. (1998). The role of feasibility and desirability considerations in near and distant future decisions: A test of temporal construal theory. *Journal of Personality and Social Psychology, 75*(1), 5–18.

Liberman, N., & Trope, Y. (2008). The psychology of transcending the here and now. *Science (New York, N.Y.), 322*(5905), 1201–1205.

Magee, J. C., Milliken, F. J., & Lurie, A. R. (2010). Power differences in the construal of a crisis: The immediate aftermath of September 11, 2001. *Personality & Social Psychology Bulletin, 36*(3), 354–370.

Milkman, K. L., Akinola, M., & Chugh, D. (2012). Temporal distance and discrimination: An audit study in academia. *Psychological Science, 23*(17), 710–717.

Paivio, A. (1990). *Mental representations*. Oxford University Press, Incorporated.

Qiu, L., Lin, H., Ramsay, J., & Yang, F. (2012). You are what you tweet: Personality expression and perception on Twitter. *Journal of Research in Personality, 46*(6), 710–718.

Reips, U.-D., & Garaizar, P. (2011). Mining twitter: A source for psychological wisdom of the crowds. *Behavior Research Methods, 43*(3), 635–642.

Rim, S., Uleman, J. S., & Trope, Y. (2009). Spontaneous trait inference and construal level theory: Psychological distance increases nonconscious trait thinking. *Journal of Experimental Social Psychology, 45*(5), 1088–1097.

Sandhaus, E. (2008). *The New York Times Annotated Corpus LDC2008T19. DVD*. Philadelphia: Linguistic Data Consortium.

Schwanenflugel, P. J., & Shoben, E. J. (1982). Differential context effects in the comprehension of abstract and concrete verbal materials. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9* (1), 82–102.

Semin, G. R., & Fiedler, K. (1988). The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of Personality and Social Psychology, 54*(4), 558–568.

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

Soderberg, C. K., Callahan, S. P., Kochersberger, A. O., Amit, E., & Ledgerwood, A. (2015). The effects of psychological distance on abstraction: Two meta-analyses. *Psychological Bulletin, 141*(3), 525.

Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review, 110*(3), 403–421.

Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review, 117*(2), 440–463.

Trope, Y., Liberman, N., & Wakslak, C. (2007). Construal levels and psychological distance: Effects on representation, prediction, evaluation, and behavior. *Journal of Consumer Psychology: The Official Journal of the Society for Consumer Psychology, 17*(2), 83–95.

Wakslak, C. J., Trope, Y., Liberman, N., & Alony, R. (2006). Seeing the forest when entry is unlikely: Probability and the mental representation of events. *Journal of Experimental Psychology: General, 135*(4), 641–653.