Living in the Past, Present, and Future:
Measuring Temporal Orientation with Language

Gregory Park[1], H. Andrew Schwartz[1,2], Maarten Sap[1], Margaret L. Kern[3],
Evan Weingarten[1], Johannes C. Eichstaedt[1], Jonah Berger[1], David J. Stillwell[4],
Michal Kosinski[5], Lyle H. Ungar[1], Martin E. P. Seligman[1]
[1] University of Pennsylvania, [2] Stony Brook University, [3] The University of Melbourne,
[4] University of Cambridge, [5] Stanford University

Author Note

Gregory Park, Department of Psychology, University of Pennsylvania; H. Andrew Schwartz, Computer & Information Science, University of Pennsylvania and Computer Science, Stony Brook University; Margaret L. Kern, Melbourne Graduate School of Education, The University of Melbourne, Australia; Maarten Sap, Department of Psychology, University of Pennsylvania; Evan Weingarten, The Wharton School, University of Pennsylvania; Johannes C. Eichstaedt, Department of Psychology, University of Pennsylvania; David J. Stillwell, Psychometrics Centre, University of Cambridge, UK; Michal Kosinski, Graduate School of Business, Stanford University; Lyle H. Ungar, Computer & Information Science, University of Pennsylvania; Martin E. P. Seligman, Department of Psychology, University of Pennsylvania

Support for this article was provided by grant #63597 from the Robert Wood Johnson Foundation (M. E. P. Seligman, PI) and by a grant from the Templeton Religion Trust (M.E.P. Seligman, H. A. Schwartz, L. H. Ungar, co-PIs).

Correspondence concerning this article should be addressed to Gregory Park, Department of Psychology, University of Pennsylvania, 3701 Market Street, 2nd floor, Philadelphia, PA 19104. Email: gregpark@sas.upenn.edu

Final accepted version, 20 Dec 2015 (*Journal of Personality*)
This preprint may differ slightly from the final, copy-edited version of record.

Abstract

Objective: Temporal orientation refers to individual differences in the relative emphasis one places on the past, present, or future, and is related to academic, financial, and health outcomes. We propose and evaluate a method for automatically measuring temporal orientation through language expressed on social media.

Method: Judges rated the temporal orientation of 4,302 social media messages. We trained a classifier based on these ratings, which could accurately predict the temporal orientation of new messages in a separate validation set (accuracy/mean sensitivity = .72; mean specificity = .77). We used the classifier to automatically classify 1.3 million messages written by 5,372 participants (50% female, aged 13-48). Finally, we tested whether individual differences in past, present, and future orientation differentially related to gender, age, Big Five personality, satisfaction with life, and depressive symptoms.

Results: Temporal orientations exhibit several expected correlations with age, gender, and Big Five personality. More future-oriented people were older, more likely to be female, more conscientious, less impulsive, less depressed, and more satisfied with life; present orientation showed the opposite pattern.

Conclusion: Language-based assessments can complement and extend existing measures of temporal orientation, providing an alternative approach and additional insights into language and personality relationships.

Keywords: temporal orientation, language, computational social science, social media, big data

Living in the Past, Present, or Future: Measuring Temporal Orientation with Language

Consider three pairs of emotions: (a) regret and nostalgia, (b) boredom and joy, and (c) dread and hope. In each pair, emotions are opposed in valence but similar in orientation towards the past (a), present (b), or future (c). Psychological research has mostly concentrated on understanding people's tendencies to express positive or negative emotions, but less attention has been given to their relative focus on the past, present, or future. One reason may be that these *temporal orientations* are hard to measure with traditional self-report methods. We introduce a method for automatically assessing temporal orientation through language expressed in social media. In addition, we explore differences across age and gender, and connections to personality, subjective well-being, and depressive symptoms.

Studies on Temporal Orientation

Most studies of temporal orientation have focused on *future-oriented* thinking and its relation to educational, health, and financial outcomes. For example, students with higher future orientation study longer and earn better grades (Horstmanshof & Zimitat, 2007; Zimbardo & Boyd, 1999), and more future-oriented adults use less alcohol and tobacco (Adams & Nettle, 2009; Daughterty & Brase, 2010; Keough, Zimbardo, & Boyd, 1999), practice safer sex (Rothspan & Read, 1996), exercise more frequently (Ouellette, Hessling, Gibbons, Reis-Bergan, & Gerrard, 2005), hold more positive attitudes towards exercise (Joireman, Shaffer, Balliet, & Strathman, 2012), control diets better (Piko & Brassai, 2009), have lower body mass indexes, (Adams & Nettle, 2009; Adams & White, 2009), save more of their income (Webley & Nyhus, 2006), and plan their finances further into the future (Adams & Nettle, 2009).

Present and future orientations also have well-established age differences. As people grow older, they report thinking less about the present and more about the future (Casey, Jones,

& Hare, 2008; Nurmi, 2005; Steinberg et al., 2009). Early childhood is characterized by a preoccupation with the immediate present, whereas weighing the consequences of today's decisions is a hallmark of maturity. According to questionnaire measures, future-oriented thinking begins in early adolescence, becomes more common throughout adolescence, and levels off in young adulthood (Steinberg et al., 2009).

Studies have also found smaller but consistent gender differences in temporal orientation. Across eight samples, Keough et al. (1999) found women were more future-oriented and men were more present-oriented. Steinberg et al. (2009) reported that women scored significantly higher than men on three measures of future orientation.

Measuring Temporal Orientation

Temporal orientation is typically measured by self-reports, such as the Zimbardo Time Preference Inventory (ZPTI; Zimbardo & Boyd, 1999) and the Consideration of Future Consequences scale (CFC; Joireman et al., 2012; Strathman, Gleicher, Boninger, & Edwards, 1994). Respondents rate statements about their thinking or planning style, and these items form subscales measuring past ("It gives me great pleasure to think about my past"; ZPTI), present ("I often follow my heart more than my head"; ZPTI), and future orientations ("When I make a decision, I think about how it might affect me in the future"; CFC). These measures are easy to administer and predict several outcomes, as noted above.

However, these self-reported items highly overlap with self-reported measures of personality traits. For example, future orientation is strongly correlated with *conscientiousness* (*r*s range from .50 to .60; Strathman et al., 1994; Zhang & Howell, 2011; Zimbardo & Boyd, 1999). It may be that conscientiousness predisposes a person to be more future-oriented, but such distinctions are complicated by the fact that questionnaire measures of conscientiousness and future orientation are also very similar. For example, the ZPTI Future scale includes the item "I make lists of things to do", while conscientiousness scales include items such as "I do things according to a plan" (Goldberg et al., 2006). A behavior-based measure of temporal orientation could provide researchers with an alternative method that has less overlap with measures of similar constructs.

Likewise, self-reports often have an implicit evaluative component, such as the ZPTI's *Past-Negative* (e.g., "Painful past experiences keep being replayed in my mind) and *Past-Positive* (e.g., "It gives me pleasure to think about my past") subscales. These two subscales correlate with measures related to subjective well-being (neuroticism, depression, and self-esteem; Zimbardo & Boyd, 1999). The evaluative aspect—the tendencies to rate experiences and memories as positive or negative—may be driving these correlations, rather than a true association with temporal orientation. If so, these measures cannot assess the unique contribution of temporal orientation on well-being.

These measurement confounds prevent researchers from clearly separating temporal orientation from other related traits. One solution lies in behavior-based measures (Roberts, Harms, Smith, Wood, & Webb, 2006). Behavior-based measures remove the shared method variance with self-reports (i.e., overlapping, similar items), reduce the influence of a respondent's evaluative style, and enable multi-method designs. Language use provides one psychologically rich and practical source of behavioral data (Kern et al., 2014; Pennebaker, Mehl, & Niederhoffer, 2003). When combined with techniques from natural language processing, statistical models can accurately predict several individual characteristics—age, gender, and personality—from language alone (Park et al., 2015; Schwartz et al., 2013b).

In the current study, we created a new language-based measure of temporal orientation. First, we developed a model to classify text as oriented towards the past, present, or future. We used this model to classify millions of Facebook *status updates* (i.e., short text messages used to describe someone's current mood, thoughts, activities, or plans), creating a person-level measure of past, present, and future orientation. We then compared orientations to age, gender, and personality—checking for consistency with patterns found using self-reports— and then extended these comparison to life satisfaction and depression.

Part 1: Message-level Temporal Classification Model

We developed a classification model on one set of language data, with the goal of automatically classifying a second set of data as past-, present-, or future-oriented on the basis of several linguistic features (Schwartz et al., 2015). This process required that we (1) obtain a set of text samples for training; (2) annotate these text samples as past-, present-, or future-oriented; (3) extract linguistic features (e.g., words, phrases, number of words) from each text sample; (4) train a statistical model to predict the text's temporal annotation based on its linguistic features; and (5) evaluate the accuracy of this model on a new set of messages.

Training Messages

For our initial set of text samples, we used 6,000 messages from Twitter and Facebook. From Twitter (a microblogging platform on which users can post short text messages, or "tweets", limited to 140 characters), we sampled 3,000 messages, drawn from a random feed provided by Twitter during September 2012. From Facebook, we sampled 3,000 status updates, drawn from users of the MyPersonality application (Kosinski, Stillwell, & Graepel, 2013) between January 2009 and October 2011. MyPersonality is a third-party application through which users can complete personality and other psychological measures and share results with friends. Users voluntarily allowed the application to access all of their Facebook status updates for research purposes. Of the 6,000 training messages, 1,489 were identified as song lyrics, famous quotations, or posts by bot (i.e., automated) accounts, and these were removed from the training sample.

Message Annotation

Three independent judges rated the temporal orientation of each of the remaining 4,511 messages, using fractions of the day in the past or future. For example, a message referring to the immediate present was rated as 0, an hour in the future was +1/24, 1 day in the future was +1, one week in the future was +7, and one day in the past was -1. Judges were instructed to mark non-interpretable messages as ‘NA’. We removed messages that were rated ‘NA’ by all three raters, which excluded an additional 209 messages (125 Twitter messages and 84 Facebook messages). Inter-rater agreement for the remaining 4,302 messages was high (intraclass correlation coefficient, ICC =.85).

We used the mean rating to classify each message into three categories: past-oriented (mean rating < 0), present-oriented (mean rating = 0), or future-oriented (mean rating > 0). Table 1 lists examples of messages, individual ratings, and final orientation classification. Of the 4,302 messages, 1,178 (27.4%) were classified as past-oriented, 2,043 (47.5%) as present-oriented, and 1,081 (25.1%) as future-oriented. Of the 2,293 Facebook messages, 659 (28.7%) were classified as past-oriented, 990 (43.2%) as present-oriented, and 644 (28.1%) as future-oriented. Of the 2,009 Twitter messages, 519 (25.8%) were classified as past-oriented, 1,053 (52.4%) as present-oriented, and 437 (21.8%) as future-oriented.

Linguistic Feature Extraction

We extracted five types of linguistic features from each message: words and phrases, time expressions, parts of speech, word categories, and length of message.

Words and phrases. We used an emoticon-aware tokenizer (*happierfuntokenizing*; Potts; 2011) to divide messages into smaller word-like units, or *tokens*. The tokenizer was sensitive to single words, punctuation, non-conventional usages and spellings (e.g., *omg*, *lol*) and emoticons (e.g., *:-]*), which are common on social media. We represented a message's constituent words, phrases, and similar features using a binary encoding. That is, for each message, if a given word or phrase appeared at least once, it was coded as 1, otherwise it was coded as 0.

Time expressions. We used the Stanford SUTime annotator (Chang & Manning, 2012) to identify time expressions (e.g., "yesterday", "next September") within each message. Once identified, time expressions were used to derive six features: the mean temporal difference (in days) between all time expressions in the message and the time of the message's creation, the log (base 2) of this difference, the absolute value of the difference, and three binary variables encoding whether any time expressions in the messages referred to the past, present, or future. We also added a feature coding that indicated the total number of time expressions that occurred in the message.

Parts of speech. We used Stanford's part-of-speech tagger (Toutanova, Klein, Manning, & Singer, 2003) to identify each token's corresponding part of speech. For each possible part of speech tag, we calculated the frequency of the tag within each message and divided the frequency by the total number of tokens in each message.

Word categories. We used the Linguistic Inquiry and Word Count (LIWC; Pennebaker, Chung, Ireland, Gonzales, & Booth, 2007) dictionaries to count the frequency of words in 64 pre-defined categories, including temporally-oriented categories such as *future* words (e.g., "will", "gonna", "might"). The frequency of words within each LIWC category was divided by the total number of tokens in the message, resulting in 64 separate features.

Message length. Two features captured message length: the mean length (i.e., number of characters) of all tokens in the message, and the total number of tokens in the message.

Temporal Classification Model

After extracting linguistic features from each message, we fit a statistical model over the set of training messages to predict their rated temporal orientation from the features. Because this task requires classification into three categories (past, present, and future), we explored four classification techniques, implemented in the *scikit-learn* Python module (Pedregosa et al., 2011): logistic regression (LR) with Lasso regularization, support vector classification with a linear kernel (lSVC), support vector classification with a radial basis kernel (rSVC), and a forest of extremely randomized trees (ERT).

ERT fits many (hundreds or more) single trees to random portions of the training data, and then combines the individual predictions to form a more stable ensembled prediction. Traditional decision tree models naturally handle non-linear relationships and interactions between predictors, but single trees are unstable and prone to overfitting (Berk, 2008). In our case, each decision tree was fit to a random subset of messages from the training data and a random subset of features. Splits at each node in the decision tree were also randomly chosen. We used the following ERT parameters: we built 1,000 trees, chose node splits using the Gini impurity measure, and used the square-root of the total number of features as the amount of randomly selected features when building each tree. To classify a new message's temporal orientation, we applied the 1,000 fitted trees to the new message (i.e., its corresponding features) and used the most frequent class as the predicted class.

Model evaluation. We evaluated the performance of all four techniques by applying it to a new independent set of messages. We randomly sampled 500 Facebook status updates from the MyPersonality data set (not included in the training set), and three independent judges rated the message orientation as either past, present, or future.[1] Agreement between raters was high (ICC = 0.83). We used the majority rating as each message's temporal orientation. The resulting orientations of the messages were 131 (26.2%) past-oriented, 250 (50.0%) present-oriented, and 105 (21.0%) future-oriented. Fourteen messages were three-way ties (one past, one present, one future), and these messages were coded as present (the most frequent class). We then applied each classification technique to these messages, comparing the agreements between model prediction and human ratings.

As benchmarks, a random classifier would have an accuracy of 0.33, and predicting the most frequent class (present) would yield an accuracy of 0.53. The resulting accuracies of the four techniques were logR = .69, lSVC = .71, rSVC = .68, and ERT = .72. We concluded that the ERT model was best for automatically classifying new messages.[2] Mean specificity of the ERT model, or how often a message was correctly *not* classified as an incorrect class, was 0.77. Of the 131 messages that were truly past (based on human judgments), 79 were predicted as past, 42 as present, and 10 as future. Of the 264 messages that were truly present, 15 were predicted as past, 232 as present, and 17 as future. Of the 105 messages that were truly future, 9 were predicted as past, 46 as present, and 50 as future.

To evaluate the relative importance of each feature type in the ERT model, we examined how model performance changed across different combinations of features. We started by using only one feature type to classify messages, resulting in the following accuracies: only message lengths (.54), only time expressions (.59), only parts of speech (.61), only word categories (.68), and only words and phrases (.69). We then tested the model performance using all *except* one feature type, resulting in the following accuracies: all except words and phrases (.67), all except word categories (.70), all except time expressions (.71), all except parts of speech (.71), and all except message lengths (.72). We concluded that all feature types but message lengths add useful information and improve performance. However, the inclusion of message length features does not reduce model performance, so we used all five feature types in the final model.

Part 2: Assessment of Person-level Temporal Orientation

After developing an accurate model, we then applied the model to a much larger set of messages from Facebook users, and compared their aggregated temporal patterns to several self-reported individual characteristics.

Participants

Participants were drawn from a pool of 72,559 users of the MyPersonality Facebook application who were not a part of the training set, who also granted access to all status messages, written between June 2009 and November 2011. This pool of users was 62% female with an average age of 23.3 years old (*SD* = 8.9; median = 20). For practical purposes, we sampled a smaller subset of users, rather than use the full pool. The pool of users wrote over 20 million messages, and extracting linguistic features, particularly the syntactic parsing needed to extract time-expressions from all of these messages is a very time-intensive process. We reasoned that a smaller sample of participants and messages (i.e., about 5,000 participants with roughly one million messages) would still yield stable estimates but also allow a much shorter development cycle (i.e., days instead of weeks).

The full MyPersonality sample had a high concentration of users between the ages of 18-22 (36% of users) and more women than men (61% of users). To ensure the subsample included adults from a large age span, we stratified our sample across age and gender, which resulted in a much more balanced sample. We also wanted to ensure that the participants in our sample had completed other relevant psychological measures. To satisfy these requirements, we sampled two subsets of participants.[3]

Subset 1 was an age- and gender-balanced sample, which was created by randomly sampling 180 participants (90 men, 90 women) from two-year age bins ranging from 13 to 48 ([13, 14], [15, 16] … [47, 48]), resulting in a sample of 3,240 participants. All participants in this stratified sample reported their age, gender, completed a self-report measure of Big Five personality factors (detailed below), and wrote at least 100 status updates. The mean and median age of the resulting subsample were 30.5.

Subset 2 included 2,132 participants who reported age, gender, wrote at least 100 status updates, and completed at least one measure of impulsivity, life satisfaction, or depressive symptoms. The subset included 754 men and 1,378 women, and had a mean age of 21.7 (*SD* = 7.6, median = 19.0).

Measures

Big Five Personality. All participants from subset 1 completed items assessing Big Five personality (openness to experience, conscientiousness, extraversion, agreeableness, and neuroticism) from the International Personality Item Pool (IPIP; Goldberg et al., 2006). All participants completed at least the 20-item version of this measure. Participants could optionally complete additional IPIP items; 636 participants completed the full 100-item version of measure.

Barratt Impulsiveness Scale. From subset two, 762 participants completed the Barratt Impulsiveness Scale (BIS; Stanford et al., 2009), a 30-item assessment of general impulsiveness. Each BIS item states a manner of acting or thinking (e.g., "I do things without thinking", "I buy things on impulse"), and participants indicate how accurately each statement describes themselves on a 4-point scale (1 = rarely/never; 4 = almost always/always). For 76 participants who were missing responses for a single item, we imputed the single missing value with the mean of the remaining items. We excluded 18 participants who were missing scores on more than one item, leaving 744 participants with BIS scores. We calculated the full-scale score as the mean across all 30 items (Cronbach's $\alpha = .83$).

Satisfaction with Life. From subset two, 1,369 participants completed the Satisfaction with Life Scale (SWLS; Diener, Emmons, Larsen, & Griffin, 1985), a five-item assessment of life satisfaction. Participants indicate their agreement with five statements (e.g., "I am satisfied with my life", "The conditions of my life are excellent") on a 7-point scale (1 = strongly disagree; 7 = strongly agree). There were no missing responses across the participants who met the inclusion criteria for subset 2. For 79 participants that completed the SWLS more than once, we only used data from the first administration. We calculated the full-scale score as the mean across the five items ($\alpha = .87$).

Center for Epidemiologic Studies Depression Scale. From subset two, 420 participants completed the Center for Epidemiologic Studies Depression Scale (CES-D; Radloff, 1977), a 20-item measure of self-reported depressive symptoms. Each item describes a symptom (e.g., "I felt depressed", "I had crying spells"), and participants indicated the frequency of experiencing each symptom on a 4-point scale (1 = rarely or none of the time; 4 = most or all of the time). For 42 participants who were missing responses for a single item, we imputed the missing item with the mean of the remaining items. We excluded nine participants who were missing scores on more than one item. We calculated the mean across all items as the total scale score for the remaining 411 participants ($\alpha = .85$).

Person-Level Evaluation

In total, participants from the two subsets wrote 1,323,939 messages (each participant individually wrote at least 100 messages). We applied the temporal classifier developed in Part 1 to every message. For each participant, we calculated the number of his/her messages that were classified as past, present, or future, and then divided these three frequencies by their total number of messages, resulting in the proportions of a person's message that were past, present, and future-oriented. On average, 19% of participants' messages were past-oriented, 65% were present-oriented, and 16% were future-oriented.

Relevant language features. To better understand which language features were relevant to classification in this new set of messages, we examined which 1-grams (i.e., single words or tokens) were most strongly correlated with classifications of past, present, and future. We chose to examine 1-grams (as opposed to two or three word phrases) because they are more easily interpreted than other features used by the model. To calculate these correlations, we first recoded every message-level classification as three binary variables (e.g., past = 0/1; present = 0/1; future = 0/1), where a 1 indicated the message's orientation. For each orientation, we correlated the message-level relative frequency of single words with the corresponding binary variable. In the resulting correlations, high positive correlations indicate that greater frequency of a given word was correlated with that temporal orientation.

For each orientation, many of the most strongly correlated 1-grams included some clear temporal information, either in verb tense (e.g., *was*, *is*, or *will*) or as a part of a temporallyrelevant phrase. For example, the 20 1-grams most strongly correlated with past orientation were (correlations shown in parentheses; all correlations are $p < .05$, Bonferroni-corrected) *was* (.37), *had* (.28), *got* (.25), *did* (.16), *went* (.15), *just* (.13), *last* (.12), *made* (.12), *been* (.11), *saw* (.11), *a* (.10), *were* (.10), c*ame* (.09), *said* (.09), *from* (.08), *found* (.08), *today* (.07), *didn't* (.07), *thought* (.07), and *he* (.06). The 1-grams most correlated with present orientation included present-tense verbs but also words likely used in interpersonal communication (e.g., second-person pronouns) and questions: were *is* (.13), *you* (.11), *love* (.09), *are* (.08), *?* (.07), *your* (.07), *happy* (.06), *don't* (.05), *life* (.05), *like* (.05), *people* (.05), *why* (.04), *want* (.04), *can* (.04), quotation marks (“”; .04), *know* (.04), ellipses (... ; .04), *you're* (.03), *right* (.03), and *do* (.03). The 1-grams correlated with future orientation included future tense verbs and time-related words: *going* (.28), *to* (.22), *tonight* (.21), *will* (.19), *wait* (.18), *be* (.12), *days* (.12), *get* (.11), *today* (.10), *go* (.10), *then* (.09), next (.08), *for* (.08), *soon* (.08), *see* (.07), *until* (.06), *excited* (.06), *can't* (.05), *watch* (.05), and *this* (.05).

Age and gender. Past and future orientation increased markedly with age; present orientation decreased markedly. Table 2 summarizes Pearson correlations (*r*) between user-level temporal orientations and age, calculated using the age-stratified subset 1. To illustrate, we standardized user-level orientations and plotted the mean standard score of each age group for each orientation (Figure 1; for an alternate display showing individual data points, see Figure A1 in Appendix A). Across all age groups, the rank order of past, present, and future orientation remained the same: present-oriented messages were always the most frequent and future-oriented were least frequent. However, there were large differences in the relative proportion of each orientation across age.

We considered the possibility that younger users may write messages more frequently than older users, and therefore younger users would be more likely to write about the present, simply because less time has passed since writing their last message. To test whether message frequency accounted for age differences in temporal orientation, we recalculated correlations between age and orientations while adjusting for each user's total number of messages. These adjusted correlations ($r_{age \times past\_adj}$ = .21; $r_{age \times present\_adj}$ = -.23; $r_{age \times future\_adj}$ = .16) were virtually identical to the unadjusted correlations ($r_{age \times past}$ = .21; $r_{age \times present}$ = -.23; $r_{age \times future}$ = .16), indicating that age differences count not be accounted for by younger users' higher message frequency.

Women were more past-oriented (overall Cohen's d = .10; 95% CI = [.03, .17]), less present-oriented (d = -.27; [-.20, -.34]), and more future-oriented than men across all ages (d = .34; [.27, .41]). We checked for changes in gender differences across age bins by calculating *d*s within each two-year age group and then regressing the *d*s on age. We found no significant trends in *d*s over age (b_{past} = .006, p = .162; $b_{present}$ = -.004, p = .397; b_{future} = -.001, p = .748).

Personality. Temporal orientation was most strongly associated with conscientiousness and openness to experience. More future-oriented people were more conscientiousness (r = .14 [.10, .17]) but less open (r = -.14 [-.17, -.10]), while the opposite pattern occurred in more present-oriented people ($r_{conscientiousness}$ = -.11, [-.14, -.07]; $r_{openness}$ = .09 [.06, .12]). Table 2 lists all *r*s and 95% confidence intervals between orientations and Big Five personality factors, calculated within subset 1.

Impulsiveness, life satisfaction, and depressive symptoms. With subset 2, we calculated Pearson correlations between each temporal orientation and impulsiveness, satisfaction with life, and depressive symptoms. We controlled for participants' age and gender by standardizing each outcome measure and temporal orientation, and then regressing temporal orientation on each outcome, with age and gender as covariates. The resulting coefficient on temporal orientation is equivalent to a Pearson correlation adjusted for age and gender. Higher future orientation was significantly correlated with lower impulsiveness (r = -.08 [-.16, -.01]), higher life satisfaction (r = .07 [.02, .13]), and fewer depressive symptoms (r = -.16 [-.29, -.03]). In contrast, higher present orientation was significantly correlated with lower life satisfaction (r = -.08 [-.13, -.02]) and more depressive symptoms (r = .16 [.04, .29]).

Self-descriptions from personality items. To complement Big Five correlations with richer psychological descriptions, we examined IPIP personality items that were significantly positively correlated with past, present, or future orientation for a subset of 636 participants who completed the 100-item IPIP measure. Significant self-descriptions are listed in Table 3, and a complete list of all items and correlations is available in Supplement 1.

Discussion

We developed a language-based measure of temporal orientation, and we applied this method to a large sample to explore associations with age, gender, personality, and well-being. This method may be a useful complement to existing methods, particularly when traditional self-report measures would not be feasible.

At the message level, our temporal classifier accurately predicted the orientation of a message, as rated by multiple human judges. At the person level, our measure of temporal orientation converged with external correlates in theoretically expected ways. Future orientation increased with age, whereas present orientation decreased with age. Women were more future-oriented than men. Future orientation correlated with higher conscientiousness, and the self-descriptions from personality items aligned with several characteristics related to different orientations.

We found several small correlations between temporal orientation and Big Five personality dimensions, but the largest were with conscientiousness; conscientious people were more future-oriented and less present-oriented. This aligns well with characterizations of the highly conscientious person, who plans, delays gratification, and controls impulses better than most (Roberts, Lejuez, Krueger, Richards, & Hill, 2014). However, the correlations between temporal orientation and the Big Five were smaller than those seen in previous mono-method, questionnaire-based studies (absolute mean r = .06, versus absolute mean r =.17 in Zimbardo & Boyd, 1999). One explanation for this attenuation is that the use of two different measurement methods (language-based and questionnaire-based) prevents shared method variance from inflating correlations (Roberts et al., 2006).

This method did replicate the expected patterns with age and gender seen in prior self-report studies. Across ages 13 to 48, people were substantially more past- and future-oriented and less present-oriented (Figure 1). This is consistent with trends found in studies of adolescents and young adults (Casey et al., 2008; Steinberg et al., 2009). Age trends were similar in women and men, but we did find a significant gender differences across all ages; women were more future-oriented and only slightly more past-oriented, while men were more present-oriented. The size of the gender difference was consistent with studies using self-reports (e.g., Keough et al., 1999).

By analyzing responses to individual personality items, we found that temporal orientation corresponded to differences in how individuals described themselves (Table 2), particularly when contrasting present and future orientation. Highly present-oriented people may be best characterized as impulsive across many domains—socially (“I cut others to pieces”), emotionally (“I have frequent mood swings”), and motivationally (“I don’t put my mind on the task at hand”)—but also more open to aesthetic experiences (“I believe in the importance of art”) and fantasy (“I enjoy wild flights of fantasy”). Highly future-oriented described a much narrower focus on practical planning (“I carry out my plans”) and getting things done (“I complete tasks successfully”), with little interest in abstract matters (“I avoid philosophical discussions” and “I am not interested in abstract ideas”).

Overall, the contrasting self-descriptions of the present-oriented and the future-oriented are similar to *stability* and *plasticity*, two higher-order traits that describe tendencies to maintain goals or engage with the world (Hirsh, DeYoung, & Peterson, 2009). Whereas stability is the capacity to resist disruption and maintain action towards future goals, plasticity is the capacity for emotional, cognitive, and environmental exploration (DeYoung, 2015). Overemphasis on the present or the future may reflect different trade-offs between these two fundamental motivations. In this framing, highly present-oriented people may be highly exploratory and engaged with the environment (high plasticity) at the cost of more stable long-term goals (low stability, or instability), while highly future-oriented people maintain a strong focus on distant goals (high stability) at the cost of exploration and information gathering from their inner and outer worlds (low plasticity, or rigidity).

More future-oriented people, however, were more satisfied with life and less depressed. Because future orientation predicts favorable educational, financial, and health outcomes (Adams & Nettle, 2009; Keough et al., 1999), it may not seem surprising that it correlates with positive evaluations of one’s life and alleviation from psychological distress. However, this pattern was not clear from prior research on orientations and well-being (Boniwell & Zimbardo, 2004; Zhang

& Howell, 2011), and our method enabled a larger study than typically possible, while removing the evaluative confounds inherent in relying solely on self-report measures.

Applications

Our method may be most valuable as a complement to ongoing studies or existing samples. Participants in a research study might be asked to voluntarily provide access to their social media language (e.g., Facebook status updates or Twitter tweets), and then the classifier can be applied to their posts, quickly adding a measure of temporal orientation or other characteristics. Given the growing popularity of social media platforms (Duggan, Ellison, Lampe, Lenhart, & Madden, 2014), language-based methods can collect large samples much faster than is feasible through other approaches. For instance, human ratings of temporal orientation requires about 90 seconds per message; at this rate, a single human judge would need to rate continuously for over three years to annotate our collection of 1.3 million messages. Our automatic classifier rated this entire set in minutes.

While our method annotated messages to characterize individuals, it can also potentially be adapted to characterize entire geographic regions. Because social media messages often contain fine-grained geographic metadata, messages from well-defined areas (e.g., U.S. counties) can be aggregated, annotated, and compared by orientation. Perceptions of time and the daily tempo of life vary substantially across regions and cultures (Banfield, 1974; Levine, 1997), and these differences may be embedded in language and related to other important outcomes. For example, a recent study of search queries found that countries differ in how much their users search for information about future dates, and that more future-oriented countries have larger per capita gross domestic product (Preis, Moat, Stanley, & Bishop, 2012). Similar social media methods have already been used to characterize regions along psychological dimensions, such as consumer confidence (O'Connor, Balasubramanyan, Routledge, & Smith, 2010), life satisfaction (Schwartz et al., 2013), and hostility (Eichstaedt et al., 2015).

Because we developed the model using a blend of Facebook and Twitter messages, it may generalize to messages written on either platform, but explicit evaluations over Twitter messages are still needed (see Sap et al., 2014 for a successful example of model building across both platforms). However, because both Facebook and Twitter are designed to elicit descriptions of a user's current status, they may be biased toward the present, and the relative proportions of past-, present-, and future-oriented messages may not hold for other online social media platforms. As users shift to other platforms, the extent to which the models need to be adjusted should be considered.

Limitations

Our study also had several limitations. We used a very coarse representation of time, splitting messages into past, present, and future categories. A fine-grained approach that distinguishes near future from the distant future would be more sensitive to the depth of one's temporal horizon. For example, thinking about the distant future may be a better predictor of health and financial behaviors than only thinking about the short-term future.

Second, we focused only on the temporal orientation of a message and ignored other qualities like emotional valence. Incorporating valence may allow distinctions between similarly-oriented emotions, such as regret or positive nostalgia, which have opposite associations with well-being (Sedikides, Wildschut, Arndt, & Routledge, 2008).

Third, our sample consisted of selected sets of social media users, who are not fully representative of the general population. However, the representativeness of social media continues to increase every year. Currently, 58% of all American adults use Facebook, and usage is spread evenly across demographic and socioeconomic lines (Duggan et al., 2014). Even if the findings only apply to the population of social media users, it still represents a considerably larger portion of the general population than small studies with U.S. undergraduates.

Fourth, while our sample spanned a large age range, it did not include adults older than 48 years old. Social media use among older adults is growing every year (31% of adults over 65 use Facebook; Duggan et al., 2014), but this demographic is still underrepresented. This is particularly limiting given our age-related findings, which contrast with the finding that “older people are mostly present-oriented” (Carstensen, Isaacowitz, & Charles, 1999, p. 168). Our sample may have been too young to detect such patterns.

Conclusion

Temporal orientation can be measured through everyday language on social media. Our language-based measure of temporal orientation replicated several theoretically expected patterns with age, gender, and personality, and allowed the discovery of new connections with well-being. As social media expands, our approach complements other measures and can help researchers study temporal orientation at large scale.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Preparation of this manuscript was supported by grants from the Robert Wood Johnson Foundation (#63597) and the Templeton Religion Trust (#TRT0048).

References

Adams, J., & Nettle, D. (2009). Time perspective, personality and smoking, body mass, and physical activity: An empirical study. *British Journal of Health Psychology*, *14*, 83-105.

Adams, J., & White, M. (2009). Time perspective in socioeconomic inequalities in smoking and body mass index. *Health Psychology*, *28*, 83-90.

Banfield, E. C. (1974). *The unheavenly city revisited.* Boston: Little, Brown and Company.

Berk, R. A. (2008). *Statistical learning from a regression perspective.* New York: Springer.

Boniwell, I., & Zimbardo, P. G. (2004). Balancing time perspective in pursuit of optimal functioning. In P. A. Linley & S. Joseph (Eds.), *Positive psychology in practice* (pp. 165-178). Hoboken, NJ: Wiley and Sons.

Carstensen, L. L., Isaacowitz, D. M., & Charles, S. T. (1999). Taking time seriously: A theory of socioemotional selectivity. *American Psychologist*, *54*, 165-181.

Casey, B. J., Jones, R. M., & Hare, T. A. (2008). The adolescent brain. *Annals of the New York Academy of Sciences*, *1124*, 111-126.

Chang, A. X., & Manning, C. D. (2012, May). SUTime: A library for recognizing and normalizing time expressions. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation*. Istanbul, Turkey. Retrieved from http://nlp.stanford.edu/pubs/lrec2012-sutime.pdf

Daugherty, J. R., & Brase, G. L. (2010). Taking time to be healthy: Predicting health behaviors with delay discounting and time perspective. *Personality and Individual Differences*, *48*, 202-207.

DeYoung, C. G. (2015). Cybernetic Big Five theory. *Journal of Research in Personality, 56*, 33–58. http://dx.doi.org/10.1016/j.jrp.2014.07.004

Diener, E., Emmons, R. A., Larsen, R. J., & Griffin, S. (1985). The satisfaction with life scale. *Journal of Personality Assessment, 49,* 71-75. http://dx.doi.org/10.1207/s15327752jpa4901_13

Duggan, M., Ellison, N. B., Lampe, C., Lenhart, A., & Madden, M. (2014). Pew Internet: Social Media Update 2014. Retrieved from: http://www.pewinternet.org/2015/01/09/social-media-update-2014/

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., ... & Seligman, M. E. P. (2015). Psychological language on Twitter predicts county-level heart disease mortality. *Psychological Science, 26*, 159-169. doi:10.1177/0956797614557867

Goldberg, L. R., Johnson, J. A., Eber, H. W., Hogan, R., Ashton, M. C., Cloninger, C. R., & Gough, H. G. (2006). The international personality item pool and the future of public domain personality measures. *Journal of Research in Personality, 40*, 84–96. http://dx.doi.org/10.1016/j.jrp.2005.08.007

Keough, K. A., Zimbardo, P. G., & Boyd, J. N. (1999). Who's smoking, drinking, and using drugs? Time perspective as a predictor of substance use. *Basic and Applied Social Psychology*, *21*, 149-164.

Kern, M. L., Eichstaedt, J. C., Schwartz, H. A., Dziurzynski, L., Ungar, L. H., Stillwell, D. J., ... & Seligman, M. E. (2014). The online social self: An open vocabulary approach to personality. *Assessment, 21*, 386-397. http://dx.doi.org/10.1177/1073191113514104

Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, *110*, 5802-5805.

Hirsh, J. B., DeYoung, C. G., & Peterson, J. B. (2009). Metatraits of the Big Five differentially predict engagement and restraint of behavior. *Journal of Personality*, *77*, 1085-1102.

Horstmanshof, L., & Zimitat, C. (2007). Future time orientation predicts academic engagement among first-year university students. *British Journal of Educational Psychology*, *77*, 703-718.

Joireman, J., Shaffer, M. J., Balliet, D., & Strathman, A. (2012). Promotion orientation explains why future-oriented people exercise and eat healthy: Evidence from the two-factor Consideration of Future Consequences-14 scale. *Personality and Social Psychology Bulletin*, *38*, 1272-1287.

Nurmi, J. E. (2005). Thinking about and acting upon the future: Development of future orientation across the life span. In A. J. Strathman & J. A. Joireman (Eds.), *Understanding behavior in the context of time: Theory, research, and application* (pp. 31-57). Mahwah, NJ: Lawrence Erlbaum.

O'Connor, B., Balasubramanyan, R., Routledge, B. R., & Smith, N. A. (2010). From tweets to polls: Linking text sentiment to public opinion time series. *Fourth International AAAI Conference on Weblogs and Social Media*, *11*, 122-129.

Ouellette, J. A., Hessling, R., Gibbons, F. X., Reis-Bergan, M., & Gerrard, M. (2005). Using images to increase exercise behavior: Prototypes versus possible selves. *Personality and Social Psychology Bulletin*, *31*, 610-620.

Owoputi, O., O'Connor, B., Dyer, C., Gimpel, K., & Schneider, N. (2012). Part-of-speech tagging for Twitter: Word clusters and other advances. Retrieved from http://www.ark.cs.cmu.edu/TweetNLP/owoputi+etal.tr12.pdf

Park, G., Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Kosinski, M., Stillwell, D. J., Ungar, L. H., & Seligman, M. E. P. (2015). Automatic personality assessment through social media language. *Journal of Personality and Social Psychology, 108,* 934-952. http://dx.doi.org/10.1037/pspp0000020

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O.,... Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830.

Pennebaker, J. W., Chung, C. K., Ireland, M., Gonzales, A., & Booth, R. J. (2007). *The development and psychometric properties of LIWC2007.* Austin, TX: LIWC.net.

Pennebaker, J. W., Mehl, M. R., & Niederhoffer, K. G. (2003). Psychological aspects of natural language use: Our words, our selves. *Annual Review of Psychology*, *54*, 547-577. http://dx.doi.org/10.1146/annurev.psych.54.101601.145041

Piko, B. F., & Brassai, L. (2009). The role of individual and familial protective factors in adolescents' diet control. *Journal of Health Psychology*, *14*, 810-819.

Potts, C. (2011). *happyfuntokenizer* (Version 1.0) [computer software]. Retrieved from: http://sentiment.christopherpotts.net/code-data/happyfuntokenizing.py

Preis, T., Moat, H. S., Stanley, H. E., & Bishop, S. R. (2012). Quantifying the advantage of looking forward. *Nature: Scientific Reports*, *2*, 350.

Radloff, L. S. (1977). The CES-D Scale: A self-report depression scale for research in the general population. *Applied Psychological Measurement, 1*, 385-401

Roberts, B. W., Harms, P. D., Smith, J. L., Wood, D., & Webb, M. (2006). Using multiple methods in personality psychology. In M. Eid & E. Diener (Eds.), *Handbook of* *psychological assessment: A multimethod perspective* (pp. 321–335). Washington, DC: American Psychological Association.

Roberts, B. W., Lejuez, C., Krueger, R. F., Richards, J. M., & Hill, P. L. (2014). What is conscientiousness and how can it be assessed? *Developmental Psychology, 50*, 1315-1330. doi: 10.1037/a0031109

Rothspan, S., & Read, S. J. (1996). Present versus future time perspective and HIV risk among heterosexual college students. *Health Psychology*, *15*, 131-134.

Sap, M., Park, G., Eichstaedt, J. C., Kern, M. L., Stillwell, D. J., Kosinski, M., … & Schwartz, H. A. (2014). Developing age and gender predictive lexica over social media. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1146-1151.

Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Agrawal, M., Park, G., ... & Lucas, R. E. (2013a, June). Characterizing geographic variation in well-being using tweets. In *Seventh International AAAI Conference on Weblogs and Social Media*, Boston, MA.

Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., …, & Ungar, L. H. (2013b). Personality, gender, and age in the language of social media: The open vocabulary approach. *PLOS ONE*, *8*, e73791. http://dx.doi.org/10.1371/journal.pone.0073791

Schwartz, H. A., Park, G., Sap, M., Weingarten, E., Eichsteadt, J. C., Kern, M. L., … & Ungar, L. H. (2015). Extracting human temporal orientation from Facebook language. *North American Chapter of the Association for Computational Linguistics (NAACL), Human Language Technologies*, Denver, CO.

Sedikides, C., Wildschut, T., Arndt, J., & Routledge, C. (2008). Nostalgia: Past, present, and future. *Current Directions in Psychological Science*, *17*, 304-307.

Stanford, M. S., Mathias, C. W., Dougherty, D. M., Lake, S. L., Anderson, N. E., & Patton, J. H. (2009). Fifty years of the Barratt Impulsiveness Scale: An update and review. *Personality and Individual Differences*, 47, 385-395. doi: http://dx.doi.org/10.1016/j.paid.2009.04.008

Steinberg, L., Graham, S., O'Brien, L., Woolard, J., Cauffman, E., & Banich, M. (2009). Age differences in future orientation and delay discounting. *Child Development*, *80*, 28-44.

Strathman, A., Gleicher, F., Boninger, D. S., & Edwards, C. S. (1994). The consideration of future consequences: Weighing immediate and distant outcomes of behavior. *Journal of Personality and Social Psychology*, *66*, 742-752. doi: 10.1037/0022-3514.66.4.742

Toutanova, K., Klein, D., Manning, C. D., & Singer, Y. (2003, May). Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology*, Edmonton, Canada.

Webley, P., & Nyhus, E. K. (2006). Parents' influence on children's future orientation and saving. *Journal of Economic Psychology*, *27*, 140-164.

Zhang, J. W., & Howell, R. T. (2011). Do time perspectives predict unique variance in life satisfaction beyond personality traits? *Personality and Individual Differences*, *50*, 1261-1266.

Zimbardo, P. G., & Boyd, J. N. (1999). Putting time in perspective: A valid, reliable individual-differences metric. *Journal of Personality and Social Psychology*, *77*, 1271-1288.

Footnotes

[1]In order to replicate how the model would be applied to the final test set, judges were not given the option to rate something as non-interpretable. We used forced-choice here because, when applying the model to messages, we cannot remove or exclude messages from classification, and this gives a more realistic assessment of how the classifier functions on a new set of text.

[2]While we selected the ERT model on the basis of test set performance, we also checked the ERT model accuracy in the training sample using 10-fold cross-validation. The average accuracy of the full ERT model over the training sample was 0.68.

[3]Although the MyPersonality sample includes participants older than 48, the sample size drops steeply with every year, and many of these users do not meet the other requirements (e.g., wrote at least 100 messages). Thus, we only included bins up to age 48.