Title
Example Generation Under Constraints Using Cascade Correlation Neural Nets

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors
Nobandegani, Ardavan S
Schultz, Thomas R

Publication Date
2018

Cognition and Emotion in Narratives of Redemption: An Automated Analysis

Eyal Sagi (esagi@stfrancis.edu)
Department of Psychology, 500 Wilcox St.
Joliet, IL 60435 USA

Brady K. Jones (bjones@stfrancis.edu)
Department of Psychology, 500 Wilcox St.
Joliet, IL 60435 USA

Abstract

Redemptive narratives are stories of challenge, failure, or adversity that in some way acknowledge the goodness or personal growth that came of the recounted difficult event. In this paper we use a corpus-statistic based approach to explore the role of cognition and emotion in these narrative arcs. In particular, we trace the shift from negative to positive sentiment (a change in the emotional valence) and vice to virtue (evidence of cognitive, moral processing) within the narrative. Our results suggest that cognitive processes, more than emotion, drive the shift to goodness and growth that is at the core of redemptive narratives. We discuss the implications of these results to both narrative psychology and cognitive psychology.

Keywords: Narrative analysis; Redemption; Latent Semantic Analysis; Sentiment analysis; Cognitive processes; Affective processes

Introduction

Narrative psychology explores the ways people tell the stories of important events in their lives and details the implications of those stories. Research in this field has linked the tendency for individuals to include particular narrative themes in their life stories to a host of consequential outcomes: prosocial behavior (McAdams, 2006), physical and psychological health (Adler, 2012; Manczak, Zapata-Gietl, & McAdams, 2014), political beliefs and behaviors (Hammack, 2006; McAdams et al., 2008; Rotella, Richeson, & McAdams, 2015), and academic performance (Jones, Destin, & McAdams, 2018) to name just a few. In this paper, we use a statistical approach to the analysis of text, based in Latent Semantic Analysis (Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990; Landauer, Foltz, & Laham, 1998) to explore cognitive factors found in such narratives and trace the arc through which the narrator develops their story.

Redemption is one of the most common and powerful themes narrative psychologists study. Redemptive themes are most frequently found in stories about low moments in individuals' pasts: personal tragedies, major failures, and difficult situations. When these stories include an acknowledgment of the goodness or growth that came out of the experienced adversity, they are considered redemptive. Because redemptive narratives show a distinctive arc – from adversity to goodness and growth – they can be thought of as following a *redemption sequence*.

Redemptive sequences can follow a variety of patterns. McAdams (2006) details six "languages of redemption" he argues are common in Americans' life stories: atonement (sin to forgiveness), emancipation (slavery to freedom), upward mobility (poverty to wealth), recovery (sickness to health), enlightenment (ignorance to knowledge), and development (immaturity to actualization). In all cases, the redemptive sequence involves two shifts that are believed to be the result of reflection and reevaluation that have taken place over time. The first is emotional; the tone of the narrative moves from negative to positive. The second is cognitive. The narrator engages in reasoning around the negative experience in a way that allows them to explain how what was once a weakness has become a strength or virtue.

The tendency to employ redemptive sequences when recounting life events is associated with a number of positive outcomes. To give just a few examples, adolescents who can name ways they learned from a personal failure show more persistence and get better grades (Jones et al., 2018), convicts who tell their stories in more redemptive terms are less likely to fall back into crime (Maruna, 2001), and midlife adults who use many redemptive sequences when they tell the stories of their lives are more meaningfully integrated into their communities and behave in more prosocial ways (McAdams, 2006).

Linguistic analysis of narrative themes

Because narratives are frequently available as texts, psychologists have also explored the possibility of using automated text analysis as a method of relating aspects of some narratives to associated outcomes. In one such influential account, Rude, Gortner, and Pennebaker (2004) demonstrate that depressed individuals are more likely to use negatively valenced words, as well as the word "I". While formerly depressed individuals still used "I" in higher prevalence than never-depressed individuals, they no longer exhibited an increased use of negatively valenced words.

More recently, Weston et al. (2016) compared predictions derived from word counting with the performance of human coders. They found that while word counts accounted for some variability, they were not consistently correlated with the human coder's identification of redemptive narratives. Moreover, word counting, while commonly used in studies of this type, is likely to be of limited utility for identifying narrative arcs, such as redemptive sequences. A better approach to analyzing such narrative arcs is to divide a given narrative into small units and measure the relevant concepts for each unit. The arc can then be observed as a shift in those measurements across the course of the narrative. When the analyzed narrative is short (e.g., the narratives described in this study, which average less than 300 words), the chances of finding the key words in each unit necessary to successfully employ the word counting method become vanishingly small. To effectively explore such narrative arcs, it is important to be able to use as many of the words found in the text as possible

Latent Semantic Analysis provides one avenue for such investigation because it is able to use patterns of word co-occurrence to provide a graded measure of similarity rather than the binary (found/not found) measure that underlies methods of word counting, such as Linguistic Inquiry and Word Count (LIWC, Tausczik & Pennebaker, 2010). In particular, LSA-based methods have been useful in exploring other temporal progressions in text, such as the formation of coalitions during negotiations (Sagi & Diermeier, 2017).

Linguistic Expression of Redemption Sequences

Our aim in this paper is to use automated text analysis to examine the properties of redemptive narratives. To this end, we need to identify the textual components that construct redemption sequences. As mentioned earlier, one of two key hallmarks of such sequences is the transition from negative to positive valence. The narrative begins by highlighting negative events and closes on a positive note.

The question of whether an utterance is positive or negative has a long history in the natural language processing literature. In particular, many studies have focused on identifying the sentiments expressed by particular texts (e.g., Hu & Liu, 2004; Kim & Hovy, 2004; cf. Pang & Lee, 2008 for a review). Borrowing from this literature, we will use lists of positively and negatively loaded terms to measure the emotional valence of sentences.

The concept of personal growth is the second hallmark of redemptive sequences. This concept suggests that the content of the narrative not merely expresses a sentiment but describes a cognitive process through which the narrator analyzed the events and reached conclusions. Such processes often employ reasoning about the causes and consequences of events, as well as the intentions of the individuals involved. In addition, the cognitive processes employed when people tell redemptive narratives usually involve themes of morality or the development of virtue. Individuals do cognitive work to transform what was once considered *bad* – some weakness, disadvantage, or shortcoming – into something *good* – a hard-won and valuable personal strength. In fact, a robust body of research suggests the individuals most likely to tell redemptive narratives are those who are exceptionally prosocial (e.g., Maruna, 2001; McAdams, 2006; McAdams & Guo, 2015).

To represent this aspect of the narratives, then, we will build on Haidt's Moral Foundation Theory (Graham et al., 2013; Haidt & Joseph, 2004) which identifies 5 different moral concerns (Harm, Fairness, Loyalty, Authority, and Purity) and argues that each is associated with a different style of reasoning. Each of these concerns has positive and negative aspects (e.g., Caring vs. Harming, Loyalty vs. Betrayal), or, in other words, aspects of virtue and vice. In redemptive narratives, we will expect virtue to predominate late in the narrative, when the narrator concludes on a positive note of personal growth.

Method

Hand Coding of Redemption in Study Dataset

McAdams pioneered the ways in which redemptive narratives have been coded for the past several decades (McAdams, Reynolds, Lewis, Patten, & Bowman, 2001), and other researchers have modified and simplified this coding procedure to fit their needs and research questions (e.g., Jones et al., 2018). For this study, we used narratives from 97 K-12 teachers and former teachers who graduated from an elite teacher preparation program. They read commonly used prompts in life story interviews and responded – in written form, through a computer survey – by recounting three difficult life events, including a life low point, a career low point, and career failure or regret (Jones, 2018). To code for redemption sequences, two coders first established inter-rater reliability. Blind to identifying information about the participants, they worked together to code the three difficult events the teachers narrated for any sense of goodness or growth that resulted from the adversity they experienced. Each narrative account received a raw score for the number of redemption sequences found. After several rounds of preliminary coding, the two coders attained an intraclass correlation coefficient (two-way consistency model, single measures) of .836. Once reliability was achieved, the narratives were split up and each coder took half. During the preliminary coding stage, all disagreements were settled by discussion. Once the two coders had split the narratives in half they re-contacted each other to discuss narratives that were particularly difficult to code. In the present study we focus exclusively on the life low point narratives as they are the narratives that have been traditionally explored in most studies of redemption.

Automated Text Analysis

The textual measure of redemptive language we use in this paper is based on the conceptual framing index approach developed by Sagi, Diermeier, and Kaufmann (2013; see also Sagi, 2018). At its core, this measure uses Latent

Semantic Analysis cosine similarity to measure the similarity between target utterances (i.e., sentences) and key terms (e.g., positive words such as *victory* and *improve*; negative words such as *danger* and *conflict*). The overall measure of positivity can then be computed as a mean of the cosine similarity of the positive terms to the target sentence.

Similar measures, based on LSA, have been used to explore a variety of psychological and linguistic phenomena in texts, including semantic priming (Chwilla & Kolk, 2002), textual coherence (Foltz, Kintsch, & Landauer, 1998), the representation of moral concerns (Sagi & Dehghani, 2014), knowledge acquisition (Wolfe et al., 1998), and coalition formation (Sagi & Diermeier, 2017).

Vectors representing individual words were generated based on the co-occurrence patterns of words in the written portion of the British National Corpus (BNC Consortium, 2007) using Infomap (2007; Schütze, 1997; Takayama, Flournoy, Kaufmann, & Peters, 1998). Each vector represents the position of a word in a 100 dimension space and vectors representing whole utterances (e.g., sentences) are computed by means of vector addition (i.e., averaging).

As mentioned above, the similarity between vectors can be measured by measuring the cosine of the angle between the vectors. In the case of unit-length vectors (as were used here), this cosine measure is mathematically equivalent to the correlation between the vectors. Consequently, it is useful to think of this measure as based on the correlation between the co-occurrence patterns of words.

Textual Measures of Redemptive Language

As discussed earlier, the arc of a redemptive narrative transforms a negative experience into a positive outlook. Consequently, we measured both positive and negative aspects of language. Furthermore, we hypothesized that both emotion (positivity/negativity) and cognition (virtue/vice) contribute to this process. This lead us to develop 4 independent measures of language: Positive sentiment, negative sentiment, positive evaluation, negative evaluation.

The words for measures of positive and negative sentiment were taken from a comprehensive collection of such terms originally compiled by Hu and Liu (2004). We identified all the words in the list that occurred at least 5000 times in the BNC. This resulted in 85 positive terms and 63 negative terms.

Our measures of cognitive evaluation were based on the Moral Foundations Dictionary (Graham, Haidt, & Nosek, 2009). Moral Foundations Theory recognizes 5 moral concerns (Harm, Fairness, Loyalty, Authority, and Purity) which represent different styles of reasoning about morality. Importantly, the dictionary identifies both positive (virtue) and negative (vice) polarity words related to each of the foundations. Since redemptive sequences often involve not only positive emotion but also personal growth, using words related to vices and virtues provides us with a good measure of the negative and positive cognitive evaluations present in the narrative. Because of the lower frequency of words in the MFD, we collected all the terms from the MFD that occurred at least 250 times in the BNC. This resulted in 128 vice terms and 201 virtue terms.

To simplify the overall analysis, we combined the 4 measures into a single measure, representing the *redemptiveness index* of each sentence. This index is computed as an average of the correlation of all the terms from the 4 lists, with the correlations of each sentence with terms from the negative lists (vice and negative sentiment) reversed (i.e., a positive correlation becomes negative and a negative correlation becomes positive). Essentially, this measure is higher for sentences that are positively correlated with positive terms and negatively correlated with negative terms and lower for sentences that are positively correlated with negative terms and negatively correlated with positive terms. As a combined measure, we expect redemptive narratives to show an increase on the redemptiveness index as they proceed (and especially at their end).

Results

The 97 narratives are of considerably varied length, with a median length of 12 sentences. Most of the stories (80%) are between 6 and 23 sentences long. Forty-six narratives were hand coded as redemptive (47.4%) and 51 were coded as non-redemptive.

We computed aggregate vector for each sentence in the narratives. Of the 4,099 sentences, 65 (1.6%) included no content words (e.g., "I got in") and were therefore not associated with a vector. These sentences were discarded and not used in the analysis. Next, we calculated the cosine between each sentence vector and the word vectors for terms from all 4 valence dimensions (Positivity, Negativity, Virtue, and Vice). We average the combined correlation as described earlier to produce our *redemptiveness index*. The overall results plotted by quintiles to control for the length of the narrative are presented in Figure 1.

To test our primary hypothesis, that the redemptiveness index will increase over the course of redemptive narratives but not non-redemptive ones, we use a general linear model, with the redemptiveness index as the dependent variable. The redemptiveness of the narrative (redemptive or non-redemptive) and the position of the sentence in the narrative (expressed in percentage to control for length) were the independent variables. The participant was included as a random effect.

The model identified no effect of redemptiveness overall, $F(1, 1330) = 1.02$, *n.s.*. There was an overall increase in positivity over the course of all narratives, $F(1, 1330) = 4.1$, $p < .05$. More importantly, the predicted interaction was observed. Redemptive narratives showed a greater increase in positivity than non-redemptive narratives, $F(1, 1330) = 8.16$, $p < .005$.

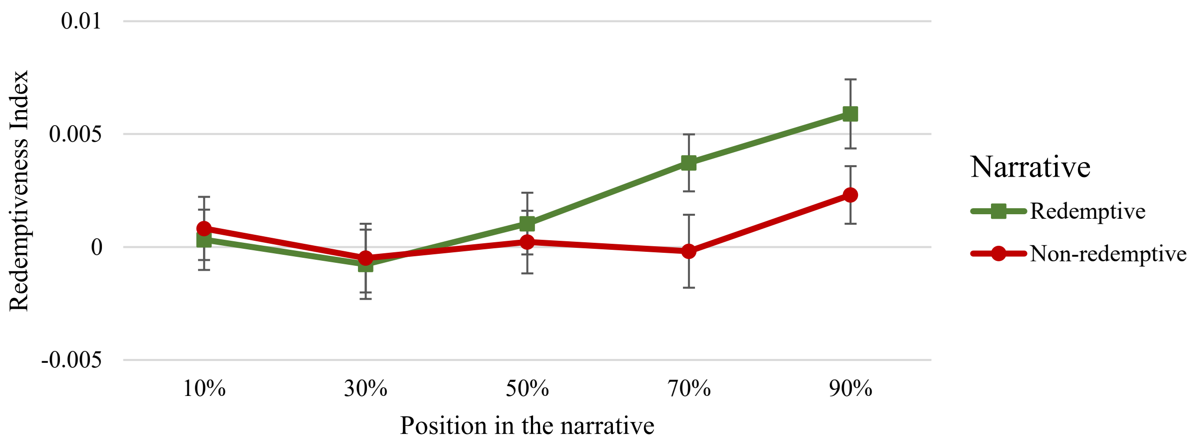

Figure 1: The mean redemptiveness index of sentences as a function of their position in the narrative. Higher values indicate more overall positivity. The actual numeric values of the index have no intrinsic meaning. Error bars represent standard error of the mean.

Similar results were obtained from a more comprehensive model, which separated the measures. The measured correlation was the dependent variable in this model, while the valence of measure (positive or negative) and its source domain (emotion or cognition) were included as additional independent variables. The primary interaction of interest, between redemptiveness, position in the narrative, and valence, was statistically significant, $F(1, 5608) = 5.78$, $p < .05$. Simplified analyses, separating the effect by domain, showed the three-way interaction for the cognitive domain ($F(1, 2756) = 6.10$, $p < .05$), but not for the domain of emotions ($F(1, 2756) = 1.89$, *n.s.*), suggesting that cognition is more important than emotion in defining a redemptive narrative arc.

It is also of interest to examine how each of the measures relates to the narrative arc independently. To explore this, we conducted a series of 4 analyses similar to the one we conducted on the overall redemptiveness index (one for each measure). Interestingly, only the negative measures significantly interacted with redemptiveness in this analysis (Negativity: $F(1, 1330) = 14.83$, $p < .001$; Vice: $F(1, 1330) = 10.62$, $p < .005$).

We now turn to examining how these measures work within individual narratives. For this, it is beneficial to examine the correlation of the redemptiveness index with each narrative separately. To better control the overall correlation, we divided each narrative into 5 parts of equal length (as measured by sentences). We then computed the correlation of the progression of the narrative with the average index of each part. Because of excluded sentences, 4 non-redemptive narratives did not have an index in all 5 parts and were excluded from the analysis.

We now turn to examining how these measures work within individual narratives. For this, it is beneficial to examine the correlation of the redemptiveness index with each narrative separately. To better control the overall correlation, we divided each narrative into 5 parts of equal length (as measured by sentences). We then computed the correlation of the progression of the narrative with the

I think the worst time in my life was while I was doing my master's. I lost my mother a month after I started, and the workload and financial worries added on to make things even worse. Midway into the program, the housing market and Lehman brothers tanked, and I could then look forward to graduating with scads of debt into a job market that would be laying off teachers. This time was horrible because it was so overwhelming. I couldn't really process my thoughts and feelings, I spent a lot of time drinking in front of the TV, zoning out. My hair fell out, I gained weight, and getting to sleep at night was difficult at best. This time tested everything in me; the support of my significant other and my friends was the only thing that kept me going. I think I've become a lot more pessimistic and cynical since then, but also more

Figure 2: A redemptive narrative. Each sentence is color coded based on its redemptiveness index (low redemptiveness indices in red, high redemptive indices in green).

average index of each part. Because of excluded sentences, 4 non-redemptive narratives did not have an index in all 5 parts and were excluded from the analysis.

The average correlation for the redemptive narratives was 0.29 whereas the average correlation for the non-redemptive narratives was 0.00. We further investigated this pattern by separating narratives with positive correlation from those with negative correlations. While the non-redemptive narratives were evenly split (23 positive and 24 negative), there were substantially more positive correlations observed for the redemptive narratives (33 positive and 13 negative). This was confirmed by a significant chi-square test, $\chi^2(1) = 5.05$, $p < .05$.

In addition to an overall analysis, it is possible to examine the arc of specific narratives and visually inspect how they progress. Figure 2 presents a redemptive narrative that is color coded based on the index. This example highlights how a narrative that begins very negatively (due to a confluence of events such as death and financial issues) can transition to a positive, redemptive, arc (how the support of friends can help an individual overcome obstacles).

Interestingly, the third sentence in the narrative would probably be regarded as fairly negative but has a relatively high redemptiveness index. This demonstrates one of the limits of the particular method we describe in this paper. Like most of the other available methods of text analysis, we ignore syntax and word order and focus on the words in isolation. These so-called "bag of words" approaches, while useful, emphasize the meaning of words over the structures in which they appear. In this case, the sentence uses some sarcasm (e.g. "look forward to graduating with scads of debt") which employs several positive words to convey a negative meaning. This negative meaning is overlooked by the statistical analysis, resulting in a higher index for the third sentence than would be expected.

Discussion

In this paper we described an automated method, based in corpus statistics, that can be used to trace narrative arcs in personal stories. We constructed a theoretically-driven measure of redemptiveness, defining a redemptive sequence as one in which both emotional and moral transformations take place. In other words, we hypothesized that in a redemptive sequence, individuals would (1) move from negative to positive affect and (2) engage in cognitive processes that indicated personal, growth, or a progression from vice to virtue. We found that our automated method tracing such changes in emotion and cognition over the course of the narratives predicted hand coded scores for redemption.

Our results also suggest that reasoning, more so than affect, is responsible for providing the positive outlook that is the hallmark of redemptive narratives. The individuals who redemptively narrated low points in this study, then, did not describe simple, serendipitous progressions from bad life events to good, lucky ones. Instead, they seemed to engage in a kind of moral reasoning as they process their low points, ending the narrative account as wiser and more virtuous people. This is an important result because it highlights the role of cognitive processes in redemption sequences, which may seem, on their face,to stem from an emotion-driven process. Moreover, it suggests that personal growth following adversity can be linked with particular styles of reasoning that are related to morality. This link provides an intriguing avenue for exploring the role of adversity and redemption in driving cognitive processes that affect future behaviors.

Interestingly, when examining each of the component measures on their own, we found that the negative measures were the primary drivers of the results. While somewhat surprising, it is not entirely unexpected. Narratives of redemption are predominantly stories of negative events and most frequently culminate in a positive outlook at their very end. Consequently, there is much more negativity in such stories than positivity.

Nevertheless, this suggests that the measure we presented here can be improved upon by taking such effects into account. In particular, a linear regression assumes that change is constant over the course of measurement (in this case, the course of the narrative) whereas a redemption sequence is expected to have a sharp transition from negative to positive towards the end of the narrative. Incorporating this information into the analysis should increase its statistical power and effectiveness.

Redemptive themes in narratives are important predictors of psychological health and positive adjustment in society, but coding for them is difficult and time-intensive. In this paper we presented a theoretically-driven method for identifying such sequences. Even though this method cannot, by itself, replace the manual coding process, it provides insight into the linguistic and cognitive factors that underlie redemptive sequences. It is therefore a promising avenue of investigation that can, in the future, lead to the development of systems that can identify similar sequences and arcs without the need for manual coding. This will provide researchers in the field with much larger sets of texts and add to our growing understanding of the importance of personal stories in peoples' lives.

References

Adler, J. M. (2012). Living into the story: Agency and coherence in a longitudinal study of narrative identity development and mental health over the course of psychotherapy. *Journal of Personality and Social Psychology*, *102*(2), 367–389.

BNC Consortium. (2007). The British National Corpus, version 3 (BNC XML Edition). Oxford University Computing Services. Retrieved from http://www.natcorp.ox.ac.uk/

Chwilla, D. J., & Kolk, H. H. J. (2002). Three-step priming in lexical decision. *Memory & Cognition*, *30*(2), 217–225. https://doi.org/10.3758/BF03195282

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, *41*(6), 391–407. https://doi.org/10.1002/(SICI)1097-4571(199009)41:6<391::AID-ASI1>3.0.CO;2-9

Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). The measurement of textual coherence with latent semantic analysis. *Discourse Processes*, *25*(2–3), 285–307. https://doi.org/10.1080/01638539809545029

Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S., & Ditto, P. (2013). Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in*

Experimental Social Psychology, *47*, 55–130. https://doi.org/10.1177/0963721412456842

Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, *96*(5), 1029–1046. https://doi.org/10.1037/a0015141

Haidt, J., & Joseph, C. (2004). Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus*, *133*(4), 55–66.

Hammack, P. L. (2006). Identity, conflict, and coexistence: Life stories of Israeli and Palestinian adolescents. *Journal of Adolescent Research*, *21*(4), 323–369.

Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 168–177). ACM.

Infomap [Computer Software]. (2007). Stanford, CA. Retrieved from http://infomap-nlp.sourceforge.net/

Jones, B. K. (2018). A special kind of ambition: The role of personality in the retention of academically elite teachers. *Teachers College Record*, *120*(9).

Jones, B. K., Destin, M., & McAdams, D. P. (2018). Telling better stories: Competence-building narrative themes increase adolescent persistence and academic achievement. *Journal of Experimental Social Psychology*, *76*, 76–80.

Kim, S.-M., & Hovy, E. (2004). Determining the sentiment of opinions. In *Proceedings of the 20th international conference on Computational Linguistics* (p. 1367). Association for Computational Linguistics.

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). An introduction to latent semantic analysis. *Discourse Processes*, *25*(2–3), 259–284.

Manczak, E. M., Zapata-Gietl, C., & McAdams, D. P. (2014). Regulatory focus in the life story: Prevention and promotion as expressed in three layers of personality. *Journal of Personality and Social Psychology*, *106*(1), 169–181.

Maruna, S. (2001). *Making Good: How Ex-Convicts Reform and Rebuild Their Lives*. Washington, DC: American Psychological Association Books.

McAdams, D. P. (2006). *The redemptive self: Stories Americans live by*. New York, NY: Oxford University Press.

McAdams, D. P., Albaugh, M., Farber, E., Daniels, J., Logan, R. L., & Olson, B. (2008). Family metaphors and moral intuitions: How conservatives and liberals narrate their lives. *Journal of Personality and Social Psychology*, *95*(4), 978–990.

McAdams, D. P., & Guo, J. (2015). Narrating the Generative Life. *Psychological Science*, *26*(4), 475–483. https://doi.org/10.1177/0956797614568318

McAdams, D. P., Reynolds, J., Lewis, M., Patten, A. H., & Bowman, P. J. (2001). When bad things turn good and good things turn bad: Sequences of redemption and contamination in life narrative and their relation to psychosocial adaptation in midlife adults and in students. *Personality and Social Psychology Bulletin*, *27*(4), 474–485.

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, *2*(1–2), 1–135.

Rotella, K. N., Richeson, J. A., & McAdams, D. P. (2015). Groups' search for meaning: Redemption on the path to intergroup reconciliation. *Group Processes & Intergroup Relations*, *18*(5), 696–715.

Rude, S., Gortner, E.-M., & Pennebaker, J. (2004). Language use of depressed and depression-vulnerable college students. *Cognition & Emotion*, *18*(8), 1121–1133.

Sagi, E. (2018). Developing a New Method for Psychological Investigation Using Text as Data. *SAGE Research Methods Cases*. https://doi.org/10.4135/9781526442604

Sagi, E., & Dehghani, M. (2014). Measuring Moral Rhetoric in Text. *Social Science Computer Review*, *32*(2), 132–144. https://doi.org/10.1177/0894439313506837

Sagi, E., & Diermeier, D. (2017). Language use and coalition formation in multiparty negotiations. *Cognitive Science*, *41*(1), 259–271.

Sagi, E., Diermeier, D., & Kaufmann, S. (2013). Identifying Issue Frames in Text. *PloS One*, *8*(7), e69185.

Schütze, H. (1997). *Ambiguity resolution in language learning: computational and cognitive models*. Stanford, CA: CSLI Publications.

Takayama, Y., Flournoy, R., Kaufmann, S., & Peters, S. (1998). *Information mapping: Concept-based information retrieval based on word associations*. Stanford, CA: CSLI Publications.

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, *29*(1), 24–54.

Weston, S. J., Cox, K. S., Condon, D. M., & Jackson, J. J. (2016). A comparison of human narrative coding of redemption and automated linguistic analysis for understanding life stories. *Journal of Personality*, *84*(5), 594–606.

Wolfe, M. B. W., Schreiner, M. E., Rehder, B., Laham, D., Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). Learning from text: Matching readers and texts by latent semantic analysis. *Discourse Processes*, *25*(2–3), 309–336. https://doi.org/10.1080/01638539809545030