

Identifying Locus of Control in Social Media Language

Masoud Rouhizadeh^{*1} Kokil Jaidka^{*2} Laura Smith³
H. Andrew Schwartz⁴ Anneke Buffone³ Lyle H. Ungar²

¹Institute for Clinical & Translational Research, Johns Hopkins University

²Department of Computer & Information Science, University of Pennsylvania

³Department of Psychology, University of Pennsylvania

⁴Department of Computer Science, Stony Brook University

mrou@jhu.edu, jaidka@sas.upenn.edu

Abstract

Individuals express their locus of control, or “control”, in their language when they identify whether or not they are in control of their circumstances. Although control is a core concept underlying rhetorical style, it is not clear whether control is expressed by *how* or by *what* authors write. We explore the roles of syntax and semantics in expressing users’ sense of control –i.e. being “controlled by” or “in control of” their circumstances– in a corpus of annotated Facebook posts. We present rich insights into these linguistic aspects and find that while the language signaling control is easy to identify, it is more challenging to label it is internally or externally controlled, with lexical features outperforming syntactic features at the task. Our findings could have important implications for studying self-expression in social media.

1 Introduction

Language is a form of action, and written occurrences constitute a performance of power and control (Fairclough, 2001). Research in natural language processing has long focused on distilling the constituents of writing that conveys authority (Danescu-Niculescu-Mizil et al., 2012), dominance (Bradley and Lang, 1999), dogmatism (Fast and Horvitz, 2016), expertise (Levy and Potts, 2015) and politeness (Danescu-Niculescu-Mizil et al., 2012). These studies have shown how authors’ use of certain lexical and syntactic patterns achieve specific rhetorical effects. In this study, we contribute to the growing literature with an analysis of how individuals express *control*, and compare the insights and predictive power obtained from both, lexical and syntactic features.

We operationalize the key aspects of locus of control described by existing psychology theories

(Rotter, 1966) to identify an external locus of control when authors express that they feel controlled by other people or the environment. On the other hand, authors communicate an internal locus of control when they ascribe the control of their decisions and circumstances to themselves. This study poses the question: *do writers rely more on content than syntax to convey their control?* Following its psychological underpinnings, we expect that overall lexical choice, self-reference, and certain verb categories would be indicative features for modeling author’s locus of control. Our study attempts to validate these assumptions by answering two research questions:

- To what extent can *lexical* and *syntactic* features predict *control relevance* and *internal and external* locus of control?
- Which verb categories are more associated with control relevance and internal vs. external locus of control?

We train a multi-stage predictive model to classify posts on Facebook as (a) control-relevant or not, and (b) internal vs. external control. We find that syntactic and semantic features work well to identify control relevance, with verb categories and pronoun use being the most predictive features. However, determining internal vs. external control is a more challenging task, where lexical features vastly outperform syntax-based features. The best performance is seen by a model that combines lexical and syntactic features along with self-reference ratio. We will be releasing our tools and models along with annotated and anonymized datasets for the benefit of the research community.

2 Background

Locus of control, or “control”, reflects the extent to which people ascribe the cause or control of events in their lives to themselves or the external

^{*}Masoud Rouhizadeh & Kokil Jaidka co-lead this work.

factors (Rotter, 1966). The language indicating internal control in a given situation signals intentionality (the author is describing an action that he/she intended) and awareness (the author is aware of the effect of the action). Internal control is often associated with causing a given event or doing something that is clearly a choice. The external control language, on the other hand, is characterized by lack of intention or awareness, or by concrete mention of being controlled by others. It is usually lined to describing an out-of-control event, or something that is not a choice

Organizations and governments are attempting to better understand issues of locus of control and self-efficacy, which are closely related with physical health and job performance (Marmot et al., 1991; Harter et al., 2003). The study by (Jaidka et al., 2018b) offered to explore the relationship between locus of control (measured through surveys) and the Big 5 personality taxonomy (John and Srivastava, 1999), and touched briefly upon the content-related linguistic signals. However, the present study is more interested in comparing the lexical (aka open-vocabulary or bag-of-words) and syntactic signals of control that are embedded in language. If deployed on a large scale with the informed consent of the authors, it may allow unobtrusive, cost-effective estimates of well-being¹.

Signals within the language should provide insight into the cognitive sense of control experienced by an individual; however, to our knowledge, no study has examined whether these signals are stronger in the *way* people express themselves, or *what* they say. Although the discourse and lexical signals that convey status and dominance have been studied in previous work, they involve controlling *other* people (Danescu-Niculescu-Mizil et al., 2012; Fast and Horvitz, 2016), while locus of control is an assessment of the writer’s own affairs – perhaps indicated through lexical features that focus on first person pronouns.

Dominance has also been qualitatively measured in terms of the affective context at the level of an utterance (Bradley and Lang, 1999); however, we found that this also did not correspond to our notion of locus of control. With reference to the syntactic features, we will test whether control may be understood through different verb forms, and active versus passive voice. Locus of control

¹<http://www.pewinternet.org/2016/06/22/social-media-and-the-workplace/>

is different from semantic role labeling: for internal control, we are mainly interested in *whether* the author is the *agent* (rather than identifying who the agent is, in a sentence), and how much they control their life. External control includes a complex mixture of semantic roles such as *patient*, *experiencer*, *theme*, *recipient*, and *beneficiary*.

3 Locus of control data

Posts on Facebook capture self-expression by a diverse audience about their daily lives, which makes it a natural starting point for exploring the linguistic features of control. For data collection, we deployed a survey on Qualtrics comprising several demographic questions, the Sense of Control facet items from the MIDUS survey (Brim et al., 2004) and a 3-item health inventory measuring poor health, general health and weekly exercise, taken from the CDCs Behavioral Risk Factor Surveillance System (BRFSS) and the European Social Survey (ESS). We invited users to share access to their Facebook status updates and randomly sampled 2000 users who had posted at least 100 words (with 839 status updates on average). We then randomly sampled 2 status updates per user to obtain 4000 statuses. We cleaned up the text by removing URLs, hashtags, user handles, etc. and split each status into sentences using our in-house pattern-based sentence splitter inspired by CMU ARK Twitter Twokenize script (Owoputi et al., 2013). Finally, from each status update, we selected one sentence that has 2 or more words.

3.1 Locus of control annotation

Building a useful computational model requires labeled training data. We labeled the Facebook dataset using three trained annotators pursuing a Master’s program in Psychology, to construct the first public corpus annotated with control. We asked the annotators to determine whether the author of the sentence is in control (internal control) or being controlled by others or circumstances (external control). To ensure quality work, we provided examples corresponding to each point on the scale. The examples provided to the annotators are also provided in the annotation scheme described in the Appendix.

In this paper, we focus on Stage 1 and Stage 2 of annotation, which distinguish control-relevant sentences and classify them as either internal or external control. We evaluate the reliability of

the annotations by computing percent-agreement. The highest pairwise inter-annotator agreement between annotators is 90.2% for Stage 1 and 79% for Stage 2.

4 Methods

To train and test the classifiers for the Stage 1 and 2, we use the Differential Language Analysis ToolKit² (DLATK) (Schwartz et al., 2017), a Python package developed for social media text analysis. Lexical features are extracted in DLATK and syntax-based features are separately generated and imported into DLATK feature sets.

4.1 Lexical features

ngrams: Normalized frequency distribution over 1-, 2- and 3-grams

LIWC: Frequency distribution of categories from the Linguistic Inquiry and (LIWC) (Pennebaker et al., 2007).

ANEW: Weighted frequency distribution of the Dominance score in the Affective Norms in English Words (ANEW) (Bradley and Lang, 1999).

Word2Vec Topics: We use word clusters built on top of Word2Vec, by Preoțiu-Pietro, Lampos and Aletras (2015). They built 200 open-ended word clusters by applying spectral clustering to the word-to-word similarity matrix from the neural embeddings Mikolov et. al (2013).

Pronouns: We clustered all pronouns (except for possessives) into 1st-, 2nd-, and 3rd-person regardless of their syntactic role.

4.2 Syntax-based features

We acquire dependency parses of our corpus by SyntaxNet (Andor et al., 2016) with Parsey McParseface model³ that produces universal dependencies in *relation, head, dependent* triples in CONLL format. We obtain *subject-verb* tuples (SVs) and *subject-verb-object* triples (SVOs) from the dependency trees. In our in-house evaluation on a random set of 100 Tweets, SyntaxNet with the Parsey McParseface model outperforms the Stanford Parser (Socher et al., 2013) on extracting SVs and SVOs from social media (P=.75, R= .68, TN Rate =.90 for the former; P=.51, R= .55, TN Rate =.80 for the latter). SyntaxNet is also a better tool for our purpose compared to the Tweepo

²<https://dlatk.wwpdb.org>

³<https://github.com/tensorflow/models/tree/master/research/syntaxnet>

Parser (Kong et al., 2014), that only provides dependency graphs and not the relations.

Verb predicate features: We identify the verb classes using (a) linguistically-driven Levin’s Classes (*Lev*) (Levin, 1993), and (b) an in-house manual verb clustering on the most frequent 130 verb into 40 semantic classes (*M*). Inspired by the previous research (Rouhizadeh et al., 2016), we extract five sets of dependency-driven verb predicate features. (1) **Pronouns-SVO:** occurrence of 1st, 2nd, and 3rd person pronouns in the subject/object positions, (2) **VerbCat-1-2-3PP** occurrence of 1st, 2nd, and 3rd person pronouns in the subject/object positions of each verb category (from the Levin’s or our own verb classes), (3) **VerbCat-1PP:** occurrence of 1st or non-1st pronouns (i.e. self-reference ratio) in the subject/object positions of each verb category, (4) **VerbCat-SVO:** all words in the subject/object positions of each verb category, and (5) **VerbCat-all** categories of all the verbs in subject-verb and subject-verb-object contexts.

POS-ngrams: We capture shallow syntactic features by constructing Penn Part-of-Speech (POS) unigrams, bigrams and trigrams after tagging every word with their POS information by using the Python NLTK package (Bird et al., 2009).

5 Results

We train multiple classifiers on different sets of lexical and syntactic features after performing feature selection based on univariate regression between each feature and the outcome. We also perform dimensionality reduction by randomized principal component analysis (PCA). After evaluating a number of classification models available in Python’s sklearn package, we report the results from the logistic regression classifier, with L2 penalization, and inverse of regularization strength set to 0.1, which obtain the best results. Although we have experimented with a variety of feature combinations which are included in the Appendix, we discuss selected classifiers (for meaningful comparisons) as well as the best-performing ones.

5.1 Predicting control relevance

We report the results for prediction performance of the main feature in Table 1a. We see that lexical ngrams are moderately effective, whereas LIWC, ANEW, and w2v features are not helpful. On the other hand, pronouns appear to be very

(a) Control Relevance				(b) Internal/External Control			
	P	R	F1		P	R	F1
Lexical features				Lexical features			
<i>ngram</i>	.854	.918	.885	<i>ngram</i>	.726	.682	.702
<i>Pronouns</i>	.858	1.00	.923	<i>Pronouns</i>	.236	.204	.218
Syntax-based features				Syntax-based features			
<i>VerbCat-all (M)</i>	.878	.997	.934	<i>VerbCat-all (M)</i>	.532	.189	.280
<i>VerbCat-all (Lev)</i>	.903	.997	.948	<i>VerbCat-all (Lev)</i>	.597	.246	.347
<i>POS-ngram</i>	.849	.913	.880	<i>POS-ngram</i>	.635	.688	.659
Best combination				Best combination			
<i>VerbCat-IPP (Lev)</i>	.917	.990	.952	<i>ngram,LIWC,w2v,VerbCat-IPP(Lev)</i>	.700	.737	.718
% Agreement for Human annotators: 90.2				% Agreement for Human annotators: 79.0			

Table 1: Precision, recall and F1-measure for predicting control-relevance and internal/external control (10-fold CV) from lexical and syntactic features.

helpful, despite their linguistic and semantic simplicity. Among the syntax-based features, using the verb categories of SV and SVOs provide the second best F1 score (Levin’s classes are better than our classes), and adding self-reference ratio to verb categories generates the best results. Interestingly, POS-ngrams perform on-par with the lexical ngrams, suggesting that the encoded syntactic information in pos-sequences is just as good as lexical information in ngrams with considerably less sparsity and more efficient computation.

5.2 Predicting Internal/External control

The best results for classifying internal vs. external control are in Table 1b. Unlike control relevance, ngrams are the most helpful features here and pronouns and verb category features result in a poor performance. LIWC, ANEW, and w2v do not noticeably improve the results when combined with ngrams, but adding verb categories and self-reference ratio to this combination creates the best result. This suggests that identifying internal/external control benefits from an ensemble of syntactic, semantic and lexical features and simple lexical features are more helpful than word embeddings. Similar to control relevance, POS-ngrams are helpful here although they are not as good as lexical ngrams.

5.3 Linguistic insights

Table 2 shows the most predictive verb categories (with subject or subject-object arguments) based on logistic regression coefficient. We see that not only does the magnitude of the relationship change, but it is possible that the direction can completely change (positive values are control relevance in Table 2a and external control in Table 2b. Interestingly, we see that verbs for audio/visual activities, demand, and start/end of pro-

cesses, are more frequently used in control-related contexts, whereas using love and like verbs is an informative signal that a given sentence is not related to control. In addition, verbs of cognition, missing, feeling, and hope are positively associated with being controlled by others or the environment, whereas verbs indicating attempt are correlated with author’s control of the situation. Although not complete, these semantic categories are intuitively well-associated with the psychological theory of the control concept. In the Appendix, we also identify the POS ngram patterns which are significantly correlated with both, control relevance and internal vs. external control.

6 Error Analysis

In observing the instances of misclassification, we observed the following general patterns:

- **False negatives:**
 - Control-relevant imperative sentences were frequently misclassified as control-irrelevant, e.g. “have a good day everyone”; “Find Jesus.”; “Stay focused!”
 - Internal LoC sentences with modals and ‘get’ were misclassified as external LoC, e.g., “Finally got a best man!”; “Join me and you can save money on all you purchase here.”
- **False positives:**
 - Control-irrelevant sentences with emotion verbs (“like”, “hate”) were often misclassified as control-relevant, e.g., “I Hate SNOW!” ; “I dont like cologne, but it smells nice.”
 - External LoC sentences with possessive pronouns and the verb “make” were misclassified as internal LoC, e.g., “It made my day.”; “Is your name math because i have a problem with you.”

(a) Control Relevance			(b) Internal/External Control		
Verb category	r	Verbs	Verb category	r	Verbs
Audio-visual	.15**	see, say, look, talk, hear, speak	Cognition	.08**	know, think, mean, understand
Demand	.10**	need, want, ask	Missing	.07**	miss
Start/end	.08**	start, stop, follow, leave	Feeling	.06**	feel
Auxiliary	.07**	be, have, do	Auxiliary	.06**	be, have, do
Buying	.06**	buy, spend, pay	Hope	.06**	hope, wish
Like	-.06**	like, love	Attempt	-.07**	try, check

Table 2: Logistic regression coefficient between verb categories and control (positive values are control relevance and external control). **: $p < 0.01$ after Benjamini-Hochberg correction. Features with positive coefficients are shaded blue and those with negative coefficients are shaded red.

7 Validation against psychological LoC

An ideal validation of language-based control against psychological LoC would find that users expressing an internal LoC in their social media posts would also be more likely to have an internal psychological LoC as measured by the Sense of Control scale (Brim et al., 2004), and poorer health as reported in previous work (Birnbaum et al., 2010). With only two observations per user, we speculate that our data does not afford us enough power to confirm this relationship. A Spearman correlation between the counts of internal and external-labeled sentences per user and their self-reported LoC and health items, revealed a weak association between authors of a higher number of sentences labeled as external LoC and their psychological LoC as per the survey questionnaire ($\rho = -.06, p < 0.05, N = 2000$), physical and general health ($-0.08 < \rho < -0.06, p < 0.01, N = 2000$). Authors with higher number of sentences labeled as internal LoC did not show a significant association with internal LoC, but were more likely to exercise regularly ($\rho = 0.04, p < 0.05, N = 2000$). These findings, although weak, correspond in direction to the findings reported in the literature (Birnbaum et al., 2010), suggesting that future work should investigate to what extent the linguistic expressions of control can extrapolate to psychological LoC.

8 Conclusion

We describe a computational linguistic approach to identify the locus of the author’s control in their social media writing. Utilizing its psychological underpinnings, we create an annotated dataset of 4000 sentences, labeled with control relevance and of internal and external control. We show that identifying control is largely dependent on syntactic information for control-relevance and lexical information for internal/external control. From

the NLP standpoint, this suggests that solely using bag-of-word features may not be sufficient for predicting specific psychological outcomes. We have also distinguished our work against dominance and thematic roles, which may be the closest approximations to LoC; however these concepts do not completely translate into each other.

Differential language analyses identified interesting associations between verbs categories and control. We found that audio/visual verbs tend to occur significantly more in the control-related text, and some emotion verbs such as “miss” or “feel” are correlated with lack of control of the surroundings. A caveat of using language-based models is that their association with traits is correlational, not causal. Furthermore, the differences in platform affordances (Jaidka et al., 2018c) and diachronic drifts in language use over time (Jaidka et al., 2018a) imply that language models to identify control may need domain adaptation before they are applied on other corpora, language from other time periods and other social media platforms, as well as posthoc domain adaptation before they can scale to measure community traits (Rieman et al., 2017).

Existing psychological theories are mainly based on self-expressed and self-perceived locus of control in questionnaires, but they may be susceptible to self-report biases. Instead, we demonstrate that some of these constructs can reliably be extracted from language samples, such as social media posts, which are unsolicited self-expressions of control.

Internal locus of control has been argued to be important for mental and physical health, and to characterize well-run (“empowered”) work teams; We hope that the model presented here can be used—with appropriate consent and privacy—for unobtrusive monitoring of LoC in many therapy and work settings.

Acknowledgement

This work was supported by a grant from the Templeton Religion Trust (ID #TRT0048). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

- Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. 2016. Globally normalized transition-based neural networks. *arXiv preprint arXiv:1603.06042*.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Howard G Birnbaum, Ronald C Kessler, David Kelley, Rym Ben-Hamadi, Vijay N Joish, and Paul E Greenberg. 2010. Employer burden of mild, moderate, and severe major depressive disorder: mental health services utilization and costs, and work performance. *Depression and anxiety*, 27(1):78–89.
- Margaret M Bradley and Peter J Lang. 1999. Affective norms for english words (ANEW): Instruction manual and affective ratings. Technical report, Citeseer.
- Orville Gilbert Brim, Carol D Ryff, and Ronald C Kessler. 2004. *The MIDUS National Survey: An Overview*. University of Chicago Press.
- Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. Echoes of power: Language effects and power differences in social interaction. In *Proceedings of the 21st international conference on World Wide Web*, pages 699–708. ACM.
- Norman Fairclough. 2001. *Language and power*. Pearson Education.
- Ethan Fast and Eric Horvitz. 2016. Identifying dogmatism in social media: Signals and models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 690–699.
- James K Harter, Frank L Schmidt, and Corey LM Keyes. 2003. Well-being in the workplace and its relationship to business outcomes: A review of the gallup studies. *Flourishing: Positive psychology and the life well-lived*, 2:205–224.
- Kokil Jaidka, Niyati Chhaya, and Lyle Ungar. 2018a. Diachronic degradation of language models: Insights from social media. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 195–200.
- Kokil Jaidka, Salvatore Giorgi, Anneke Buffone, Johannes Eichstaedt, Masoud Rouhizadeh, and Lyle Ungar. 2018b. Modeling and visualizing locus of control with facebook language. In *Proceedings of the International AAAI Conference on Web and Social Media*.
- Kokil Jaidka, Sharath Chandra Guntuku, Anneke Buffone, H. Andrew Schwartz, and Lyle Ungar. 2018c. Facebook vs. twitter: Cross-platform differences in self-disclosure and trait prediction. In *Proceedings of the International AAAI Conference on Web and Social Media*.
- Oliver P John and Sanjay Srivastava. 1999. The big five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research*, 2(1999):102–138.
- Lingpeng Kong, Nathan Schneider, Swabha Swayamdipta, Archana Bhatia, Chris Dyer, and Noah A Smith. 2014. A dependency parser for tweets.
- Beth Levin. 1993. *English verb classes and alternations: A preliminary investigation*. University of Chicago press.
- Roger Levy and Christopher Potts. 2015. Negotiating lexical uncertainty and expertise with disjunction. In *Poster presented at the 89th Annual Meeting of the Linguistic Society of America, Portland*.
- Michael G Marmot, Stephen Stansfeld, Chandra Patel, Fiona North, Jenny Head, Ian White, Eric Brunner, Amanda Feeney, and G Davey Smith. 1991. Health inequalities among british civil servants: the whitehall ii study. *The Lancet*, 337(8754):1387–1393.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Olutobi Owoputi, Brendan O’Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A Smith. 2013. Improved part-of-speech tagging for online conversational text with word clusters. Association for Computational Linguistics.
- James W Pennebaker, Roger J Booth, and Martha E Francis. 2007. Linguistic inquiry and word count: LIWC [computer software]. *Austin, TX: liwc. net*.
- Daniel Preoțiuc-Pietro, Vasileios Lampos, and Nikolaos Aletras. 2015. An analysis of the user occupational class through twitter content. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1754–1764.
- Daniel Rieman, Kokil Jaidka, H Andrew Schwartz, and Lyle Ungar. 2017. Domain adaptation from user-level facebook models to county-level twitter predictions. In *Proceedings of the Eighth International*

Joint Conference on Natural Language Processing (Volume 1: Long Papers), volume 1, pages 764–773.

Julian B Rotter. 1966. Generalized expectancies for internal versus external control of reinforcement. *Psychological monographs: General and applied*, 80(1):1.

Masoud Rouhizadeh, Lyle Ungar, Anneke Buffone, and H Andrew Schwartz. 2016. Using syntactic and semantic context to explore psychodemographic differences in self-reference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2054–2059.

H Andrew Schwartz, Salvatore Giorgi, Maarten Sap, Patrick Crutchley, Lyle Ungar, and Johannes Eichstaedt. 2017. Dlatk: Differential language analysis toolkit. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 55–60.

Richard Socher, John Bauer, Christopher D Manning, et al. 2013. Parsing with compositional vector grammars. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 455–465.