

An Approach to the CLPsych 2018 Shared Task Using Top-Down Text Representation and Simple Bottom-Up Model Selection

Micah Iserman, Molly E. Ireland, Andrew K. Littlefield, Tyler Davis, and Sage Maliepaard

Department of Psychological Sciences, Texas Tech University, Lubbock, Texas

{micah.iserman, molly.ireland}@ttu.edu

Abstract

The Computational Linguistics and Clinical Psychology (CLPsych) 2018 Shared Task asked teams to predict cross-sectional indices of anxiety and distress, and longitudinal indices of psychological distress from a subsample of the National Child Development Study (Brown and Goodman, 2014), started in the United Kingdom in 1958. Teams aimed to predict mental health outcomes from essays written by 11-year-olds about what they believed their lives would be like at age 25. In the hopes of producing results that could be easily disseminated and applied, we used largely theory-based dictionaries to process the texts, and a simple data-driven approach to model selection. This approach yielded only modest results in terms of out-of-sample accuracy, but most of the category-level findings are interpretable and consistent with existing literature on psychological distress, anxiety, and depression.

The CLPsych Shared Task¹ this year asked a question with relevance to the nature of continuity and change in mental health: Can we predict concurrent and future mental health from childhood writing samples? Aggregate results from this task may have the potential for near-future applied value, especially for clinicians working with children. If we find usable signals in the linguistic data that have not been uncovered by the extensive prior research on the NCDS (Davie et al., 1972; Elliott, 2010; Richardson et al., 1976), the writing task that served as the basis of our analyses (asking 11-year-old children to imagine what

their lives would be like at age 25) could be easily adapted into individual-level clinical practice or group-level school counseling programs.

Interpretability. Our team’s overarching goal in this analysis was to produce interpretable results. To this end, we used dictionaries and latent Dirichlet allocation (LDA; Blei et al., 2003) to make sense of the texts, and focused on reducing these features further in our modeling to arrive at a tractable number of variables. Before describing our methods, we briefly outline our thought processes as we worked through the Shared Task’s prediction problems, and foreshadow results that we believe may be of particular interest to psychologists and clinicians.

Perhaps because of our backgrounds in psychology, our bias is typically to rely on dictionary-based processing (such as done by Linguistic Inquiry and Word Count; LIWC; Pennebaker et al., 2015) as a first step. However, we felt that taking a relatively simple, dictionary-based approach to the task was a particularly good fit with the applied focus of this year’s CLPsych Workshop: “From Keyboard to Clinic, Talk to Walk” (Loveys et al., 2018). Clinicians and practitioners with limited computational linguistic experience may be more willing to adopt methods—or, at the minimum, consider results—that are more transparent and face valid.

Although our ultimate model selection was data-driven, many of LIWC’s dictionaries are theory-driven, and tend to be face valid as a result. That is, they were designed by psychologists to measure psychological constructs, such as future focus or tentativeness, and the words in those dictionaries tend to directly reflect the constructs they aim to measure. In contrast, data-driven approaches to dictionary development may include some words that statistically predict a particular state of mind (e.g., temporal focus) but are not in-

¹This paper uses data from the NCDS, based on a dataset prepared for the CLPsych Shared Task 2018. We are grateful to the Centre for Longitudinal Studies (CLS), University College London Institute of Education for the use of these data. The CLS does not bear any responsibility for the analysis or interpretation of these data.

Topic	Words
1	<i>came, day, going, got, left, man, next, said, started, went</i>
2	<i>bed, clock, come, comes, day, dinner, get, go, home, make, morning, o, o'clock, past, put, ready, start, tea, work</i>
3	<i>best, car, cars, catch, city, club, cop, fishing, football, good, live, old, play, playing, team, time, week, wife, years</i>
4	<i>air, beat, going, job, leave, new, people, place, places, see, time, world</i>
5	<i>age, books, college, enjoy, family, flat, happy, home, hope, interested, interests, life, live, married, parents, quite, scare, swimming, think, time</i>
6	<i>day, days, friend, friends, go, home, night, saturday, see, sometimes, stay, sunday, take, times, town, week, work</i>
7	<i>baby, boy, child, children, class, clothes, girl, hospital, house, husband, little, look, married, nice, nurse, school, take, teach, teacher, work</i>
8	<i>act, arm, big, car, country, food, garden, house, live, room, side</i>
9	<i>buy, dad, get, give, good, help, hordes, hourse, house, job, make, money, mother, mum, nice, people, shop, try, want, week</i>
10	<i>boy, brother, called, father, female, girl, going, hair, live, male, married, mother, name, old, sister, year, years</i>

Table 1: Latent Dirichlet allocation topics.

tuitively related to how psychologists typically operationalize that mental state (Garten et al., 2018; Schwartz et al., 2013). Likewise, approaches that regress some outcome on every possible combination of two or more characters (e.g., moving windows) or every word (e.g., the baseline unigram models) run the risk of not providing clinicians with actionable insights, or generalizing well to other samples.

Another way to think about interpretability is as a source of insensitivity to data, or stability. Following something like a mental models perspective (and specifically Yufik and Friston, 2016), we might think of individuals’ engagement with the world as a modeling process. Understanding in this view can then be seen as a mechanism of conservation; when we feel we understand something,

we can stop collecting data, or process data less deeply (maybe only enough to check for disagreement with our understanding). The stable models associated with a sense of understanding also allow us to go about predicting future states—something that would not be possible if our models were too much in flux, if only because they would not offer any stable prediction.

A more traditional (and perhaps fallibilist) perspective on understanding might relate it to something like causal explanations, where understandability is taken to be associated with truth. For example, a theory-based approach might be seen as focused on the discovery of truths (understanding), where a more exploratory approach might be focused on fleeting but temporarily useful results (prediction). The thinking here may be that there are stable, causal underpinnings to the world, but there is considerable noise hiding those underpinnings, so we need to develop theories, and test those theories in a triangular fashion, to cut through the noise and arrive at those causes. An equivalent perspective that does not appeal to stable, discoverable causes is to see the whole understanding-seeking process as statistical modeling, where theories (and even the concept of truth) are biasing factors that work to stabilize conceptualizations of the world, allowing for prediction (and, therefore, long term action) within it.

Dictionary-based approaches can be thought of in much the same way, at least in terms of their understandability and biasing influence on modeling. Dictionaries work toward understandability by simplifying the representation of the data (reducing its dimensions). In this reduction process, dictionaries also smooth the raw text data insensitively by effectively unit-weighting words (at least as we generally apply them). Theory-based dictionaries are additionally biased as they draw on other assimilative, simplifying efforts.

Predictions. Although our approach was very exploratory, with no specific predictions (or, rather, a large number of ad hoc, speculative predictions based on our reading of the essays and our background knowledge of relevant research), there were some categories that we paid particular attention to. Some of these predictors were paralinguistic rather than standard word lists. For example, we were especially interested in the total percentage of words captured by our dictionary before and after automated text cleaning for misspellings—

which, in texts written by 11-year-olds, were predictably quite common. Those predictors were considered not based on existing theories that we are aware of, but on our reading of the training essays, which varied widely in spelling and coherence. Although we did not find any evidence that cognitive complexity (variously measured using LIWC’s cognitive mechanism and analytic categories) related to outcomes, the two variables representing misspellings (dictionary percentage captured before and after text processing) were robust predictors of present and future mental distress.

The essays also varied widely in adherence to the writing instructions, which were to write about life at age 25 as though you were currently that age. That variance led us to predict that focusing on the future—that is, *not* following the implicit instructions to use the present tense—would be positively rather than negatively related to distress. As the results show, that prediction was somewhat supported in the cross-sectional data.

Finally, based on previous findings on early life writing predicting distant future outcomes (e.g., positivity in novice nuns’ autobiographies predicting longer lives; Danner et al., 2001), we expected LIWC’s emotion categories (posemo, negemo, anx, anger, and sad) to predict outcomes either alone or in combination with sex or personal pronouns as moderators. Those predictions were not supported, perhaps partly due to low base-rates of emotional language in the essays.

This analysis of deidentified archival data is considered exempt according to federal standards for human subjects research in the United States and has been approved by the Institutional Review Board at Texas Tech University, Lubbock, TX.

1 Text Processing

The full training set included 9,217 essays written at age 11. The first step in processing the texts was to account for regular aspects that would make word boundaries less clear. For example, asterisks marked uncertain transcriptions, and line breaks were marked by characters. Illegible words were also filled in with variable numbers of asterisks or xs, which we standardized so they would all be treated as the same (as an “illegible” category). After this initial, more text specific cleaning, we translated the texts into a unigram document-term matrix. This involved more generic processes that aimed to identify word boundaries (such as at-

tempting to account for unusual punctuation or formatting; using an R package currently in development²).

This preprocessing resulted in a matrix with each essay in a row, and counts of each unique word-form in columns. To roughly account for misspellings (which were intentionally preserved in the transcription process), we first looked for matches to dictionary terms in the unique word-forms. The complete dictionary we used consisted of 121 categories from a few dictionaries: the LIWC 2015 dictionary, the revised Moral Foundations dictionary (Frimer et al., 2018), and an internal lab dictionary (Ireland and Iserman, 2018). We compared words that did not match any dictionary term to those that did; if the unmatched word was within 1 edit distance (optimal string alignment, calculated with the stringdist package; van der Loo, 2014) of any matched word, we added it to that matched word. Once we checked all unmatched words and included them if there was a close match, we calculated category scores from the matched words.

After dictionary matching and edit-distance reduction, we split the official training sample into 2/3 (6,145) train, and 1/3 (3,072) test samples. Using the 2/3 training sample, we extracted LDA topics (using the topicmodels package; Grün and Hornik, 2011) from the augmented matched words (excluding function words; Table 1), and calculated topic scores based on the words in each topic, in the same way as the dictionary categories. Then we converted category scores to percents of the author’s total word count. Once the category scores were calculated and weighted, we calculated a few composites (z -scored, averaged combinations of categories; Table 2). We also calculated the percent of words captured before and after edit-distance reduction: Mean dictionary capture before reduction = 90.24%, after = 95.83%.

2 Modeling and Results

The models we ended up using for our predictions were linear mixed-effects regression models which estimated intercepts for each sex and class group (fit with the lme4 package; Bates et al., 2015). We selected predictors from the full set of variables by calculating the correlation (Pearson’s r) of each with each outcome within the 2/3 training sample; any variable with an absolute correla-

²<https://github.com/miserman/lingmatch>

Composite	Categories
LIWC	
Feminine style	<i>i + shehe + auxverb + adverb + conj - article - prep - quant</i>
Feminine content	<i>feel + social + posemo + anx - Sixltr - number</i>
Extraversion	<i>family + friend + posemo + sexual + social + we</i>
Emotional Stability	<i>-anger - anx - i - negemo - sad</i>
Agreeableness	<i>family + posemo + time + we - anger - negate - swear</i>
Conscientiousness	<i>time - anger - negate - negemo - swear</i>
Openness	<i>article + death - family - home - pronoun - time</i>
Categorical-Dynamic Index	<i>article + prep + ppron - ipron - auxverb - conj - adverb - negate</i>
Exclusive present focus	<i>focuspresent - focusfuture - focuspast</i>
Distress	<i>negate + swear + negemo + i + death + tentat + discrep - posemo - we</i>
Revised social dictionary	
Social: near	<i>communication + family + friend + romance</i>
Social: far	<i>association + humans</i>

Table 2: Dictionary category composites.

	Age 11 Bristol Social Adjustment Guide			Future Psychological Distress		
	Total	Anxiety	Depression	Age 23	Age 33	Age 42
Intercept	-.14***	.00*	-.14*	-.17	-.10	-.10*
Control: class II	-.06	-.07	-.03	.06	.02	.01
Control: class III manual	.18	.02	.16	.21*	.13	.10
Control: class III non-manual	.04	-.02	.10	.05	-.04	-.01
Control: class IV	.26	.04	.24	.25*	.13	.17
Control: class V	.33	-.03	.39 [†]	.37**	.22	.34*
Control: sex (0 = male)	-.10	.01	-.04	.26	.17	.16
WC	-.08***	-.03*	-.08***	-.01	-.01	-.01
Dic before	-.24***	-.13***	-.20***	-.12***	-.08**	-.02
Dic after	.09***	.05**	.09***	.06*	.05	-.02
informal	.03	.03	.05*	-.01	.01	-.02
netspeak	.00	-.02	-.04*	.02	.02	.01
prep	-.05***	-.03 [†]	-.04**	-.01	-.02	-.01
focusfuture	.06***	.02	.02 [†]	.01	-.01	-.01
conj	.06***	.00	.06***	.05**	.04*	.03 [†]
fairness.vice	.05***	.02 [†]	.05***	.03 [†]	.03*	.01
Topic 7	.02	.05***	.00	.04*	.03 [†]	.03

Table 3: Standardized beta weights for each selected variable predicting each outcome. Markings denote unadjusted p-values: [†] < .1, * < .05, ** < .01, *** < .001

tion over .1 with any of the outcomes was added as a predictor in all models. These variables included word count, and both forms of dictionary capture percent (Dic; before and after reduction); the informal, netspeak, prep, focusfuture, and conj LIWC categories; the fairness.vice Moral Foundations category; and the 7th LDA topic. The estimates from each model are reported in Table 3. Once we settled on these models, we refit them to the full official training set, and calculated predictions for submission. The results of these models on the held-out sample are reported in Table 4.

One motivation for these models was the generally poor results of the other methods we attempted. We considered recursive partitioning algorithms and elastic net regularization for their effective reduction of the number of variables. We applied these to a the full variable set, as well as sets which variably included individual words and interaction terms with the controls in addition to the dictionary categories, and we allowed models to vary between outcomes. Even though these attempts were more sophisticated in that they had more potential cues, and attuned to each outcome, they did not outperform our ultimate, blunt approach.

Looking at the relationship between our predictions and the actual outcomes in the full training sample suggests we may have been able to improve some of our predictions by better accounting for sex. For example, our model predicts psychological distress at age 23 for women better than it does for men (Figure 1). For other outcomes, the model performs equally well between sexes, such as when predicting total Bristol Social Adjustment Guide (BSAG) scores at age 11 (Figure 2). These figures also point to the loose relationships between predictions and outcomes in general, and to the very low rates of maladjustment and distress. The bulk of actual and predicted outcome scores tended to be low on each scale, and some of the positive relationship between them seems to be driven by only a relatively small number of more extreme scores (those few points that scatter out toward higher predicted scores).

3 Discussion

Altogether our methods performed better for future than present psychological distress (relative to other teams). Within the cross-sectional predictions, anxiety was least reliably related to our

Outcome	MAE	dis r	Rank
Age 11 BSAG Total	6.050	.461	8
Age 11 BSAG Anxiety	0.704	.142	6
Age 11 BSAG Depression	1.055	.330	9
Average future distress	1.176	.314	2
Age 23 distress	1.087	.457	1
Age 33 distress	1.092	.277	3
Age 42 distress	1.350	.208	3

Table 4: Mean Absolute Error (MAE) and disattenuated Pearson’s r ($r/.7341662$; dis r) of the predictions with the held-out sample outcomes. Rankings are out of 9 teams, based on dis r .




Figure 1: Actual versus predicted age 23 psychological distress in the full training sample, separated by sex.

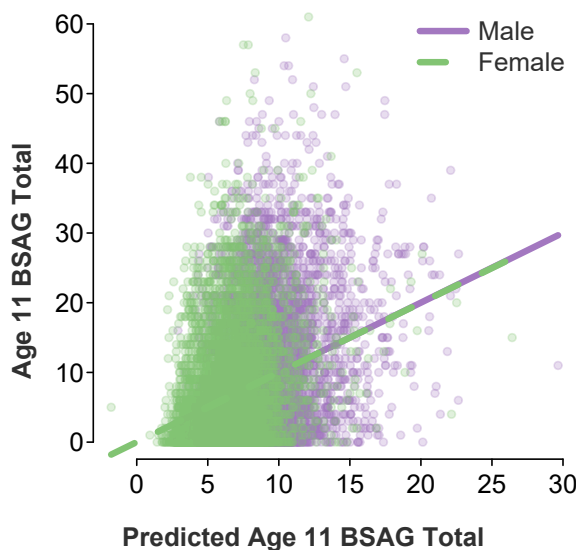


Figure 2: Actual versus predicted age 11 BSAG total score in the full training sample, separated by sex.

linguistic predictors. In the future predictions, we particularly struggled to identify predictors of age 42 psychological distress. Below, we will attempt to make sense of our exploratory (that is, mostly unpredicted) findings in the context of existing literature from clinical psychology and computational linguistics, focusing primarily on findings that were relatively robust across concurrent and future mental health outcomes.

Anxiety. Although the explanation for our limited success in predicting distress at the latest time point (age 42) is relatively obvious—that is, predicting mental states and behavior across time becomes more challenging as the amount of lapsed time increases—the difficulty of predicting present anxiety is less clear. As the existing literature on the linguistic correlates of mental health focuses more on depression than anxiety (De Choudhury et al., 2013; Rude et al., 2004; Tackman et al., 2018), we are reluctant to speculate on why the signal for depression should be stronger than for anxiety. Indeed, it may be the case that depression is studied more frequently in psychology and computational linguistics specifically because it has a stronger or more reliable linguistic signal than anxiety or related mental health conditions (a file-drawer effect).

On the other hand, a more promising explanation (than simply arguing that anxiety is hard to measure) could be that there were differences in measurement error between the anxiety and depression measures; perhaps depressive symptoms were easier for teachers to accurately rate than anxiety symptoms, in general or in this particular population of British children. We will be able to more confidently interpret our weak prediction accuracy for cross-sectional anxiety after exploring the other teams' summaries of their successes and failures in this same subtask.

Dictionary percent capture. One of the most robust predictors of present and future psychological distress was, somewhat surprisingly, rate of words captured by the dictionary before and after our automated cleaning process, described above. Having more words captured by dictionaries—in other words, using more commonly used and recognizably spelled words—in the raw (unprocessed) texts predicted less anxiety and depression at age 11, and less general psychological distress at ages 23 and 33. These effects were all significant in the full model, suggesting a useful signal

above and beyond other related predictors, such as two LIWC measures of less conscientious or formal language use (netspeak and informal categories).

Results suggest that misspelling—above and beyond its moderate association with socioeconomic status (Sénéchal and LeFevre, 2002)—reflects general psychological distress that is not limited to a specific disorder or class of symptoms. If the associations we observed are causal, the relation between misspellings and distress could be bidirectional, with distress leading to cognitive load and distraction from class, and poor academic engagement or performance exacerbating existing psychological vulnerabilities via academic stress.

In contrast, the dictionary percent captured after text processing (automated spelling correction) was positively correlated with anxiety, depression, and distress at all but the latest time point. That is, after accounting for misspellings in the original, people who use higher frequency words tend to be more distressed. Again, these results could reflect a bidirectional relationship between distress and the various aspects of academic performance that dictionary percentage may reflect—such as creativity or vocabulary level. That is, distress may limit academic performance, and poor academic performance increases stress for most children.

Word count. Another modest but reliable predictor of cross-sectional mental health was word count. Children who wrote fewer words in their essays had more severe behavioral and psychological symptoms, as measured by teachers' observations. Verbosity or word count has not played a major role in most past LIWC research. When it appears in analyses at all, word count is often treated as a nuisance variable to be partialled out of predictive models (Ireland et al., 2011).

However, recent evidence suggests that saying fewer words in daily life is a robust predictor of general psychological distress (Mehl et al., 2017). That finding and our present results dovetail nicely with earlier theories, partly arising from the expressive writing paradigm (Pennebaker, 2018), that inhibition is a key predictor of future physical illness and psychological distress (Pennebaker, 1989). The rationale is that chronic inhibition (e.g., when keeping secrets or concealing stigmatized identities, such as sexual orientation) not only requires constant vigilance, greatly increasing stress and allostatic load (Meyer, 2003), but

also by definition limits individuals' agency and self-efficacy, or ability to freely pursue personal goals (Bandura, 1982).

More parsimoniously, decreased word count in these essays could simply reflect less academic engagement or poorer attentional control, perhaps resulting from higher impulsivity (Stevens et al., 2018). Along the same lines, future focus (i.e., future tense verbs and references to the future) was positively correlated with overall behavioral problems, suggesting that failing to follow the instructions—which were to write as though you were *currently* age 25—may have reflected academic defiance or disengagement (refusing to follow instructions), poor reading comprehension (not understanding the instructions), or language impairments (not being able to follow the instructions).

Other explanations for the associations between concurrent mental health and word count may have to do with the nature of the task. Thoughts about the future—or any area of life that involves uncertainty—are often a primary source of anxiety for people with anxiety disorders (Grupe and Nitschke, 2013). Writing as little as possible when asked to think and write about the future could represent avoidant coping with an anxiety-inducing task (Herman-Stabl et al., 1995). Future analyses may benefit from taking a finer-grained approach to measuring temporal orientation or prospecting, perhaps differentiating between various aspects of thinking about the future, including affective forecasting, episodic simulation, and autobiographical planning (Szpunar et al., 2014).

Conjunctions. Conjunctions were positively correlated with all indicators of psychological distress except for concurrent anxiety. These results add to the impression that academic engagement and conformity to academic norms may have been the primary predictors of both present and later life distress for these participants. Conjunctions are a key part of a language style sometimes referred to as dynamic or narrative thinking; the opposite is categorical or analytic thinking, which involves more nouns, prepositions, and articles, and fewer conjunctions, pronouns, adverbs, and verbs, among a few other categories (Jordan and Pennebaker, 2017). Dynamic thinking is more conversational and informal—better suited for social interactions than an essay writing assignment, perhaps. Analytic thinking—again, mathemati-

cally the opposite of dynamic thinking—in students' college admissions essays predicts higher GPAs throughout college (GPA and graduation rates; Pennebaker et al., 2014).

Fairness (vice). Finally, children who discussed the “vice” side of fairness (e.g., *unfair, unequal, bias*) experienced more psychological distress concurrently and in the future (Graham et al., 2009). The simplest explanation could be that children talked about what they had experienced in life, and people who experience chronic maltreatment or unfairness have more stress and are therefore at higher risk for distress and mental illness (Shonkoff et al., 2012).

Alternately, discussing unfairness in the future—that is, expecting your future to be as unfair as your present—could represent hopelessness (Van Allen et al., 2015) or pessimism (Plomin et al., 1992). Hopelessness in particular has emerged recently as a factor that leads to poorer adherence to prescribed healthcare regimens (e.g., in type 1 diabetes mellitus) and worse health outcomes longitudinally in children (Van Allen et al., 2015). In other words, hopelessness leads to maladaptive coping strategies, such as disengagement coping or drinking to cope, and impedes goal-congruent behavior, thus exacerbating existing mental health vulnerabilities (Carver and Connor-Smith, 2010).

4 Limitations and Future Directions

If we think of people as agents who act and make decisions within their relatively immediate environments, their lives are Markovian processes, and so, prediction of their distant positions is bound to be limited. This perspective is comforting at least in terms of a sense of free will—it allows for (what at least feel like) meaningful decisions even within an effectively deterministic system.

Of course, there are strong considerations of starting position within the system, and their associated levels of adversity. Random walks or those biased toward known regions will generally hover around the same position. Moving in a directed way may also be difficult given any uncertainty about where a step will lead, so even without randomness or bias, starting positions can be informative. This is where the sex and class controls come in.

Directed movement from those initial positions can be thought of as long term, prospective action,

which draws back on the notion of understanding as a modeling process and means of prediction, as it allows for such action. This may be one way of interpreting the LDA topic from the final model: The 7th LDA topic focuses on family and work, and may be reflective of an expected future—a future potentially modeled and referred to explicitly by parents, family, and society in general. This topic is positively related to age 11 anxiety, which makes sense if seen as an expectation others have for the child’s future, that the child is aware of and applies to their own image of their future. The topic is also related to age 23 distress, which might make sense if we imagine the child continues to direct themselves toward others’ expectations. Cues to models such as this may be a route to prediction of future states within a stochastic system inasmuch as they speak to the directed behavior of the agent. This would stand in contrast to some theoretical psychological feature outside of the agent’s control (such as mental health vulnerabilities), and to features of their more immediate environment, potentially observable in less psychologically meaningful linguistic patterns.

At a more quotidian level, any generalizations that we or interested clinicians can draw from the current results are limited by our modest effect sizes—which, as noted above, are partly a consequence of the low base-rates of distress at any time point in the NCDS sample. As with any prediction of low base-rate behaviors (such as spree killing or suicide; e.g., Pokorny, 1983; Iserman and Ireland, 2017; Walsh et al., 2017) based on relatively noisy behavioral data, the clinical utility of our results is limited. Any attempt to use the language patterns that we have identified in clinical practice as diagnostic tools or prospective predictors of clients’ future depression or anxiety may lead to a large number of false positives (Mitchell et al., 2009), which in some cases may be more ethically troubling than false negatives.

We have no easy solution to our results’ various statistical and methodological shortcomings. Small effect sizes are a common limitation of text analytic approaches to understanding human psychology, particularly when attempting to predict low base-rate events or diagnoses (Pennebaker and King, 1999). Still, text analysis could have practical value; for example, a clinician might take a rubber mallet approach, analyzing text (perhaps that they have already collected, or have ready

access to via social media) for a low impact, low precision tool to supplement their more intensive and refined tool set. Working with language in this way, and focusing on subtle linguistic cues may also positively carry over into the clinician’s other methods (such as hearing the client differently in interviews or sessions). Along the same lines, the current linguistic results—and similar interpretable findings uncovered by other Shared Task teams—could help fine-tune (rather than solely determine) treatment regimens on a client-by-client basis.

5 Conclusion

One takeaway from this task is that current maladjustments and future distress are not readily predictable from largely unrelated writing tasks. We believe this to be more encouraging than discouraging. The only real discouraging aspect of this perspective is the limit it suggests on the accurate detection of such issues. The encouraging aspects are that—in the near term—forms of illadjustment do not always and overwhelmingly pervade every aspect of a child’s life (they can imagine their future without obvious distress), and that—in the longer term—children who experience these issues are not destined for inordinate future distress. Reading through some of the essays and comparing with the adjustment and distress scores seem to support this perspective as well.

Judging by our rankings, the simplicity of our approach to the texts may have harmed our age 11 predictions, but it may also have improved our longer term predictions (or perhaps just failed to actively harm them). This may be due to the insensitivity of dictionary-based processing; it is limited in its ability to capitalize on idiosyncrasies (of individuals or datasets), which may tend to make it more modest.

Future approaches to similar tasks may benefit from more seamlessly integrating top-down and bottom-up approaches to dictionary-based prediction. We are encouraged by new strategies that improve theory-driven dictionaries using data-driven methods (e.g., distributed representations; Garten et al., 2018) and hope that additional work in that vein will bolster computational linguists’ ability to provide clinicians and other practitioners with actionable insights about mental health.

References

- Albert Bandura. 1982. Self-efficacy mechanism in human agency. *American psychologist*, 37(2):122.
- Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. [Fitting linear mixed-effects models using lme4](#). *Journal of Statistical Software*, 67(1):1–48. Version 1.1-17.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Matt Brown and Alissa Goodman. 2014. National child development study (or 1958 birth cohort). *Open Health Data*, 2(1).
- Charles S Carver and Jennifer Connor-Smith. 2010. Personality and coping. *Annual review of psychology*, 61:679–704.
- Deborah D Danner, David A Snowdon, and Wallace V Friesen. 2001. Positive emotions in early life and longevity: findings from the nun study. *Journal of personality and social psychology*, 80(5):804.
- Ronald Davie, Neville Butler, and Harvey Goldstein. 1972. From birth to seven: the second report of the national child development study.(1958 cohort). *London Longmans 1972*. 198 p. 1 ref.
- Munmun De Choudhury, Michael Gamon, Scott Counts, and Eric Horvitz. 2013. Predicting depression via social media. *ICWSM*, 13:1–10.
- Jane Elliott. 2010. Imagining a gendered future: Childrens essays from the national child development study in 1969. *Sociology*, 44(6):1073–1090.
- Jeremy Frimer, Jonathon Haidt, Jesse Graham, Morteza Dehgani, and Reihane Boghrati. 2018. [Moral foundations dictionary 2.0](#).
- Justin Garten, Joe Hoover, Kate M Johnson, Reihane Boghrati, Carol Iskiwitch, and Morteza Dehgani. 2018. Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior research methods*, 50(1):344–361.
- Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5):1029.
- Bettina Grün and Kurt Hornik. 2011. [topicmodels: An R package for fitting topic models](#). *Journal of Statistical Software*, 40(13):1–30. Version 0.2-7.
- Dan W Grupe and Jack B Nitschke. 2013. Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nature Reviews Neuroscience*, 14(7):488.
- Mindy A Herman-Stabl, Mark Stemmler, and Anne C Petersen. 1995. Approach and avoidant coping: Implications for adolescent mental health. *Journal of Youth and Adolescence*, 24(6):649–665.
- Molly E. Ireland and Micah Iserman. 2018. [Lusi lab development dictionaries](#).
- Molly E Ireland, Richard B Slatcher, Paul W Eastwick, Lauren E Scissors, Eli J Finkel, and James W Pennebaker. 2011. Language style matching predicts relationship initiation and stability. *Psychological science*, 22(1):39–44.
- Micah Iserman and Molly Ireland. 2017. A dictionary-based comparison of autobiographies by people and murderous monsters. In *Proceedings of the Fourth Workshop on Computational Linguistics and Clinical Psychology—From Linguistic Signal to Clinical Reality*, pages 74–84.
- Kayla N Jordan and James W Pennebaker. 2017. The exception or the rule: Using words to assess analytic thinking, donald trump, and the american presidency. *Translational Issues in Psychological Science*, 3(3):312.
- Kate Loveys, Kate Niederhoffer, Emily Prudhommeaux, Philip Resnik, and Rebecca Resnik. 2018. [Call for papers! fifth workshop on computational linguistics and clinical psychology: From keyboard to clinic, talk to walk](#).
- Matthias R Mehl, Charles L Raison, Thaddeus WW Pace, Jesusa MG Arevalo, and Steve W Cole. 2017. Natural language indicators of differential gene regulation in the human immune system. *Proceedings of the National Academy of Sciences*, 114(47):12554–12559.
- Ilan H Meyer. 2003. Prejudice, social stress, and mental health in lesbian, gay, and bisexual populations: conceptual issues and research evidence. *Psychological bulletin*, 129(5):674.
- Alex J Mitchell, Amol Vaze, and Sanjay Rao. 2009. Clinical diagnosis of depression in primary care: a meta-analysis. *The Lancet*, 374(9690):609–619.
- James W Pennebaker. 1989. Confession, inhibition, and disease. In *Advances in experimental social psychology*, volume 22, pages 211–244. Elsevier.
- James W Pennebaker. 2018. Expressive writing in psychological science. *Perspectives on Psychological Science*, 13(2):226–229.
- James W. Pennebaker, Ryan L. Boyd, Kayla Jordan, and Kate Blackburn. 2015. [The development and psychometric properties of liwc2015](#). *UT Faculty/Researcher Works*.
- James W Pennebaker, Cindy K Chung, Joey Frazee, Gary M Lavergne, and David I Beaver. 2014. When small words foretell academic success: The case of college admissions essays. *PloS one*, 9(12):e115844.

- James W Pennebaker and Laura A King. 1999. Linguistic styles: Language use as an individual difference. *Journal of personality and social psychology*, 77(6):1296.
- Robert Plomin, Michael F Scheier, Cindy S Bergeman, Nancy L Pedersen, John R Nesselroade, and Gerald E McClearn. 1992. Optimism, pessimism and mental health: A twin/adoption analysis. *Personality and individual differences*, 13(8):921–930.
- Alex D Pokorny. 1983. Prediction of suicide in psychiatric patients: report of a prospective study. *Archives of general psychiatry*, 40(3):249–257.
- Ken Richardson, M Calnan, J Essen, and L Lambert. 1976. The linguistic maturity of 11-year-olds: Some analysis of the written compositions of children in the national child development study. *Journal of Child Language*, 3(1):99–115.
- Stephanie Rude, Eva-Maria Gortner, and James Pennebaker. 2004. Language use of depressed and depression-vulnerable college students. *Cognition & Emotion*, 18(8):1121–1133.
- Hansen Andrew Schwartz, Johannes Eichstaedt, Eduardo Blanco, Lukasz Dziurzyński, Margaret L Kern, Stephanie Ramones, Martin Seligman, and Lyle Ungar. 2013. Choosing the right words: Characterizing and reducing error of the word count approach. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, volume 1, pages 296–305.
- Monique Sénéchal and Jo-Anne LeFevre. 2002. Parental involvement in the development of children's reading skill: A five-year longitudinal study. *Child development*, 73(2):445–460.
- Jack P Shonkoff, Andrew S Garner, Benjamin S Siegel, Mary I Dobbins, Marian F Earls, Laura McGuinn, John Pascoe, David L Wood, Committee on Psychosocial Aspects of Child, Family Health, Adoption Committee on Early Childhood, Dependent Care, et al. 2012. The lifelong effects of early childhood adversity and toxic stress. *Pediatrics*, 129(1):e232–e246.
- Angela K Stevens, Brittany E Blanchard, and Andrew K Littlefield. 2018. Impulsive dispositions and alcohol: what we know, how we know it, and where to go from here. *Borderline personality disorder and emotion dysregulation*, 5(1):4.
- Karl K Szpunar, R Nathan Spreng, and Daniel L Schacter. 2014. A taxonomy of prospection: Introducing an organizational framework for future-oriented cognition. *Proceedings of the National Academy of Sciences*, 111(52):18414–18421.
- Allison M Tackman, David A Sbarra, Angela L Carey, M Brent Donnellan, Andrea B Horn, Nicholas S Holtzman, To'Meisha S Edwards, James W Pennebaker, and Matthias R Mehl. 2018. Depression, negative emotionality, and self-referential language: A multi-lab, multi-measure, and multi-language-task research synthesis. *Journal of personality and social psychology*.
- Jason Van Allen, Ric G Steele, Michael B Nelson, James Peugh, Anna Egan, Mark Clements, and Susana R Patton. 2015. A longitudinal examination of hope and optimism and their role in type 1 diabetes in youths. *Journal of pediatric psychology*, 41(7):741–749.
- Mark P.J. van der Loo. 2014. [The stringdist package for approximate string matching](#). *The R Journal*, 6:111–122. Version 0.9.4.7.
- Colin G Walsh, Jessica D Ribeiro, and Joseph C Franklin. 2017. Predicting risk of suicide attempts over time through machine learning. *Clinical Psychological Science*, 5(3):457–469.
- Yan M. Yufik and Karl Friston. 2016. [Life and understanding: The origins of understanding in self-organizing nervous systems](#). *Frontiers in Systems Neuroscience*, 10:98.