The Centroid Cannot Hold¹: Comparing Sequential and Global Estimates of Coherence as Indicators of Formal Thought Disorder

Weizhe Xu, BS¹, Jake Portanova, BA, BS¹, Ayesha Chander, MRes², Dror Ben-Zeev, PhD², Trevor Cohen, MBChB, PhD¹

¹Biomedical Informatics and Medical Education; ²Behavioral Research in Technology and Engineering (BRiTE) Center, Psychiatry and Behavioral Sciences; University of Washington, Seattle, WA

Abstract

Thought disorder (TD) as reflected in incoherent speech is a cardinal symptom of schizophrenia and related disorders. Quantification of the degree of TD can inform diagnosis, monitoring, and timely intervention. Consequently, there has been an interest in applying methods of distributional semantics to quantify incoherence of spoken language. Prior studies have generally involved few participants and utilized speech data collected in on-site structured interviews. In this paper we conduct a comprehensive evaluation of approaches to quantify incoherence using distributional semantics, including a novel variant that measures the global coherence of text. This evaluation is conducted in the context of “audio diaries” collected from participants experiencing auditory verbal hallucinations using a smartphone application. Results reveal our novel global coherence metric using the centroid (weighted vector average) outperforms established approaches in their agreement with human annotators, supporting their preferential use in the context of short recordings of unstructured and largely spontaneous speech.

Introduction

Coherent discourse is characterized by an orderly and interconnected flow of ideas, and coherence has been defined as the “semantic similarity” between these ideas in previous work¹. In psychiatry, speech in which it is difficult to perceive such connections is thought to reflect a type of underlying thought disorder (TD), which has long been recognized as an important diagnostic feature of schizophrenia in particular^{2,3,4}. Schizophrenia is a serious mental illness that has an estimated prevalence of 4.6 per 1000 persons⁵. It has been shown that patients with schizophrenia have worse quality of life than the general population and other physically ill patients⁶. TD in particular has been shown to strongly correlate with the impairment of work performance⁷ and poor functionality in the community⁸. Therefore, quantifying TD is important for schizophrenia prognosis⁹, as well as for diagnosis⁴ which is clinically important because early detection allows for timely intervention to mitigate the condition¹⁰.

TD manifests as abnormalities in speech pattern, ranging from loose association of contents (derailment) to entirely incomprehensible speech¹¹. Traditionally, TD is evaluated clinically, but it can also be more objectively measured by certain scoring constructs such as a self-reporting scales¹² and clinically administered rating scales^{13,14}. However, inherent problems are associated with these scale measures: self-scoring measures lack the identification of objective symptoms associated with TD; while the administration of clinical scales is time-consuming, requires specific training and expertise, and even when used regularly can only provide intermittent measures during office visits. To address these issues, there is a burgeoning interest in the development and evaluation of automated methods to quantify TD by measuring coherence of speech^{1,15,16,17}. Motivated by prior work on automated assessment of the coherence of written text¹⁸, Latent Semantic Analysis (LSA)^{19,20} has been a prominent component of efforts to quantify speech coherence using automated metrics^{1,15}. LSA is a method of distributional semantics (for reviews see^{21,22}) that represents semantic units (words, paragraphs, or sentences) as vectors that capture patterns of co-occurrence across a large text corpus. Words that occur in similar documents will have similar vectors, which can be used to compose representations of larger text units. The semantic similarity of such units is then estimated by calculating the cosine of the angle between their vector representations. Coherence can thus be quantified by the semantic similarity between consecutive units^{1,18}.

Previous work has demonstrated the utility of LSA-based metrics as a means to estimate the coherence of speech in the context of psychotic-spectrum disorders. The automated coherence estimates resulting from this work correspond well

¹To appear in the Proceedings of AMIA 2020 Annual Symposium. Please cite the canonical version once it is available.

to clinical assessment of the degree of TD using a standardized instrument¹; have been used to distinguish between transcribed speech samples produced by patients with schizophrenia, their family members and healthy controls²³; and provide a key feature for machine learning methods used to predict the onset of psychosis in high-risk subjects^{15,16}. Although these studies provide insight into the utility of quantifying speech coherence, they share some limitations. First, they have a limited participant pool size and are confined to experimental rather than naturalistic settings. When considering the seminal studies establishing the validity of LSA-based coherence metrics¹ and their utility as features for prediction of the onset of psychosis¹⁵, both have under 40 participants whose speech samples were collected in a research setting and elicited using either extended clinical interviews or tasks such as story-telling selected on account of their perceived utility as a means to reveal TD. Smaller participants pools may lead to undercoverage bias, with exclusion of certain case variants, as may be indicated by a substantial drop in the performance of machine-learning approaches to the prediction of psychosis onset when evaluated with a larger sample¹⁶ (although accuracy nonetheless remained promising at $\approx 80\%$ even when generalizing across sample populations). With respect to experimental setting, while evaluation under controlled conditions using standardized tasks or extended interviews has advanced the science of automated measures of coherence considerably, the translational impact of these methods is contingent upon their being readily deployable under naturalistic conditions, such that they can capture fluctuation in symptom severity without the need for additional clinic visits.

Additionally, and of particular importance to the current work, the usage of LSA in previous studies focuses exclusively on semantic similarity between juxtaposed units. This includes both comparison of the subsequent units (termed “first order” coherence¹⁵) and comparison of gapped units with an intervening unit in between (“second order” coherence¹⁵). For example, in the 3-word “ $w_1 w_2 w_3$ ” sequence, word-based variants using first order coherence will compare $w_1 : w_2$ and $w_2 : w_3$ for coherence calculation, whereas the second order coherence calculation will compare $w_1 : w_3$. This method may capture local coherence characteristics but it does not consider coherence globally and *global coherence* - the ability to sustain a topic throughout spoken discourse - is an important consideration for normal speech capabilities²⁴. This is a somewhat surprising omission given that the vector average or *centroid* of words in a document has been used to represent larger units of text with LSA since the method was first introduced for information retrieval¹⁹, and provides a convenient geometric approximation of the central topic of a document that could also be used to estimate global coherence. Finally, previous studies usually focus on one semantic unit of either words or phrases for coherence analysis, without systematic comparison of the utility of different semantic units as a basis for LSA-derived estimates of coherence. A comprehensive analysis using words, phrases, and sentences may provide much-needed guidance to modelers as to which unit type best supports distributional estimates of coherence.

In this study, we aim to address these issues, while presenting and evaluating a novel approach to quantify global speech coherence. To overcome the limitation of participant pool size and address the need to estimate coherence in naturalistic settings, we evaluate our methods in the context of speech samples from a smartphone application that collects “audio diaries” describing participant experiences of auditory verbal hallucinations (AVH) recorded in naturalistic environments. AVH, like TD, are an important diagnostic consideration in schizophrenia and may be caused by the disordered monitoring of inner speech²⁵. The sample collection process did not involve structured questions, and the recordings were limited to three minutes in duration, which enhances the scalability of the data collection procedure. However distributional semantics based estimates of coherence have yet to be validated in the context of relatively unstructured speech samples of this length. The speech samples used in the current study were collected using this smartphone application from more than 150 participants, with the potential to scale to much larger participant pools. As such, validation of coherence metrics in the context of these data would provide a solution to the scalability limitations inherent in laboratory-based evaluations using lengthy interviews or structured instruments that require specialized expertise to use. To analyze transcribed speech samples for their coherence, we evaluated a range of established and novel distributional semantics based approaches that collectively not only compare consecutive semantic units, but also compute a mean vector (the *centroid*) to estimate global coherence. We also explored the utility of different semantic units by conducting a comprehensive analysis comparing them. The coherence scores generated by each approach were compared to human annotations for validation and comparative evaluation.

Table 1: Characteristics of participant pool.

Gender	Number	Percentage	Age	Number	Percentage
Male	56	39.4%	19-29	24	16.9%
Female	82	57.8%	30-39	52	36.7%
Transgender (MTF)	3	2.1%	40-49	37	26.1%
Transgender (FTM)	1	0.7%	≥ 50	29	20.4%

Method

Participants: Data were obtained from a study of participant experience of AVH, which uses a smartphone platform to capture a range of ecological momentary assessment and sensor-derived variables. The study was approved by the institutional review boards at the University of Washington and Dartmouth College. Participants experiencing AVH were recruited via both in-person and online means. Informed consents from participants were obtained through a rigorous procedure involving triple confirmations from a screening questionnaire. All participants were asked to install a mobile application, which had the capability of recording and uploading audio diaries, and were prompted to describe their experiences of AVH, as well as anything else they would like to share or think it would be helpful for the research team to know, with prompts for audio diaries following the collection of other data, and the option to record an entry directly on demand. Although no monetary incentives were offered for the audio diary component, most participants submitted their recorded audio diaries. We used data collected up to October 18th 2019, consisting of 1868 recordings from 202 users. As short recordings seldom contained interpretable language, we restricted this set to recordings of length 30 seconds or more (maximum three minutes), leaving 909 recordings from 154 users. We randomly sampled up to three recordings per user, leaving 355 recordings which were professionally transcribed. After manual inspection, we retained 310 transcripts with interpretable content, covering 142 participants (Table 1).

Transcripts: Each transcript was labeled by two human annotators with a score between 0 and 4 to indicate the degree of derailment, which is an indicator of TD^{4,26}, and was selected as a construct for the current study because it does not concern deviation from the topic of a question, and audio diary prompts were open-ended in nature. Annotation was guided by the definitions and training materials for the Thought and Language Disorder (TALD) rating scale¹⁴, a validated instrument for the assessment of TD. A score of 0 indicates that derailment is not present. A score of 4 indicates that speech is incomprehensible. Scores from 1 to 3 represent intermediate degrees of derailment, in which the connections between sentences grow less recognizable as the score increases. The raters each rated all transcripts. Any transcripts with a disagreement of two or more units on the scale ($n=22$) were re-evaluated independently, to reach a quadratically-weighted Kappa of 0.71. Note that quadratically-weighted Kappa scores penalize larger differences between scale categories more than smaller ones, which we deemed appropriate given the subtle distinctions between neighboring TALD categories. The average score of the two raters for each transcript was calculated to be used for further analysis. Table 2 shows the number of transcripts by average rater score.

Preprocessing: Transcripts were pre-processed by stop-word removal and term tokenization. Stop-words are filler words with little or no semantic content (such as “a”, “the”, and “on”). These words were defined by the stopword list distributed with the natural language tool kit (NLTK²⁷), an open source tool, and their occurrences were removed from the transcripts to reduce noise. Tokenization is a process of extracting semantic units from documents so that they can be represented by vectors for similarity comparison. For example, a word tokenizer extracts individual words from a document while maintaining their sequential order. In this study, because we were interested in conducting a comprehensive analysis of various semantic units, we tokenized the transcripts into three different units: words, noun phrases and sentences. The word and sentence tokenizations were performed using the NLTK word and sentence tokenizers. The noun-phrase tokenization involved a different tool, `textblob`²⁸, which is also publicly available.

Semantic vectors (word embeddings): Most previous work modeling coherence using distributional semantics has employed LSA^{20,18} to generate semantic vectors for words (see for example^{1,15,16,23}). LSA creates vectors based on the distributional statistics of words in a corpus (usually the Touchstone Applied Science Associates (TASA) corpus, which was used in these previous studies^{1,15,16,23}). However, *neural word embeddings*, distributed representations of words derived from neural networks trained to predict words in proximity to an observed word (such as the popular *skipgram* and *continuous bag of words* architecture²⁹ embodied in the widely used `word2vec`³⁰ and `fastText`³¹

Table 2: Transcripts by mean rater score. Line ruled between 2.5 and 3 indicates categorization threshold.

Score ($\bar{x}$)	Number	Percentage	TALD category (paraphrased and abridged)
0	35	11.29%	<i>not present</i> : no derailment
0.5	62	20.00%	
1	93	30.00%	<i>doubtful</i> : connections still obvious
1.5	53	17.10%	
2	24	7.74%	<i>moderate</i> : sometimes disconnected from prior speech
2.5	25	8.06%	
3	8	2.58%	<i>severe</i> : no meaningful connection between ideas
3.5	8	2.58%	
4	2	0.65%	<i>extreme</i> : interview is incomprehensible
Total	310	100.00%	

software packages), have been shown to outperform matrix decomposition-based approaches like LSA³², especially on novel tasks involving solving proportional analogy problems using geometric operators³³. While some of these improvements in performance have subsequently been shown to be contingent upon the selection of task-specific optimized hyperparameters³⁴, it remains true that the efficient algorithms used to train neural embeddings allow for training on much larger corpora in a relatively short time. Neural word embeddings have received little attention in related work on quantifying coherence, except in a recent study³⁵, where neural embeddings performed better than LSA in distinguishing participants with schizophrenia from controls in some experiments. Thus, in this study, neural embedding was used as a technique to generate the vector space for automated analysis of the transcripts. Publicly available *fastText* pre-trained word embeddings³⁶ were selected for this study. These vectors were trained on a large corpus derived from Common Crawl³⁷, comprised of approximately 600 billion word-level tokens, as compared with approximately 12 million in the TASA corpus, and without the use of subword embeddings.

While we used the aforementioned *fastText*-derived space for the majority of our experiments, we used four additional vector spaces - two Wikipedia-derived, and two trained on the TASA corpus - to evaluate two methodological variants applied in prior studies. Firstly, the utility of lemmatization of words was evaluated as this has been used in prior work on coherence¹⁵. Lemmatization refers to the process of reducing various forms of a word to a canonical form, for example converting “does”, “did”, “doing” and “done” to “do”. Lemmatization has been used as a normalization procedure to accommodate morphological variants in distributional semantics²². For the purpose of comparison, we trained neural word embeddings using the open source *Gensim*^{38,39} implementation of the skipgram-with-negative-sampling algorithm⁴⁰ (which is also a component of *word2vec* and *fastText*), to generate a vector space from lemmatized and non-lemmatized versions of a Wikipedia-derived corpus. We generated a 100-dimensional vector space without imposing frequency thresholds (i.e. including all terms), using a window size of 5, a subsampling threshold of 10^{-3} and five iterations of training across the corpus. The transcripts were also lemmatized when using the vectors trained on the lemmatized corpus. In addition, we trained word vectors on the TASA corpus using both LSA and neural word embeddings, both using *Semantic Vectors*⁴¹ which implements a number of distributional semantics algorithms in a manner conducive to comparative evaluation (e.g. with consistent pre-processing). Both spaces were 300-dimensional. With LSA we used log-entropy weighting of terms. Neural embeddings were trained using the skipgram-with negative sampling algorithm with five negative samples per observed term, a subsampling threshold of 10^{-3} , and ten iterations of training across the corpus. For both models we excluded terms that occurred fewer than five times or more than 15,000 times. The latter constraint approximates a stopwords list, which in our experience is important for the quality of LSA vectors in particular.

Semantic units: One goal of the current work was to conduct a comprehensive analysis of the utility of modeling coherence using differently sized semantic units. The semantic units considered in this study were words, noun phrases, and sentences. Words as a unit are straightforward, in that individual embeddings can be retrieved directly from the lookup table of a vector space. For noun phrases, once extracted from documents using *textblob*, vectors were calculated by summing the vectors of individual words that composed a phrase. Similarly, a sentence vector was also calculated by summing of vectors representing component words. We explored one additional sentence vector variant by multiplying each component word vector by the relevant word’s inverse document frequency (IDF)⁴². The

Figure 1: Vector computation comparison: suppose a transcript is tokenized into three units, which are then represented by vectors $V1$, $V2$, and $V3$. The coherence metrics will be calculated as shown above.

IDF of a word is derived from the total number of documents N , and the number of documents that contain the word of interest n as $\log \frac{N}{n}$. The higher the IDF, the rarer the word, and it has been argued on theoretical grounds that IDF is an optimal measure of a word’s importance for information retrieval⁴³. The IDF of each word was obtained from distributional statistics derived from the Wikipedia corpus, and used to scale each corresponding word vector before summation. In all cases, the resulting vectors were normalized to unit length.

Centroid-derived metrics: In addition to implementing previously developed methods involving computing the vector similarity between consecutive (or “gapped”) units, we also developed and evaluated a novel method of estimating global coherence. For this method, we computed each vector’s similarity to the mean vector, or *centroid*, of a transcript. Similarity was calculated as the cosine of the angle between two vectors - one representing the centroid and the other representing a semantic unit - with the centroid calculated as the vector average of the individual unit vectors. The idea underlying this approach is that the dispersion of units from the centroid gives a measure of the extent to which they diverge in meaning from the central topic of a transcript. As this central topic may evolve as speech proceeds, we also developed and evaluated a *cumulative centroid* coherence metric where centroid of a document changes as more vectors are considered in sequence. The cumulative centroid is therefore sensitive to the position of each semantic unit within the document, and measures whether what has been said is consistent with what was said previously.

Aggregation: The sequential, gap, centroid and cumulative centroid metrics were then applied to words (all metrics), noun phrases and sentences (sequential and centroid metrics only) with and without IDF weighting for a total of 13 coherence metrics. Each metric produces a series of similarity calculations, one for each comparison it makes. For example, the sequential word-level metric produces a cosine value for each pair of neighboring words, and the centroid-based metrics produce a cosine value for the comparison between each independent unit and the centroid. Thus, the output for every metric was an array of cosine values. We then calculated the *minimum* and *mean* value of the array to evaluate their utility as transcript-level coherence scores. Our motivation for doing so was that previous studies suggested the minimum and mean of the cosine array were effective in representing coherence¹⁵.

Evaluation: For each of the metrics, an area under the curve (AUC) of a receiver operating characteristic (ROC) curve was calculated, using $1 - \text{coherence}(t)$ as an estimate of incoherence, and comparison against average human annotations with derailment score ≥ 3 labeled 1, and derailment score < 3 labeled 0. Our choice of this threshold was motivated by the immediate clinical implications of severe to extreme degrees of TD, and the likely utility of a downstream application that could detect deterioration to this point. The coherence metrics’ performance was further evaluated by computing their Spearman Rho correlation coefficient with the average of the scores assigned by human annotators. The Spearman Rho correlation is a ranked-based correlation metric that evaluates the monotonic relationship between two continuous or ordinal variables. Consequently, the Spearman Rho does not require that the two variables under consideration change together in a linear fashion, which makes it a suitable metric to evaluate the relationship between automatically generated coherence scores and human ratings. We measured the correlation between the average human rating and $1 - \text{coherence}(t)$ for each method.

Results

Aggregation: The *mean* and *minimum* aggregation methods were evaluated by comparing the number of coherence metrics that performed best in terms of ROC curve AUC and Spearman correlation with each method. With ROC curve AUC, nine out of thirteen coherence metrics performed better when summarized by the *minimum* method. The 4 exceptions were the *centroid* metric at phrase level and the *cumulative centroid* metric at word, phrase, and weighted sentence levels. With Spearman correlation, the *mean* performed better than the *minimum* for only two of thirteen coherence metrics: *centroid* and *cumulative centroid* both at phrase level. Because of the generally better performance of the *minimum*, we report results with this approach to aggregation for the remainder of the paper.

Coherence metrics: The results of our experiments comparing coherence metrics are shown in Table 3. Across both metrics (AUC and Spearman Rho) and all unit types, the best-performing metric is always one of the centroid variants, with the *cumulative centroid* ($CTRD_{cuml}$) predominating in two of the eight configurations, the *static centroid* ($CTRD_{stat}$) variant predominating in three, and these two metrics tied for best performance in the remaining three. The sentence level sequential (SEQ) model performs well with respect to AUC, but relatively poorly when considering correlation, suggesting that it is effective in identifying severe TD, but less well equipped to identify subtler manifestations of this condition. IDF weighting did not improve sentence vector performance.

Table 3: ROC Curve AUC (left) and Spearman Rho (right) for each of the metrics. **Boldface** indicates best performance across models, and underscored text indicates best performance across unit types. SEQ : sequential, GAP :gapped, $CTRD_{stat}$ and $CTRD_{cuml}$: static and cumulative variants of the centroid respectively.

AUC					Spearman Rho				
	SEQ	GAP	$CTRD_{stat}$	$CTRD_{cuml}$		SEQ	GAP	$CTRD_{stat}$	$CTRD_{cuml}$
Word	0.67	<u>0.55</u>	0.70	0.68	Word	0.21	<u>0.26</u>	<u>0.50</u>	0.51
Noun-phrase	0.69	-	0.78	0.77	Noun-phrase	<u>0.38</u>	-	0.49	0.50
Sentence	0.83	-	0.84	<u>0.83</u>	Sentence	0.26	-	0.44	0.44
Sentence with IDF	0.74	-	0.76	0.76	Sentence with IDF	0.21	-	0.41	0.41

Lemmatization: Out of thirteen coherence metrics, only two performed better with the vectors trained on the lemmatized Wikipedia corpus: the *centroid* and *cumulative centroid* at phrase level. The performance of the remaining eleven metrics with vectors trained on the original unlemmatized Wikipedia corpus was better in terms of the AUC of the ROC curve. The Spearman Rho coefficient revealed a different ratio of eight to five (instead of 11:2), but the original wikipedia-trained vectors still predominated over those trained on a lemmatized corpus. The five exceptions were *sequential* at phrase level, *centroid* at phrase, sentence and weighted sentence level, and *cumulative centroid* at weighted sentence level. Overall, lemmatization did not improve performance on the task of quantifying coherence.

Distributional models: To evaluate the influence of the underlying method of distributional semantics, LSA (a matrix-decomposition-based method) and skipgram-with-negative sampling (SGNS) (a neural network based method) were compared. Both models were trained on the TASA corpus. Of thirteen coherence metrics, nine performed better when implemented with LSA vectors in terms of ROC AUC. The four exceptions were *sequential* and *cumulative centroid* at word level and *centroid* and *cumulative centroid* at phrase level. Similar results were observed with Spearman correlation, aside from that instead of *sequential* at word level, it was *sequential* at phrase level where SGNS performed better. These results show LSA outperforming neural embeddings when trained on a relatively small corpus, a finding consistent with previous research⁴⁴. However, we note that the performance of either TASA-trained space in the majority of metrics was exceeded by the performance of models using neural embeddings trained on Common Crawl. Thus, while LSA appears to offer advantages when restricted to smaller corpora, the capacity of neural embedding models to scale to much larger corpora appears advantageous for automated estimates of coherence.

Discussion

In this paper we present a comprehensive study of automated methods of measuring speech coherence in the context of transcripts of short (≤ 3 minutes) responses to an open-ended prompt. When evaluated for their agreement with human annotators, our results show strong performance for two novel coherence metrics: the *centroid* and *cumulative centroid*. When considering the consistency with speech coherence level rated by humans, the two novel

metrics outperformed the established *sequential* and *gap* metrics of coherence in terms of both their ability to detect severe cases (as estimated by the AUC of the ROC curve) and their correlation with average annotator scores across all categories of severity (as estimated by the Spearman Rho coefficient). This observation holds true for all semantic units considered in this study: words, noun phrases, sentences and IDF-weighted sentences. For detection of severe cases, the best performing metric of coherence is the *centroid*, while the *cumulative centroid* performed slightly better with respect to overall correlation. The *centroid* measure attained a AUC of the ROC curve above 0.7 for all semantic unit types with some above 0.8. For overall correlation, this metric attained a Spearman Rho coefficient larger than 0.4, with in some cases larger than 0.5. While not presented in detail on account of space constraints, we note that the centroid-based methods performed best across all of the vector spaces generated during the course of this research, whether derived from TASA, Wikipedia or Common Crawl, and irrespective of whether neural embeddings or LSA were used. These findings suggest that in the context of short unstructured speech samples, coherence metrics using distributional similarity perform better when modeling global coherence with a centroid vector.

When considering different semantic units, using a sentence as a unit performed best on the task of identifying severe cases, with AUC values above 0.8. For overall correlation, using a word as a unit led to best performance with *centroid* metrics, while noun phrase units performed best with the *sequential* metric. The disparity between these findings may be due to the nature of the tasks concerned. The AUC measures the ability of the model to predict positive cases at relatively low false positive rate, with “positive” in our case indicating severe derailment. Thus, the AUC measure focuses on the coherence metric’s ability to identify severely incoherent speech. On the other hand, the Spearman Rho coefficient is a rank-based correlation measure that takes into consideration all coherence categories. It does not require the imposition of a dichotomous classification threshold like the AUC and thus, it measures the overall prediction quality of the coherence metrics. Therefore, the sentence semantic unit appears best for identifying severely incoherent cases, and word or phrase units appear best used to model subtler distinctions in coherence.

To focus on detection of manifestations of severe TD, we set a threshold at an average human rating of three to calculate the AUC. However, detection of milder degrees of TD is also of interest, and previous work, albeit with a different rating scale based on clinical observation rather than text, has employed a threshold of two on a five point scale to identify TD¹. To verify the consistency of our findings at a different threshold level, we also computed the AUC of the ROC curve with threshold of two, with a more than threefold increase in the number of positive examples. Our main findings were consistent at this threshold: the *centroid* measures still outperformed the *sequential* and *gap* measures with every type of semantic unit. Of note, a difference is that the *cumulative centroid* measures now have higher AUCs than their static counterpart (best AUC of 0.78 at phrase level), which is consistent with this metric’s better performance for correlation across all levels of severity. The effects of aggregation on metrics of coherence have rarely been explicitly compared. Bedi et al¹⁵ used the *minimum* aggregate as a predictor for a classification model, whereas Elvevag et al¹ used the *mean* aggregate to generate LSA-derived coherence scores. Experiments with aggregates in this study suggest the *minimum* of a set of similarity values performs better than the *mean* in most cases. The few exceptions are from the *centroid* metrics (best AUC of 0.83 for *cumulative centroid* at phrase level), indicating the *minimum* aggregate does not impair performance of *sequential* or *gap* coherence metrics.

We also examined the effects of lemmatization of the corpus used to train word embeddings (as well as the transcripts themselves). Our findings with word embedding methods suggest lemmatization does not improve the performance of most coherence methods, which is consistent with previous work evaluating the utility of lemmatization for automated grading of summaries for coherence⁴⁵. There is some degree of disparity between the AUC and Spearman Rho when considering the ratios of number of best performing measures (unlemmatized:lemmatized - 11:2 for AUC, 8:5 for Spearman Rho). However, this is likely due to the different aspects of performance these metrics focus on, as discussed earlier. As coherence metrics without lemmatization generally outperform their lemmatized counterparts, the use of lemmatization is not recommended. In addition, the comparison between LSA and neural word embeddings trained on the TASA corpus found LSA vectors performed better with most coherence measures. This finding suggests that neural embeddings are not inherently better than LSA for automated estimates of coherence. LSA remains an robust alternative to generate word vectors, especially with a small corpus, which is consistent with prior work comparing these models⁴⁴. The main advantage of neural embeddings over LSA appears to be attributable to the availability of neural embeddings trained on larger and more comprehensive corpora. This advantage was substantial with the best-performing sentence-level *centroid* metrics where improvements in performance of $\approx 11\%$ in AUC were observed

with the vectors trained on Common Crawl as compared with those trained on the TASA corpus.

A limitation of this study is that only one construct (i.e. derailment) was used to evaluate the degree of TD. As argued previously, we considered derailment to be the TD construct most applicable to the evaluation of responses to open-ended prompts. Nonetheless, there are a number of other TD-related constructs that have yet to be modeled using *centroid*-based methods. These constructs include tangentiality, logorrhoea and poverty of speech¹⁴, which if accurately recognized would result in a more granular automated system for the characterization of linguistic manifestations of disordered thinking. Cosine-based coherence metrics alone may be inadequate for this task; suggesting a need to incorporate more automated measurements such as graph-based measures¹⁷ and measures considering verbosity in relation to topical breadth⁴⁶. However, this is beyond the scope of the current paper. In future work we plan to evaluate the utility of the coherence measures from this study as features for predictive models of clinically important outcomes, such as use of mental health services. Studies suggest that key differences between AVH patients with and without the need for care include whether normal functioning is affected, and whether AVH content is negative^{47,48}. The coherence metrics from this study may serve as a proxy for level of function as coherence of speech is prerequisite to functional communication. Combining coherence with sentiment analysis to measure the negativity of speech content may be of value in anticipating need for care. We also plan to develop automated models of emotional content on account of the prevalence of TD and AVH in severe mood disorders⁴⁹.

Conclusion

In this paper, we compared the performance of novel *centroid*-based estimates of speech coherence to established *sequential* measures in the context of transcribed recordings of responses to an open-ended prompt. The novel methods agreed better with human annotation both for detection of severe cases, and in terms of overall coherence. In addition, we evaluated a number of methodological alternatives, providing guidance for future efforts toward automated detection of linguistic manifestations of disordered thinking.

Acknowledgments

This research was supported by National Institute of Mental Health grant 3R01MH112641-03S2.

References

- [1] Brita Elvevåg et al. “Quantifying incoherence in speech: An automated methodology and novel application to schizophrenia”. In: *Schizophrenia Research* 93.1 (2007), pp. 304–316. ISSN: 0920-9964. DOI: <https://doi.org/10.1016/j.schres.2007.03.001>. URL: <http://www.sciencedirect.com/science/article/pii/S092099640700117X>.
- [2] S. Hossein Fatemi and Paula J. Clayton. *The medical basis of psychiatry*. Springer, 2016.
- [3] Nancy C Andreasen and GARY J TUCKER. “Introductory textbook of psychiatry”. In: *American Journal of Psychiatry* 148.5 (1991), pp. 670–670.
- [4] Nancy C Andreasen and William M Grove. “Thought, language, and communication in schizophrenia: diagnosis and prognosis”. In: *Schizophrenia bulletin* 12.3 (1986), pp. 348–359.
- [5] Sukanta Saha et al. “A systematic review of the prevalence of schizophrenia”. In: *PLoS medicine* (May 2005). URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1140952/>.
- [6] Julio Bobes et al. “Quality of life in schizophrenic patients”. In: *Dialogues in clinical neuroscience* (2007). URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3181847/#ref2>.
- [7] J. Meg Racenstein et al. “Thought Disorder and Psychosocial Functioning in Schizophrenia”. In: *The Journal of Nervous Mental Disease* 187.5 (1999), pp. 281–289. DOI: 10.1097/00005053-199905000-00003.
- [8] Ross M.G. Norman et al. *Symptoms and Cognition as Predictors of Community Functioning: A Prospective Analysis*. Mar. 1999. URL: <https://ajp.psychiatryonline.org/doi/full/10.1176/ajp.156.3.400>.
- [9] Nancy C Andreasen and William M Grove. “Thought, language, and communication in schizophrenia: diagnosis and prognosis”. In: *Schizophrenia bulletin* 12.3 (1986), pp. 348–359.

- [10] Paolo Fusar-Poli, Patrick D McGorry, and John M Kane. “Improving outcomes of first-episode psychosis: an overview”. In: *World psychiatry* 16.3 (2017), pp. 251–265.
- [11] Nancy C Andreasen. “Scale for the assessment of thought, language, and communication (TLC).” In: *Schizophrenia bulletin* 12.3 (1986), p. 473.
- [12] Alvaro Barrera, Peter J. McKenna, and German E. Berrios. “Two new scales of formal thought disorder in schizophrenia”. In: *Psychiatry Research* 157.1 (2008), pp. 225–234. ISSN: 0165-1781. DOI: <https://doi.org/10.1016/j.psychres.2006.09.017>.
- [13] Peter F. Liddle et al. “Thought and Language Index: an instrument for assessing thought and language in schizophrenia”. In: *British Journal of Psychiatry* 181.4 (2002), pp. 326–330. DOI: 10.1192/bjp.181.4.326.
- [14] Tilo Kircher et al. “A rating scale for the assessment of objective and subjective formal Thought and Language Disorder (TALD)”. In: *Schizophrenia Research* 160.1 (2014), pp. 216–221. ISSN: 0920-9964. DOI: <https://doi.org/10.1016/j.schres.2014.10.024>.
- [15] Gillinder Bedi et al. “Automated analysis of free speech predicts psychosis onset in high-risk youths”. In: *npj Schizophrenia* 1.1 (2015), p. 15030. ISSN: 2334-265X. DOI: 10.1038/npjrschz.2015.30. URL: <https://doi.org/10.1038/npjrschz.2015.30>.
- [16] Cheryl M Corcoran et al. “Prediction of psychosis across protocols and risk cohorts using automated language analysis”. In: *World Psychiatry* 17.1 (2018), pp. 67–75.
- [17] Natalia B Mota et al. “Speech graphs provide a quantitative measure of thought disorder in psychosis”. In: *PloS one* 7.4 (2012).
- [18] Peter W Foltz, Walter Kintsch, and Thomas K Landauer. “The measurement of textual coherence with latent semantic analysis”. In: *Discourse processes* 25.2-3 (1998), pp. 285–307.
- [19] Scott Deerwester et al. “Indexing by latent semantic analysis”. In: *Journal of the American society for information science* 41.6 (1990), pp. 391–407.
- [20] Thomas K Landauer, Peter W Foltz, and Darrell Laham. “An introduction to latent semantic analysis”. In: *Discourse processes* 25.2-3 (1998), pp. 259–284.
- [21] Trevor Cohen and Dominic Widdows. “Empirical distributional semantics: methods and biomedical applications”. In: *Journal of biomedical informatics* 42.2 (2009), pp. 390–405.
- [22] Peter D Turney and Patrick Pantel. “From frequency to meaning: Vector space models of semantics”. In: *Journal of artificial intelligence research* 37 (2010), pp. 141–188.
- [23] Brita Elvevåg et al. “An automated method to analyze language use in patients with schizophrenia and their first-degree relatives”. In: *Journal of neurolinguistics* 23.3 (2010), pp. 270–284.
- [24] Charles Ellis et al. “Global coherence during discourse production in adults: A review of the literature”. In: *International journal of language & communication disorders* 51.4 (2016), pp. 359–367.
- [25] P. K. McGuire et al. “The Neural Correlates of Inner Speech and Auditory Verbal Imagery in Schizophrenia: Relationship to Auditory Verbal Hallucinations”. In: *British Journal of Psychiatry* 169.2 (1996), pp. 148–159. DOI: 10.1192/bjp.169.2.148.
- [26] Marcia Radanovic et al. “Formal Thought Disorder and language impairment in schizophrenia”. In: *Arquivos de neuro-psiquiatria* 71.1 (2013), pp. 55–60.
- [27] Edward Loper and Steven Bird. “NLTK: the natural language toolkit”. In: *arXiv preprint cs/0205028* (2002).
- [28] Steven Loria et al. “Textblob: simplified text processing”. In: *Secondary TextBlob: Simplified Text Processing* 3 (2014).
- [29] Tomas Mikolov et al. “Efficient estimation of word representations in vector space”. In: *arXiv preprint arXiv:1301.3781* (2013).
- [30] *Word2Vec software package*. URL: <https://github.com/tmikolov/word2vec>.

- [31] *Fasttext software package*. URL: <https://fasttext.cc/>.
- [32] Marco Baroni, Georgiana Dinu, and Germán Kruszewski. “Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors”. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2014, pp. 238–247.
- [33] Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. “Linguistic Regularities in Continuous Space Word Representations”. In: *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Atlanta, Georgia: Association for Computational Linguistics, June 2013, pp. 746–751. URL: <https://www.aclweb.org/anthology/N13-1090>.
- [34] Omer Levy, Yoav Goldberg, and Ido Dagan. “Improving distributional similarity with lessons learned from word embeddings”. In: *Transactions of the Association for Computational Linguistics* 3 (2015), pp. 211–225.
- [35] Dan Iter, Jong Yoon, and Dan Jurafsky. “Automatic Detection of Incoherent Speech for Diagnosing Schizophrenia”. In: *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*. New Orleans, LA: Association for Computational Linguistics, June 2018, pp. 136–146. DOI: 10.18653/v1/W18-0615. URL: <https://www.aclweb.org/anthology/W18-0615>.
- [36] *Fasttext pretrained vectors*. URL: <https://fasttext.cc/docs/en/english-vectors.html>.
- [37] *Common crawl corpus*. URL: <https://commoncrawl.org/>.
- [38] *Gensim software package*. URL: <https://radimrehurek.com/gensim/models/word2vec.html>.
- [39] Radim Řehůřek and Petr Sojka. “Software Framework for Topic Modelling with Large Corpora”. English. In: *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. <http://is.muni.cz/publication/884893/en>. Valletta, Malta: ELRA, May 2010, pp. 45–50.
- [40] Tomas Mikolov et al. “Distributed representations of words and phrases and their compositionality”. In: *Advances in neural information processing systems*. 2013, pp. 3111–3119.
- [41] *Semantic vectors software package*. URL: <https://github.com/semanticvectors/semanticvectors>.
- [42] Karen Sparck Jones. “A statistical interpretation of term specificity and its application in retrieval”. In: *Journal of documentation* (1972).
- [43] Kishore Papineni. “Why Inverse Document Frequency?”. In: *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics on Language Technologies*. NAACL ’01. Pittsburgh, Pennsylvania: Association for Computational Linguistics, 2001, pp. 1–8. DOI: 10.3115/1073336.1073340. URL: <https://doi.org/10.3115/1073336.1073340>.
- [44] Edgar Altszyler et al. “Comparative study of LSA vs Word2vec embeddings in small corpora: a case study in dreams database”. In: *arXiv preprint arXiv:1610.01520* (2016).
- [45] Iraide Zipitria, Ana Arruarte, and Jon Ander Elorriaga. “Observing Lemmatization Effect in LSA Coherence and Comprehension Grading of Learner Summaries”. In: *Intelligent Tutoring Systems*. Ed. by Mitsuru Ikeda, Kevin D. Ashley, and Tak-Wai Chan. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 595–603. ISBN: 978-3-540-35160-3.
- [46] Neguine Rezaii, Elaine Walker, and Phillip Wolff. “A machine learning approach to predicting psychosis using semantic density and latent content analysis”. In: *npj Schizophrenia* 5.1 (2019), pp. 1–12.
- [47] Louise C. Johns et al. “Auditory Verbal Hallucinations in Persons With and Without a Need for Care”. In: *Schizophrenia Bulletin* 40.Suppl_4 (June 2014), S255–S264. ISSN: 0586-7614. DOI: 10.1093/schbul/sbu005. eprint: https://academic.oup.com/schizophreniabulletin/article-pdf/40/Suppl_4/S255/16704110/sbu005.pdf. URL: <https://doi.org/10.1093/schbul/sbu005>.
- [48] David Baumeister et al. “Auditory verbal hallucinations and continuum models of psychosis: A systematic review of the healthy voice-hearer literature”. In: *Clinical Psychology Review* 51 (2017), pp. 125–141. ISSN: 0272-7358. DOI: <https://doi.org/10.1016/j.cpr.2016.10.010>. URL: <http://www.sciencedirect.com/science/article/pii/S0272735816301064>.

- [49] C. Raymond Lake. "Disorders of thought are severe mood disorders: the selective attention defect in mania challenges the Kraepelinian dichotomy a review". In: *Schizophrenia bulletin* 34.1 (Jan. 2008). sbm035[PII], pp. 109–117. ISSN: 0586-7614. DOI: 10.1093/schbul/sbm035. URL: <https://pubmed.ncbi.nlm.nih.gov/17515440>.