Automated Scoring of the Autobiographical Interview with Natural Language Processing

Ruben D.I. van Genugten and Daniel L. Schacter

Department of Psychology, Harvard University, Cambridge, MA, USA

Corresponding Author:

Ruben D.I. van Genugten (ruben_vangenugten@g.harvard.edu)

Acknowledgments

We thank Alea Devitt, Signy Sheldon, and Peggy St. Jacques for sharing data with us. We thank Patrick Mair for discussions related to this manuscript. Preparation of this manuscript was supported by National Institute on Aging [grant number R01 AG008441] awarded to DLS.

Disclosure of interest

The authors report no conflict of interest.

Abstract

The Autobiographical Interview has been used in more than two hundred studies to assess the content of autobiographical memories. In a typical experiment, participants recall memories, which are then scored manually for internal details (episodic details from the central event) and external details (largely non-episodic details). Scoring these narratives requires a significant amount of time. As a result, large studies with this procedure are often impractical, and even conducting small studies is time-consuming. To reduce scoring burden and enable larger studies, we developed an approach to automatically score responses with natural language processing. We fine-tuned an existing language model (distilBERT) to identify the amount of internal and external content in each sentence. These predictions were aggregated to obtain internal and external content estimates for each narrative. We evaluated our model by comparing manual scores with automated scores in five datasets. We found that our model performed well across datasets. In four datasets, we found a strong correlation between internal detail counts and the amount of predicted internal content. In these datasets, manual and automated external scores were also strongly correlated, and we found minimal misclassification of content. In a fifth dataset, our model performed well after additional preprocessing. To make automated scoring available to other researchers, we provide a Colab notebook that is intended to be used without additional coding.

Key words: Autobiographical Interview, Autobiographical Memory, Automated Scoring, Natural Language Processing, distilBERT

The Autobiographical Interview (AI) (Levine et al., 2002) is a widely used method to study the contents of participants' autobiographical memories. In a typical experiment using the AI, participants are asked to recall a specific event for each cue, making sure to report as much detail as they can. Several human raters then use a manual to identify and count internal details (central episodic details) and external details (mostly non-episodic details) in the narratives. Levine et al. (2002) first developed the AI and its scoring manual to study age-related differences in memory. With this procedure, Levine et al. (2002) showed that older adults provide fewer internal details than young adults when they were asked to retrieve autobiographical memories that are more than a year old, despite being able to provide the same number of external details. For memories that were less than a year old, older adult provided more external details than younger adults. These findings extended our understanding of the effect of aging on autobiographical memory and in so doing, also highlights that the AI can be used to study group differences in memory.

Studies conducted since Levine et al. (2002) have made it clear that the AI enables researchers to test specific theories about memory. For example, Addis et al. (2008) hypothesized that if episodic retrieval supports the construction of specific future events, age-related decreases in internal details during retrieval of past events should be accompanied by age-related decreases in internal details when imagining future events. Addis et al. (2008) adapted the AI to ask participants about imagined future events and remembered events. In addition to finding support for their hypothesis, they found a positive correlation between the number of internal details participants provided when remembering the past and imagining the future. These results are consistent with the constructive episodic simulation hypothesis, which suggests that we imagine future experiences by recombining elements from episodic memories

(Schacter & Addis, 2007). Other researchers have similarly adapted the AI to study future thinking. For example, Race et al. (2011) studied amnesic individuals with the AI, and showed that damage to the medial temporal lobe led to reductions in central episodic details on both episodic memory and future thinking tasks, despite intact descriptive abilities (i.e., normal performance when describing pictures).

Because the AI can be used to study features of autobiographical memory and future thinking, the procedure has seen widespread use across domains of psychology. In addition to the aforementioned studies of aging and amnesia, it has been used to study Alzheimer's disease (e.g., Irish et al., 2011), depressive disorders (e.g., Söderlund et al., 2014), and the contributions of episodic retrieval to other domains of cognition, such as means-end problem solving (e.g., Madore & Schacter, 2014). As of November 2021, over two hundred studies have used this interview (Levine, 2021), and the paper that first described the interview has over 1400 citations listed on Google Scholar.

The AI is widely used in psychology research despite the fact that scoring the AI takes a lot of time and effort (typically 10 minutes per memory, and each participant may provide ten or twenty memories). Having to manually annotate hundreds of pages of memories potentially limits the AI's usefulness and breadth of applications. Because large studies using the AI are impractical, past research typically studied only effects that could be detected with small samples (e.g., approximately thirty participants).

Here we introduce an automated scoring procedure for the AI that can reduce experimenter burden and help researchers to conduct larger experiments and study smaller effects. We believe that this new procedure will broaden the scope of research questions the field can address. Our automated scoring approach could also make online data collection more

practical. Online data collection is rarely used with the AI, likely because large numbers of participants are needed to compensate for noisy data. Online studies would allow researchers to gain access to a larger and more diverse population than the typical samples that have been used.

In the remainder of this paper, we will describe how narratives are typically scored, how researchers have attempted to streamline scoring, and our new approach for automating AI scoring.

Current approach for manually scoring memory details

Researchers follow a set of rules from the AI to score narratives. These rules explain how to classify pieces of text as internal or external, and how to identify bits of information within these segments that count as details.

Internal details refer to episodic details, and external details refer to non-episodic details (or episodic details that do not correspond to the central event being remembered or imagined). Internal details describe components of an event that are specific to time and place. The event's location and time, the people, objects, actions, thoughts, and perceptual details involved are all internal details. External details, on the other hand, are largely non-episodic details. These are any details that do not belong in the internal details category, and largely consist of factual information that does not require the participant to remember or imagine a specific event (e.g., "I've always enjoyed going to the beach for my birthday"). Participants sometimes provide information about events other than the central event they are being asked to describe. These details, while episodic in nature, are considered external details as well. Lastly, repetitive information (e.g., someone describes the same thing twice), and information unrelated to the event the participant is trying to describe (e.g., 'sorry for that cough!') are also considered external.

The manual for the AI provides clear rules on how to divide up segments of text into individual details. For internal segments, each piece of information that tells us something more about the event is generally counted as a detail. For example, “he had a hat” is considered one detail. Any additional descriptors count as additional details, e.g., “he had a brown hat” is considered two details. This description is illustrative of the general approach; for an exhaustive list of rules and exceptions, see the manual available upon request from Dr. Brian Levine (blevine@research.baycrest.org). A brief scoring example is provided in the appendix.

Existing automation approaches for memories

Several researchers have streamlined scoring of the AI, yet no group has fully automated the scoring of details in narratives. Previous work consists of two approaches: speeding up the processes involved in scoring, and predicting the number of internal and external details. For example, Wardell, Esposito, et al. (2021) automated the process of transcribing spoken narratives to text with Dragon NaturallySpeaking software. The researchers also reduced the time necessary for scoring by setting up keyboard shortcuts in Microsoft Word. Once details were manually scored, their software automatically counted the scored details in each memory. After implementing a protocol of this kind, a research group would be able to score more rapidly (see, e.g., Wardell, Madan, et al., 2021). However, much of the work remains to be done by hand: identifying internal and external content and separating the narratives into details are both still done manually.

To the best of our knowledge, only one paper reports an attempt to automatically generate AI detail scores for each narrative. Peters et al. (2017) first extracted eighty-three features from each narrative, such as the number of emotion words, the valence of these emotion words, the number of words in the story, and the number of nouns in the story. Peters et al. then

used these features in several regression models (e.g., principal component regression) to predict the number of internal and semantic details each participant provided (summed across five or twelve narratives per participant). When Peters et al. predicted the number of internal details provided by each participant, Root Mean Square Error (RMSE) was approximately .5 for internal details and .65 for semantic details. Peters et al. report one model built to predict the number of internal details in individual narratives. RMSE for this model was approximately .75 for episodic future thinking narratives and .85 for autobiographical memory narratives. No model was reported that predicted semantic details in individual narratives. To contextualize these results, a simple model that predicts the mean number of internal or semantic details for every narrative in this dataset would result in $RMSE = 1$, while a model with perfect predictions would result in $RMSE = 0$. Importantly, word count was a significant predictor for models predicting internal and semantic details counts. Since predictions were driven in part by the total amount of content, these regression models are presumably misclassifying internal content as semantic, and vice versa. So, while these researchers took an important first step by attempting to automatically score the AI, their predicted memory scores differed significantly from the actual memory scores, and additional work is needed to automate AI scoring.

Related work (Takano et al., 2017, 2018, 2019) has automated scoring of the Autobiographical Memory Test (Williams & Broadbent, 1986) with more success. When using the Autobiographical Memory Test, human raters classify memories as specific or general. Takano et al. used word frequencies and parts of speech frequencies to train a classifier to determine which memories were specific and which memories were general. Across studies, Takano et al. report good classification results, with high accuracy, frequent correct identification of specific memories, and frequent correct identification of general memories. For

example, for narratives from English-speaking adults reported in Takano et al. (2019), classification was 81.1% accurate. Specific memories were correctly identified in 81.8 % of cases, and general memories were correctly identified in 80.3% of cases. These results suggest that natural language processing provides a promising path for automated memory interview scoring.

Our automated scoring approaches

To improve automated scoring accuracy of the AI, we relied on advances in natural language processing to identify the amount of internal and external content in each sentence of an AI narrative. After classifying each sentence, we counted the amount of internal and external content in each narrative and validated these counts against detail counts obtained through manual scoring.

We trained our classifier by using data scored according to the AI. Specifically, we used these data to fine-tune weights at the end of an existing neural network, which had previously been trained on different natural language tasks. This procedure allowed us to take advantage of the language representations that the neural network had previously learned. This process, known as *transfer learning*, is a standard approach for classifying language content according to new labels, especially when few training examples are available (for introduction, see e.g. Azunre, 2021). Specifically, we fine-tuned distilBERT (Sanh et al., 2019) with the ‘huggingface’ library (Wolf et al., 2020).

We trained and evaluated our model with five datasets, which involved data scored according to the standard or adapted AI. We found that our code accurately identified internal and external content, with minimal misclassification of internal content as external, and minimal misclassification of external content as internal.

Methods

Model Training and Evaluation Data

To train our model to classify the amount of internal and external content in sentences, we requested data from several different researchers. All data we used were previously scored on a computer using standard or adapted AI scoring manuals. These data spanned several different tasks. Three datasets contained autobiographical memories (King et al., 2021; Sheldon et al., 2020; Strikwerda-Brown et al., 2021), and one of these contained data from both younger and older adults (Sheldon et al., 2020). Another dataset contained future simulation data from younger and older adults (Devitt & Schacter, 2018, 2019). We also included data from a study on creative writing (van Genugten et al., 2021) that was scored using an adapted AI scoring manual. These data were included to test whether the model would generalize to non-memory or future simulation paradigm that used adapted AI scoring. Last, one dataset included a picture description task and an open-ended thoughts description task (Strikwerda-Brown et al., 2021). These data were scored with guidelines that were different from the adapted or standard AI manuals. So, these data were included for exploratory analyses, without the expectation that our model would perform well on them. Because they were scored differently from all the other datasets, they were never included in the training sets. Each of these datasets is described in more details below.

Dataset 1: Autobiographical memories (King, Romero, Schacter, & St. Jacques, 2021)

King et al. (2021) examined how retrieving memories from an observer perspective (as opposed to a first-person perspective) changed the narratives. In the first session of this study, participants were asked to elaborate on a subset of memories in which they rated the event as occurring through their own eyes (at least a 5 on a 7-point scale measuring self-perspective). We

used these memories for our analyses. These data were in written form and were scored according to the AI (Levine et al., 2002).

This study generated a dataset of 40 individuals (25 female). Participants were, on average, 23.33 years old ($SD = 3.17$). All participants indicated that they were not previously diagnosed with a mood or cognitive disorder, nor taking any medication that could affect performance on the study. All participants were recruited from the Harvard study pool and the community.

Data Set 2: Autobiographical Memories (Sheldon et al., 2020)

Sheldon et al. (2020) collected autobiographical memories to test whether cue valence and arousal affected subsequent retrieval and elaboration of memories. In this experiment, participants listened to a series of 24 musical excerpts, which served as retrieval cues. After each retrieval cue, participants wrote down a caption to describe the memory they had retrieved. In a second session, participants were presented with the captions they had previously written down, given 30 seconds to remember the memory, and then used two minutes to describe what they remembered. Responses were audio-recorded and transcribed. Responses were then scored using the standard scoring guidelines from the AI (Levine et al., 2002).

Participants were recruited from McGill University's study pool. Each of the 42 participant was fluent in English and free of major neurological or psychiatric disorders. Participants were on average 20 years old ($SD = 1.4$) and had 14.6 years of education ($SD = 1.1$). 37 of the participants were female.

Data Set 3: Autobiographical Memories, Thoughts, and Picture Descriptions. (Strikwerda-Brown et al., 2021)

Strikwerda-Brown et al. (2021) investigated age-related changes in memory on a cued retrieval task and an open-ended task. Participants also completed a picture description task (cf., Gaesser et al., 2011). On each trial of the experiment, the participants saw a picture, were asked to retrieve a memory related to the image (memory task), to describe what was present in the image as if to someone who could not see the image (description task), or to describe the thoughts that arise when viewing the picture (thoughts task). All narratives were transcribed after being verbally reported by the participants.

Scoring of the memory task followed guidelines developed in Levine et al. (2002), with modified scoring guidelines for external details as described in Strikwerda-Brown et al. (2019). The picture description task was scored by following guidelines developed by Gaesser et al. (2011). Perceptual details in the picture were scored as internal details, and all other details (e.g., inferences about the picture, general comments about the picture) were scored as external. Details in the thoughts task were considered internal if they described any past event; all other details were considered external.

24 older adults and 25 younger adults were included in the analysis of this study. Younger adults were, on average, 21.7 years old ($SD = 2.4$). Participants reported no neurological or psychiatric impairments that would affect the study. Older adults were recruited from an existing database of older adults in the Montreal area. Younger adults were recruited from the McGill University study pool and surrounding areas.

Data Set 4: Future Simulation: Young Adult and Older Adult Data (Devitt & Schacter, 2018, 2019)

Devitt & Schacter (2018, 2019) examined how episodic simulation of an event before learning of its outcome affected the subsequent memory of that outcome. In their studies,

participants were presented with a series of cues for future events and were instructed to imagine the events going well or poorly for 3 minutes. Participants were instructed that each imagined event should occur within the next year. Afterwards, participants were given descriptions of how the events happened. In a second session, participants were tested for their memory of how the event happened.

Across two studies, future simulations from older and younger adults were audio-recorded, transcribed, and scored according to the AI. Data from these experiments include 27 younger adults (mean age = 22.59 years, SD = 3.18, 12 male) and 25 older adults (mean age = 72.24, SD = 6.49; 7 male). Participants indicated no history of neurological or psychiatric impairment. These participants were recruited from Harvard University and the surrounding community, using the Harvard psychology study pool.

Data set 5: Creative Writing Narratives. (van Genugten et al., 2021)

van Genugten et al. tested whether episodic retrieval contributes to creative writing performance. Specifically, van Genugten et al. used the Episodic Specificity Induction (Madore et al., 2014; for review, see Schacter & Madore, 2016) to manipulate episodic retrieval prior to a creative writing task. Detail counts after the ESI were compared to detail counts after two control inductions.

In the creative writing task, participants read a series of excerpts from literature and were asked to continue writing each story in a style that felt natural to them. Each story was scored according to scoring guidelines from the ESI studies of Madore et al. (2014) and Jing et al. (2016), which were adapted from the standard AI scoring (Levine et al., 2002). In their scoring, Van Genugten et al. also considered all event details as internal details. This procedure differs from previous guidelines, which only considered details from the central events to be internal.

This change was made to ensure that no episodic details were marked as external. Data from the first experiment were scored by hand, and as such were not used in the training and evaluation of our model. Data from the second experiment of this study were used since scoring was done on the computer.

Data used in this paper come from 32 participants, who each wrote 10 stories. Participants were young (18-30 years old, $M = 24.03$ years, $SD = 3.51$; 21 female, 11 male) and recruited from the Harvard University study pool. No participant reported neurological or psychiatric impairment at the time of the study.

Data Preparation

Data were read in from various sources, including text files and SciTos (Wickner, Englert, & Addis, 2015) html exports. Any prompting by the researcher (e.g., ‘tell me more about that’) was removed. Data were manipulated so that formatting was identical across datasets. Because our approach classifies individual sentences, narratives were split into sentences using pySBD (Sadvilkar & Neumann, 2020). pySBD splits text into sentences based on 48 rules that rely in part on punctuation.

Additional preprocessing was necessary after we noticed that some exceptionally long sentences contained a large majority of the narrative they came from (or even the full narrative). These narratives were transcribed with little or no punctuation, leading to few sentence splits by pySBD. To mimic narratives transcribed with full punctuation, we removed sentences that contained more than 8 details, since these sentences are likely missing punctuation. Detail counts associated with the narratives, which we used for validation, were updated to reflect the removal of this content. This preprocessing step is not included in the code we make available that other researchers can use to automatically score their own narratives. Researchers who want to use our

model in their own research should add punctuation as they are transcribing, to accommodate sentence splitting by pySBD. Alternatively, participants can be asked to type narratives, so that researchers do not have to add punctuation as they transcribe.

Each sentence was classified as belonging to one of four categories: containing 0% internal content (i.e., 100 % external content), 50% internal content, 75% internal content, or 100% internal content. We modified training datasets such that there were an equal number of sentences in all four categories. We did this by identifying the category with the greatest number of sentences, and upsampling data from all other categories. So, for example, if a training dataset were to contain 10,000 fully internal sentences, and 8,000 sentences from each of the three remaining categories, we would sample 2,000 sentences with replacement from each of those three categories, then add those sentences to the dataset so that we have 10,000 training examples in each category. We used training data with an equal number of examples for each label because this procedure prevented the model from learning to use relative frequencies of internal and external details to improve prediction accuracy. This step is necessary because if we did not upsample our training sets, and our narratives contained many more internal details than external details, the model could obtain relatively high accuracy by classifying all details as internal.

Model Training and Evaluation

We trained and evaluated the performance of our classification model with five datasets that are described in more detail in *Model Training and Evaluation Data*. We iteratively left out one dataset for evaluation, using the other four for training. For some datasets, data from multiple tasks or experiments were available. When these datasets served as the testing set, performance on each task was separately evaluated. For example, one dataset involved future

simulation for older adults and younger adults. When this dataset was left out for evaluation, the model was trained on the other four datasets and was then separately evaluated on the older adult data and the younger adult data. We report performance of all evaluation sets separately. Picture description and thoughts tasks from Strikwerda-Brown et al. (2021) were never included in training data because they were not scored with the adapted or standard AI.

As we trained and evaluated our model, we aimed to mimic the AI scoring, which involves two primary steps: annotating information as internal or external, and separating text segments into individual details. Below, we outline our approach to both steps.

Separating text segments into details

The AI scoring manual provides guidelines on how to split internal and external text segments into individual details, which allows researchers to quantify how much content is present in these segments. We chose not to train a model to split segments into individual details, since we have near-perfect proxies for internal detail counts and external detail counts: the number of words in internal segments, and the number of words in external segments.

To determine if internal word count adequately captures the number of internal details in narratives, we examined the correlation between these two variables across our datasets. The correlation between the number of internal details and internal word count ranged from .86 to .98 (mean = .92). We repeated this process for external details. The correlation between the number of external details and external word count ranged from .87 to .98 (mean = .94). Together, these extremely high correlations suggest that we can use internal and external word counts to quantify the amount of internal and external information. That is, we do not need to split sentences into individual details for the purposes of this project. An example of scoring with internal and external details and internal and external word counts is provided in the appendix.

We should note that internal and external word counts would not be adequate for all circumstances. For example, splitting content into individual details may be helpful for researchers that want to assign each detail to a subcategory (e.g., place, time, perceptual etc.). However, because the purpose of this project is simply to quantify the amount of internal and external content in narratives, we used internal and external word count as excellent approximations of internal and external details.

Classifying Information as Internal or External: Overview

To identify internal and external information, we adapted a common approach for classifying text. We fine-tuned weights at the end of an existing neural network with new data. We trained our model to classify sentences as containing only external content, 50% internal content, 75% internal content, or 100% internal content. To select classification labels, we calculated the percent of internal content in each sentence in the first dataset we obtained (Devitt & Schacter, 2018, 2019). A histogram of these percentages showed clusters at approximately 0, 50, 75 and 100%; hence, our labels.

Classifying Information as Internal/External: Model Specifics

To identify internal and external information, we used a model designed to be fine-tuned on sentences for classification. Specifically, we used distilBERT (Sanh et al., 2019), which is a language representation model that can be fine-tuned to new tasks by adding a classification head. The classification head contains a single linear layer for classification at the end of the network's pooled output. Fine-tuning this model involves changing the weights of this last layer to improve predictions on the fine-tuning data. distilBERT provides state-of-the-art performance on a range of natural language processing benchmark tests, while using fewer parameters than its ancestor BERT (Devlin et al., 2019). distilBERT has been trained to mimic BERT's performance

on two tasks: masked word prediction and next sentence prediction. Training with next sentence prediction involves providing the model with pairs of sentences. For each pair of sentences, the model must determine whether the second sentence followed the first sentence in the source text, or whether that pair of sentences is randomly paired. Training with masked word prediction involves randomly masking a subset of words in each sentence (e.g., ‘the [MASK] gave the soccer player a yellow card’), then training the model to predict what the masked words are (‘referee’ in this case). Both types of learning require no human annotation but allow the network to acquire language knowledge that can then be taken advantage of in subsequent fine-tuning. Training data for these prediction tasks come from English Wikipedia text and the BookCorpus (a dataset of 11,038 unpublished books). We chose to use distilBERT instead of BERT or RoBERTa (Liu et al., 2019) because of its rapid training, as our model had to be trained six times: five times for our leave-one-dataset-out cross validation, and once for training on all datasets together.

We used Huggingface Transformers (Wolf et al., 2020) to fine-tune distilBERT on our classification task. We used accuracy as our evaluation criterion when training. We used Huggingface’s default training arguments for fine-tuning. We used three training epochs, a batch size of 16 per device during training, a batch size of 64 for evaluation, 500 warmup steps, and a .01 weight decay.

Evaluating Model Performance

After classifying each of our sentences in the evaluation sets, we aggregated all sentences for each narrative and obtained an estimate of the amount of internal and external content in each narrative. We correlated these estimates with internal and external detail counts for validation. If our model were successful, we would expect the predictions of internal and external content to

match actual internal and external detail counts. Our expectations, then, were that 1) the number of internal details in the narratives would correlate with the amount of predicted internal content; 2) the number of external details in the narratives would correlate with the amount of predicted external content. Because the purpose of the AI is to correctly label content as internal or external, we also expected the model to not misclassify internal content as external. In other words, we also expected that 3) the number of internal details would be unrelated to the amount of predicted external content. Likewise, we should not misclassify external content as internal, so we further expected that 4) the number of external details would be unrelated to the amount of predicted internal content.

These four predictions are displayed graphically below in Figure 1. For each evaluation dataset, we report results in a similar format.

Figure 1. Target model performance. Internal (panel A) and external (panel B) content are accurately identified, with no misclassification of internal details as external (panel C) and external details as internal (panel D).

Results

To evaluate how well our model scored narratives, we examined internal and external detail scores as a function of predicted internal and external content. We evaluated performance on each left-out dataset separately.

Narratives Scored with the Standard or Adapted AI

Results: Future Simulation: Older Adult Data (Devitt & Schacter, 2018, 2019)

We examined whether our model correctly identified internal content. We found a strong relationship between predicted internal content and the number of internal details in future simulation narratives (fig 2a, $r = .81$, $p < .001$). We also examined the extent to which our model correctly identified external content. We found a strong relationship between predicted external content and the number of external details in future simulation narratives (fig 2b, $r = .65$, $p < .001$). We expected to find lower correlations when we examined the extent to which our model misclassified data. We examined how much internal content was misclassified as external content by our model. We did not find a significant relationship between internal details and predicted external content (fig 2c, $r = .12$, $p = .151$). We also examined how much external content was misclassified as internal content. We found a weak negative relationship between external details and predicted internal content (fig 2d, $r = -.22$, $p = .012$). To summarize, we found greater correct classification of internal content than misclassification ($R^2 = 0.65$ vs $R^2 = 0.02$). We also found greater correct classification of external content than misclassification ($R^2 = 0.42$ vs $R^2 = 0.05$). These results are summarized in figure 2 below.

Automated Scoring Performance: Older Adult Episodic Simulation (Devitt & Schacter, 2018; 2019)

Figure 2. Model performance on older adult episodic simulation data from Devitt & Schacter (2018, 2019). Internal (panel A) and external (panel B) content are accurately identified, with minimal misclassification of internal details as external (panel C) and external details as internal (panel D).

Results: Future Simulation: Young Adult Data (Devitt & Schacter, 2018, 2019)

Our model correctly identified much of the internal content (fig 3a; $r = .71$, $p < .001$) and also correctly identified much of the external content (fig 3b, $r = .40$, $p < .001$). As expected, we observed less misclassification than correct classification. Internal content was not significantly misclassified as external content (fig 3c, $r = .09$, $p = .27$), and external content was not significantly misclassified as internal content (fig 3d, $r = .08$, $p = .33$). To summarize, we found greater correct classification of internal content than misclassification ($R^2 = 0.47$ vs $R^2 = 0.01$). We also found greater correct classification of external content than misclassification ($R^2 = 0.16$ vs $R^2 = 0.01$). These results are summarized in the figure below.

Automated Scoring Performance: Young Adult Episodic Simulation (Devitt & Schacter, 2018; 2019)

Figure 3. Model performance on young adult episodic simulation data from Devitt & Schacter (2018, 2019). Internal (panel A) and external (panel B) content are accurately identified, with minimal misclassification of internal details as external (panel C) and external details as internal (panel D).

Results: Autobiographical Memory (King et al., 2021)

Once again, our model correctly identified much of the internal content (fig 4a; $r = .97$, $p < .001$) and also correctly identified much of the external content (fig 4b, $r = .54$, $p < .001$). As expected, we observed less misclassification than correct classification. Internal content was not often misclassified as external content (fig 4c, $r = .18$, $p < .001$), and external content was not significantly misclassified as internal content (fig 4d, $r = -.01$, $p = .791$). To summarize, we found greater correct classification of internal content than misclassification ($R^2 = 0.94$ vs $R^2 = 0.03$). We also found greater correct classification of external content than misclassification ($R^2 = 0.29$ vs $R^2 = 0.00$). These results are summarized in the figure below.

Automated Scoring Performance: Autobiographical Memory Task (King et al., 2021)

Figure 4. Model performance on autobiographical memory data from King et al. (2021). Internal (panel A) and external (panel B) content are accurately identified, with minimal misclassification of internal details as external (panel C) and external details as internal (panel D).

Results: Autobiographical Memories (Strikwerda-Brown et al., 2021)

Our model correctly identified much of the internal content (fig 5a; $r = .90$, $p < .001$) and external content in autobiographical memories (fig 5b, $r = .79$, $p < .001$). Internal content was not significantly misclassified as external content (fig 5c, $r = .08$, $p = .366$), and external content was also not significantly misclassified as internal content (fig 5d, $r = .08$, $p = .35$). To summarize, we found greater correct classification of internal content than misclassification ($R^2 = 0.80$ vs $R^2 = 0.01$). We also found greater correct classification of external content than misclassification ($R^2 = 0.68$ vs $R^2 = 0.01$). These results are summarized in the figure below.

Automated Scoring Performance: Autobiographical Memory Task (Strikwerda-Brown et al., 2021)

Figure 5. Model performance on autobiographical memory data from Strikwerda-Brown et al. (2021). Internal (panel A) and external (panel B) content are accurately identified, with minimal misclassification of internal details as external (panel C) and external details as internal (panel D).

Results: Creative Writing Narratives. (van Genugten et al., 2021)

Our model correctly identified much of the internal content (fig 6a; $r = .80$, $p < .001$) and external content (fig 6b, $r = .77$, $p < .001$) in creative writing narratives. As expected, we observed less misclassification than correct classification. Internal content was rarely misclassified as external content (fig 6c, $r = -.17$, $p = .003$), and external content was not significantly misclassified as internal content (fig 6d, $r = .01$, $p = .848$). To summarize, we found greater correct classification of internal content than misclassification ($R^2 = 0.64$ vs $R^2 = 0.03$). We also found greater correct classification of external content than misclassification ($R^2 = 0.59$ vs $R^2 = 0.00$). These results are summarized in the figure below.

Automated Scoring Performance: Creative Writing Task (van Genugten et al., 2021)

Figure 6. Model performance on autobiographical memory data from van Genugten et al. (2021). Internal (panel A) and external (panel B) content are accurately identified, with minimal misclassification of internal details as external (panel C) and external details as internal (panel D).

Results: Autobiographical Memories (Sheldon et al., 2020)

Consistent with the previous analyses, our model correctly identified much of the internal content (fig 7a; $r = .86$, $p < .001$) and also correctly identified much of the external content (fig 7b, $r = .67$, $p < .001$). In contrast to previous results, we observed significant misclassification. Internal content was often misclassified as external content (fig 7c, $r = .40$, $p < .001$) and external content was often misclassified as internal content (fig 7d, $r = .33$, $p < .001$). Even though we found greater correct classification of internal content than misclassification ($R^2 = 0.76$ vs $R^2 = 0.16$), misclassification rates were high. We also found greater correct classification of external content than misclassification ($R^2 = 0.46$ vs $R^2 = 0.11$), but misclassification of external content is frequent. These results are summarized in the figure below.

Automated Scoring Performance: Autobiographical Memory Task (Sheldon et al., 2020)

Figure 7. Model performance on autobiographical memory data from Sheldon et al. (2020). Internal (panel A) and external (panel B) content are accurately identified, with significant misclassification of internal details as external (panel C) and external details as internal (panel D).

Narratives Scored with Alternative Scoring Procedures

Results: Picture Description Task (Strikwerda-Brown et al., 2021).

Our model correctly identified much of the internal content (fig 8a; $r = .97$, $p < .001$) and external content (fig 8b, $r = .67$, $p < .001$) in picture descriptions. However, internal content was often misclassified as external content (fig 8c, $r = .41$, $p = .035$) and external content was often misclassified as internal content (fig 8d, $r = .34$, $p = .161$). While we found greater correct classification of internal content than misclassification ($R^2 = 0.95$ vs $R^2 = 0.16$), and we found greater correct classification of external content than misclassification ($R^2 = 0.45$ vs $R^2 = 0.11$), misclassification is frequent. These results are summarized in the figure below.

Automated Scoring Performance: Picture Description Task (Strikwerda-Brown et al., 2021)

Figure 8. Model performance on picture description data from Strikwerda-Brown et al. (2021). Internal (panel A) and external (panel B) content are accurately identified, with significant misclassification of internal details as external (panel C) and external details as internal (panel D).

Results: Thoughts Task (Strikwerda-Brown et al., 2021)

When participants describe their unconstrained thoughts, our model correctly identified much of the internal content (fig 9a; $r = .89$, $p < .001$) and external content (fig 9b, $r = .83$, $p < .001$) in the resulting transcribed narratives. Significant misclassification was present. Internal content was often misclassified as external content (fig 9c, $r = .29$, $p < .001$), and external content was often misclassified as internal content (fig 9d, $r = .50$, $p < .001$). To summarize, even though we found greater correct classification of internal content than misclassification ($R^2 = 0.80$ vs $R^2 = 0.09$) and greater correct classification of external content than misclassification ($R^2 = 0.68$ vs $R^2 = 0.25$), we observed significant misclassification. These results are summarized in the figure below.

Automated Scoring Performance: Thoughts Task (Strikwerda-Brown et al., 2021)

Figure 9. Model performance on thoughts task data from Strikwerda-Brown et al. (2021). Internal (panel A) and external (panel B) content are accurately identified, with significant misclassification of internal details as external (panel C) and external details as internal (panel D).

Discussion

We have described and tested a model-based approach that can automatically score memories, imagined events, and related narrative output for internal and external content. In general, we found that our model performs well across datasets with a variety of tasks in both older and younger populations. The amount of predicted internal content is highly correlated with actual internal detail counts in narratives. Likewise, the amount of predicted external content is highly correlated with actual external detail counts. Importantly, in most of the datasets content misclassification is relatively low: the number of internal details has little relationship to the amount of predicted external content. Likewise, there is no strong relationship between the number of external details and predicted internal content.

However, we found that model performance differed across datasets. Model performance was very good for future simulation narratives from younger adults, future simulation narratives from older adults (Devitt & Schacter, 2018, 2019), creative writing narratives (van Genugten et al., 2021), and memory narratives from Strikwerda-Brown et al (2021) and King et al. (2021). For these datasets, there was little misclassification of content. For musically cued memories (Sheldon et al., 2020), however, we found rates of misclassification that were higher than in the other datasets.

We believe that model performance on data from Sheldon et al. (2020) differed from performance on other datasets because many narratives in this dataset, in contrast to the others, were transcribed 1) without much punctuation and 2) without removing uninformative speech (e.g. ‘I don’t know, we went to, I guess, well,’). Uninformative speech is not a problem for manual scoring, but it is problematic for our model because this text will still be classified as internal or external content. Narratives with little punctuation are also problematic for our approach because it leads pySBD to split our narratives into sentences incorrectly. Resulting segments of text contained multiple sentences, which makes accurate prediction difficult for the model. While our preprocessing removed very long pieces of text without punctuation, some pieces of text still contained multiple sentences.

To explore whether our model would work on this dataset if it had been transcribed with more punctuation, we manually added punctuation to a random subset of 100 narratives. We did not remove any sentences from these newly punctuated narratives before classification. We found that adding punctuation to the transcriptions led to correct identification of internal content with no detectable misclassification of internal content as external ($R^2 = .55$ vs $R^2 = .00$). Our model also correctly identified external content with less misclassification of external content as

internal ($R^2 = .49$ vs $R^2 = .07$). We would expect even less content misclassification for this dataset if meaningless text were also removed. Together, these results suggest that performance of our model is very good across datasets when punctuation is present in narratives.

We expected performance to be worse on the two tasks scored with different guidelines, because the model was not trained to mimic those guidelines. Indeed, performance was comparatively poor on the thoughts task, in which participants provided the thoughts that came to mind as they looked at a picture, and also on the related picture description task. These results suggest that our automated scoring procedure should not be used, or used with caution, on tasks that are not optimized for scoring with the AI (Levine et al., 2002) or the adapted AI (e.g., Addis, et al., 2008).

Optimal Setting & Potential Limitations

Our model will likely perform best when used with data that are similar to our training data, i.e., when scoring internal and external data from future simulation and autobiographical memory tasks. We believe that the model can also be used for scoring other narrative data, as evidenced by strong performance on the creative writing dataset. However, we do not know how well this code will perform under new circumstances. For example, while the model seems to work well with data from both healthy young and older adults with relatively intact speech, its use with patient populations and populations with more rambling speech is untested. Researchers who want to use this automated scoring approach for new populations should manually score a subset of narratives to verify reliability. We are also unsure how well this model will perform for different dialects and for different language usage more generally. The datasets that we used for fine-tuning and evaluation were collected in the United States and Canada and thus were presumably from WEIRD (Henrich et al., 2010) populations. The data used for training

distilBERT (the model we fine-tuned) comes from English Wikipedia text and the bookCorpus (a set of unpublished English books). Accordingly, narratives that use language that is significantly different from the English text found in our fine-tuning datasets or in the pretraining data may or may not be scored as accurately. Researchers who want to use this procedure in new populations can manually score a subset of narratives to confirm accuracy.

There are several situations in which we expect the model to score text differently from the standard AI procedure. First, we expect our code to improperly score narratives that do not have punctuation to mark sentence boundaries, as discussed earlier. Second, we expect narratives that contain several events to be scored differently from the standard AI procedure. With the typical AI scoring, researchers identify a central event, and mark all details in non-central events as external. The current code is not able to identify which details belong to central versus peripheral events. As a result, the model is likely to identify event-related details as internal, regardless of which event the details came from. Depending on the research question, this feature may or may not matter. Researchers interested in the total amount of episodic and non-episodic content in narratives can use the code as is, while researchers interested in only the central event may have to manually read the narratives and score a subset of them by hand.

Recommendations for Study Design

We provide several design recommendations to maximize the power to detect an effect when using our automated scoring approach. First, we recommend using written narratives to maximize scoring accuracy. Splitting the narrative into sentences will be easiest with written text, so scoring is more likely to be accurate. In addition, memories written by a participant will have less meaningless content (e.g. ‘I don’t know, we went to, I guess, well, ’) than transcribed memories. If transcribed memories include meaningless text, this text will get scored, which adds

noise to our internal and external content measures. Further, distilBERT was pretrained on written text, so we expect performance to be best when the evaluation data is also written. While we provide this recommendation, we do not expect performance to drop a great deal for transcribed narratives with enough punctuation. Second, to achieve the same power as a manually scored study, studies using our procedure should recruit larger sample sizes. Third, prompting by the researcher (e.g., ‘Is there anything more you can tell me?’) should not be transcribed, as this text will be automatically scored as details.

Future Directions

Model Modifications

In our work, we retrained distilBERT for classification. We opted to use distilBERT because of its training speed. Other neural network architectures that perform slightly better on transfer learning tasks, such as BERT and RoBERTa could be used in future work. Future work could also systematically search for different hyperparameters to improve classification accuracy. We used default hyperparameters for fine-tuning distilBERT.

Scoring Internal and External Detail Subcategories

The approach that we have presented here is useful for automatically calculating the amount of internal and external content in narratives. A second function of the AI is to sort internal and external content into further subcategories. Internal details can be further classified as perceptual, event, time, place, and thought/emotion related. External details can be classified as event (event details from non-central events), semantic (general knowledge or facts), repetition, and other content such as metacognitive statements and editorializing (for more detailed external subcategories, see Strikwerda-Brown et al., 2019).

The same approach that we used for classifying internal and external content could be extended for classifying detail subcategories. That is, researchers could train models to determine what percentage of content in each sentence belongs to each detail subcategory. In an alternative approach, researchers could first split sentences into individual clauses and train a model to classify the resulting clauses into subcategories. Classifying at the clause levels may lead to fewer misclassifications. If researchers are unable to find adequate training data, since only a subset of AI studies score for subcategories, an alternative approach may be to use zero-shot classifiers. These classifiers can attempt to predict the topic, from a list of given topics (e.g. action, time, perceptual, etc.), that a piece of text belongs to (e.g. Yin et al., 2019; for demo, see <https://huggingface.co/zero-shot/>). These models do not require training data to fine-tune predictions.

Conclusion and Application

We believe that the tool presented here will enable researchers to conduct studies with considerably larger sample size than typically used, and thus perhaps capture smaller effect sizes, using the AI. We believe this tool will also facilitate internet-based research with the AI. This type of research has often been impractical because of the scoring burden that comes from collecting more participants to offset data quality. Internet-based research would allow researchers to study more diverse populations and would allow memory researchers to take advantage of strategies used in other areas of psychology, such as rapidly piloting multiple experiments online. Importantly, the automated scoring procedure will enable research groups that have fewer resources for scoring narratives to also conduct large AI studies.

To accompany this paper, we provide a Colab notebook that researchers can open in their web browser. Researchers can use this notebook to automatically score memories by providing a

spreadsheet with narratives. The notebook is intended to be useable out-of-the-box without any additional coding required. This notebook and instructions for using it can be found at <https://github.com/rubenvangenugten/>. The final model used in this notebook has been retrained on all adapted and standard AI-scored datasets.

Appendix

Example scoring

During the AI, a participant might provide a response such as the following:

We went to the beach because it was my birthday. I was sitting on the beach with a beer.

We had driven up last night. We ran into the surf.

After applying the rules described in the scoring manual, we obtained the scored response below.

We used forward slashes to separate details and we have surrounded the external details with brackets.

[We went to the beach because it was my birthday/]. I was sitting on the beach/ with a beer /. [We had driven up last night/]. We ran into the surf/.

We then counted the details to summarize the narrative: 2 external details and 3 internal details.

If we were using our automated approach, we would summarize the narrative by counting the number of words in external segments (16 words in the brackets) and internal segments (14 words outside of the brackets).

References

- Addis, D. R., Wong, A. T., & Schacter, D. L. (2008). Age-related changes in the episodic simulation of future events. *Psychological Science*, *19*, 33–41.
<https://doi.org/10.1111/j.1467-9280.2008.02043.x>
- Azunre, P. (2021). *Transfer learning for natural language processing*. Simon and Schuster.
- Devitt, A. L., & Schacter, D. L. (2018). An optimistic outlook creates a rosy past: the impact of episodic simulation on subsequent memory. *Psychological Science*, *29*, 936–946.
<https://doi.org/10.1177/0956797617753936>
- Devitt, A. L., & Schacter, D. L. (2020). Looking on the bright side: aging and the impact of emotional future simulation on subsequent memory. *The Journals of Gerontology: Series B*, *75*, 1831–1840. <https://doi.org/10.1093/geronb/gbz041>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *ArXiv*.
<http://arxiv.org/abs/1810.04805>
- Gaesser, B., Sacchetti, D. C., Addis, D. R., & Schacter, D. L. (2011). Characterizing age-related changes in remembering the past and imagining the future. *Psychology and Aging*, *26*, 80–84. <https://doi.org/10.1037/a0021054>
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*, 61-83.
- Irish, M., Hornberger, M., Lah, S., Miller, L., Pengas, G., Nestor, P. J., Hodges, J. R., & Piguet, O. (2011). Profiles of recent autobiographical memory retrieval in semantic dementia, behavioural-variant frontotemporal dementia, and Alzheimer's disease. *Neuropsychologia*, *49*, 2694–2702.

<https://doi.org/10.1016/j.neuropsychologia.2011.05.017>

Jing, H. G., Madore, K. P., & Schacter, D. L. (2016). Worrying about the future: An episodic specificity induction impacts problem solving, reappraisal, and well-being. *Journal of Experimental Psychology: General*, 145, 402.

King, C.I., Romero, A.S.L., Schacter, D.L., & St. Jacques, P.L. (2021). The influence of shifting perspective on episodic and semantic details during autobiographical memory recall.

Submitted for publication.

Levine, B. (2021) Memory. The Levine Lab. Retrieved November 9, 2021 from

<https://levinelab.weebly.com/memory.html>

Levine, B., Svoboda, E., Hay, J. F., Winocur, G., & Moscovitch, M. (2002). Aging and autobiographical memory: Dissociating episodic from semantic retrieval. *Psychology and Aging*, 17, 677–689. <https://doi.org/10.1037/0882-7974.17.4.677>

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A robustly optimized bert pretraining approach. *ArXiv*. <http://arxiv.org/abs/1907.11692>

Madore, K. P., Gaesser, B., & Schacter, D. L. (2014). Constructive episodic simulation: Dissociable effects of a specificity induction on remembering, imagining, and describing in young and older adults. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40, 609.

Madore, K. P., & Schacter, D. L. (2014). An episodic specificity induction enhances means-end problem solving in young and older adults. *Psychology and Aging*, 29, 913.

Peters, J., Wiehler, A., & Bromberg, U. (2017). Quantitative text feature analysis of autobiographical interview data: Prediction of episodic details, semantic details and

- temporal discounting. *Scientific Reports*, 7, 14989. <https://doi.org/10.1038/s41598-017-14433-6>
- Race, E., Keane, M. M., & Verfaellie, M. (2011). Medial temporal lobe damage causes deficits in episodic memory and episodic future thinking not attributable to deficits in narrative construction. *Journal of Neuroscience*, 31, 10262–10269. <https://doi.org/10.1523/JNEUROSCI.1145-11.2011>
- Sadvilkar, N., & Neumann, M. (2020). PySBD: Pragmatic sentence boundary disambiguation. *ArXiv*. <http://arxiv.org/abs/2010.09657>
- Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. *ArXiv*. ArXiv:1910.01108.
- Schacter, D. L., & Addis, D. R. (2007). The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362, 773-786.
- Schacter, D. L., & Madore, K. P. (2016). Remembering the past and imagining the future: Identifying and enhancing the contribution of episodic memory. *Memory Studies*, 9, 245–255.
- Sheldon, S., Williams, K., Harrington, S., & Otto, A. R. (2020). Emotional cue effects on accessing and elaborating upon autobiographical memories. *Cognition*, 198, 104217. <https://doi.org/10.1016/j.cognition.2020.104217>
- Söderlund, H., Moscovitch, M., Kumar, N., Daskalakis, Z., Flint, A., Herrmann, N., & Levine, B. (2014). Autobiographical episodic memory in major depressive disorder. *Journal of Abnormal Psychology*, 123, 51–60. <https://doi.org/10.1037/a0035610>
- Strikwerda-Brown, C., Mothakunnel, A., Hodges, J. R., Piguet, O., & Irish, M. (2019). External

details revisited – A new taxonomy for coding ‘non-episodic’ content during autobiographical memory retrieval. *Journal of Neuropsychology*, 13, 371–397.

<https://doi.org/10.1111/jnp.12160>

Strikwerda-Brown, C., Williams, K., Lévesque, M., Brambati, S., & Sheldon, S. (2021). What are your thoughts? Exploring age-related changes in episodic and semantic autobiographical content on an open-ended retrieval task. *Memory*, 29(10), 1375–1383.

<https://doi.org/10.1080/09658211.2021.1987476>

Takano, K., Gutenbrunner, C., Martens, K., Salmon, K., & Raes, F. (2018). Computerized scoring algorithms for the Autobiographical Memory Test. *Psychological Assessment*, 30(2), 259–273. <https://doi.org/10.1037/pas0000472>

Takano, K., Hallford, D. J., Vanderveren, E., Austin, D. W., & Raes, F. (2019). The computerized scoring algorithm for the autobiographical memory test: Updates and extensions for analyzing memories of English-speaking adults. *Memory*, 27, 306–313.

<https://doi.org/10.1080/09658211.2018.1507042>

Takano, K., Ueno, M., Moriya, J., Mori, M., Nishiguchi, Y., & Raes, F. (2017). Unraveling the linguistic nature of specific autobiographical memories using a computerized classification algorithm. *Behavior Research Methods*, 49, 835–852.

<https://doi.org/10.3758/s13428-016-0753-x>

van Genugten, R. D., Beaty, R. E., Madore, K. P., & Schacter, D. L. (2021). Does episodic retrieval contribute to creative writing? an exploratory study. *Creativity Research Journal*, 1–14.

Wardell, V., Esposito, C. L., Madan, C. R., & Palombo, D. J. (2021). Semi-automated transcription and scoring of autobiographical memory narratives. *Behavior Research*

- Methods*, 53, 507–517. <https://doi.org/10.3758/s13428-020-01437-w>
- Wardell, V., Madan, C. R., Jameson, T. J., Cocquyt, C. M., Checknita, K., Liu, H., & Palombo, D. J. (2021). How emotion influences the details recalled in autobiographical memory. *Applied Cognitive Psychology*, 35, 1454–1465. <https://doi.org/10.1002/acp.3877>
- Wickner, C., Englert, C., Addis, D.R. (2015). Developing a tool for autobiographical interview scoring. Kiwicam Conference, Wellington, New Zealand. <https://github.com/scientific-tool-set/scitos>
- Williams, J. M., & Broadbent, K. (1986). Autobiographical memory in suicide attempters. *Journal of Abnormal Psychology*, 95, 144–149. <https://doi.org/10.1037/0021-843X.95.2.144>
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., ... Rush, A. (2020). Transformers: State-of-the-Art natural language processing. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>
- Yin, W., Hay, J., & Roth, D. (2019). Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. *ArXiv*. <http://arxiv.org/abs/1909.00161>