



**APPLYING ATTENTION BASED MODELS FOR
DETECTING COGNITIVE PROCESSES AND
MENTAL HEALTH CONDITIONS**

Esaú VILLATORO-TELLO^a Shantipriya Parida^b
Sajit Kumar Petr Motlicek

Idiap-RR-01-2022

JANUARY 2022

^aIdiap

^bIdiap Research Institute

Applying Attention Based Models for Detecting Cognitive Processes and Mental Health Conditions

Esau Villatoro-Tello · Shantipriya Parida · Sajit Kumar · Petr Motlicek

Received: date / Accepted: date

Abstract Objective: According to the psychological literature, implicit motives allow for the characterization of behavior, subsequent success, and long-term development. Contrary to personality traits, implicit motives are often deemed to be rather stable personality characteristics. Normally, implicit motives are obtained by Operant Motives, unconscious intrinsic desires measured by the Operant Motive Test (OMT). The OMT test requires participants to write freely descriptions associated with a set of provided images and questions. In this work, we explore different recent machine learning techniques and various text representation techniques for facing the problem of the OMT classification task.

Methods: We focused on advanced language representations (e.g, BERT, XLM, and DistilBERT) and deep Supervised Autoencoders for solving the

Corresponding author: Esau Villatoro-Tello, Ph.D., E-mail: evillatoro@cua.uam.mx

E. Villatoro-Tello
Universidad Autónoma Metropolitana,
Unidad Cuajimalpa,
Mexico City, Mexico.
E-mail: evillatoro@cua.uam.mx
Idiap Research Institute
Rue Marconi 19, 1920 Martigny
Switzerland.
E-mail: esau.villatoro@idiap.ch

S. Parida, P. Motlicek
Idiap Research Institute
Rue Marconi 19, 1920 Martigny
Switzerland.
E-mail: {shantipriya.parida, petr.motlicek}@idiap.ch

S. Kumar
GreatLearning
Bangalore, Karnataka
India.
E-mail: kumar.sajit.sk@gmail.com

OMT task. We performed an exhaustive analysis and compared their performance against fully connected neural networks and traditional support vector classifiers.

Results: Our comparative study highlights the importance of BERT which outperforms the traditional machine learning techniques by a relative improvement of 7.9%. In addition, we performed an analysis of how the BERT attention mechanism is being modified. Our findings indicate that the writing style features acquire higher importance at the moment of accurately identifying the different OMT categories.

Conclusions: This is the first time that a study to determine the performance of different transformers-based architectures in the OMT task is performed. Similarly, our work propose, for the first time, using Deep Supervised Autoencoders in the OMT classification task. Our experiments demonstrate that transformers-based methods exhibit the best empirical results, obtaining a relative improvement of 7.9% over the competitive baseline suggested as part of the GermEval 2020 challenge. Additionally, we show that features associated with the writing style are more important than content-based words. Some of these findings show strong connections to previously reported behavioral research on the implicit psychometrics theory.

Keywords Operant Motive Test · Psycholinguistics · BERT · Supervised Autoencoder · Deep Learning · Natural Language Processing

1 Introduction

The idea that language use reveals information about personality has long circulated in the social and medical sciences. Previous research has demonstrated that the way people use words convey a great deal of information about themselves and their mental health conditions [5, 31, 34, 35], including academic success [32]; however, much of the previous research has focused on the analysis of self-reports or essays. In contrast, implicit motives, which are indicators used by professional psychologists during the aptitude diagnosis, are not readily accessible to the conscious mind and, therefore, not detected using self-reports of personal needs, or through essay writing [12]. Instead, they are primarily assessed using indirect measures that rely on projective techniques that instruct individuals to produce imaginative stories based on ambiguous pictures stimuli that depict people in different situations (Examples of these stimuli pictures are shown in Fig. 1). These pictures aim to influence the content of the subjects' fantasy and to provoke that such fantasy is projected onto the characters through a short (textual) story (see Table 1 for some examples of the type of produced stories for the top-left image from Figure 1). Consequently, this motivational response emerges through the contents of the written imaginative material and can be coded for its motive imagery using standardized and validated content coding systems.

The most frequently used measures of implicit motives are the Picture Story Exercise (PSE) [29], the Thematic Apperception Test (TAT) [30], the

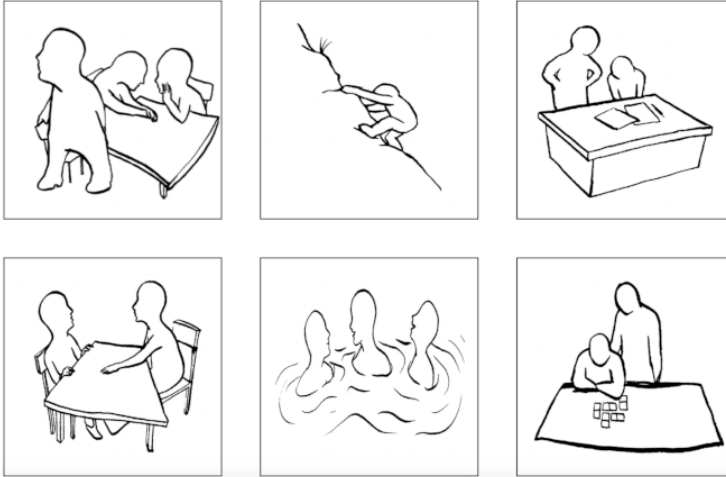


Fig. 1: Sample images that are shown to subjects during the OMT test. Credits of the image correspond to organizers of the GermEval 2020 shared task.²

Motive	Sample text (<i>Closest English translation</i>)
A	sie nimmt am Gespräch nicht teil und wendet sich ab. gelangweilt. es interessiert sie nicht, worüber die anderen beiden reden. schlecht. (<i>she does not take part in the conversation and turns away. bored. they don't care what the other two talk about. bad.</i>)
M	weicht ängstlich zurück. unterlegen. wird zurechtgewiesen. Gelegenheit den Fehler zu korrigieren (<i>withdraws fearfully. inferior. is rebuked. Opportunity to correct the mistake</i>)

Table 1: Example answers produced to the stimuli image on the top-left side of Figure 1.

Multi-Motive Grid (MMG) [43], and the Operant Motive Test (OMT) [9, 22]. Generally speaking, these tests are based on the operant methods, i.e., participants are asked ambiguous questions or are shown simple images, which they have to describe. Specifically, the OMT test is a projective instrument in which participants are presented ambiguous pictures (e.g., sketched scenarios) and ask to think of a story that describes what is happening in the picture. Thus, participants are asked to first pick the main protagonist, think of a story involving this person, and then answer the following three questions as spontaneously as possible: “*What is important to this person in this situation and what is the person doing?*”, “*How does the person feel?*”, and “*Why does the person feel that way?*” [3, 22]. Then, trained psychologists label these textual answers with one of five motives, namely M-power, A-affiliation, L-achievement, F-freedom, and 0-zero; and each motive is associated with its corresponding level or emotion (from 0 to 5), resulting in a total of 30 (5-

Type	Definition
Power (M)	Subjects concerned about having impact, control, or influence on another person, group, or the world at large.
Freedom (F)	Subjects concerned about stories that include the themes self-joy, increases in self-esteem through praise and attention, self-growth and self-actualization, integration of negative experiences into the self, rigid self-protection, or expressions of the fear of self-devaluation.
Affiliation (A)	Subjects concerned about establishing, maintaining, or restoring friendship or friendly relations among persons, groups, etc.
Achievement (L)	Subjects concerned about a standard of excellence.

Table 2: Brief outline of the motives (imagery types)[3, 22].

motives \times *6-levels*) different OMT categories. In Table 2 we briefly describe the meaning of the operant motives; interested reader is referred to [3, 22].

Even though nowadays there is a huge demand for psychological data and its automated analysis, see for example works presented at forums such as the CLPsych workshop¹, until recently, little research has been performed on the Operant Motive classification task [17, 18, 19, 40, 6, 52]. The primary reason is the lack of available labeled psychological text data, as [14] point out, and the difficulty of capturing psychological traits from text data, especially from very short texts.

Accordingly, in this paper, we aim to mitigate the lack of research on this issue and we explore how the recent advances in Natural Language Processing (NLP) can be applied in the task of automatically identify psychological traits from short textual data, specifically, we performed an exhaustive evaluation on the classification of the Operant Motives from the text. We evaluate the impact of very recent deep learning architectures such as Transformers [55] (BERT, XLM, DistilBert), recent generalization techniques as supervised autoencoders [24], traditional classification methods, e.g., fully connected neural networks and support vector machines. To perform our experiments we use the dataset provided during the "GermEval 2020 Task on the Classification and Regression of Cognitive and Emotional Style from Text" [20].²

The present paper represents an important extension of our previous reported participation [52] at the GermEval 2020 [20]. The main addition relies upon the adaptation and evaluation of recent generalization techniques, namely, deep supervised autoencoders (SAE). To the best of our knowledge, this is the very first time such an exhaustive evaluation of SAE techniques is performed on the OMT classification task. Additionally, we perform a thorough analysis of how the attention-mechanism from the transformers architectures is affected when solving the OMT task.

In summary, the main contributions of this paper are as follows:

¹ Computational Linguistics and Clinical Psychology Workshop (<https://clpsych.org/>)

² <https://www.inf.uni-hamburg.de/en/inst/ab/lt/resources/data/germeval-2020-cognitive-motive.html>

1. To the best of our knowledge, this paper represents the very first systematic exploration, and comparison, of several recent NLP technologies as well as machine learning techniques on the OMT classification task;
2. We propose a supervised autoencoder architecture for solving the OMT classification task, which as per our literature review, none of the recent research has applied this type of technology on the posed task. At the same time, we evaluate the impact of different feature types, ranging from char ngrams to recent contextual embeddings, as inputs to the proposed SAE;
3. Finally, we conducted an analysis of how the attention-mechanism of the transformer-based architectures is adapted during the OMT classification task. This type of analysis provides, to some extent, transparency on how the classifier is making its decisions. During this process, we observed some strong connections between our obtained results and psychometrics research.

The rest of the paper is organized as follows. Related work is discussed in Section 2. Section 3 describes the dataset used in this work and its main characteristics. Details of our applied methodology are given in Section 4. The experimental setup, results, and analysis are provided in Section 5. Finally, we share our main conclusions and future work directions in Section 6.

2 Related Work

Nowadays, there is an acknowledged necessity for digital solutions for addressing the burden of mental health diagnosis and treatment. It is recognized that won't be possible to treat people by professionals alone, and even if possible, some people might require to use alternative modalities to receive mental health support [56]; such situation has become more evident with the current COVID-19 pandemic. Examples of recent efforts building technology towards this direction are the Woebot [11] and Wysa [16] dialog systems for health and therapy support for patients that have depression symptoms; Expressive Interviewing [53], which is a conversational agent aiming at support users to cope with COVID-19 issues.

The underlying hypothesis of most these works relies on the notion of the language as a powerful indicator about our personality, social, or emotional status, and mental health [34, 48]. Accordingly, the NLP community has focused on proposing several methods to identify different psychological traits from texts, and to examine the connection between language and mental health. As a few examples of this type of research, we can mention dementia identification [27, 47], depression detection [1, 26, 48], crisis counselling [8], suicide risks identification [25, 26, 28], mental illnesses classification [58, 51], anxiety detection [41], personality traits identification [36, 44], etc.

Although plenty of research has been done in the field of mental disorders detection and personality traits detection, there has been very little research for identifying motivation, success, or similar characteristics from psychological projective tests. One representative work, that brought back the discussion

of how these traits could be automatically detected through traditional NLP techniques, is the research reported in [19]. Authors performed a process of features engineering to train a logistic model tree (LMT) [23] to classify a reduced set of implicit motives (O, M, A, and L). An LMT is a decision tree, which performs logistic regression at its leaves. In their research, authors found that the perplexity of the language models for each motive, closed-class words, and ratios (words per sentence ratio, type-token ratio) were the most discriminating features during the classification process. In [18], authors proposed using a Long Short-Term Memory (LSTM) neural network combined with an attention mechanism for classifying OMT motives (O, M, A, and L) from text data. For their experiments, authors employed pre-trained german fastText word embeddings [4] to represent tokens in the OMT data. Authors mention that when reviewing tokens that have high associated attention weights and compared with the Linguistic Inquiry and Word Count (LIWC) tool [33], they found a weak connection between LIWC categories and the OMT theory.

More recently, during the 5th SwissText & 16th KONVENS Joint Conference 2020³, it was organized a shared task on the classification and regression of cognitive and motivational style from the text, GermEval 2020². This represented the very first time a shared task for detecting and classifying OMT motives was organized, providing a huge dataset which was carefully curated for this purpose [20]. An important characteristic of the provided dataset is the fact of including the F-motive, and the labeling of 6 emotions (or levels), representing the first time a dataset with these characteristics is ever released. Three different research teams participated during the shared task, showing as the main results, the pertinence of using recent Bidirectional Encoder Representations from Transformers (ie., BERT) for solving the posed task. The winning approach employed the pre-trained Digitale Bibliothek Munchener Digitalisierungszentrum (DBMDZ) German model, achieving an F-score of 70% in the prediction of motives, and levels [40]. Generally speaking, and based on the shared task submissions [20], it is possible to observe that the performance of systems using pre-trained BERT embeddings and attention-based models tend to perform better than linear models. We referred the interested reader to the system description papers presented during the shared task [6, 40, 52].

Hence, and with the aim of providing a more detailed analysis of the performance of different Natural Language Processing techniques and Machine Learning approaches in the recently released GermEval 2020 dataset, this paper will positively impact future research done on the comprehension of the implicit motive theory and its automatic detection through recent techniques. Accordingly, in this paper, we present a substantial extension of our system description paper [52], presented during the GermEval 2020, with the following main differences: *i*) we introduce and perform a series of experiments using novel generalization techniques, namely deep supervised autoencoders; which have never been tested on the OMT task before; *ii*) we evaluate the performance of three different transformer-based architectures; and, *iii*) we

³ <https://swisstext-and-konvens-2020.org/shared-tasks/>

	0	1	2	3	4	5	Total:
O	7610	0	1	1	2	4	7,618
A	13	2838	9574	1362	7556	6798	28,141
M	21	9482	6870	10797	27175	14251	68,596
L	6	2396	12863	6289	7544	3747	32,845
F	8	1766	5539	4302	9059	9313	29,987
Total:	7,658	16,482	34,847	22,751	51,336	34,113	167,187

Table 3: Distribution of the *train* partition across the OMT’s motives and levels values.

	0	1	2	3	4	5	Total:
O	1004	0	0	0	0	0	1,004
A	0	386	1198	169	939	855	3,547
M	4	1203	886	1321	3335	1787	8,536
L	0	295	1531	833	951	493	4,103
F	0	200	647	541	1121	1201	3,710
Total:	1,008	2,084	4,262	2,864	6,346	4,336	20,900

Table 4: Distribution of the *dev* partition across the OMT’s motives and levels values.

	0	1	2	3	4	5	Total:
O	963	0	0	0	0	1	964
A	0	367	1182	164	936	888	3,537
M	0	1224	821	1383	3396	1766	8,590
L	0	315	1525	815	906	437	3,998
F	1	192	715	555	1109	1239	3,811
Total:	964	2,098	4,243	2,917	6,347	4,331	20,900

Table 5: Distribution of the *test* partition across the OMT’s motives and levels values.

perform an analysis on the attention mechanism of transformer-based technologies which helps to understand why this type of technologies are well suited for this particular problem, showing strong connections with previous findings from the psychometrics research field.

3 Dataset

To perform our experiments, we employed the dataset available in the GermEval 2020 shared task on the “Classification and Regression of Cognitive and Motivational style from the text”.¹ The provided data, mostly written in standard German language, has been collected from around 14000 subjects that participated in the OMT test.⁴ Each answer was manually labeled with the motives (O, A, L, M, F) and the levels (from 0 to 5). This annotation was performed by expert psychologists, trained by the OMT manual as described

⁴ There are a few cases where answers were given in English or French.

in [22]. The size of the data set is 209,000 texts, from which 167,200⁵ are part of the training (*train*) partition, 20,900 are part of the development (*dev*), and partition and 20,900 for testing (*test*). Table 3, 4, and 5 shows the distribution of the instances across the different classes for the *train*, *dev*, and *test* partitions respectively.

As can be observed in tables (Table 3, 4, and 5), the dataset is highly unbalanced, making the classification task more challenging. The majority of the instances ($\approx 41\%$) are from the power motive (M), followed by the achievement (L) motive ($\approx 20\%$). Regarding the levels, most of the instances are grouped among classes 4 ($\approx 30.7\%$), 5 ($\approx 20.4\%$), and 2 ($\approx 20.84\%$). It is important to mention that the same distribution remains in the *dev* and *test* partitions (Table 4 and 5). Table 6 shows some statistics of the GermEval 2020 dataset, for *train*, *dev*, and *test* partitions. We compute the average number of tokens, vocabulary, and lexical richness of each text in the dataset. Lexical richness (LR), also known as “type/token ratio” is a value that indicates how the terms from the vocabulary are used within a text. LR is defined as the ratio between the vocabulary size and the number of tokens from a text ($LR = |V|/|T|$). Thus, a value close to 1 indicates a higher LR, which means vocabulary terms are used only once, while values near to 0 represent a higher number of tokens used more frequently (i.e., more repetitive).

Two main observations can be done at this point. On the one hand, notice that for the three partitions (i.e., *train*, *dev*, and *test*), textual descriptions are very short, on average 20 tokens with a vocabulary of 18 words, resulting in a very high LR (0.92). The high LR value means that very few words are repeated within each textual description, i.e., very few redundancies. On the other hand, globally speaking, the complete dataset has a low LR (0.08 for *train* and 0.13 for *dev* and *test*). Although these values are not directly comparable due to the size of each partition, they indicate, to some extent, that information across texts is very repetitive, i.e., similar types of words are being used by tested subjects for describing different images, even though they belong to different classes (motives and levels). Overall, this initial analysis helped us to envision the complexity and nature of the data. Finally, we measure the coverage ratio of the German word embeddings (EmbC)⁶ into our dataset, resulting in a 68.26% of coverage for the *training* partition, 83.93% for the *dev* set, and 84.10% for the *test* set. Similar to the LR score, we can not directly compare EmbC results due to the size of each partition. However, it was expected not to have a 100% coverage due to the noise contained in the OMT data, e.g., the many spelling and grammar errors present. Nevertheless, it is relevant to highlight the low coverage in the training set ($\sim 68\%$). As it will be explained in Section 5.1, this low coverage has an important impact on the experiments based on these word embeddings.

⁵ A total of 13 instances were removed from the training partition due to its lack of label, leaving 167,187 instances.

⁶ For some of our experiments we used word embeddings trained with FastText on 2 million German Wikipedia articles, available at: <https://www.spinningbytes.com/resources/wordembeddings/>

Partition	Metric	Average (σ)	Total
<i>Training</i>	Tokens	20.27 (± 12.08)	3,389,945
	Vocabulary	18.07 (± 9.82)	63,133
	LR	0.92 (± 0.08)	0.08
	EmbC		43,095 (68.26%)
<i>Development</i>	Tokens	20.38 (± 12.17)	425,880
	Vocabulary	18.17 (± 9.94)	19,571
	LR	0.92 (± 0.08)	0.13
	EmbC		16,426 (83.93%)
<i>Test</i>	Tokens	20.24 (± 12.01)	423,018
	Vocabulary	18.05 (± 9.76)	19,780
	LR	0.92 (± 0.08)	0.13
	EmbC		16,636 (84.10%)

Table 6: Statistics of the OMT dataset in terms of number of tokens, vocabulary size and lexical richness. The minimum length of the texts is 1 token, while the maximum length is 99, 90, and 96 tokens for *train*, *dev*, and *test* partitions respectively. In all partitions, the 75% of the data has a length of 27 tokens.

4 Methodology

Figure 2 shows the general view of the applied methodology for performing our experiments in this paper. As shown, given a textual description, we extract different types of features: character n-grams, word n-grams, non-contextual word embeddings (FastText), and pre-trained contextual embeddings (BERT). Then, depending on the selected learning strategy, computed features are feed to a specific learning technique, e.g., a supervised autoencoder (SAE), a fully connected network (FC), or into a transformer-based architecture (ST); that is trained to detect motives and levels from text, i.e., the OMT task. It is important to mention that instead of facing the OMT task as a 30 class classification problem, we split the problem into two separate classification tasks: motives (5 classes), and levels detection (6 classes). This decision was made in accordance with the operant motive (OMT) theory [22], which states that motives and levels are disjoint orthogonal and thus not directly connected. Thus, at the end of our methodology, we fuse the predicted labels in order to get the motive-level combination of the given instance. Notice the dashed lines that go from BERT to the FC and SAE; these lines represent a series of experiments done after fine-tuning BERT on the posed task, where the newly computed embeddings are used as features to the FC and SAE learning techniques (See section 4.2 to know the details of the fine-tuning process). Finally, the solid line that goes directly from the input textual description to the transformers block, it represents a series of experiments where the ST’s are configured as as classifiers by adding a simple dense layer at the end.

Following sections describe in detail the applied methodology for performing all our experiments. The proposed supervised autoencoder is detailed in

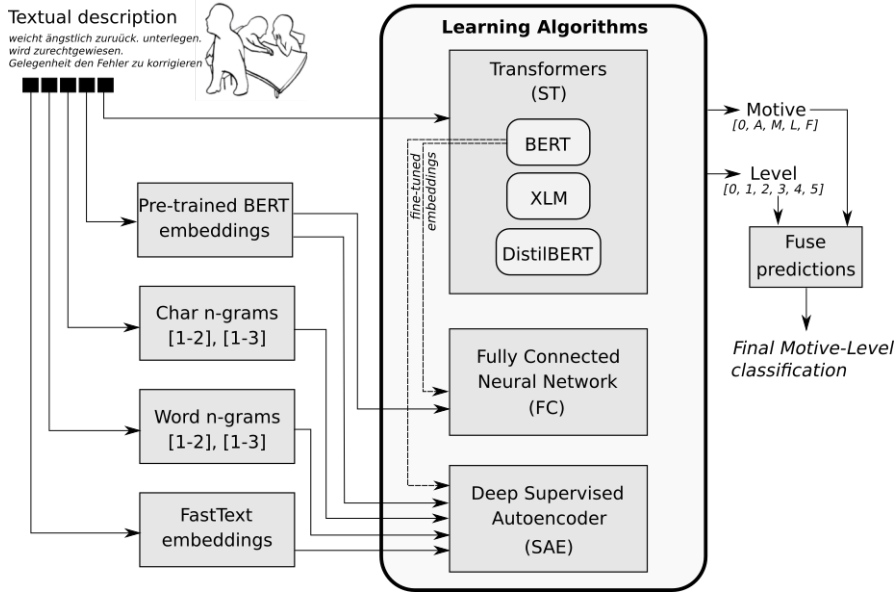


Fig. 2: Proposed methodology. Given a textual description, we evaluated several text representations as input for different recent learning algorithms, including simple transformers, fully connected neural networks, and our proposed deep supervised autoencoder. In the end, predicted motive and level are combined to produce the final motive-level classification output.

Section 4.1. The description of the fine-tuning of transformers based architectures is depicted in Section 4.2, and the configuration of the traditional fully-connected neural network is described in Section 4.3. Section 4.4 details the pre-processing operations to the dataset. Section 4.5 describe the considered evaluation metrics. An finally, Section 4.6 defines the considered baseline and employed validation strategies.

4.1 Supervised Autoencoders

An autoencoder (AE) is a neural network that learns a representation (encoding) of input data and then learns to reconstruct the original input from the learned representation. The autoencoder is mainly used for dimensionality reduction or feature extraction [57]. Normally, AE are used in an unsupervised learning fashion, meaning that we leverage the neural network for the task of representation learning. By learning to reconstruct the input, the AE extracts the underlying abstract attributes that facilitate accurate prediction of the input.

Thus, a supervised autoencoder (SAE) is an autoencoder with the addition of a supervised loss on the representation layer. For the case of a single hidden

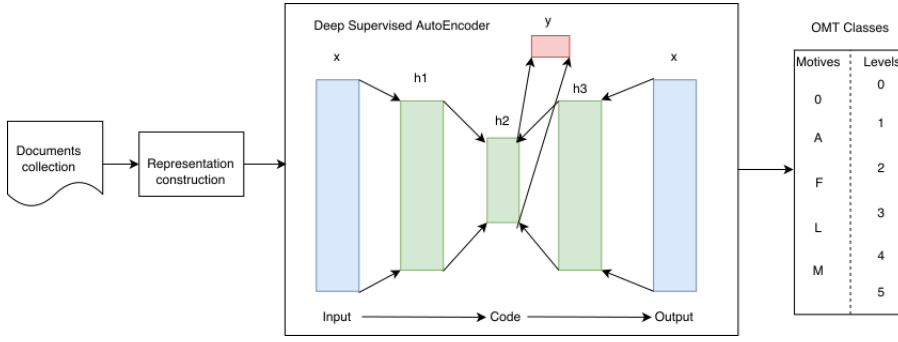


Fig. 3: An example of a Supervised Autoencoder, where the supervised component (y -labels) is included. The input of the SAE is any type of pre-defined features computed over the document collection, e.g., character n-grams, word embeddings, or sentence encodings.

layer, a supervised loss is added to the output layer and for a deeper autoencoder, the innermost (smallest) layer would have a supervised loss added to the bottleneck layer that is usually transferred to the supervised layer after training the autoencoder.

In supervised learning, the goal is to learn a function for a vector of inputs $\mathbf{x} \in \mathbb{R}^d$ to predict a vector of targets $\mathbf{y} \in \mathbb{R}^m$. Consider a SAE with a single hidden layer of size k , and with weights for the first layer defined as $\mathbf{F} \in \mathbb{R}^{k \times d}$. The function is trained on a finite batch of independent and identically distributed (i.i.d.) data, $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_t, \mathbf{y}_t)$, with the goal of a more accurate prediction on new samples generated from the same distribution. The weight for the output layer consists of weights $\mathbf{W}_p \in \mathbb{R}^{m \times k}$ to predict \mathbf{y} and $\mathbf{W}_r \in \mathbb{R}^{d \times k}$ to reconstruct \mathbf{x} . Let L_p be the supervised loss and L_r be the loss for the reconstruction error. In the case of regression, both losses might be represented by a squared error, resulting in the objective:

$$\frac{1}{t} \sum_{i=1}^t \left[L_p(\mathbf{W}_p \mathbf{F} \mathbf{x}_i, \mathbf{y}_i) + L_r(\mathbf{W}_r \mathbf{F} \mathbf{x}_i, \mathbf{x}_i) \right] = \frac{1}{2t} \sum_{i=1}^t \left[\|\mathbf{W}_p \mathbf{F} \mathbf{x}_i - \mathbf{y}_i\|_2^2 + \|\mathbf{W}_r \mathbf{F} \mathbf{x}_i - \mathbf{x}_i\|_2^2 \right] \quad (1)$$

The addition of supervised loss to the autoencoder loss function acts as regularizer and results (as shown in equation 1) in the learning of the better representation for the desired task [24]. Although SAE have been extensively evaluated on image classification tasks [24], its pertinence on thematic and non-thematic text classification tasks has not been extensively evaluated, being this an important contribution of this work.

Consequently, in order to perform a broad evaluation of this approach, we passed as input features to the SAE different types of text representation,

namely pre-trained BERT encodings and also fine-tuned BERT encodings, in both cases using as representation the information extracted from the last hidden layer (LastHL), and the concatenation of the 4 last hidden layers (Concat4LHL).⁷ Additionally, we also tested several traditional text representation techniques: word and char n-grams (with ranges 1-2 and 1-3). Finally, we also evaluate the performance of the SAE architecture using as a representation type non-contextual embeddings, in particular, we tried the German FastText embeddings trained on 2 million German Wikipedia articles.⁸ All these variations can be observed at the bottom part of Figure 2. For all our performed experiments, the overall configuration of the SAE model was done using non-linear activation function (ReLU) with 3 hidden layers, the number of nodes in the representation layer was set to 300, and we trained to a maximum of 100 epochs.⁹

4.2 Simple Transformer

The transformer model [50] introduces an architecture that is solely based on attention mechanism and does not use any recurrent networks but yet produces results superior in quality to Seq2Seq [46] models, incorporating the advantage of addressing the long term dependency problem found in Seq2Seq model.

For our experiments using Simple Transformers (ST) architectures, we setup three different state-of-the-art configurations:

1. Bert [10]: we use a pre-trained model referred as `bert-base-german-cased`, with 12-layer, 768-hidden, 12-heads, 110M parameters.¹⁰ The model is pre-trained on German Wikipedia dump (6GB of raw text files), the OpenLegalData dump (2.4 GB), and news articles (3.6 GB). We refer to this configuration as ST-Bert in our experiments.
2. XLM [7]: for this configuration we use a model with 6-layer, 1024-hidden, 8-heads, which is an English-German model trained on the concatenation of English and German Wikipedia documents (`bert-base-german-cased`). We refer to this configuration as ST-XLM in our experiments.
3. DistilBert [38]: for this model we used a model with 6-layer, 768-hidden, 12-heads, 66M parameters (`distilbert-base-german-cased`). We refer to this configuration as ST-DistilBert in our experiments.

For all the previous configurations, in order to perform the fine-tuning of the ST architecture as a classifier, a simple dense layer with softmax activation is added on top of the final hidden state h of the first token [CLS], through a weight matrix W , and we predict the probability of label c the following way:

$$p(c|\mathbf{h}) = \text{softmax}(W\mathbf{h}) \quad (2)$$

⁷ More details on how these encodings are extracted in Section 4.3

⁸ <https://www.spinningbytes.com/resources/wordembeddings/>

⁹ The implementation of our proposed SAE method can be found at https://github.com/idiap/sae_lang_detect

¹⁰ <https://deepset.ai>

Hyper Parameter	Range
number of layers	3
number of hidden layers	1
nodes in hidden layer	16
activation function	ReLU

Table 7: Fully connected neural network configuration parameters. Notice that the number of input neurons depends on the representation type, while the number of output neurons is determined by the classification task, e.g., for the motive task, there are 5 output neurons.

Then, all weights, including the model’s ones and W , are adapted, in order to maximize the log-probability of the correct label. The training is done using a Cross-Entropy loss function. To perform these experiments, we used the Simple Transformers library which allows us to easily implement this setup.¹¹

As main configuration parameters, we set the `max_length` parameter to 90,¹² and we re-trained the models up to 2 epochs. From here after, we refer as Bert-(FT) to the fine-tuned experiments. It is important to mention that the considered models, i.e., BERT, XLM, and DistilBERT, represent state-of-the-art language models available in the German Language. Although there are many other recent technologies, none of these are trained in German. Further details of employed models can be found at huggingface web page.¹³

4.3 Fully Connected Neural Network

As an additional classification method, we configured a fully connected neural network (FC). This type of artificial neural network is configured such that all the nodes, or neurons, in one layer, are connected to all neurons in the next layer. The topology of the employed network and its configuration parameters are mentioned in Table 7.

For the performed experiments using FCs, we passed as input features to the FC the sentence representation generated using BERT encodings. Thus, to generate the representation of the input text, we evaluate several configurations, namely: last hidden layer (LHL), concatenation of the 4 last hidden layers (Concat4LHL), min, max and mean pool of the last hidden layers. However, we only report the best performances obtained during the validation stage, i.e., LHL and Concat4LHL configurations. On the one hand, for generating the Concat4LHL representation we concatenate the last four layers values from the token [CLS]. As known, the [CLS] token at the beginning of the sentence is treated as the sentence representation. On the other hand, for

¹¹ <https://pypi.org/project/simpletransformers>

¹² For deciding this value, we consider the information obtained in the data analysis described in Section 3.

¹³ https://huggingface.co/transformers/pretrained_models.html

the LHL configuration, we preserve as the text representation the values of the last hidden layer from the token [CLS].

For the reported experiments under the FC method (See Figure 2), two configurations of BERT were tested for generating the LHL and Concat4LHL representation: i) pre-trained German encodings of BERT (distilbert-base-german-cased), referred as Bert(pre-trained); and ii) resultant fine-tuned BERT encodings from the re-training we explained in Section 4.2, referred as Bert(fine-tuned).

4.4 Pre-processing

For performing our experiments, we perform different pre-processing operations depending on the selected representation type, i.e., word/char n-grams or contextual/non-contextual word embeddings (see Figure 2).

In particular, for all the experiments performed with the pre-trained BERT embeddings, or in the fine-tuning process of the transformers architectures, we did not perform any type of pre-processing operation. Contrastly, when char/word n-grams or FastText embeddings are employed, we apply the following pre-processing operations to the input text: 1) we remove all non-alphabetical symbols, e.g., numbers, strange symbols; 2) every word is lower cased. Other pre-processing techniques, like removing stopwords, punctuation or replacing German umlauts (ä, Ä, ö, Ö, ü, Ü) and ligatures (e.g., ß) were not applied as previous research indicates that no improvement is obtained from doing it [40].

Finally, it is also worth mentioning that we did not apply any type of spelling or grammar corrector. We decide not to do it given that such types of errors have shown to be important style markers in several tasks of authorship analysis [45].

4.5 Evaluation metrics

For measuring the overall effectiveness of the classification process, we use standard set-based evaluation measures, such as precision, recall and macro F-score. This decision was based on agreement with previous work and the official OMT classification task that reports and rank results with these metrics [20], specifically using the macro F1.

Generally speaking, when evaluating a classification task, there are four types of outcomes that occur:

1. True positives (TP) refer to the case when the classifier predicts an observation belongs to class c and it actually belongs to that class.
2. True negatives (TN) refers to the case when the classifier predicts an observation not belonging to class c and it actually does not belong to that class.

3. False positives (TP) occur when the classifier predicts an observation belongs to class c when in reality it does not.
4. False negatives (FN) occur when the classifier predicts an observation as not belonging to class c when in fact it does.

Thus, precision (P) and recall (R) are defined as shown in expression Eq. 3 and Eq. 4 respectively.

$$P = \frac{TP}{TP + FP} \quad (3)$$

$$R = \frac{TP}{TP + FN} \quad (4)$$

The F-score (or F1), also known as the harmonic mean of P and R , is computed as follows:

$$F1 = 2 \times \frac{P \times R}{P + R} \quad (5)$$

Although the F1 score is a good metric to compare the performance of classifiers, the macro F1-score (F1-macro) is recommended to assess the quality of problems with multiple binary labels or multiple classes. Accordingly, the F1-macro is defined as the mean of the class-wise F1-scores (Eq. 6):

$$F1\text{-macro} = \frac{1}{N} \sum_{i=0}^N F1_i \quad (6)$$

where i is the class index and N the total number of classes. Notice that the F1-macro is not affected by the classes imbalance.

Although the Accuracy ($Acc = \frac{TP}{TP+FP+TN+FN}$) score is a common metric to compare performance results, the Acc is not recommended in classification problems where there is a large class imbalance. In such particular scenario, it is very likely that a model tends to predict the value of the majority class for all predictions and achieve a high classification accuracy, however, this does not mean that such model is useful in the posed task, a phenomenon known as the accuracy paradox [49].

4.6 Baseline and validation approaches

As a baseline, we replicated the approach proposed by the GermEval 2020 OMT task organizers, i.e., a linear Support Vector Classifier (SVC) using as a form of representation of the documents a traditional *tf-idf* strategy. Proposed baseline consists of a 30 (combined motive/level labels) binary SVCs (one-vs-all) classifiers.

In order to report robust and stable results, we implemented two different validation strategies. On the one hand, we performed a stratified k cross-fold-validation strategy with $k = 5$ using the entire dataset (*train+dev+test*); we

Method	Configuration type	Configuration sub-type	F1-macro $\mu(\pm\sigma)$	Precision $\mu(\pm\sigma)$	Recall $\mu(\pm\sigma)$
ST [†]	Bert	base-german-cased	0.701 (± 0.003)	0.701 (± 0.003)	0.702 (± 0.004)
ST [†]	XLM	mlm-ende-1024	0.691 (± 0.002)	0.693 (± 0.003)	0.693 (± 0.002)
ST [†]	Distilbert	base-german-cased	0.695 (± 0.004)	0.696 (± 0.003)	0.696 (± 0.004)
FC	Bert	LHL	0.589 (± 0.002)	0.603 (± 0.004)	0.586 (± 0.003)
FC	Bert	Concat4LHL	0.610 (± 0.006)	0.621 (± 0.003)	0.608 (± 0.007)
FC [†]	Bert-(FT)	LHL	0.693 (± 0.002)	0.689 (± 0.002)	0.697 (± 0.003)
FC [†]	Bert-(FT)	Concat4LHL	0.687 (± 0.002)	0.682 (± 0.001)	0.693 (± 0.003)
SAE	char	n-grams(1-2)	0.556 (± 0.001)	0.555 (± 0.004)	0.569 (± 0.003)
SAE	char	n-grams(1-3)	0.628 (± 0.006)	0.624 (± 0.006)	0.634 (± 0.005)
SAE	word	n-grams(1-2)	0.630 (± 0.004)	0.628 (± 0.006)	0.634 (± 0.005)
SAE	word	n-grams(1-3)	0.634 (± 0.002)	0.632 (± 0.002)	0.638 (± 0.002)
SAE	FastText	German-Wiki	0.614 (± 0.004)	0.614 (± 0.006)	0.618 (± 0.003)
SAE	Bert	LHL	0.558 (± 0.010)	0.558 (± 0.011)	0.567 (± 0.008)
SAE	Bert	Concat4LHL	0.563 (± 0.011)	0.572 (± 0.009)	0.568 (± 0.010)
SAE [†]	Bert-(FT)	LHL	0.674 (± 0.020)	0.669 (± 0.021)	0.683 (± 0.021)
SAE [†]	Bert-(FT)	Concat4LHL	0.656 (± 0.007)	0.651 (± 0.007)	0.664 (± 0.008)
Baseline	SVM	tf-idf	0.646(± 0.003)	0.650(± 0.003)	0.644(± 0.003)

Table 8: Average performance (μ) obtained across the 5-cross-fold-validation strategy (5CFV); the number between parenthesis represents the standard deviation (σ). Symbol ‘†’ indicates that the obtained performance is statistically significant according to the Wilcoxon test with $p < 0.05$.

refer to this configuration as ‘5CFV’ experiments. And, on the other hand, we report results on the *dev* and *test* partitions, which allows direct comparisons with the GermEval 2020 shared task participants.

5 Results and Discussion

The results of each of the considered approaches (see Figure 2) are reported in Tables 8 and 9 for the 5-fold validation strategy and for the *dev* partition respectively. Results are reported in terms of F1-macro, precision and recall metric.

For the results reported in Table 8, we can observe that the proposed baseline is able to reach an F1-macro of 64.5%, even though this baseline faces the OMT task as a 30 class problem. A similar behaviour is observed in Table 9, where the SVC baseline classifier yields good performance on the *dev* partition (F1 = 63.9%). Hence, we can conclude that the proposed SVC represents a hard baseline, showing stable results in both validation strategies.

In general, we can observe that our proposed supervised autoencoder is not able to generalize well in comparison to ST and FC methods. Observe that while in the 5CFV configuration the best result is achieved when fine-tuned encodings are used as text representation technique (Table 8, SAE(Bert-FT)) F1 = 67.4%; in the *dev* partition the best performance is obtained when input features are defined by word n-grams from range 1 to 2 (F1 = 63.4%). As known, word n-grams are useful to capture the identity of a word and its context.

Method	Configuration type	Configuration sub-type	F1-macro	Precision	Recall
ST [†]	Bert	base-german-cased	0.698	0.699	0.699
ST [†]	XLM	mlm-ende-1024	0.688	0.689	0.690
ST [†]	Distilbert	base-german-cased	0.692	0.694	0.693
FC	Bert	LHL	0.589	0.601	0.585
FC	Bert	Concat4LHL	0.616	0.620	0.617
FC	Bert-(FT)	LHL	0.673	0.675	0.673
FC [†]	Bert-(FT)	Concat4LHL	0.675	0.673	0.678
SAE	char	n-grams(1-2)	0.554	0.556	0.564
SAE	char	n-grams(1-3)	0.629	0.624	0.636
SAE	word	n-grams(1-2)	0.634	0.631	0.639
SAE	word	n-grams(1-3)	0.630	0.627	0.635
SAE	FastText	German-Wiki	0.474	0.485	0.476
SAE	Bert	LHL	0.540	0.555	0.547
SAE	Bert	Concat4LHL	0.571	0.569	0.575
SAE	Bert-(FT)	LHL	0.633	0.631	0.636
SAE	Bert-(FT)	Concat4LHL	0.624	0.623	0.629
<i>Baseline</i>	SVM	<i>tf-idf</i>	0.639	0.644	0.638

Table 9: Obtained results on the *dev* partition of the OMT classification task. Symbol ‘†’ indicates that the obtained performance is statistically significant according to the Wilcoxon test with $p < 0.05$.

Thus, these results indicate, to some extent, that the SAE attempts to exploit this information when solving the classification task. A similar performance is obtained when character n-grams are used as input features, specifically n-grams of size 1-3. These results are also interesting, as they are aligned with previous research findings, demonstrating the relevance of character n-grams in different non-thematic classification tasks [39, 37]. Char n-grams are capable of providing an excellent trade-off between sparseness and word’s identity, while at the same time they combine different types of information: punctuation, morphological makeup of a word, lexicon, and even context. As main observations of the SAE performance we can highlight that, using fine-tuned BERT encodings produced the best results under the 5CFV strategy (outperforming the proposed baseline), but word and character n-grams are not capable to improve the baseline performance. Similarly, for the experiments on the *dev* partition (Table 9) where even though the results obtained with the fine-tuned BERT encodings are similar to those obtained with word and char n-grams, none of these configurations were able to improve the SVC baseline (63.9%).

Regarding the performance of the FCs, the best performance is obtained when we use as features the fine-tune BERT encodings extracted from the last hidden layer (F1 = 69.3%) for the 5CFV experiments, and when the concatenation of the 4 last hidden layers is used (F1 = 67.5%), in the *dev* partition. In both cases, Table 8 and 9, we can observe an important difference on the performance of the FC when pre-trained or fine-tuned encodings are used.

Generally speaking, the impact of the fine-tuning allows a better performance of the neural networks (as expected), outperforming the proposed baseline in both cases.

Overall, based on these experiments (Table 8 and 9), the best performance (in terms of classification F1 score) was obtained by a simple transformer using BERT embedding. The attention-based architecture was found effective in comparison to FC and SAE methods. Consequently, during GermEval 2020 competition, we submitted a subset of what we found were the most effective configurations. Table 10 shows the performance of our submitted systems.

As can be observed in Table 10, our best performing system was the simple transformer architecture using BERT encodings. Specifically, this was our configuration that obtained the second place during the GermEval 2020 competition [20]. As a reference, we put at the bottom of the table the performance of the baseline system, and the performance obtained by the first and second places. As expected, the SAE were not able to improve the baseline system. However, our configuration based on the fully-connected network, using the fine-tuned BERT encodings was able to outperform the proposed baseline. It is worth mentioning that the winning approach during GermEval 2020 is based on a BERT methodology as well [40], with the pre-trained Digitale Bibliothek Münchener Digitalisierungszentrum (DBMDZ¹⁴) German model, validating the positive impact of transformers based methods. Although the employed methodology by the winning approach [40] and our best configuration is the same, there are a few variations that are important. We consider as base model the BERT pre-trained on German Wikipedia, Open-Legal-Data and news articles. On the contrary, [40] used the DBMDZ¹⁴ model. We did not apply any data correction process, while [40] made a data exploration to find all non-German texts, they applied an automatic translation process of all of these into German, and applied a spellchecker to correct spelling mistakes. Nevertheless, in spite of these extra effort, obtained results are very close for both techniques.

5.1 Results analysis

In this section we present a more detailed analysis of the obtained results by our best configuration. Accordingly, Figure 4 and Figure 5 show detailed classification results between the SVC baseline, and our best configuration (ST-BERT). For ease of understanding we split the problem into detection of motives (Figure 4) and detection of levels (Figure 5). It can be observed that the ST architecture significantly increases the number of correctly classified instances in motives M (+2.5%), L (+6.2%), F (+14.38%), and A (+5.9%), however, this situation is not the same for motive zero (-4.5%). A similar situation occurs in the levels detection task (Figure 5). For all the levels' categories, the ST is able to increase the number of correctly classified instances: 5 (+9.7%), 4 (+2.0%), 3 (+29.5%), 2 (+0.07%), 1 (+7.4%), 0 (+2.3%).

¹⁴ <https://github.com/dbmdz/berts>

Method	Configuration type	Configuration sub-type	F1-macro
ST [†]	Bert	base-german-cased	0.698
ST [†]	XLM	mlm-ende-1024	0.686
ST [†]	Distilbert	base-german-cased	0.688
FC	Bert-(FT)	LHL	0.671
SAE	char	n-grams(1-3)	0.628
SAE	word	n-grams(1-2)	0.634
<i>Baseline</i>	SVM	<i>tf-idf</i>	0.644
1st place [40]	-	-	0.704
3rd place [6]	-	-	0.678

Table 10: Obtained results on the *test* partition of the OMT classification task. Performance results are reported as given by the GermEval 2020 organizers [20]. Symbol ‘†’ indicates that the obtained performance is statistically significant according to the Wilcoxon test with $p < 0.05$.

As mentioned in section 4, instead of solving a 30 class problem, we split the OMT classification task in two separated problems, i.e., motives and levels classification problems. Thus, obtained results are aligned with the OMT theory [22], since according to our experiments it is possible to detect motives and levels separately, reinforcing the fact that motives and levels are not directly connected. Nevertheless, some of the methods presented during the GermEval 2020 did face the OMT problem as a 30 class classification task [20], which indicates that the OMT theory has to be revised and compared against what the NLP community has found.

In addition to the previous analysis, and given that a common concern is the lack of transparency of many deep learning architectures, we perform an analysis of what is the attention mechanism focusing on when solving the OMT task. The result of this analysis is shown in Table 12. The main intention of this type of analysis is to provide a better understanding of the connection between machine learning algorithms and language usage. To perform this analysis, we randomly select 5 sample texts produced by evaluated subjects in the *dev* partition. We show the results only for the distinct motives (A, F, L, M, O). For a fair comparison, we selected textual samples belonging to the same level class, in this case, all samples belong to level 4.

To visualize the attention-mechanism, we extract the attention weights based on [CLS] token from the last layer, average each token weight across all attention heads, and finally normalize weights across all tokens so that each weight is a value between 0.0 (very low attention) and 1.0 (very high attention)¹⁵. The highlight criteria of the words are shown in Table 11.

Table 12 demonstrates the visualization of how the attention-mechanism works in the Operant Motive Test classification task. Attention weights are

¹⁵ We adapted the implementation from: <https://github.com/uzaymacar/comparatively-finetuning-bert/tree/f579ad55bf7afee5292f40a0943eb0ef018abe83>

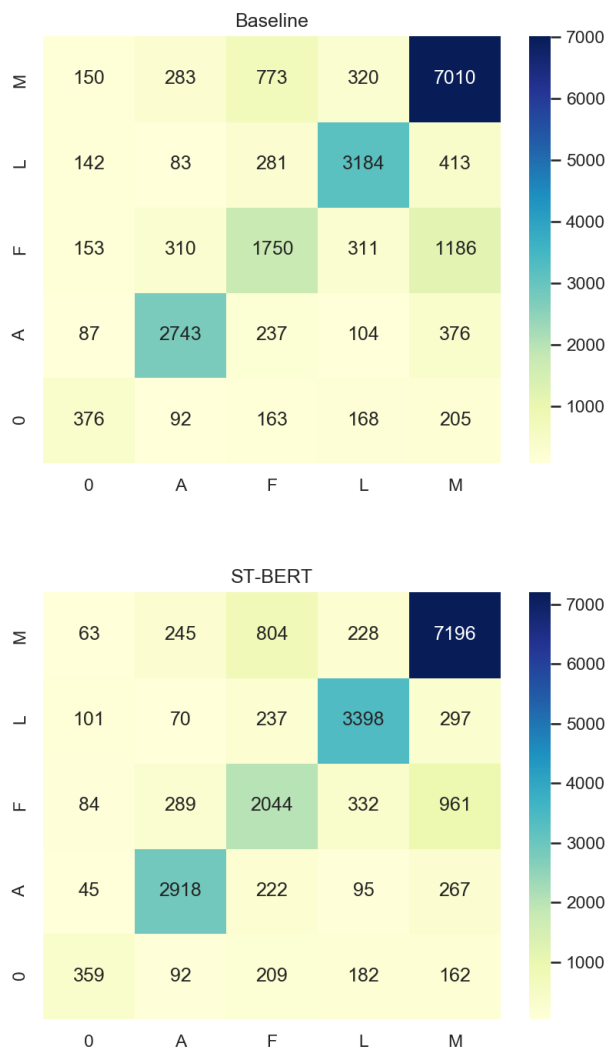


Fig. 4: Confusion matrices for the MOTIVES classification task: top-baseline performance; bottom-ST performance.

extracted after the fine-tuning of the BERT method. As an important observation, notice the attention of functional words: *ist, und, an, der, das, zu, sie, sind, ein, einer, von* (*is, and, at, the, the, to, they, are, a, one, from*). This indicates that for the simple transformer architecture the writing style is becoming more relevant at the moment of solving the classification task.

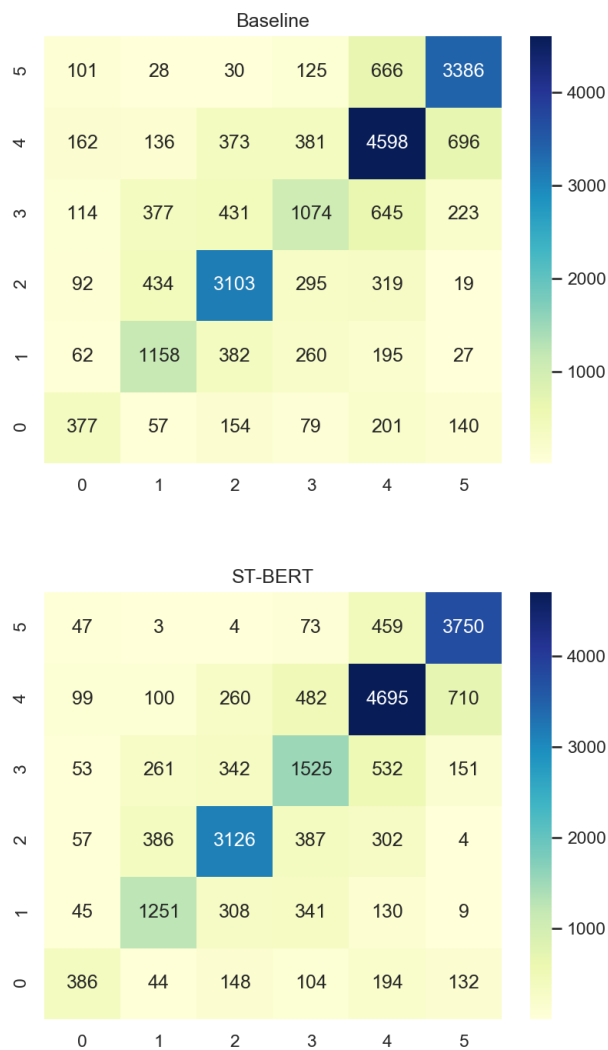


Fig. 5: Confusion matrices for the LEVELS classification task: top-baseline performance; bottom-ST performance.

And additional observation is the attention paid to the word's ending. As known, during the tokenization process, unknown words are split into smaller tokens. When this is the case, the symbol '#' is added to the generated tokens. Especially for motive L, we can notice many cases where tokens with the symbol '#' are receiving the attention from the simple transformer architecture.

Attention weights	Description	Highlight color
0.00 - 0.20	Very Low Attention	N/A
0.20 - 0.40	Low Attention	■
0.40 - 0.60	Medium Attention	■
0.60 - 0.80	High Attention	■
0.80 - 1.00	Very High Attention	■

Table 11: Followed criteria of highlight colors used in the visualization of the attention-mechanism

Motive	Sample text (closest English translation)
A	sie braucht Verständnis und wendet sich an jemand der ihr zu ##hört und sie versteht. sie fühlt sich geborgen und angenommen und erzählt , was sie belastet. sie ist angenommen so wie sie ist. (she needs understanding and turns to someone who listens to and understands her; she feels safe and accepted and tells what is burdensome; she is accepted as she is.)
F	alles richtig gemacht zu haben und dem gegenüber Par ##oli bietet , trotz ##ig. weil das gegenüber unein ##sichtig scheint. (to have done everything correctly and to stand up to the other person, despite the fact that the other party seems unreasonable.)
L	sie kl ##etter ##t ein berg hoch, und darf das gleich ##gewicht nicht verlieren. ein bisschen in einer unangen ##hem ##ne position. sie dar ##chte nicht dass es so schwer wäre. sie sch ##af ##t es mit mut h ##ich zu kommen, und ist erstaun ##t das sie es gesch ##af ##t hat. glücklich eine neue seit ##e von ihr selbst zu finden. (she climbs a mountain and must not lose her balance. a little in an uncomfortable position. she does not think that it would be so difficult. she manages to do it with courage and is amazed that she did it. fortunately a new one to find her side.)
M	sie möchte der anderen Person zeigen, dass sie enttäuscht ist. enttäuscht über die andere Person. weil die andere Person nicht ihren Erwartungen stand gehalten hat. (she wants to show the other person that she is disappointed, disappointed with the other person, because the other person has not lived up to her expectations.)
O	die Personen überlegen nach einer Lösung. verzwei ##felt und pa ##nisch. weil sie kurz vor dem ert ##rin ##ken sind?. gut . . . alle drei werden gerettet. (the people are thinking of a solution, desperate and panicked, because they are about to drown? ... good ... all three are saved.)

Table 12: Attention-mechanism visualization for the OMT classification task

Notice that negation words are resulting very important, e.g., *nicht*, as well as some punctuation marks, e.g., ‘.’, ‘?’.

Furthermore, in Figure 6 we show the usage given to the top 25 terms with higher attention values. For performing this analysis, we obtained the top most important words, i.e., words with higher attention values, for each motive category (A, F, L, M). Then, to obtain these 25 words, we intersected the corresponding sets. The figure illustrates the relative frequency given to each of these words according to the motive class. As can be seen in the figure, subjects

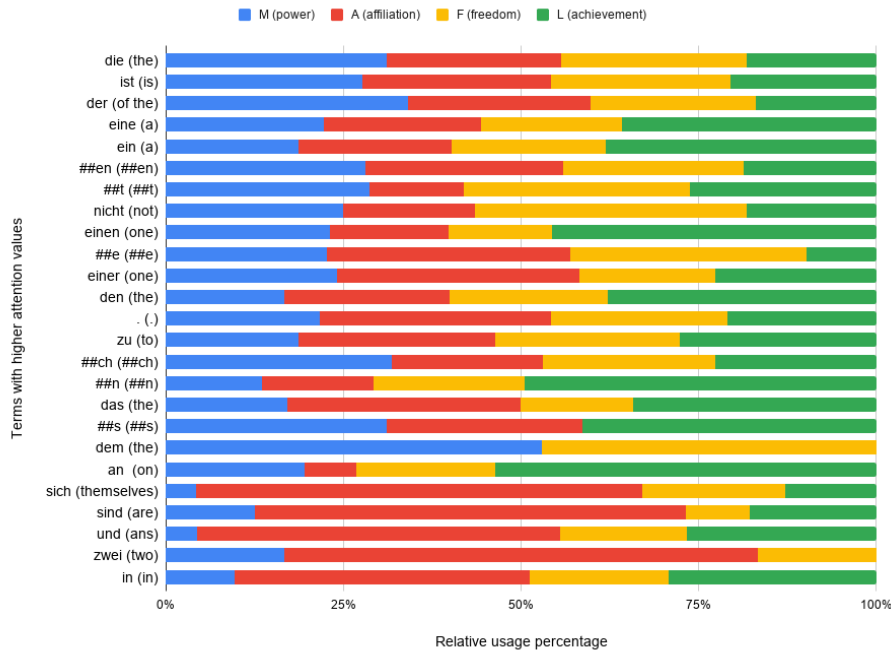


Fig. 6: Top 25 terms with higher attention across MOTIVES categories.

from different categories use these words with different frequency values. This frequency analysis also explains the good performance of the SVM classifier; which is based on a traditional *tf-idf* vectorial representation. However, even though frequency counts are helping the SVM to accurately separate among classes, the context in which these words appear is important, i.e., how users are employing these lexical units is relevant for solving the task.

Accordingly, in Figure 7 and Figure 8 we illustrate the context in which words *nicht* (*not*) and *sind* (*are*) are employed in our dataset. For this analysis, we took all the text generated by users from the same category (i.e., M, A, F, and L), and perform a collocation analysis fixing a target word (in this case: *nicht/sind*). As known, a collocation is a sequence of words that co-occur with high frequency within some corpus. Thus, for generating the visualization of each tree, we kept the most frequent collocations from each category. From this analysis, it is possible to observe that, even though these target words are frequently used by subjects, the employed contexts by each category are very different from each other. For example, subjects labeled with motive M (power) use the *nicht* words in an imperative/control fashion, e.g., *kommen nicht auf* (do not come up), while subjects categorized with A (affiliation) motive use it to show concern about others, e.g., *sie nicht alleine* (they are not alone). Similarly, for F (freedom) class, subjects use this negation to indi-

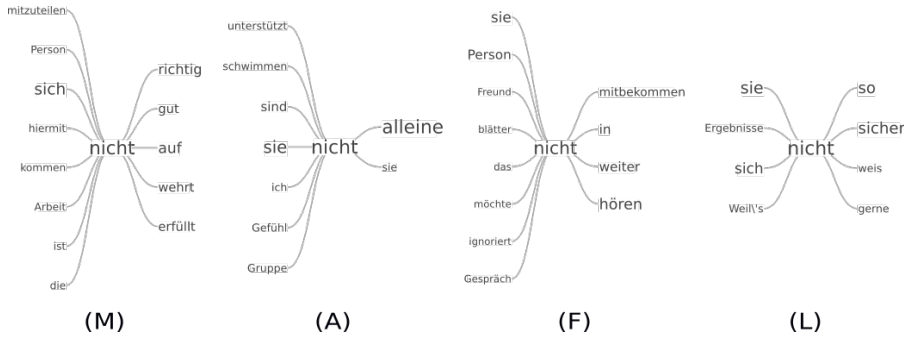


Fig. 7: Contextual tree of word *nicht*. Words in the leaves represent the most frequent words appearing next to the target word, in this case, *nicht*. Below each tree, its corresponding motive class is mentioned.

cate concern about themselves, e.g., *möchte nicht mitbekommen* (don't want to notice); and in the L (achievement) motive, the common context denotes insecurity, e.g., *sie nicht weiss* (she doesn't know). Hence, the good performance of transformer-based NN architectures is explained by this analysis, as the attention mechanism of BERT is able to learn contextual relations between words (and sub-words) from the input text. Notice that these findings are aligned with previous psychometrics research (see Table 2).

This analysis provided interesting insights that we can summarize as: ST architecture pays higher attention to the use of personal pronouns, stop words, negation, punctuation marks, unknown words, and some conjugation styles, filtering out most of the unimportant elements such as content words. But not just isolated words, the context in which these words appear are providing important information to the transformer-based NN methods at the moment of detecting motives-levels. In other words, the writing style (*how we write*) is more relevant than content words (*what we write*) for solving the OMT classification task. Additionally, it is particularly interesting how the usage of negations words (*nicht*, and *##t* which correspond to “-n't” contractions) are frequently used by the power (M) and the freedom (F) motives. This finding is partially aligned with previously reported from the psychological theory [54], where it has been showed that so-called activity inhibition (AI) trait is mainly described as negations in combination with the power motive. Finally, our performed collocations analysis helped to understand and visualize the type of context that is helping the transformer-based NN architecture to solve the problem more accurately. Overall, these findings could foster implicit psychometrics theory, and consequently, advanced aptitude diagnostics supported by NLP technologies.

The aptitude test (i.e., OMT) is a type of psychological test that could affect the subjects' lives, specially if performed automatically without any human intervention [21]. Hence, there is an important urgency for understanding how this type of automatic decisions are being done by recent machine learning

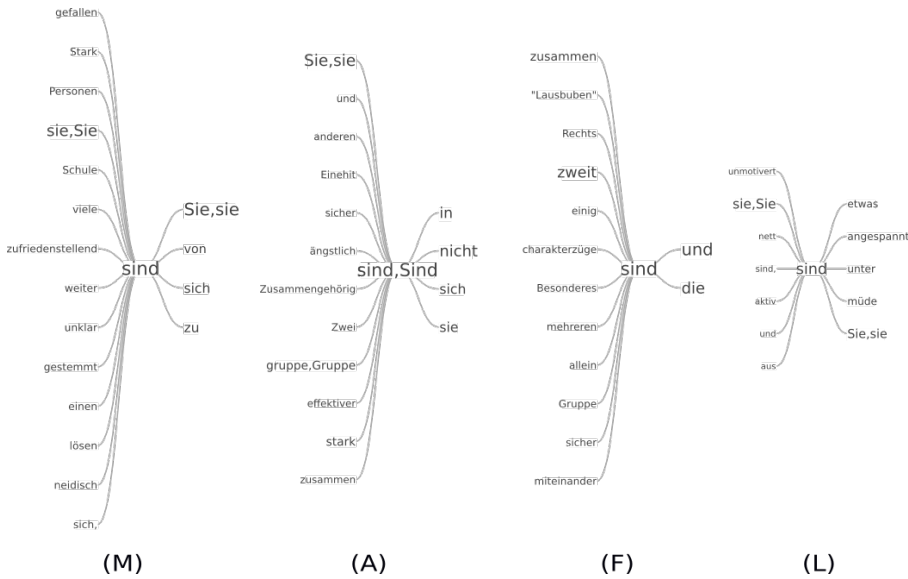


Fig. 8: Contextual tree of word *sind*. Words in the leaves represent the most frequent words appearing next to the target word, in this case, *sind*. Below each tree, its corresponding motive class is mentioned.

technologies. As stated in [15], explainable methods are becoming more relevant, particularly in the health-care domain. Thus, it is necessary to consider many aspects when designing explainable ML methods, e.g., *who is the domain expert?*, *who are the affected users?*, among others [2, 13]. Accordingly, and as part of our future work, we plan to extend our interpretability analysis towards the design of responsible artificial intelligence algorithms (i.e., explainable and transparent) in the context of mental health automated analysis by applying some of the proposed recommendations in [2].

6 Conclusion

This paper represents a first step towards the analysis of recent NLP technologies for solving the OMT classification task. To this end, we performed a comparative analysis among state-of-the-art Simple Transformers based architectures, e.g., BERT, XLM, and DistilBert, very recent generalization techniques as Supervised Autoencoders and traditional machine learning techniques. Notably, transformers based methods exhibit the best empirical results, obtaining a relative improvement of 7.9% over the baseline suggested as part of the GermEval 2020 challenge [20]. We performed an exploration on how the attention-mechanism is working in this particular task, and obtained results revealed that features associated with the writing style are more important than content-based words. Some of these findings shown strong connections to

behavioral research made on the implicit psychometrics theory. For example, as the result of our performed analysis, we observed that the usage of negations in combination with the power motive it is supported by the research made by [54]. As future work, we plan to evaluate the impact of hyperparameter tuning through optimization methods, such as Bayes optimizer [42], evaluate the impact of early-fusion strategies in the performance of the SAE, and to perform further analysis on how the attention-mechanism from the transformers architecture is working in the OMT task.

Finally, we would like to emphasize the importance of the ethical necessity of carefully understanding the research being done in the field of NLP & psychology. Although NLP technologies indicate that solving this type of tasks is, to some extent, possible, further research needs to be conducted to carefully explain the relation between psychological tests and subjects aptitudes. The authors would like to clearly state that we are against the use of this type of technology to discriminate against people in any type of our daily life situations. Even though we believe that this research is important, as can be useful for psychologists professionals, claiming that the NLP/ML community is able to accurately classify users according to their professional aptitudes and personality traits is not something we agree on. We support the idea that this type of research can help to validate previous theories as well as to support mental health care practitioners to evaluate or get important insights from closed and controlled studies.

Acknowledgements Esaú Villatoro-Tello was partially supported by Idiap Research Institute, SNI-CONACyT Mexico, and UAM-Cuajimalpa Mexico during the elaboration of this work. We would like to thank the GERMEVAL 2020 organizers, in particular, Dirk Johannßen for all the provided help.

Compliance with Ethical Standards

Ethical approval: This article does not contain any studies with human participants or animals performed by any of the authors. The authors declare that they have no conflict of interest. The authors do not see any significant ethical or privacy concerns that would prevent the processing of the data used in the study. All datasets containing personal data are processed in compliance with the GDPR and national law.

References

1. Aragón, M.E., López-Monroy, A.P., González-Gurrola, L.C., Montes, M.: Detecting depression in social media using fine-grained emotions. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pp. 1481–1486 (2019)

2. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F.: Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion* **58**, 82–115 (2020). DOI <https://doi.org/10.1016/j.inffus.2019.12.012>. URL <https://www.sciencedirect.com/science/article/pii/S1566253519308103>
3. Baum, I.R., Baumann, N.: Autonomous creativity: The implicit autonomy motive fosters creative production and innovative behavior at school. *Gifted and Talented International* **33**(1-2), 15–25 (2018)
4. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* **5**, 135–146 (2017)
5. Burdick, L., Mihalcea, R., Boyd, R.L., Pennebaker, J.W.: Analyzing connections between user attributes, images, and text. *Cognitive Computation* pp. 1–20 (2020)
6. Çöltekin, C.: Predicting educational achievement using linear models. In: *Proceedings of the GermEval 2020 Task 1 Workshop in conjunction with the 5th SwissText & 16th KONVENS Joint Conference*, pp. 23–29 (2020)
7. Conneau, A., Lample, G.: Cross-lingual language model pretraining. In: *Advances in Neural Information Processing Systems*, pp. 7057–7067 (2019)
8. Demasi, O., Hearst, M.A., Recht, B.: Towards augmenting crisis counselor training by improving message retrieval. In: *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pp. 1–11. Association for Computational Linguistics, Minneapolis, Minnesota (2019). DOI 10.18653/v1/W19-3001. URL <https://www.aclweb.org/anthology/W19-3001>
9. Denzinger, F., Brandstätter, V.: Stability of and changes in implicit motives. a narrative review of empirical studies. *Frontiers in psychology* **9**, 777 (2018)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186 (2019). DOI 10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>
11. Fitzpatrick, K.K., Darcy, A., Vierhile, M.: Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (woebot): a randomized controlled trial. *JMIR mental health* **4**(2), e19 (2017)
12. Gawronski, B., De Houwer, J.: Implicit measures in social and personality psychology. *Handbook of research methods in social and personality psychology* **2**, 283–310 (2014)
13. Holzinger, A., Malle, B., Saranti, A., Pfeifer, B.: Towards multi-modal causability with graph neural networks enabling information fusion for explainable ai. *Information Fusion* **71**, 28–37 (2021). DOI <https://doi.org/>

- 10.1016/j.inffus.2021.01.008. URL <https://www.sciencedirect.com/science/article/pii/S1566253521000142>
14. Husseini Orabi, A., Buddhitha, P., Husseini Orabi, M., Inkpen, D.: Deep learning for depression detection of twitter users. In: Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic, pp. 88–97. Association for Computational Linguistics, New Orleans, LA (2018). DOI 10.18653/v1/W18-0609. URL <https://www.aclweb.org/anthology/W18-0609>
 15. Ieracitano, C., Mammone, N., Hussain, A., Morabito, F.C.: A novel explainable machine learning approach for eeg-based brain-computer interface systems. *Neural Computing and Applications* pp. 1–14 (2021)
 16. Inkster, B., Sarda, S., Subramanian, V.: An empathy-driven, conversational artificial intelligence agent (wysa) for digital mental well-being: real-world data evaluation mixed-methods study. *JMIR mHealth and uHealth* **6**(11), e12106 (2018)
 17. Johannßen, D., Biemann, C.: Between the Lines: Machine Learning for Prediction of Psychological Traits - A Survey **LNCS-11015**, 192–211 (2018). DOI 10.1007/978-3-319-99740-7_13. URL <https://hal.inria.fr/hal-02060047>. Part 2: MAKE-Text
 18. Johannßen, D., Biemann, C.: Neural classification with attention assessment of the implicit-association test omt and prediction of subsequent academic success. In: Proceedings of the 15th Conference on Natural Language Processing (KONVENS 2019): Long Papers, pp. 68–78. German Society for Computational Linguistics & Language Technology, Erlangen, Germany (2019)
 19. Johannßen, D., Biemann, C., Scheffer, D.: Reviving a psychometric measure: Classification and prediction of the operant motive test. In: Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology, pp. 121–125 (2019)
 20. Johannßen, D., Biemann, C., Remus, S., Baumann, T., Scheffer, D.: Germeval 2020 task 1: Classification and regression of cognitive and motivational style from text. In: Proceedings of the GermEval 2020 Task 1 Workshop in conjunction with the 5th SwissText & 16th KONVENS Joint Conference, pp. 1–10 (2020)
 21. Johannßen, D., Biemann, C., Scheffer, D.: Ethical considerations of the germeval20 task 1. iq assessment with natural language processing: Forbidden research or gain of knowledge? In: Proceedings of the GermEval 2020 Task 1 Workshop in conjunction with the 5th SwissText & 16th KONVENS Joint Conference, pp. 30–44 (2020)
 22. Kuhl, J., Scheffer, D.: Der operante multi-motiv-test (omt): Manual [the operant multi-motive-test (omt): Manual]. Germany: University of Osnabrück (1999)
 23. Landwehr, N., Hall, M., Frank, E.: Logistic model trees. *Machine Learning* **59**, 161–205 (2005). DOI 10.1007/s10994-005-0466-3
 24. Le, L., Patterson, A., White, M.: Supervised autoencoders: Improving generalization performance with unsupervised regularizers. In: Advances in

- Neural Information Processing Systems, pp. 107–117 (2018)
25. Losada, D.E., Crestani, F., Parapar, J.: Overview of erisk: Early risk prediction on the internet. In: International Conference of the Cross-Language Evaluation Forum for European Languages, pp. 343–361. Springer (2018)
 26. Losada, D.E., Crestani, F., Parapar, J.: erisk 2020: Self-harm and depression challenges. In: J.M. Jose, E. Yilmaz, J. Magalhães, P. Castells, N. Ferro, M.J. Silva, F. Martins (eds.) *Advances in Information Retrieval*, pp. 557–563. Springer International Publishing, Cham (2020)
 27. Masrani, V., Murray, G., Field, T., Carenini, G.: Detecting dementia through retrospective analysis of routine blog posts by bloggers with dementia. In: *BioNLP 2017*, pp. 232–237. Association for Computational Linguistics, Vancouver, Canada, (2017). DOI 10.18653/v1/W17-2329. URL <https://www.aclweb.org/anthology/W17-2329>
 28. Matero, M., Idnani, A., Son, Y., Giorgi, S., Vu, H., Zamani, M., Limbachiya, P., Guntuku, S.C., Schwartz, H.A.: Suicide risk assessment with multi-level dual-context language and BERT. In: *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pp. 39–44. Association for Computational Linguistics, Minneapolis, Minnesota (2019). DOI 10.18653/v1/W19-3005. URL <https://www.aclweb.org/anthology/W19-3005>
 29. McClelland, D.C., Koestner, R., Weinberger, J.: How do self-attributed and implicit motives differ? *Psychological review* **96**(4), 690 (1989)
 30. Murray, H.: *Thematic Apperception Test*. Harvard University Press (1943). URL <https://books.google.co.in/books?id=r9a0CsRpEzYC>
 31. Pennebaker, J.W.: The secret life of pronouns. *New Scientist* **211**(2828), 42–45 (2011)
 32. Pennebaker, J.W., Chung, C.K., Frazee, J., Lavergne, G.M., Beaver, D.I.: When small words foretell academic success: The case of college admissions essays. *PloS one* **9**(12), e115844 (2014)
 33. Pennebaker, J.W., Francis, M.E., Booth, R.J.: *Linguistic inquiry and word count: Liwc 2001*. Mahway: Lawrence Erlbaum Associates **71**(2001), 2001 (2001)
 34. Pennebaker, J.W., Mehl, M.R., Niederhoffer, K.G.: Psychological aspects of natural language use: Our words, our selves. *Annual review of psychology* **54**(1), 547–577 (2003)
 35. Reece, A.G., Reagan, A.J., Lix, K.L., Dodds, P.S., Danforth, C.M., Langer, E.J.: Forecasting the onset and course of mental illness with twitter data. *Scientific reports* **7**(1), 1–11 (2017)
 36. Ramírez-de-la Rosa, G., Villatoro-Tello, E., Jiménez-Salazar, H.: Txpi-u: A resource for personality identification of undergraduates. *Journal of Intelligent & Fuzzy Systems* **34**(5), 2991–3001 (2018)
 37. Sánchez-Vega, F., Villatoro-Tello, E., Montes-y Gómez, M., Rosso, P., Stamatatos, E., Villaseñor-Pineda, L.: Paraphrase plagiarism identification with character-level features. *Pattern Analysis and Applications* **22**(2), 669–681 (2019)

38. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arXiv preprint arXiv:1910.01108 (2019)
39. Sapkota, U., Bethard, S., Montes, M., Solorio, T.: Not all character n-grams are created equal: A study in authorship attribution. In: Proceedings of the 2015 conference of the North American chapter of the association for computational linguistics: Human language technologies, pp. 93–102 (2015)
40. Schäfer, H., Idrissi-Yaghir, A., Schimanowski, A., Bujotzek, M.R., Damm, H., Nagel, J., Friedrich, C.M.: Predicting cognitive and motivational style from german text using multilingual transformer architectures. In: Proceedings of the GermEval 2020 Task 1 Workshop in conjunction with the 5th SwissText & 16th KONVENS Joint Conference, pp. 17–22 (2020)
41. Shen, J.H., Rudzicz, F.: Detecting anxiety through Reddit. In: Proceedings of the Fourth Workshop on Computational Linguistics and Clinical Psychology — From Linguistic Signal to Clinical Reality, pp. 58–65. Association for Computational Linguistics, Vancouver, BC (2017). DOI 10.18653/v1/W17-3107. URL <https://www.aclweb.org/anthology/W17-3107>
42. Snoek, J., Larochelle, H., Adams, R.P.: Practical bayesian optimization of machine learning algorithms. In: Advances in neural information processing systems, pp. 2951–2959 (2012)
43. Sokolowski, K., Schmalt, H.D., Langens, T.A., Puca, R.M.: Assessing achievement, affiliation, and power motives all at once: The multi-motive grid (mmg). *Journal of personality assessment* **74**(1), 126–145 (2000)
44. Souri, A., Rahmani, A., Hosseinpour, S.: Personality classification based on profiles of social networks’ users and the five-factor model of personality. *Human-centric Computing and Information Sciences* **8**, 8–24 (2018). DOI 10.1186/s13673-018-0147-4
45. Stamatatos, E.: A survey of modern authorship attribution methods. *Journal of the American Society for information Science and Technology* **60**(3), 538–556 (2009)
46. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: Advances in neural information processing systems, pp. 3104–3112 (2014)
47. Szatloczki, G., Hoffmann, I., Vincze, V., Kalman, J., Pakaski, M.: Speaking in alzheimer’s disease, is that an early sign? importance of changes in language abilities in alzheimer’s disease. *Frontiers in aging neuroscience* **7**, 195 (2015)
48. Tackman, A.M., Sbarra, D.A., Carey, A.L., Donnellan, M.B., Horn, A.B., Holtzman, N.S., Edwards, T.S., Pennebaker, J.W., Mehl, M.R.: Depression, negative emotionality, and self-referential language: A multi-lab, multi-measure, and multi-language-task research synthesis. *Journal of personality and social psychology* **116**(5), 817 (2019)
49. Valverde-Albacete, F.J., Carrillo-de Albornoz, J., Peláez-Moreno, C.: A proposal for new evaluation metrics and result visualization technique for sentiment analysis tasks. In: P. Forner, H. Müller, R. Paredes, P. Rosso,

- B. Stein (eds.) *Information Access Evaluation. Multilinguality, Multimodality, and Visualization*, pp. 41–52. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
50. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Advances in neural information processing systems*, pp. 5998–6008 (2017)
 51. Villatoro-Tello, E., Dubagunta, P., Fritsch, J., Ramírez-de-la Rosa, G., Motlicek, P., Magimai-Doss, M.: Late fusion of the available lexicon and raw waveform-based acoustic modeling for depression and dementia recognition. In: *Proceedings of the INTERSPEECH 2021. ISCA-International Speech Communication Association 2021*. (2021)
 52. Villatoro-Tello, E., Parida, S., Kumar, S., Motlicek, P., Zhan, Q.: Idiap & uam participation at germeval 2020: Classification and regression of cognitive and motivational style from text. In: *Proceedings of the GermEval 2020 Task 1 Workshop in conjunction with the 5th SwissText & 16th KONVENS Joint Conference*, pp. 11–16 (2020)
 53. Welch, C., Lahmala, A., Perez-Rosas, V., Shen, S., Seraj, S., An, L., Resnicow, K., Penebaker, J., Mihalcea, R.: Expressive interviewing: A conversational system for coping with COVID-19. In: *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020. Association for Computational Linguistics, Online* (2020). DOI 10.18653/v1/2020.nlpCOVID19-2.6. URL <https://www.aclweb.org/anthology/2020.nlpCOVID19-2.6>
 54. Winter, D.G.: The role of motivation, responsibility, and integrative complexity in crisis escalation: comparative studies of war and peace crises. *Journal of personality and social psychology* **92**(5), 920 (2007)
 55. Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Brew, J.: Hugging-face’s transformers: State-of-the-art natural language processing. *ArXiv abs/1910.03771* (2019)
 56. Wykes, T., Lipshitz, J., Schueller, S.M.: Towards the design of ethical standards related to digital mental health and all its applications. *Current Treatment Options in Psychiatry* **6**(3), 232–242 (2019)
 57. Zhu, Q., Zhang, R.: A classification supervised auto-encoder based on predefined evenly-distributed class centroids. *arXiv preprint arXiv:1902.00220* (2019)
 58. Zomick, J., Levitan, S.I., Serper, M.: Linguistic analysis of schizophrenia in Reddit posts. In: *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pp. 74–83. Association for Computational Linguistics, Minneapolis, Minnesota (2019). DOI 10.18653/v1/W19-3009. URL <https://www.aclweb.org/anthology/W19-3009>