

BERT meets LIWC: Exploring State-of-the-Art Language Models for Predicting Communication Behavior in Couples' Conflict Interactions

Jacopo Biggiogera
University of Surrey
Guildford, U.K.
jb01493@surrey.ac.uk

George Boateng
ETH Zurich
Zurich, Switzerland
gboateng@ethz.ch

Peter Hilpert
University of Lausanne
Lausanne, Switzerland
peter.hilpert@unil.ch

Matthew Vowels
University of Surrey
Guildford, U.K.
m.j.vowels@surrey.ac.uk

Guy Bodenmann
University of Zurich
Zurich, Switzerland
guy.bodenmann@psychologie.uzh.ch

Mona Neysari
University of Zurich
Zurich, Switzerland
m.neysari@psychologie.uzh.ch

Fridtjof Nussbeck
University of Konstanz
Konstanz, Germany
fridtjof.nussbeck@uni-konstanz.de

Tobias Kowatsch
tkowatsch@ethz.ch
ETH Zurich
Zurich, Switzerland
University of St. Gallen
St. Gallen, Switzerland

ABSTRACT

Many processes in psychology are complex, such as dyadic interactions between two interacting partners (e.g., patient-therapist, intimate relationship partners). Nevertheless, many basic questions about interactions are difficult to investigate because dyadic processes can be within a person and between partners, they are based on multimodal aspects of behavior and unfold rapidly. Current analyses are mainly based on the behavioral coding method, whereby human coders annotate behavior based on a coding schema. But coding is labor-intensive, expensive, slow, focuses on few modalities, and produces sparse data which has forced the field to use average behaviors across entire interactions, thereby undermining the ability to study processes on a fine-grained scale. Current approaches in psychology use LIWC for analyzing couples' interactions. However, advances in natural language processing such as BERT could enable the development of systems to potentially automate behavioral coding, which in turn could substantially improve psychological research. In this work, we train machine learning models to automatically predict positive and negative communication behavioral codes of 368 German-speaking Swiss couples during an 8-minute conflict interaction on a fine-grained scale (10-seconds sequences) using linguistic features and paralinguistic features derived with openSMILE. Our results show that both simpler TF-IDF features as well as more complex BERT features performed better

than LIWC, and that adding paralinguistic features did not improve the performance. These results suggest it might be time to consider modern alternatives to LIWC, the de facto linguistic features in psychology, for prediction tasks in couples research. This work is a further step towards the automated coding of couples' behavior which could enhance couple research and therapy, and be utilized for other dyadic interactions as well.

KEYWORDS

Couples, predicting observer ratings, multimodal fusion, behavioral signal processing, BERT, LIWC, SVM

1 INTRODUCTION

There are many processes in the field of psychology that are very complex such as dyadic interactions — interactions between two people [23]. These processes are difficult to investigate because each person's behavior is multimodal, both persons influence each other's behavior mutually, and this process unfolds rapidly [21]. Such dynamic processes are relevant for a large number of human interactions (e.g., romantic partners, patient-therapist, student-teacher, buyer-seller).

Of the different human interactions, conflict interactions in intimate relationships have been well studied over the last decades [11]. Results indicate two principal types of communication behaviors: functional and dysfunctional. For example, contempt and criticism are a reliable predictor for later divorce and therefore seen as negative or dysfunctional, whereas providing appreciation and taking responsibility are considered to be functional and are associated with stable relationships [19, 20, 22]. It is therefore important to understand conflict interactions better as divorces are not only often emotionally and financially difficult for partners, but also have long-term negative consequences on the children involved [2].

The major reason for the disappointing progress of understanding behavioral processes during conflict interactions is the lack of methods that enable an automated approach for a fine-grained understanding of behavior. Traditionally, analyses in interaction research are mainly undertaken using data obtained from observer rating methods which are labor-intensive, expensive and time-consuming [27]. Consequently, codes are generally assigned on a global scale (e.g., one rating for 8-10 minutes sessions) rather than on a fine-grain scale (e.g., every talk turn or 10 seconds) resulting in sparse data. While observer ratings provide a means to capture global aspects of behavior (e.g., positive behavior), the analysis of such global behavioral aspects and sparse data has forced the field to focus on predictions based on average behaviors across entire interactions, thereby undermining the ability to study intra- and inter-individual processes [23].

Beyond observer rating methods, psychology has also included technology to extract linguistic (i.e., *what* was said) and paralinguistic features (i.e., *how* it was said). Various paralinguistic features have been extracted mainly using Praat [10] and openSMILE [18] which are software tools that compute various acoustic features over audio signals (e.g., pitch, fundamental frequency) over sequential time segments (e.g., 25 ms). They have been used in various works for example to show that the fundamental frequency of oscillation of the vocal folds is a valid proxy for emotional arousal [25] and a larger range in fundamental frequency is associated with more conflict interactions [4, 5]. Furthermore, a specific set of 88 features computed called eGeMAPS have been shown to be a minimalist feature set that performs well for affective recognition tasks [17].

Linguistic features have mainly been extracted through word-count-based programs like Linguistic Inquiry and Word Count (LIWC) [36] which is a software for extracting the count of words using an existing list of words and categories (e.g., positive/negative words, personal pronouns, social process). Its usage in couples research for example has shown the words partners utilize during conflict significantly affect interaction and overall marital quality. Findings indicate that greater first-person plural pronoun usage ('we'), compared to first-person singular pronoun usage ('I') produces more positive resolutions to conflicts [34, 38]. Tools such as LIWC however are not without their limitations. In fact, they depend on the accuracy and comprehensiveness of the dictionary they are based upon, together with not being able to take into account both the context that words are placed in and the different meanings they might hold [3]. In a context such as conflict interactions where specific word choices and their meanings are important in affecting

how the conflict unfolds [38], such limitations hold great significance for the validity and accuracy of applications that use such tools. However, recent advances in natural language processing such as Bidirectional Encoder Representations and Transformations (BERT) [16] based on the Transformer architecture [44] have been shown to set new state-of-the-art records in various natural language understanding tasks such as natural language inference, question answering, and sentiment analysis. Some prior works have evaluated the predictive capability of BERT relative to LIWC in psychotherapy and mental health classification with BERT outperforming models based on LIWC features in populations with mental health diagnosis [24, 39]. Yet, BERT features have not yet been used in couples' interaction research for prediction tasks.

Some studies have used linguistic and/or paralinguistic features specifically to predict behavioral codes for interacting romantic partners with the goal of automating behavioral coding. Most of these works have focused on session-level prediction — predicting one code for the whole 8-10 minute session [7, 8, 12, 13, 26, 30–32, 40–42, 46] with a scarcity of works focused on prediction of fine-grained behavioral codes such as at the speaker turn level or every few seconds. One such work is that of Chakravarthula et al [14] in which they trained machine learning models to predict 3 behavioral codes on a speaker turn level of 85 couples' 10-minute conversations using paralinguistic features (from openSMILE) and linguistic features (custom sentence embedding model) and achieved 57.4% unweighted average recall (balanced accuracy) for 3-class classification. Leveraging advanced sentence embedding methods such as BERT could potentially improve performance and increase the potential of automating behavioral coding. Yet, it has not been investigated in the context of couples research. Furthermore, including paralinguistic features could potentially enable better recognition.

In order to overcome current limitations, we utilized a data set collected from 368 couples (N = 736 participants) who were recorded during an 8-minute conflict interaction. Our main goal is to examine how linguistic and paralinguistic features in 10-second sequences can be used to predict how the same sequence is perceived and rated by human coders as positive or negative communication behavior. We aim to answer the following research questions (RQs):

RQ1: Which linguistics features — LIWC or BERT — are better for predicting sequences-to-sequences-rated positive and negative communication behavior of partners?

RQ2: Given that the raters focused on coding linguistic aspects of behavior, how does adding openSMILE's eGeMAPS paralinguistic features affect the prediction performance?

Our contributions are (1) an evaluation of the predictive capability of BERT vis-à-vis LIWC in the context of the automatic recognition of couples' communication behavioral codes on a fine-grained time-scale (every 10 seconds) (2) an investigation into how the addition of paralinguistic features affects prediction performance (3) the use of a unique dataset — spontaneous, real-life, speech data collected from German-speaking Swiss couples (n=368 couples, N=736 participants), and the largest ever such dataset used in the literature for automatic coding of couples' behavior. The insights from our work would enable the usage of new technologies to potentially automate the behavioral coding of couples, which could substantially improve the efficiency of couples research.

2 METHODOLOGY

Data Collection and Preprocessing: This work used data from a larger dyadic interaction laboratory project conducted at the premises of the University of Zurich, Switzerland over 10 years with 368 heterosexual German-speaking, Swiss couples (N=736 participants; age 20-80) [29, 43]. The inclusion criterion was to have been in the current relationship for at least 1 year. Couples had to choose one problematic topic for the conflict interaction from a list of common problems, and participants were then videotaped as they discussed the selected issue for 8 minutes. The data used in this work had one interaction from each couple and consequently, 368 8-minute interactions.

Two research assistants were trained to code communication behaviors using an adapted version of the Specific Affect Coding System (SPAFF) [15, 29]. Both raters practiced coding for at least 60 hours on videotapes that were not part of the study, with Cohen's kappa indicating that they had achieved an acceptable interobserver agreement ($k = 0.9$). Each interaction was rated by both raters, with one rater focusing on the male partner and the other rater focusing on the female partner. Ratings were produced every 10 seconds to account for the behavior unfolding during each sequence, resulting in 48 sequences for each interaction. *Positive communication*: (1) careful listening, interest, curiosity, (2) recognition, approval, factual praise, (3) affective communication, caring, (4) constructive criticism, and *negative communication*: (1) blaming, criticism, (2) defensiveness, (3) domineering, (4) withdrawal, stonewalling, (5) formally negative interaction, (6) contempt, (7) provocation, beligerence. For each 10-second sequence, raters would thus assign the code representing the communication behavior that was most prevalent out of the ones listed above. Raters were asked to focus on the verbal aspect of the behavior in assigning the codes. Due to the vast variety of codes present, we categorized all types of positive and negative as 1 and 2 respectively, and then passed them to machine learning models in the form of a binary classification problem.

The speech was manually annotated with the start and end of each speaker's turn, along with pauses and noise. The speech was manually transcribed in 15-second chunks separately for each partner. Given that Swiss German is mostly spoken with different dialects across Switzerland, the spoken words were written as the corresponding German word equivalent. Transcripts and audio recordings acquired from the interactions were divided along the same 10-second sequence to match the 10-second sequences used for behavioral coding. This process was done separately for each partner's transcript and speech data. Consequently, we dropped 10-second matched transcript-audio-code sequences that contained no speech and transcribed words.

Of the original 368 Swiss heterosexual couples that took part in the study, we could only use 345 because some couples requested their data to be removed and some data were missing arising from technical problems in data collection. In addition, while the original dataset presented instances where behaviors had been coded as neutral/no communication, these were dropped from the analyses since no accurate description for what constituted neutral communication was given in the codebook, and no differentiation with instances of no communication was provided. This thus resulted

in a total of 9930 10-second speech sequences with their matching behavioral codes. Out of that number, 6978 were instances where communication had been coded as positive, while 2952 were instances where it had been coded as negative, highlighting a significant class imbalance that is characteristic of real-world datasets and partners' behavior as seen in other works (e.g., [14]).

Linguistic Features: We extracted linguistic features from each 10-second transcript sequence using the LIWC software for German [33]. It utilizes an existing list of words and categories (e.g., positive/negative words, personal pronouns, social process) to count the number of the corresponding words in the transcript sequences and categorize them across 97 different features. The internal German LIWC dictionary was used to analyze the transcript and extract the features. We normalized each transcript sequence's feature vector by dividing the value of all the other features by the "word count" feature which represents the number of words present in each transcript sequence. We then dropped the word count feature. This procedure thus left 96 normalized features that were passed as input to the machine learning models.

Also, we extracted features from each 10-second sequence using a pretrained Sentence-BERT (SBERT) model [37]. Sentence-BERT is a modification of the BERT architecture with siamese and triplet networks to compute sentence embeddings such that semantically similar sentences are close in vector space. Sentence-BERT has been shown to outperform the mean and CLS token outputs of regular BERT models for semantic similarity and sentiment classification tasks. Given that the text is in German, we used the German BERT model [1] as SBERT's Transformer model and the mean pooling setting. The German BERT model was pretrained using the German Wikipedia dump, the OpenLegalData dump, and German news articles. The extraction resulted in a 768-dimensional feature vector.

Paralinguistic Features: We extracted acoustic features from the voice recordings. For each 10-second sequence, we first used the speaker annotations to get the acoustic signal for each partner separately. Next, we used openSMILE [17] to extract the 88 eGeMAPS acoustic features which have been shown to be a minimalist set of features for affective recognition tasks [18]. The original audio was encoded with 2 channels. As a result, we extracted the features for each channel resulting in a 176-dimensional feature vector.

3 EXPERIMENTS AND EVALUATION

We performed multiple experiments using the support vector machine (SVM) algorithm with the radial basis function (RBF) kernel and the scikit-learn library [35]. We used RBF since it performed the best in as our initial explorations in comparison to random forests, XGBoost, and linear SVM. We trained models to perform binary classification of the behavioral codes for positive and negative communication using different feature sets. Specifically, we used features from LIWC, BERT, openSMILE. Also, we explored multimodal fusion at the features level of BERT and openSMILE by concatenating features from both groups. We used TF-IDF unigram and bigram features (using the most frequent 1000 features) of the transcripts as a linguistic baseline. To train and evaluate the models, we used nested K-fold cross-validation (CV). The nested procedure consisted of utilizing an "inner" run of 3-fold CV for hyperparameter tuning, followed by an "outer" run of 5-fold CV which utilizes

the best values for each hyperparameter found by the “inner” run. We prevented data from the same couple from being in both the train and test folds, thereby evaluating the model’s performance on data from unseen couples. As the data was imbalanced, we utilized the metric balanced accuracy which is the unweighted average of the recall of each class, and confusion matrices for evaluation. We used different values of the hyperparameter “C”, presenting results for the hyperparameter that produced the best results. We used the “balanced” hyperparameter for all the SVM models to mitigate the class imbalance while training. Standard errors were computed by running models repeatedly, randomizing the groups used for the K fold CV and thus gathering a set of 20 accuracy measures for each model.

4 RESULTS AND DISCUSSION

Table 1 presents the results of the best performing models for each of the feature modalities. The model that used only the BERT features performed the best with 69.4% accuracy, compared to the LIWC model with 65.4% accuracy. The BERT only approach also performed better than combining paralinguistic features to the BERT features. The latter performed closely with 69.2% accuracy, but still significantly worse than the BERT only model ($p < .001$ on a Wilcoxon signed rank test). Indeed, the paralinguistic baseline approach using openSMILE features performed the worst, 61.3%, being outperformed by TF-IDF linguistic baseline (65.6%) which also slightly outperformed LIWC. The worse performance of the paralinguistic features is expected given that raters in the study were instructed to focus on the verbal aspect of the interaction rather than nonverbal behavior in assigning codes.

Our results indicate that LIWC features do not have as much discriminative potential for prediction tasks compared both with even simpler approaches such as TF-IDF, as well as more advanced methods such as BERT. One likely explanation for BERT’s superior performance is its ability to capture the semantics of text via contextualized embeddings. These results have also been shown in a similar work using emotion psychotherapy or mental health data [24, 39]. The performance gain afforded by BERT notwithstanding, the simpler approaches did perform better than expected, with the performance improvement between TF-IDF and BERT being less than 4%. It is worth noting that this BERT model is used out-of-the-box and it is out-of-domain without any customization on the couples conversational text. The results indicate that researchers in social psychology ought to consider alternatives to LIWC, such as BERT, for extracting features for prediction tasks such as automated behavioral coding and emotion recognition. Although LIWC features (and indeed, TF-IDF) have the advantage of being simpler and more easily interpretable, various approaches are being developed to make BERT features more interpretable via its multi-head attention mechanism [45] and Shapley [28]. Finally, the result of the BERT model performing better than the multimodal approach is consistent with other works that found a similar result for emotion [9] and behavioral recognition [14]. Including paralinguistic features did not seem to add any more predictive information especially considering the context of the study in which assigning codes focused on verbal behavior. Further approaches need to be

Table 1: Results across 20 runs (with standard errors) for models using LIWC, BERT, openSMILE and multimodal input features

Input Features	Balanced Accuracy (% +/- S.E.)
openSMILE	61.28 ± .07
TF-IDF + ngrams	65.61 ± .08
LIWC	65.41 ± .05
BERT	69.39 ± .06
BERT and openSMILE	69.18 ± .06

explored to better combine the openSMILE and BERT features for improved results.

5 LIMITATIONS AND FUTURE WORK

In this work, we used manual transcripts. To accomplish true automated behavioral coding, our approach needs to use and work for automated transcriptions. Current speech recognition systems do not work for this unique dataset given that couples speak Swiss German, which is (1) a spoken dialect and not written, and (2) varies across different parts of the German-speaking regions of Switzerland. Further work is needed to develop automatic speech recognition systems for Swiss German.

Also, we only used the BERT model as a feature extractor to make a fair comparison with the LIWC features. Fine-tuning the BERT model on this task and domain to update the weights of the model would potentially improve the prediction results. This approach will be explored in future work. Finally, BERT models have been shown to encode gender and racial bias because of the data they are trained on. This consideration needs to be factored in for the specific prediction task and context [6].

6 CONCLUSION

In this work, we investigated the predictive potential of BERT features for automated coding of couples’ communication behavior compared to LIWC features, the de facto linguistic features in social psychology. We extracted and compared LIWC and BERT features, used openSMILE features as a paralinguistic baseline and TF-IDF with ngrams as a linguistic baseline. We trained an RBF SVM to classify positive and negative communication behavior of each romantic partner on a 10-second granularity. Our results showed that both simple TF-IDF features as well as more complex BERT features both outperform LIWC, indicating that it might be time for researchers to consider alternatives to LIWC for predictive tasks in couples interactions. Additionally, adding paralinguistic features did not perform better than the BERT-only approach. Our work is a further step towards better approaches in automating the coding of couples’ behavior which could enhance couples research and assessments in couples therapy.

ACKNOWLEDGMENTS

Funding was provided by the Swiss National Science Foundation: CR12I1_166348/1; CRSI11_133004/1; P3P3P1_174466; P300P1_164582

REFERENCES

- [1] [n.d.]. Open Sourcing German BERT. <https://deepset.ai/german-bert>. Accessed: 2020-05-1.
- [2] Paul R Amato. 2001. Children of divorce in the 1990s: an update of the Amato and Keith (1991) meta-analysis. *Journal of family psychology* 15, 3 (2001), 355.
- [3] Erin O'Carroll Bantum and Jason E Owen. 2009. Evaluating the validity of computerized content analysis programs for identification of emotional expression in cancer narratives. *Psychological assessment* 21, 1 (2009), 79.
- [4] Brian R Baucom, David C Atkins, Kathleen Eldridge, Pamela McFarland, Mia Sevier, and Andrew Christensen. 2011. The language of demand/withdraw: verbal and vocal expression in dyadic interactions. *Journal of Family Psychology* 25, 4 (2011), 570.
- [5] Brian R Baucom, Darby E Saxbe, Michelle C Ramos, Lauren A Spies, Esti Iturralde, Sarah Duman, and Gayla Margolin. 2012. Correlates and characteristics of adolescents' encoded emotional arousal during family conflict. *Emotion* 12, 6 (2012), 1281.
- [6] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [7] Matthew Black, Athanasios Katsamanis, Chi-Chun Lee, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Automatic classification of married couples' behavior using audio features. In *Eleventh annual conference of the international speech communication association*.
- [8] Matthew P Black, Athanasios Katsamanis, Brian R Baucom, Chi-Chun Lee, Adam C Lammert, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2013. Toward automating a human behavioral coding system for married couples' interactions using speech acoustic features. *Speech communication* 55, 1 (2013), 1–21.
- [9] George Boateng and Tobias Kowatsch. 2020. Speech Emotion Recognition among Elderly Individuals using Transfer Learning and Multimodal Fusion. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion)*, October 25–29, 2020, Virtual event, Netherlands.
- [10] Paul Boersma and Vincent Van Heuven. 2001. Speak and unSpeak with PRAAT. *Glott International* 5, 9/10 (2001), 341–347.
- [11] T.N. Bradbury and B.R. Karney. 2010. *Intimate Relationships*. W.W. Norton & Company. <https://books.google.ch/books?id=YMTeHAAACAAJ>
- [12] Sandeep Nallan Chakravarthula, Brian Baucom, and Panayiotis Georgiou. 2018. Modeling Interpersonal Influence of Verbal Behavior in Couples Therapy Dyadic Interactions. *arXiv preprint arXiv:1805.09436* (2018).
- [13] Sandeep Nallan Chakravarthula, Rahul Gupta, Brian Baucom, and Panayiotis Georgiou. 2015. A language-based generative model framework for behavioral analysis of couples' therapy. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2090–2094.
- [14] Sandeep Nallan Chakravarthula, Haoqi Li, Shao-Yen Tseng, Maija Reblin, and Panayiotis Georgiou. 2019. Predicting Behavior in Cancer-Afflicted Patient and Spouse Interactions Using Speech and Language. *Proc. Interspeech 2019* (2019), 3073–3077.
- [15] James A Coan and John M Gottman. 2007. The specific affect coding system (SPAFF). *Handbook of emotion elicitation and assessment* (2007), 267–285.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [17] Florian Eyben, Klaus R Scherer, Björn W Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y Devillers, Julien Epps, Petri Laukka, Shrikanth S Narayanan, et al. 2015. The Geneva minimalistic acoustic parameter set (GeMAPS) for voice research and affective computing. *IEEE transactions on affective computing* 7, 2 (2015), 190–202.
- [18] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*. 1459–1462.
- [19] John Gottman. 2018. *The Seven Principles for Making Marriage Work*. Hachette UK.
- [20] John Mordechai Gottman. 1994. *What predicts divorce?: The relationship between marital processes and marital outcomes*. Lawrence Erlbaum Associates, Inc.
- [21] John Mordechai Gottman. 2005. *The mathematics of marriage: Dynamic nonlinear models*. MIT Press.
- [22] John M Gottman, James Coan, Sybil Carrere, and Catherine Swanson. 1998. Predicting marital happiness and stability from newlywed interactions. *Journal of Marriage and the Family* (1998), 5–22.
- [23] Peter Hilpert, Timothy R Brick, Christoph Flückiger, Matthew J Vowels, Eva Ceulemans, Peter Kuppens, and Laura Sels. 2020. What can be learned from couple research: Examining emotional co-regulation processes in face-to-face interactions. *Journal of Counseling Psychology* 67, 4 (2020), 475.
- [24] Zheng Ping Jiang, Sarah Ita Levitan, Jonathan Zomick, and Julia Hirschberg. 2020. Detection of Mental Health from Reddit via Deep Contextualized Representations. In *Proceedings of the 11th International Workshop on Health Text Mining and Information Analysis*. 147–156.
- [25] Patrik N Juslin and Klaus R Scherer. 2005. *Vocal expression of affect*. Oxford University Press.
- [26] Athanasios Katsamanis, James Gibson, Matthew P Black, and Shrikanth S Narayanan. 2011. Multiple instance learning for classification of human behavior observations. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 145–154.
- [27] Patricia K Kerig and Donald H Baucom. 2004. *Couple observational coding systems*. Taylor & Francis.
- [28] Enja Kokalj, Blaž Škrli, Nada Lavrač, Senja Pollak, and Marko Robnik-Šikonja. 2021. BERT meets Shapley: Extending SHAP Explanations to Transformer-based Classifiers. In *Proceedings of the EACL Hackshop on News Media Content Analysis and Automated Report Generation*. 16–21.
- [29] Monika Kuster, Katharina Bernecker, Sabine Backes, Veronika Brandstätter, Fridtjof W Nussbeck, Thomas N Bradbury, Mike Martin, Dorothee Sutter-Stickel, and Guy Bodenmann. 2015. Avoidance orientation and the escalation of negative communication in intimate relationships. *Journal of Personality and Social Psychology* 109, 2 (2015), 262.
- [30] Chi-Chun Lee, Matthew Black, Athanasios Katsamanis, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Quantification of prosodic entrainment in affective spontaneous spoken interactions of married couples. In *Eleventh Annual Conference of the International Speech Communication Association*.
- [31] Chi-Chun Lee, Athanasios Katsamanis, Matthew P Black, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2014. Computing vocal entrainment: A signal-derived PCA-based quantification scheme with application to affect analysis in married couple interactions. *Computer Speech & Language* 28, 2 (2014), 518–539.
- [32] Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2016. Sparsely connected and jointly trained deep neural networks for low resource behavioral annotation: Acoustic classification in couples' therapy. *arXiv preprint arXiv:1606.04518* (2016).
- [33] Tabea Meier, Ryan L Boyd, James W Pennebaker, Matthias R Mehl, Mike Martin, Markus Wolf, and Andrea B Horn. 2019. "LIWC auf Deutsch": The Development, Psychometrics, and Introduction of DE-LIWC2015. *PsyArXiv a* (2019).
- [34] Mona Neysari, Guy Bodenmann, Matthias R Mehl, Katharina Bernecker, Fridtjof W Nussbeck, Sabine Backes, Martina Zemp, Mike Martin, and Andrea B Horn. 2016. Monitoring pronouns in conflicts. *GeroPsych* (2016).
- [35] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
- [36] James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates* 71, 2001 (2001), 2001.
- [37] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).
- [38] Rachel A Simmons, Peter C Gordon, and Dianne L Chambliss. 2005. Pronouns in marital interaction: What do "you" and "I" say about marital health? *Psychological science* 16, 12 (2005), 932–936.
- [39] Michael J Tanana, Christina S Soma, Patty B Kuo, Nicolas M Bertagnolli, Aaron Dembe, Brian T Pace, Vivek Srikumar, David C Atkins, and Zac E Imel. 2021. How do you feel? Using natural language processing to automatically rate emotion in psychotherapy. *Behavior Research Methods* (2021), 1–14.
- [40] Shao-Yen Tseng, Brian R Baucom, and Panayiotis G Georgiou. 2017. Approaching Human Performance in Behavior Estimation in Couples Therapy Using Deep Sentence Embeddings.. In *INTERSPEECH*. 3291–3295.
- [41] Shao-Yen Tseng, Sandeep Nallan Chakravarthula, Brian R Baucom, and Panayiotis G Georgiou. 2016. Couples Behavior Modeling and Annotation Using Low-Resource LSTM Language Models.. In *INTERSPEECH*. 898–902.
- [42] Shao-Yen Tseng, Haoqi Li, Brian Baucom, and Panayiotis G Georgiou. 2018. "Honey, I Learned to Talk" Multimodal Fusion for Behavior Analysis. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*. 239–243.
- [43] UZH University of Zurich. [n.d.]. PASEZ Project-Impact of stress on relationship development of couples and children. <http://www.dynage.uzh.ch/en/newsevents/news/news25.html>. Accessed: 2021-05-1.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762* (2017).
- [45] Jesse Vig. 2019. A multiscale visualization of attention in the transformer model. *arXiv preprint arXiv:1906.05714* (2019).
- [46] Wei Xia, James Gibson, Bo Xiao, Brian Baucom, and Panayiotis G Georgiou. 2015. A dynamic model for behavioral analysis of couple interactions using acoustic features. In *Sixteenth Annual Conference of the International Speech Communication Association*.