

“You made me feel this way”: Investigating Partners’ Influence in Predicting Emotions in Couples’ Conflict Interactions using Speech Data

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

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

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

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

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

ABSTRACT

How romantic partners interact with each other during a conflict influences how they feel at the end of the interaction and is predictive of whether the partners stay together in the long term. Hence understanding the emotions of each partner is important. Yet current approaches that are used include self-reports which are burdensome and hence limit the frequency of this data collection. Automatic emotion prediction could address this challenge. Insights from psychology research indicate that partners’ behaviors influence each other’s emotions in conflict interaction and hence, the behavior of both partners could be considered to better predict each partner’s emotion. However, it is yet to be investigated how doing so compares to only using each partner’s own behavior in terms of emotion prediction performance. In this work, we used BERT to extract linguistic features (i.e., what partners said) and openSMILE to extract paralinguistic features (i.e., how they said it) from a data set of 368 German-speaking Swiss couples (N = 736 individuals) who were videotaped during an 8-minutes conflict interaction in the laboratory. Based on those features, we trained machine learning models to predict if partners feel positive or negative after the conflict interaction. Our results show that including the behavior of the other partner improves the prediction performance. Furthermore, for men, considering how their female partners spoke is most important and for women considering what their male partner said is most important in getting better prediction performance. This work is a step towards automatically recognizing each partners’ emotion based on the behavior of both, which would enable a better understanding of couples in research, therapy, and the real world.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ICMI '21 Companion, October 18–22, 2021, Montréal, QC, Canada
© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8471-1/21/10...\$15.00
<https://doi.org/10.1145/3461615.3485424>

CCS CONCEPTS

• Applied computing → Psychology.

KEYWORDS

couples, emotion recognition, multimodal fusion, linguistic, paralinguistic, conflict

ACM Reference Format:

George Boateng, Peter Hilpert, Guy Bodenmann, Mona Neysari, and Tobias Kowatsch. 2021. “You made me feel this way”: Investigating Partners’ Influence in Predicting Emotions in Couples’ Conflict Interactions using Speech Data. In *Companion Publication of the 2021 International Conference on Multimodal Interaction (ICMI '21 Companion)*, October 18–22, 2021, Montréal, QC, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3461615.3485424>

1 INTRODUCTION

Understanding the emotions partners feel during and after conflict interactions is important because of its long-term effects on couples’ relationship quality and stability [13]. Happy couples, for example, experience more positive and less negative emotions during conflict interactions compared to unhappy couples [21]. The current study focuses on one fundamental aspect of the conflict mechanism — how the emotional experience *within* each partner is influenced by the behavioral exchange *between* partners.

A crucial aspect of conflict interaction in couples is how the behavioral exchange makes each person feel during and after the interaction [25]. But although both partners experience the same interaction, they can feel very differently about it. For example, if we assume that partner A shows contempt and criticizes partner B, we can assume that partner A might feel angry or superior whereas partner B might feel hurt or humiliated. Thus, the experience can be very different for the partner who communicates something compared to the partner who perceives it [7]. This differentiation allows us to reflect on another mechanism, namely that the emotions a person experiences are the results of two kinds of influences. Obviously, a person’s emotional experience is constantly influenced by a partner’s behavior as a kind of co-regulating force, talk turn by talk turn. In addition, however, each person has the ability to

regulate one's own emotional response (e.g., cognitive appraisal, emotion regulation) [16], which then affects one's own subsequent behavioral response. Thus, what partners experience emotionally during and at the end of a conflict interaction is a reflection of the co-regulation and self-regulation processes [6].

To better understand emotions in couples and their impact on relationships, often self-report assessments are used in which each partner is asked to provide a rating of their own emotions right after an interaction, or partners are asked to watch a video recording of the interaction and provide continuous ratings using a joystick, for example, [17, 25]. Self-reports are burdensome to complete and may not be collected frequently. This means that the relationship between behavior and emotions cannot be studied often. Thus, an automatic emotion recognition system would allow scaling of couples research.

Various works have used linguistic features (i.e., *what* has been said) and paralinguistic features (i.e., *how* it was said) to predict the emotions of each partner in couples interactions more broadly [3, 4, 8, 9, 19, 20, 22, 28–30, 32] and in conflict interactions in particular [5, 10]. Most of these works have used observer ratings (perceived emotions) as labels rather than self-reports (one's actual emotions). Hence, the prediction task becomes that of recognizing external individuals' perception of each partner's emotion rather than each partner's emotion per their own assessment. Though similar, the latter is more challenging than the former for a number of reasons. First, the rating might be biased and may not reflect their actual emotions over the period the rating is for (e.g., the past 5 minutes). Whereas for observer ratings, coders are generally trained over several weeks, it is done by more than one person and various approaches are employed to resolve ratings that are not in agreement and ensure the validity of the labels. Second, the self-reported emotion may not be reflected in that partner's behavior in comparison to observer ratings which are purely based on behavioral observation.

Despite these challenges, insights from psychology research could be leveraged to make the prediction task easier. Specifically, given that partners' behaviors influence each other's emotions in conflict interaction, the behavior of both partners could be considered to better predict each partner's end-of-conversation emotion. However, it is yet to be investigated how doing so compares to using each partner's own behavior only in terms of emotion prediction performance. In this work, we used a dataset collected from 368 couples who were recorded during an 8-minute conflict interaction, extracted linguistic and paralinguistic features, and used machine learning approaches to predict how each partner felt directly after the conflict interaction (self-reported emotion). We answer the following research questions (RQs)

RQ1: *How well can the end-of-conversation emotion of each partner be predicted by their own behavior — a combination of linguistic and paralinguistic data? (self-regulation)*

RQ2: *How does the prediction performance change when including the other partner's behavior — (a) linguistic only, (b) paralinguistic only, and (c) combination of linguistic and paralinguistic data? (co-regulation)*

Our contributions are (1) an evaluation of how well a partner's own linguistic and paralinguistic features predict one's own end-of-conversation emotion (2) an investigation of how the prediction

performance changes when including one's partner's features (linguistic, paralinguistic, and both) (3) the use of a unique dataset — spontaneous, real-life, speech data collected from German-speaking, Swiss couples (n=368 couples, N=736 participants), which is the largest ever such dataset used in the literature for automatic recognition of partners' end-of-conversation emotion. The insights from our work would advance the use of methods to automatically recognize the emotions of each partner which could enable research and applications to better understand couples' relationships in therapy and the real world.

The rest of the paper is organized as follows: In Section 2, we describe our data collection, preprocessing and feature extraction, in Section 3, we describe our experiments and evaluation, in Section 4 we present and discuss our results, in Section 5, we present limitations and future work, and we conclude in Section 6.

2 METHODOLOGY

2.1 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) [18, 31]. 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.

After each conversation, each partner provided self-reported responses to the Multidimensional Mood questionnaire [27] of their emotions on four bipolar dimensions — namely “good mood versus bad mood,” “relaxed versus angry,” “happy versus sad” and “calm versus stressed” — with the scale: 1 — very much, 2 — much, 3 — a little, 4 — a little, 5 — much, 6 — very much. In this work, we sort to focus on predicting emotional valence (positive or negative) based on Russell's circumplex model of emotions [26]. Hence, we used an average of the “good mood versus bad mood” and “happy versus sad” scales which enables us to get a more valid score since several dimensions that measure similar constructs are combined. We did not use the other two scales because their polarity also could represent the arousal dimension of emotion (low vs high arousal). We then binarized the averaged values similar to prior works (e.g., [3, 5]) such that values greater than or equal to 3.5 were negative (0) and the rest were positive (1). Binarization enables us to map the data into Russell's circumplex model of emotions which has 4 quadrants for emotions, further enabling its real-world utility — easily being able to tell which group of emotions are being felt by each partner using binarized valence and arousal dimensions.

The speech data were manually annotated with the start and end of each speaker's turn, along with pauses and noise. This was necessary in order to later be able to extract linguistic and paralinguistic features for each partner separately. In addition, speech content of both partners was manually transcribed for each partner separately and stored in 15-second chunks. Given that Swiss German is mostly spoken with different dialects across Switzerland, the spoken words were written as the corresponding German word equivalent.

Some couples requested their data to be removed and some data were missing due to technical problems in data collection. Of the original 368 couples that took part in the study, we could use 338 samples for females (46 negative labels) and 341 samples for males (32 negative labels). The distribution highlights a significant class imbalance that is characteristic of real-world datasets and partners' emotions as seen in other similar works (e.g., [5]).

2.2 Linguistic Features

We extracted linguistic features from the transcripts of the whole 8-minute interaction using a pre-trained model – Sentence-BERT (SBERT) [24]. 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 pre-trained using the German Wikipedia dump, the OpenLegalData dump, and German news articles. The extraction resulted in a 768-dimensional feature vector.

2.3 Paralinguistic Features

We extracted acoustic features from the voice recordings. First, we used the speaker annotations to get the acoustic signal for each gender separately. Next, we used openSMILE [11] to extract the 88 eGeMAPS acoustic features which have been shown to be a minimalist set of features for affective recognition tasks [12]. The features are extracted in 25 ms sequences and then various functions (e.g., mean, median, range, etc.) are computed over the sequences resulting in 88 features for the whole 8-minute audio. 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.

2.4 Multimodal and Dyadic Feature Fusion

Given that emotions are reflected in what and how people say things, we performed multimodal fusion (early fusion) by concatenating the linguistic and paralinguistic features resulting in a 944-dimensional feature. We did this separately for each partner. This feature vector was used as the baseline approach to answer research question (1).

Additionally, we fused features from both partners to answer research question (2). Specifically, for partner A, we concatenated their multimodal feature vector with the features of partner B and used it to predict partner A's emotion label. This process was done for partner B as well. In order to investigate which behavioral data of the interacting partner was most important in the prediction of the emotions, we included the features in the following order (1) linguistic only, (2) paralinguistic only, and (3) multimodal fusion of both.

Consequently, we had four feature sets: (1) Multimodal fusion (baseline – own features), (2) Multimodal + Dyadic Fusion (with partner's linguistic only), (3) Multimodal + Dyadic Fusion (with partner's paralinguistic features only) (4) Multimodal + Dyadic

Table 1: Prediction Results of the Best Models for the Fusion Approaches

Approach	Balanced Accuracy (%)	
	Male	Female
Multimodal fusion (baseline)	49.8	59
Multimodal + Dyadic Fusion (with partner's combined linguistic and paralinguistic only)	52.3	63.2
Multimodal + Dyadic Fusion (with partner's linguistic features only)	53.5	64.8
Multimodal + Dyadic Fusion (with partner's paralinguistic only)	56.1	59.9

Fusion (with partner's combined linguistic and paralinguistic only). These were passed to machine learning models to answer the two research questions.

3 EXPERIMENTS AND EVALUATION

We run experiments using scikit-learn [23] the following machine learning models: support vector machine (SVM) algorithm with linear and radial basis function kernel, and random forests. We trained models to perform binary classification of each partner's self-reported positive and negative emotion using different feature sets. We used the four feature sets described in the previous section. To train and evaluate the models, we used nested K-fold cross-validation (CV). The nested procedure consisted of utilising an "inner" run of 5-fold CV for hyperparameter tuning, followed by an "outer" run of 10-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 used the metric balanced accuracy which is the unweighted average of the recall of each class and confusion matrices for evaluation. We used the "balanced" hyperparameter for all the models to mitigate the class imbalance while training. We compare to a random baseline of 50% balanced accuracy.

4 RESULTS AND DISCUSSION

Our results are shown in Table 1. The baseline approach with multimodal fusion was not better than chance in predicting men's emotions at the end of the conflict interaction (49.8%). This is unexpected as this indicates that men's own behaviors during the interaction are not related to how they feel at the end of the interaction. This might be due to self-regulation processes or not showing much emotions during the interaction. This is clearly different from women as their emotions can be predicted by their own behavior (59%). Thus, it seems that women seem to express their emotions more clearly in their behavior. These results of poorer prediction performance for men compared to women is consistent with the results of [5].

Including features from the interacting partner improved the results for both men (52.3%) and women (63.2%). These results are consistent with psychology research that the behavior of partners'




Figure 1: Confusion matrix for best performing model for male partners (Multimodal + Dyadic Fusion with partner’s linguistic features only)

have an effect on each other’s emotions in conflict interaction [7, 14]. This is a crucial finding because (i) previous research shows that the behavior of one person influences the behavior of the other person [13] (ii) as well as that the emotional changes of one person affect the emotions of the other [7]. However, this is the first study showing that behavioral features assessed during the conflict interaction can be used to predict one partner’s emotion at the end of the conversation. In addition, the improvement in women’s emotion prediction at the end of interaction is greater when including their partner’s linguistic data (64.8%) whereas there is hardly any difference when including partner’s paralinguistic features. This is a surprising finding as women generally pay more attention to paralinguistic cues [15]. Notably, the results are different for men. The prediction for men’s emotions slightly increases when including their partner’s linguistic features but the prediction improves substantially when including women’s paralinguistic features. Although we do not know which specific paralinguistic features are the main drivers for predicting the emotions, this finding is in line with prior findings – when women “nag”, men experience strong negative physiologic reactions and tend to withdraw [13, 15]. Future research is needed to investigate which aspects of one partner’s behavior exactly cause their own emotion prediction performance to decline. In addition, these results have implications for the kind of behavioral information to consider to best predict each partner’s end-of-conversation emotions.

We show the confusion matrices for the best models for the men and women in Figures 1 and 2 respectively. They reveal the models’ challenges at recognizing positive emotions (1), resulting in them being misclassified as negative emotions (0).

5 LIMITATIONS AND FUTURE WORK

Further work is needed to investigate if these results are the same for couples in a different cultural context and also explore the effect of the interacting partners’ behavior at a more granular level such as on a talk-turn basis. More fine-grained emotion ratings may be needed to investigate that. Other fusion approaches like late fusion can be explored.

We used BERT as a feature extractor in this work. Generating domain-specific sentence embeddings via fine-tuning the BERT




Figure 2: Confusion matrix for best performing model for female partners (Multimodal + Dyadic Fusion with partner’s paralinguistic features only)

model and exploring deep transfer learning models for the paralinguistic features may improve the results. BERT models have been shown to encode gender and racial bias because of the models they are trained on. Further investigations are needed in the future of potential biases in prediction [2].

Finally, we used manual annotations and transcripts. To accomplish true automatic emotion prediction, speaker annotations need to be done automatically and 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.

6 CONCLUSION

In this work, we investigated one’s own and the partner’s behavior in predicting end-of-conversation emotions in the context of conflict interactions in German-speaking Swiss couples. We extracted linguistic features using BERT and paralinguistic features using openSMILE. We fused both features in a multimodal approach for each partner. We also fused the features of both partners to predict the emotions of each partner. Our results show that including the behavior of the other partner improves the prediction performance. Furthermore, for men, considering *how* their female partners spoke is most important, and for women considering *what* their male partners said is most important in getting better prediction performance. These insights have implications for the behavioral information to (not) include to better predict each partner’s end-of-conversation emotions which will enable a better understanding of couples relations in research, therapy, and the real world.

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] 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.
- [3] 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*.
- [4] 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.
- [5] George Boateng, Laura Sels, Peter Kuppens, Peter Hilpert, and Tobias Kowatsch. 2020. Speech Emotion Recognition among Couples using the Peak-End Rule and Transfer Learning. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion), October 25–29, 2020, Virtual event, Netherlands*.
- [6] Steven M Boker and Jean-Philippe Laurenceau. 2006. Dynamical systems modeling: An application to the regulation of intimacy and disclosure in marriage. *Models for intensive longitudinal data* 63 (2006), 195–218.
- [7] Emily A Butler. 2011. Temporal interpersonal emotion systems: The "TIES" that form relationships. *Personality and Social Psychology Review* 15, 4 (2011), 367–393.
- [8] 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).
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] John Mordechai Gottman. 1994. *What predicts divorce?: The relationship between marital processes and marital outcomes*. Lawrence Erlbaum Associates, Inc.
- [14] John Mordechai Gottman. 2014. *What predicts divorce?: The relationship between marital processes and marital outcomes*. Psychology Press.
- [15] John M Gottman and Robert W Levenson. 1992. Marital processes predictive of later dissolution: behavior, physiology, and health. *Journal of personality and social psychology* 63, 2 (1992), 221.
- [16] James J Gross. 2014. Emotion regulation: Conceptual and empirical foundations. (2014).
- [17] 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.
- [18] 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.
- [19] 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*.
- [20] 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.
- [21] Robert W Levenson, Laura L Carstensen, and John M Gottman. 1994. Influence of age and gender on affect, physiology, and their interrelations: A study of long-term marriages. *Journal of personality and social psychology* 67, 1 (1994), 56.
- [22] Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2016. Sparsely connected and disjointly trained deep neural networks for low resource behavioral annotation: Acoustic classification in couples' therapy. *arXiv preprint arXiv:1606.04518* (2016).
- [23] 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.
- [24] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).
- [25] Anna Marie Rief and Robert W Levenson. 2007. Continuous measurement of emotion. *Handbook of emotion elicitation and assessment* (2007), 286–297.
- [26] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161.
- [27] Rolf Steyer, Peter Schwenkmezger, Peter Notz, and Michael Eid. 1997. Der Mehrdimensionale Befindlichkeitsfragebogen MDBF [Multidimensional mood questionnaire]. *Göttingen, Germany: Hogrefe* (1997).
- [28] 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.
- [29] 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.
- [30] Shao-Yen Tseng, Haoqi Li, Brian Baucom, and Panayiotis 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.
- [31] 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.
- [32] 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*.