



AMII: Adaptive Multimodal Inter-personal and Intra-personal Model for Adapted Behavior Synthesis

Jieyeon Woo, Mireille Fares, Catherine Pelachaud, Catherine Achard

► To cite this version:

Jieyeon Woo, Mireille Fares, Catherine Pelachaud, Catherine Achard. AMII: Adaptive Multimodal Inter-personal and Intra-personal Model for Adapted Behavior Synthesis. 2023. hal-04293258

HAL Id: hal-04293258

<https://hal.science/hal-04293258>

Preprint submitted on 21 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

AMII: ADAPTIVE MULTIMODAL INTER-PERSONAL AND INTRA-PERSONAL MODEL FOR ADAPTED BEHAVIOR SYNTHESIS

Jieyeon Woo
ISIR - Sorbonne University
France
woo@isir.upmc.fr

Mireille Fares
ISIR - Sorbonne University
France
fares@isir.upmc.fr

Catherine Pelachaud
CNRS - ISIR - Sorbonne University
France
pelachaud@isir.upmc.fr

Catherine Achard
ISIR - Sorbonne University
France
achard@isir.upmc.fr

May 22, 2023

ABSTRACT

Socially Interactive Agents (*SIAs*) are physical or virtual embodied agents that display similar behavior as human multimodal behavior. Modeling *SIAs*' non-verbal behavior, such as *speech* and *facial gestures*, has always been a challenging task, given that a *SIA* can take the role of a *speaker* or a *listener*. A *SIA* must emit appropriate behavior adapted to its own *speech*, its previous behaviors (*intra-personal*), and the *User*'s behaviors (*inter-personal*) for both roles. We propose *AMII*, a novel approach to synthesize *adaptive facial gestures* for *SIAs* while interacting with *Users* and acting interchangeably as a *speaker* or as a *listener*. *AMII* is characterized by *modality memory encoding schema* - where *modality* corresponds to either *speech* or *facial gestures* - and makes use of *attention mechanisms* to capture the *intra-personal* and *inter-personal* relationships. We validate our approach by conducting objective evaluations and comparing it with the state-of-the-art approaches.

Keywords Adapted Behavior Synthesis · Reciprocal Adaptation · Conversational AI · Multimodal modeling

1 Introduction

Socially Interactive Agents (SIAs) are embodied agents emitting human-like non-verbal behavior such as *facial*, *body*, and *hand* gesturing during *speech* [1]. In many applications such as medical therapy [2] and educational assistance [3], they can carry out interactive and natural conversations with humans (*Users*) by sending and receiving multimodal non-verbal signals [4, 5] in addition to the verbal message. To render *SIA-User* communication effective and engaging, *SIAs* must play active roles of *speakers* and *listeners*. For instance, in successful and engaging human-human conversations, *speaker*'s and *listener*'s behaviors are constantly coordinated and adapted to each other [6], through the perception and response of the multimodal emitted signals such as linguistic [7], facial [5], and body [8] features. The behavior adaptation can be: (1) *intra-personal* which is between the multimodal signals of a single person or (2) *inter-personal* which is between *listener*'s and *speaker*'s signals adapting to the other's verbal and non-verbal behavior. This adaptation is *continuous*, *reciprocal* and *dynamical* and is referred to as *reciprocal adaptation* [9]. The display of such adaptation capability can enable *SIAs* to be perceived as *social* and *engaging* [10].

Various works have generated *SIAs*' non-verbal behavior focusing on only the modeling of *intra-personal* relationship for a single person [11, 12, 13]. Other works looked into capturing multimodal *inter-personal* information in dyadic settings. In the work of Feng et al. [14] and Dermouche et al. [15], *SIA facial gestures* are synthesized based on past gestures of both *SIA* and *User* without considering the existing relation with *audio modality* [16, 17], and do not ensure

the *motion continuity*. Grafsgaard et al. [18] synthesize interlocutor’s *gestures* based on both the interlocutors’ *audio* and their *facial modalities*. In the work of Jonell et al. [19], *SIA*’s *facial gestures* are generated based on *SIA*’s *speech* and the *User*’s *speech* and *facial gestures*. However, these models [18, 19] are prone to produce *non-continuous gestures*. The works presented in Woo et al. [20, 21] and Ng et al. [22] ensure *SIA*’s *behavior continuity* by employing autoregressive online inference while modeling *SIA*’s and *User*’s multimodal features. Only the *listener*’s behavior is modeled in [22]. [20, 21] and all the aforementioned models for dyadic setting do not explicitly model the *intra-personal* relationship.

Our overall aim is to create a *social* and *engaging SIA* by modeling its *behavior adaptation*. Our hypothesis is that *behavior adaptation* is captured by *intra-personal* and *inter-personal* synchronizations. We propose **Adaptive Multimodal Inter-personal and Intra-personal (AMII)** model, a novel method to synthesize *adaptive* and *engaging facial* gesturing for *SIAs*. We explicitly model the *intra-personal* relationship by encoding the prior emitted multimodal signals (*modality memory*) while ensuring *motion continuity*. We model the *inter-personal* relationship to generate *SIA* behavior for both roles of *speaker* and *listener*. We explore the best way to capture the *reciprocal adaptation* to generate *adaptive* non-verbal *SIA* behavior within a dyadic setting. *Intra-personal* and *inter-personal* relations are learned through *attention mechanisms*. We demonstrate that our *AMII* model outperforms the state-of-the-art models, in terms of both *behavior appropriateness* and *reciprocal adaptation resemblance*, by conducting objective evaluations. We perform ablation studies and show the influence of *AMII*’s key encoders.

Our contributions are as follows:

- We propose *AMII*, an approach to capture the *reciprocal adaptation* of *SIA* behaviors.
- We explicitly model the *intra-personal* and *inter-personal* relationships by encoding prior emitted multimodal signals via our *modality memory* encoders.

The paper is organized as follows. Section 2 describe the proposed *AMII* model architecture. The dataset is presented in Section 3 and we detail the training regime in Section 4. Then, we report our experiments in Section 5 and discuss our results in Section 6. We finally conclude and talk about our future works in Section 7.

2 Adaptive Multimodal Inter-personal and Intra-personal (AMII) Model Architecture

We focus on generating *adaptive* non-verbal *SIA* behavior as both a *speaker* and *listener*. For this purpose, we propose an approach trained on real human-human interactions, to learn human-human *inter-personal* and *intra-personal* relations for *SIA* and simulate our predictions on a *SIA*. We propose a new model architecture, the Adaptive Multimodal Inter-personal and Intra-personal model (*AMII*), to synthesize *adaptive* and *engaging facial gestures* for *SIAs*. It takes as input *speech* and *facial gestures* of both *SIA* (*A*) and *User* (*U*), corresponding to their past behavior, and predicts *SIA*’s and *User*’s *facial gestures* at the next time step.

Facial gestures are represented by:

- **Gaze movements:** represented by G_x and G_y which are the gaze angles w.r.t. the x and y axis.
- **Head movements:** represented by R_x , R_y , and R_z which are the Euler head rotations w.r.t. the x , y and z axis.
- **Facial movements:** represented by facial Action Units (AUs) [23] which are facial muscle movements defined by the Facial Action Coding Systems (FACS) [24]. We use *AU1* (inner brow raiser), *AU2* (outer brow raiser), *AU4* (brow lowerer), *AU6* (cheek raiser), and *AU12* (lip corner puller for smile).

Speech and facial gestures are highly tied together [16, 17]. With this relation, we decide to use *speech* information to drive *SIA*’s *facial gestures*. *Speech features* are listed as follows:

- **Pitch:** represented by the fundamental frequency f_0 , which is the main prosodic feature correlated with *facial gestures* [25].
- **Loudness:** speech intensity from the auditory spectra.
- **Voicing probability:** speech presence probability expressed as a probability score in the range of 0 to 1.
- **Mel-Frequency-Cepstral Coefficients (MFCC):** represented by 13 *MFCC* features (0-12).

AMII model operates as follows. It takes as input the 100 past frames ($t - 99 : t$), where t is the current frame, of the:

1. *Speech features* of *A* (X_{speech}^A) and those of *U* (X_{speech}^U),
2. *Facial features* of *A* (X_{face}^A) and those of *U* (X_{face}^U).

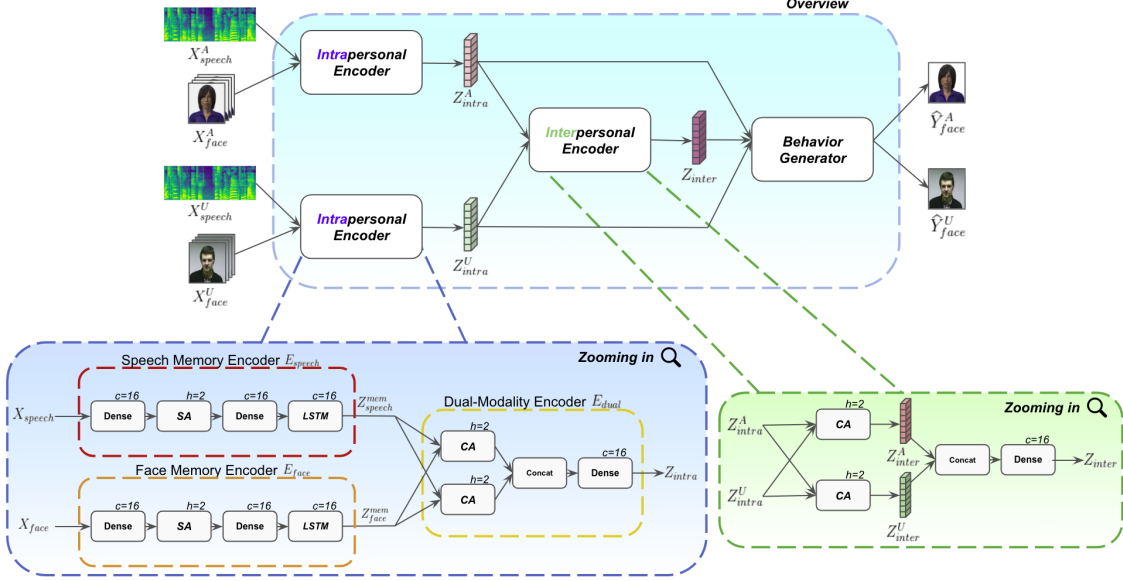


Figure 1: **AMII** (Adaptive Multimodal Intra-personal and Inter-personal) model architecture. The intra-personal encoder (E_{intra}) takes the *speech* X_{speech} and the *facial gestures* X_{face} of the previous 100 frames of either the *SIA* (A) or the *User* (U) to encode the corresponding *intra-personal* relationship Z_{intra} . The inter-personal encoder (E_{inter}) learns from *intra-personal* relationships Z_{intra}^A and Z_{intra}^U to encode the *inter-personal* relationship between them Z_{inter} . The behavior generator (G_{face}) takes Z_{intra}^A , Z_{intra}^U , and Z_{inter} to generate the sequence of *facial gestures* for the next frame $t+1$ \hat{Y}_{face}^A and \hat{Y}_{face}^U . At *training time*, **AMII** is trained with human-human (U_1 - U_2) interactions (U_1 for A and U_2 for U) and predicts both of humans' facial gestures ($\hat{Y}_{face}^{U_1}$ and $\hat{Y}_{face}^{U_2}$). At *inference time*, **AMII** renders the facial gestures of A and U . To infer the next A 's behavior, we feed back the predicted A 's behavior and the ground truth of U .

For each prediction of the next frame ($t + 1$), the model predicts:

1. A 's facial gestures (\hat{Y}_{face}^A),
2. U 's facial gestures (\hat{Y}_{face}^U).

AMII consists of three main components, as illustrated in Figure 1. The first component is the intra-personal encoder E_{intra} , which explicitly encodes the *intra-personal* relations via a *modality memory schema*. This schema consists of encoding each *modality* - *speech features* and *facial features* - corresponding to the past 100 frames. The second component is the inter-personal encoder E_{inter} , which encodes the *inter-personal* relations by applying cross-attentions between A 's and U 's features' embeddings. The last component is the *behavior generator* which generates A 's and U 's *facial gestures* of the next frame. These components are detailed in the following.

2.1 Intra-personal Encoder (E_{intra})

As shown in Figure 1, E_{intra} takes as input X_{speech} and X_{face} of either A or U and generates the corresponding intra-personal embedding Z_{intra} . It consists of two sub-encoders. The first is the modality memory encoder (E_{speech} or E_{face}). The second is the dual-modality encoder E_{dual} .

Modality Memory Encoder (E_{speech} or E_{face}) Both E_{speech} and E_{face} takes its corresponding *modality* - X_{speech} or X_{face} respectively - as input and renders the modality memory embedding Z_{speech}^{mem} and Z_{face}^{mem} representing the past 100 frames. Each corresponding modality memory encoder firstly learns the *modality specific* information by applying *self-attention*, with a $h = 2$ (where h is the head size), preceded and followed by dense layers (E_d) with $c = 16$ (where c is the cell size). Then, it embeds the memory sequence of the chosen *modality* via a LSTM layer (E_m) with $c = 16$, as depicted in Figure 1. It takes X_{mod} - where *mod* represents either *speech* or *facial gestures* - and outputs Z_{mod}^{mem} , and can be expressed as:

$$Z_{mod}^{mem} = E_m(E_d(SA(E_d(X_{mod})))) \quad (1)$$

where $SA(\cdot)$ denotes self-attention layer.

Dual-modality Encoder (E_{dual}) E_{dual} captures the relationship between the multimodal signals by applying *cross-attention mechanisms* on the corresponding modalities: CA^{speech} and CA^{face} with $h = 2$ followed by E_d with $c = 16$, as shown in Figure 1. CA^{speech} has a query Q equals to Z_{speech}^{mem} with key K and value V equal to Z_{face}^{mem} . CA^{face} has a query Q equals to Z_{face}^{mem} with key K and value V equal to Z_{speech}^{mem} . E_{dual} takes Z_{speech}^{mem} and Z_{face}^{mem} as inputs and generates Z_{intra} . It can be written as:

$$Z_{intra} = E_d([CA^{speech}(Q_{speech}, K_{face}, V_{face}), CA^{face}(Q_{face}, K_{speech}, V_{speech})]) \quad (2)$$

where $CA(Q, K, V)$ denotes cross-attention layer and $[\cdot]$ denotes concatenation layer.

2.2 Inter-personal Encoder (E_{inter})

As illustrated in Figure 1, E_{inter} takes as input Z_{intra}^A and Z_{intra}^U , which are the *intra-personal* representations of A and U respectively. It renders Z_{inter} , a representation of *inter-personal* relation between A and U . E_{inter} applies *cross-attention mechanisms* on the both *intra-personal* representations: CA^A and CA^U with $h = 2$ followed by E_d with $c = 16$. CA^A has a query Q equals to Z_{intra}^A with key K and value V equal to Z_{intra}^U . CA^U has a query Q equals to Z_{intra}^U with key K and value V equal to Z_{intra}^A . It can be written as:

$$Z_{inter} = E_d([CA^A(Q^A, K^U, V^U), CA^U(Q^U, K^A, V^A)]) \quad (3)$$

where $CA(Q, K, V)$ denotes cross-attention layer and $[\cdot]$ denotes concatenation layer.

2.3 Behavior Generator (G_{face})

G_{face} takes as input the:

1. A 's *intra-personal* representation (Z_{intra}^A),
2. U 's *intra-personal* representation (Z_{intra}^U),
3. *Inter-personal* representation of A and U (Z_{inter}).

It generates the corresponding *facial gestures* \hat{Y}_{face}^A and \hat{Y}_{face}^U by decoding with a dense layer (D_d) with $c = 20$, as depicted in Figure 1. The final outputs \hat{Y}_{face}^P can be written as:

$$\hat{Y}_{face}^P = D_d(Z_{intra}^P, Z_{inter}) \quad (4)$$

where P represents either A or U . At training time, $AMII$ synthesizes \hat{Y}_{face}^U to better learn *inter-personal* and *intra-personal* relations. \hat{Y}_{face}^U is disregarded at inference time since the aim is to predict only A .

2.4 Training and Inference Modes

We train our model on real human-human (U_1 - U_2) interactions. Our model learns to synthesize adapted gestures of U_1 and U_2 . During inference, $AMII$ synthesizes the behavior of A and U . \hat{Y}_{face}^A is inferred using the previous prediction of A and the ground truth of U . We apply *adaptive online prediction* to generate *continuous* A 's behavior in an *autoregressive* fashion.

3 NoXi Dataset

For our experiments, we choose the French *NoXi* dataset [26] consisting of 21 human-human dyadic interactions performed by 28 participants with a total duration of 7h22min. We extract non-verbal behavior features for each time-step: openSMILE [27] for *speech* and OpenFace [28] for *facial gestures*. Data preprocessing - median filter and linear interpolation - is applied on both extracted features and they are adjusted to 25fps. We split our dataset into 3 sets: training (70%), validation (10%), and test (20%). The test set does not include data of *speakers* and *listeners* that are seen during training. The aim is to test $AMII$'s capacity to extrapolate on new unseen *speakers* and *listeners* and therefore its capability to generalize.

4 Training Regime

To train our model, we use the Mean Squared Error (MSE) as our loss function and the Adam optimizer [29] with Cyclical Learning Rate (CLR) [30] (*triangular* learning rate policy, *base_lr* of $1e-7$, *max_lr* of $1e-3$, and step size factor of 10). The training was done for 300 epochs (with an average runtime of 125h) on a 2.2GHz Intel Xeon Linux server with NVIDIA GeForce GTX TITAN X and 64GB RAM with a batch size of 32. The best set of hyperparameters is chosen after manual optimization based on the validation set.

5 Experiments

To assess our model, we conduct objective evaluation to check its performance against the state-of-the-art approaches, which we select as our baselines, and to verify the effectiveness of each *AMII*'s key components through ablation studies.

5.1 Objective Metrics

We want to assess whether the generated behavior is *appropriate* (*intra-personal*) and *reciprocally adaptive* (*inter-personal*). To do so, we employ the metrics used in previous works [14, 15, 18, 20, 21, 22]. We measure *behavior appropriateness* of A 's predictions (\hat{A}) against its ground truth (GT) behavior (A). The metrics are as follows:

- **MAE and RMSE/L2:** distance between the predictions and GT to measure the generated error, which are used in [14, 15, 20, 21, 22].
- **Kolmogorov-Smirnov two-sample test (KS test)** [31]: density probability difference between A 's corresponding predictions and GT to check for the distribution similarities between them, which is used in [21].

For *reciprocal adaptation resemblance*, we measure the resemblance between \hat{A} and U 's GT data (U). We choose to assess only smile (*AU12*) as it is a key socio-emotional signal [8]. The metrics are as follows:

- **Time lagged cross correlation coefficient (TLCC)** [32]: linear relationship invariant to speed, which is used to quantify *global synchrony* like PCC. TLCC is computed in chunks of 8sec with a time lag of 2sec as in [22].
- **DTW** [33]: proximity/resemblance check of $\hat{A}\&U$ against that of the GT interaction (between two humans) to evaluate *reciprocal adaptation*. DTW is computed in chunks of 1min with a stride of 30sec as in [21].
- **Synchrony (Sync) and Entrainment Loop (EL)** [9]: synchrony and entrainment loop measures proposed in [9] to evaluate the *reciprocal adaptation* between \hat{A} and U .

Lower values denote better performance for MAE, RMSE, and KS test. For the *resemblance metrics* (PCC, TLCC, DTW, Sync, and EL), the closer the value of the metric is to the GT, the better the model performs in generating *adaptive* A 's behaviors.

5.2 Baseline Models and Ablation Studies

We consider the following *baseline models* (*base*):

- **IL-LSTM** [15]: models only the *inter-personal* based on a single modality of *facial gestures* of A and U ,
- **Symmetrized IL-LSTM with online LSTM (sym-IL-LSTM)** [20] and **ASAP** [21]: model the *inter-personal* relation of A and U , *multimodality* of *speech* and *facial gestures*, and assure *motion continuity*.

To check for the effectiveness and influence of each of *AMII*'s key encoders, we conducted additional ablation studies. We performed the ablations of:

- Modality Memory Encoder E_m (*no* E_m),
- Dual-modality Encoder E_{dual} (*no* E_{dual}),
- Inter-personal Encoder E_{inter} (*no* E_{inter}).

Table 1: Objective evaluation of AMII against the baselines along with ablations using the selected metrics. GT denotes ground truth interaction. Best results are highlighted in **bold**. Δ_{base} represents the change in performance over the best performing baseline approach of each metric. Δ_{base} entries in **green** when AMII outperforms best baseline, in **red** when it is **not** the case.

	MAE	RMSE	KS test	TLCC	DTW	Sync	EL
GT				0.334	1317.5	132.4	1172.3
IL-LSTM [15]	0.304	0.415	0.329	0.343	1216.2	45.3	323.3
sym-IL-LSTM [20]	0.180	0.227	0.284	0.335	1281.3	33.3	232.3
ASAP [21]	0.185	0.254	0.282	0.317	1399.3	142.0	1890.5
AMII- noE_m (ours)	0.099	0.132	0.515	0.271	1228.4	79.1	603.8
AMII- noE_{dual} (ours)	0.143	0.186	0.396	0.300	1352.6	262.3	2255.8
AMII- noE_{inter} (ours)	0.136	0.178	0.406	0.261	1127.8	82.3	586.3
AMII (ours)	0.156	0.197	0.437	0.291	1319.6	137.4	989.0
Δ_{base}	↓ 0.024	↓ 0.030	↑ 0.155	↑ 0.042	↓ 79.7	↓ 4.6	↓ 534.9

6 Results and Discussion

The evaluation results are listed in Table 1. Δ_{base} represents the change in performance over the best performing baseline approach for each metric.

Comparing with Baselines We remark that *AMII* outperforms the *baselines* in terms of *behavior appropriateness*. This is reflected through low errors of MAE and RMSE represented by Δ_{base} (↓ 0.024 and ↓ 0.020 respectively). For the density distribution, via the KS test, we observe that *AMII* performs comparatively less than the *baselines* indicating that *AMII* possesses the least similar density distribution compared to that of the GT. In detail, ASAP performs the best in terms of having the most similar density w.r.t. GT (0.282) and *AMII* the worst (0.437) with Δ_{base} of ↑ 0.155. This low performance of *AMII* does not imply that it generates wrong *SIA* behavior but that it has either a smaller or a wider range of behavior variety than that of the GT. The focus of this study is not to produce a variety of behaviors but to generate *SIA* behaviors that are *adaptive* to its interlocutor. Thus, this weak performance of the KS test metric is not critical for our aim. Moreover, *AMII* performs the best in terms of the *reciprocal adaptation resemblance* metrics as seen in the Table 1. DTW, synchrony, and entrainment loop of $\hat{A}\&U$ show that *AMII* resembles the GT the most with Δ_{base} of ↓ 79.7, ↓ 4.6, and ↓ 534.9 respectively. With TLCC, we remark that the sym-IL-LSTM is the closest to the GT while *AMII* is the farthest one with Δ_{base} of ↑ 0.042. As DTW considers the variation of sequence length while being invariant to speed unlike TLCC, it represents better the *global correlation*. Thus, for the interpretation, we can put more emphasis on the DTW results compared to that of TLCC. This comparative study shows that the inclusion of explicit modeling of *intra-personal* relation leverages the quality of produced gestures in terms of both *behavior appropriateness* and *reciprocal adaptation resemblance*.

Ablation studies The ablation of each of *AMII* key encoder - E_m , E_{dual} , and E_{inter} - results in the improvement of the *reciprocal adaptation resemblance*. This was seen by an increase in DTW (87.0, 33.0, and 187.6 respectively), synchrony (48.3, 124.9, and 45.1 respectively), and entrainment loop resemblance (385.2, 900.2, and 402.7 respectively). TLCC shows that the insertion of E_m improves the *AMII* by 0.020 along with E_{inter} by 0.030. However, E_{dual} slightly deteriorates the performance by 0.009. As in the *baseline comparison study*, it is better to concentrate on the other *reciprocal adaptation resemblance* metrics as DTW is a more dynamic measure of synchrony than TLCC. However, this enhancement of *reciprocal adaptation resemblance* is at the expense of lowering its *behavior appropriateness* performance. This is observed via MAE (0.057, 0.013, 0.020 respectively) and RMSE (0.065, 0.011, 0.019 respectively). The same conclusion can be drawn by looking at KS test result. The fall of performance is seen for the additions of E_{dual} (0.041) and E_{inter} (0.031) while E_m improves (0.078). This compromise of losing *behavior appropriateness* to gain an *adaptive* one may be a good exchange. It is more valuable to generate *SIA* behaviors with *adaptation capacity* than to reproduce the same GT behavior. In fact, in a human-human interaction there could be multiple possible behaviors and generation timings facing a same interacting partner’s behavior. This might vary depending on the various factors such as the context, situation, and interlocutor’s personality and mood.

We can conclude that it is important to model the *intra-personal* relation with the encodings of E_m and E_{dual} , and the *inter-personal* relation with E_{inter} to synthesize *adaptive non-verbal facial gestures* for both roles as *speaker* and *listener*.

7 Conclusions and Future works

In this paper, we propose a new approach to generate *adaptive SIA* behavior as both *speaker* and *listener* by encoding the multimodality and *intra-personal* and *inter-personal* relationships. We conclude that *AMII* model achieves state-of-the-art performance notably in terms of *reciprocal adaptation resemblance*. Our approach still has some limitations. We do not model the inter-personal relation memory which could improve our results. As the next step, we plan to add this component. We also plan to integrate and assess *AMII* within a real-time *SIA-User* interaction. Our overall aim is to model *SIA* that can *socially engage* users in an interaction through its behaviors that capture *intra-personal* and *inter-personal* relations. In the next future, we will subjectively evaluate the dimensions of social attitudes and engagement to validate this hypothesis.

References

- [1] Justine Cassell, Tim Bickmore, Lee Campbell, Hannes Vilhjalmsson, and Hao Yan. Human conversation as a system framework: Designing embodied conversational agents. *Embodied conversational agents*, pages 29–63, 2000.
- [2] Kazuhiro Shidara, Hiroki Tanaka, Hiroyoshi Adachi, Daisuke Kanayama, Yukako Sakagami, Takashi Kudo, and Satoshi Nakamura. Automatic thoughts and facial expressions in cognitive restructuring with virtual agents. *Frontiers in Computer Science*, 4:8, 2022.
- [3] Yanghee Kim, Jeffrey Thayne, and Quan Wei. An embodied agent helps anxious students in mathematics learning. *Educational Technology Research and Development*, 65(1):219–235, 2017.
- [4] Joseph N Cappella. Mutual adaptation and relativity of measurement. *Studying interpersonal interaction*, 1:103–117, 1991.
- [5] Judee K Burgoon, Laura K Guerrero, and Valerie Manusov. Nonverbal signals. *The SAGE handbook of interpersonal communication*, pages 239–280, 2011.
- [6] Judee K Burgoon, Lesa A Stern, and Leesa Dillman. *Interpersonal adaptation: Dyadic interaction patterns*. Cambridge University Press, 1995.
- [7] Martin J Pickering and Simon Garrod. Toward a mechanistic psychology of dialogue. *Behavioral and brain sciences*, 27(2):169–190, 2004.
- [8] Mark L Knapp, Judith A Hall, and Terrence G Horgan. *Nonverbal communication in human interaction*. Cengage Learning, 2013.
- [9] Jiyeon Woo, Catherine Pelachaud, and Catherine Achard. Reciprocal adaptation measures for human-agent interaction evaluation. In *ICAART*, 2023.
- [10] Beatrice Biancardi, Soumia Dermouche, and Catherine Pelachaud. Adaptation mechanisms in human-agent interaction: Effects on user’s impressions and engagement. *Frontiers in Computer Science*, 3:696682, 2021.
- [11] Simon Alexanderson, Gustav Eje Henter, Taras Kucherenko, and Jonas Beskow. Style-controllable speech-driven gesture synthesis using normalising flows. In *Computer Graphics Forum*, volume 39, pages 487–496. Wiley Online Library, 2020.
- [12] Uttaran Bhattacharya, Nicholas Rewkowski, Abhishek Banerjee, Pooja Guhan, Aniket Bera, and Dinesh Manocha. Text2gestures: A transformer-based network for generating emotive body gestures for virtual agents** this work has been supported in part by aro grants w911nf1910069 and w911nf1910315, and intel. code and additional materials available at: <https://gamma.umd.edu/t2g>. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pages 1–10. IEEE, 2021.
- [13] Mireille Fares, Catherine Pelachaud, and Nicolas Obin. Zero-shot style transfer for multimodal data-driven gesture synthesis. In *2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–4. IEEE, 2023.
- [14] Will Feng, Anitha Kannan, Georgia Gkioxari, and C Lawrence Zitnick. Learn2smile: Learning non-verbal interaction through observation. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4131–4138. IEEE, 2017.
- [15] Soumia Dermouche and Catherine Pelachaud. Generative model of agent’s behaviors in human-agent interaction. In *2019 International Conference on Multimodal Interaction*, pages 375–384, 2019.
- [16] Marc D Pell. Prosody–face interactions in emotional processing as revealed by the facial affect decision task. *Journal of Nonverbal Behavior*, 29:193–215, 2005.

- [17] Hani C Yehia, Takaaki Kuratate, and Eric Vatikiotis-Bateson. Linking facial animation, head motion and speech acoustics. *Journal of phonetics*, 30(3):555–568, 2002.
- [18] Joseph Grafsgaard, Nicholas Duran, Ashley Randall, Chun Tao, and Sidney D’Mello. Generative multimodal models of nonverbal synchrony in close relationships. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 195–202. IEEE, 2018.
- [19] Patrik Jonell, Taras Kucherenko, Gustav Eje Henter, and Jonas Beskow. Let’s face it: Probabilistic multi-modal interlocutor-aware generation of facial gestures in dyadic settings. In *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*, pages 1–8, 2020.
- [20] Jieyeon Woo, Catherine Pelachaud, and Catherine Achard. Creating an interactive human/agent loop using multimodal recurrent neural networks. In *WACAI 2021*, 2021.
- [21] Jieyeon Woo, Catherine Pelachaud, and Catherine Achard. Asap: Endowing adaptation capability to agent in human-agent interaction. In *28th International Conference on Intelligent User Interfaces*, 2023.
- [22] Evonne Ng, Hanbyul Joo, Liwen Hu, Hao Li, Trevor Darrell, Angjoo Kanazawa, and Shiry Ginosar. Learning to listen: Modeling non-deterministic dyadic facial motion. *arXiv preprint arXiv:2204.08451*, 2022.
- [23] Paul Ekman and Wallace V Friesen. Measuring facial movement. *Environmental psychology and nonverbal behavior*, 1(1):56–75, 1976.
- [24] Paul Ekman and Wallace V Friesen. Facial action coding system. *Environmental Psychology & Nonverbal Behavior*, 1978.
- [25] Dwight L Bolinger. A theory of pitch accent in english. *Word*, 14(2-3):109–149, 1958.
- [26] Angelo Cafaro, Johannes Wagner, Tobias Baur, Soumia Dermouche, Mercedes Torres Torres, Catherine Pelachaud, Elisabeth André, and Michel Valstar. The noxi database: Multimodal recordings of mediated novice-expert interactions. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, page 350–359. Association for Computing Machinery, 2017.
- [27] Florian Eyben, Martin Wöllmer, and Björn Schuller. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 1459–1462, 2010.
- [28] Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. Openface: an open source facial behavior analysis toolkit. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–10. IEEE, 2016.
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [30] Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*, pages 464–472. IEEE, 2017.
- [31] Frank J Massey Jr. The kolmogorov-smirnov test for goodness of fit. *Journal of the American statistical Association*, 46(253):68–78, 1951.
- [32] Steven M Boker, Jennifer L Rotondo, Minquan Xu, and Kadiah King. Windowed cross-correlation and peak picking for the analysis of variability in the association between behavioral time series. *Psychological methods*, 7(3):338, 2002.
- [33] Meinard Müller. Dynamic time warping. *Information retrieval for music and motion*, pages 69–84, 2007.