

Face-to-Face Contrastive Learning for Social Intelligence Question-Answering

Alex Wilf^{1,*}, Martin Q. Ma^{1,*}, Paul Pu Liang², Amir Zadeh^{3,1}, Louis-Philippe Morency¹

¹ Language Technologies Institute, Carnegie Mellon University

² Machine Learning Department, Carnegie Mellon University

³ Amazon Alexa AI

Abstract—Creating artificial social intelligence – algorithms that can understand the nuances of multi-person interactions – is an exciting and emerging challenge in processing facial expressions and gestures from multimodal videos. Recent multimodal methods have set the state of the art on many tasks, but have difficulty modeling the complex face-to-face conversational dynamics across speaking turns in social interaction, particularly in a self-supervised setup. In this paper, we propose Face-to-Face Contrastive Learning (F2F-CL), a graph neural network designed to model social interactions using *factorization nodes* to contextualize the multimodal face-to-face interaction along the boundaries of the speaking turn. With the F2F-CL model, we propose to perform contrastive learning between the factorization nodes of different speaking turns within the same video. We experimentally evaluate our method on the challenging Social-IQ dataset and show state-of-the-art results.

I. INTRODUCTION

Effectively modeling face-to-face social interaction is an important task that presents key challenges for automatically understanding body language [7], in part because the benchmarks for the task are relatively low-resource [37]. Recent years have shown the effectiveness of self-supervised learning combined with expressive model architectures such as graph neural networks and transformers [14] to improve state-of-the-art performance on low-resource tasks [27]. In this paper, we consider one self-supervised approach in particular, contrastive learning, and ask *How can we train models to understand face-to-face social interaction using self-supervised contrastive learning?*

Contrastive learning is an approach to self-supervised learning in which the model is encouraged to encode “positive pairs” of inputs in a similar way and “negative pairs” differently [23], [4], [13]. An effective contrastive learning setup for this task will require two components: a way of defining positive and negative pairs that provides a strong learning signal and an architecture that can effectively model face-to-face social interaction. To do the latter, a model must both understand the conversational dynamics of each speaking turn and be able to understand interactions across speaking turns. However, despite efforts to improve contrastive learning for videos [42], [17], there has not been a contrastive framework that effectively models social interaction by using the rich conversational dynamics of videos. A good approach should contextualize information within speaking turns, meanwhile

* Equal contribution

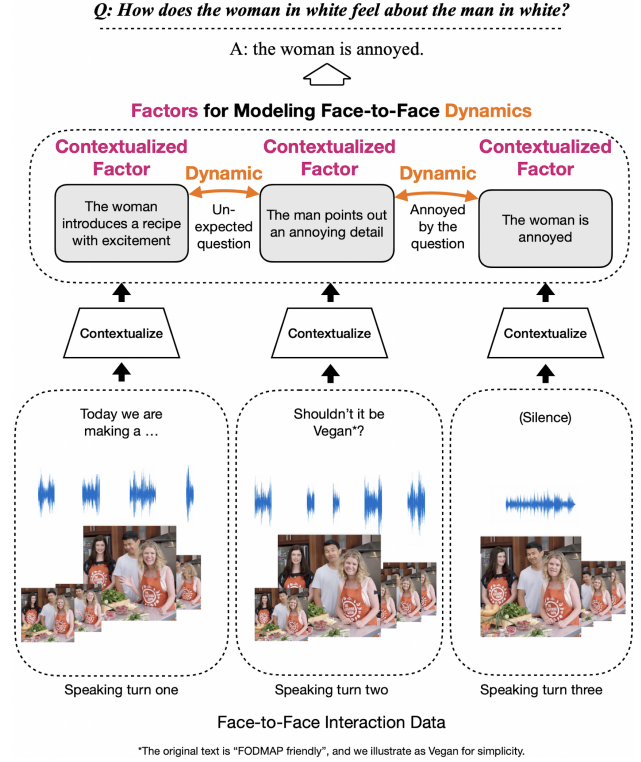


Fig. 1. An example approach for modeling social interaction. It first contextualizes face-to-face interaction within each speaking turns into some contextualized latent factors, and then models the cross speaking-turn dynamics. This approach nicely captures the face-to-face conversational dynamics by leveraging speaking turns.

modeling the dynamics across turns. We illustrate the essence of this approach in Figure 1, where each contextualized factor could capture speaking-turn information and interact with other contextualized factors to model conversational dynamics.

In this paper, motivated by the idea of contextualization and interaction, we propose **Face-to-Face Contrastive Learning**, a graph-based approach to model multimodal social interaction using contrastive learning. Our method first encodes the multimodal videos into fully connected graphs based on speaking turn boundaries, and then contextualizes each speaking turn using a novel node structure, which we call factorization nodes. With this structure, we perform contrastive learning by explicitly modeling the interactions

between factorization nodes. We treat the factorization nodes from the same speaking turn as the positive pair, and those from different speaking turns as the negative pairs. The factorization nodes must contextualize information within each speaking turn to perform well on the contrastive learning objective, and the interactions between different factorization nodes can effectively capture face-to-face conversational dynamics across speaking turns.

Empirically, we evaluate on the publicly available Social-IQ dataset [37] comparing with both supervised and self-supervised baselines, and achieve state-of-the-art results on the social interaction question-answering task.

II. RELATED WORK

a) Social Intelligence: A benchmark called Social-IQ (Social Intelligence Queries) [37] provides a diverse set of videos and evaluates models’ “social intelligence” through many “why” and “how” questions related to social interaction. We run our experiments on this benchmark because it is more focused on social interaction than other video QA datasets such as TVQA [21] and MovieQA [26]. Graph neural networks have not been used for this task before, so to thoroughly evaluate our claims, we provide results not only our method, but also for ablations of our method that represent common GNN approaches to other video QA tasks.

b) Multimodal Graph Network: Graph neural networks (GNN) [11], [24] have been proposed to encode graph-structured data, and have been applied on multimodal tasks such as Visual Dialog [12] and VQA [15]. [14] show that multimodal transformers can be considered a special case of fully connected GNNs, including large multimodal transformers such as MERLOT [41] and UNITER [6]. Most similar to our work is [34], which uses a graph neural network to model cross-modality, within-modality, and temporal connections, but differs from our strategy in our factorization strategy and contrastive learning setup.

c) Self-Supervised Contrastive Learning: Self-supervised learning attempts to create “labels” from data (e.g. masking) or the training process itself (e.g. contrasting) to pretrain representations [16] and uses the learned representations for downstream tasks [5], [13], [20], [23]. A well-studied method of self-supervised learning is contrastive learning [23], [13], [4], which tries to construct pairs of samples that should be represented similarly (termed positive pairs) and pairs of samples that should be represented differently (termed negative pairs) and learn a model that distinguishes the positive pair from the negative pairs. Contrastive learning methods are widely used in graph representation learning. For example, [29] treats node-level and graph-level representations of the same graph as positive pairs, and different graphs as negative pairs. [35] stochastically perturbs the edges and nodes on the same graph to create a positive pair and treats random graphs as negative pairs.

The most similar to our work on contrastive learning for videos are [42] and [17]. [42] learns representations of videos by defining a loss considering positive pairs as different

augmentations of the same temporal segment within a single video, and negative pairs as different temporal segments. Our work differs from this approach in that instead of defining segments with a fixed window size, we assign segment boundaries based on speaking turns. [17] also considers temporal segments within videos, but instead of contrasting at the temporal segment level as we do, contrasts by collating frames drawn at random from different temporal segments. There are many other works that perform contrastive learning on videos [30], [36], [9], [32], but we have found none that contrasts based on dynamic and semantically meaningful temporal boundaries (such as speaking turn).

d) Factorization in Multimodal Networks: [25] groups neighboring frames of sequences that have similar semantic information and uses group-level representation for event summarization. [22] models the hierarchical relations between sentences and phrases and between whole images and image regions to learn a joint embedding for image captioning. Many more networks have used factorization effectively on a variety of multimodal tasks, such as [3], [31], [28], [40].

III. FACE-TO-FACE CONTRASTIVE LEARNING

In this section, we describe the proposed Face-to-Face Contrastive Learning (F2F-CL) to model social interaction. Our method consists of two main steps: (1) encoding multimodal video into graphs and encoding speaking-turn level information into factorization nodes and (2) performing contrastive learning using the factorization nodes.

A. Problem Definition

We define the task of modeling social interactions given the multimodal video, including textual input T , visual input V and acoustic input A , and questions and answers regarding the social interactions in the video. The task is to model inter-speaker interaction to correctly answer the questions. Following prior work [33], [37], we feed each modality into a modality-specific encoder to produce unimodal features with the same hidden dimensions (details in Section IV-A.0.c). Then, we partition the processed features into groups based on speaking turn boundaries [18]. We leave an investigation of the relationship between speaking turn boundary quality and our algorithm’s performance to future work.

B. Graph Encoding and Factorization

We create speaking-turn graphs based on unimodal features. Formally, we first create nodes denote the visual (v), textual (t) and acoustic (a) nodes of the speaking turn s as $\{v_i^s\}_{i=1}^{N_v}$, $\{t_i^s\}_{i=1}^{N_t}$ and $\{a_i^s\}_{i=1}^{N_a}$, respectively, where each node is a vector representing one modality at one temporal step. The numbers of modality nodes N_v, N_t, N_a are possibly different for each modality. For each pair of nodes, there is a connecting edge representing their interaction. We denote the edge connecting t_i^s to v_j^s by e_{t_i, v_j}^s . Each edge is directed (e_{t_i, v_j}^s is different from e_{v_j, t_i}^s). Since the graph is fully connected within each turn, there will be $N_e = \binom{N_v + N_t + N_a}{2}$ edges for each turn. For each edge, there will be a normalized attention

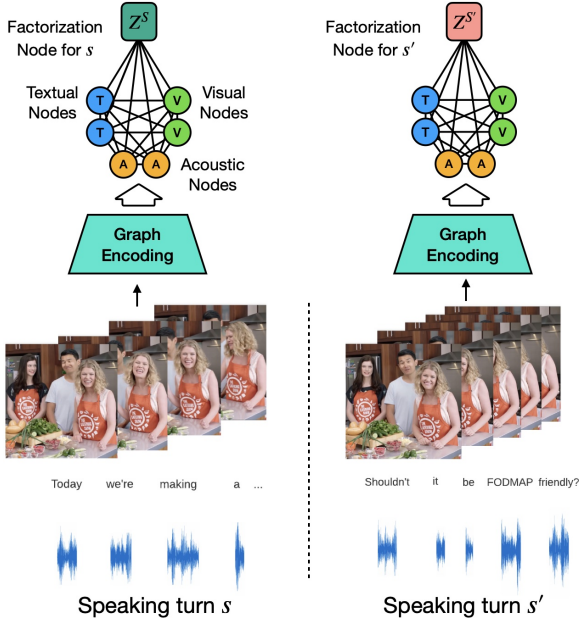


Fig. 2. Constructing speaking-turn graphs. We first segment the videos by speaking turn, and construct a fully connected graph for every speaking turn where a node is a modality at one time step. Then the factorization node extracts information from all modalities in the speaking turn. Each factorization node represents one speaking turn.

score α and an unnormalized attention score β , whose details we will introduce soon.

Next, we connect each node to a single “factorization node” z^s , which pools representation information from the whole speaking-turn graph by fully connecting to all other nodes in the graph. The edge connecting z^s and v_j^s , for example, will be $e_{z^s, v_j^s}^s$. We illustrate the overall architecture in Figure 2.

Then, we perform attention updates using the node embeddings and edge weights in the graphs. There are three steps: the first is to compute the raw attention score $\beta^{[h]}$ between two nodes (we use t_i and v_j as an example, neglecting the s superscript for simplicity because all operations are within the same speaking turn s):

$$\beta_{t_i, v_j}^{[h]} = \text{LeakyRelu}(e_{t_i, v_j}^s \cdot [t_i || v_j]), \quad (1)$$

where $[||\cdot]$ denotes the concatenation of node representations. The $[h]$ index is the head index of multi-head attention. The next step is then to calculate the normalized $\alpha^{[h]}$ of $\beta_{t_i, v_j}^{[h]}$ on all neighbors of t_i using softmax:

$$\alpha_{t_i, v_j}^{[h]} = \frac{\exp(\beta_{t_i, v_j}^{[h]})}{\sum_{k \in N_{t_i}} \exp(\beta_{t_i, k}^{[h]})}, \quad (2)$$

where N_{t_i} denotes all the neighboring nodes of t_i . In our fully connected case, N_{t_i} includes all nodes except t_i . In the last step, node t_i is updated according to the attention scores:

$$t_i = \text{concat}(\sum_{h=1}^H \alpha_{t_i, k}^{[h]} t_i) \quad (3)$$

where H is the number of total attention heads. After this step, t_i encodes heterogeneous modal-temporal interactions

between t_i and its neighbors. We perform such updates for all other modality nodes, $\{v_i^s\}_{i=1}^{N_v}$, $\{t_i^s\}_{i=1}^{N_t}$ and $\{a_i^s\}_{i=1}^{N_a}$. More importantly, the same also applies to the factorization node z , where its update will be based on the neighboring nodes N_z and the attention scores between z and its neighbors:

$$z = \text{concat}(\sum_{h=1}^H \alpha_{z, k}^{[h]} z) \quad (4)$$

Because N_z contains all the modality nodes in speaking turn s , z contextualizes all the modality information in speaking turn s . Such updates are performed on all speaking turns so that z^s contextualizes the speaking turn s .

C. Contrastive Learning with Factorization Nodes

Next, we introduce the self-supervised contrastive learning step. Contrastive learning [4], [13] creates positive and negative pairs from data and uses a contrastive objective to pull together embeddings of positive pairs and push away embeddings of negative pairs. The first step is to create positive and negative pairs. We demonstrate the overall approach in Figure 3.

a) *Creating positive and negative pairs:* Following [4], [35], we define *stochastic data augmentations* on graphs as creating realistic new graphs by applying a certain transformation without affecting the semantics label of the graph. We want to augment each speaking turn graph to create two different but related graphs. Given a speaking turn graph \mathcal{G} , we discuss the following four types of graph augmentation technique to create *two* graphs for each speaking turn:

- 1) **Node Dropping** drops a ratio of the nodes in the graph \mathcal{G} , together with their connecting edges with other nodes in the graph. It assumes that the missing part of the nodes in graph \mathcal{G} does not affect its semantic meaning.
- 2) **Edge Perturbation** randomly changes the edges in \mathcal{G} by adding or dropping edges by a certain ratio. It assumes that the semantics of \mathcal{G} is reasonably robust to edge connection variations.
- 3) **Node Masking** sets the embeddings of some nodes to zero vectors. The underlying assumption is that missing nodes do not affect the overall graph information much.
- 4) **Subgraph Sampling** samples a subgraph from \mathcal{G} using a random walk [35]. It assumes that the semantics of \mathcal{G} can be much preserved in its local (partial) structure.

Node dropping, edge perturbation, and node masking follow uniform distributions to select nodes or edges. Note that we do not perform augmentation on the factorization nodes or their corresponding edges, as they need to extract information from all other nodes.

Next, we perform contrastive learning. After graph augmentation, every speaking turn now has two augmented graphs, each with a unique factorization node. We define positive pairs as two randomly augmented graphs from one speaking turn graph, while negative pairs are graphs from different speaking turns. We use a contrastive objective which takes into account only the factorization nodes, not the whole graph. Assuming that we have N_S speaking turns, we denote

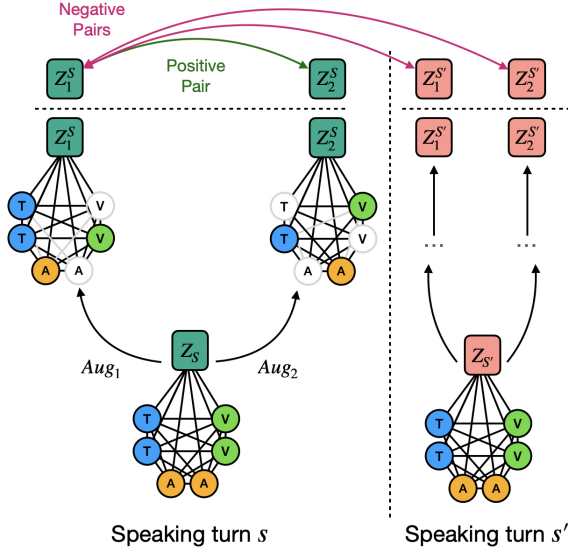


Fig. 3. Contrastive learning example. Aug_1 and Aug_2 are two random augmentations. One speaking-turn graph is augmented to create two augmented graphs, and contrastive learning is performed on the factorization nodes. Nodes from the same speaking turn are a positive pair, and nodes from different speaking turns are negative pairs.

factorization nodes from a positive pair as z_1^s, z_2^s . Inspired by the popular InfoNCE loss [23], [4], we formulate our loss below for the s -th speaking turn:

$$\mathcal{L}_{ssl} = -\log \frac{e^{\text{sim}(z_1^s, z_2^s)/\tau}}{\sum_{s' \neq s} \sum_{i=1}^2 e^{\text{sim}(z_1^s, z_i^{s'})/\tau}}, \quad (5)$$

where $\text{sim}(\cdot)$ is a cosine similarity function and τ is a temperature hyperparameter. The final loss is computed for all positive pairs of factorization nodes at different speaking turns.

b) Why does factorization help?: The key benefit of factorization is improving contrastive learning to model social interactions. By factorizing based on speaking turns, we can augment each speaking turn to create multiple positive pairs and negative pairs within a single video. Each video could generate many positive and negative pairs, compared to the case without factorization where only multiple videos can generate positive and negative pairs. This scales the number of positive and negative pairs by the average number of speaking turns per video, increasing the data efficiency of our method.

D. Supervised Fine-tuning on Question-Answering

To evaluate the representation quality of the model after self-supervised pretraining we perform a supervised fine-tuning, where we freeze the learned model and fine-tune a two-layer network that considers the hidden representation of the graph neural network. In our case, because we evaluate on Social-IQ [37], which is a data set of questions and answers, we use the questions and answers as labels to perform the fine-tuning.

To incorporate the QA representations, the question and answer sequences (both correct and incorrect) are first encoded by separate LSTM networks as in [37], and the

representations of questions, correct answers, incorrect answers, and multimodal input are all contextualized in a graph convolution operation. We use the updated Graph Attention Layer (GATv2) [2] as our graph convolution to operate on all edges, including edges connecting two modality nodes, edges connecting question and answer (correct and incorrect) nodes, and edges connecting modality and Q/A nodes. We run a fixed number of convolutional layers over the graph to create a contextualized representation vector for each pair of question and answer. Our fine-tune layer predicts, using contextualized representations for each question and answer pair, a number between zero and one (zero means the answer to the question is wrong, while one means the answer to the question is correct) by a sigmoid activation. We randomize which position the correct answer is in and train with a mean squared error loss between the predicted and correct answer indexes.

IV. EXPERIMENTAL METHODOLOGY

Our experiments evaluate the quality of the representations learned through our novel contrastive approach. In the following, we describe the dataset as well as the baselines.

A. Experimental details

a) Dataset: We are using the Social-IQ dataset [37]. Social-IQ proposes using a diverse set of videos to reason about human social interaction. Videos often contain a social scene, where different people interact with each other. The dataset evaluates models' performance on certain types of "why" and "how" questions about the interaction between people in the video based on the social scene, which is well below human performance (68.8% vs. 95.08%) [19], even though the models perform well in other multimodal QA datasets. We split the dataset along the standard splits defined in the Social-IQ dataset SDK and evaluate our algorithm on the held out validation set.

b) Data Preprocessing: We use as input modalities BERT [8] representations for text, COVAREP representations for audio, and features from a large pretrained image transformer called BEIT [1]. We use the "aligned" version of the Social-IQ dataset, meaning that we average representations of higher frequency modalities across timestamps denoting the start and end of each word boundary, as in [38], [37]. The number of modality nodes for each graph is equivalent across modalities, where, for example, the first vision node represents the average of the BEIT vectors for the vision frames corresponding to the timestamps for the first word.

c) Implementation Details: Our hyperparameter choices are that we use the AdamW optimizer [8] with learning rate .001, shared dimensionality across all graph nodes of 80, and 2 graph attention layers, each containing graph convolutions for each modal-temporal pair. We train with a batch size of 15, sequence length 25, and 25 max epochs (although the network often converges well before then), reporting the max validation accuracy.

For supervised training (baselines we will introduce next), we model the modality-modality components of each graph

similar to the fashion in [34], and implemented in the Pytorch-Geometric package [10]. We add the following edges and associated graph convolutions to the graph: the edge connecting questions, the edge connecting answers, the edge connecting questions and answers, and the edge connecting factorization nodes z^s to the question and answer nodes. Our training takes 3 hours on a single NVIDIA 3060Ti GPU. The code will be released soon.

B. Baselines

In the following, we introduce two types of baselines: self-supervised baselines and supervised baselines. The self-supervised baselines are contrastive baselines without factorization nodes. The supervised baselines are the prior state-of-the-art baselines on Social-IQ.

1) *Self-Supervised Baselines*: The self-supervised baselines we test perform the same graph augmentations and contrastive pre-training but with two different ways of creating factorization nodes: none (one video-level factorization node created by averaging all node representations as in [34]), and mean (segment the graph, create factorization nodes for each segment using mean readout of nodes in the subgraph).

Video-Level Contrastive, No factorization nodes. This approach creates positive and negative pairs from multiple videos. Each positive pair has two graphs created based on the whole *video*, not on a single speaking turn. Negative pairs are graphs created from different videos. In other words, this approach uses a fully graph to represent each video, augments the graph to create two slightly different graphs for each video, and performs contrastive learning on augmented video-level graphs.

Speaking-Turn-Level Contrastive, Mean of nodes. This approach is an ablation of our method. This baseline first segments the graph based on speaking turns, similar to ours, and then directly averaging all node embeddings of the speaking turn graphs as the representations of the graph. The representation is similar to our proposed factorization node in the sense that both contextualize the whole graph. The difference is that the baseline does not parametrize the connections between the graph representation and the modality nodes, whereas the proposed factorization nodes do.

2) *Supervised Baselines*: Below we also introduce several other baselines that achieved prior state-of-the-art results on Social-IQ. All of them are supervised methods. We consider three published baselines for the Social-IQ task as well as two versions of our model. The baseline models that we compare are listed below.

TMFN [37], which uses Tensor Fusion [38] for multimodal fusion in the recurrent stages of the MFN network [39].

MCQA [19], which proposes to perform cross-modal alignment and multimodal context-query alignment to solve the social intelligence task.

MTAG [34], which uses a fully connected modal-temporal graph neural network to contextualize nodes, then uses a mean readout operation to arrive at a single latent scene vector which is used for classification. MTAG was originally applied to classification tasks; to apply it to QA, we reimplement

it by fusing the graph readout representation with the $q/a/i$ nodes and use a two-layer MLP as in our model.

The supervised baselines, which are versions of our model trained from scratch in a supervised way, are as follows.

Video-Level Supervised: This is an ablation where the question and answer nodes are fully connected to the multimodal graphs at the video level.

F2F-CL Supervised: The F2F-CL model we proposed, but trained directly with question and answer labels. The difference between F2F-CL Supervised and the Video-Level Supervised baseline is that F2F-CL Supervised connects speaking-turn level graphs to the question and answer nodes, while Video-Level Supervised connects Video-Level graphs to the questions and answer nodes.

V. RESULTS

TABLE I
RESULTS ON SOCIAL-IQ DATASET. THE PROPOSED F2F-CL OUTPERFORMS ALL OTHER SELF-SUPERVISED AND SUPERVISED BASELINES AND ACHIEVES STATE-OF-THE-ART RESULT.

Model	Accuracy (%)
Social IQ QA (Self-Supervised)	
Video-Level Contrastive	67.90
Speaking-Turn-Level Contrastive	70.76
F2F-CL	72.03
Social IQ QA (Supervised)	
MFN [39]	64.82
MCQA [19]	68.80
MTAG [34]	65.58
Video-Level Supervised	70.66
F2F-CL Supervised	72.01

First, we find that segmenting the videos into speaking turns helps improve performance. This is shown in Table I, as the overall accuracy of contrastive learning at the level of speaking turn (segment level contrastive and F2F-CL) is greater than contrastive learning at video level, even when factorization nodes are defined without parametrization as the mean of nodes in their subgraph. Then we find that the proposed F2F-CL, which relies on factorization on speaking turns with parametrized factorization nodes, achieves the state-of-the-art results on both baselines.

This result also holds in the supervised case, and surprisingly, we find that the self-supervised setup performs on par with the supervised setup. This is an interesting result, suggesting that the contrastive pretraining within each video with the proposed factorization node structure is a particularly effective method for learning useful multimodal graph representations in a self-supervised way from relatively small amounts of data (compared to many pretraining regimes).

VI. QUANTITATIVE ANALYSIS

In this section, we consider a number of questions to support or challenge our intuition for why F2F-CL learns strong representations from self-supervision.

a) Factorization Nodes Contain Speaker Information:

We expect that our speaking-turn factorization nodes will learn high level information about the speaking turns they represent. We constructed a dataset to ask the question: *Can we discern, better than random, which speaker a given factorization node refers to?* To do this, we constructed a balanced dataset as follows: for each video in Social-IQ that contains at least two factorization nodes of speaking turn of one speaker (z_1^a , z_2^a) and one speaking turn of another (z_1^b), we add two data points to the dataset: (z_1^a , z_2^a ; 1) and (z_1^a , z_1^b ; 0). We attempt to learn a simple linear classifier that takes the concatenated input vectors and attempts to distinguish whether they are from the same speaker or not. We train an MLP with two linear layers, one with dimensions [160, 80], the other with dimensions [80, 1], with ReLU and Dropout between the layers and a Sigmoid layer afterwards for max 30 epochs with early stopping and MSE loss, then we report the validation accuracy. Training and validation sets are the factorization nodes saved from the trained F2F-CL model on the Social-IQ training and validation sets. We find that this model achieved 60.29% accuracy, above the 50% accuracy that we would expect randomly.

b) The Fully Connected GNN Pays Less Attention Across Speaking Turns: As we describe in Section I, we expect that the fully connected MTAG [34] will pay less attention across speaking turns, as it will be trying to learn the “general idea” of each speaking turn before reasoning across them. We found support for this idea – attentions for nodes across speaking turns on the validation set were, on average, 36.17% higher than attentions between nodes within the same speaking turn.

c) Factorization Reduces the Number of Edges: We expect that factorization will substantially reduce the number of edges in the graph, which would support our intuition that contextualizing graphs across speaker turns is an effective inductive bias such that cross-speaking-turn interactions can be efficiently modeled by the interactions of factorization nodes. We find that factorization reduces the number of edges in multi-speaking-turn graphs. We considered all the videos in Social-IQ for this analysis, as the number of edges is relevant both for the train and validation sets. Our results are visualized in Figure 4. For videos with 4, 5, and ≥ 6 speaking turns, we found a reduction in the number of edges of 5.45%, 7.08%, and 9.96%.

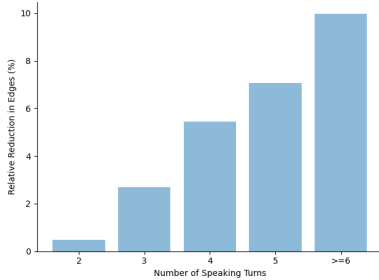


Fig. 4. Factorization reduces the number of edges in the graph from fully connected GNN.

d) Moderate levels of graph augmentations help the

most: Lastly, we investigate how graph augmentation affects self-supervised learning results. We vary the ratio of each of the four graph augmentations, Node Dropping, Edge Perturbation, Node Masking, and, Subgraph Sampling, while fixing the ratios of the other three augmentations. The default augmentation ratio is set to 0.5. From Table II, we observe that F2F-CL benefits from a smaller ratio of augmentations (0.25) on the nodes (Node Dropping and Node Masking) and a moderate ratio of augmentations (0.5) on the edges (Edge Perturbation and Subgraph Sampling) when creating the positive and negative pairs. We attribute this to the fact that the node representation embeddings are much richer, so a medium-to-large (0.5 and 0.75) ratio to drop or mask nodes will likely hurt performances. On the other hand, the edges contains attention weights, so randomly dropping edges will effectively enforcing the model to learn a sparser connections between the factorization nodes and the modality nodes.

TABLE II

AN ABLATION ANALYSIS OF OUR MODEL REVEALS THAT F2F-CL BENEFITS FROM A SMALLER RATIO OF AUGMENTATIONS (0.25) ON NODES AND A MODERATE RATIO OF AUGMENTATION (0.5) ON EDGES TO CREATE POSITIVE AND NEGATIVE PAIRS IN CONTRASTIVE LEARNING. THE NUMBERS BELOW ARE ACCURACY (%). OUR FULL METHOD (WHICH COMBINES AUGMENTATIONS) ACHIEVES 72.03% ACCURACY.

Augmentation	Ratio		
	.25	.5	.75
Node Dropping	66.81	66.79	66.76
Edge Perturbation	66.14	66.22	66.18
Node Masking	68.33	68.12	68.09
Subgraph Sampling	68.01	67.98	67.92

VII. CONCLUSION

In this paper, we proposed a novel method, Face-to-Face Contrastive Learning (F2F CL), to perform contrastive learning in face-to-face interactions. We leverage the natural conversational dynamics from speaking turns by contextualizing information within each speaking turn into factorization nodes and using the factorization nodes to define positive and negative pairs for contrastive learning. Our method effectively improves the performance of social interaction question-answering and achieves state-of-the-art results on the Social-IQ benchmark.

VIII. ACKNOWLEDGEMENTS

This material is based upon work partially supported by BMW, National Science Foundation awards 1722822 and 1750439, and National Institutes of Health awards R01MH125740, R01MH096951 and U01MH116925. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsors, and no official endorsement should be inferred.

REFERENCES

- [1] H. Bao, L. Dong, and F. Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [2] S. Brody, U. Alon, and E. Yahav. How Attentive are Graph Attention Networks? *arXiv:2105.14491 [cs]*, Jan. 2022. arXiv: 2105.14491.
- [3] R. J. Chen, C. Chen, Y. Li, T. Y. Chen, A. D. Trister, R. G. Krishnan, and F. Mahmood. Scaling vision transformers to gigapixel images via hierarchical self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16144–16155, 2022.
- [4] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [5] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [6] Y.-C. Chen, L. Li, L. Yu, A. El Kholy, F. Ahmed, Z. Gan, Y. Cheng, and J. Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020.
- [7] E. De Stefani and D. De Marco. Language, gesture, and emotional communication: An embodied view of social interaction. *Frontiers in Psychology*, 10:2063, 2019.
- [8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [9] S. Ding, M. Li, T. Yang, R. Qian, H. Xu, Q. Chen, J. Wang, and H. Xiong. Motion-aware contrastive video representation learning via foreground-background merging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9716–9726, 2022.
- [10] M. Fey and J. E. Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.
- [11] M. Gori, G. Monfardini, and F. Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pages 729–734. IEEE, 2005.
- [12] D. Guo, H. Wang, H. Zhang, Z.-J. Zha, and M. Wang. Iterative context-aware graph inference for visual dialog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10055–10064, 2020.
- [13] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- [14] Z. Hu, Y. Dong, K. Wang, and Y. Sun. Heterogeneous Graph Transformer. *arXiv:2003.01332 [cs, stat]*, Mar. 2020. arXiv: 2003.01332.
- [15] D. A. Hudson and C. D. Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019.
- [16] A. Jaiswal, A. R. Babu, M. Z. Zadeh, D. Banerjee, and F. Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2021.
- [17] H. Kuang, Y. Zhu, Z. Zhang, X. Li, J. Tighe, S. Schwertfeger, C. Stachniss, and M. Li. Video contrastive learning with global context. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3195–3204, 2021.
- [18] O. Kuchaiev, J. Li, H. Nguyen, O. Hrinchuk, R. Leary, B. Ginsburg, S. Kriman, S. Beliaev, V. Lavrukhin, J. Cook, et al. Nemo: a toolkit for building ai applications using neural modules. *arXiv preprint arXiv:1909.09577*, 2019.
- [19] A. Kumar, T. Mittal, and D. Manocha. Mcqa: Multimodal co-attention based network for question answering. *arXiv preprint arXiv:2004.12238*, 2020.
- [20] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- [21] J. Lei, L. Yu, M. Bansal, and T. L. Berg. Tvqa: Localized, compositional video question answering. *arXiv preprint arXiv:1809.01696*, 2018.
- [22] Z. Niu, M. Zhou, L. Wang, X. Gao, and G. Hua. Hierarchical multimodal lstm for dense visual-semantic embedding. In *Proceedings of the IEEE international conference on computer vision*, pages 1881–1889, 2017.
- [23] A. v. d. Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [24] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.
- [25] Y. Song, L.-P. Morency, and R. Davis. Action recognition by hierarchical sequence summarization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3562–3569, 2013.
- [26] M. Tapaswi, Y. Zhu, R. Stiefelhofen, A. Torralba, R. Urtasun, and S. Fidler. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4631–4640, 2016.
- [27] C. Tran, Y. Tang, X. Li, and J. Gu. Cross-lingual retrieval for iterative self-supervised training. *Advances in Neural Information Processing Systems*, 33:2207–2219, 2020.
- [28] Y.-H. H. Tsai, P. P. Liang, A. Zadeh, L.-P. Morency, and R. Salakhutdinov. Learning factorized multimodal representations. *arXiv preprint arXiv:1806.06176*, 2018.
- [29] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm. Deep graph infomax. *arXiv preprint arXiv:1809.10341*, 2018.
- [30] H. Wu and X. Wang. Contrastive learning of image representations with cross-video cycle-consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10149–10159, 2021.
- [31] N. Wu, X. W. Zhao, J. Wang, and D. Pan. Learning effective road network representation with hierarchical graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 6–14, 2020.
- [32] J. Xu and X. Wang. Rethinking self-supervised correspondence learning: A video frame-level similarity perspective. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10075–10085, 2021.
- [33] J. Yang, Y. Wang, R. Yi, Y. Zhu, A. Rehman, A. Zadeh, S. Poria, and L.-P. Morency. Mtag: Modal-temporal attention graph for unaligned human multimodal language sequences. *arXiv preprint arXiv:2010.11985*, 2020.
- [34] J. Yang, Y. Wang, R. Yi, Y. Zhu, A. Rehman, A. Zadeh, S. Poria, and L.-P. Morency. Mtag: Modal-temporal attention graph for unaligned human multimodal language sequences. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1009–1021, 2021.
- [35] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems*, 33:5812–5823, 2020.
- [36] L. Yuan, R. Qian, Y. Cui, B. Gong, F. Schroff, M.-H. Yang, H. Adam, and T. Liu. Contextualized spatio-temporal contrastive learning with self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13977–13986, 2022.
- [37] A. Zadeh, M. Chan, P. P. Liang, E. Tong, and L.-P. Morency. Social-iq: A question answering benchmark for artificial social intelligence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8807–8817, 2019.
- [38] A. Zadeh, M. Chen, S. Poria, E. Cambria, and L.-P. Morency. Tensor fusion network for multimodal sentiment analysis. *arXiv preprint arXiv:1707.07250*, 2017.
- [39] A. Zadeh, P. P. Liang, N. Mazumder, S. Poria, E. Cambria, and L.-P. Morency. Memory fusion network for multi-view sequential learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [40] A. Zadeh, C. Mao, K. Shi, Y. Zhang, P. P. Liang, S. Poria, and L.-P. Morency. Factorized multimodal transformer for multimodal sequential learning. *arXiv preprint arXiv:1911.09826*, 2019.
- [41] R. Zellers, J. Lu, X. Lu, Y. Yu, Y. Zhao, M. Salehi, A. Kusupati, J. Hessel, A. Farhadi, and Y. Choi. MERLOT Reserve: Neural Script Knowledge through Vision and Language and Sound. *arXiv:2201.02639 [cs, eess]*, Jan. 2022. arXiv: 2201.02639.
- [42] Z. Zeng, D. McDuff, Y. Song, et al. Contrastive learning of global and local video representations. *Advances in Neural Information Processing Systems*, 34:7025–7040, 2021.