

Can Reasons Help Improve Pedestrian Intent Estimation? A Cross-Modal Approach

Vaishnavi Khindkar^{*}, Vineeth Balasubramanian¹, Chetan Arora[‡], Anbumani Subramanian[†], C.V. Jawahar[†]

Abstract—With the increased importance of autonomous navigation systems has come an increasing need to protect the safety of Vulnerable Road Users (VRUs) such as pedestrians. Predicting pedestrian intent is one such challenging task, where prior work predicts the binary cross/no-cross intention with a fusion of visual and motion features. However, there has been no effort so far to hedge such predictions with human-understandable reasons. We address this issue by introducing a novel problem setting of exploring the intuitive reasoning behind a pedestrian’s intent. In particular, we show that predicting the ‘WHY’ can be very useful in understanding the ‘WHAT’. To this end, we propose a novel, reason-enriched PIE++ dataset consisting of multi-label textual explanations/reasons for pedestrian intent. We also introduce a novel multi-task learning framework called MINDREAD, which leverages a cross-modal representation learning framework for predicting pedestrian intent as well as the reason behind the intent. Our comprehensive experiments show significant improvement of 5.6% and 7% in accuracy and F1-score for the task of intent prediction on the PIE++ dataset using MINDREAD. We also achieved a 4.4% improvement in accuracy on a commonly used JAAD dataset. Extensive evaluation using quantitative/qualitative metrics and user studies shows the effectiveness of our approach.

I. INTRODUCTION

Autonomous navigation and Advanced Driver Assistance Systems (ADAS) have seen significant interest and developments in recent years from the computer vision research community in both industry and academia. The development of newer deep learning architectures and models, as well as the availability of datasets for various tasks such as object detection and semantic segmentation for road scenes [3]–[5], have catalyzed the development of deployable solutions for these tasks in autonomous navigation systems. However, an important aspect of such systems and vehicles is how they understand and interact with Vulnerable Road Users (VRUs) like pedestrians [6]. There have been extensive studies in recent years by government and custodial organizations to protect the safety of VRUs, especially pedestrians, as we develop more advanced autonomous navigation systems [7], [8]. In this work, we focus on this key aspect of such systems – understanding pedestrian behavior in road scenes.

Prior efforts in recent years [1], [9]–[14] have attempted the problem of predicting pedestrian intent using motion models, poses and trajectories. While each of these efforts have been vital to progress in this field, an important

aspect in understanding pedestrian behavior that has not been considered in these methods is – reasoning. Studies on human behavior have shown reasoning to be an important link between user belief, intention and behavior [15]. We specifically seek to address this perspective in this work by asking: “Can reasoning help improve pedestrian intent prediction?” We answer the WHAT, and study the importance of WHY in predicting the WHAT. We illustrate this with a sample result of our method in Fig 1. As shown in the figure, the pedestrian has an intention to cross the road (also called ‘crossing’ intention in this work and related literature). This is supported by the reason that ‘*the pedestrian is waiting for the signal to turn red, and is waiting with a group of pedestrians*’. Current state-of-the-art [2] incorrectly predicts the intent as ‘*no-crossing*’ (no intention to cross), while our method – which considers the reason in its prediction – predicts the correct pedestrian intent. Providing a window into such reasoning can also be helpful to a driver in increasing the trust in the model that predicts pedestrian intent. We use the terms reasons and explanations interchangeably at times in this work, although explanations refer to the interpretation of pedestrian intent in our context and not model interpretability.

Motivated by such a reasoning-driven understanding of pedestrian intent, we make two significant contributions in this work: (i) We study the usefulness of reason in predicting pedestrian intent on road scenes, and develop a new methodology based on cross-modal representation learning and attention that uses vision and language modules towards the stated objective; and (ii) We obtain enriched annotations on reasons for pedestrian intent on the benchmark PIE dataset [1], and thus develop the first pedestrian reason+intent prediction dataset which we call PIE++ (we describe in detail why we choose PIE as the dataset for enrichment with reason annotations in Sec III). PIE++ consists of human-referenced, multi-label explanation annotations for all the 1842 pedestrians, and could potentially serve as a useful resource for researchers in the community. Based on a meticulous study (discussed in Sec III), we include commonly understood reasons for pedestrian intent while crossing or not crossing the streets in PIE++, such as: waiting for safe passage, pedestrians doing their work, and so on. These reasons, beyond explaining pedestrian intent, may also provide a driver with a better understanding of the road scene.

From a method perspective, we hypothesize that cross-modal representation learning that includes vision and language modules helps capture pedestrian intent better, com-

^{*}[†][‡] The authors are with the Center for Visual Information Technology (CVIT) Lab, IIIT Hyderabad, IIT Delhi and IIT Hyderabad, India. * vkhindkar@gmail.com, ¹ vineethnb@cse.iith.ac.in, [†] {anbumani, jawahar}@iiit.ac.in, [‡] chetan@cse.iitd.ac.in

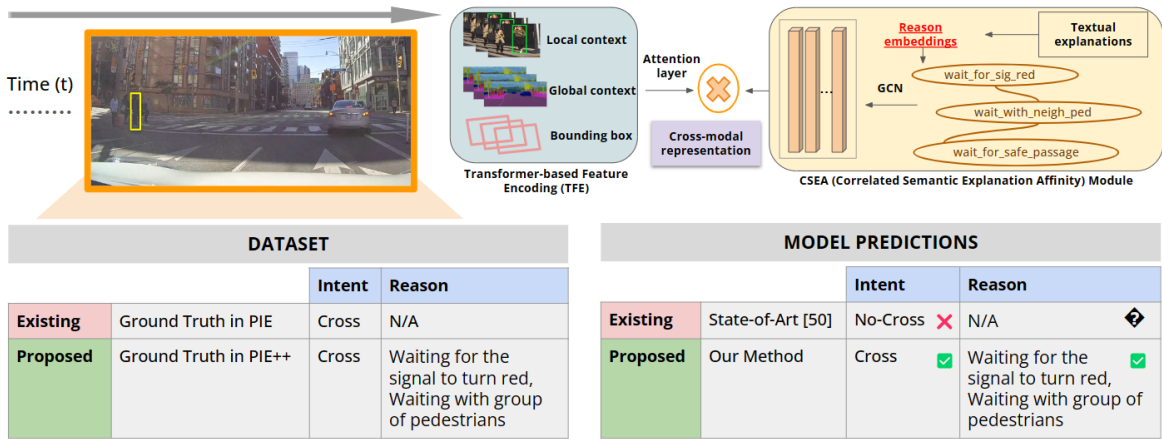


Fig. 1. Illustration of our overall objectives. Given a sample scene from the PIE dataset [1], we study the usefulness of “WHY” (reason) for analyzing “WHAT” (pedestrian intent) by creating a reason-enriched PIE++ dataset. While state-of-the-art [2] makes an incorrect prediction on this scene, MINDREAD, our cross-modal representation learning approach (architecture on top), correctly predicts the crossing intent with the support of a corresponding reason.

pared to prior works that use only visual features. In particular, textual explanations (annotated reasons in PIE++) that are rich in semantic descriptions provide more information than binary labels of crossing intention of a pedestrian, and thus help support better intent prediction. We hence introduce a multi-task formulation based on cross-modal representation learning to predict both pedestrian intent as well as the reason behind the intent. We propose a novel semantic correlation module to capture the relationships between the reason text embeddings, and fuse this with the visual spatio-temporal features through an attention-based mechanism to obtain cross-modal representations. Our overall method, which we call MINDREAD (cross-Modal representation learning moDel for REASONing peDESTrian-intent), consists of three modules: (i) Correlated Semantic Explanation Affinity (CSEA); (ii) Transformer-based Feature Encoding (TFE); and (iii) Attention-based cross-modal representation learner. CSEA assesses the co-existence of explanations to obtain an affinity matrix, and then uses a Graph Convolutional Network to obtain embeddings for textual explanations. These are fused with spatiotemporal features obtained from local visual context, global visual context and bounding boxes in the TFE module using an attention mechanism. The textual-explanation affinity modeling is a critical component of our framework as it helps leverage the reason for the pedestrian intent, thereby improving intent prediction. Following are the key contributions of our work:

- We present a novel perspective of considering reason as a means to predict pedestrian intent and a multi-task formulation to simultaneously predict pedestrian intent and its corresponding reason.
- We provide PIE++, enrichment of a benchmark dataset for intuitively reasoning pedestrian intent, comprising human-referenced multi-label explanation/reason annotations rich in semantic descriptions, providing more information than the binary crossing intention labels alone.
- We propose a new cross-modal representation learning framework, MINDREAD, that exploits semantic correlations between textual embeddings and attention-based

- fusion with spatio-temporal features for intent prediction.
- We perform a comprehensive suite of experiments to study our method and dataset. Our results show significant improvement of 5.6% and 7% in accuracy and F1-score for intent prediction on the PIE++ dataset using MINDREAD. We also achieve 4.4% improvement over state-of-the-art in accuracy on the JAAD dataset [16].

II. RELATED WORK

We herein discuss earlier efforts related to pedestrian intent prediction, the focus area of this work. Other related work such as action prediction [2], [17]–[22] have different objectives such as generating future frames, predicting action types, measuring confidence in event occurrence, and forecast object motion moreover focusing on future action anticipation. Textual explanations [23]–[30] have sometimes been used in earlier work for obtaining insights on a neural network’s understanding of images or scenes. Cross-modal representation learning [31]–[33] works learn latent semantic representations using multiple modalities and is applicable in diverse deep-learning tasks including image captioning, cross-modal retrieval, visual question-answering classification, and detection.

Intent vs. Action? Existing action prediction methods [2], [17]–[22] have different objectives such as generating future frames, predicting action types, measuring confidence in event occurrence, and forecast object motion. We hence note that pedestrian action prediction and pedestrian intent prediction (which is understanding the event before the action) are two different streams of work as defined in the literature and followed by the community in this space. The results from action prediction works are not directly comparable with intent prediction results, given the differences between the two predefined tasks. We also note that our work does not change any pre-existing definitions of pedestrian intent defined in [1]. Also, following [1], we consider the intent only up to a critical point (i.e. frames preceding the crossing point).

Pedestrian Intent Prediction. Prior works [1], [9]–[14], [34], [35] typically model pedestrian intent as motion and use poses or trajectories to predict future goals. PIEint [1] encodes past visual features, and then concatenates them with bounding boxes to predict intent. A framework based on a pedestrian-centric graph is proposed in [9] to uncover spatio-temporal relationships in the scene. Other works such as [11] study early, late, and combined (early+late) fusion mechanisms to exploit skeletal features and improve intent prediction. A conditional random field-based approach is used in [12] for early prediction of pedestrian intent. In the current state-of-the-art for this problem, the recently proposed method in [2] predicts future pedestrian actions and uses predicted action to detect present intent. While each of these efforts have shown promising results, we approach this problem from a different perspective of reasoning. In particular, we employ rich semantic descriptions (reasons) to evaluate the potential of obtaining better prediction of pedestrian intent, and thereby also assist an end-user with a human-interpretable reason. Also, unlike the abovementioned works that utilize only vision features, ours is a first effort in this problem space to use cross-modal, vision-language features for improved pedestrian intent prediction. Besides, while earlier methods largely follow combinations of VGG [36] and GRU [37] architectures, we build on contemporary models such as transformers (Swin Transformer V2 [38], Sentence-BERT [39], and Transformer [40]) for encoding spatio-temporal features and textual explanation embeddings in this work.

III. PIE++: REASON-ENRICHED DATASET FOR PEDESTRIAN INTENT ESTIMATION

Existing datasets for pedestrian intent prediction provide rich behavior and intent annotations; however, as stated earlier, understanding a pedestrian intent along with its reason enhances reliability in the developed system. Besides, it helps predict the intent itself better, as we show in this work. To address this, we propose PIE++, an enrichment of the PIE dataset [1] with reasons for pedestrian intent. We hypothesize that with the significant advancements of text embeddings in recent years, textual representations of such reasons will support the better prediction of pedestrian intent. We annotate different reasons for both crossing and no-crossing pedestrian intents, as shown in Tables II and III. For pedestrians that: (i) have no crossing intention, there could be reasons such as “*Pedestrians are just doing their work*” or “*Two pedestrians just interacting with each other on the road-side*”; and (ii) have crossing intention, reasons could include “*Pedestrian acknowledges the ego-car to stop/slow-down with a hand-ack gesture since the pedestrian is intending to cross*”. Our list of reasons indicate that it is possible for more than one reason to be relevant for a given pedestrian’s intent in a scene. We hence provide *multi-label annotations* for every pedestrian for their crossing vs no-crossing intent.

We compare PIE++ with other existing datasets in Table I. It is evident that considering the reason behind a pedestrian’s intention has been unexplored in prior datasets (including

Dataset	Year	Len (mins)	Frames	Peds	Intent	Reason
KITTI [5]	2012	90	80,000	12,000	×	×
JAAD [16]	2017	46	82,000	337,000	✓	×
BDD100k [4]	2018	60,000	100,000	86,047	×	×
PedX [41]	2019	-	10,152	14,091	×	×
PIE [1]	2019	360	909,480	738,970	✓	×
NuScenes [42]	2020	330	1,400,000	1,400,000	×	×
STIP [9]	2020	923.48	1,108,176	3,500,000	✓	×
PIE++	2023	360	909,480	738,970	✓	✓

TABLE I

COMPARISON OF PIE++ WITH OTHER DATASETS.

PIE). (Datasets with a × in the “Intent” column have no pedestrian intent annotations.) While STIP and JAAD provide pedestrian intent, We extend the PIE dataset [1] in this work, since it is reasonably large and has behavioral annotations that were helpful in validating our reason annotations at times.

Sourcing Reason Annotations in PIE++: Before obtaining the multi-label reason annotations in PIE++, we provided videos for all the 1842 pedestrians in PIE to 5 annotation professionals with experience in data annotation in the mobility industry. They were provided with all relevant information, including pedestrian intent annotations and pedestrian attributes (e.g. looking, walking, standing) to help decide on the plausible reasons leading to the pedestrian’s intent in a given scene. A set of plausible reasons for crossing/no-crossing (C/NC) intent was first developed after viewing the scenes in the dataset. Our reason/explanation categories were decided in a 3-stage process: (i) we first came up with a preliminary set of reason categories based on a user survey; (ii) then thoroughly examined a significant part of the PIE dataset to refine the reason categories; and (iii) added reasons for corner cases (e.g. “pedestrian neglecting ego-car”, “pedestrian giving right-of-way to ego-car”). Each of the annotation professionals subsequently went through a given video sequence as many times as required to understand the scene. We asked the annotators to choose as many options as relevant from the set of reason categories for a particular C/NC intent of the pedestrian. We also included an additional option - ‘other’ if a subject considered a different reason. The pedestrian of interest was highlighted with a bounding box in each video frame for consistency of the study. Since we do not have access to ground truth reason data, we analyzed the agreement among annotators to validate our results with this study. First, we computed the intraclass correlation coefficient (ICC), a measure of inter-rater consistency, commonly used to analyze subjective responses from a population of raters in the absence of ground truth [43]. Our measured ICC was 0.98, which suggests a very high degree of agreement among our raters for the proposed reason annotations (ICC = 1 for absolute agreement). Also notably, the “other” option was never chosen by the annotators in this study, corroborating the sufficiency of the reason categories.

Validating Reason Annotations in PIE++: We also validated the goodness of the reason annotations (on parameters of completeness, reliability, trust, correctness, and usefulness) using an independent study in a lab setting. This study involved watching video sequences for each pedestrian in

Reasons/Explanations
Waiting to cross with a neighbouring pedestrian
Waiting for a safe passage to cross
Waiting to cross with a group of pedestrians
Waiting for the signal to turn red
Waiting since the ego-vehicle speed is high
Waiting since the vehicle speed is high
Waiting for vehicles to slow down
Waiting while giving right-of-way to ego-vehicle
Pedestrian acknowledges ego-vehicle to stop
Pedestrian intends to cross since the signal is red
Pedestrian intends to cross since it's a safe passage
Pedestrian intends to cross since ego-vehicle speed is slow
Pedestrian intends to cross since vehicle speed is slow
Neglects the ego-vehicle

TABLE II

REASONS USED FOR 'Cross' PEDESTRIAN INTENT

Reasons/Explanations
Two pedestrians just interacting (on road-side)
Group of pedestrians just interacting (on road-side)
Pedestrians doing their work on road-side

TABLE III

REASONS USED FOR 'No-Cross' PEDESTRIAN INTENT

the PIE++ dataset. Along with this, we provided the intent and reason annotations to 5 human subjects who were asked to watch each video sequence twice. We then asked the subjects to rate the appropriateness of annotations in PIE++ for the given road scene on a 3-point scale: *Yes*, *Confusing* and *No*. The specific questions for each subject included: (1) Are the provided reasons complete for the (No-)cross intent (*completeness*)?; (2) Would the reasons be useful if this was provided to you while driving (*usefulness*)?; (3) How correct are the reasons for the pedestrian's perceived intent (*correctness*)?; (4) If you were the driver of a car with such predictions, would this improve your trust in the model prediction (*trust*)?; (5) How reliable is the provided reason for the prediction (*reliability*)? Table IV summarizes the results of our study, and shows strong support for the reason annotations included in PIE++ across the considered factors.

We note that our work does not change any pre-existing definitions of pedestrian intent defined in [1]. Also, following [1], we consider the intent only up to a critical point (i.e. frames preceding the crossing point). PIE++ consists of 4950 multi-label reason annotations for 1410 pedestrians with crossing intent and 657 multi-label reason annotations for 432 pedestrians with no-crossing intent.

IV. MINDREAD: OUR PROPOSED METHODOLOGY

We propose a multi-task learning architecture with novel components, which we call MINDREAD (cross-Modal representation learning model for REasoning peDestrian-intent), to predict pedestrian intent as well as the reasons behind the intent. Our overall architecture, which seeks to learn a cross-modal representation for the aforementioned objective, is shown in Fig 2. The cross-modal representation is learned from vision and language modalities, where we use an attention mechanism to fuse vision-based spatio-temporal features and reason text embeddings.

Factors	Yes	Confusing	No
Completeness	100%	-	-
Usefulness	98%	1%	1%
Correctness	99%	1%	-
Trust	98%	2%	-
Reliability	97%	2%	1%

TABLE IV

RESULTS FROM OUR USER STUDY ON THE UTILITY OF REASON ANNOTATIONS IN PIE++

In this section, we discuss architecture details of our MINDREAD framework. Given a video dataset (e.g. a driving dataset containing 2-D RGB traffic video clips), let $[F_1, F_2, \dots, F_t]$ denote t time steps of past observations (image frames) in a sequence. Our goal is to detect the probability of a pedestrian's crossing intent $I_t \in \{0, 1\}$ ($0 = NC, 1 = C$) as well as the multi-label textual reasons $E_t \in \{0, 1\}^n$ (Tables II and III) that support the intent prediction at time t ; (n represents the total reason classes). MINDREAD consists of three modules: Correlated Semantic Explanation Affinity (CSEA), Transformer-based Feature Encoding (TFE) and an attention-based cross-modal representation learning module. The CSEA module receives the textual embeddings of the reasons, and represents them as a learned correlated semantic embedding matrix based on their co-occurrence in the dataset. The TFE module extracts and encodes spatio-temporal features from video frames using a fusion of local features, global features and bounding box information of pedestrians. An attention mechanism brings together the outputs of the TFE and CSEA modules to provide our final cross-modal representation. These cross-modal features are finally fed to two classifier heads that output reason and pedestrian intent predictions (as shown in Fig 2). We explain each module of the proposed method in the following subsections.

A. Correlated Semantic Explanation Affinity

We model the list of textual reasons as a directed graph where each node is a reason (rather, its text embedding), and the edges denote their co-occurrence in the dataset. We note that since each pedestrian intent instance is annotated with reasons in a multi-label manner, multiple reasons are plausible for a given intent, thus making their co-occurrence a rich piece of input information to the model. While some reasons may or may not appear together, some of them can never appear together. For e.g., reasons like "pedestrian neglects ego-vehicle" can never appear with "pedestrians doing their work on the road-side". On the other hand, reasons like "pedestrian waiting since ego-vehicle speed is high" have a higher chance of appearing with "pedestrian waiting for a safe passage". Such semantic explanation affinity modeling helps improve reason prediction thereby improving intent prediction. We use the widely used Sentence-BERT [39] to generate textual embeddings for the graph nodes. To obtain the co-occurrence of reasons within the dataset, we compute the adjacency matrix using the conditional probability $P(E_j|E_i)$, which denotes the probability of occurrence of reason E_j when reason E_i appears. Given the above explanation node embeddings $X_I \in R^{n \times d}$ (where

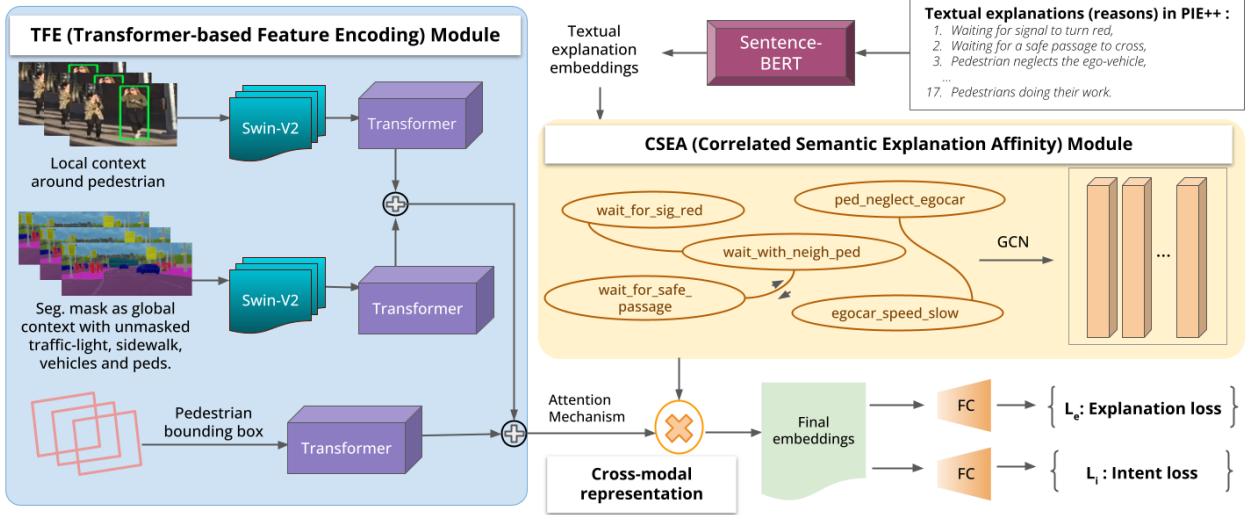


Fig. 2. Overall architecture of MINDREAD (cross-Modal representation learning model for REasoning peDestrian-intent)

n is the number of reasons and d is the node embedding dimensionality) and adjacency matrix $A \in \mathbb{R}^{n \times n}$, we use a Graph Convolution Network (GCN) [44] to learn the final reason embeddings. Every GCN layer can be written as a non-linear function $X_{l+1} = f(X_l, A)$, which takes the textual reason representations from the previous layer (X_l) as inputs and outputs new correlated reason representations X_{l+1} . For the first layer, the input is an $\mathbb{R}^{n \times d}$ matrix, where d is the dimensionality of the reason-level sentence embedding. For the last layer, we choose output of dimension $\mathbb{R}^{n \times D}$ where D represents the dimensionality of the output of the TFE module with visual spatio-temporal features. This gives the final learned correlated semantic explanation embeddings (output of the CSEA module).

B. Transformer-based Feature Encoding

Following earlier work in intent prediction, the visual spatiotemporal features in MINDREAD are obtained using a combination of local visual context, global visual context and pedestrian bounding box coordinates in the scene (as shown in Fig 2). The local context consists of a region around the target pedestrian across the given video sequence (we use a small square patch around the pedestrian in this work). Unlike previous efforts that use pedestrian information alone, we note that global context is useful as it offers visual features that account for socio-environmental interactions. Our global context includes the pixels corresponding to key socio-environmental information in the scene such as traffic-light, vehicles, crosswalks, and other pedestrians (obtained from a segmented scene with masking).

In order to obtain visual spatio-temporal features from the video frames, we use transformer models [40] (in contrast to earlier work that use older architectures such as VGG [36] and GRU [37] to extract spatial and temporal features). We extract both local and global vision features using two Swin-V2 transformers [38] pre-trained on the ImageNet-22K [45] dataset. The final dense layer of the pre-trained model is replaced by a global average pooling layer to extract visual features. Similar to [46], we use a standard transformer

model for temporal encoding of the features obtained from the Swin-V2 transformers. The transformer encoder consists of alternating layers of multiheaded self-attention (MSA) and multi-layer perceptron (MLP) blocks. LayerNorm is applied before every block, and residual connections after every block following [47], [48]. The MLP contains two layers with a GELU non-linearity. Before feeding the sequence features into the the transformer MSA block, they are combined with positional embedding to capture the sequential properties of input sequence features. As in Fig 2, both local and global context features are encoded by extracting spatial features from the Swin-V2 transformer followed by temporal features using the Transformer. Bounding box features are also encoded similarly. We concatenate these three kinds of feature to get the final fused spatiotemporal features.

C. Attention-based Cross-modal Representations

The final module of our framework, an attention-based mechanism to fuse the outputs of the TFE and CSEA modules, helps leverage the reason relationships to better predict pedestrian intent. This is a key element of our framework. The TFE output features are fed to an attention block for selectively focusing on the appropriate feature information for our task. The temporal sequence of features (output of transformer-based encoding from the TFE module) are represented as hidden states $h = h_1, h_2, \dots, h_e$, where e represents the end hidden state. Following [49], the attention weight is computed as: $\alpha = \frac{\exp(\text{score}(h_e, \tilde{h}_s))}{\sum_s \exp(\text{score}(h_e, \tilde{h}_s))}$, where $\text{score}(h_e, \tilde{h}_s) = h_e^T W_s \tilde{h}_s$ and W_s is a learned weight matrix. Such an attention mechanism provides information on how the current video frame (i.e. the end hidden state h_e) should leverage the immediately preceding frames given in \tilde{h}_s . The final output of the attention module is obtained as $F = \tanh(W_c[h_c; h_e])$, where W_c is a weight matrix, h_c is the sum of all attention-weighted hidden states given as $h_c = \sum_s \alpha \tilde{h}_s$ and $F \in \mathbb{R}^{B \times 3076}$, where B represent batch size. Finally, our cross-modal representation is given by, $C = F \cdot X^T$, where F represents final attention-based feature representations and

Method	Intent Accuracy	F1-score	Precision	AUC
PIE _{int} (ICCV' 19) [1]	79.0 %	87.0 %	88.0 %	79.0 %
STIP (ICRA' 20) [9]	80.0 %	88.0 %	91.0 %	81.0 %
CIA (IJCAI' 21) [2]	82.0 %	88.0 %	94.0 %	83.0 %
MINDREAD (Ours)	87.6 ± 0.1%	95.0 ± 0.1%	96.0 ± 0.1%	89.0 ± 0.2%

TABLE V

PEDESTRIAN INTENT PREDICTION ON PIE++ (SAME AS PIE FOR THIS TASK)

Method	Reason Accuracy	Intent Accuracy	F1-score	Precision	AUC
PIE _{int} *	48.0 ± 0.2%	79.5 ± 0.1%	87.0 ± 0.1%	88.0 ± 0.1%	79.0 ± 0.2%
STIP*	54.0 ± 0.1%	80.7 ± 0.1%	88.0 ± 0.2%	91.0 ± 0.1%	81.0 ± 0.1%
CIA*	63.0 ± 0.2%	82.8 ± 0.1%	88.0 ± 0.3%	94.0 ± 0.2%	83.0 ± 0.1%
MINDREAD (Ours)	72.4 ± 0.2%	87.6 ± 0.1%	95.0 ± 0.1%	96.0 ± 0.1%	89.0 ± 0.2%

TABLE VI

PEDESTRIAN INTENT AND REASON PREDICTION RESULTS ON PIE++ DATASET. METHODS WITH * INDICATE AN ADAPTATION OF EXISTING BASELINES WITH A REASON-BASED CLASSIFIER HEAD FOR A FAIR COMPARISON.

Method	Reason Acc	Intent Acc
PV-LSTM*	85.3 ± 0.1 %	91.6 ± 0.1 %
MINDREAD (Ours)	91.0 ± 0.1 %	95.4 ± 0.1 %

TABLE VII

PEDESTRIAN INTENT AND REASON PREDICTION ON JAAD. METHODS WITH * INDICATE ADAPTATION OF EXISTING BASELINES WITH A REASON-BASED CLASSIFIER HEAD FOR FAIR COMPARISON.

Method	Reason Acc	Intent Acc	F1
MINDREAD_w/o_CM	64.0 ± 0.1%	83.4 ± 0.1%	90.0%
MINDREAD (Ours)	72.4 ± 0.2%	87.6 ± 0.1%	95.0%

TABLE VIII

MINDREAD WITH AND WITHOUT CROSS-MODALITY

X represents the learned correlated semantic explanation embeddings. These final cross-modal encodings are fed to two classifiers giving textual explanation E_t and intent I_t prediction outputs. We use standard Binary Cross-Entropy (BCE) losses, L_e for explanation and L_i for intent, to learn the final parameters of our model.

V. EXPERIMENTS AND RESULTS

We evaluate the effectiveness of our method for intent and reason prediction on PIE++. To study the generalizability of this approach, we also annotate the JAAD [16] dataset with reason explanations, although this dataset has very short sequences with limited variety in user behavior. We first present our results on the pedestrian intent prediction task, where we compare our method's performance with state-of-the-art methods, followed by results for the novel problem of reasoning the pedestrian intent. Following earlier methods in this space, we measure performance using intent accuracy, explanation accuracy, and F1 score. Qualitative results on video sequences are provided in the supplementary video.

We compare our method with well-known methods for pedestrian intent prediction: PIE_{int} [1], STIP [9] and CIA [2] on PIE++, and PV-LSTM [50] on the JAAD dataset. Table V shows the intent prediction results on PIE++ dataset. MINDREAD outperforms the state-of-the-art for intent prediction, with an improvement in accuracy and F1-score by 5.6% and 7% respectively. Table IX shows results on JAAD. Once again, our method outperforms all earlier works with an improvement of 4.4% over the state-of-the-art on intent prediction accuracy.

Textual Emb.	Reason Acc	Intent Acc	F1
GloVe	68.2 ± 0.1%	86.4 ± 0.1%	93.0 %
Sentence-BERT	72.4 ± 0.2%	87.6 ± 0.1%	95.0 %

TABLE X

STUDY COMPARING WORD-LEVEL (GLOVE) VS SENTENCE-LEVEL (SENTENCE-BERT) MODELS FOR REASON EMBEDDINGS

Backbone	Exp. Acc	Intent Acc	F1
VGG + GRU	69.2 ± 0.1%	87.0 ± 0.1%	94.0%
Swin-V2 + Transformer	72.4 ± 0.2%	87.6 ± 0.1%	95.0%

TABLE XI

ABLATION USING VGG + GRU VERSUS NOVEL SWIN-V2 + TRANSFORMER BACKBONE FOR SPATIO-TEMPORAL VISUAL FEATURE MODELING.

A. Results on Reasoning Pedestrian Intent

Since there is no existing method for simultaneous intent and reason prediction, we adapt existing methods by adding an explicit reason head and retraining them on PIE++ for fair comparison. We denote the corresponding variants of these methods as PIE_{int}*, STIP*, CIA* for the PIE++ dataset, and PV-LSTM* for the JAAD dataset (enriched with our reason annotations). Table VI shows pedestrian intent and reason prediction results on PIE++. MINDREAD

Method	Intent Acc
PV-LSTM [50]	91.0 %
MINDREAD (Ours)	95.4 %

TABLE IX

INTENT PREDICTION ON JAAD outperforms all the baseline methods, including state-of-the-art method CIA* with significant improvement of 9.4% in reason prediction accuracy, 4.8% in intent prediction accuracy and 7% in intent F1-score. We hypothesize that this improvement in performance is due to our cross-modal representation learning as is also evident from our ablation studies (Sec VI Table VIII). Table VII shows results of similar experiments on the JAAD dataset. MINDREAD improves over the state-of-the-art method by 5.7% and 3.8% improvement in reason and intent accuracy. For JAAD, we use the same train-test split as [14] and F1-score, precision, and AUC are 97%, 95% and 91%.

VI. ANALYSIS AND DISCUSSION

Word vs. Sentence-level Embeddings. We studied the relevance of different text embedding methods, as shown in

Method	Runtime	Intent Acc
PV-LSTM [50]	4.7 ms	91.0 %
MINDREAD	4.3 ms	95.4 %

TABLE XII
RUN TIME ANALYSIS ON JAAD DATASET

γ_R	γ_I	Reason Acc	Intent Acc
0.5	1.0	57.4 %	81.2 %
1.0	0.5	63.6 %	74.1 %
1.0	1.0	72.4 %	87.6 %

TABLE XIII

RESULTS USING DIFFERENT WEIGHTS; γ_R = WEIGHT FOR REASON LOSS;
 γ_I = WEIGHT FOR INTENT LOSS

the results in Table X. Specifically, we compare word vs sentence embedding models. For word-level embedding, we use a 300-dim Glove [51] model trained on the Wikipedia dataset. We consider reason-specific words and average the embeddings. For sentence-level embeddings, we use Sentence-BERT [39] (as in our method). The results show that sentence-level embeddings outperform word-level embeddings. This is expected since the context of the words in the reason matter, which is considered in sentence-level embeddings.

Effect of Different Spatio-temporal Backbones. Table XI shows the results of our study using different spatio-temporal visual feature extraction backbones in MINDREAD. As in Sec IV, we use contemporary transformer models instead of older models used in earlier work. Specifically we use Swin-V2-L + Transformer backbone instead of the VGG + GRU backbone. In order to study this choice, we used the VGG +GRU backbone in MINDREAD and obtained the results. Table XI shows that transformer-based models outperform VGG + GRU models by 3.2% in reason prediction accuracy and 0.6% in intent prediction accuracy.

MINDREAD without Cross-modality. In order to understand the usefulness of the cross-modal representation learning framework proposed in MINDREAD, we conducted experiments to study whether visual features alone can address the new task with an added reason classification head (in addition to the intent prediction head). As shown in Table VIII, we can see performance drop by 4.2% and 5% in intent accuracy and F1 using MINDREAD without cross-modality. We observe that the use of reason embeddings and cross-modal representation learning plays a significant role for reason prediction accuracy in particular, and also helps improve intent prediction (as we hypothesized).

Loss, Metrics and Runtime Analysis. Our main results used a weight of 1 for each loss term. Our initial studies, reported in Tab XIII, supported this choice. Weighting one loss term more than the other did not increase performance. Tab XII shows the runtimes, compared with those reported on JAAD. We note that our model is faster than SOTA PV-LSTM [50] while achieving higher accuracy. Runtime for our model on PIE++ was 4.1 ms, of the same range as JAAD. It’s due to the significance of *learned* reason embeddings as well as spatiotemporal features that speed up inference in real time.

VII. CONCLUSIONS AND FUTURE WORK

We introduced a novel problem of reasoning pedestrian intent. To address this problem, we made two significant contributions: (1) We proposed a reason-enriched PIE++ dataset that contains explanations/textual reasons for pedestrian intent prediction; and (2) We proposed a multi-task formulation called MINDREAD based on cross-modal representation learning that obtains strong results on this new problem setting. Our comprehensive experiments show we achieve significantly more accurate and reasonable predictions than prior works, and that such explanations improve trust in users of such systems. Exploring uncertainty estimates for reasons and collision avoidance are interesting future directions of our work.

REFERENCES

- [1] A. Rasouli, I. Kotscheruba, T. Kunic, and J. K. Tsotsos, “Pie: A large-scale dataset and models for pedestrian intention estimation and trajectory prediction,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 6262–6271.
- [2] Y. Yao, E. Atkins, M. J. Roberson, R. Vasudevan, and X. Du, “Coupling intent and action for pedestrian crossing behavior prediction,” *International Joint Conferences on Artificial Intelligence (IJCAI)*, 2021.
- [3] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The Cityscapes Dataset for Semantic Urban Scene Understanding,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [4] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell, “BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [5] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 3354–3361.
- [6] A. Reyes-Muñoz and J. Guerrero-Ibáñez, “Vulnerable road users and connected autonomous vehicles interaction: A survey,” *Sensors*, vol. 22, no. 12, p. 4614, 2022.
- [7] “Vulnerable Road Users, Position/Policy Statement, NATIONAL SAFETY COUNCIL,” <https://www.nsc.org/getattachment/d5babe66-582d-4e66-804f-8d06f9b021a4/t-vulnerable-road-users-147>, 2022, [Online; accessed 10-November-2022].
- [8] “VULNERABLE ROAD USER (VRU) PROTECTION,” <https://www.euroncap.com/en/vehicle-safety/the-ratings-explained/vulnerable-road-user-vru-protection/>, 2022, [Online; accessed 10-November-2022].
- [9] B. Liu, E. Adeli, Z. Cao, K.-H. Lee, A. Sheno, A. Gaidon, and J. C. Niebles, “Spatiotemporal relationship reasoning for pedestrian intent prediction,” *IEEE Robotics and Automation Letters (R-AL)*, vol. 5, no. 2, pp. 3485–3492, 2020.
- [10] J. Lorenzo, I. Parra, and M. Sotelo, “Intformer: Predicting pedestrian intention with the aid of the transformer architecture,” *arXiv preprint arXiv:2105.08647*, 2021.
- [11] F. Piccoli, R. Balakrishnan, M. J. Perez, M. Sachdeo, C. Nunez, M. Tang, K. Andreasson, K. Bjurek, R. D. Raj, E. Davidsson, *et al.*, “Fussi-net: Fusion of spatio-temporal skeletons for intention prediction network,” in *2020 54th Asilomar Conference on Signals, Systems, and Computers*. IEEE, 2020, pp. 68–72.
- [12] S. Neogi, M. Hoy, K. Dang, H. Yu, and J. Dauwels, “Context model for pedestrian intention prediction using factored latent-dynamic conditional random fields,” *IEEE Transactions on Intelligent Transportation Systems (T-ITS)*, 2020.
- [13] K. Saleh, M. Hossny, and S. Nahavandi, “Real-time intent prediction of pedestrians for autonomous ground vehicles via spatio-temporal densenet,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 9704–9710.

- [14] —, “Intent prediction of vulnerable road users from motion trajectories using stacked lstm network,” in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, 2017, pp. 327–332.
- [15] J. D. Westaby, “Behavioral reasoning theory: Identifying new linkages underlying intentions and behavior,” *Organizational Behavior and Human Decision Processes*, vol. 98, no. 2, pp. 97–120, 2005.
- [16] A. Rasouli, I. Kotseruba, and J. K. Tsotsos, “Are they going to cross? a benchmark dataset and baseline for pedestrian crosswalk behavior,” in *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCV-W)*, 2017, pp. 206–213.
- [17] A. Rasouli, M. Rohani, and J. Luo, “Bifold and semantic reasoning for pedestrian behavior prediction,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 15 600–15 610.
- [18] A. Rasouli, I. Kotseruba, and J. K. Tsotsos, “Pedestrian action anticipation using contextual feature fusion in stacked rnns,” *arXiv preprint arXiv:2005.06582*, 2020.
- [19] T. Yau, S. Malekmohammadi, A. Rasouli, P. Lakner, M. Rohani, and J. Luo, “Graph-sim: A graph-based spatiotemporal interaction modelling for pedestrian action prediction,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 8580–8586.
- [20] I. Kotseruba, A. Rasouli, and J. K. Tsotsos, “Benchmark for evaluating pedestrian action prediction,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2021, pp. 1258–1268.
- [21] X. Zhai, Z. Hu, D. Yang, L. Zhou, and J. Liu, “Social aware multi-modal pedestrian crossing behavior prediction,” in *Proceedings of the Asian Conference on Computer Vision*, 2022, pp. 4428–4443.
- [22] A. Rasouli, T. Yau, M. Rohani, and J. Luo, “Multi-modal hybrid architecture for pedestrian action prediction,” in *2022 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2022, pp. 91–97.
- [23] R. Vedantam, S. Bengio, K. Murphy, D. Parikh, and G. Chechik, “Context-aware captions from context-agnostic supervision,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 251–260.
- [24] L. A. Hendricks, R. Hu, T. Darrell, and Z. Akata, “Grounding visual explanations,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 264–279.
- [25] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, “Show, attend and tell: Neural image caption generation with visual attention,” in *International Conference on Machine Learning (ICML)*. PMLR, 2015, pp. 2048–2057.
- [26] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [27] J. Kim, A. Rohrbach, T. Darrell, J. Canny, and Z. Akata, “Textual explanations for self-driving vehicles,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 563–578.
- [28] T. Chen, R. Tian, Y. Chen, J. Domeyer, H. Toyoda, R. Sherony, T. Jing, and Z. Ding, “Psi: A pedestrian behavior dataset for socially intelligent autonomous car,” *arXiv preprint arXiv:2112.02604*, 2021.
- [29] Y. Ge, Y. Xiao, Z. Xu, M. Zheng, S. Karanam, T. Chen, L. Itti, and Z. Wu, “A peek into the reasoning of neural networks: Interpreting with structural visual concepts,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 2195–2204.
- [30] W. Bao, Q. Yu, and Y. Kong, “Drive: Deep reinforced accident anticipation with visual explanation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 7619–7628.
- [31] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, “Learning transferable visual models from natural language supervision,” in *International Conference on Machine Learning (ICML)*. PMLR, 2021, pp. 8748–8763.
- [32] X. Yuan, Z. Lin, J. Kuen, J. Zhang, Y. Wang, M. Maire, A. Kale, and B. Faieta, “Multimodal contrastive training for visual representation learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 6995–7004.
- [33] C. Choi, J. H. Choi, J. Li, and S. Malla, “Shared cross-modal trajectory prediction for autonomous driving,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 244–253.
- [34] K. Kim, Y. K. Lee, H. Ahn, S. Hahn, and S. Oh, “Pedestrian intention prediction for autonomous driving using a multiple stakeholder perspective model,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 7957–7962.
- [35] Z. Sui, Y. Zhou, X. Zhao, A. Chen, and Y. Ni, “Joint intention and trajectory prediction based on transformer,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 7082–7088.
- [36] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *CoRR*, vol. abs/1409.1556, 2015.
- [37] K. Cho, B. van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning phrase representations using rnn encoder-decoder for statistical machine translation,” in *Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [38] Z. Liu, H. Hu, Y. Lin, Z. Yao, Z. Xie, Y. Wei, J. Ning, Y. Cao, Z. Zhang, L. Dong, F. Wei, and B. Guo, “Swin transformer v2: Scaling up capacity and resolution,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 12 009–12 019.
- [39] N. Reimers and I. Gurevych, “Sentence-bert: Sentence embeddings using siamese bert-networks,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 3982–3992.
- [40] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [41] W. Kim, M. S. Ramanagopal, C. Barto, M.-Y. Yu, K. Rosaen, N. Goumas, R. Vasudevan, and M. Johnson-Roberson, “Pedx: Benchmark dataset for metric 3-d pose estimation of pedestrians in complex urban intersections,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 4, no. 2, pp. 1940–1947, 2019.
- [42] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nusscenes: A multimodal dataset for autonomous driving,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, 2020, pp. 11 621–11 631.
- [43] P. E. Shrout and J. L. Fleiss, “Intraclass correlations: uses in assessing rater reliability,” *Psychological bulletin*, vol. 86, no. 2, p. 420, 1979.
- [44] M. Welling and T. N. Kipf, “Semi-supervised classification with graph convolutional networks,” in *J. International Conference on Learning Representations (ICLR)*, 2016.
- [45] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, *et al.*, “Imagenet large scale visual recognition challenge,” *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [46] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [47] Q. Wang, B. Li, T. Xiao, J. Zhu, C. Li, D. F. Wong, and L. S. Chao, “Learning deep transformer models for machine translation,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019, pp. 1810–1822.
- [48] A. Baevski and M. Auli, “Adaptive input representations for neural language modeling,” in *International Conference on Learning Representations (ICLR)*, 2018, pp. 1810–1822.
- [49] M.-T. Luong, H. Pham, and C. D. Manning, “Effective approaches to attention-based neural machine translation,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2015, pp. 1412–1421.
- [50] S. A. Bouhsain, S. Saadatnejad, and A. Alahi, “Pedestrian intention prediction: A multi-task perspective,” *arXiv preprint arXiv:2010.10270*, 2020.
- [51] J. Pennington, R. Socher, and C. D. Manning, “Glove: Global vectors for word representation,” in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.