

MMUTF: Multimodal Multimedia Event Argument Extraction with Unified Template Filling

Philipp Seeberger, Dominik Wagner, Korbinian Riedhammer

Technische Hochschule Nürnberg Georg Simon Ohm

{philipp.seeberger, dominik.wagner, korbinian.riedhammer}@th-nuernberg.de

Abstract

With the advancement of multimedia technologies, news documents and user-generated content are often represented as multiple modalities, making Multimedia Event Extraction (MEE) an increasingly important challenge. However, recent MEE methods employ weak alignment strategies and data augmentation with simple classification models, which ignore the capabilities of natural language-formulated event templates for the challenging Event Argument Extraction (EAE) task. In this work, we focus on EAE and address this issue by introducing a unified template filling model that connects the textual and visual modalities via textual prompts. This approach enables the exploitation of cross-ontology transfer and the incorporation of event-specific semantics. Experiments on the M2E2 benchmark demonstrate the effectiveness of our approach. Our system surpasses the current SOTA on textual EAE by +7% F1, and performs generally better than the second-best systems for multimedia EAE.

1 Introduction

Traditional Event Extraction (EE) consists of two tasks: Event Detection (ED) and Event Argument Extraction (EAE), which address event type identification and argument role extraction, respectively. Traditional methods focus on individual modalities, such as text, images, or videos (Huang et al., 2023; Yatskar et al., 2016; Khan et al., 2022). However, focusing on single modalities results in incomplete event understanding, as up to 33% of visual arguments are absent in the accompanying text (Li et al., 2020) (cf. Figure 1).

As the creation of multimodal EE datasets is time-consuming and costly, recent advanced MEE methods focus on (weak) image-text alignment based on available unimodal task-dependent and task-independent vision-language datasets (Li et al., 2020; Liu et al., 2022; Du et al., 2023; Liu et al., 2024). However, current multimedia EAE models

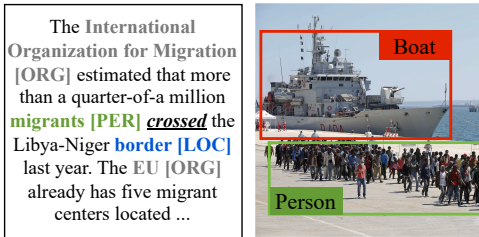
<p>The International Organization for Migration [ORG] estimated that more than a quarter-of-a million migrants [PER] <u>crossed</u> the Libya-Niger border [LOC] last year. The EU [ORG] already has five migrant centers located ...</p>		
		
Movement:Transport (<u>crossed</u>)		
Visual Arguments	Artifact	Person
	Vehicle	Boat
Textual Arguments	Artifact	Migrants
	Destination	Border

Figure 1: Example of Multimedia Event Extraction from the M2E2 benchmark, where event argument roles are extracted from both the textual and visual modality.

are still based on simple classification techniques while ignoring cross-ontology transfer capabilities and event template semantics.

To address this gap, we reformulate existing query-based (template filling) methods (Wang et al., 2022; Zhang et al., 2022; Ma et al., 2022; Zheng et al., 2023) into a unified template filling framework, which enables us to utilize event templates as natural language prompts for different input modalities. First, our **Multimodal Multimedia Event Argument Extraction with Unified Template Filling (MMUTF)** model exploits candidate structures, such as textual entities and visual objects, and connects them via query representations (i.e., argument roles) in a unified latent space. Finally, these representations are used to match event argument roles with the corresponding candidates via queries extracted from event templates.

Contributions We summarize our main contributions as follows: 1) We introduce MMUTF, a unified template filling framework that addresses the EAE task for multiple input modalities. 2) We

validate our approach on the widely used M2E2 benchmark and demonstrate superior performance compared to the majority of baselines. 3) We analyse the transfer learning capabilities for multimedia EAE by employing FrameNet as a resource-rich data source.

2 Method

For event argument roles, the extraction task can be framed as a candidate-query matching problem, where argument roles serve as the queries. We can utilize the matching between candidates and queries to address the multimedia EAE task with a unified template filling model. Specifically, we aim to create candidate representations for either textual entities or image objects and compute a matching score for each query (i.e., argument role) in a given event template.¹

2.1 Multimodal Unified Template Filling

Our unified template filling model is depicted in Figure 2. We use both the sentences and images from a multimedia document as input contexts, with the written descriptions of event templates and argument roles acting as placeholders. We approach the EAE task in multiple steps: 1) Given a predicted event trigger, we select the corresponding event template. 2) We encode the input sentences and images with modality-specific encoder models and compute candidate representations. 3) We feed the event template into the query model and obtain query representations for each role. 4) Lastly, we compute the matching scores and treat the EAE task as a binary classification problem.

Candidate Representations In order to obtain modality-specific candidates, we recognize entities and objects with off-the-shelf systems and extract contextualized representations using a Transformer model for each modality (Vaswani et al., 2017; Raffel et al., 2020; Radford et al., 2021). The encoding process yields a sequence of N token embeddings $t = \{t_i\}_1^N \in \mathbb{R}^{N \times H}$ and M patch embeddings $p = \{p_i\}_1^M \in \mathbb{R}^{M \times H}$ with dimensionality H for each input sentence S and image I , respectively. For entities tokenized into multiple subwords, we use mean pooling and concatenate the feature vector with the representation of the

event trigger word. Similarly, for images, we obtain object feature vectors using the max pooling² operation over all patches containing the detected object, and concatenate these with the image CLS embeddings for further processing. Here, the CLS embeddings serve as event triggers represented as image scenes. The concatenated vectors are then transformed via their corresponding mapping networks, which will be described in more detail later.

Query Representations For event template processing, we employ the decoder of a Transformer Encoder-Decoder model (Raffel et al., 2020) as shown in Figure 2. This approach has shown promising results without compromising on the length of the input context, as seen in encoder architectures (Ma et al., 2022). This configuration enables the prompt to interact with the input contexts at the cross-attention layers of every decoder block, resulting in context-dependent query representations. Similar to the generation of candidate representations, we use the mean pooling operation for argument roles that span multiple subwords and feed the features into a subsequent mapping network. We denote the query representations as $q = \{q_i\}_1^R \in \mathbb{R}^{R \times H}$, where H is the dimensionality, and R is the number of argument roles.

Mapping Network We aim to align the modality-specific candidate and query representations into a unified latent space, based on the candidate and query feature vectors. However, during training, we face the challenge of unifying the rich and diverse representations of the different modalities and Transformer models. The mapping networks are designed to bridge the gap between the latent spaces and transform all candidate and query features into H -sized representations. Each mapping network shares the same architecture and consists of a two-layer feedforward network with $4 \times H$ hidden units, ReLU (Agarap, 2018) activation, and is trained with a dropout probability of 40%.

2.2 Training and Inference

For each event, we evaluate whether a candidate c represents an argument role r by computing a matching score $\phi(c, r) = \sigma(\mathbf{h}_c^T \mathbf{q}_r) \in \mathbb{R}$. Here, \mathbf{h}_c denotes the representation vector of the candidate c , \mathbf{q}_r indicates the query vector of argument role r , and σ represents the sigmoid activation. The model parameters and representations are optimized to

¹We can also consider text spans, image segments, etc., as candidates. However, entities and objects follow the present annotation schemes and suit the EAE task best.

²Max pooling achieved the best results among other techniques such as mean and RoI pooling.

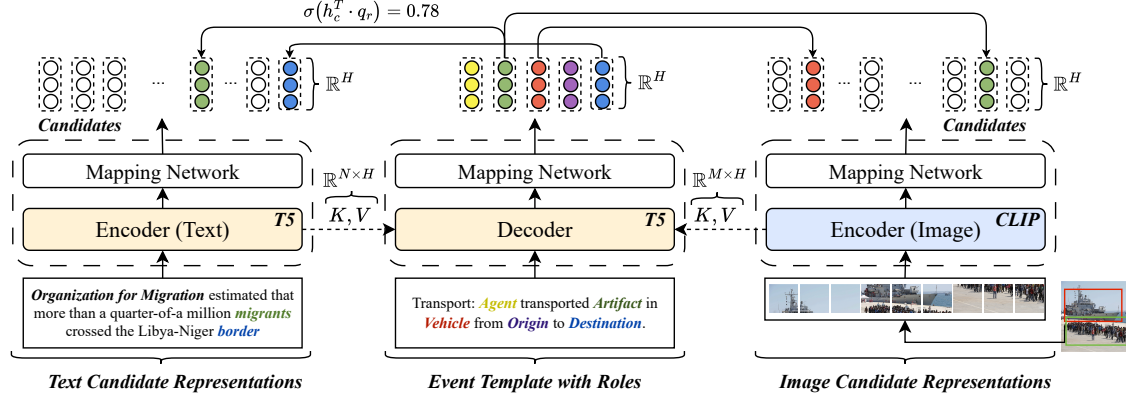


Figure 2: The overall architecture of MMUTF. Given a text or image event, our model encodes the corresponding event template and argument roles as a textual prompt. The textual or image context interacts with the prompt via the cross-attention mechanism, resulting into candidate and query representation vectors. Here, the candidates correspond to entities or objects, depending on the modality. A matching score then assigns the candidates to the argument roles based on a predefined threshold.

enhance the matching scores for positive and to reduce it for negative candidate-query pairs:

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{|C|} \sum_{i=1}^{|C|} \sum_{j=1}^{|A|} y_{ij} \log \phi(i, j) \quad (1)$$

where \mathcal{L}_{BCE} denotes the binary cross-entropy loss, y_{ij} the ground truth label, C the collection of candidates and A the set of possible argument roles. At inference time, we use a threshold τ with $\phi(c, r) \geq \tau$ to decide whether candidate c corresponds to argument role r . For each event and candidate, we assign the role with the highest score that exceeds the threshold τ . Otherwise, the candidate is assigned to the *None* class.

3 Experiments

Datasets We conduct our experiments on the widely used M2E2 multimedia event extraction benchmark (Li et al., 2020) with 8 event types and 15 argument roles. The dataset contains 245 multimedia documents with 6,167 sentences and 1,014 images. It encompasses 1,297 textual events and 391 visual events, with 192 textual events and 203 visual events aligning to 309 multimedia events (cf. Figure 1). Since the M2E2 dataset does not provide a training split, we follow previous work (Li et al., 2020; Liu et al., 2022; Du et al., 2023) and employ ACE2005 and SWiG for the training phase. ACE2005 (Walker, Christopher et al., 2006) consists of 33 textual event types and 36 argument roles. SWiG (Pratt et al., 2020) annotates visual events with 504 activity verbs and 1,788 semantic

roles. Following Li et al. (2020), we align both datasets with the M2E2 event ontology.

Metrics For evaluation, we follow previous studies on EAE (Huang et al., 2023) and use Precision (P), Recall (R), and F1-Score (F1) as our evaluation metrics. In this work, we focus on the more challenging EAE task but also report event classification results for completeness.

Baselines We compare our proposed approach with a wide range of state-of-the-art models. These models include FLAT (Li et al., 2020), WASE (Li et al., 2020), CLIP_{EVENT} (Li et al., 2022), UNICL (Liu et al., 2022), CAMEL (Du et al., 2023), MGIM (Liu et al., 2024), and UMIE (Sun et al., 2024). The majority of baselines set the focus on weak image-text alignment but neglect more sophisticated classification approaches w.r.t. EAE. We provide a more comprehensive overview in A.1.

Experimental Setup To ensure a fair comparison with the baselines, we utilize the CLIP-base model (Radford et al., 2021) as visual encoder with a 16×16 patch size, and the T5-base (Raffel et al., 2020) model as textual encoder-decoder. We generate object candidates with YOLOv8 (Jocher et al., 2023) trained on the COCO (Lin et al., 2014) dataset and use entities as textual span candidates.

Our system is trained jointly for 5 epochs on the visual task and 25 epochs on the textual task with a batch size of 16 and 32, respectively. The varying number of epochs reduces the imbalance of modalities in the training datasets. For optimization, we

Model	Textual Events						Visual Events						Multimedia Events					
	Event Mention			Argument Role			Event Mention			Argument Role			Event Mention			Argument Role		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1	P	R	F1
FLAT _{ATT}	34.2	63.2	44.4	20.1	27.1	23.1	27.1	57.3	36.7	4.3	8.9	5.8	33.9	59.8	42.2	12.9	17.6	14.9
FLAT _{OBJ}	38.3	57.9	46.1	21.8	26.6	24.0	26.4	55.8	35.8	9.1	6.5	7.6	34.1	56.4	42.5	16.3	15.9	16.1
WASE _{ATT}	37.6	66.8	48.1	27.5	33.2	30.1	32.3	63.4	42.8	9.7	11.1	10.3	38.2	67.1	49.1	18.6	21.6	19.9
WASE _{OBJ}	42.8	61.9	50.6	23.5	30.3	26.4	43.1	59.2	49.9	14.5	10.1	11.9	43.0	62.1	50.8	19.5	18.9	19.2
CLIP _{EVENT}	-	-	-	-	-	-	41.3	72.8	52.7	21.1	13.1	17.1	-	-	-	-	-	-
UNI _{CL}	49.1	59.2	53.7	27.8	34.3	30.7	54.6	60.9	57.6	16.9	13.8	15.2	44.1	67.7	53.4	24.3	22.6	23.4
MGIM	50.1	66.5	55.8	28.2	34.7	<u>31.2</u>	55.7	64.4	58.5	24.1	14.1	17.8	46.3	69.6	55.6	25.2	21.7	24.6
UMIE*	-	-	-	-	-	-	-	-	-	-	-	-	-	-	62.1	-	-	24.5
CAMEL	45.1	71.8	55.4	24.8	41.8	31.1	52.1	66.8	58.5	21.4	28.4	24.4	55.6	59.5	57.5	31.4	35.1	33.2
CAMEL [†]	48.5	65.0	55.5	30.3	33.6	31.9	55.1	59.1	57.0	21.7	22.1	21.9	47.9	63.4	54.6	25.9	30.4	27.9
MMUTF	48.5	65.0	55.5	33.6	44.2	38.2	55.1	59.1	57.0	23.6	18.8	<u>20.9</u>	47.9	63.4	54.6	39.9	20.8	<u>27.4</u>

Table 1: Comparison with SOTA methods on the M2E2 benchmark for textual, visual, and multimedia events. **Bold** numbers indicate the best EAE model whereas underlined metrics denote the second best. [†] denotes our reproduced results and * is not directly comparable due to lack of evaluation details.

Strategy	Textual F1	Visual F1	MM F1
TEXT + IMAGE	38.2	20.9	27.4
TEXT → IMAGE	38.2	19.0	26.9
IMAGE → TEXT	36.6	20.4	26.5
IMAGE LOCKED	38.7	20.1	27.7

Table 2: Comparison of multimodal training strategies.

use AdamW (Loshchilov and Hutter, 2019) with an initial learning rate of 3×10^{-5} , weight decay of 1×10^{-3} , and a linear scheduler. We select the best model checkpoint based on F1 of the ACE2005 test dataset and set the default threshold $\tau = 0.5$ for inference. We report threshold tuning and implementation details in A.2 and A.3.

Since MMUTF focuses on EAE and requires event predictions for a realistic comparison, we ran the CAMEL event detection model to obtain event predictions. We achieve comparable results with a F1 of 55.5%, 57.0%, and 54.6% for textual, visual, and multimedia event detection, respectively. However, we observe an absolute performance drop of 1.5% and 2.9% for the visual and multimedia event detection tasks.

3.1 Results

The results of our main experiments can be found in Table 1. We find that our model surpasses the SOTA in terms of textual EAE by +7% F1. For visual and multimedia EAE, we generally perform better than the second-best systems, achieving improvements of +3.1% and +2.8% F1, respectively. These results demonstrate the efficacy of our approach by utilizing cross-ontology transfer capabilities and event semantics via natural language prompts. However, we observe a gap of -5.8%

F1 in multimedia EAE compared to CAMEL, primarily due to the visual modality. With the same set of predicted event mentions, the gap narrows to -0.5% F1 in multimedia EAE. We hypothesize that CAMEL’s better visual performance is related to its additional data augmentation pipeline and the integration of textual clues. Note that we do not utilize any additional image-language paired datasets or cross-modal data augmentation in this work. Incorporating cross-modal information from the document’s context and external datasets, as well as enhanced event detection, might further improve our model’s performance.

3.2 Analysis

Multimodal Training Strategy Table 2 presents results for different multimodal training strategies, focusing on cross-modality alignment for EAE. TEXT + IMAGE denotes the training strategy that jointly trains the textual and visual input modalities. TEXT → IMAGE represents the sequential procedure. Here, we first train our model on the text input modality, then freeze the text and query model parameters before training the visual input modality. With IMAGE → TEXT, we follow the same procedure as before, but in reverse. Lastly, we follow recommendations in vision-language alignment (Zhai et al., 2022) and freeze the visual encoder during joint training (IMAGE LOCKED). The results suggest that joint training achieves superior performance, while the sequential procedure experiences performance drops of up to -1.9% F1 for the aligned modality. The IMAGE LOCKED approach yields the best textual and multimodal F1 scores, while also representing the most memory-efficient alternative during training.

Model	Textual F1	Visual F1	MM F1
MMUTF	38.2	20.9	27.4
1 w/o Cross-Attention	38.5	17.8	26.5
2 w/o Joint-Prompts	37.0	17.5	24.4
3 w/o Prompts	37.9	15.9	23.3

Table 3: Ablations for the different components. (1) without cross-attention between candidate and query representations. (2) Replacement of joint query model with modality-specific models. (3) Replacement of textual query representations with trainable prototypes.

Ablation Study To better illustrate the effectiveness of different components, we conduct ablation studies and present the results in Table 3. In line 1, we remove the cross-attention mechanism between the candidate and query representations. This results in a performance reduction for visual and multimedia EAE, while we observe a slight increase for textual EAE. In line 2, we replace the joint model for query representations with two modality-specific models. The performance of all EAE cases significantly declines by up to +3.4% in terms of F1 scores. These observations suggest the effectiveness of unified template filling, particularly for the visual and multimedia EAE tasks. Lastly, to assess the influence of natural language-formulated event templates, we replace the textual query representations with trainable argument role prototypes in line 3. The performance declines mostly for visual and multimedia EAE, indicating that event template semantics provide beneficial information to enhance cross-ontology transfer capabilities. We provide backbone and prompt template ablations in A.5 and A.6, respectively.

Transfer Learning from SRL Previous work has shown that semantic role labeling (SRL) can serve as a valuable resource for EAE (Zhang et al., 2022). Analogous task similarities are also present in image activity recognition, which utilizes FrameNet verbs and frame elements as semantic roles (Yatskar et al., 2016). However, no work has yet explored the transfer learning capabilities for multimedia EAE. We use FrameNet 1.7 (Baker et al., 1998) and SWiG (Pratt et al., 2020) as our training datasets and merge their ontologies to 155 frames. Due to the large number of frames, the event templates are constructed with a semi-automatic approach (Zhang et al., 2022) (cf. A.4). Table 4 shows the results for zero-shot prediction on the M2E2 benchmark with gold and predicted event triggers. We observe remarkable transfer

Model	Textual F1	Visual F1	MM F1
<i>Gold Triggers</i>			
FN17+SWiG	36.9	34.7	24.5
ACE+SWiG _{M2E2}	55.1	36.9	31.7
<i>Pred Triggers</i>			
FN17+SWiG	25.8	17.4	20.7
ACE+SWiG _{M2E2}	38.2	20.9	27.4

Table 4: Evaluation of FrameNet transfer learning. *Gold Triggers* represent EAE based on annotated event mentions. *Pred Triggers* represent EAE based on predicted event mentions.

learning capabilities without any training on the M2E2 ontology. In fact, this simplistic approach outperforms even the supervised baselines FLAT and WASE by up to +11.6% and +5.8% F1 for visual and multimedia EAE, respectively.

4 Conclusion

In this work, we investigate the multimedia EAE task and propose MMUTF, a simple yet effective unified template filling approach that assigns candidates to corresponding argument roles. We encode entity candidates and object candidates from textual and visual inputs, respectively. Then, we feed a natural language event template into the query decoder, which facilitates interactions with document contexts via cross-attention mechanisms. Finally, candidate and query representations serve as the basis for computing candidate-query matching scores. Our results demonstrate the effectiveness of the model and highlight the benefits of cross-ontology transfer and event semantic incorporation. This assumption is supported by SRL transfer learning experiments, which we adapt to the multimedia EAE task. Future work will focus on incorporating additional textual and visual cues from input documents, integrating event detection, and exploring bi-encoder architectures to reduce inference time.

Limitations

In this work, we propose a unified template-filling approach for solving multimodal multimedia event argument extraction. Our method relies on manually crafted prompts that represent the event templates for each event type. The time required to construct such prompts was reasonable for the M2E2 benchmark, which covers only 8 event types and 15 argument roles for solving this task. Other event extraction ontologies cover over 100 event types and complex argument relations (Huang et al., 2023).

Although we used a semi-automatic approach for the more comprehensive FrameNet, we still identified erroneous event templates, indicating that this method is not fully reliable. The manual prompt construction for such large event ontologies also poses challenges in terms of required time and complexity. In addition, our system relies on event predictions, textual entities, and visual objects generated by preceding systems. This approach is prone to error propagation (cf. A.7) and requires a holistic assessment of all system components with appropriate evaluation data.

References

- Abien Fred Agarap. 2018. [Deep Learning using Rectified Linear Units \(ReLU\)](#). *arXiv preprint*. Version Number: 2.
- Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. 2022. [data2vec: A general framework for self-supervised learning in speech, vision and language](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 1298–1312. PMLR.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. [The Berkeley FrameNet Project](#). In *Proceedings of the 36th annual meeting on Association for Computational Linguistics -*, volume 1, page 86, Montreal, Quebec, Canada. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. [Scaling instruction-finetuned language models](#). *Journal of Machine Learning Research*, 25(70):1–53.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *International Conference on Learning Representations*.
- Zilin Du, Yunxin Li, Xu Guo, Yidan Sun, and Boyang Li. 2023. [Training Multimedia Event Extraction With Generated Images and Captions](#). In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5504–5513, Ottawa ON Canada. ACM.
- Kuan-Hao Huang, I-Hung Hsu, Tanmay Parekh, Zhiyu Xie, Zixuan Zhang, Premkumar Natarajan, Kai-Wei Chang, Nanyun Peng, and Heng Ji. 2023. [TextEE: Benchmark, Reevaluation, Reflections, and Future Challenges in Event Extraction](#). *arXiv preprint*. Version Number: 3.
- Glenn Jocher, Ayush Chaurasia, and Jing Qiu. 2023. [Ultralytics YOLO](#).
- Zeeshan Khan, C. V. Jawahar, and Makarand Tapaswi. 2022. [Grounded Video Situation Recognition](#). In *Advances in Neural Information Processing Systems*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Manling Li, Ruochen Xu, Shuohang Wang, Luowei Zhou, Xudong Lin, Chenguang Zhu, Michael Zeng, Heng Ji, and Shih-Fu Chang. 2022. [CLIP-Event: Connecting Text and Images with Event Structures](#). In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16399–16408, New Orleans, LA, USA. IEEE.
- Manling Li, Alireza Zareian, Qi Zeng, Spencer Whitehead, Di Lu, Heng Ji, and Shih-Fu Chang. 2020. [Cross-media Structured Common Space for Multimedia Event Extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2557–2568, Online. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft COCO: Common Objects in Context](#). In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, volume 8693, pages 740–755. Springer International Publishing, Cham. Series Title: Lecture Notes in Computer Science.
- Jian Liu, Yufeng Chen, and Jinan Xu. 2022. [Multimedia Event Extraction From News With a Unified Contrastive Learning Framework](#). In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1945–1953, Lisboa Portugal. ACM.

- Yang Liu, Fang Liu, Licheng Jiao, Qianye Bao, Long Sun, Shuo Li, Lingling Li, and Xu Liu. 2024. [Multi-grained Gradual Inference Model for Multimedia Event Extraction](#). *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Yubo Ma, Zehao Wang, Yixin Cao, Mukai Li, Meiqi Chen, Kun Wang, and Jing Shao. 2022. [Prompt for Extraction? PAIE: Prompting Argument Interaction for Event Argument Extraction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6759–6774, Dublin, Ireland. Association for Computational Linguistics.
- Sarah Pratt, Mark Yatskar, Luca Weihs, Ali Farhadi, and Aniruddha Kembhavi. 2020. [Grounded Situation Recognition](#). In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, volume 12349, pages 314–332. Springer International Publishing, Cham. Series Title: Lecture Notes in Computer Science.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning Transferable Visual Models From Natural Language Supervision](#). *arXiv preprint*. Version Number: 1.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Lin Sun, Kai Zhang, Qingyuan Li, and Renze Lou. 2024. [UMIE: Unified Multimodal Information Extraction with Instruction Tuning](#). In *Proceedings of AAAI*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pages 6000–6010, Red Hook, NY, USA. Curran Associates Inc. Event-place: Long Beach, California, USA.
- Walker, Christopher, Strassel, Stephanie, Medero, Julie, and Maeda, Kazuaki. 2006. [ACE 2005 Multilingual Training Corpus](#). Artwork Size: 1572864 KB Pages: 1572864 KB.
- Sijia Wang, Mo Yu, Shiyu Chang, Lichao Sun, and Lifu Huang. 2022. [Query and Extract: Refining Event Extraction as Type-oriented Binary Decoding](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 169–182, Dublin, Ireland. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Trans-formers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Mark Yatskar, Luke Zettlemoyer, and Ali Farhadi. 2016. Situation Recognition: Visual Semantic Role Labeling for Image Understanding. In *Conference on Computer Vision and Pattern Recognition*.
- Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. 2022. [LiT: Zero-Shot Transfer with Locked-image text Tuning](#). In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18102–18112, New Orleans, LA, USA. IEEE.
- Zhisong Zhang, Emma Strubell, and Eduard Hovy. 2022. [Transfer Learning from Semantic Role Labeling to Event Argument Extraction with Template-based Slot Querying](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2627–2647, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ce Zheng, Yiming Wang, and Baobao Chang. 2023. [Query Your Model with Definitions in FrameNet: An Effective Method for Frame Semantic Role Labeling](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):14029–14037.

A Appendix

A.1 Baselines

We compare our proposed approach with a wide range of SOTA models. 1) WASE (Li et al., 2020) uses graph neural networks (GNN) together with abstract meaning representations and situation graphs for text and image, respectively. These models additionally include the VOA image-caption dataset for weak cross-modality alignment, while the 2) FLAT (Li et al., 2020) baseline removes the GNN components. 3) CLIP_{EVENT} (Li et al., 2022) combines the previous approaches with pre-trained CLIP model weights and introduces a Optimal Transport strategy for event structure representation and modality alignment. 4) UniCL (Liu et al., 2022) proposes a unified contrastive learning framework, uses a multimodal contrastive learning mechanism and incorporates textual and visual clues. 5) CAMEL (Du et al., 2023) is based on cross-modal data augmentation, integrates an adapter module and applies a specialized training strategy. 6) MGIM (Liu et al., 2024) is a multi-grained gradual inference model with fine-grained alignment between texts and images by multiple rounds of gradual inference. 7) UMIE (Sun et al., 2024) utilizes instruction tuning for unified multimodal information extraction based on text generation and Flan-T5 (Chung et al., 2024) in various model sizes.

A.2 Threshold Tuning

We investigate the influence of an automatic or manually chosen threshold for both the textual and visual matching scores. In our work, we select the default threshold $\tau = 0.5$ and obtain reasonable results. Figure 3 shows the metrics for different thresholds and modalities. Additionally, we analyze threshold tuning based on the best F1-scores for the ACE and SWiG test datasets. This procedure yields $\tau^T = 0.2$ and $\tau^V = 0.3$ for text and vision, respectively. For automatic threshold tuning, we obtain a multimedia F1-score of 28.3%, improving our default τ by +0.9%. In addition, we find the most performance degradation above a threshold of 0.6 while achieving the best results between 0.3 and 0.5. This suggests a weak discriminative power between positive and negative candidate-query pairs. Contrastive learning and hard negative mining may address this issue, but is not within the scope of this work.




Figure 3: EAE F1 results for our MMUTF model with varying thresholds for each modality.

A.3 Implementation Details

In our experiments, we use the implementation of the *Transformers* (Wolf et al., 2020) (v4.41.0) library in conjunction with PyTorch (v2.3.0). For all runs, we evaluated the performance at each epoch on the ACE2005 test set and selected the best-performing checkpoint. Unless otherwise mentioned, we use T5-base (Raffel et al., 2020) with 222M parameters as text encoder-decoder and CLIP-base (Radford et al., 2021) with 85M parameters as vision encoder model. All models are trained with mixed precision and executed on A100 GPUs with 40GB HBM using the same compute node running CUDA 12.3 and NVIDIA device driver version 545. The average training time for each run did not exceed 3h.

A.4 Event Templates

All manually crafted event templates for the M2E2 benchmark are shown in Table 5. The words in brackets correspond to the argument roles we aim to extract. These event templates are selected based on the predicted event type, then tokenized and fed into the query model to obtain the final query representations. For FrameNet, it is infeasible to manually construct templates for each frame and corresponding frame element. Therefore, we follow (Zhang et al., 2022) and adopt their semi-automatic method to generate templates for our transfer learning experiments. We merge the ontologies of FrameNet 1.7 (Baker et al., 1998) and SWiG (Pratt et al., 2020) and use the FrameNet templates for both datasets. We show some excerpts in Table 5.

A.5 Backbone Models

We compare different text and vision model configurations in Table 6. For the pre-trained text models, we compare the following encoder and encoder-

Event Type	Event Template
M2E2	
Movement:Transport	<i>[Agent]</i> transported <i>[Artifact]</i> in <i>[Vehicle]</i> from <i>[Origin]</i> to <i>[Destination]</i> .
Conflict:Attack	<i>[Attacker]</i> attacked <i>[Target]</i> using <i>[Instrument]</i> as <i>[Place]</i> .
Conflict:Demonstrate	<i>[Entity]</i> protested besides <i>[Police]</i> using <i>[Instrument]</i> at <i>[Place]</i> .
Justice:Arrest-Jail	<i>[Person]</i> was sent to jail or arrested by <i>[Agent]</i> using <i>[Instrument]</i> at <i>[Place]</i> .
Contact:Phone-Write	<i>[Entity]</i> called or texted messages using <i>[Instrument]</i> at <i>[Place]</i> .
Contact:Meet	<i>[Entity]</i> met at <i>[Place]</i> .
Life:Die	<i>[Agent]</i> killed <i>[Victim]</i> with <i>[Instrument]</i> at <i>[Place]</i> .
Transaction:Transfer-Money	<i>[Giver]</i> paid <i>[Recipient]</i> with <i>[Money]</i> for the benefit of <i>[Beneficiary]</i> at <i>[Place]</i> .
FrameNet	
Abandonment	<i>[Agent]</i> abandon <i>[Theme]</i> in <i>[Place]</i> .
Motion	<i>[Theme]</i> move in <i>[Area]</i> <i>[Direction]</i> from <i>[from]</i> to/into <i>[Goal]</i> in <i>[Place]</i> .
Losing track of theme	<i>[Perceiver]</i> lose <i>[Theme]</i> in <i>[Place]</i> .
Adjusting	<i>[Agent]</i> adjust <i>[Part]</i> <i>[Feature]</i> in <i>[Place]</i> with <i>[Instrument]</i> .
Killing	<i>[Cause]</i> <i>[Killer]</i> kill <i>[Victim]</i> with <i>[Instrument]</i> <i>[Means]</i> in <i>[Place]</i> .
Appointing	<i>[Selector]</i> appoint to <i>[Body]</i> <i>[Official]</i> as <i>[Role]</i> <i>[Function]</i> in <i>[Place]</i> .
Arranging	<i>[Agent]</i> arrange <i>[Theme]</i> in <i>[Configuration]</i> with <i>[Instrument]</i> .
Fining	<i>[Speaker]</i> fine <i>[Payer]</i> <i>[Fine]</i> <i>[Reason]</i> with <i>[Instrument]</i> in <i>[Place]</i> .

Table 5: M2E2 and FrameNet event templates for each event type serve as natural language prompts for EAE.

Model		Textual F1	Visual F1	MM F1
BERT	ViT	37.9	19.2	26.6
	CLIP	37.6	19.7	26.6
	DATA2VEC	38.4	18.3	26.1
BART	ViT	38.1	19.2	24.9
	CLIP	38.5	20.4	27.9
	DATA2VEC	39.0	18.4	27.0
T5	ViT	38.3	19.5	27.6
	CLIP	38.2	20.9	27.4
	DATA2VEC	38.5	18.5	27.0

Template	Textual F1	Visual F1	MM F1
w/ event type			
Concatenation	38.0	18.9	27.9
Standard	38.3	17.9	25.3
Enriched	38.5	16.8	25.2
w/o event type			
Concatenation	38.7	19.0	28.0
Standard	38.2	20.9	27.4
Enriched	38.0	19.7	28.1

Table 6: Results for text and vision models.

Table 7: Results for different prompt templates.

decoder variants: BERT³ (Devlin et al., 2019), BART⁴ (Lewis et al., 2020), and T5⁵ (Raffel et al., 2020). Regarding BERT, we use two distinct models for candidate and query representations and add cross-attention to the query encoder. For the pre-trained vision models, we rely on ViT⁶ (Dosovitskiy et al., 2021), CLIP⁷ (Radford et al., 2021), and DATA2VEC⁸ (Baevski et al., 2022). All models represent their corresponding base variants, as we did not observe substantial improvements with larger model sizes.

Overall, the results demonstrate robust performance across different model configurations. How-

ever, we observe that CLIP generally outperforms the other vision models in terms of visual and multi-media F1 scores. We hypothesize that pre-training on large image-caption datasets is beneficial for the M2E2 task.

A.6 Prompt Analysis

In Table 7, we investigate prompt template variations and conduct experiments with and without event type prefixes. For event type prefixes, we prepend the event type name to the event template (e.g., "Conflict Attack: <Event Template>"). We compare three different prompt variations: *Concatenation* simply concatenates the argument role names from the event templates. *Standard* represents the manually crafted event templates used in this work (cf. Table 5). *Enriched* adds argument role definitions after each argument role name.

³<https://huggingface.co/google-bert/bert-base-uncased>

⁴<https://huggingface.co/facebook/bart-base>

⁵<https://huggingface.co/google-t5/t5-base>

⁶<https://huggingface.co/google/vit-base-patch16-224>

⁷<https://huggingface.co/openai/clip-vit-base-patch16>

⁸<https://huggingface.co/facebook/data2vec-vision-base>

The results show that event templates without event type prefixes generally outperform their counterparts across all variations. We conjecture this is because similar arguments appear in multiple events, and event type prefixes may impede the transfer of knowledge between roles. All variations perform similarly for textual EAE, while event type prefixes slightly impair model performance for visual and multimedia EAE.

A.7 Error Analysis

While our method achieves significant improvements across all SOTA models for the textual EAE task, we still recognize shortcomings for the visual modality. Therefore, we analyse the impact of visual event mentions and visual candidates predicted by YOLOv8 (Jocher et al., 2023). With gold triggers (i.e., annotated visual event mentions), we achieve an absolute improvement of 16% in visual F1 which highlights the importance of reliable ED. However, the inclusion of gold candidates (i.e., annotated bounding boxes) demonstrates an absolute improvement of 35.5% visual F1 which further enhances the gold triggers by 19.5%. These results suggests that the correct selection of candidates is crucial and the joint modeling of candidate generation and EAE poses a promising direction.