

# Tensorized Optical Multimodal Fusion Network

Yequan Zhao<sup>1</sup>, Xian Xiao<sup>2</sup>, Geza Kurczveil<sup>2</sup>, Raymond G. Beausoleil<sup>2</sup>, and Zheng Zhang<sup>1</sup>

<sup>1</sup> Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106, USA

<sup>2</sup> Hewlett Packard Labs, Hewlett Packard Enterprise, 820 N. McCarthy Blvd., Milpitas, CA 95305, USA

yequan\_zhao@ucsb.edu, xian.xiao@hpe.com, zhengzhang@ece.ucsb.edu

**Abstract:** We propose the first tensorized optical multimodal fusion network architecture with a self-attention mechanism and low-rank tensor fusion. Simulation results show  $51.3\times$  less hardware requirement and  $3.7\times 10^{13}$  MAC/J energy efficiency. © 2023 The Author(s)

## 1. Introduction

Modern machine learning (ML) applications often include heterogeneous data sources, e.g., visual, audio, and text data, in virtual reality (VR) interactions. Multimodal fusion networks offer a solution to robust performance under data uncertainty/attacks/corruptions by fully leveraging the complementary information from different modalities. However, modern edge devices' limited memory and power budget constrain their real-time performance. Optical ML accelerators, such as optical neural networks (ONNs), are competitive solutions for edge computing due to their wide bandwidth, high speed, and low energy consumption. Moreover, recent progress in tensorized optical neural networks (TONN) enables the implementation of a large-scale (e.g.,  $1024\times 1024$ ) ONN with cascaded small-scale (e.g.,  $8\times 8$ ) photonic tensor cores [1], which significantly improved the scalability of ONN.

In this paper, we propose a tensorized optical multimodal fusion network (TOMFN) to further explore TONN's capability for processing large-scale and versatile ML data. By utilizing both tensor-train (TT) [2] and CP decompositions [4], the proposed compact and energy-efficient architecture can be implemented on integrated photonic platforms with small-scale photonic tensor cores as the building blocks. Moreover, compared to full-size counterparts, our proposed network maintains a compatible inference accuracy in multimodal sentiment analysis tasks while requiring  $51.3\times$  fewer hardware resources and excellent power reduction (i.e., 7.9nJ per inference).

## 2. Principles and Architecture

Fig.1 (c) depicts the conceptual architecture of the proposed TOMFN. It starts with three subnetworks that extract visual, audio, and text embeddings from corresponding input sources. The extracted vectors then go through a low-rank multimodal fusion network (LMF) [5], which maintains the expressiveness and effectiveness of tensor fusion network (TFN) [6] while directly computing a rich and joint multimodal representation  $h$  without explicitly creating the tensor. Lastly, the sentiment inference network consists of simply a softmax layer conditioned on  $h$ .

In particular, the processing of text embeddings is realized by the tensorized self-attention mechanism instead of the conventional LSTM layers. The benefit of self-attention is that it is a feed-forward structure without any

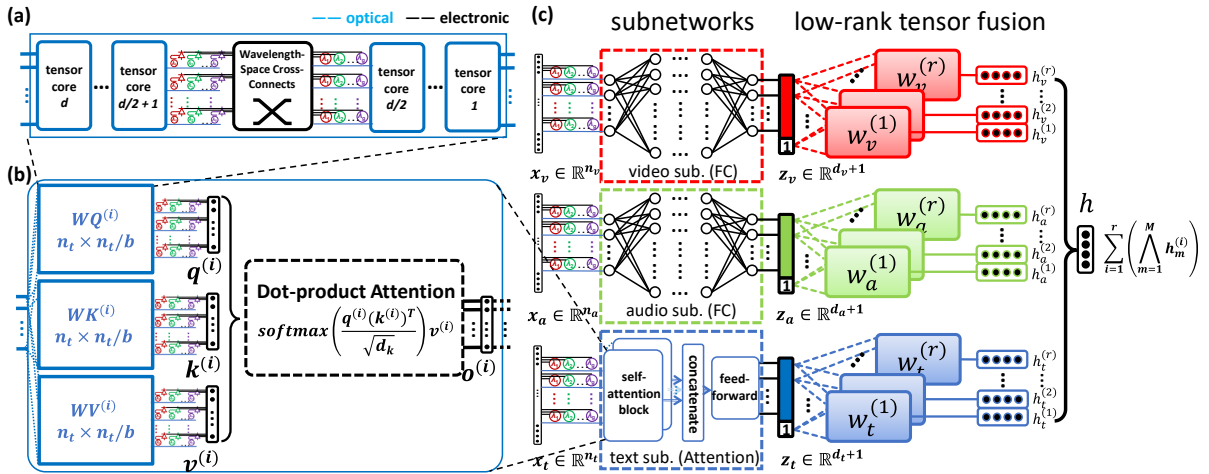


Fig. 1: (a) Device implementation of a tensorized optical layer. (b) Detailed structure of a self-attention block. (c) Architecture of the tensorized optical multimodal fusion network (TOMFN).

recurrent layer, thus requiring no memory for intermediate states, and it has better parallelism. Fig. 1 (b) shows the detailed structure of the tensorized self-attention block. A single self-attention block computes one set of a dot-product attention score containing three optical tensorized layers that transform the input sequence  $x_i$  into query vector  $\mathbf{q}$ , key vector  $\mathbf{k}$ , and value vector  $\mathbf{v}$  in parallel. By concatenating multiple blocks, *multi-head attention*, the model jointly attends to information from different representation subspaces at different positions [7]. The low-rank fusion network is also feed-forward. It comprises three modality-specific tensorized optical layers followed by element-wise product over the corresponding entries of each modality output and summation over fusion rank  $r$ . In practice, The TT and CP-format tensor ranks of all blocks could be learnt automatically in the training process [3]. The element-wise product can be realized by two cascaded stages of balanced homodyne detection.

The photonic implementation of the large-scale weight matrices in fully-connected (FC) layers, fusion layers, and self-attention layers are all based on TONN architecture [1]. TONN architecture takes advantage of the TT-decomposition algorithm and wavelength division multiplexing (WDM) technology, as shown in Fig. 1 (a). Such architecture can realize large-scale tensorized weight matrices by cascading small-radix photonic tensor cores implemented by Mach-Zehnder interferometer (MZI) meshes. This way, the number of MZIs and cascaded stages of MZIs are significantly reduced, leading to a significantly more compact footprint, smaller optical loss budget, and more manageable control complexity. More importantly, TONN architecture eliminates O/E/O conversions and intermediate memory, significantly improving energy efficiency and latency.

### 3. Simulation Results and Conclusion

The sentiment analysis result on the interactive emotional dyadic motion capture (IEMOCAP) dataset [8] shows that our proposed TOMFN model has a competitive accuracy compared with the full-size LMF model. At the same time, our architecture exhibits  $92.8\times$  fewer model parameters, thus requiring  $51.3\times$  fewer photonic devices (i.e., MZIs). Note that our tensorized attention-based model outperforms its counterpart with a tensorized LSTM layer with the exact hardware requirements. In detail, we utilized three FC layers ( $80\times 32$ ,  $32\times 32$ , and  $32\times 32$ ) for the visual subnetwork, three FC layers ( $36\times 32$ ,  $32\times 32$ , and  $32\times 32$ ) for the audio subnetwork, and two blocks of self-attention layers (each contains three  $300\times 150$ ) and a feed-forward layer ( $300\times 64$ ) for the text subnetwork. In total, thirty-eight  $4\times 4$ , seventy-four  $6\times 6$ , and sixteen  $8\times 8$  photonic tensor cores based on MZI meshes are needed for implementing all the weight matrices. Assuming the modulation speed is 10 Gb/s, the system power consumption is estimated at 79.87W, corresponding to 7.9 nJ per inference and  $3.7\times 10^{13}$  MAC/J.

	Model size			IEMOCAP F1-result				Power estimation @10GHz
	# Param.*	# MZI	# stage	Happy	Sad	Angry	Neutral	Efficiency
TFN [6]	758176	607570	2841	83.6	82.8	84.2	65.4	N/A
LMF [5]	106912	86802	921	<b>85.8</b>	<b>85.9</b>	<b>89</b>	<b>71.7</b>	N/A
TOMFN(LSTM)	<b>844</b>	<b>1540</b>	204	81.3	78.2	83.5	61.6	$1.9\times 10^{13}$ MAC/J
TOMFN(Attention)	1152	1691	<b>166</b>	<b>83.4</b>	<b>82.7</b>	<b>85.7</b>	<b>66.7</b>	<b><math>3.7\times 10^{13}</math> MAC/J</b>

Table 1: Test Accuracy and Power Consumption Analysis (\* Only including the parameters of weight matrices)

In conclusion, given the success of attention-based architecture (e.g., Transformer [7]) in versatile applications, including natural language processing, object detection, and reinforcement learning, our proposed TOMFN lays the foundations for scalable, energy-efficient, and robust optical ML accelerators for heterogeneous data sources.

### References

1. Xiao X, On M B, Van Vaerenbergh T, et al. Large-scale and energy-efficient tensorized optical neural networks on III–V-on-silicon MOSCAP platform[J]. *APL Photonics*, 2021, 6(12): 126107.
2. Oseledets, Ivan V. "Tensor-train decomposition." *SIAM Journal on Scientific Computing* 33.5 (2011): 2295-2317.
3. Hawkins, Cole, Xing Liu, and Zheng Zhang. "Towards compact neural networks via end-to-end training: A Bayesian tensor approach with automatic rank determination." *SIAM Journal on Mathematics of Data Science* 4.1 (2022): 46-71.
4. Bro, Rasmus. "PARAFAC. Tutorial and applications." *Chemometrics and intelligent laboratory systems* 38.2 (1997): 149-171.
5. Liu Z, Shen Y, Lakshminarasimhan V B, et al. Efficient low-rank multimodal fusion with modality-specific factors. *arXiv preprint arXiv:1806.00064*, 2018
6. Zadeh A, Chen M, Poria S, et al. Tensor fusion network for multimodal sentiment analysis[J]. *arXiv preprint arXiv:1707.07250*, 2017.
7. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. *Advances in neural information processing systems*, 2017, 30.
8. Busso, Carlos, et al. "IEMOCAP: Interactive emotional dyadic motion capture database." *Language resources and evaluation* 42.4 (2008): 335-359.