

Compact Tensor Pooling for Visual Question Answering

Yang Shi
University of California, Irvine
shiy4@uci.edu

Tommaso Furlanello
University of Southern California
tfurlanello@gmail.com

Anima Anandkumar
Amazon
anima@amazon.com

Abstract

Performing high level cognitive tasks requires the integration of feature maps with drastically different structure. In Visual Question Answering (VQA) image descriptors have spatial structures, while lexical inputs inherently follow a temporal sequence. The recently proposed Multimodal Compact Bilinear pooling (MCB) forms the outer products, via count-sketch approximation, of the visual and textual representation at each spatial location. While this procedure preserves spatial information locally, outer-products are taken independently for each fiber of the activation tensor, and therefore do not include spatial context. In this work, we introduce multi-dimensional sketch (MD-sketch), a novel extension of count-sketch to tensors. Using this new formulation, we propose Multimodal Compact Tensor Pooling (MCT) to fully exploit the global spatial context during bilinear pooling operations. Contrarily to MCB, our approach preserves spatial context by directly convolving the MD-sketch from the visual tensor features with the text vector feature using higher order FFT. Furthermore we apply MCT incrementally at each step of the question embedding and accumulate the multi-modal vectors with a second LSTM layer before the final answer is chosen.

1. Introduction

Visual question answering tasks require a joint representation of both visual and textual input [9]. Image and text features are respectively extracted with convolutional neural networks [7] and recurrent neural networks [8, 11]. Information of each modality are then combined together, using concatenation, element-wise product or sum operations. Bilinear pooling[2] exploits the outer product operation to capture richer information from the two sets of features. However, its computational/memory complexity ($O(n^2)$, if each feature is $O(n)$) is prohibitive in practice. Compact bilinear pooling [1] has the same discriminative power as the bilinear pooling while requiring far less computation/memory storage ($O(d)$, $d \ll n^2$).

In this paper, we first propose the MD-sketch that generalizes count-sketch to tensors of arbitrary orders. We then leverage this formulation of the MD-sketch in a new compact tensor pooling method. This new pooling, which we

coin *MCT*, directly estimates the outer product between tensor descriptors of different orders, therefore preserving context. We present novel VQA architectures with MCT and also aggregate MCT within LSTM. Specifically, we apply the MCT model at each time-step of the language LSTM to combine the image features with each word representation.

2. Related work

Tensor sketch Count Sketch [4] estimates the data frequency in a stream. Considering that as a specific random projection in high dimensional space, Pham and Pagh propose tensor sketch for polynomial kernel estimation in [1]. Wang et al. [6] computes tensor contraction via tensor sketch. The computational burden is reduced while performing the operations by sketching the tensor as a vector. However, this ignores the natural data structure. We generalize count-sketch method to tensors of arbitrary orders while preserving spatial context and avoiding expensive computation in Section 4.

Pooling methods Pooling methods are widely used in visual tasks to combine information for various streams into one final feature representation. Common pooling methods are average pooling and bilinear pooling. Bilinear pooling requires taking the outer product between two features. However the result is high dimensional. Count sketch is applied as a feature hashing operator to avoid such increase in dimension. Gao et al. [2] uses convolution layers from two different neural networks as the local descriptor extractors of the image and combine them using count sketch. Simon et al. [3] presents a new " α -pooling" which allows the network to learn the pooling strategy: a continuous transition between linear and polynomial pooling. They show that higher α gives larger gain for fine-grained image recognition tasks. However, as α goes up, the computation complexity increases in polynomial order. We propose up with compact polynomial pooling to reduce the computational burden. Details in Section 3.

3. Multimodal compact bilinear pooling(MCB)

Fukui et al. [5] uses count sketch as a pooling method in VQA tasks and obtains the best results on VQA dataset v1.0. We denote the count sketch operator as *cs*, the tensor

product (outer product in vector case) as \otimes and convolution operator as \star . The key idea in MCB is to avoid directly computing the outer product and to reduce the output dimension by writing $cs(x \otimes y) = cs(x) \star cs(y)$. According to convolution theorem, this can be further simplified by doing element-wise product in the frequency domain before transforming the result back to the time domain using inverse FFT(IFFT). Here, we argue that the IFFT step is redundant in the VQA task. FFT itself preserves the correlation between the two data streams. We can therefore accelerate the training by skipping the IFFT step. By padding the representation vectors with $\bar{1}$ (vector with all elements 1) before the count sketch, we additionally preserve the first order information along with the joint correlation. Assume we have two input features $x, \bar{1}_2 \in \mathcal{R}^{n_1}, y, \bar{1}_1 \in \mathcal{R}^{n_2}$, set $\tilde{x} = [x, \bar{1}_1], \tilde{y} = [y, \bar{1}_2]$ such that $\tilde{x}, \tilde{y} \in \mathcal{R}^{n_1+n_2}$, $cs(\tilde{x} \otimes \tilde{y})$ will contain information of $x \otimes y, x$ and y . By repeating the same operation multiple times it is possible to obtain estimates of higher polynomials of x and y .

4. Multimodal compact tensor pooling(MCT)

While MCB shows great performance in [5], it loses the spatial context in the image as the sketch is done independently between the 1-dimensional text feature vector and each fiber of the image tensor. We instead propose to preserve that context in VQA tasks, by applying MD-sketch on both the image tensor and the text features and convolving the resulting tensors.

MD-sketch of a tensor Even though the MD-sketch is defined for tensors of arbitrary orders, we limit ourselves to third order tensors for clarity. Given a feature tensor $\mathcal{I} \in \mathcal{R}^{C \times H \times W}$, a random hash functions $h_1: [C] \rightarrow [d_1], h_2: [H] \rightarrow [d_2], h_3: [W] \rightarrow [d_3]$, and a random sign function $s_1: [C] \rightarrow \{\pm 1\}, s_2: [H] \rightarrow \{\pm 1\}, s_3: [W] \rightarrow \{\pm 1\}$, for $i \in [C], j \in [H], k \in [W]$, the MD-sketch of \mathcal{I} is defined as:

$$ss(\mathcal{I})_{t_1, t_2, t_3} = \mathcal{X}(t_1, t_2, t_3) = \sum_{h_1(i)=t_1, h_2(j)=t_2, h_3(k)=t_3} s_1(i) s_2(j) s_3(k) \mathcal{I}(i, j, k)$$

The resulting tensor $\mathcal{X} \in \mathcal{R}^{d_1 \times d_2 \times d_3}$ is an estimation of \mathcal{I} with preserved spatial information. We apply count sketch on text feature $v \in \mathcal{R}^L$. Given hash function $h_4: [L] \rightarrow [d_4]$ and sign function $s_4: [L] \rightarrow \{\pm 1\}$:

$$cs(v)_{t_1} = w(t_1) = \sum_{h_4(i)=t_1} s_4(i) v(i)$$

To compute the bilinear pooling between \mathcal{I} and v , we propose MCT as follows:

$$\begin{aligned} MCT(\mathcal{I} \otimes v)_{t_1, t_2, t_3} &= \mathcal{Y}(t_1, t_2, t_3) \\ &= (ss(\mathcal{I}) \star cs(v))_{t_1, t_2, t_3} \\ &= 3dFFT(ss(\mathcal{I}))_{t_1, t_2, t_3} \cdot 1dFFT(cs(v))_{t_4} \end{aligned}$$

where $t_4 = (t_1 + t_2 + t_3) \bmod d_4$. We show the MCT module in Figure 1.



Figure 1: MCT in VQA

Figure 2: MCT-group in VQA

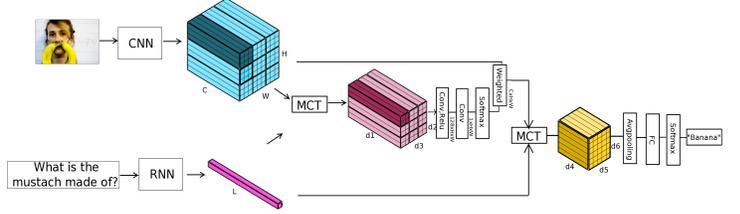


Figure 3: Our architecture for VQA with MCT

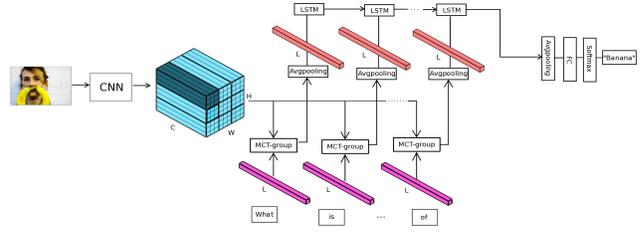


Figure 4: Our architecture for VQA with LSTM-MCT, MCT-group is as per in Figure 2

The intuition here is that the visual representation at each grid location is not fully isolated from grids around it. When applying compact bilinear pooling along each pixel, spatial context information is lost. By combining these surrounding information using MCT, we preserve the natural tensor structure of the image. In practice, we can do local tensor spatial pooling: divide the image tensor in several sub tensors and apply MCT on each of them.

5. VQA architecture

MCT The first extension that we propose in Figure 3 is the use of MCT instead of MCB during the feature integration steps. In the attention pathway we use MD-sketch to project the image feature to a lower dimensional space. After 2 layers of convolution we compute a soft attention map of the original image size. After filtering the image features with the attention map the resulting tensor is combined through MCT with the LSTM vector again before a final choice is taken.

LSTM-MCT We propose in Figure 4 a sequential model to answer the question incrementally by giving a glimpse to the image for each word. We apply the MCT model at each time-step of the language LSTM and obtain a sequence of multimodal vectors. A second LSTM is used to aggregate and answer from the sequence of multimodal vectors.

References

- [1] Ninh Pham and Rasmus Pagh, Fast and Scalable Polynomial Kernels via Explicit Feature Maps, In KDD'13 [1](#)
- [2] Yang Gao and Oscar Beijbom and Ning Zhang and Trevor Darrell, Compact Bilinear Pooling, Computer Vision and Pattern Recognition (CVPR), 2016 [1](#)
- [3] Marcel Simon and Erik Rodner and Yang Gao and Trevor Darrell and Joachim Denzler, Generalized orderless pooling performs implicit salient matching [1](#)
- [4] M. Charikar and K. Chen and M. Farach-Colton, Finding frequent items in data streams, In Proceedings of ICALP02 [1](#)
- [5] Akira Fukui and Dong Huk Park and Daylen Yang and Anna Rohrbach and Trevor Darrell and Marcus Rohrbach, Multimodal Compact Bilinear Pooling for Visual Question Answering and Visual Grounding, EMNLP 2016 [1, 2](#)
- [6] Yining Wang and Hsiao-Yu Tung and Alexander Smola and Animashree Anandkumar, Fast and Guaranteed Tensor Decomposition via Sketching, NIPS 2015 [1](#)
- [7] Jeff Donahue and Yangqing Jia and Oriol Vinyals and Judy Hoffman and Ning Zhang and Eric Tzeng and Trevor Darrell, DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition, CoRR 2013 [1](#)
- [8] Ilya Sutskever and Oriol Vinyals and Quoc V. Le Sequence to Sequence Learning with Neural Networks, CoRR 2014 [1](#)
- [9] Stanislaw Antol and Aishwarya Agrawal and Jiasen Lu and Margaret Mitchell and Dhruv Batra and C. Lawrence Zitnick and Devi Parikh, VQA: Visual Question Answering, International Conference on Computer Vision (ICCV) 2015 [1](#)
- [10] Sutskever, Ilya and Vinyals, Oriol and Le, Quoc V., Sequence to Sequence Learning with Neural Networks, NIPS'14
- [11] Hochreiter, Sepp and Schmidhuber, Jürgen, Long short-term memory, Neural computation 1997 [1](#)