Convolution, attention and structure embedding

Jean-Marc Andreoli*
Naverlabs Europe, Grenoble, France
http://www.europe.naverlabs.com

Abstract

Deep neural networks are composed of layers of parametrised linear operations intertwined with non linear activations. In basic models, such as the multi-layer perceptron, a linear layer operates on a simple input vector embedding of the instance being processed, and produces an output vector embedding by straight multiplication by a matrix parameter. In more complex models, the input and output are structured and their embeddings are higher order tensors. The parameter of each linear operation must then be controlled so as not to explode with the complexity of the structures involved. This is essentially the role of convolution models, which exist in many flavours dependent on the type of structure they deal with (grids, networks, time series etc.). We present here a unified framework which aims at capturing the essence of these diverse models, allowing a systematic analysis of their properties and their mutual enrichment. We also show that attention models naturally fit in the same framework: attention is convolution in which the structure itself is adaptive, and learnt, instead of being given a priori.

1 A generic framework for convolution on arbitrary structures

Convolution is a powerful operator, which is widely used in deep neural networks in many different flavours: [12, 11, 8, 6, 9, 17, 14]. It allows to express in a compact form operations on a structured bundle of similarly shaped data instances (embeddings of nodes in a network, of instants in a time series, of pixels in an image, etc.) taking into account some known structural dependencies between them (edges between nodes, or temporal relations between instants, or positional relations between pixels). In spite of their apparent diversity, these structures can be formalised as *families* of weighted graphs, where each graph in a family captures one aspect of the structure. We develop a generic model of convolution over such structures.

1.1 Some useful properties of tensors

A tensor is characterised by its shape $S{=}\langle S_1 \cdots S_{|S|}\rangle$, which is a sequence of integers, its index set which is the cartesian product $\bar{S} \triangleq \prod_{i=1:|S|}\{1 \cdots S_i\}$ of cardinality $|\bar{S}|{=}\prod_{i=1:|S|} S_i$, and its value which is a mapping from its index set into the set of scalars. If S and T are shapes, we let ST denote their concatenation. The following classical concepts are defined in Appendix A.1: tensor *slicing*, *flattening* (and its instances: *matricisation* and *vectorisation*) [15] using the *canonical* bijection, *outer product* (a.k.a. tensor product). We make use of a specific operator, *Composition*: if $\boldsymbol{a}, \boldsymbol{b}$ are tensors of shape $\langle K\rangle S$ and $\langle K\rangle T$, respectively, where K is an integer, their composition denoted $\boldsymbol{a} \circ \boldsymbol{b}$ is a tensor of shape ST defined by:

$$\boldsymbol{a} \circ \boldsymbol{b} \;\triangleq\; \sum\nolimits_k \boldsymbol{a}_k \otimes \boldsymbol{b}_k \tag{1}$$

Operator $\circ$ combines features of both the inner and outer products. When S and T are both of length 1, then $\langle K\rangle S$, $\langle K\rangle T$, ST are all of length 2, so $\boldsymbol{a}$, $\boldsymbol{b}$ and $\boldsymbol{a} \circ \boldsymbol{b}$ are all matrices and $\boldsymbol{a} \circ \boldsymbol{b} = \boldsymbol{a}^\top \boldsymbol{b}$.

* jean-marc.andreoli@naverlabs.com

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

Proposition 1 (Decomposition). *Let S, T be arbitrary shapes, K an integer, and $\boldsymbol{a}$ a tensor of shape $\langle K\rangle S$ s.t. the family $(\boldsymbol{a}_k)_{k=1:K}$ be a basis of the space of tensors of shape S (hence $K{=}|\bar{S}|$). Then for any tensor $\boldsymbol{\Phi}$ of shape ST there exists a unique tensor $\boldsymbol{\Theta}$ of shape $\langle K\rangle T$ such that $\boldsymbol{\Phi} = \boldsymbol{a} \circ \boldsymbol{\Theta}$.*

Proof in Appendix A.1. When S and T are of length 1, then $\boldsymbol{\Phi}$ is a matrix, and the decomposition simply becomes $\boldsymbol{\Phi}{=}\boldsymbol{a}^\top\boldsymbol{\Theta}$, which is uniquely realised by $\boldsymbol{\Theta}{=}\boldsymbol{a}^{-1\top}\boldsymbol{\Phi}$ ($\boldsymbol{a}$ invertible by assumption).

1.2 A generic convolution model

In a convolution layer, the input does not consist of a simple embedding vector, as in a standard linear layer. Instead, it is a matrix $\boldsymbol{x}$ of shape $\langle M, P\rangle$, representing a bunch of M entries encoded as vectors of shape $\langle P\rangle$. Similarly, the output $\boldsymbol{y}$ is a matrix of shape $\langle N, Q\rangle$ (N entries encoded with shape $\langle Q\rangle$). For example, in image convolutions, M, P are the number of pixels and channels, respectively, of the input image, while N, Q are those of the output image. More generally $\boldsymbol{x}$ and $\boldsymbol{y}$ could be tensors — e.g. images are ternary tensors — but a tensor can always be flattened into a matrix, or even a vector (see Section 1.1). Matricisation, rather than full vectorisation, is used here in order to keep separate the uncontrolled, structural dimensions (width and height in images, of size M in input and N in output) from the controlled ones (channels, of size P in input and Q in output). By analogy with a simple linear layer, the most general form of a convolution layer is an arbitrary linear transform, given by

$$\boldsymbol{y}_{nq} \;=\; \sum\nolimits_{mp} \boldsymbol{x}_{mp}\boldsymbol{\Phi}_{mnpq} \tag{2}$$

Tensor $\boldsymbol{\Phi}$, of shape $\langle M, N, P, Q\rangle$, induces (linear) dependencies between each component of each input entry in $\boldsymbol{x}$ and each component of each output entry in $\boldsymbol{y}$. Using an arbitrary $\boldsymbol{\Phi}$ directly as parameter of the convolution is not satisfactory. First, its shape depends on the numbers M, N of input and output entries: M, N may vary for different instances of the data, or may be too large to be involved in the size of a parameter[2]. Furthermore, in Equation (2), the structural dependencies between the M input and N output entries are not captured. We propose to capture this structure as a tensor $\boldsymbol{A}$ of shape $\langle K, M, N\rangle$, for some integer K, and to constrain $\boldsymbol{\Phi}$ to be of the form:

$$\boldsymbol{\Phi} \;=\; \boldsymbol{A} \circ \boldsymbol{\Theta} \qquad \left(= \sum\nolimits_k \boldsymbol{A}_k \otimes \boldsymbol{\Theta}_k\right) \tag{3}$$

where $\boldsymbol{\Theta}$ is a tensor of shape $\langle K, P, Q\rangle$. Integer K is assumed to be a hyper-parameter controlled by the model, so $\boldsymbol{\Theta}$ has a fully controlled shape and is chosen as parameter of the convolution. Tensor $\boldsymbol{A}$ on the other hand characterises the structure underlying the convolution, and can be viewed as a family $(\boldsymbol{A}_k)_{k=1:K}$ of matrices (weighted graphs between input and output entries). The variety of existing convolution mechanisms derives from various choices for K and $\boldsymbol{A}$ (called resp. the *size* and *basis* of the convolution), which obey different intuitions in different domains. Examples are given below. But in general, combining Equations (2) and (3) together, we obtain a formula for convolution over arbitrary structures:

$$\boldsymbol{y} \;=\; \sum\nolimits_k \boldsymbol{A}_k^\top \boldsymbol{x} \boldsymbol{\Theta}_k \tag{4}$$

Note that our model of structural dependencies is flexible. If $(\boldsymbol{A}_k)_{k=1:K}$ is taken to be a basis of the whole space of matrices of shape $\langle M, N\rangle$, then by Proposition 1 any $\boldsymbol{\Phi}$ can be written as $\boldsymbol{A} \circ \boldsymbol{\Theta}$, and the resulting class of convolutions is the class of arbitrary linear transforms. But of course, this assumes $K{=}MN$, which is uncontrolled. At the other end of the spectrum, if $K{=}1$ and $\boldsymbol{A}_1$ is the identity matrix, the input entries are processed identically and fully independently, leading to a degenerate class of convolutions also known as $1{\times}1$ convolutions in the image domain. In fact, Equation (3) can be viewed as a truncated version of the factorisation of $\boldsymbol{\Phi}$ defined by Proposition 1 where family $(\boldsymbol{A}_k)_{k=1:K}$ is seen as a subset of a basis (of the whole space of matrices of shape $\langle M, N\rangle$), of which the other members are ignored. $(\boldsymbol{A}_k)_{k=1:K}$ act as “principal components”.

2 Some examples

2.1 Grid convolutions

A grid is the index set $\bar{S}$ associated with a given sequence of integers S. In the case of images, the archetypal grids, S is the sequence $\langle \text{width}, \text{height}\rangle$ of length $|S|{=}2$. We assume given some bijective

[2] The dependence on P, Q is not problematic, since these are hyper-parameters controlled by the model.

mapping $\omega{:}\bar{S}{\mapsto}\{1\cdots N\}$ where $N{=}|\bar{S}|$, for example the canonical bijection (see Section 1.1). In this way, an embedding of the whole grid, which would naturally be represented by a tensor of shape $S\langle L\rangle$ where each node in the grid is encoded as a vector of shape $\langle L\rangle$, can be matricised (see Section 1.1) into a matrix of shape $\langle N, L\rangle$ as required by our model. Let's first consider convolutions which preserve the grid, hence $M{=}N$.

Definition 1. *For any integer valued vector $\boldsymbol{d}{\in}\mathbb{Z}^{|S|}$, the* shift matrix *$\mathcal{A}_{\boldsymbol{d}}$ of shape $\langle N, N\rangle$ is*

$$(\mathcal{A}_{\boldsymbol{d}})_{mn} \triangleq \mathbb{I}[\omega^{-1}n - \omega^{-1}m = \boldsymbol{d}]$$

A grid convolution *of size K and basis $\boldsymbol{A}$ is one s.t. for each $k{\in}1{:}K$, $\boldsymbol{A}_k{=}\mathcal{A}_{\boldsymbol{\Delta}_k}$ for some $\boldsymbol{\Delta}_k{\in}\mathbb{Z}^{|S|}$.*

Thus, $\mathcal{A}_{\boldsymbol{d}}$ is the adjacency matrix of the relation: "node n is obtained from node m by a shift of $\boldsymbol{d}$ in the grid". With some padding conventions, Equation (4) for a grid convolution becomes,

$$\boldsymbol{y}_{(\omega(s))} = \sum_k \boldsymbol{x}_{(\omega(s-\boldsymbol{\Delta}_k))}\boldsymbol{\Theta}_k \qquad \forall s{\in}\bar{S}$$

The traditional grid (image) convolutions of "Convolutional Neural Networks" (CNNs) [12] are exactly obtained by choosing $\boldsymbol{\Delta}$ to be a regular right cuboid with possibly different strides and offsets in the different grid dimensions. Variants of grid convolutions which do not necessarily preserve the grid can also be captured in our framework using different choices of $\boldsymbol{\Delta}$, and variants of the shift matrices. This includes average pooling and dilated convolutions, where the output grid is a sub-sample of the input one (hence $M{=}N$ does not hold). However, max pooling is out of scope since, in our framework, convolutions are linear operators.

2.2 Graph convolutions

Let $\mathcal{G}$ be a graph over $\{1\cdots N\}$ given a priori. We assume $M{=}N$ (graph convolutions usually preserve the graph).

Definition 2. *A* graph convolution *of size K and basis $\boldsymbol{A}$ is one s.t. for each $k{\in}1{:}K$, matrix $\boldsymbol{A}_k$ is constructed from $\mathcal{G}$ by some procedure dependent on k.*

The traditional "Graph Convolution Networks" (GCNs) [9] are exactly obtained by choosing $K{=}1$ and $\boldsymbol{A}_1$ to be a normalised form of the adjacency matrix defined by $\mathcal{G}$. Constraining the size to 1 yields a very simple, efficient architecture, at the price of some expressiveness. For example, although grids can be represented as graphs, grid convolutions cannot be expressed as graph convolutions with a size restricted to 1.

In alternative definitions of graph convolution, the size is possibly greater than 1, and each $\boldsymbol{A}_k$ is computed from $\mathcal{G}$ in different ways. For example, in the full spectral analysis of graph convolution [4], each $\boldsymbol{A}_k$ is a Chebyshev polynomial of the scaled Laplacian matrix of $\mathcal{G}$, up to order K. In a simpler version [13], Chebyshev polynomials are replaced by elementary monomials, and $\boldsymbol{A}_k$ is simply the adjacency matrix of $\mathcal{G}$ raised to the power of k, capturing the random walks of length k through the graph. This notion can be extended to knowledge graphs [16].

3 Attention as content-based convolution

3.1 Content-based vs index-based convolution

In the previous examples of convolution, the basis tensor captures prior knowledge about the structural relationships between input and output entries through their indices. This is not the only option. Instead of relying solely on indices, the basis tensor of a convolution can also be computed from any content associated with the inputs and output entries. We propose a generic model to achieve this, and claim that it captures the essence of many attention mechanisms: [18, 7, 19, 2, 3, 10].

Definition 3. *An* attention mechanism *of type $\langle M, P, M', P'\rangle$ is a parametrised mapping with two input matrices, of shape $\langle M, P\rangle$ and $\langle M', P'\rangle$, respectively, yielding an output matrix of shape $\langle M, M'\rangle$. The input matrices represent M and M' entries encoded as vectors of shape $\langle P\rangle$ and $\langle P'\rangle$, respectively, and the output matrix represents an influence graph of the former on the latter, based on their encodings.*

Attention mechanisms can be added or multiplied term-wise, or transformed by term-wise, row-wise or column-wise normalisation. A common normalisation is column-wise softmax.

Definition 4. *An* attention convolution *of size K and basis $\boldsymbol{A}$ is one s.t. for each k=1:K, $\boldsymbol{A}_k$=$a(\boldsymbol{x}, \boldsymbol{z}; \boldsymbol{\Lambda}_k)$ for some attention mechanism a of type $\langle M, P, N, P' \rangle$ and some $\boldsymbol{\Lambda}_k$ in the parameter space of a. Here, $\boldsymbol{z}$ is an auxiliary input matrix of shape $\langle N, P' \rangle$, in addition to the main input $\boldsymbol{x}$ of shape $\langle M, P \rangle$. A* self-attention *convolution is an attention convolution where the auxiliary input $\boldsymbol{z}$ is taken to be a copy of the main input $\boldsymbol{x}$. This assumes N=M and P'=P.*

Equation (4) for an attention convolution becomes:

$$\boldsymbol{y} = \sum\nolimits_k a(\boldsymbol{x}, \boldsymbol{z}; \boldsymbol{\Lambda}_k)^\top \boldsymbol{x} \boldsymbol{\Theta}_k \tag{5}$$

Graph and grid convolutions described in the previous sections can be seen as degenerate cases of attention convolutions, in which the output of the mechanism does not depend on its input $\boldsymbol{x}, \boldsymbol{z}$, but solely on its parameter, given a priori (not learnt). As a result, they are linear in their main input $\boldsymbol{x}$ and independent of their auxiliary input $\boldsymbol{z}$, which is ignored. On the other hand, in the non degenerate case, the mechanism uses both its inputs, and the resulting convolutions are not linear in either of them. A commonly used attention mechanism is bi-affine attention [5], and in particular its bi-linear variant:

Definition 5. *Let $\boldsymbol{\Lambda}$ be a matrix of shape $\langle P, P' \rangle$. The* bi-linear *attention mechanism $\mathcal{A}$ of type $\langle M, P, M', P' \rangle$ and parameter $\boldsymbol{\Lambda}$ is defined by $\mathcal{A}(\boldsymbol{x}, \boldsymbol{x}'; \boldsymbol{\Lambda}) \triangleq \boldsymbol{x} \boldsymbol{\Lambda} \boldsymbol{x}'^T$*

3.2 Attention in Transformer

We now show how the attention model described by Equation (5) encompasses the scaled dot product attention used in the Transformer model of [18]. Attention is used in three distinct layers of the Transformer architecture. Two of them are instances of self-attention (on the source sequence and on the target sequence, respectively) while the third one is a cross-attention (the main input is the source sequence and the auxiliary input is the target sequence).

In all three cases, the scaled dot product attention mechanism used in Transformer essentially consists of a bi-linear attention, followed by a column-wise softmax. Furthermore, parameter $\boldsymbol{\Lambda}$ of the bilinear attention is constrained to be of the form

$$\boldsymbol{\Lambda}_k \;=\; \boldsymbol{\Lambda}_k^{(\text{key})} \boldsymbol{\Lambda}_k^{(\text{query})\top} \qquad \left(= \boldsymbol{\Lambda}_k^{(\text{key})\top} \circ \boldsymbol{\Lambda}_k^{(\text{query})\top}\right) \tag{6}$$

where both matrices $\boldsymbol{\Lambda}_k^{(\text{key})}, \boldsymbol{\Lambda}_k^{(\text{query})}$ are of shape $\langle P, D \rangle$. This can be viewed as a simple dimension reduction technique, since only $2PD$ parameters are required instead of P^2 for an arbitrary $\boldsymbol{\Lambda}_k$.

Now, Transformer attention introduces a seemingly richer mechanism to combine the different heads. Instead of simply summing them together as in Equation (5), it combines them with yet another linear layer:

$$\boldsymbol{y} = [\boldsymbol{h}_1, \ldots, \boldsymbol{h}_K] \boldsymbol{\Theta}^{(\text{O})\top} \qquad \text{where} \qquad \boldsymbol{h}_k \triangleq \boldsymbol{A}_k^\top \boldsymbol{x} \boldsymbol{\Theta}_k^{(\text{value})}$$

where $\boldsymbol{\Theta}^{(\text{O})}$ is a matrix of shape $\langle Q, KD \rangle$. In fact, this expression can be rewritten, splitting $\boldsymbol{\Theta}^{(\text{O})}$ into K blocks $(\boldsymbol{\Theta}_k^{(\text{O})})_{k=1:K}$ of shape $\langle Q, D \rangle$, as

$$\boldsymbol{y} = \sum\nolimits_k \boldsymbol{h}_k \boldsymbol{\Theta}_k^{(\text{O})\top} = \sum\nolimits_k \boldsymbol{A}_k^\top \boldsymbol{x} \boldsymbol{\Theta}_k^{(\text{value})} \boldsymbol{\Theta}_k^{(\text{O})\top}$$

In other words, it is strictly equivalent to the sum model of Equation (5), only with the constraint

$$\boldsymbol{\Theta}_k \;=\; \boldsymbol{\Theta}_k^{(\text{value})} \boldsymbol{\Theta}_k^{(\text{O})\top} \qquad \left(= \boldsymbol{\Theta}_k^{(\text{value})\top} \circ \boldsymbol{\Theta}_k^{(\text{O})\top}\right) \tag{7}$$

This constraint is not even specific to an attention model and could apply to any convolution. In fact, Equations (7) and (6) are meant to reduce the dimensionality of the parameters ($\boldsymbol{\Theta}$, and, in the case of attention, $\boldsymbol{\Lambda}$) by factorisation. Formally, they apply the exact same recipe as applied to $\boldsymbol{\Phi}$ in Equation (3), with the same purpose.

Finally, Transformers take the extreme approach of relying exclusively on content-based convolution ("attention is all you need"), so that any index-based information such as the relative position of the tokens must be incorporated into the content. They propose a smart but not completely intuitive scheme to achieve that, called "positional encoding". Alternatively, one or several additional heads with purely index-based basis matrices (e.g. shift matrices as in grid convolutions) could also be used, in complement to the attention heads.

References

[1] François Chollet. "Xception: Deep Learning with Depthwise Separable Convolutions". In: *arXiv:1610.02357 [cs]* (Oct. 7, 2016). arXiv: `1610.02357`.

[2] Yagmur Gizem Cinar et al. "Period-aware content attention RNNs for time series forecasting with missing values". In: *Neurocomputing* 312 (Oct. 27, 2018), pp. 177–186.

[3] Yagmur G. Cinar et al. "Position-based Content Attention for Time Series Forecasting with Sequence-to-sequence RNNs". In: *arXiv:1703.10089 [cs]* (Mar. 29, 2017). arXiv: `1703.10089`.

[4] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. "Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering". In: NIPS. Barcelona, Spain, 2016, p. 9.

[5] Timothy Dozat and Christopher D. Manning. "Deep Biaffine Attention for Neural Dependency Parsing." In: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. 2017.

[6] Vincent Dumoulin and Francesco Visin. "A guide to convolution arithmetic for deep learning". In: *arXiv:1603.07285 [cs, stat]* (Mar. 23, 2016). arXiv: `1603.07285`.

[7] Maha Elbayad, Laurent Besacier, and Jakob Verbeek. "Pervasive Attention: 2D Convolutional Neural Networks for Sequence-to-Sequence Prediction". In: *arXiv:1808.03867 [cs]* (Aug. 11, 2018). arXiv: `1808.03867`.

[8] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. 785 pp.

[9] Thomas N. Kipf and Max Welling. "Semi-Supervised Classification with Graph Convolutional Networks". In: *arXiv:1609.02907 [cs, stat]* (Sept. 9, 2016). arXiv: `1609.02907`.

[10] Wouter Kool, Herke van Hoof, and Max Welling. "Attention Solves Your TSP, Approximately". In: *arXiv:1803.08475 [cs, stat]* (Mar. 22, 2018). arXiv: `1803.08475`.

[11] Y. LeCun, K. Kavukcuoglu, and C. Farabet. "Convolutional networks and applications in vision". In: *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. Proceedings of 2010 IEEE International Symposium on Circuits and Systems. May 2010, pp. 253–256.

[12] Yann LeCun and Yoshua Bengio. "Convolutional networks for images, speech, and time series". In: *The Handbook of Brain Theory and Neural Networks*. Ed. by Michael A. Arbib. Cambridge, MA, USA: MIT Press, 1998, pp. 255–258.

[13] Yaguang Li et al. "Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting". In: *arXiv:1707.01926 [cs, stat]* (July 6, 2017). arXiv: `1707.01926`.

[14] Stéphane Mallat. "Understanding Deep Convolutional Networks". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374.2065 (Apr. 13, 2016), p. 20150203. arXiv: `1601.04920`.

[15] Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. "Introduction to Tensor Decompositions and their Applications in Machine Learning". In: *arXiv:1711.10781 [cs, stat]* (Nov. 29, 2017). arXiv: `1711.10781`.

[16] Michael Schlichtkrull et al. "Modeling Relational Data with Graph Convolutional Networks". In: *arXiv:1703.06103 [cs, stat]* (Mar. 17, 2017). arXiv: `1703.06103`.

[17] Xingjian Shi et al. "Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting". In: *arXiv:1506.04214 [cs]* (June 12, 2015). arXiv: `1506.04214`.

[18] Ashish Vaswani et al. "Attention Is All You Need". In: *arXiv:1706.03762 [cs]* (June 12, 2017). arXiv: `1706.03762`.

[19] Petar Veličković et al. "Graph Attention Networks". In: *arXiv:1710.10903 [cs, stat]* (Oct. 30, 2017). arXiv: `1710.10903`.

A Supplementary material

A.1 Some useful properties of tensors

A.1.1 Definitions of the main concepts

Slicing: if $\boldsymbol{a}$ is a tensor of shape ST, then its slice at $s{\in}\bar{S}$ denoted $\boldsymbol{a}_s$ is the tensor of shape T defined by

$$(\boldsymbol{a}_s)_t \;\triangleq\; \boldsymbol{a}_{st}$$

Flattening: if $\omega{:}\bar{S}{\mapsto}\{1\cdots K\}$ is a bijective mapping (hence $K{=}|\bar{S}|$) and $\boldsymbol{a}$ is a tensor of shape ST, then its ω-flattening denoted $\boldsymbol{a}^{[\omega]}$ is the tensor of shape $\langle K\rangle T$ defined by

$$\boldsymbol{a}^{[\omega]}_{\langle k\rangle t} \;\triangleq\; \boldsymbol{a}_{(\omega^{-1}k)t}$$

When T is of length 1 (resp. 0), then $\boldsymbol{a}^{[\omega]}$ is a matrix (resp. a vector) and flattening is then called *matricisation* (resp. *vectorisation*) [15]. We often use the *canonical bijection* ω_S (see [15]) defined for each $s{\in}\bar{S}$ by

$$\omega_S(s) \;\triangleq\; 1 + \sum\nolimits_{i=1:|S|}(s_i-1)\prod\nolimits_{j=i+1:|S|}S_j$$

Outer product: if $\boldsymbol{a},\boldsymbol{b}$ are tensors of shape S and T, respectively, their outer product (a.k.a. tensor product) denoted $\boldsymbol{a}\otimes\boldsymbol{b}$ is a tensor of shape ST defined by:

$$(\boldsymbol{a}\otimes\boldsymbol{b})_{st} \;\triangleq\; \boldsymbol{a}_s\boldsymbol{b}_t$$

A.1.2 Proof of Proposition 1

Proof. Observe that $\boldsymbol{\Theta}\mapsto\boldsymbol{a}\circ\boldsymbol{\Theta}$ is a linear mapping from the space of tensors of shape $\langle K\rangle T$ into the space of tensors of shape ST. The assumption ($\boldsymbol{a}$ is a basis) implies that it is injective, and since the two spaces have the same dimension, the mapping is an isomorphism. □

A.2 Composition of convolutions

Proposition 2. *Given two convolutions of size K',K'', basis $\boldsymbol{A}',\boldsymbol{A}''$, parameter $\boldsymbol{\Theta}',\boldsymbol{\Theta}''$, respectively, their composition, when the dimensions match (i.e. $\langle N',Q'\rangle{=}\langle M'',P''\rangle$), is a convolution of size K, basis $\boldsymbol{A}$, parameter $\boldsymbol{\Theta}$ where*

$$K = K'K'' \qquad \boldsymbol{A}_{\omega(k',k'')} = \boldsymbol{A}'_{k'}\boldsymbol{A}''_{k''} \qquad \boldsymbol{\Theta}_{\omega(k',k'')} = \boldsymbol{\Theta}'_{k'}\boldsymbol{\Theta}''_{k''}$$

and ω is a bijective mapping $\{1\cdots K'\}{\times}\{1\cdots K''\}{\mapsto}\{1\cdots K\}$, e.g. the canonical bijection $\omega_{\langle K',K''\rangle}$.

Proof. Simple application of Equation (4). □

A.3 Separable convolutions

The parameter $\boldsymbol{\Theta}$ of a convolution, of shape $\langle K,P,Q\rangle$, although controlled, may still be too large and it may be useful to constrain it further, e.g. by imposing it to be of the factorised form

$$\boldsymbol{\Theta} \;=\; \boldsymbol{\Theta}^{(\text{basis})}\circ\boldsymbol{\Theta}^{(\text{channel})} \tag{8}$$

where $\boldsymbol{\Theta}^{(\text{basis})}$ is a matrix of shape $\langle H,K\rangle$ (for some integer H) and $\boldsymbol{\Theta}^{(\text{channel})}$ a tensor of shape $\langle H,P,Q\rangle$. The resulting convolutions are said to be *separable*. This generalises the so called depth-wise separable convolutions, common in the image domain [1], which are the special case $H{=}1$. The parameter size of a separable convolution is $H(K{+}PQ)$ instead of KPQ for an arbitrary one. An alternative form of dimension reduction is discussed in Section 3.2.

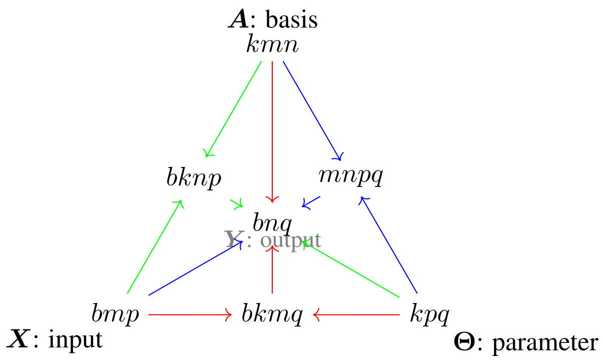

Figure 1: A representation of three alternatives (red-green-blue, each starting from one side of the triangle) to compute a convolution $\boldsymbol{Y}$ (order-3 tensor at the centre). The vertices of the triangle are the order-3 tensors involved ($\boldsymbol{X}$: input, $\boldsymbol{\Theta}$: parameter, $\boldsymbol{A}$: basis) with their respective indices (b: batch index, m/n: input/output entry, p/q: input/output channel, k: basis index). The arrows represent sum-product operations in the so called Einstein's notation. For example, the two arrows $bmp, kpq \to bkmq$ (bottom) represent an operation yielding the order-4 tensor $R_{bkmq} = \sum_p \boldsymbol{X}_{bmp}\boldsymbol{\Theta}_{kpq}$.

A.4 A note on the computation of convolutions

In practice, the input and output entries of convolutions are usually batched. Batched input $\boldsymbol{X}$ and output $\boldsymbol{Y}$ are given by tensors of shape $\langle B, M, P\rangle$ and $\langle B, N, Q\rangle$, respectively, where B is the batch size. Figure 1 shows the three alternatives to compute $\boldsymbol{Y}$ (at the centre of the triangle) as a function of $\boldsymbol{A}, \boldsymbol{X}, \boldsymbol{\Theta}$ (on the vertices of the triangle), according to the convolution formula of Equation (4) extended to batches:

$$\boldsymbol{Y}_b = \sum\nolimits_k \boldsymbol{A}_k^\top \boldsymbol{X}_b \boldsymbol{\Theta}_k$$

Which of these alternatives should be used essentially depends on the respective dimensions M, P, N, Q, B, K. In any case, operations involving the basis tensor $\boldsymbol{A}$ may require a specific treatment, since it is usually very sparse, and sometimes possesses a regularity which can be exploited for optimal computation, as in the case of the "shift matrices" of grid convolution (see below).

A.5 Illustration of the shift matrices in grid convolutions

Shift matrices constitute the basis of grid convolutions. For example, a shift by 2 in a 1-D grid of dimension 10 is captured by the matrix:

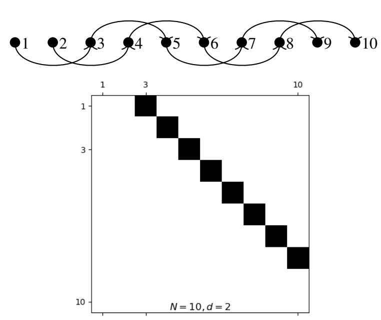

A shift by $(2, 4)$ in a 2-D grid of dimensions $8{\times}10$ (flattened by the canonical mapping into indices 1:80) is captured by the matrix:

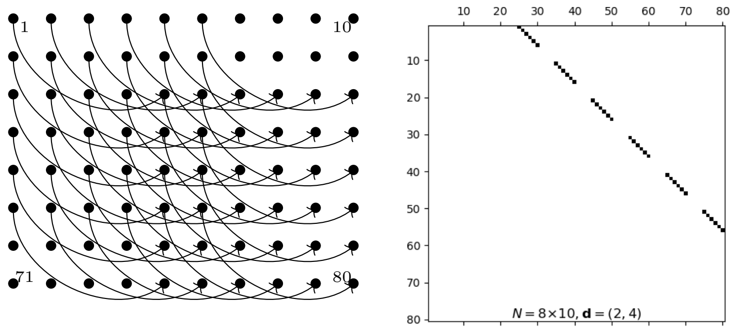

A.6 Attention in Graph Attention Networks

Graph attention networks [19] are based on self-attention convolutions which use a variant of Equation (5): the attention mechanism a is passed the term $\boldsymbol{x}\boldsymbol{\Theta_k}$ instead of $\boldsymbol{x}$ as both main and auxiliary input. That term is already computed in Equation (4) when starting from the bottom of the triangle in Figure 1. The attention mechanism proposed in [19] starts with a simplified variant of bi-affine mechanism (without bi-linear term):

$$a(\boldsymbol{x}, \boldsymbol{x}'; \boldsymbol{\lambda}, \boldsymbol{\lambda}', \xi) \;\triangleq\; (\boldsymbol{x}\boldsymbol{\lambda}) \otimes \mathbf{1}_{M'} + \mathbf{1}_M \otimes (\boldsymbol{x}'\boldsymbol{\lambda}') + \xi \mathbf{1}_M \otimes \mathbf{1}_{M'} \tag{9}$$

where ξ is a scalar, and $\boldsymbol{\lambda}, \boldsymbol{\lambda}'$ are vectors of shapes, respectively, $\langle P\rangle, \langle P'\rangle$. The output is then masked by a graph given a priori, limiting the zone of influence on each node to a neighbourhood of that node, which is more realistic in the case of large structures such as publication networks (up to 50,000 nodes in their experiments). The masked output is then normalised by a term-wise "leaky ReLU" followed by a column-wise softmax. These choices can be motivated to some extent by properties of their simplified bi-affine model. Indeed, observe that the masked values are still masked after the leaky ReLU (which would not be the case with a plain ReLU), and, after the softmax, the masked values become 0, cancelling the influence of the corresponding inputs. Skipping ReLU altogether before the softmax would make the term involving $\boldsymbol{\lambda}'$ in the sum of Equation (9) redundant: it is constant along each column, and softmax is invariant to an additive constant.