

HexCNN: A Framework for Native Hexagonal Convolutional Neural Networks

Yunxiang Zhao[†], QiuHong Ke[†], Flip Korn[‡], Jianzhong Qi[†], Rui Zhang^{†*}

[†]The University of Melbourne, Australia, [‡]Google Research, USA

{yunxiangz@student., qiuhong.ke@, jianzhong.qi@, rui.zhang@ }unimelb.edu.au, flip@google.com

Abstract—Hexagonal CNN models have shown superior performance in applications such as IACT data analysis and aerial scene classification due to their better rotation symmetry and reduced anisotropy. In order to realize hexagonal processing, existing studies mainly use the ZeroOut method to imitate hexagonal processing, which causes substantial memory and computation overheads. We address this deficiency with a novel native hexagonal CNN framework named *HexCNN*. HexCNN takes hexagon-shaped input and performs forward and backward propagation on the original form of the input based on hexagon-shaped filters, hence avoiding computation and memory overheads caused by imitation. For applications with rectangle-shaped input but require hexagonal processing, HexCNN can be applied by padding the input into hexagon-shape as preprocessing. In this case, we show that the time and space efficiency of HexCNN still outperforms existing hexagonal CNN methods substantially. Experimental results show that compared with the state-of-the-art models, which imitate hexagonal processing but using rectangle-shaped filters, HexCNN reduces the training time by up to 42.2%. Meanwhile, HexCNN saves the memory space cost by up to 25% and 41.7% for loading the input and performing convolution, respectively.

Index Terms—Hexagonal Convolution, Convolutional Neural Networks, Deep Learning

I. INTRODUCTION

Recent studies show that compared with traditional rectangle-based CNN models, CNN models with hexagon-shaped filters achieve better performance in applications such as Imaging Atmospheric Cherenkov Telescope (IACT) data analysis [2], [19], [20], [23], Hex move-prediction [27], and IcdCube data analysis [8]. Applying hexagonal filters in group CNNs can even surpass the performance of traditional CNN models with image classification tasks on data sets such as CIFAR-10 [6], [22], [26].

To realize hexagonal processing, most existing studies apply rectangle-shaped filters with the ZeroOut method to imitate hexagonal processing [6], [17], [11]. We refer to these models as hexagon-imitation models. These models, however, require a padding strategy as a pre-processing for the input. Fig. 1b and Fig. 1c illustrate the idea of such padding on the hexagon and rectangle-shaped input, and resampling is required if the input is not hexagonal grids such as Fig 1e. The padded area (dark elements in Fig. 1b and Fig. 1c) is not needed for the output but is computed as part of the input throughout the hexagon-imitation models. This is due to the restrictions of

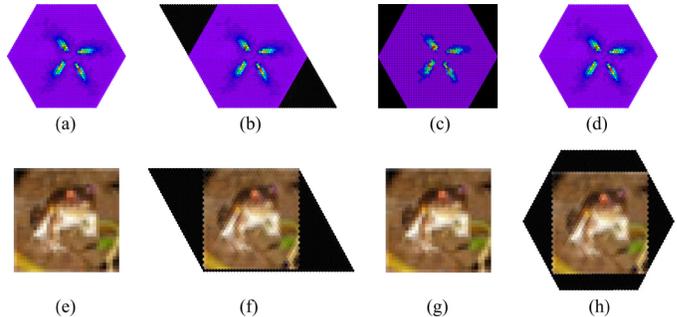


Fig. 1: The original inputs (a, e), and the results after pre-processing using ZeroOut [6] (b, f), Quasi-H [26] (c, g), and HexCNN (d, h). (a) a hexagonal telescope image of gamma-ray events [18]. (e) a rectangle-shaped image from CIFAR-10 [13] (best view in color).

existing deep learning frameworks where data is represented in a rectangular way. After padding the input, the rectangle-shaped filters in the network can be used to process the input without having to accommodate the boundary of hexagonal grids. To achieve hexagonal processing in convolution and pooling, specific positions of the rectangle-shaped filter are set to zero to eliminate their influence on the output (e.g., the two blue-colored elements with “0” in Fig. 2a). These zeroed-out elements, however, lead to unnecessary memory and computation costs in hexagon-imitation models. To sum up, the limitation of hexagon-imitation models is that the padding strategy and the ZeroOut manner in convolution and pooling cause significant memory and computation overhead for both hexagon and rectangle-shaped input.

In this paper, we introduce a new framework called HexCNN to address this limitation. HexCNN takes hexagon-shaped input and performs hexagonal forward and backward propagation on the original form of the input based on hexagon-shaped filters. We refer to this method as “native hexagonal processing”. Compared with previous methods, the proposed native hexagonal processing eliminates the memory and computation overhead from the padding and ZeroOut operations. As shown in Fig. 1d, for hexagon-shaped input, the proposed HexCNN does not perform padding as done in hexagon-imitation models. The convolution is achieved using hexagon-shaped filters, which bypasses the unnecessary computation of the zeroed-out elements, as shown in Fig. 2b. Therefore, HexCNN saves memory and computation costs significantly when loading the input and performing

*Rui Zhang is the corresponding author

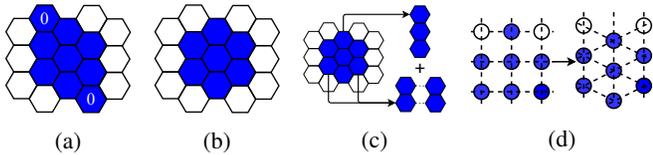


Fig. 2: Hexagonal convolution via (a) ZeroOut [6], (b) HexCNN, (c) HexagDLy [25], and (d) Quasi-H [26]. Elements in blue denote the filter (best view in color).

filter related operations such as convolution, pooling, and backpropagation. For rectangle-shaped input such as images, HexCNN incurs resampling and padding, and the incurred unnecessary area has a similar size to those in hexagon-imitation models, as shown in Fig. 1f, and Fig. 1h. In this case, HexCNN still has less overhead due to the reduced space and computation cost on native hexagonal convolution. We detail the comparison between HexCNN and hexagon-imitation models on both hexagon and rectangle-shaped input in Section V. We summarize our contributions as follows:

- We propose a framework for realizing native hexagonal processing named HexCNN. HexCNN provides native hexagonal CNN operation algorithms for both forward and backward propagation to achieve superior time and space efficiency.
- We propose an algorithm for transforming hexagonal convolution into matrix multiplication so as to exploit matrix optimization algorithms for reducing the time consumption of hexagonal convolution.
- We implement HexCNN on TensorFlow and perform extensive experiments to evaluate its efficiency. The experimental results show that HexCNN outperforms the state-of-the-art models, which imitate hexagonal processing, by reducing up to 42.2% of the training time. Meanwhile, HexCNN saves the memory cost by up to 25% and 41.7% for loading the input and performing convolution, respectively. For applications with rectangle-shaped input but require hexagonal processing, HexCNN still outperforms existing hexagonal CNN methods by reducing more than 20% of the training time.

We organize the rest of this paper as follows. We review related work in Section II and present the fundamentals of HexCNN in Section III. We describe the algorithm for transforming hexagonal convolution to matrix multiplication in Section IV, and report the experimental results in Section V. We conclude the paper in Section VI.

II. RELATED WORK

In this section, we review studies that imitate hexagonal processing. Existing hexagon-imitation models only work on rectangle-shaped input. The rectangle-shaped input can be hexagonal grids or square grids, as shown in Fig. 2a and Fig. 2d. We classify existing hexagon-imitation models into two categories. The first category is for input that is rectangle-shaped hexagonal grids (Section II-A) and the second category is for input that is rectangle-shaped square grids (Section II-B).

A. Hexagonal CNN Models on Hexagonal Grids

In order to realize hexagonal processing on hexagonal grids, the ZeroOut method has been widely used [6], [17]. Young et al. [27] propose a hexagon-imitation model for a chess game Hex and apply ZeroOut to achieve hexagonal processing. Ke et al. [11] propose a hexagon-imitation model for forecasting the ride-sourcing supply and demand, which splits the input into a hexagonal lattice and then realizes hexagonal processing by applying the rectangle-shaped filters together with transform matrices. The concept of transform matrices is similar to that of ZeroOut. A similar method was proposed by Jacquemont [9], where the values of specific positions in the filter are eliminated by a mask matrix. For convolution, the models above mainly use rectangle-shaped filters together with ZeroOut methods to achieve hexagonal processing [12]. We use **ZeroOut** to represent above models.

Steppa and Holch [25] propose **HexagDLy**, which implements hexagonal convolution and pooling by combining multiple rectangle-shaped filters with different dilations and sizes, as shown in Fig. 2c. HexagDLy does not use the ZeroOut method during the convolution and achieves better efficiency along with the increase of filter size due to the advantage of divide-and-conquer. However, HexagDLy only supports convolution and pooling.

B. Hexagonal CNN Models on Square Grids

To perform hexagonal CNNs on square grids such as images and videos [6], [12], [26], it is required to first transform the input from square grids to hexagonal grids. Resampling and zero-padding have been widely used for this transformation, which introduces unnecessary computation and space costs, as illustrated in Fig. 1f. For convolution, existing work mainly uses rectangle-shaped filters together with the ZeroOut method to achieve hexagonal processing [12], as shown in Fig. 2a. We consider them as **ZeroOut**, which is the same as that in Section II-A. The only difference is that the ZeroOut here takes square grids as the input, and the one in Section II-A takes hexagonal grids.

Sun et al. [26] propose quasi-hexagonal kernels (**Quasi-H**) for different convolution layers. Quasi-H is not hexagonal processing in the strict sense and can only be considered as hexagonal processing on deep CNN models with multiple convolutional layers (each layer has multiple filters). Applying Quasi-H on rectangle-shaped input for simulating hexagonal convolution does not incur resampling or unnecessary padding, as shown in Fig 1g. Applying Quasi-H to hexagon-shaped input requires resampling and padding for pre-processing, which the same as ZeroOut, as shown in Fig. 1c.

III. HEXCNN

Existing deep learning frameworks such as TensorFlow, PyTorch and Caffe are efficient and maintained by open source developers to keep optimizing their performance [1], [10], [21], [28]. Instead of implementing HexCNN from scratch, we have implemented it based on the TensorFlow by adapting the core libraries and third-party Eigen library of

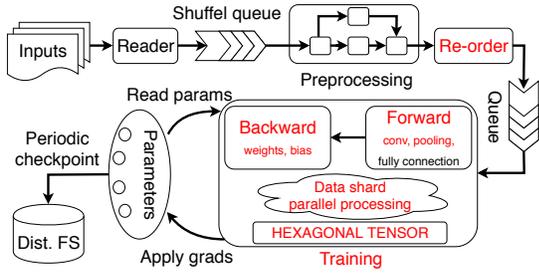


Fig. 3: The HexCNN dataflow graph inherited from TensorFlow [1]. HexCNN achieves native hexagonal processing for forward and backward propagation, and the low-level parallel and data sharding algorithms, which are marked in red (best view in color).

TensorFlow, so that HexCNN retains the general applicability of the framework. This way, implementing hexagonal CNN models of various neural network structures on HexCNN is similar to implementing rectangular CNN models of various structures on TensorFlow. Therefore, HexCNN is a framework for implementing hexagonal CNN models, i.e., CNN models that directly take hexagonal inputs. The tensor in TensorFlow is rectangle-shaped, and all the tensor-related steps need to be rewritten for hexagonal processing. In TensorFlow, the input is rectangular. For example, an RGB image is represented as a 3D tensor [height, width, channel], while in HexCNN, the input is hexagonal, and the corresponding tensor is represented as [side_length, channel]. In this section, we first give an overview of HexCNN followed by its data representation and fundamental operations.

A. Overview

Fig. 3 is the schematic dataflow graph of HexCNN, the “Re-order” operation feeds an input into a hexagon-shaped tensor. For optimization purposes, TensorFlow performs parallel computing and data sharding, which are highly related to the shape of the input, filters, and how they are stored in the memory. HexCNN adapts these algorithms according to the hexagonal data representation, which is detailed in Section III-B. HexCNN applies hexagon-shaped filters for convolution and pooling during forward-propagation. For backpropagation, HexCNN converts the gradient computation into convolution, and hence it is also based on hexagonal convolution. We detail the hexagonal operations in Section III-C.

B. Data Representation and Storage

One of the fundamentals of HexCNN is data representation, which determines the way to perform hexagonal processing. HexCNN takes the idea of the Axial coordinate system [5], [6], and moves the origin to the top-left corner, as shown in Fig. 4a. Suppose the side length of the hexagon-shaped input is k . Then the length of the first k rows increases by one for each row and reduces by one after the k^{th} row.

To reduce the memory and computation costs of hexagonal processing, HexCNN transforms the hexagon-shaped input into a column-based vector during the data storage phase. Compared with hexagon-imitation models, HexCNN does not

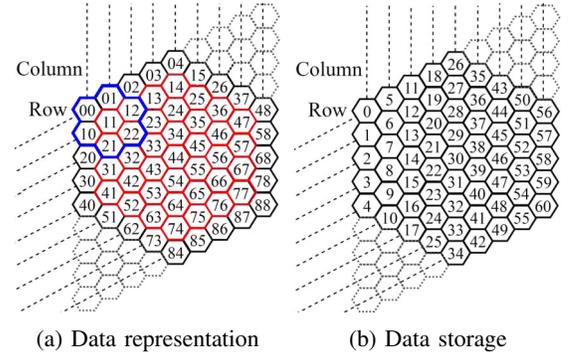


Fig. 4: Data representation (a) and storage (b) of HexCNN. HexCNN is based on 2D Axial coordinate system, and transform the input into a column major vector for data storage (best view in color).

require unnecessary spaces for storing both the input and convolution filters. Thus, it has lower memory and computation costs. The dashed elements in Fig. 4 illustrates the space that HexCNN saves for storing the input and filters in comparison to existing hexagon-imitation models.

In convolution, existing hexagon-imitation models mainly use ZeroOut to realize hexagonal convolution. This approach, requires extra space for packing the filter, so as to run on existing deep learning frameworks, as shown in Fig. 2a. In pooling, both HexCNN and hexagon-imitation models iteratively go through the input, and they have the same space complexity. However, extra computation is required by hexagon-imitation models due to the padding strategy applied to the input.

C. Fundamental Operations

For rectangle-shaped tensors, each column has the same number of elements, and hence it is easy to perform CNN operations such as convolution and pooling. For hexagon-shaped tensors, different columns have different length and need to be considered separately. Based on the hexagonal data representation and storage in Fig. 4, HexCNN modifies the fundamental operations of CNN, including both forward and backward propagation algorithms [4], [14], [15].

Forward-propagation: Forward-propagation mainly contains convolution, pooling, and fully connection operations. After flattening the output of convolution and pooling layers, the fully connected layer works the same as that for traditional rectangle-based CNN models. The convolution of HexCNN is represented as Equation 1,

$$O_{u,v} = \sum_{j=v}^{v+2 \cdot l_k - 1} \sum_{i=i_s}^{i_e} I_{i,j} \cdot k_{i-i_s, j-v}$$

$$i_s = \begin{cases} j - l_I + 1 + u, & j \geq l_I \\ u, & j < l_I \end{cases} \quad (1)$$

$$i_e = \begin{cases} i_s + 2 \cdot l_k - 1, & j \geq v + l_k \\ i_s + l_k + (j - v), & j < v + l_k \end{cases}$$

where input I and filter k are two-dimensional tensors in this example. l_k and l_I denote the side length of the filter and the input, respectively. i_s and i_e denote the start and the end indices for each column in the sliding window, which are u

and v dependant. i_s increases with a step length of one after the first l_I columns, and the range for index j increases with a step length of one for the first l_k columns and then decreases for the onward columns.

For max-pooling and average-pooling, we either pick the maximum or average the value of all elements within a sliding window to obtain an output element. Take maxpooling as example, we iteratively fetch the elements from the red bounded area (output) in Fig. 4a and obtain the corresponding sliding window from the input. More formally, the procedure is represented as:

$$O_{u,v} = \max_{v \leq j < v+2 \cdot l_k - 1, i_s \leq i < i_e} I_{i,j} \quad (2)$$

Backpropagation: In backpropagation, the operations for dense layers are the same as those for rectangle-based frameworks. For convolutional and pooling layers, we backpropagate the error and update the weights of each layer.

For back propagating the error in convolutional layers, HexCNN upsamples the error δ^c to $\hat{\delta}^c$ to get the same size as the input of the current layer I . HexCNN then performs full convolution on the upsampled $\hat{\delta}^c$ and the filter k that has been rotated for 180 degrees. After the full convolution, the former layer's error δ^p is the Hadamard product [7] with the derivative of the former layer's activation function σ^p :

$$\delta^p = \hat{\delta}^c * \text{rot}180(k) \odot \sigma'^p \quad (3)$$

where function $\text{rot}180$ rotates a matrix by 180 degrees, operator $*$ denotes the full convolution, and operator \odot computes the Hadamard product. For pooling layers, we take max-pooling as an example to illustrate how HexCNN works. HexCNN first upsamples δ^c to $\hat{\delta}^c$, and figures out which position in I corresponds to a specific value in δ^c (the perceptive field of δ^c). HexCNN then transfers the value in δ^c to that position in $\hat{\delta}^c$, and sets the rest positions in $\hat{\delta}^c$ as zeros. As shown in Equation 4, each value in the former layer's error $\delta_{u,v}^p$ is the element-wise product of $\hat{\delta}_{u,v}^c$ and the sum of those values in δ^c that come from $I_{u,v}$.

$$\delta_{u,v}^p = \hat{\delta}_{u,v}^c \cdot \sum_{I_{u,v} \rightarrow \delta_{i,j}} \delta_{i,j} \quad (4)$$

where $I_{u,v} \rightarrow \delta_{i,j}$ denotes elements in I that are within the perceptive field of $\delta_{i,j}$, “ \cdot ” computes the element-wise product.

For updating the weights, we propose derivative computation algorithms for native hexagonal processing. HexCNN computes the partial derivative for both input and filters in convolutional layers, except the initial input of the network. The partial derivative of filters is for updating the weights, and the partial derivative of the input is for updating the weights of the former layer. (i) The partial derivative for a filter k can be represented as the convolution between the input I and the error δ (with the same size as the output O) according to the chain rule. We assume that “VALID” padding is applied. Then the partial derivative for the filter k can be represented as:

$$\frac{\partial E}{\partial k_{u,v}} = \sum_{j=v}^{v+2 \cdot l_O - 1} \sum_{i=i_s}^{i_e} I_{i,j} \cdot \delta_{i-i_s, j-v} \quad (5)$$

where i_s and i_e are computed similar to Equation 1. (ii) Com-

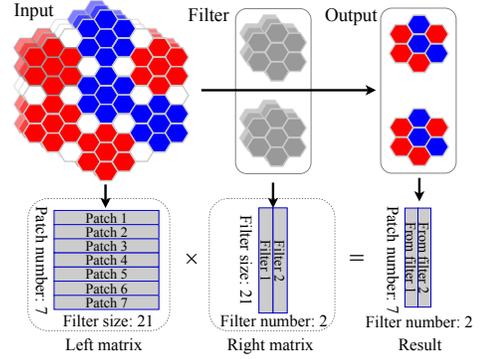


Fig. 5: Hexagonal convolution to matrix multiplication. The red and blue color are used for locating differentiating different filters (sliding windows) only (best view in color).

puting the partial derivative of the input can be transformed into the full convolution between the error δ (with the same size of the output O) and the filter k . If the output size is smaller than the input size, HexCNN pads zeros around the δ to achieve the same size as the input before the full convolution. We summarize the procedure as:

$$\begin{aligned} \frac{\partial E}{\partial I_{u,v}} &= \sum_{j=v-l_k+1}^{v+l_k} \sum_{i=i_s}^{i_e} \delta_{i,j} \cdot k_{i-i_s, j-v+l_k-1} \\ i_s &= \begin{cases} j - l_O + u - l_k + 2, & j \geq l_O \\ u - l_k + 1, & j < l_O \end{cases} \\ i_e &= \begin{cases} i_s + 2 \cdot l_k - 1, & j \geq v \\ i_s + 2 \cdot l_k + j - v - 1, & j < v \end{cases} \end{aligned} \quad (6)$$

where i_s and i_e are column index i and filter size l_k related. $\delta_{i,j}$ returns zero when i_s and i_e out of the range of δ and this is the reason that we pad zeros around δ .

IV. TRANSFORMATION FROM HEXAGONAL CONVOLUTION TO MATRIX MULTIPLICATION

Existing deep learning frameworks transform convolution into matrix multiplication to speed up the computation. In this section, we present how HexCNN transforms hexagonal convolution into matrix multiplication. HexCNN splits the input into multiple patches according to the filter size and stride. Suppose the hexagon-shaped input has a side length l_I and a channel c . The hexagonal convolution filters have a side length l_k , and the stride during convolution is s . HexCNN computes the number of patches by $l_p \cdot (l_p - 1) + 1$, where $l_p = (l_I - l_k) / s + 1$. Take Fig. 5 as an example, where $l_I = 5$, $c = 3$, $l_k = 2$, and $s = 3$. The value of l_p is two, and hence there are seven patches in the input (the red and blue colors are used for differentiation only).

To transform the hexagonal convolution into matrix multiplication, HexCNN first obtains the number of patches according to the input shape, the stride, and padding strategy. It then re-orders the elements in each patch as a vector in column-major to obtain the left matrix in Fig. 5. For the right matrix in Fig. 5, the width is the number of filters, and the height is the filter size. After transforming the input and

filters into the left and right matrices, HexCNN obtains the convolutional output via matrix multiplication. To speed up the matrix multiplication, HexCNN then applies the GotoBLAS library [3] for acceleration.

V. EXPERIMENTS

Hexagon-imitation models have shown better accuracy than traditional rectangular CNN models on various applications and data sets [6], [11], [17], [26]. HexCNN has the same output as hexagon-imitation models, and we focus on the time and space efficiency of HexCNN in this section. We first evaluate the time efficiency of HexCNN and then evaluate the space efficiency. We implement HexCNN by adapting the third-party Eigen library for the TensorFlow CPU version. We are unable to test the TensorFlow GPU version because the GPU version is based on cuDNN, which is a closed source library (same for Caffe and PyTorch).

Experimental setup: The PC that we used is equipped with Intel(R) Core(TM) i7-7700 @ 3.60 GHz \times 8 with 32GB memory. We implement two hexagon-imitation models based on TensorFlow. The first is the widely used ZeroOut [6], and the second is the Quasi-H [26]. Note that Quasi-H simulates hexagonal processing on deep CNN models, and it is not hexagonal processing in the strict sense (see Section II-B for more details). We consider it as one of the baselines because it simulates hexagonal processing in certain circumstances.

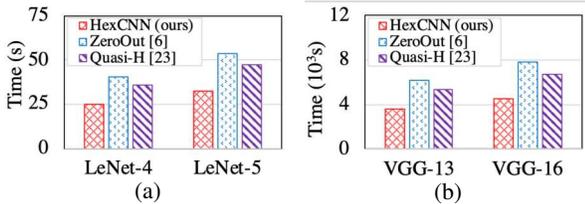


Fig. 6: Time consumption of HexCNN, ZeroOut and Quasi-H on (a) shallow and (b) deep hexagonal CNN models (best view in color).

A. Time Efficiency Evaluation

To evaluate the time efficiency of HexCNN, we compare the training time of a single batch (batch size as 40) in LeNet-4 [16], LeNet-5 [16], VGG-13, and VGG-16 [24]. We keep the structure of these CNN models but perform hexagonal processing. The synthetic data set that we used contains 2,000 hexagon-shaped input with a channel of three and a side length of 256 (511×511 for ZeroOut, 511×443 Quasi-H to keep the equivalent information). We obtain the data set from MNIST [16] by resizing images into a size of 511×511 followed by resampling them to hexagonal grids. We then crop them to the size needed.

Fig. 6a shows the time consumption of different methods on two shallow hexagonal CNN models LeNet-4 and LeNet-5. HexCNN saves the training time by 38.7% and 39.7% compared with ZeroOut, respectively. Comparing with Quasi-H, HexCNN saves the training time by 30.8% and 31.4%, respectively. For deep hexagonal CNN models VGG-13 and VGG-16, HexCNN saves the training time by 41.3% and 42.2% compared with ZeroOut, respectively. In comparison

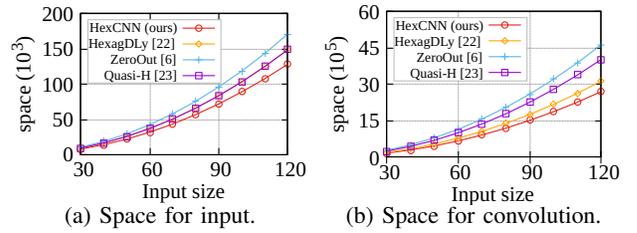


Fig. 7: The space required by HexCNN, HexagDLy, ZeroOut, and Quasi-H for (a) loading one input and (b) preparing the corresponding matrices for convolution (best view in color).

with Quasi-H, HexCNN saves the training time by 32.6% and 33.2%, respectively. We find that on the same network structure, the more layers the models have, the more time that HexCNN saves. This is because the start-up cost is relatively stable, and the more layers, the more relative time that HexCNN saves.

B. Space Efficiency Evaluation

We compute the memory consumption of HexCNN, HexagDLy, ZeroOut, and Quasi-H for loading the input and the matrices for performing convolution to evaluate the space efficiency of HexCNN. We set the input channel as three, and the side length x varies from 30 to 120. In convolution, we set the filter number as one, the side length of hexagon-shaped filters as two, and keep the stride as one.

When loading the input, (i) ZeroOut method pads the input into a parallelogram with the side length of $2x - 1$ to keep the equivalent information. Fig. 1b illustrates the input after padding, where the side length of the parallelogram is twice as that for the hexagonal input in Fig. 1a. (ii) Quasi-H re-samples and pads the input into a rectangle with a size of $2x - 1, \sqrt{3}x$ to keep the equivalent information, as shown in Fig. 1c. (iii) HexagDLy first rotates the hexagonal array in Cartesian coordinates to achieve a vertical alignment of neighboring elements and then aligns horizontally by shifting every second column upwards by half the distance between neighboring pixels. HexagDLy does not incur resampling because it takes hexagonal grids directly, and the input becomes a rectangle-shaped hexagonal grid with a size of $2x - 1$ under the offset coordinate system. In convolution, ZeroOut and Quasi-H have a filter size of three to keep the equivalent information as HexCNN. For HexagDLy, it divides the hexagon-shaped filters into multiple rectangle-shaped filters with different sizes, as shown in Fig. 2c. Therefore, the space for filters in HexagDLy is the same as that for HexCNN.

As shown in Fig. 7, HexCNN saves the space cost significantly compared with existing hexagon-imitation models. With the increase of input size, HexCNN reduces the space consumption for storing input and performing convolution by 13.8% compared with HexagDLy. Compared with ZeroOut, HexCNN reduces the space consumption for storing input and performing convolution by 25% and 41.7%, respectively. Compared with Quasi-H, HexCNN reduces the space consumption for storing the input and performing convolution by 13.8% and 32.9%, respectively.

C. HexCNN for Rectangle-Shaped Input

HexCNN achieves superior efficiency for hexagon-shaped input due to the seamless representation. For rectangle-shaped input, HexCNN incurs resampling and padding. Fig. 1h shows the padding cost in HexCNN for square input. Suppose the side length of a square input is x , the side length of the corresponding hexagon that exactly covers the square input is y . Then y has the following restriction: $y \geq \lceil (3x + 1)/4 \rceil$, and the padded area when x is approaching infinity is $0.563x^2$.

For ZeroOut, the padded area is $0.577x^2$, as illustrated in Fig. 1f. Although the extra areas incurred by HexCNN and ZeroOut are similar, HexCNN still outperforms ZeroOut on the training time with both shallow and deep CNN models, because the averaged training time for each element is reduced, as illustrated in Section V-A. Meanwhile, HexCNN achieves superior space efficiency in convolution compared with ZeroOut, as illustrated in Section V-B.

Quasi-H takes rectangle-shaped input directly without padding, as shown in Fig. 1g. In convolution, Quasi incurs unnecessary padding for the filters, so as to apply existing deep learning frameworks. Although the padded area for the filters in convolution leads to deficiency, the overall efficiency of Quasi-H on rectangle-shaped input is still superior to HexCNN. However, Quasi-H simulates hexagonal processing on deep CNN models. It is not hexagonal processing in the strict sense (see Section II-B for more details).

VI. CONCLUSIONS AND FUTURE WORK

We proposed a native hexagonal CNN framework named HexCNN. HexCNN directly takes hexagon-shaped input and performs native hexagonal processing for both forward and backward propagation based on hexagon-shaped filters. It transforms hexagonal convolution to matrix multiplication to achieve high time efficiency. Experimental results on complete CNN models show that HexCNN outperforms hexagon-imitation models by reducing up to 42.2% of the training time. Meanwhile, HexCNN saves the space by up to 25% for loading the input, and up to 41.7% for performing convolution. HexCNN shows great performance from our experimental results on CPU. In the future, we will implement complete forward and backward propagation of HexCNN on GPU to obtain higher applicability of the framework.

REFERENCES

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, pages 265–283, 2016.
- [2] Martin Erdmann, Jonas Glombitza, and David Walz. A deep learning-based reconstruction of cosmic ray-induced air showers. *Astroparticle Physics*, 97:46–53, 2018.
- [3] Kazushige Goto and Robert A Geijn. Anatomy of high-performance matrix multiplication. *ACM TOMS*, 34(3):199–223, 2008.
- [4] Robert Hecht-Nielsen. Theory of the backpropagation neural network. In *Neural networks for perception*, pages 65–93, 1992.
- [5] Innchyn Her. Geometric transformations on the hexagonal grid. *IEEE TIP*, 4(9):1213–1222, 1995.
- [6] Emiel Hoogeboom, Jorn WT Peters, Taco S Cohen, and Max Welling. Hexaconv. *arXiv preprint arXiv:1803.02108*, 2018.

- [7] Roger A Horn. The hadamard product. In *Proc. Symp. Appl. Math.*, volume 40, pages 87–169, 1990.
- [8] Mirco Huennefeld. Deep learning in physics exemplified by the reconstruction of muon-neutrino events in icecube. *Verhandlungen der Deutschen Physikalischen Gesellschaft*, 2017.
- [9] Mikaël Jacquemont, Luca Antiga, Thomas Vuillaume, Giorgia Silvestri, Alexandre Benoit, Patrick Lambert, and Gilles Maurin. Indexed operations for non-rectangular lattices applied to convolutional neural networks. 2019.
- [10] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM MM*, pages 675–678, 2014.
- [11] Jintao Ke, Hai Yang, Hongyu Zheng, Xiquan Chen, Yitian Jia, Pinghua Gong, and Jieping Ye. Hexagon-based convolutional neural network for supply-demand forecasting of ride-sourcing services. *IEEE TITS*, 2018.
- [12] Dermot Kerr, Sonya A Coleman, Thomas Martin McGinnity, Qingxiang Wu, and Marine Clogenson. A novel approach to robot vision using a hexagonal grid and spiking neural networks. In *IJCNN*, pages 1–7, 2012.
- [13] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, 2009.
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436, 2015.
- [16] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [17] Junren Luo, Wanpeng Zhang, Jiongming Su, and Fengtao Xiang. Hexagonal convolutional neural networks for hexagonal grids. *IEEE Access*, 7:142738–142749, 2019.
- [18] Salvatore Mangano, Carlos Delgado, María Isabel Bernardos, Miguel Lallena, Juan José Rodríguez Vázquez, CTA Consortium, et al. Extracting gamma-ray information from images with convolutional neural network methods on simulated cherenkov telescope array data. In *IAPR Workshop on Artificial Neural Networks in Pattern Recognition*, pages 243–254, 2018.
- [19] Nikita Moriakov, Ashwin Samudre, Michela Negro, Fabian Gieseke, Sydney Otten, and Luc Hendriks. Inferring astrophysical x-ray polarization with deep learning. *arXiv preprint arXiv:2005.08126*, 2020.
- [20] D Nieto, A Brill, Q Feng, M Jacquemont, B Kim, T Miener, and T Vuillaume. Studying deep convolutional neural networks with hexagonal lattices for imaging atmospheric cherenkov telescope event reconstruction. *arXiv preprint arXiv:1912.09898*, 2019.
- [21] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [22] Tobias Schlosser, Michael Friedrich, and Danny Koworko. Hexagonal image processing in the context of machine learning: Conception of a biologically inspired hexagonal deep learning framework. *arXiv preprint arXiv:1911.11251*, 2019.
- [23] Idan Shilon, Manuel Kraus, Matthias Büchele, Kathrin Egberts, Tobias Fischer, Tim Lukas Holch, Thomas Lohse, Ullrich Schwanke, Constantin Steppa, and Stefan Funk. Application of deep learning methods to analysis of imaging atmospheric cherenkov telescopes data. *Astroparticle Physics*, 105:44–53, 2019.
- [24] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [25] Constantin Steppa and Tim L Holch. Hexagdy-processing hexagonally sampled data with cnns in pytorch. *SoftwareX*, 9:193–198, 2019.
- [26] Zhun Sun, Mete Ozay, and Takayuki Okatani. Design of kernels in convolutional neural networks for image classification. In *ECCV*, pages 51–66, 2016.
- [27] Kenny Young, Gautham Vasan, and Ryan Hayward. Neurohex: A deep q-learning hex agent. In *Computer Games*, pages 3–18, 2016.
- [28] Yunxiang Zhao, Jianzhong Qi, and Rui Zhang. Cbhe: Corner-based building height estimation for complex street scene images. In *WWW*, pages 2436–2447, 2019.