

OrthCaps: An Orthogonal CapsNet with Sparse Attention Routing and Pruning

Xinyu Geng¹ Jiaming Wang¹ Jiawei Gong¹ Yuerong Xue¹
Jun Xu^{1†} Fanglin Chen^{1†} Xiaolin Huang²

¹ Harbin Institute of Technology, Shenzhen ² Shanghai Jiao Tong University
{22s153095, 21s153144}@stu.hit.edu.cn, {jiawei.gong, xyuerong}@163.com,
xujunqgy@hit.edu.cn, chenfanglin@gmail.com, xiaolinhuang@sjtu.edu.cn

Abstract

Redundancy is a persistent challenge in Capsule Networks (CapsNet), leading to high computational costs and parameter counts. Although previous studies have introduced pruning after the initial capsule layer, dynamic routing’s fully connected nature and non-orthogonal weight matrices reintroduce redundancy in deeper layers. Besides, dynamic routing requires iterating to converge, further increasing computational demands. In this paper, we propose an Orthogonal Capsule Network (OrthCaps) to reduce redundancy, improve routing performance and decrease parameter counts. Firstly, an efficient pruned capsule layer is introduced to discard redundant capsules. Secondly, dynamic routing is replaced with orthogonal sparse attention routing, eliminating the need for iterations and fully connected structures. Lastly, weight matrices during routing are orthogonalized to sustain low capsule similarity, which is the first approach to use Householder orthogonal decomposition to enforce orthogonality in CapsNet. Our experiments on baseline datasets affirm the efficiency and robustness of OrthCaps in classification tasks, in which ablation studies validate the criticality of each component. OrthCaps-Shallow outperforms other Capsule Network benchmarks on four datasets, utilizing only 110k parameters – a mere 1.25% of a standard Capsule Network’s total. To the best of our knowledge, it achieves the smallest parameter count among existing Capsule Networks. Similarly, OrthCaps-Deep demonstrates competitive performance across four datasets, utilizing only 1.2% of the parameters required by its counterparts.

1. Introduction

While convolutional Neural Networks (CNNs) excel in computer vision tasks, certain challenges remain, which include information loss in pooling layers, low robustness, and poor spatial feature correlation [8, 30]. To address

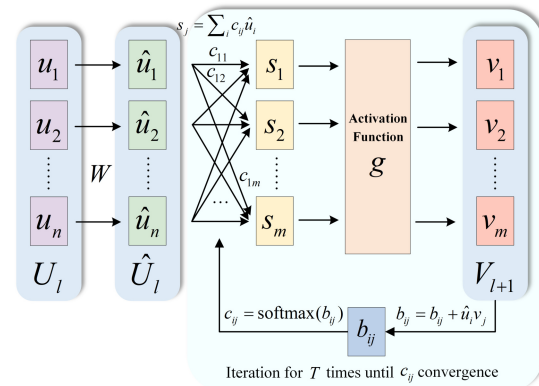


Figure 1. Dynamic routing mechanism. u_i, v_j are the lower-level capsule, and higher-level capsule, respectively. W is the weight matrix to produce the pose prediction \hat{u}_i of u_i for next level. b_{ij} is a temporary variable to calculate the coupling coefficient c_{ij} .

these limitations, Capsule Network (CapsNet) was proposed, using capsule vectors instead of traditional neurons. In CapsNet, each capsule vector’s length represents the presence probability of specific entities in the input image, and its direction encodes the captured features [30]. This setup allows the capsule vectors to capture features related to corresponding entities. CapsNet’s architecture includes a primary capsule extraction layer, a digit capsule layer, dynamic routing, and class-conditioned reconstruction. As a key component of CapsNet, dynamic routing aligns lower-level capsules with higher-level ones, which is described in Fig. 1. First, lower-level capsules U_l (in matrix form) predict poses \hat{U}_l for higher-level capsules V_{l+1} via weight matrix W . Then, the routing process iteratively clusters to adjust the coupling coefficients c_{ij} of each lower-level capsule $u_{l,i}$ to all higher-level capsules, with more crucial capsules receiving larger c_{ij} .

Recent studies have mentioned that CapsNet contains redundant capsules [1, 28, 31]. As evidence, Fig. 2 shows 48.2% of primary capsule pairs exhibit cosine similarities above 0.65, indicating significant redundancy. Although certain studies have implemented pruning techniques at

[†] Corresponding authors.




Figure 2. **Left:** In the primary capsule layer of CapsNet, 48.2% of capsule pairs have cosine similarities greater than 0.65, indicating significant redundancy among capsules. **Right:** After introducing the Pruned Layer, capsule similarities effectively decrease.

the primary capsule layer [29], deeper layers still show considerable over-similar issues, as shown in Fig. 6. We attribute this persistent redundancy in deeper layers to dynamic routing. On the one hand, since $v_{l+1,j} = g(\sum_i c_{ij} \hat{u}_{l,i})$, each higher-level capsule is essentially a weighted sum of lower-level capsules, indicating a fully connected relationship between lower and higher layers in CapsNet [14]. This full connection leads to a potential transmission of redundant information. On the other hand, considering $\hat{U}_l = WU_l$, we can express higher-level capsules as $V_{l+1} = g[(C*W)U_l]$. Non-orthogonal matrices $C*W$ in routing may increase the capsule similarity, which impairs routing performance and reintroduces redundancies in subsequent layers. Additionally, dynamic routing requires multiple iterations to repeatedly update c_{ij} until convergence, further straining computational resources.

Inspired by the successful use of orthogonality in CNNs [38] and Transformer [10] to reduce filter overlaps, we introduce the Orthogonal Capsule Network(OrthCaps), which has two versions: the lightweight OrthCaps-Shallow (**OrthCaps-S**) and the OrthCaps-Deep (**OrthCaps-D**). OrthCaps addresses the above problems of **the fully connected structure of dynamic routing, increasing similarity in deep layers, and the need for iteration**, detailed as follows:

Firstly, we introduce a pruned capsule layer after the primary capsule layer, which eliminates redundant capsules and retains only essential and representative ones. Here, capsules are firstly ordered by importance, then their cosine similarity is calculated to identify redundant capsules. Beginning with the least important, the process consistently prunes capsules that exceed the similarity threshold, proceeding through the entire set of capsules.

Secondly, to solve the iteration issue, dynamic routing is replaced with attention routing, which is a straightforward routing mechanism. To solve the fully connected problem, we leverage sparsemax-based attention [3] to produce an attention map, which selectively amplifies relevant feature groups corresponding to specific entities while downplaying irrelevant ones. For OrthCaps-S, a simplified attention-routing is adopted, optimizing parameter counts.

Thirdly, to address the issue of increased capsule similarity in deeper layers, we introduce orthogonality into capsule networks. By applying Householder orthogonal decomposition, we enforce orthogonality in the weight matrices during attention routing. Orthogonal weight matrices sustain low inter-capsule correlation, which encourages fewer capsules to represent more features during backpropagation, thereby enhancing accuracy while effectively reducing the number of parameters.

Contributions. To summarize our work, we make the following contributions:

- 1) To our knowledge, this approach addresses the issue of deep redundancy of capsules for the first time. A novel pruning strategy is implemented to alleviate capsule redundancy and an orthogonal sparse attention routing mechanism is proposed to replace dynamic routing.
- 2) It is the first time that Householder orthogonal decomposition is used to enforce orthogonality in CapsNet. This simple, penalty-free orthogonalization method is also adaptable to other neural networks.
- 3) Two OrthCaps versions are created: OrthCaps-S and OrthCaps-D. OrthCaps-S sets a new benchmark in accuracy with just 1.25% of CapsNet’s parameters on datasets of MNIST, SVHN, smallNORB, and CIFAR10. OrthCaps-D excels on CIFAR10, CIFAR100 and FashionMNIST while keeping parameters minimal.

2. Related Work

Capsule Networks. Dynamic routing was first introduced by Sabour et al.[30]. Though numerous studies have used attention strategies to improve dynamic routing, the issue of the fully connected structure and reintroduction of redundancy remains unaddressed [9, 21, 23]. Choi et al. incorporated attention into the capsule routing via a non-iterative feed-forward operation [2]. Tsai et al. introduced a parallel iterative routing, which did not address the complexity of iterative requirements [35]. Furthermore, many works focused on pruning but did not mention new redundancies introduced by dynamic routing. Jeong et al. established a ladder structure, using a pruning algorithm based on encoding [14]. Sharifi et al. created a pruning layer based on Taylor Decomposition [31]. Renzulli et al. used LOBSTER to create a sparse tree [29]. Different from existing works, we incorporate pruning, orthogonality and sparsity to effectively reduce redundancy.

Orthogonality. Various methods are proposed to introduce orthogonality into neural networks, which can be categorized into hard and soft orthogonality. Hard orthogonality maintains matrix orthogonality throughout training by either optimizing over the Stiefel manifold [11, 19], or parameterizing a subset of orthogonal matrices [33, 34, 37]. These methods incur computational overhead

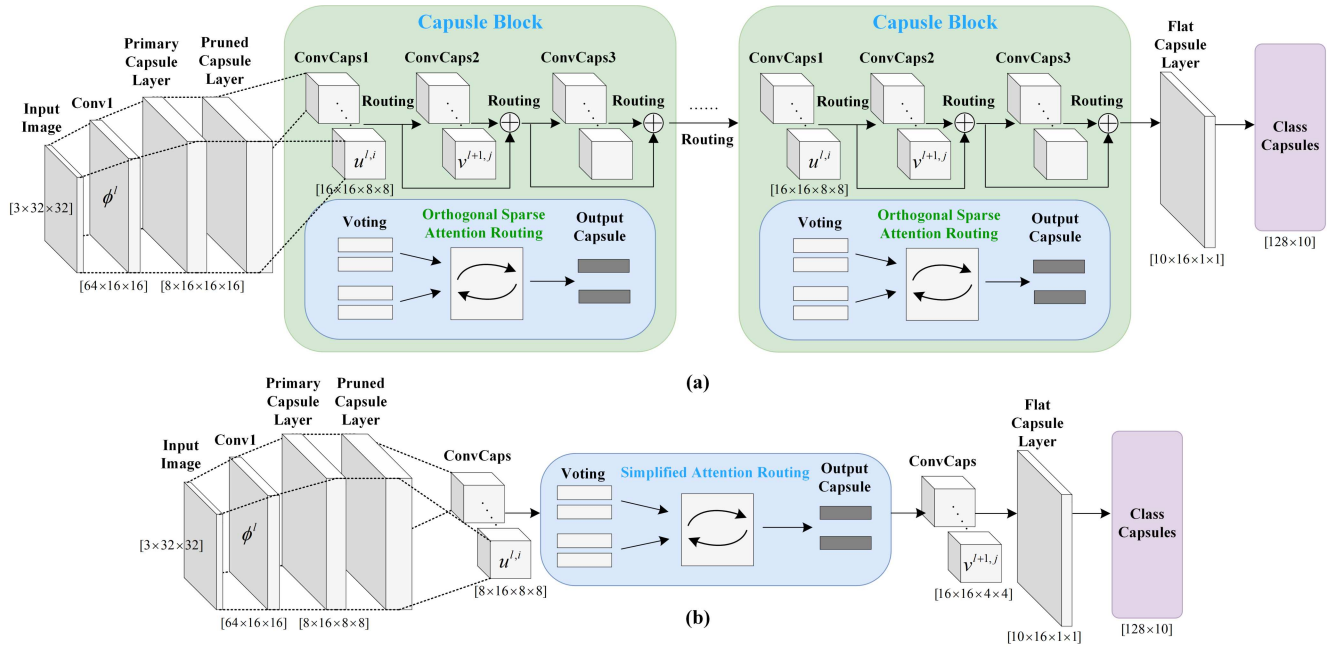


Figure 3. (a): In CIFAR10 classification task, the OrthCaps-D model comprises seven capsule blocks, each with three capsule layers, interconnected via shortcut connections and orthogonal sparse attention routing. (b): The OrthCaps-S model contains two capsule layers coping with CIFAR10 and does not use any capsule layer with MNIST. These layers are linked through simplified attention routing.

and result in vanishing or exploding gradients. Soft orthogonality, on the other hand, employs a regularization term in the loss function to encourage orthogonality among column vectors of weight matrix without strict enforcement [12, 16, 26]. Yet, strong regularization overshadows the primary task loss, while weak regularization fails to encourage orthogonality. We leverage Householder orthogonal decomposition [20, 36] to achieve strict matrix orthogonality, minimizing computational complexity and obviating the need for additional regularization terms.

3. Methodology

3.1. Overall Architecture

We introduce OrthCaps, offering both shallow (OrthCaps-S) and deep (OrthCaps-D) architectures to minimize parameter counts while exploring the potential for deep multi-layer capsule networks.

As illustrated in Fig. 3a, OrthCaps-D comprises five key components: a convolutional layer, a primary capsule layer, a pruned capsule layer, seven capsule blocks and a flat capsule layer. Given input images $x \in \mathbb{R}^{(B,3,W,H)}$, initial features $\Phi^0 \in \mathbb{R}^{(B,C_n,W^0,H^0)}$ are extracted through a convolutional layer. The primary capsule layer generates initial capsules $u^1 \in \mathbb{R}^{(B,n,d,W^1,H^1)}$ with a kernel size of 3 and stride of 2. B, n, d, C_n represent the batch size, number of capsules, capsule dimensions and channels, respectively. A pruned capsule layer is then placed to remove redundant capsules. OrthCaps-D has seven capsule blocks, each containing three depthwise convolutional capsule

layers(ConvCaps Layers) linked by shortcut connections to prevent vanishing gradient. Within each block, lower-level capsules u^l are routed to the next layer v^{l+1} via orthogonal sparse attention routing. Blocks are also connected through routing, allowing for stacking to construct deeper capsule networks. The flatcaps layer is employed to map capsules into classification categories for final classification tasks.

OrthCaps-S, as illustrated in Fig. 3b, replaces the complete attention routing with a simplified version and has a single block within capsule layers. The number of layers can be adjusted as needed. Convolutional capsules in the primary layer utilize a 9×9 kernel with a stride of 1, and other layers are consistent with OrthCaps-D. The subsequent introduction of the network's components follows the order of data flow.

3.2. Pruned Capsule Layer

The generation of capsules starts with the primary capsule layer. At this initial stage, it is crucial to generate low-correlated capsules, which ensures efficient feature representation and reduces feature overlap during subsequent layers. Therefore, we introduce an efficient capsule pruning Algorithm 1, including the following parts:

Redundancy Definition. Redundancy occurs when two capsules capture identical or similar features. Given that the direction of each capsule vector encodes specific features, capsules with closer directions(or angles) indicate that they capture similar features and entities. Thus, we employ the cosine similarity of capsule angles to measure redundancy.

Capsule Importance Ordering. For redundant capsule

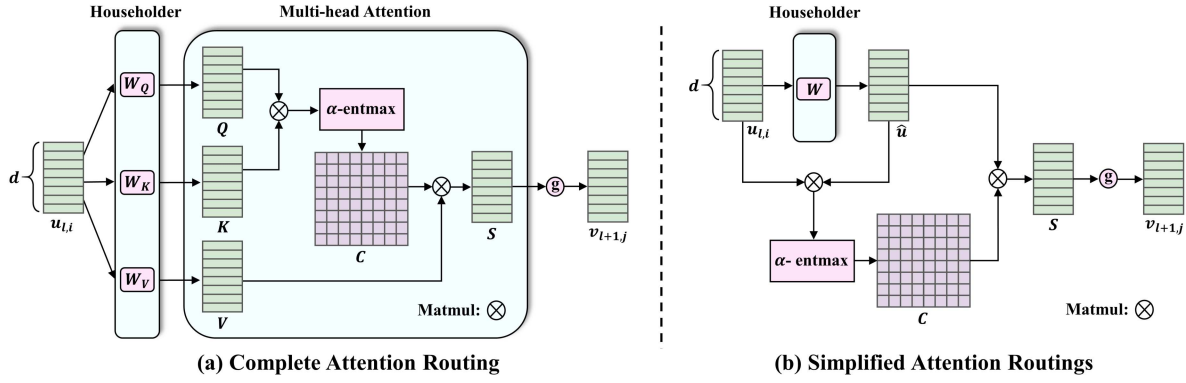


Figure 4. Orthogonal self-attention routing.

pairs, random pruning may result in losing capsules vital for accurate classification. To ensure that the less crucial capsule is pruned first when the similarity between a pair is high, capsules are sorted in an order based on $\|u_{flat}\|_2$. We employ L_2 -norm as it calculates the length of capsule vectors, indicating the existence probability of encoded entities, which shows the activeness of capsules [14].

Pruning. After ordering, a mask matrix $M \in \mathbb{R}^{(1,n,1)}$ is initialized to all-ones. Starting with the least active capsule, the process computes the cosine similarity between less active capsule $u_{ordered,i}$ with more active capsule $u_{ordered,j}$. When the similarity exceeds the threshold θ , the corresponding column in the mask for $u_{ordered,i}$ is set to 0, indicating that less active capsule is pruned. In this way, only the active capsules are retained all along. The final step is applying M to $u_{ordered}$, producing pruned capsules u_{pruned} . n' is the number of remaining capsules after pruning.

Algorithm 1 Efficient Capsule Pruning

Require: Capsules $u \in \mathbb{R}^{(B,n,d,W,H)}$, threshold θ

Ensure: $u_{pruned} \in \mathbb{R}^{(B,n,d,W,H)}$

- 1: Reshape $u \rightarrow u_{flat} \in \mathbb{R}^{(B,n,(d \times W \times H))}$
- 2: Compute L_2 -norm: $\|u_{flat}\|_2$
- 3: Order capsules by L_2 -norm: $u_{flat} \rightarrow u_{ordered}$
- 4: Initialize M : all-ones matrix
- 5: **for** $i < j$ **do**
- 6: $t_{ij} = \text{cosine_similarity}(u_{ordered,i}, u_{ordered,j})$
- 7: $m_i = 0$ where $t_{ij} > \theta$
- 8: Prune using M : $u_{pruned} = u_{ordered} \odot M$
- 9: Reshape u_{pruned} to $u_{pruned} \in \mathbb{R}^{(B,n',d,W,H)}$
- 10: **return** u_{pruned}

Notably, we compute the cosine similarity matrix T by broadcasting, which avoids explicit for-loop iteration and reduces computational complexity. The 5D capsule tensor of u_l with dimensions d and number n , $[B, n, d, W, H]$, is reshaped to $[B, n, d \times W \times H]$ to suit for broadcasting.

3.3. Routing Algorithm

We introduce the orthogonal sparse attention routing to replace dynamic routing. This approach eliminates the

need for iteration and leverages sparsity to reduce redundant feature transmission.

Let $u_{l,i}$ and $v_{l+1,j}$ represent capsules at layer l and $l + 1$ respectively, each with dimension d . We employ three weight matrices $W_Q, W_K, W_V \in \mathbb{R}^{d \times d}$ to derive queries Q , keys K , and values V from $u_{l,i}$, $Q = W_Q \times u_{l,i}$, $K = W_K \times u_{l,i}$, $V = W_V \times u_{l,i}$. Specifically, W_Q, W_K , and W_V are designed as orthogonal matrices, enabling them to project $u_{l,i}$ into a d -dimensional orthogonal subspace.

As shown in Fig. 4a, attention routing aims to produce coupling coefficient c_{ij} , which serves as the weight during routing from lower-level to higher-level capsules. The coupling coefficient matrix C is derived from the attention map, generated through the dot product of Q and K , $C = \alpha\text{-Entmax}(QK^T/\sqrt{d})$. Here, we replace the softmax of the original attention mechanism with $\alpha\text{-Entmax}$ in Eq. (1), enhancing the sparsity of the attention map. $\alpha\text{-Entmax}$ adaptively sets small c_{ij} to zero, thereby encouraging routing to prioritize more important capsules while minimizing irrelevant information transfer.

$$\alpha\text{-Entmax}(x)_i = \max\left(\frac{x_i - \tau}{\alpha}, 0\right)^{\frac{1}{\alpha-1}} \quad (1)$$

τ is a self-adaption threshold and α is a hyperparameter controlling the sparsity of the attention map.

The vote $s_{l+1,j}$ is computed as the product of V and C . In Eq. (2), higher-level capsule $v_{l+1,j}$ is generated by $s_{l+1,j}$ from a multi-head self-attention mechanism with 16 heads, using the non-linear activation function g .

$$v_{l+1,j} = g(s_{l+1,j}) = g(\alpha\text{-Entmax}(QK^T/\sqrt{d}) \times V) \quad (2)$$

For simplified attention-routing in Fig. 4b, we condense prediction matrices W from three to one and replace K, Q, V with $u_{l,i}$ [7]. The $\hat{u}_{l,i}$ is the prediction for $v_{l+1,j}$. The attention map C is obtained using $\alpha\text{-entmax}$ with the dot product to produce the vote $s_{l+1,j} = \hat{u}_{l,i} \times C = \hat{u}_{l,i} \times (\alpha\text{-Entmax}(\hat{u}_{l,i}u_{l,i}^T/\sqrt{d}))$. $s_{l+1,j}$ is processed through g to produce $v_{l+1,j}$. Notably, standard convolutions are supplanted by depthwise convolutions to minimize parameter count. Without any iteration, attention routing reduces computational complexity.

3.4. Orthogonalization

In Sec. 3.2, we reduce redundancy by pruning highly similar capsules. To preserve pruning effect in subsequent layers, it's vital to maintain low similarity among capsules all along. Here, capsules with diverse angles span a broader multi-dimensional feature subspace, enabling the network to capture a wider array of features with fewer capsules, which boosts accuracy and reduces parameter count.

We utilize orthogonalization to achieve this. An orthogonal matrix, signifying a rotation or reflection transformation, keeps vector lengths and inter-vector angles constant during multiplication with a vector set. Considering that cosine similarity quantifies angles between vectors, applying an orthogonal matrix W to a vector set retains the mutual cosine similarity among all vector pairs in the set. Let $\{v_{l+1,j} \mid j = 1, 2, \dots, m\}$ be a set of capsule vectors at layer $l + 1$, we derive Lemma 1:

Lemma 1: For any $i, j \in \{1, 2, \dots, m\}$, if W is orthogonal, the cosine similarity between $v_{l+1,i}$ and $v_{l+1,j}$ remains unchanged after multiplication by W .

Sec. 3.4.1 discusses the selection of targets for orthogonalization, while Sec. 3.4.2 details the method.

3.4.1 Orthogonalization of Weight Matrices

The goal of orthogonalization is to maintain low similarity among higher-level capsules. Following Sec. 3.3, we represent higher-level capsules into matrix multiplication:

$$V_{l+1} = g(S_l) = g(CV) = g[(C \times W_V)U_l] \quad (3)$$

As the activation function g in CapsNet preserves the capsule vector's direction [30], in line with Lemma 1, ensuring orthogonality of C and W_V can maintain low similarity. However, orthogonalizing C directly is hard, so we delve deeper into its calculation process in Eq. (4):

$$C = \alpha\text{-Entmax}(QK^T/\sqrt{d}) = \alpha\text{-Entmax}(W_Q U_l U_l^T W_K^T/\sqrt{d}) \quad (4)$$

Since d is a constant, it requires no additional analysis. The impact of $U_l U_l^T$ on orthogonalizing C is mitigated by pruning, which removes capsules with short lengths and high cosine similarity. Consequently, the remaining capsules were updated to approximate unit length and low correlation, akin to a standard orthonormal basis. Thus, $U_l U_l^T$ gradually becomes more orthogonal as the network trains, minimizing its impact on the orthogonality of C . Although $\alpha\text{-Entmax}$, a nonlinear function, may not preserve the orthogonality of inputs, it renders C sparse. This sparsity directs lower-level capsules effectively toward their respective higher-level targets, which reduces interference during routing, thus encouraging a relatively low similarity among higher-level capsules.

Our above analysis indicates that orthogonalizing W_Q, W_K , and W_V is essential for maintaining low cosine similarity among higher-level capsules. While C is not fully orthogonalized, our experiments demonstrate considerable improvements in both accuracy and parameter efficiency.

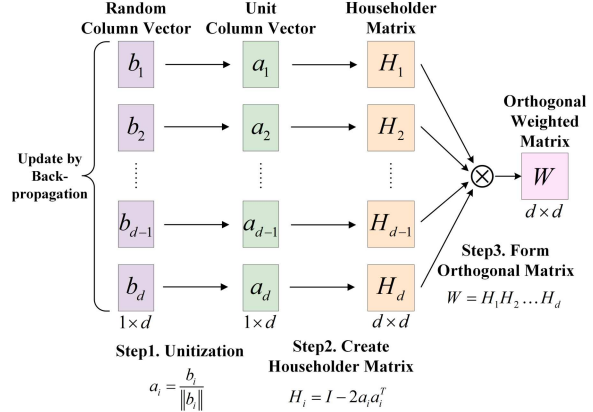


Figure 5. The computing process of HouseHolder method.

3.4.2 Householder Orthogonalization

Let W be the weight matrix requiring orthogonalization. As shown in Fig. 5, the Householder orthogonal decomposition theorem is employed to formulate an endogenously optimizable orthogonal matrix. The essence of this approach is in the following algebraic lemma [36]:

Lemma 2: Any orthogonal $n \times n$ matrix is the product of at most n orthogonal Householder transformations.

Based on Lemma 2, an orthogonal matrix $W \in \mathbb{R}^{d \times d}$ can be formulated in Eq. (5):

$$W = H_0 H_1 \dots H_{d-1} \quad (5)$$

Each H_i represents a Householder transformation, defined as $H_i = I - 2a_i a_i^T$, where a_i is a unit column vector. We utilize a set of randomly generated column vectors $\{b_i \mid i = 0, \dots, d - 1\}$ instead of a_i to construct H_i as detailed in Eq. (6). During training, b_i is optimized through gradient backpropagation. W inherently preserves its orthogonality during training.

$$W = \prod_{i=0}^{d-1} \left(I - \frac{2b_i b_i^T}{\|b_i\|^2} \right) \quad (6)$$

Lemma 3: W_Q, W_K , and W_V constructed using Equation Eq. (6) are Orthogonal.

Following Eq. (6), W_Q, W_K , and W_V could easily be orthogonalized, where the proof is provided in supplement material ???. Householder orthogonalization enables computationally efficient transformation of arbitrary coefficient matrices into orthogonal matrices without any additional penalty terms in the loss function.

4. Experiments

4.1. Experimental Setup

Implementation Details and Datasets

OrthCaps is developed using PyTorch 1.12.1, running on Python 3.9, and the training process is accelerated using

Table 1. **(a):** OrthCaps-S ranks as the top or second best across five datasets, standing out as being resource-efficient with only 105.5K parameters and 673.1M FLOPS. **(b):** OrthCaps-D shows competitive performance with fewer parameters and less computational cost.

Shallow Networks	Param↓	FLOPS[M]↓	MNIST	SVHN	smallNORB	CIFAR10
OrthCaps-S ¹	105.5K	673.1	99.68	96.26	98.30	86.84
Efficient-Caps	162.4K	631.1	99.58	93.12	97.46	81.51
CapsNet	8388K	803.8	99.52	91.36	95.42	68.72
Matrix-CapsNet with EM routing	450K	949.6	99.56	87.42	95.56	81.39
AR CapsNet	9.1M	2562.7	99.46	85.98	96.47	85.39
DA-CapsNet	7M*	-	99.53*	94.82*	98.26*	85.47*
AA-CapsNet	6.6M*	-	99.34*	91.23*	89.72*	79.41*
CapProNet	-	-	94.98	93.41	92.01	80.84
Baseline CNN	4.6M	1326.9	99.22	91.28	87.11	72.20

(a)

Deep Networks	Param ↓	FLOPS[M]↓	CIFAR10	CIFAR100	MNIST	FashionMNIST
OrthCaps-D ¹	574K	3345	90.56	70.56	99.59	94.60
AR-CapsNet(7 ensembled)	6.3M	16657.5	88.94	56.53	99.49	91.73
CapsNet(7 ensembled)	5.8M*	5137.4*	89.4*	-	-	-
Inverted Dot-Product	1.4M	5340.9	84.98	57.32	99.35	92.85
RS-CapsNet	5.0M*	-	89.81*	64.14*	-	93.51*
DeepCaps	13.5M	2687	91.01	69.72	99.46	92.52
ResNet-18	11.7M	5578.8	95.10	77.60	99.29	93.32
VGG-16	147.3M	15143.1	93.57	73.10	99.21	92.21

(b)

four GTX-3090 GPUs. We adopt the margin loss as defined in [30]. Observing minimal performance benefits in our experiments, we decide to exclude the reconstruction loss. Our model utilizes the AdamW optimizer with a cosine annealing learning rate scheduler and a 5-cycle linear warm-up. We set learning rate at $5e-3$, weight decay at $5e-4$, and batchsize at 512. Experiments are conducted on SVHN [22], smallNORB [18], CIFAR10, and MNIST [17] for OrthCaps-S. OrthCaps-D is evaluated on CIFAR10, CIFAR100 [15], Fashion-MNIST [39], and MNIST. We resize SmallNORB from 96×96 to 64×64 and cropped it to 48×48 like [30]. All other datasets retained original sizes. For data augmentation, we adopt the methods outlined in [8]. For reproducibility, we detailed hyperparameters and setups in supplement material ??.

Comparison Baselines

For OrthCaps-S, we compare it with Efficient-Caps[21], CapsNet[30], Matrix-CapsNet with EM routing[8], AR-CapsNet[2], AA-CapsNet[25], DA-CapsNet[13], CapProNet[41] and standard 7-layer CNN. We compare OrthCaps-D with CapsNet (7 ensembles), AR-CapsNet (7 ensembles), RS-CapsNet[40], Inverted Dot-Product[35], DeepCaps[27], ResNet-18[6], and VGG-16[32]. All CapsNet variants use a backbone of 4 convolutional layers and undergo training for 500 epochs, and results are derived from running official codes with our hyperparameters.

4.2. Classification Performance Comparison

Tab. 1 illustrates the classification performance of OrthCaps-S and OrthCaps-D, with model sizes denoted

¹<https://github.com/ornamentt/OrthCap>.

by Param and computational demands represented as FLOPS[M]. The Param and FLOPS[M] of each table are tested on MNIST and CIFAR10, respectively. An asterisk (*) signifies that no official code is available, so we refer to the model performance stated in the original papers.

As shown in Tab. 1a, OrthCaps-S achieves superior efficiency with merely 105.5K parameters, outperforming CNN, CapsNet, and many variants. For instance, Efficient-Caps, a state-of-the-art model on efficiency, has about 50% more parameters. Furthermore, OrthCaps-S either outperforms or matches the performance of other models across all four datasets. On the SVHN and CIFAR10, OrthCaps-S achieves accuracies of 96.26% and 86.84%, respectively, surpassing CapsNet which has 80 times more parameters. With a computational demand of 673.1M FLOPS, it's worth noting that the slight increase in FLOPS compared with Efficient-Caps is due to the additional computations from the pruned capsule layer and orthogonal transformations. Given the substantial decrease in parameter count and the enhanced accuracy, this FLOPS trade-off is warranted.

For OrthCaps-D, as illustrated in Tab. 1b, it exhibits competitive performance with fewer parameters on complex datasets. Although convolution-based networks such as ResNet-18 and VGG-16 perform well on CIFAR10 and CIFAR100, OrthCaps-D offers competitive performance using just 1.41% and 0.11% of their parameters as well as 56% and 20.8% of their FLOPS, respectively. The efficiency of OrthCaps becomes evident when compared with DeepCaps. Although DeepCaps achieves a 91.01% accuracy on CIFAR10, its parameter count of 13.42M is 23.4 times higher than OrthCaps-D. Both OrthCaps variants maintain high performance with fewer parameters.

Table 2. Comparison of Orthogonal sparse attention routing and dynamic routing algorithms. FPS is tested under MNIST dataset.

Variants	FPS \uparrow	MNIST	CIFAR10
Attention routing & α -entmax & orthogonality	1639	99.68	86.84
Attention routing & softmax	1785	99.62	83.44
Dynamic routing & α -entmax & orthogonality	1232	99.51	70.01
Dynamic routing & softmax	1339	99.49	68.72

Table 3. CapsNets are compared with and without the pruning layer, with the similarity threshold set to 0.7. Param[K] is tested on MNIST.

Variant	Param[K] \downarrow	MNIST	CIFAR10
OrthCaps-S with pruning	105	99.68	86.84
OrthCaps-S	157	99.53	85.32
Capsnet with pruning	7492	99.51	71.08
Capsnet	8388	99.42	68.72

4.3. Ablation Study

4.3.1 Orthogonal Attention Routing

Through a cross-comparison of frames-per-second(FPS) and accuracy on two datasets with different complexity, as shown in Tab. 2, we compare attention routing with dynamic routing[30] and sparse softmax with standard softmax, respectively. Additionally, α is settled to 1.5 in our experiments according to [24].

Attention routing consistently outperforms dynamic routing in both classification accuracy and processing speed, achieving a 25.8% speed enhancement on average. Even with a faster softmax, dynamic routing only reaches 1339 FPS, indicating its inherent computational inefficiencies. Although α -entmax’s complexity and the additional computational demands from orthogonality slightly reduce processing speed, this trade-off is justified by a substantial increase in accuracy and robustness. Our attention routing combined with α -entmax and orthogonality balances performance and computational efficiency.

4.3.2 Pruned Capsule Layer

Fig. 2 illustrates that by integrating the pruned layer, the average capsule similarity decreases due to redundant capsule elimination. Consequently, as the capsule count reduces, the dimensions of the associated prediction matrix diminish, thereby lowering the parameter count. This is proved in Tab. 3, where the pruned OrthCaps-S reduces parameters from 127K to 105K without sacrificing performance. In fact, accuracy improves from 99.53% to 99.68% and from 85.32% to 86.84% on MNIST and CIFAR10, respectively. Similarly, applying pruning to CapsNet results in higher accuracy with reduced parameters (7492K from 8388K). This shows our pruning method’s efficacy in streamlining the model and enhancing performance.

Fig. 6 illustrates the necessity of incorporating pruning with orthogonality. Capsule similarity is gauged with

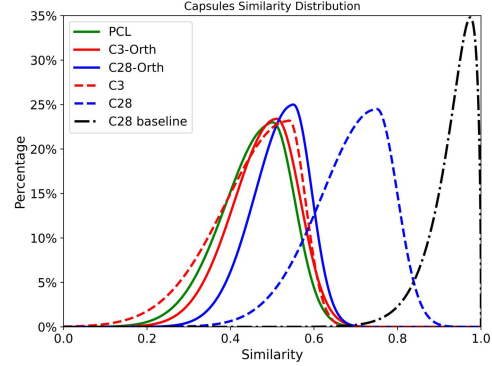


Figure 6. Redundancy comparison between different pruning strategies. **(The more to the left, the better.)** The x-axis shows capsule similarity; the y-axis indicates capsule count percentage. PCL, C3, and C28 mark the primary, third, and twenty-eighth capsule layers. Solid lines (C3-Orth, C28-Orth) and dashed lines represent pruning with and without orthogonality, respectively; the dash-dot line indicates the absence of both pruning and orthogonality. Tests are on OrthCaps-D with CIFAR10 dataset.

cosine similarity to measure the redundancy as mentioned above. As the network goes deeper, the dashed line (indicating pruning without orthogonality) shifts rightward, suggesting an increase in capsule similarity. This shift proves that non-orthogonal weight matrices reintroduce redundancy. However, the solid line (indicating pruning with orthogonality) shows consistently low capsule similarity. Even at layer 28, considered quite deep for CapsNets, the similarity remains low. This affirms the efficacy of orthogonality in preserving capsule directions to maintain low inter-capsule correlations. The black dash-dot line denotes similarity without orthogonality and pruning, exhibiting the highest redundancy, further evaluating the effectiveness of our method.

4.4. Similarity Threshold

To find the optimal similarity threshold θ , we evaluated the classification accuracy on three datasets and capsule number after pruning (N) for thresholds ranging from 0.3 to 1.0. At $\theta = 1.0$, the pruning layer becomes ineffective as it targets capsules with similarities above one.

As Tab. 4 shows, the optimal accuracy occurs at $\theta = 0.7$. A threshold below 0.7 leads to excessive pruning, over-reducing capsule numbers and causing feature loss, thus impacting the accuracy. Notably, at $\theta = 0.4$, all three datasets show a marked accuracy drop, indicating a critical point where key information is lost. CIFAR10 is the most affected, likely due to its complex background and rich features, making it more sensitive to excessive pruning. Conversely, a higher θ weakens pruning effectiveness. As redundant information accumulates, the classification capsules become disrupted, slightly diminishing performance. With overall consideration, we set θ into 0.7.

Table 4. Comparison of different similarity thresholds. Num is the number of capsules after pruning.

θ	Num	MNIST	SVHN	CIFAR10
0.3	58	38.25	18.96	10.25
0.4	116	43.60	24.87	11.03
0.5	303	90.09	86.12	63.35
0.6	676	97.46	94.66	81.50
0.7	952	99.68	96.26	86.79
0.8	1093	99.63	96.13	86.84
0.9	1139	99.61	95.96	86.41
1.0(without pruning)	1152	99.53	95.25	85.32

Table 5. Comparison of OrthCaps, CapsNet, OCNN and baseline CNN under PGD attack. The CIFAR10 dataset is used without any data augmentation. Our results are an average of 5 test runs.

Variants	AT(s) \uparrow	QC[K] \uparrow	ACC \uparrow
OrthCaps	345.92	69K	23.52
CapsNet	198.93	48K	14.62
OCNN	136.7	46K	-
baseline CNN	16.65	10K	0.35

4.5. Robustness to Adversarial Attacks

Capsule networks have demonstrated exceptional performance in terms of robustness [8]. Considering OrthCaps as it eliminates redundant capsules to suppress low L_2 -norm capsules, which we consider as noise capsules [4]. It can enhance better robustness against small perturbations. To evaluate this, we conduct a robustness comparison between OrthCaps, Capsule Networks, Orthogonal CNNs (OCNN) and 7-layer CNNs with CIFAR10. We employ the Projected Gradient Descent (PGD) white-box attack method [5], setting the maximum iteration count at 40, step size at 0.01, and the maximum perturbation at 0.1. We assess the robustness using three metrics: attack time (AT), model query count (QC), and accuracy after attacks (ACC). As shown in Tab. 5, OrthCaps outperforms in all three metrics.

4.6. Orthogonality

This experiment demonstrates the effectiveness of the HouseHolder orthogonalization method and its advantages over other methods. We define an orthogonality metric $O = \|K^T K - I\|$. In Tab. 6a, the metric O decreases from 0.02 to 0.01 during training, showing the effectiveness of the orthogonalization method. In Tab. 6b, we compare the orthogonalization of weight matrices in attention routing, and the results show that orthogonalizing all W_Q , W_K and W_V matrices achieves the best performance.

We further demonstrate Householder’s role as a regularization technique for neural networks. In Fig. 7, our method achieves better orthogonality and loss decay than OCNN [38]. The baseline ResNet18, without any orthogonal regularization, is depicted by the blue line, while the green and red lines stand for OCNN and our method, respectively. The near-flat trajectory of the red line shows Householder’s consistent orthogonality preservation across

Table 6. (a): The proof of orthogonality in weight matrices. (b): Orthogonalizing all weight matrices performs best.

EPOCH	SVHN	$O \downarrow$	Orthogonalization	MNIST	CIFAR10
1	83.75	0.0236	None	99.62	84.25
10	98.58	0.0215	$W_Q \& W_K$	99.32	81.40
100	99.42	0.0153	W_V	99.47	83.33
300	99.56	0.0120	$W_Q \& W_K \& W_V$	99.68	86.84

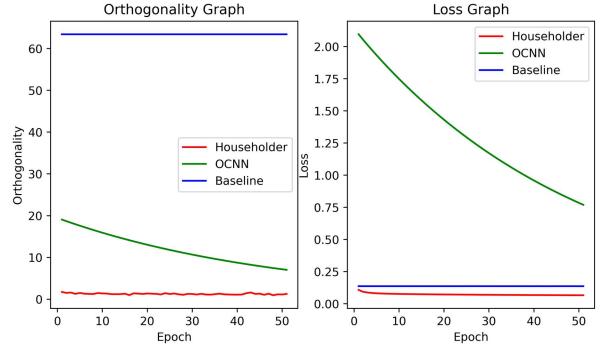


Figure 7. Different Orthogonal regularization on MNIST dataset. HouseHolder orthogonalization method reaches the best orthogonality and loss decay.

the training. Furthermore, our method registers a smaller loss than OCNN, due to its better training performance.

5. Conclusions and Future Work

This study has introduced a novel capsule network with orthogonal sparse attention routing and pruning. Specifically, Householder orthogonal decomposition is used to ensure orthogonality in attention routing without additional penalty terms. With pruning, orthogonalization and sparse routing, OrthCaps minimizes capsule redundancies. Experiments show that OrthCaps has lower parameters and reduces computational overhead, overcoming the challenges of computational expense and redundancy in dynamic routing. On image classification tasks, OrthCaps outperforms state-of-the-art methods and demonstrate improved robustness. We look forward to future research in this area.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62173113 and 62376155, the Science and Technology Innovation Committee of Shenzhen Municipality under Grant GXWD20231129101652001, Natural Science Foundation of Guangdong Province of China under Grant 2022A1515011584, the Guangdong Basic and Applied Basic Research Foundation under Grant 2024A1515011706, the Shenzhen Key Technical Project under Grant KJZD20230923115117033, Shanghai Municipal Science and Technology Research Program under Grant 22511105600 and Major Project under Grant 2021SHZDZX0102.

References

- [1] Hao Chen, Xian-bo Wang, and Zhi-Xin Yang. Fast robust capsule network with dynamic pruning and multiscale mutual information maximization for compound-fault diagnosis. *IEEE/ASME Transactions on Mechatronics*, 2022.
- [2] Jaewoong Choi, Hyun Seo, Suii Im, and Myungjoo Kang. Attention routing between capsules. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 2019.
- [3] Gonalo M Correia, Vlad Niculae, and Andr  FT Martins. Adaptively sparse transformers. *arXiv preprint arXiv:1909.00015*, 2019.
- [4] Fabio De Sousa Ribeiro, Georgios Leontidis, and Stefanos Kollias. Introducing routing uncertainty in capsule networks. In *Advances in Neural Information Processing Systems*, pages 6490–6502, 2020.
- [5] Chuan Guo, Jacob Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Weinberger. Simple black-box adversarial attacks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2484–2493. PMLR, 2019.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [7] Geoffrey Hinton. How to represent part-whole hierarchies in a neural network. *Neural Computation*, 35(3):413–452, 2023.
- [8] Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with EM routing. In *International Conference on Learning Representations*, 2018.
- [9] Assaf Hoogi, Brian Wilcox, Yachee Gupta, and Daniel L Rubin. Self-attention capsule networks for object classification. *arXiv preprint arXiv:1904.12483*, 2019.
- [10] Huaibo Huang, Xiaoqiang Zhou, and Ran He. Orthogonal transformer: An efficient vision transformer backbone with token orthogonalization. *Advances in Neural Information Processing Systems*, 35:14596–14607, 2022.
- [11] Lei Huang, Xianglong Liu, Bo Lang, Adams Yu, Yongliang Wang, and Bo Li. Orthogonal weight normalization: Solution to optimization over multiple dependent stiefel manifolds in deep neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), 2018.
- [12] Lei Huang, Li Liu, Fan Zhu, Diwen Wan, Zehuan Yuan, Bo Li, and Ling Shao. Controllable orthogonalization in training DNNs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6429–6438, 2020.
- [13] Wenkai Huang and Fobao Zhou. Da-capsnet: Dual attention mechanism capsule network. *Scientific Reports*, 10(1): 11383, 2020.
- [14] Taewon Jeong, Youngmin Lee, and Heeyoung Kim. Ladder capsule network. In *Proceedings of the 36th International Conference on Machine Learning*, pages 3071–3079. PMLR, 2019.
- [15] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [16] Shubhranil Kundu and B Gagana. Orthogonalizing weights in capsule network architecture. In *ICT Analysis and Applications: Proceedings of ICT4SD 2019, Volume 2*, pages 77–84. Springer, 2020.
- [17] Yann LeCun. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [18] Yann LeCun, Fu Jie Huang, and L on Bottou. Learning methods for generic object recognition with invariance to pose and lighting. Technical report, Courant Institute of Mathematical Sciences, New York University, 2004.
- [19] Jun Li, Li Fuxin, and Sinisa Todorovic. Efficient riemannian optimization on the stiefel manifold via the cayley transform. *arXiv preprint arXiv:2002.01113*, 2020.
- [20] Alexander Mathiasen, Frederik Hvilsh  j, Jakob R  dsgaard J  rgensen, Anshul Nasery, and Davide Mottin. What if neural networks had SVDs? In *Advances in Neural Information Processing Systems*, pages 18411–18420, 2020.
- [21] Vittorio Mazzia, Francesco Salvetti, and Marcello Chiaberge. Efficient-capsnet: Capsule network with self-attention routing. *Scientific Reports*, 11(1):14634, 2021.
- [22] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011.
- [23] Dunlu Peng, Dongdong Zhang, Cong Liu, and Jing Lu. Bg-sac: Entity relationship classification model based on self-attention supported capsule networks. *Applied Soft Computing*, 91:106186, 2020.
- [24] Ben Peters, Vlad Niculae, and Andr  FT Martins. Sparse sequence-to-sequence models. *arXiv preprint arXiv:1905.05702*, 2019.
- [25] Rita Pucci, Christian Micheloni, and Niki Martinel. Self-attention agreement among capsules. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 272–280, 2021.
- [26] Haozhi Qi, Chong You, Xiaolong Wang, Yi Ma, and Jitendra Malik. Deep isometric learning for visual recognition. In *Proceedings of the 37th International Conference on Machine Learning*, pages 7824–7835. PMLR, 2020.
- [27] Jathushan Rajasegaran, Vinoj Jayasundara, Sandaru Jayasekara, Hirunima Jayasekara, Suranga Seneviratne, and Ranga Rodrigo. Deepcaps: Going deeper with capsule networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10725–10733, 2019.
- [28] Riccardo Renzulli and Marco Grangetto. Towards efficient capsule networks. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 2801–2805. IEEE, 2022.
- [29] Riccardo Renzulli, Enzo Tartaglione, and Marco Grangetto. Rem: Routing entropy minimization for capsule networks. *arXiv preprint arXiv:2204.01298*, 2022.
- [30] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, page 3856–3866, 2017.
- [31] Ramin Sharifi, Pouya Shiri, and Amirali Baniasadi. Pruned-caps: A case for primary capsules discrimination. In *2021*

- 20th IEEE International Conference on Machine Learning and Applications (ICMLA), pages 1437–1442. IEEE, 2021.
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
 - [33] Sahil Singla and Soheil Feizi. Skew orthogonal convolutions. In *Proceedings of the 38th International Conference on Machine Learning*, pages 9756–9766. PMLR, 2021.
 - [34] Asher Trockman and J Zico Kolter. Orthogonalizing convolutional layers with the cayley transform. *arXiv preprint arXiv:2104.07167*, 2021.
 - [35] Yao-Hung Hubert Tsai, Nitish Srivastava, Hanlin Goh, and Ruslan Salakhutdinov. Capsules with inverted dot-product attention routing. *arXiv preprint arXiv:2002.04764*, 2020.
 - [36] Frank Uhlig. Constructive ways for generating (generalized) real orthogonal matrices as products of (generalized) symmetries. *Linear Algebra and its Applications*, 332:459–467, 2001.
 - [37] Aladin Virmaux and Kevin Scaman. Lipschitz regularity of deep neural networks: Analysis and efficient estimation. In *Advances in Neural Information Processing Systems*, pages 3835–3844, 2018.
 - [38] Jiayun Wang, Yubei Chen, Rudrasis Chakraborty, and Stella X. Yu. Orthogonal convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11505–11515, 2020.
 - [39] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
 - [40] Shuai Yang, Feifei Lee, Ran Miao, Jiawei Cai, Lu Chen, Wei Yao, Koji Kotani, and Qiu Chen. Rs-capsnet: An advanced capsule network. *IEEE Access*, 8:85007–85018, 2020.
 - [41] Liheng Zhang, Marzieh Edraki, and Guo-Jun Qi. Cappronet: Deep feature learning via orthogonal projections onto capsule subspaces. *Advances in Neural Information Processing Systems*, 31, 2018.