



No routing needed between capsules

Adam Byerly^{a,b,*}, Tatiana Kalganova^a, Ian Dear^a

^a Department of Electronic and Electrical Engineering, Brunel University London, Uxbridge UB8 3PH, UK

^b Department of Computer Science and Information Systems Bradley University, Peoria, IL 61615, USA

ARTICLE INFO

Article history:

Received 9 April 2021

Revised 9 August 2021

Accepted 16 August 2021

Available online 19 August 2021

Communicated by Zidong Wang

Keywords:

Capsules

Convolutional Neural Network (CNN)

Homogeneous Vector Capsules (HVCs)

MNIST

ABSTRACT

Most capsule network designs rely on traditional matrix multiplication between capsule layers and computationally expensive routing mechanisms to deal with the capsule dimensional entanglement that the matrix multiplication introduces. By using Homogeneous Vector Capsules (HVCs), which use element-wise multiplication rather than matrix multiplication, the dimensions of the capsules remain unentangled. In this work, we study HVCs as applied to the highly structured MNIST dataset in order to produce a direct comparison to the capsule research direction of Geoffrey Hinton, et al. In our study, we show that a simple convolutional neural network using HVCs performs as well as the prior best performing capsule network on MNIST using $5.5\times$ fewer parameters, $4\times$ fewer training epochs, no reconstruction sub-network, and requiring no routing mechanism. The addition of multiple classification branches to the network establishes a new state of the art for the MNIST dataset with an accuracy of 99.87% for an ensemble of these models, as well as establishing a new state of the art for a single model (99.83% accurate).

© 2021 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction and related work

Capsules (vector-valued neurons) have become a more active area of research since [1], which demonstrated near state of the art performance on MNIST [2] classification (at 99.75%) by using capsules and a routing algorithm to determine which capsules in a previous layer feed capsules in the subsequent layer. MNIST is a classic image classification dataset of hand-written digits consisting of 60,000 training images and 10,000 validation images. Studying MNIST, due to the more highly structured content as compared to many other image datasets, allows for the use of more informed data augmentation techniques and when using capsules, the ability to investigate the capsules' interpretability. In [3], the authors extended their work by conducting experiments with an alternate routing algorithm. Research in capsules has since focused mostly on various computationally expensive routing algorithms [4,5]. In [6], we proposed a capsule design that used element-wise multiplication between capsules in subsequent layers and relied on backpropagation to do the work that prior capsule designs were relying on routing mechanisms for. We referred to this capsule design as homogeneous vector capsules (HVCs). In this work, we directly extend the work of [7,1] on capsules applied to MNIST by applying HVCs to MNIST. By using this capsule design, we avoid

the computationally expensive routing mechanisms of prior capsule work and we surpass the performance of [1] on MNIST, all while requiring $5.5\times$ fewer parameters, $4\times$ fewer epochs of training, and using no reconstruction sub-network.

Many of the best performing convolutional neural networks (CNNs) of the past several years have explored multiple paths from input to classification [8–13]. The idea behind multiple path designs is to enable one or more of the following to contribute to the final classification: (a) different levels of abstraction, (b) different effective receptive fields, and (c) valuable information learned early to flow more easily to the classification stage.

In [10] (and subsequent extensions [14–17]) the authors added extra paths through the network with *residual blocks* which are meta-layers that contained one or more convolutional operations as well as a “skip connection” that allowed information learned earlier in the network to skip over the convolutional operations. Similarly, in [8,9], the authors presented a network architecture that made heavy use of *inception blocks*, which are meta-layers that branch from a previous layer into anywhere from 3 to 6 branches of varying layers of convolutions. Then the branches were merged back together by concatenating the filters of those branches. Let n be the average number of branches of different length (in terms of successive convolutions) and m be the number of successive inception blocks. Then $n \times m$ effective receptive fields and levels of abstraction are present at the output of the final inception block. Additionally, the designs presented in both of these papers included two output stems (one branching out before going

* Corresponding author at: Department of Computer Science and Information Systems Bradley University, 1501 WPeoria, Peoria, IL 61615, USA.

E-mail address: abyerly@fsmail.bradley.edu (A. Byerly).

through additional inception blocks and the other after all inception blocks) each producing classification predictions. These classifications were combined via static weighting to produce the final prediction. In contrast to the aforementioned work, in this work, we present a network design that produces 3 output stems, each coming after a different number of convolutions, and thus representing different effective receptive fields and levels of abstraction. We conduct experiments that include statically weighted combinations as in [8,9]. We then go further and investigate learning the branch weights simultaneously with all of the other network parameters via backpropagation. Again, in contrast to the aforementioned work, in these experiments, each of the separate classifications were performed with capsules rather than simple fully connected layers.

Our analysis of the existing literature shows that of the many branching methods explored, those that produced multiple final classifications merged those classifications via static weighting, which presupposes the relative importance of each output. In this work we include and compare the results of both statically weighting the classification branches and learning the weights of the classification branches via backpropagation.

1.1. Our contribution

Our contribution is as follows:

1. We present a novel method for branching a CNN that allows for multiple effective receptive fields and levels of abstraction where each branch makes it's own classification prediction. These classifications are then merged together, each contributing a "vote". We present the results of experiments that include and compare both statically weighting the votes and learning the weights of the votes via backpropagation simultaneously with the rest of the network parameters.
2. We do classification without any fully connected layers, but rather with HVCs. HVCs are simpler, less computationally expensive, and our network design requires $5.5\times$ fewer parameters and $4\times$ fewer training epochs compared to the previously best performing capsule network, all while using no reconstruction sub-network and no computationally expensive routing mechanism.
3. This design, in combination with a domain-specific set of randomly applied augmentation techniques, establishes a new state of the art for the MNIST dataset with an accuracy of **99.87%** for an ensemble of these models, as well as establishing a new state of the art for a single model (**99.83%** accurate).

2. Proposed network design

The starting point for the network design was a conventional convolutional neural network following many widely used practices. These include stacked 3×3 convolutions, each of which with ReLU [18] activation preceded by batch normalization [19]. We also followed the common practice of increasing the number of filters in each subsequent convolutional operation relative to the previous one. Specifically, our first convolution uses 32 filters and each subsequent convolution uses 16 more filters than the previous one. Additionally, the final operation before classification was to softmax the logits and to use categorical cross entropy for calculating loss.

One common design element found in many convolutional neural networks which we intentionally avoided was the use of any pooling operations. We agree with Geoffrey Hinton's assessment [20] of pooling (a method of down-sampling) as an operation to be avoided due to the information it "throws away". With the MNIST data being only 28×28 , we have no need to down-

sample. In choosing not to down-sample, we face the potential dilemma of how to reduce the dimensionality as we descend deeper into the network. This dilemma is solved by choosing not to zero-pad the convolution operations and thus each convolution operation by its nature reduces the dimensionality by 2 in both the horizontal and vertical dimensions. We deem choosing not to zero-pad as preferable in its own right in that zero padding effectively adds information not present in the original sample.

Rather than having a single monolithic design such that each operation in our network feeds into the next operation and only the next operation, we chose to create multiple branches. After the first two sets of three convolutions, in addition to feeding to the subsequent convolution, we also branched off the output to be forwarded on to an additional operation (detailed next). Thus, after all convolutions have been performed, we have three branches in our network.

- 1) The first of which has been through three 3×3 convolutions and consists of 64 filters each having an effective receptive field of 7 of the original image pixels.
- 2) The second of which has been through six 3×3 convolutions and consists of 112 filters each having an effective receptive field of 11 of the original image pixels.
- 3) The third of which has been through nine 3×3 convolutions and consists of 160 filters each having an effective receptive field of 15 of the original image pixels.

For each branch, rather than flattening the outputs of the convolutions into scalar neurons, we instead transformed each filter into a vector to form the first capsule in a pair of homogeneous vector capsules. This operation is represented by "Caps 1(a)", "Caps 2(a)" and "Caps 3(a)" in Fig. 1.

We then performed element-wise multiplication of each of those capsules with a set of weight vectors (one for each class) of the same length. This results in n times m weight vectors where n

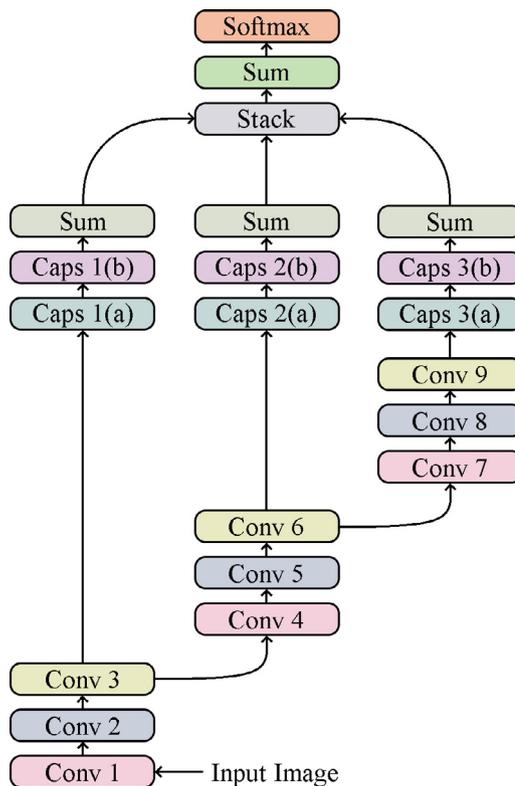


Fig. 1. The proposed network from input to classification.

is the number of capsules transformed from filter maps and m is the number of classes. We summed, per class (m), each of the n vectors to form the second capsule in each pair of homogeneous vector capsules. After this that we applied batch normalization and then ReLU activation. The process elucidated in this paragraph is represented by “Caps 1(b)”, “Caps 2(b)” and “Caps 3(b)” in Fig. 1.

After the pairs of capsules for each breach, the second capsule vector in each pair is reduced to a single value per class by summing the components of the vector. These values can be thought of as the branch-level logits.

Before classifying, the three branch-level sets of logits need to be reconciled with the fact that each image only belongs to one class. This is accomplished by stacking each class’s branch-level logits into vectors of length 3. Then, each vector is reduced by summation to a single value to form the final set of logits to be classified from. Fig. 1 shows the high-level view of the entire network.

In [6], we experimented with a variety of methods for constructing the first layer of capsules out of the preceding filter maps. In this work, we limited our experiments to 2 of these methods (see Fig. 2). The first method constructs each capsule from each distinct feature map (a method that, for brevity, we will refer to as *XY-Derived Capsules* in this work), whereas the second method constructs each capsule from each distinct x and y coordinate of the combination of all of the feature maps (a method that, for brevity, we will refer to as *Z-Derived Capsules* in this work).

We used no weight decay regularization [21], a staple regularization method that improves generalization by penalizing the emergence of large weight values. Nor did we use any form of dropout regularization [22,23] which are regularization methods designed to stop the co-adaptation of weights. We also did not use a reconstruction sub-network as in [1]. These decisions were made in order to investigate the generalization properties of our novel network design elements in the absence of other techniques associated with good generalization. In addition, we intentionally left out any form of “routing” algorithm as in [1,3], preferring to rely on traditional trainable weights and backpropagation.

3. Experimental setup

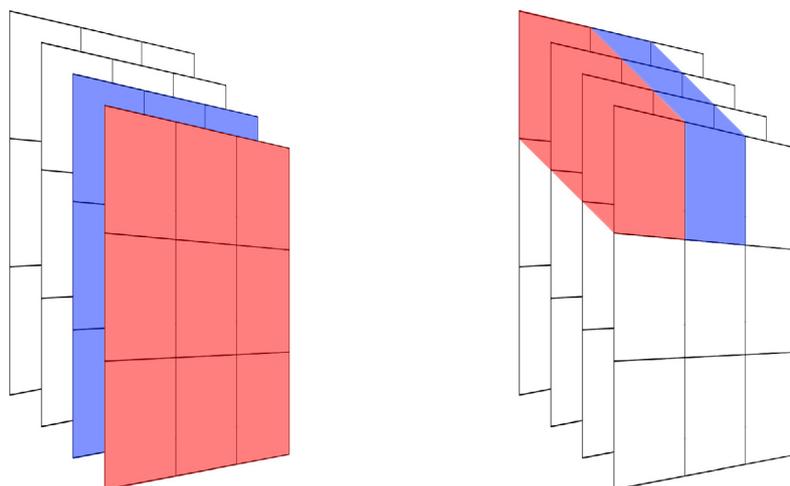
3.1. Merge strategies

In [8,9], the authors chose to give static, predetermined weights to both output branches and then added them together. In our case, for both capsules configurations from Fig. 2, we conducted three separate experiments of 32 trials each in order to investigate the effects of predetermined equal weighting of the branch outputs compared to learning the branch weights via backpropagation:

- 1) **Not learnable.** For this experiment, we merged the three branches together with equal weighting in order to investigate the effect of disallowing any one branch to have more impact than any other.
- 2) **Learnable with randomly initialized branch weights.** (Abbreviated as **Random Init.** subsequently.) For this experiment, we allowed randomly initialized weights to be learned via backpropagation.
- 3) **Learnable with branch weights initialized to one.** (Abbreviated as **Ones Init.** subsequently.) For this experiment, we also allowed the weights to be learned via backpropagation. The difference with the **Random Init.** experiment being that we initialized the weights to 1. We conducted this experiment in addition to the **Random Init.** experiment in order to understand the difference between starting with random weights and starting with equal weights that are subsequently allowed to diverge during training.

3.2. Data augmentation

Most (but not all [24,25]) of the state of the art MNIST results achieved over the past decade have used data augmentation [26,23,13]. In addition to the network design, a major part of our work involved applying an effective data augmentation strategy



(a) A visualization of *XY-Derived Capsules*. Shown are 4 filter maps. The first 2 capsules formed from the filter maps are highlighted in red and blue respectively.

(b) A visualization of *Z-Derived Capsules*. Shown are 4 filter maps. The first 2 capsules formed from the filter maps are highlighted in red and blue respectively.

Fig. 2. Illustrating the construction of capsules from 4 3×3 filter maps. These are the processes denoted by “Caps 1(a)”, “Caps 2(a)”, and “Caps 3(a)” in Fig. 1.

that included transformations informed specifically by the domain of the data. For example, we wanted to be sure we did not rotate our images into being more like a different class (e.g. rotating an image of the digit 2 by 180 degrees to create something that would more closely resemble a malformed 5). Nor did we want to translate the image content off of the canvas and perhaps cut off the left side of an 8 and thus create a 3. Choosing data augmentation techniques specific to the domain of interest is not without precedent (see for example [13,1], both of which used data augmentation techniques specific to MNIST).

By modern standards, in terms of dataset size, MNIST has a relatively low number of training images. As such, judicious use of appropriate data augmentation techniques is important for achieving a high level of generalizability in a given model. In terms of structure, hand-written digits show a wide variety in their rotation relative to some shared true “north”, position within the canvas, width relative to their height, and the connectedness of the strokes used to create them. Throughout training for all trials, every training image in every epoch was subjected to a series of four operations in order to simulate a greater variety of the values for these properties.

- 1) **Rotation.** First, we randomly rotated each training image by up to 30 degrees in either direction. Whether to actually apply this rotation was chosen by drawing from a Bernoulli distribution with probability p of 0.5 (a fair coin toss).
- 2) **Translation.** Second, we randomly translated each training image within the available margin present in that image. In [1], the authors limited their augmentation to shifting the training images randomly by up to 2 pixels in either or both directions. The limit of only 2 pixels for the translation ensured that the translation is label-preserving. As the MNIST training data has varying margins of non-digit space in the available 28×28 pixel canvas, using more than 2 pixels randomly, would be to risk cutting off part of the digit and effectively changing the class of the image. For example, a 7 that was shifted too far left could become more appropriately classed as a 1, or an 8 or 9 shifted far enough down could be more appropriately classed as a zero. The highly structured nature of the MNIST training data allows for an algorithmic analysis of each image that will provide the translation range available for that specific image that will be guaranteed to be label-preserving. Fig. 3 shows an example of an MNIST training image that has an asymmetric translation range that, as long as any translations are performed such that the digit part of the image is not moved by more pixels than are present in the margin, will be label

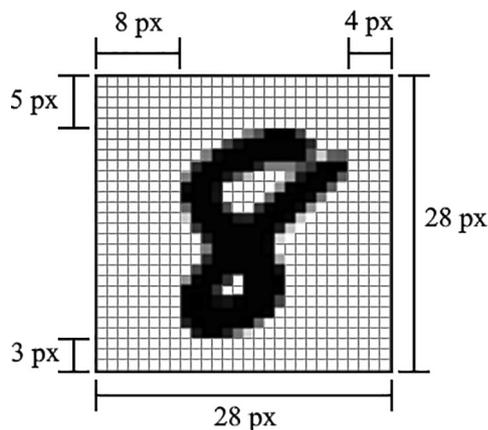


Fig. 3. Example MNIST digit w/annotated margins.

preserving. In other words, the specific training example shown in Fig. 3 could be shifted by up to 8 pixels to the left or 4 to the right and up to 5 up or 3 down, and after doing so, all of the pixels belonging to the actual digit will still be in the resulting translated image. The amount within this margin to actually translate a training image was chosen randomly. Whether to translate up or down and whether to translate left or right were drawn independently from a Bernoulli distribution with probability p of 0.5 (a fair coin toss).

- 3) **Width.** Third, we randomly adjusted each training image's width. MNIST images are normalized to be within a 20×20 central patch of the 28×28 canvas. This normalization is ratio-preserving, so all images are 20 pixels in the height dimension but vary in the number of pixels in the width dimension. This variance not only occurs across digits, but intra-class as well, as different peoples' handwriting can be thinner or wider than average. In order to train on a wider variety of these widths, we randomly compressed each image's width and then added equal zero padding on either side, leaving the digit's center where it was prior. This was inspired by a similar approach adopted in [13]. In our work, we compressed the width of each sample randomly within a range of 0–25%.
- 4) **Random Erasure.** Fourth, we randomly erased (setting to 0) a 4×4 grid of pixels chosen from the central 20×20 grid of pixels in each training image. The X and Y coordinates of the patch were drawn independently from a random uniform distribution. This was inspired by the random erasing data augmentation method in [27]. The intention behind this method was to expose the model to a greater variety of (simulated) connectedness within the strokes that make up the digits. An alternative interpretation would be to see this as a kind of feature-space dropout.

3.3. Training

We followed the training methodology from [6] and trained with the Adam optimizer [28] using all of the default/recommended parameter values, including the base learning rate of 0.001. Also, as in both [6,1], we exponentially decayed the base learning rate. For our experiments, which trained for 300 epochs, we applied an exponential decay to the learning rate at a rate of 0.98 per epoch.

Test accuracy was measured using the exponential moving average of prior weights with a decay rate of 0.999. [29].

4. Experimental results

4.1. Individual models

For both of the capsule construction methods (see Fig. 2) and each of the three merge strategies (see subsection 3.1) we ran 32 trials. Each trial had weights randomly initialized prior to training and, due to the stochastic nature of the data augmentation, a different set of training images. As a result, training progressed to different points in the loss surface resulting in a range of values for the top accuracies that were achieved on the test set. See Table 1.

4.2. Ensembles

Ensembling multiple models together and predicting based on the majority vote among the ensembled models routinely outperforms the individual models' performances. Ensembling can refer to either completely different model architectures with different weights or the same model architecture after being trained multi-

Table 1
Test accuracy of the individual models.

HVC Config.	Experiment	Min	Max	Mean	SD
XY-Derived	Not Learnable	99.71%	99.79%	0.997500	0.0002190
	Random Init.	99.72%	99.78%	0.997512	0.0001499
	Ones Init.	99.70%	99.77%	0.997397	0.0001885
Z-Derived	Not Learnable	99.74%	99.81%	0.997731	0.0001825
	Random Init.	99.73%	99.80%	0.997684	0.0002023
	Ones Init.	99.72%	99.83%	0.997747	0.0002509

In all cases, using the Z-Derived Capsules was superior to using the XY-Derived Capsules. For Z-Derived Capsules, no merge strategy produced statistically significantly superior test accuracy. For XY-Derived Capsules, the only statistically significant test accuracy result was that the Ones Init. strategy produced inferior accuracy. It should be noted that, though no strategy produced statistically significantly superior *test accuracies*, when branches were allowed to learn their weights, the weights learned were statistically significant. (Bold indicates a surpassing of the previous state of the art for individual models on MNIST.) (SD abbreviates Standard Deviation).

ple times and finding different sets of weights that correspond to different locations in the loss surface. The previous state of the art of 99.82% was achieved using an ensemble of 30 different randomly generated model architectures [30]. Our ensembling method used the same architecture but with different weights. We calculated the majority vote of the predictions for all possible combinations of the weights produced by the 32 trials. See Table 2.

4.3. Branch weights

What follows are visualizations of the final branch weights (after 300 epochs of training) for each of the branches in all 32 trials of the experiment wherein the branch weights were initialized to one for both HVC configurations.

In Fig. 4a, we see that for all trials, the ratio between the all three learned branch weights is consistent, demonstrating that the amount of contribution from each branch plays a significant role. In Fig. 4b, we see a similar, though less pronounced consistency between the first branch’s weight and the other two branches, however, branches two and three show no significant difference. Strikingly, when using XY-Derived Capsules we see that branch three (the one having gone through all nine convolutions) has learned to be a more significant contributor. When using Z-Derived Capsules, branch one (the one having gone through only three convolutions) has learned to be a more significant contributor, but only slightly. Indeed, in the latter configuration, the contributions from all three branches is much more equal.

The experiments with randomly initialized branch weights showed the same relative weight of the branches for the magnitude of the weights learned. However, when the initial random branch weight was a negative number, it learned the negative value of that magnitude, and backpropagation took care of flipping the signs of weights as needed further up the network.

Because the models using Z-Derived Capsules are clearly superior to XY-Derived Capsules, unless otherwise stated, all analyses throughout the remainder of this work will restrict attention to these 96 trials, and thus, when the text reads “all 96 trials”, it

Table 2
Test accuracy of the ensembles.

	99.87%	99.86%	99.85%	99.84%	99.83%	99.82%
Using XY-Derived Capsules						
Not Learnable	0	0	0	0	4	1,183
Random Init.	0	0	0	0	21	2,069
Ones Init.	0	0	0	1	19	1,292
Using Z-Derived Capsules						
Not Learnable	184	4,029	89.3K	1.58M	17.7M	121M
Random Init.	0	1,226	533K	17.3M	148M	554M
Ones Init.	64	9,920	1.11M	34.6M	426M	1.27B

Shown here are the number of ensembles that were generated that either matched the previous state of the art of 99.82% or exceeded it.

should be understood that this refers to all 96 trials using Z-Derived Capsules.

4.4. Troublesome digits

Across all 96 trials there was total agreement on 9,912 out of the 10,000 test samples. There were only 14 digits that were misclassified more often than not across all 96 trials. This shows that although the accuracies of the models in the three experiments were quite similar, the different merge strategies of the three experiments did have a significant effect on classification. Across all 96 trials, only 5 samples were misclassified in all models. Those samples, as numbered by the order they appear in the MNIST test dataset (starting from 0) are 1901, 2130, 2597, 3422, and 6576.

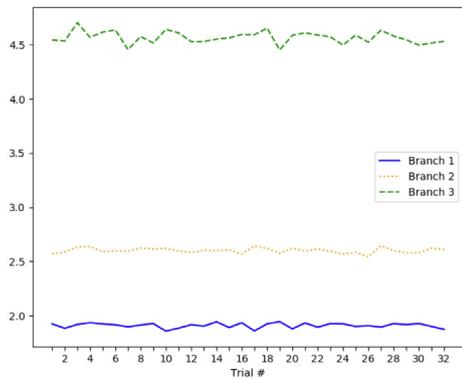
4.5. MNIST state of the art

In Table 3 we present a comparison of previous state of the art MNIST results for both single model evaluations and ensembles along with the results achieved in our experiments.

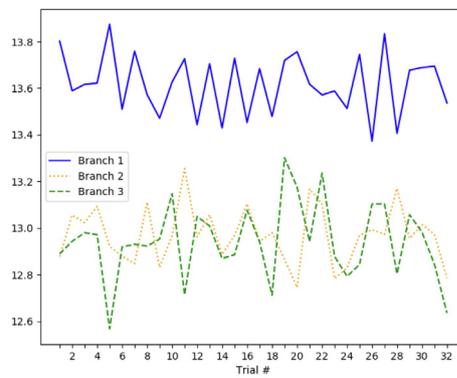
How long a model takes to train is an important factor to consider when evaluating a neural network. Indeed, it is an enabling factor during initial experimentation as faster training leads to a greater exploration of the design space. In Table 4 we present a comparison of the number of epochs of training used in experiments for the results achieved in the networks shown in Table 3. Across all 96 trials, the design achieved peak accuracy in an average of 168 epochs, with a minimum peak achieved in 38 epochs and a maximum peak achieved at epoch 296. Since, all trials were allowed to run for up to 300 epochs, that is the number reported in Table 4.

4.6. Interpreting capsules’ dimensions

By adding a reconstruction sub-network to the overall network, it can be trained not just to classify the input digits, but also to reconstruct them. Then, by following the method in [1], we can



(a) Final branch weights for the experiment using XY-Derived Capsules.



(b) Final branch weights for the experiment using Z-Derived Capsules.

Fig. 4. Final branch weights (after 300 epochs) for 32 trials of the experiments for which the branch weights were initialized to one.

Table 3
Current and previous MNIST state of the art results.

Paper	Year	Accuracy
Single Models		
Dynamic Routing Between Capsules [1]	2017	99.75%
Lets keep it simple, Using simple architectures to outperform deeper and more complex architectures [24]	2016	99.75%
Batch-Normalized Maxout Network in Network [25]	2015	99.76%
APAC:Augmented PAttern Classification with Neural Networks [26]	2015	99.77%
Multi-Column Deep Neural Networks for Image Classification [13]	2012	99.77%
The method proposed in this work	2021	99.83%
Ensembles		
Regularization of Neural Networks using DropConnect [23]	2013	99.79%
RMDL:Random Multimodel Deep Learning for Classification [30]	2018	99.82%
An ensemble of the method proposed in this work	2021	99.87%

Table 4
Epochs of training.

Paper	Epochs
Dynamic Routing Between Capsules [1]	1,200
APAC:Augmented PAttern Classification with Neural Networks [26]	15,000
Multi-Column Deep Neural Networks for Image Classification [13]	800
Regularization of Neural Networks using DropConnect [23]	1,200
RMDL:Random Multimodel Deep Learning for Classification [30]	120
The method proposed in this work	300

Neither [24] nor [25] report on how many epochs their designs were trained for.

examine the effects of perturbing individual dimensions of the second set of capsules in a pair of HVCs. The experiments using Z-Derived Capsules had capsules with 64, 112, and 160 dimensions. When perturbing only one of that many dimensions the changes to the resulting constructed images are very subtle. So we ran another experiment with no branches, reconstruction, and using multiple 8-dimensional capsules for each distinct x and y coordinate of the feature maps. By perturbing one of only eight dimensions the effects are more visible and allows us to interpret the meaning of values in the digits’ capsules (see Table 5).

Table 5
Dimensional perturbations.

Rightward tilt	
Top curl and height of lower loop	
Length of lower stroke	
Angle of the top part of one stroke	
Sharpness of the angle of the lower two curves	
Width of entire digit	
“Hook” in the initial part of the stroke	
Width of lower loop	
Lean angle	

4.7. Ablation experiments

In each of the following set of experiments, we compared the first 10 trials of the 32 trials for the Ones Init. merge strategy with 10 trials each of the additional experiments.

In [1], the authors used a custom loss function they called *margin* loss combined with the mean squared error of the difference between the input images and the result of reconstructing them. In our work and with our design, we chose to rely solely on categorical cross-entropy and not to use a reconstruction loss, as reconstruction adds a considerable number of parameters to the model (2.1M). We ran two additional experiments to understand the

Table 6
Comparison of loss methods.

Loss Method	Mean Accuracy	SD
Categorical Cross-Entropy (no reconstruction)	99.7741%	0.000186455
Categorical Cross-Entropy (with reconstruction)	99.7740%	0.000245764
Margin Loss (with reconstruction)	99.7820%	0.000198997

effect of our choice of loss strategy (which used categorical cross-entropy and no reconstruction). The first used margin loss and reconstruction, and the second used categorical cross-entropy and reconstruction. There was no statistically significant difference among the three loss methods (see Table 6).

In order to understand the relative importance of using HVCs vs. a fully connected layer and 3 branches vs. a single branch, we ran a series of experiments that ablated these components of the architecture. Table 7 shows that HVCs are statistically significantly superior to a fully connected layer for both 1 and 3 branches, and shows that 3 branches are superior to 1 branch for both HVCs and a fully connected layer.

In [1], the authors used translation, by a maximum of 2-pixels, as the only data augmentation method. In our work, we devised a method for translating by up to the full margin available in any given direction. We compared the effect of using only 2-pixel translation, only maximum margin translation, and our full suite of data augmentation methods. Using the full suite of data augmentation methods was shown to be statistically superior to either of the other two methods. Much to our surprise, we found that the 2-pixel translation method just barely crossed the threshold of being statistically significantly superior to the full margin translation method (see Table 8).

The result we obtained by when using 2-pixel translation as the only data augmentation strategy allows for a direct comparison to the work of [1]. We obtained the same level of accuracy as they did, but using $5.5\times$ fewer parameters, $4\times$ fewer training epochs, no reconstruction sub-network, and requiring no routing mechanism.

4.8. Additional datasets

In order to better understand the effect of the Z-Derived HVCs and additional branches, we ran additional sets of paired experiments for several additional datasets. The first set of experiments in each pair used the network design as described in this work, using Z-Derived HVCs and 3 branches (labeled M1 in Table 9) and the second set of experiments excluded the Z-Derived HVCs and additional branches (labeled M2 in Table 9). These second sets of experiments thus use a very small and typical convolutional neural network with $9\ 3\times 3$ convolutions and a final fully connected layer.

For MNIST and Fashion-MNIST we used the data augmentation strategy discussed in subsection 3.2. For CIFAR-10 and CIFAR-100, this data augmentation strategy is inappropriate, so we used a very typical strategy of randomly flipping the images horizontally and applying random adjustments to brightness, contrast, hue, and saturation.

For all four datasets, the model that included Z-Derived HVCs and 3 branches achieved the higher mean accuracy with statistical significance (see Table 9).

The fact that the accuracies for Fashion-MNIST [31], CIFAR-10, and CIFAR-100 [32] were not competitive with current state of the art for those datasets is not especially surprising for several reasons. First, our network was designed for optimal accuracy on classification of Arabic numerals which are highly structured and

Table 7
Comparison of network structures.

Network Structure	Mean Accuracy	SD
Using HVCs and 3 branches	99.7741%	0.000186455
Using HVCs and 1 branch	99.7140%	0.000185472
Using a fully connected layer and 3 branches	99.7550%	0.000111803
Using a fully connected layer and 1 branch	99.6870%	0.000141774

Table 8
Comparison of data augmentation strategies.

Data Augmentation Strategy	Mean Accuracy	SD
Translation (full margin), rotation, width adjustment, and random erasure	99.7741%	0.000186455
Translation only (max. 2 pixels) (as in [1])	99.7570%	0.000195192
Translation only (using full margin)	99.7430%	0.000118743

significantly simpler than the types of data in the other three datasets. Second, due to the significantly simpler nature of MNIST, we used a small number of parameters for our network (1.5M). For comparison, models competitive with state of the art for CIFAR-10 and CIFAR-100 use 10s and even 100s of millions of parameters. Finally, models competitive with state of the art for CIFAR-10 and CIFAR-100 use additional training data beyond the canonical set for each, and we used no additional training data.

5. Conclusion

In this work, we proposed using a simple convolutional neural network and established design principles as a basis for a network architecture. We then presented a design that branched out of the series of stacked convolutions at different points to capture different levels of abstraction and effective receptive fields, and from these branches, rather than flattening to individual scalar neurons, used Homogeneous Vector Capsules instead.

We also investigated three different methods of merging the output of the branches back into a single set of logits. Each of the three merge strategies generated models that could be ensembled to create new state of the art results.

Beyond the network architecture, we proposed a robust and domain specific data augmentation strategy aimed at simulating a wider variety of renderings of the digits.

In doing this work, we established new MNIST state of the art accuracies for both a single model and an ensemble. In addition to the network design and augmentation strategy, the ability to use an adaptive gradient descent method [6] allowed us to achieve this on consumer hardware (2x NVIDIA GeForce GTX 1080 Tis in an otherwise unremarkable workstation) and was an enabling factor in both initial explorations and the training of all 322 trials of experiments referenced in this work.

The code used for all experiments and summary level data is publicly available on GitHub at: <https://github.com/AdamByerly/BMCNNwHFCs>.

CRedit authorship contribution statement

Adam Byerly: Conceptualization, Methodology, Software, Investigation, Writing - original draft, Visualization. **Tatiana Kalganova:** Conceptualization, Writing - review & editing, Supervision. **Ian Dear:** Writing - review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A

A.1. Digits disagreed upon

What follows is the complete set of 88 digits that were predicted correctly by at least one model and incorrectly by at least one model. These in combination with the digits from Fig. 5 represent

Table 9
Effects of Z-Derived HVCs and branching on additional datasets.

Dataset	Network	Max	Mean	SD	p-value
MNIST	M1	99.81%	99.7741%	0.0001864	1.824×10^{-7}
	M2	99.71%	99.6870%	0.0001417	
Fashion-MNIST	M1	93.89%	93.6850%	0.0016391	5.243×10^{-6}
	M2	93.36%	93.0410%	0.0014616	
CIFAR-10	M1	89.23%	88.9290%	0.0015514	0.020898
	M2	89.06%	88.7500%	0.0017515	
CIFAR-100	M1	64.15%	63.8260%	0.0026743	6.859×10^{-6}
	M2	62.96%	62.3760%	0.0035046	

M1 denotes the network design as described in this work, using Z-Derived HVCs and 3 branches. M2 denotes a network design that excludes the Z-Derived HVCs and additional branches. MNIST results come from the same experiments detailed in Table 7 and are repeated here to facilitate ease of comparison. We conducted 10 trials of each unique type of experiment in order to establish statistical significance.

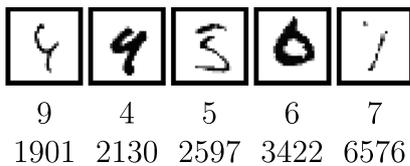
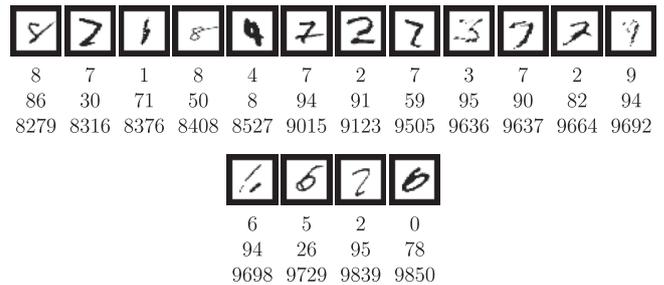
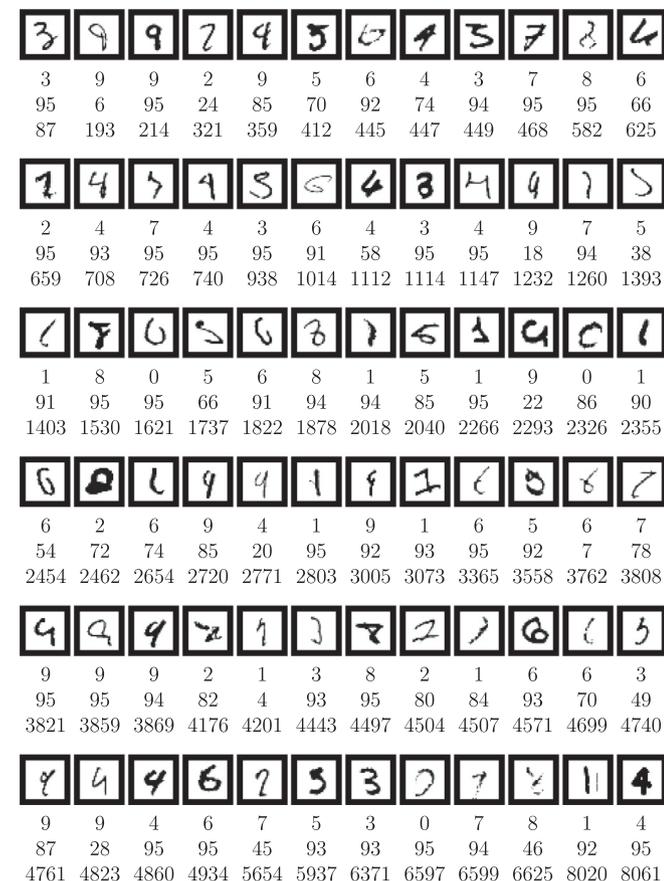


Fig. 5. The most troublesome digits.

the complete set of digits that were not predicted correctly by all 96 trials. Each image is captioned first by the class label in the test data set associated with the image, then the number of trials that predicted it correctly, and last the index of the digit in the test data. For example, the first image presented below has a class label of 3, 95 trials predicted that correctly, and it exists at index 87 in the MNIST test data.



References

- [1] Sara Sabour, Nicholas Frosst, Geoffrey E. Hinton, Dynamic routing between capsules, in: Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017), 2017.
- [2] Yann LeCun, Corinna Cortes, C.J. Burges, MNIST handwritten digit database, ATT Labs [Online] 2 (2010), URL: <http://yann.lecun.com/exdb/mnist>.
- [3] Geoffrey E. Hinton, Sara Sabour, Nicholas Frosst, Matrix Capsules with EM Routing, in: Proceedings of the 6th International Conference on Learning Representations (ICLR 2018), 2018.
- [4] Sai Raam Venkataraman, S. Balasubramanian, R. Raghunatha Sarma, Building Deep Equivariant Capsule Networks., in: International Conference on Learning Representations, 2020. .
- [5] Mohammed Amer, Tomás Maul, Path Capsule Networks 52 (2020) 545–559. doi: 10.1007/s11063-020-10273-0. .
- [6] Adam Byerly, Tatiana Kalganova, Homogeneous vector capsules enable adaptive gradient descent in convolutional neural networks, IEEE Access 9 (2021) 48519–48530, <https://doi.org/10.1109/ACCESS.2021.3066842>.
- [7] Geoffrey E. Hinton, Alex Krizhevsky, Sida D. Wang, Transforming auto-encoders, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 6791 LNCS, 2011, pp. 44–51. doi: 10.1007/978-3-642-21735-7_6..
- [8] Christian Szegedy et al., Going Deeper with Convolutions, in: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 1–9.
- [9] Christian Szegedy et al., Rethinking the Inception Architecture for Computer Vision, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 2818–2826.
- [10] Kaiming He et al., Deep Residual Learning for Image Recognition, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 770–778.
- [11] B. Zhou, et al., BBN: Bilateral-Branch Network With Cumulative Learning for Long-Tailed Visual Recognition, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 9716–9725. doi: 10.1109/CVPR42600.2020.00974..
- [12] C. Wang et al., CSPNet: A New Backbone that can Enhance Learning Capability of CNN, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 2020, pp. 1571–1580, <https://doi.org/10.1109/CVPRW50498.2020.00203>.
- [13] Dan C. Ciresan, Ueli Meier, Jürgen Schmidhuber, Multi-Column Deep Neural Networks for Image Classification, in: 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2012, pp. 3642–3649.
- [14] Rupesh Kumar Srivastava, Klaus Greff, Jürgen Schmidhuber, Highway Networks, 2015, arXiv: 1505.00387 [cs.LG].
- [15] S. Xie, et al., Aggregated Residual Transformations for Deep Neural Networks, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 5987–5995. doi: 10.1109/CVPR.2017.634. .
- [16] S. Jégou et al., The One Hundred Layers Tiramisu: Fully Convolutional DenseNets for Semantic Segmentation, in: 2017 IEEE Conference on

Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 1175–1183.

- [17] Hang Zhang et al., ResNeSt: Split-Attention Networks, ArXiv abs/2004.08955 (2020).
- [18] Xavier Glorot, Antoine Bordes, Yoshua Bengio, Deep Sparse Rectifier Neural Networks, in: Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS 2011), 2011.
- [19] Sergey Ioffe, Christian Szegedy, Batch Normalization, Accelerating Deep Network Training by Reducing Internal Covariate Shift, in: Proceedings of the 32nd International Conference on Machine Learning (ICML 2015), 2015.
- [20] Geoffrey E. Hinton, What's wrong with convolutional nets? MIT Tech TV, 2018. URL: <https://techtv.mit.edu/collections/bcs/videos/30698-what-s-wrong-with-convolutional-nets>.
- [21] Geoffrey E. Hinton, Learning translation invariant recognition in a massively parallel networks, in: PARLE Parallel Architectures and Languages Europe, Springer, Berlin Heidelberg, 1987, pp. 1–13, isbn: 978-3-540-47144-8.
- [22] Geoffrey E. Hinton, et al., Improving Neural Networks by Preventing Co-Adaptation of Feature Detectors, 2012. arXiv: 1207.0580v1 [cs.NE].
- [23] Li Wan et al., Regularization of Neural Networks using DropConnect, in: Proceedings of the 30th International Conference on Machine Learning (ICML 2013), 2013.
- [24] Seyyed Hossein Hasanpour, et al., Lets keep it simple, Using simple architectures to outperform deeper and more complex architectures, 2016, arXiv: 1608.06037 [cs.CV].
- [25] Jia-Ren Chang, Yong-Sheng Chen, Batch-Normalized Maxout Network in Network, 2015. arXiv: 1511.02583 [cs.CV].
- [26] Ikuro Sato, Hiroki Nishimura, Kensuke Yokoi, APAC: Augmented PAttern Classification with Neural Networks, 2015. arXiv: 1505.03229 [cs.CV].
- [27] Zhun Zhong, et al., Random Erasing Data Augmentation, 2017. arXiv: 1708.04896 [cs.CV].
- [28] Diederik P. Kingma, Jimmy Ba, Adam: A Method for Stochastic Optimization, in: Proceedings of the 6th International Conference on Learning Representations (ICLR 2014), 2014.
- [29] Pavel Izmailov et al., Averaging weights leads to wider optima and better generalization, in: Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence (UAI 2018), 2018.
- [30] Kamran Kowsari et al., RMDL: Random multimodel deep learning for classification, in: Proceedings of the 2nd International Conference on Information System and Data Mining (ICISDM 2018), 2018, pp. 19–28.
- [31] Han Xiao, Kashif Rasul, Roland Vollgraf, Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms, 2017. arXiv: 1708.07747 [cs.LG].
- [32] A. Krizhevsky, Learning Multiple Layers of Features from Tiny Images, Tech. rep. (2009).



Adam Byerly received the M.S. degree in Computer Science from Bradley University, Peoria, IL, USA, in 2015 and is currently pursuing the Ph.D. degree with the Department of Electronic and Electrical Engineering at Brunel University London, Uxbridge, UK. He now lectures in the Department of Computer Science and Information Systems at Bradley University, Peoria, IL, USA. Prior to transitioning to academia, he accumulated over 20 years of industry experience in software engineering and architecture in various industries including marketing, financial, insurance, and manufacturing.



Tatiana Kalganova (BUL:PI:TK) BSc (Hons), PhD, is a Reader in Intelligent Systems and Postgraduate Research Director in Electronic and Electrical Engineering and Director of the AI Center at Brunel University London, Uxbridge, UK. She has over 20 years of experience in design and implementation of applied Intelligent Systems.



Ian D. Dear BSc (Hons), MSc, PhD in Digital VLSI Test strategies, is a Senior Tutor in Electronic and Electrical Engineering at Brunel University London, Uxbridge, UK. He has over 35 years of experience in design and testing of complex Digital Systems.