

# StitchNet: Composing Neural Networks from Pre-Trained Fragments

Surat Teerapittayanon  
National Nanotechnology Center  
Pathum Thani, Thailand  
surat.tee@nanotec.or.th

Marcus Comiter  
Harvard University  
Cambridge, MA, USA  
marcuscomiter@post.harvard.edu

Bradley McDanel  
Franklin & Marshall College  
Lancaster, PA, USA  
bmcdanel@fandm.edu

H.T. Kung  
Harvard University  
Cambridge, MA, USA  
kung@harvard.edu

**Abstract**—We propose StitchNet, a novel neural network creation paradigm that stitches together fragments (one or more consecutive network layers) from multiple pre-trained neural networks. StitchNet allows the creation of high-performing neural networks without the large compute and data requirements needed under traditional model creation processes via backpropagation training. We leverage Centered Kernel Alignment (CKA) as a compatibility measure to efficiently guide the selection of these fragments in composing a network for a given task tailored to specific accuracy needs and computing resource constraints. We then show that these fragments can be stitched together to create neural networks with accuracy comparable to that of traditionally trained networks at a fraction of computing resource and data requirements. Finally, we explore a novel on-the-fly personalized model creation and inference application enabled by this new paradigm. The code is available at <https://github.com/steerapi/stitchnet>.

**Index Terms**—StitchNet, Neural Networks, Deep Learning, Centered Kernel Alignment (CKA), Reusable Network Components

## I. INTRODUCTION

AI models have become increasingly more complex to support additional functionality, multiple modalities, and higher accuracy. While the increased complexity has improved model utility and performance, it has imposed significant model training costs. Therefore, training complex models is often infeasible for resource-limited environments, such as those at the cloud edge and in fully or often disconnected environments.

In response to these challenges, this paper proposes a new paradigm for creating neural networks: rather than training networks from scratch or retraining existing networks, we create neural networks through composition by *stitching together* fragments of existing pre-trained neural networks. A fragment is one or more consecutive layers of a neural network. We call the resulting neural network composed of one or more fragments a “StitchNet” (Fig. 1). By significantly reducing the amount of computation and data resources needed to create neural networks, StitchNet enables an entire new set of applications, such as the rapid generation of personalized neural networks at the edge or in fully disconnected environments.

StitchNet’s model creation mechanism is fundamentally different from today’s predominant backpropagation-based method for creating neural networks. Given a dataset and a task as input, the traditional training method uses backpropagation with stochastic gradient descent (SGD) or other

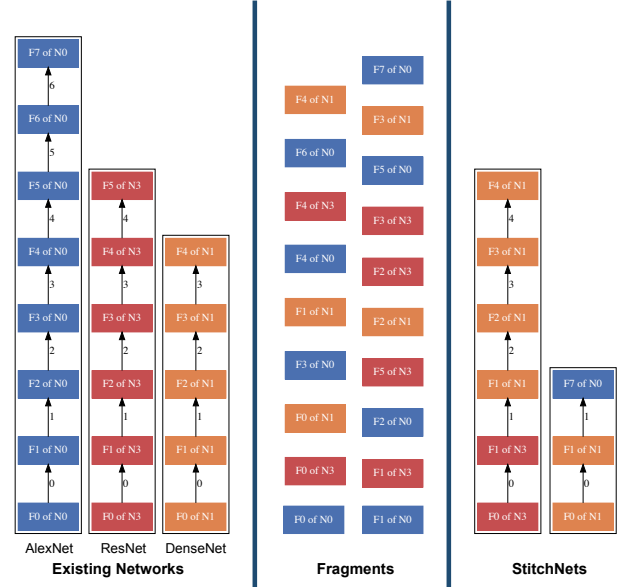


Fig. 1: Overview of the StitchNet approach. Existing networks (left) are cut into fragments (middle), which are composed into StitchNets (right) created for a particular task. No retraining is needed in this process.

optimization algorithms to adjust the weights of the network. This training process iterates through the full dataset multiple times, and therefore requires compute resources that scale with the amount of data and the complexity of the network. Training large models in this way also requires substantial amounts of data to mitigate overfitting. While successful, this traditional paradigm for model creation is not without its limitations, especially as AI moves out of the data center and into highly resource-constrained and disconnected environments. Creating complex neural networks without access to large amounts of data and compute resources is a growing challenge. In the extreme case (e.g., for very large language models (LLMs) and computer vision models), only a few companies with access to unrivaled amounts of data and compute resources are able to create such models.

StitchNet solves this problem by creating new neural networks using fragments of already existing neural networks. The new approach takes advantage of the growing amount

of neural networks that already exist. StitchNet enables the efficient reuse of the learned knowledge resident in those pre-trained networks, which has been distilled from large amounts of data, rather than having to relearn it over and over again for new tasks as is done with traditional model creation paradigms. StitchNet’s ability to reuse existing pre-trained fragments, rather than recreating from scratch or retraining for every task, will help accelerate the growth and application of neural networks for solving more and more complex tasks in heterogenous environments. For example, StitchNets can be created on-the-fly to serve as classifiers specially tuned for local conditions and classes of interest in a given environment or task.

However, compositing these existing fragments into a coherent and high-performing neural network is non-trivial. To reuse the knowledge of pre-trained neural network fragments, we need a way to 1) measure the compatibility between any two fragments, and 2) compose compatible fragments together. In the past, Centered Kernel Alignment (CKA) [1]–[3] has been used to measure similarity between neural network representations. We leverage CKA to assess the compatibility of any two fragments from any neural networks, and compose new neural networks from fragments of existing pre-trained neural networks to create high-performing networks customized for specific tasks without the costs of traditional model creation methods. The CKA score is used to reduce the search space to identify compatible fragments and guide the fragment selection process.

We present empirical validations on benchmark datasets, comparing the performance of StitchNet to that of the original pre-trained neural networks. We demonstrate that StitchNet achieves comparable or higher accuracy on personalized tasks compared with off-the-shelf networks and has significantly lower computational and data requirements than training networks from scratch or by fine-tuning.

Our contributions are:

- The StitchNet paradigm: a novel neural network creation method with versatile applications.
- Innovative use of CKA to assess fragment compatibility.
- Technique for seamlessly combining compatible fragments in both linear and convolutional layers.

## II. COMPOSING FRAGMENTS

The core mechanism to create StitchNets is to identify reusable fragments from a pool of existing networks and to compose them into a coherent neural network model capable of performing a given task. To this end, we need a way to determine how compatible any two candidate fragments are with each other. In previous work, Kornblith, et al. [1] presented centered kernel alignment (CKA) [2], [3] as a way to measure the similarity between neural network representations. Rather than looking at the neural network as a whole, we adopt and use CKA to as a measure of compatibility between any two *fragments* of any neural networks.

In this section, we first define CKA as a way to measure how compatible any two fragments are with one another

and therefore their ability to be composed. Using CKA, we then present a technique to stitch different fragments together. Finally, we describe the algorithm to generate StitchNets.

### A. Centered Kernel Alignment (CKA)

We define  $\mathbf{X} \in \mathbb{R}^{p \times n}$  as outputs of a fragment  $F_A$  of model  $A$  and  $\mathbf{Y} \in \mathbb{R}^{q \times n}$  as inputs of a fragment  $F_B$  of model  $B$  of the same dataset  $\mathbf{D}$ , where  $n$  is the number of samples in the dataset,  $p$  is the output dimension of  $F_A$ , and  $q$  is the input dimension of  $F_B$ . Given a target dataset  $\mathbf{D}$ , we define the compatibility score between any two fragments as  $\text{CKA}(\mathbf{X}, \mathbf{Y})$  (specifically linear CKA) of fragment  $F_A$  and fragment  $F_B$ :

$$\frac{\|\text{cov}(\mathbf{X}^T \mathbf{X}, \mathbf{Y}^T \mathbf{Y})\|_F^2}{\sqrt{\|\text{cov}(\mathbf{X}^T \mathbf{X}, \mathbf{X}^T \mathbf{X})\|_F^2 \|\text{cov}(\mathbf{Y}^T \mathbf{Y}, \mathbf{Y}^T \mathbf{Y})\|_F^2}}. \quad (1)$$

To reduce memory usage for a large  $\mathbf{D}$ , CKA can be approximated by averaging over minibatches as presented in [4].

### B. Stitching Fragments

Once we have determined compatible fragments based on a configurable compatibility threshold, the next step in creating a StitchNet is to *stitch* the two fragments together. To do so, we find a projection tensor  $\mathbf{A}$  that projects the output space of one fragment to the input space of the other fragment we are composing. We now describe this.

Without loss of generality, we assume the output and input tensors are 2D tensors, where the first dimension is the sample dimension. If the tensors are not 2D tensors, we first flatten all other dimensions with the exception of the sample dimension.

We use Einstein summation notation, where  $i$  represents the sample dimension,  $j$  the output dimension of the incoming fragment, and  $k$  the input dimension of the outgoing fragment. Given an output tensor  $\mathbf{X}_{ij}$  of the incoming fragment and an input tensor  $\mathbf{Y}_{ik}$  of the outgoing fragment, we seek to find  $\mathbf{A}$  such that  $\mathbf{Y}_{ik} = \mathbf{A}_{kj} \mathbf{X}_{ij}$ . We can then solve for  $\mathbf{A}$  using the Moore-Penrose pseudoinverse:

$$\mathbf{A}_{kj} = \mathbf{Y}_{ik} \mathbf{X}_{ij}^T (\mathbf{X}_{ij} \mathbf{X}_{ij}^T)^{-1} \quad (2)$$

Once  $\mathbf{A}$  is found, we fuse  $\mathbf{A}$  with the weight of the first layer of the outgoing fragment. For linear layers, we simply do the following:  $\mathbf{W}'_{lk} = \mathbf{W}_{lj} \mathbf{A}_{kj}$ , where  $l$  is the dimension of the output feature of the outgoing layer.

For convolutional layers, we first upsample or downsample the spatial dimension to match each other, and then adjust the weight along the input channel dimension as follows:  $\mathbf{W}'_{okmn} = \mathbf{W}_{ojmn} \mathbf{A}_{kj}$ , where  $o$  is the output channel dimension,  $j$  is the original input channel dimension,  $k$  is the new input channel dimension, and  $m$  and  $n$  are the spatial dimensions.

For stitching a convolutional layer with an output tensor  $\mathbf{X}$  and a linear layer with an input tensor  $\mathbf{Y}$ , we first apply adaptive average pooling so that the spatial dimension is 1x1 and flatten  $\mathbf{X}$  into a 2D tensor. Then, we follow Eq. 2 to find  $\mathbf{A}$  and fuse it with the  $\mathbf{W}$  of the linear layer.

### C. StitchNet Generation

We now describe the main algorithm for creating StitchNet networks (“StitchNet” for short), shown in Algorithm 1. A StitchNet network is created by joining a set of pre-trained network fragments drawn from a pool  $\mathbf{P} = \{F_{ijk}\}$ . We use the notation  $F_{ijk}$  to denote a fragment of a neural network  $i$  from layer  $j$  to layer  $k$ , and the notation  $N_{ik}$  to denote the computation performed by the portion of the neural network from which the fragment was taken up to layer  $k$ . Other than the fragment pool  $\mathbf{P}$  and creation process hyperparameters ( $K, T, L$ ), the only other input to the StitchNet creation process is a dataset  $\mathbf{D}$  for which the StitchNet will be optimized. The dataset  $\mathbf{D}$  can be chosen to include relevant samples and classes of interest for the specific task at hand, as illustrated in Sec. III-B, III-F, and III-G.

---

#### Algorithm 1 StitchNet( $\mathbf{P}, \mathbf{D}, K, T, L, \mathbf{R}, Q, s$ )

---

```

1: Input: fragment pool  $\mathbf{P} = \{F_{ijk}\}$ , network  $i$  in  $\mathbf{P}$ 
   up to layer  $k$   $N_{ik}$ , fragment starting with layer  $j$  and
   ending in layer  $k$  of network  $i$   $F_{ijk}$ , target dataset  $\mathbf{D}$  with
    $M$  samples, span  $K$ , threshold  $T$ , maximum number of
   fragments  $L$ , result set of StitchNets and their associated
   scores  $\mathbf{R}$ , current StitchNet  $Q$ , current score  $s$ 
2: Output: result set of StitchNets and their associated
   scores  $\mathbf{R}$ 
3: if  $Q$  is empty then
4:    $\{F_{ijk}\} = \text{select starting fragments in } \mathbf{P}$ 
5:   for  $F_{ijk}$  in  $\{F_{ijk}\}$  do
6:     StitchNet( $\mathbf{P}, \mathbf{D}, K, T, L, \mathbf{R}, F_{ijk}, 1$ )
7: if the number of fragments in  $Q \geq L$  then
8:   return  $\mathbf{R}$ 
9:  $\{F_{ijk}\} = \text{select } K \text{ middle or terminating fragments in } \mathbf{P}$ 
10: for  $F_{ijk}$  in  $\{F_{ijk}\}$  do
11:    $\mathbf{X} = Q(\mathbf{D}); \mathbf{Y} = N_{ik}(\mathbf{D})$ 
12:    $s_n = s \times \text{CKA}(\mathbf{X}, \mathbf{Y})$  (see section II-A)
13:   if  $s_n > T$  then
14:      $Q = \text{Stitch}(Q, F_{ijk}, \mathbf{X}, \mathbf{Y})$  (see section II-B)
15:     if  $F_{ijk}$  is a terminating fragment then
16:        $\mathbf{R} = \mathbf{R} \cup \{Q, s_n\}$ 
17:     else
18:       StitchNet( $\mathbf{P}, \mathbf{D}, K, T, L, \mathbf{R}, Q, s_n$ )
19: return  $\mathbf{R}$ 

```

---

We now describe the creation of the pool of network fragments  $\mathbf{P}$  derived from a set of pre-trained off-the-shelf networks. These pre-trained networks are divided into one of three types of fragments: *starting* fragments for which the input is the original network input, *terminating* fragments for which the output is the original network output, and *middle* fragments that are neither starting nor terminating fragments.

The first step in the StitchNet creation process is to choose the set of starting fragments. This could include all starting fragments in  $\mathbf{P}$ , or a subset based on certain criteria, e.g., the smallest, biggest or closest starting fragment.

Once a set of starting fragments is selected, a StitchNet is built on top of each starting fragment that has a current starting

score of 1. First, a set of  $K$  candidate fragments are selected from  $\mathbf{P}$ . These fragments can be selected based on CKA scores (i.e.,  $K$  fragments with the highest CKA scores), the number of parameters of the fragments (i.e.,  $K$  fragments with the least number of parameters in  $\mathbf{P}$ ), the closest fragments (i.e.,  $K$  fragments with the least latency in  $\mathbf{P}$  in a distributed setting), or other selection methods.

For each of the candidate fragments, we then compute two intermediate neural network computations in order to derive the compatibility score for further stitching. First, we pass the dataset  $\mathbf{D}$  through the candidate StitchNet in its current form, resulting in the value  $\mathbf{X}$ . Second, we pass the same dataset  $\mathbf{D}$  through the neural network from which the candidate fragment  $F_{ijk}$  was selected, resulting in the value  $\mathbf{Y} = N_{ik}(\mathbf{D})$ .

After running these computations, we compute  $\text{CKA}(\mathbf{X}, \mathbf{Y})$  as in Sec. II-A. We then multiply the CKA with the current score  $s$  to obtain the new current score  $s_n$ . If  $s_n$  is still greater than a set threshold  $T$ , the candidate fragment is selected and the process continues recursively. Otherwise, the candidate fragment is rejected. The threshold can be set to balance the amount of exploration allowed per available compute resources.

This process continues until a terminating fragment is selected, the maximum number of fragments  $L$  is reached, or all recursive paths are exhausted. At this point, the completed StitchNets and their associated scores  $\mathbf{R}$  are returned for user selection.

### III. RESULTS

We now demonstrate that StitchNets can perform comparably with traditionally trained networks but with significantly reduced computational and data requirements at both inference and creation time. Through these characteristics, StitchNet enables the immediate on-the-fly creation of neural networks for personalized tasks without traditional training.

#### A. Fragment pool

To form the fragment pool  $\mathbf{P}$ , we take five off-the-shelf networks pre-trained on the ImageNet-1K dataset [5] from Torchvision [6]: *alexnet*, *densenet121*, *mobilenet\_v3\_small*, *resnet50* and *vgg16* with IMAGENET1K\_V1 weights.<sup>1</sup>

These pre-trained networks are cut into fragments at each convolution and linear layer that has a single input. As shown in Fig. 2, there are 8 fragments for *alexnet*, 5 fragments for *densenet121*, 13 fragments for *mobilenet\_v3\_small*, 6 fragments for *resnet50* and 16 fragments for *vgg16*. This results in the creation of a fragment pool  $\mathbf{P}$  of 48 fragments consisting of 5 starting fragments, 38 middle fragments, and 5 terminating fragments. We use this fragment pool in all experiments.

#### B. Dataset

The dataset used to evaluate StitchNet in the first batch of experiments is the “Dogs vs. Cats” dataset [7]. This dataset includes 25,000 training images of dogs and cats, and we

<sup>1</sup>In practice, many more or specialized pre-trained networks can be included in the pool, including smaller and larger networks.

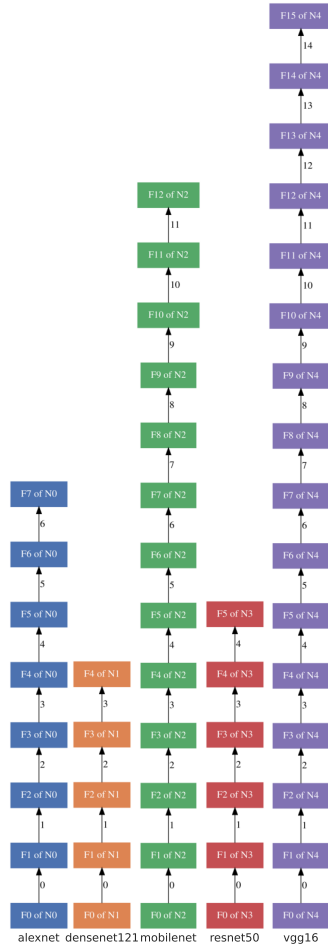


Fig. 2: Five pre-trained networks are fragmented into a fragment pool  $P$ . These fragments will be stitched together to form StitchNets.

use an 80:20 train:test split. We map ImageNet-1K class labels into cat and dog labels (class IDs 281-285 and 151-250, respectively). To form the target dataset  $D$  for use in the stitching process of Algorithm 1, we randomly select  $M$  samples from the training set as the target dataset  $D$ . We choose this task as it is characteristic of the type of task where StitchNet would be used: a user needs a custom classifier for a particular task and desired set of classes, but any general neural network problem formulation can also be used.

### C. StitchNet Generation

We generate StitchNets with Algorithm 1 using the fragment pool and the dataset described in Sec. III-A and III-B. We set  $K = 2$ ,  $T = 0.5$  and  $L = 16$ . The number of samples  $M$  in  $D$  used for the stitching process is 32.

Given these hyperparameters, a total of 89 StitchNets are generated. We evaluate them on the test set of completely unseen test samples. Summary statistics for the generated StitchNets are shown in Fig. 3, including accuracy (3a), number of fragments per StitchNet (3b), CKA score (3c), and number of parameters per StitchNet (3d).

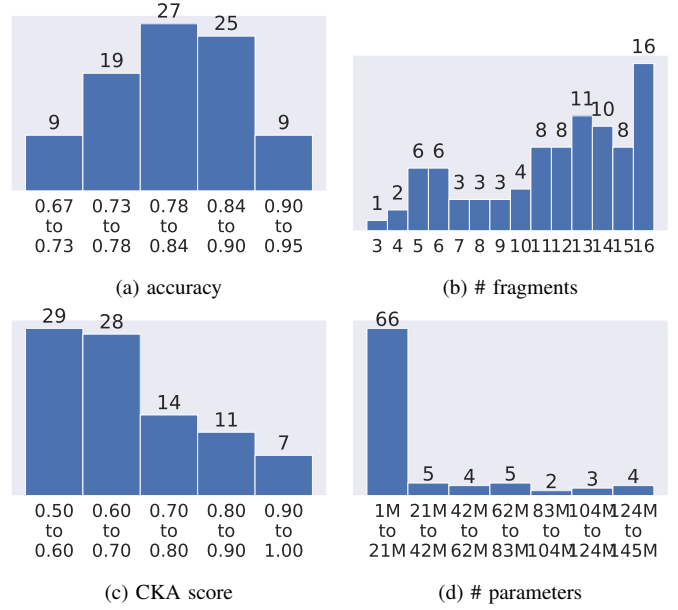


Fig. 3: Histogram of (a) accuracy, (b) # fragments, (c) CKA score, (d) # parameters in the generated batch of StitchNets.

### D. Reduction in Inference Computation

We now demonstrate how StitchNet significantly reduces inference-time computational requirements over traditional neural network training paradigms by studying StitchNet accuracy as a function of parameters.

Fig. 4 shows the resulting accuracy of the generated StitchNets as a function of the overall CKA score (top) and Multiply-Accumulate Operations (MACs) (middle) for each generated StitchNet and the number of parameters (proportional to marker size) as a proxy for inference-time computation cost. We find a number of StitchNets outperform the pre-trained network while realizing significant computational savings. For example, StitchNet27 (denoted by an orange star) achieves an accuracy of 0.86 with 3.6M parameters compared with the 0.70 accuracy of the pre-trained alexnet with 61.1M parameters. Therefore, StitchNet achieves a 16-point absolute increase in accuracy with a 94.1% reduction in number of parameters for the given task when compared with those of the pre-trained alexnet. Similar trends hold in comparison of StitchNets with the other pre-trained models. Fig. 5 shows the composition of some of these high-performing StitchNets, demonstrating the diversity in fragment use, ordering, and architectures.

These results crystallize one of the core benefits of StitchNet: without significant data and computation requirements of traditional training procedures, the method can discover networks that are personalized for the task, outperform the original pre-trained networks, and do so while significantly reducing inference-time compute requirements. This is due to the fact that these pre-trained networks are not trained to focus on these two specific classes, while our StitchNets are stitched together specifically for the task. In the next section, we will

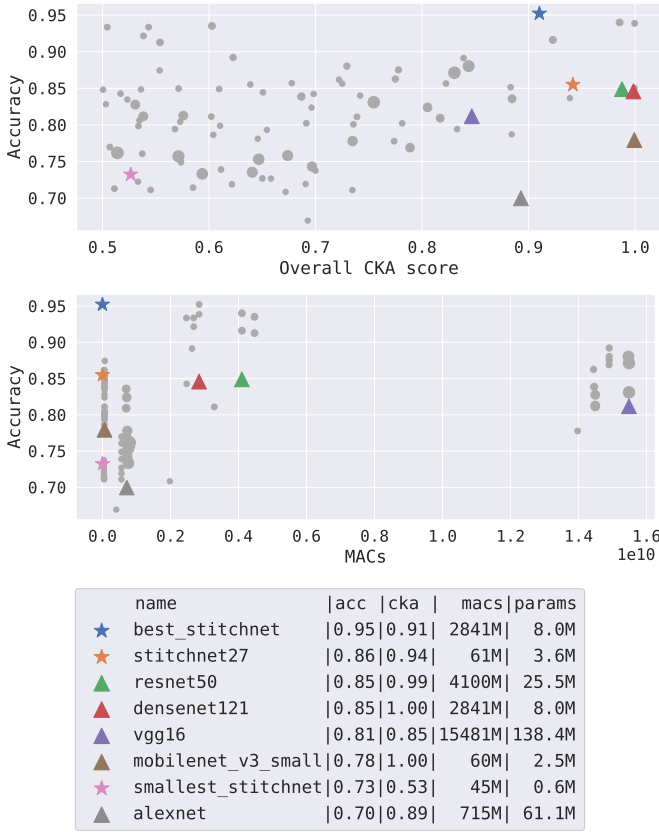


Fig. 4: Accuracy vs Overall CKA score (top) and MACs (middle) on “Cat vs. Dogs.” *macs* is the number of Multiply-Accumulate Operations (MACs), *acc* is the accuracy and *cka* is the overall CKA score. The gray dots are the StitchNets generated. The best StitchNet (*acc*=0.95) has 10-point higher accuracy than the best pre-trained model(s) (densenet121 and resnet50 with *acc*=0.85).

see that very little data is required for the stitching process.

We also validate the effectiveness of using CKA to guide the stitching procedure. StitchNets with high CKA scores (especially above 0.9) also have high accuracy. Therefore CKA can be a proxy to measure compatibility between connecting fragments.<sup>2</sup> But while high CKA scores imply high accuracy, having high accuracy does not necessarily imply equally high CKA scores. Regardless, CKA is still useful as a heuristic because it removes highly incompatible choices (i.e., those with very low CKA scores), therefore saving computation time in the stitching process. These savings are substantial as there are more candidates with low CKAs than high CKAs.

### E. Reduction in Network Creation Computation

We now demonstrate that StitchNet can be created without significant data and computation requirements. Specifically,

<sup>2</sup>Note that there exist high accuracy StitchNets with low overall CKA scores. This is because neural networks are highly redundant and invariant to small changes, making them able to tolerate a certain amount of errors while still providing quality predictions (see Sec. IV-A).

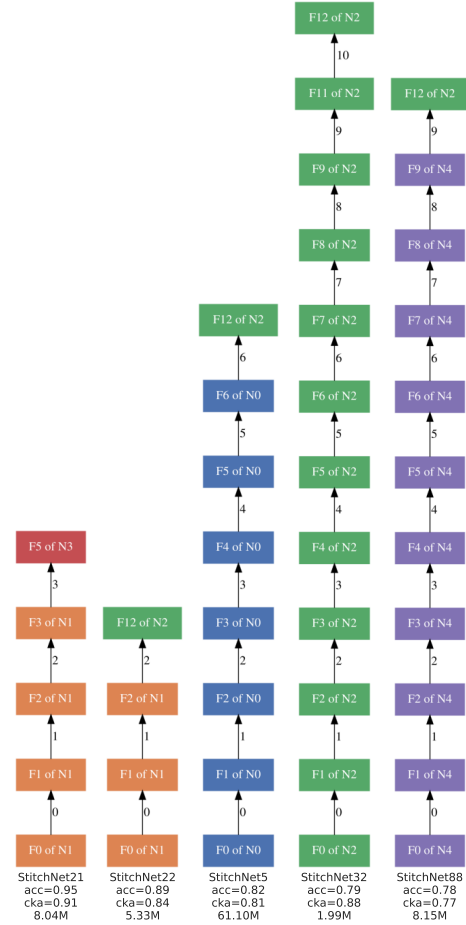


Fig. 5: Examples of generated StitchNets.

we compare StitchNet21 (generated in Fig. 5 on the target dataset of  $M = 32$  samples) with fine-tuning the same five off-the-shelf networks (retraining them using the training portion of dataset of Sec. III-B). For fine-tuning, we replace and train only the last layer of the pre-trained network using SGD with batch size 32, learning rate 0.001 and momentum 0.9. The results shown are averaged over 10 runs. For ease of comparison, we normalize the computation cost in terms of the number of samples processed through a neural network. In practice, fine-tuning requires backpropagation, which incurs additional computation per sample than StitchNet generation.

Fig. 6 compares the accuracy of StitchNet21 (denoted by the red star) with the traditionally fine-tuned networks as a function of the number of training samples processed. For a given accuracy target, StitchNet processes a substantially smaller number of data samples than traditionally fine-tuned networks. Specifically, to reach an accuracy of 0.95, fine-tuning of *alexnet*, *densenet121*, and *mobilenet\_v3\_small* requires processing more than 320 samples, while StitchNet requires only 32 samples used to stitch the fragments together (realizing over a 90% reduction).

Therefore, only a small amount of training samples and computation are required for StitchNet to achieve comparable

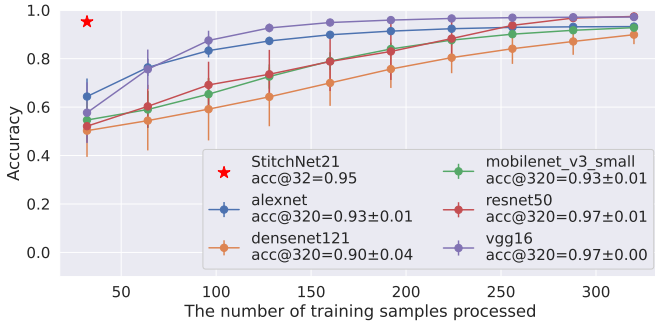


Fig. 6: Accuracy vs the number of training samples processed (i.e., data and computation required). StitchNet requires only a fraction of the computation of traditional training methods to achieve comparable performance.

accuracy. This demonstrates that StitchNet effectively reuses the information already captured in the fragments to bootstrap network creation. This allows for personalization of tasks and on-the-fly training without substantial data requirements.

#### F. StitchNet on New Tasks

We now show that the StitchNet approach of leveraging existing pre-trained network fragments is highly general. It achieves high performance on new specific tasks and domains not explicitly used in training the pre-trained network fragments. To demonstrate this, we use an entirely different dataset for evaluation that shares no overlap with the ImageNet-1K dataset used to train the pre-trained network fragments. For this experiment, we use the “Beans” dataset [8], which includes images of diseased and healthy beans leaves. This dataset is divided into three categories: *angular\_leaf\_spot*, *bean\_rust*, and *healthy*. We use 32 samples from the validation split of the dataset to generate StitchNets. Then, we use all the validation split of 133 samples to train a KNN classifier that uses the final 1,000-dimensional vectors obtained from the networks as feature inputs. To evaluate network accuracy, we use the 128-sample test split. The results are shown in Fig. 7. We see that the top-performing StitchNet with an accuracy (*acc*) of 0.84 and 3296M MACs outperforms the original pre-trained model (vgg16) in both accuracy and efficiency, 0.79 and 15481M MACs respectively. This shows again that with StitchNet, we can generate the most efficient and smallest neural network tailor-made for this specific task, eliminating the need for retraining the neural networks for different tasks.

#### G. On-the-fly Network Creation and Inference

We now delve into a novel set of applications and scenarios that are made possible by StitchNet: on-the-fly neural network creation and inference. In this scenario, we use a batch of unlabeled images that we wish to use for a particular task (e.g., classification or detection) as our target dataset in the generation of StitchNets. With only a minor modification to the StitchNet algorithm to additionally return task results, the StitchNet generation process can return the inference outputs along with the generated StitchNets.

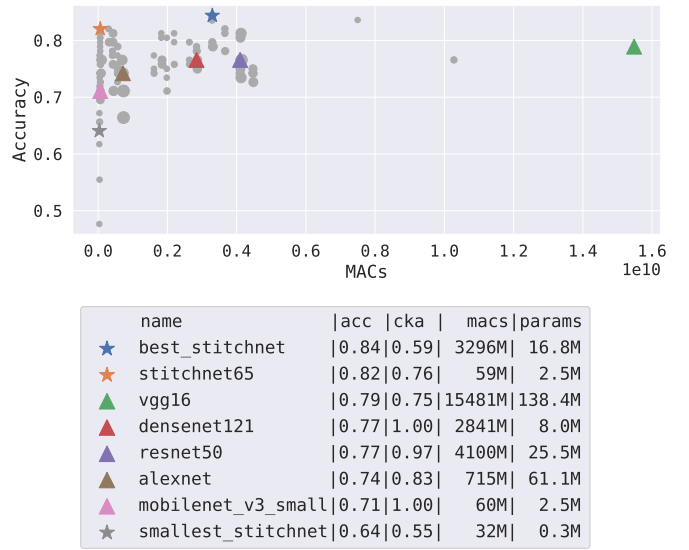


Fig. 7: Accuracy vs MACs on “Beans.” The best StitchNet (*acc*=0.84) with only 16.8M parameters and 3296M MACs performs better than the best pre-trained model(s) (vgg16 with *acc*=0.79 and 15481M MACs).

Let us illustrate how this could be applied in practice. Imagine a world where fragments of pre-trained neural networks for different tasks are indexed and distributed on the Internet. Compatible fragments can be found and organized into a pool from which candidate fragments can be selected and quickly composed to form a new StitchNet for a certain task. Now, imagine that we want to create a neural network for classifying local food images with only a limited number of labeled images.

Without StitchNet, we either need to train a network from scratch (which may fail due to our limited amount of training data), or find an existing pre-trained neural network, label the dataset, and finetune the network. If the existing pre-trained network is too big or too slow for our use, we will then have to train a new one from scratch. But, with limited amount of data, this task seems impossible.

With StitchNet, we can instead generate a set of candidate StitchNets on-the-fly with the small batch of unlabeled local food images. These StitchNets are created from the pool of existing neural network fragments that have been indexed and distributed over the Internet. The proper fragments can be identified with a given search criteria (e.g., the depth of the network should be less than 5 for computational efficiency reasons, etc.). With only the relatively limited computation needed to find stitches rather than the normal high-cost network training methods, we can generate various StitchNets capable of detecting and classifying local food images at the desired computation requirements.

To demonstrate this new paradigm, we use the “Food101” dataset [9] and method described in Sec. III-F to demonstrate the feasibility of this use case. In this experiment, we explore 5 different classification tasks. Each task consists of 3 classes

randomly picked out of 101 classes from the dataset. The classes for each task are listed in the x-axis of Fig. 8. Each task has around 300 samples for training and 30 samples for testing. 32 samples of the training set are used to generate StitchNets. As shown in Fig. 8, in all tasks top-performing StitchNets surpass all original pre-trained networks. This underscores StitchNet’s efficacy in generating efficient and high-performing networks instantly in response to new tasks.

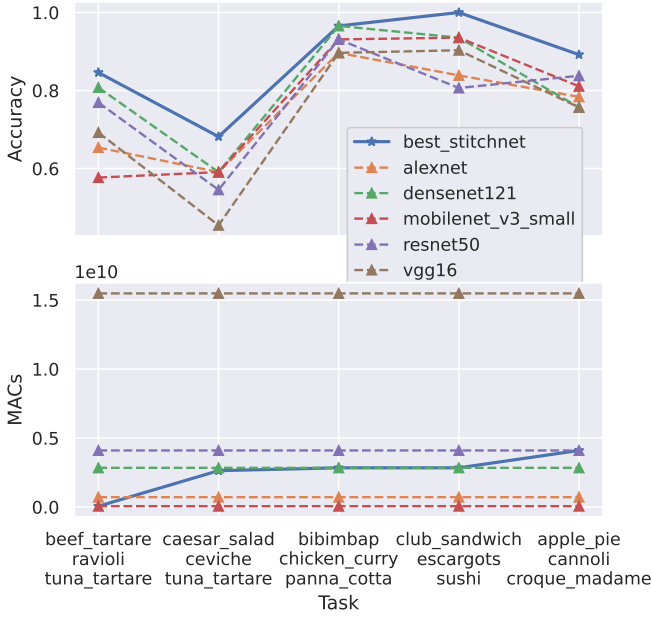


Fig. 8: Accuracy and MACs as a function of various tasks derived from the “Food101” dataset.

#### IV. DISCUSSION

##### A. Why does *StitchNet* work?

We first discuss why we can reuse existing fragments of networks to create new neural networks without retraining. One core reason is that neural networks tend to learn fundamental and universal features. Studies [1], [10]–[15] have shown that neural networks learn fundamental features, such as edges for different tasks. Since these learned features are fundamental, they should be reusable rather than relearned. The challenge, however, is that although these features may be universal, they may not be compatible with one another “out of the box.” Therefore, we require the stitching process introduced in Sec. II-B to project the fragments into a compatible space.

Beyond the reuse of universal features and compatibility transformations, *StitchNet* benefits from the inherent redundancy in networks. This redundancy, arising from nonlinear activations and redundant features, allows networks to tolerate certain levels of error. Consequently, individual fragments do not need to be perfectly compatible to form a network that operates effectively as a whole.

##### B. Complexity Comparison

We now compare the complexity of the traditional training process using backpropagation with the *StitchNet* generation process. Traditional training complexity is  $O(ndp)$ , where  $n$  is the number of parameters in the network,  $p$  is the number of epochs used to train, and  $d$  is the size of the dataset.

The complexity of *StitchNet* generation is  $O(nqm) + O(K^L) + O(nq)$ . The first term  $O(nqm)$  is the evaluation cost of the target dataset of size  $q$  on  $m$  networks in the pool, where  $q \ll d$  and  $n$  is the number of parameters in the network (assuming that the networks have the same number of parameters). The second term  $K^L$  is the search cost, where  $K$  is the span value we search at each level and  $L$  is the maximum depth of search. Using a high threshold cutoff  $T$  on the overall CKA score keeps search cost  $K^L$  small. The third term  $O(nq)$  is the cost of the stitching (See Sec. II-B). Therefore, for a reasonable setting of hyperparameters ( $K, T, L$ ) in Algorithm 1, *StitchNet* realizes substantial computation gains over traditional training methods since  $q \ll d$  and  $m \ll p$ .

##### C. Limitations

While a large pool of network fragments can lead to higher applicability and quality of the generated *StitchNets*, it can also result in high search costs. Indexing large quantities of networks to form the fragment pool will necessitate the development of novel search methods. We see this as analogous to indexing web pages on the World Wide Web, suggesting a “Google for Fragments.” Much like the web search needed to index written content, large amounts of neural network “content” need to be indexed for their value to be unlocked. Early indexing efforts can tag fragments based on dataset characteristics, computational characteristics, etc. More advanced efforts can look at inward and outward connections of each fragment to determine its rank in results. Once a narrowed set of fragments is coarsely identified, the efficient procedure introduced in this paper can generate the *StitchNets*. Future work will address indexing and distribution that will enable *StitchNet* to operate at scale.

#### V. RELATED WORK

Current methods to adapt existing networks to target tasks are transfer learning, fine-tuning [16], distillation [17] and neural architecture search (NAS) [18]. A related concept is unsupervised domain adaptation [19]–[25], where an existing network is modified using an unlabeled target dataset. *StitchNet* works similarly by stitching fragments using an unlabeled target dataset to create a neural network for the target task. What distinguishes *StitchNet* from most of the previous works is that it eliminates the need for retraining the network.

*StitchNet* takes advantage of the assumption that fragments have shareable representations. This assumption helps explain why fragments can be stitched together into a coherent high-performing network: dissimilar yet complementary fragments once projected into a similar space are compatible with one another. Several existing works including [1], [10]–[15], [26] have studied this shareable representation assumption.

Gygli, et al. [27] reuse network components by training networks to produce compatible features by adding regularization during training to make networks compatible. StitchNet, however, focuses on creating neural networks without retraining the network. It is therefore more generally applicable. Several works [14], [28]–[31] combine network components by adding a stitching layer and training the recombined network with supervised loss for several epochs. StitchNet differs by approaching the topic through an applied lens focusing on minimizing training costs by adding a parameter-less stitching mechanism, and therefore does not require any retraining. Instead, weights are adapted to be compatible (See Sec. II-B).

## VI. CONCLUSION

StitchNet is a new paradigm that can take advantage of a growing global library of neural networks to fundamentally change the way neural networks are created. By reusing fragments of these networks to efficiently compose new neural networks for a given task, StitchNet addresses two of the most fundamental issues that limit the creation and use of neural networks: large data and computation requirements.

StitchNet does this by leveraging CKA as a compatibility measure that guides the selection of network fragments, tailored to specific accuracy needs and computing resource constraints. Our work has shown that neural networks can be efficiently created from compatible network fragments of different models at a fraction of computing resources and data requirements while achieving comparable accuracy. We also explore a novel on-the-fly efficient neural network creation and inference application unlocked by StitchNet. Future work will extend StitchNet to LLMs and additional applications.

## ACKNOWLEDGMENT

This research was supported in part by the Air Force Research Laboratory under contract number FA8750-22-1-0500 and in part by a generous gift to the Center for Research on Computation and Society at Harvard University in support of research on applied cryptography and society.

## REFERENCES

- [1] S. Kornblith, M. Norouzi, H. Lee, and G. Hinton, “Similarity of neural network representations revisited,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 3519–3529.
- [2] C. Cortes, M. Mohri, and A. Rostamizadeh, “Algorithms for learning kernels based on centered alignment,” *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 795–828, 2012.
- [3] N. Cristianini, J. Kandola, A. Elisseeff, and J. Shawe-Taylor, “On kernel target alignment,” in *Innovations in machine learning*. Springer, 2006, pp. 205–256.
- [4] T. Nguyen, M. Raghu, and S. Kornblith, “Do wide and deep networks learn the same things? uncovering how neural network representations vary with width and depth,” *arXiv preprint arXiv:2010.15327*, 2020.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [6] S. Marcel and Y. Rodriguez, “Torchvision the machine-vision package of torch,” in *Proceedings of the 18th ACM international conference on Multimedia*, 2010, pp. 1485–1488.
- [7] Kaggle, “Dogs vs. cats,” 2013. [Online]. Available: <https://www.kaggle.com/competitions/dogs-vs-cats/data>
- [8] M. A. Lab. (2020, January) Bean disease dataset. [Online]. Available: <https://github.com/AI-Lab-Makerere/ibean/>
- [9] L. Bossard, M. Guillaumin, and L. Van Gool, “Food-101 – mining discriminative components with random forests,” in *European Conference on Computer Vision*, 2014.
- [10] Y. Li, J. Yosinski, J. Clune, H. Lipson, J. E. Hopcroft *et al.*, “Convergent learning: Do different neural networks learn the same representations?” in *FE@ NIPS*, 2015, pp. 196–212.
- [11] Q. Lu, P.-H. Chen, J. W. Pillow, P. J. Ramadge, K. A. Norman, and U. Hasson, “Shared representational geometry across neural networks,” *arXiv preprint arXiv:1811.11684*, 2018.
- [12] A. S. Morcos, M. Raghu, and S. Bengio, “Insights on representational similarity in neural networks with canonical correlation,” *arXiv preprint arXiv:1806.05759*, 2018.
- [13] L. Wang, L. Hu, J. Gu, Y. Wu, Z. Hu, K. He, and J. Hopcroft, “Towards understanding learning representations: To what extent do different neural networks learn the same representation,” *arXiv preprint arXiv:1810.11750*, 2018.
- [14] K. Lenc and A. Vedaldi, “Understanding image representations by measuring their equivariance and equivalence,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 991–999.
- [15] S. Tang, W. J. Maddox, C. Dickens, T. Diethe, and A. Dami-anou, “Similarity of neural networks with gradients,” *arXiv preprint arXiv:2003.11498*, 2020.
- [16] K. Weiss, T. M. Khoshgoftaar, and D. Wang, “A survey of transfer learning,” *Journal of Big data*, vol. 3, no. 1, pp. 1–40, 2016.
- [17] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [18] H. Pham, M. Guan, B. Zoph, Q. Le, and J. Dean, “Efficient neural architecture search via parameters sharing,” in *International conference on machine learning*. PMLR, 2018, pp. 4095–4104.
- [19] M. Wang and W. Deng, “Deep visual domain adaptation: A survey,” *Neurocomputing*, vol. 312, pp. 135–153, 2018.
- [20] W. Zhang, W. Ouyang, W. Li, and D. Xu, “Collaborative and adversarial network for unsupervised domain adaptation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3801–3809.
- [21] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, “Deep domain confusion: Maximizing for domain invariance,” *arXiv preprint arXiv:1412.3474*, 2014.
- [22] A. Kumar, R. Sattigeri, K. Wadhawan, L. Karlinsky, R. Feris, W. T. Freeman, and G. Wornell, “Co-regularized alignment for unsupervised domain adaptation,” *arXiv preprint arXiv:1811.05443*, 2018.
- [23] R. Shu, H. H. Bui, H. Narui, and S. Ermon, “A dirt-t approach to unsupervised domain adaptation,” *arXiv preprint arXiv:1802.08735*, 2018.
- [24] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan, “A theory of learning from different domains,” *Machine learning*, vol. 79, no. 1, pp. 151–175, 2010.
- [25] K. Saito, Y. Ushiku, and T. Harada, “Asymmetric tri-training for unsupervised domain adaptation,” in *International Conference on Machine Learning*. PMLR, 2017, pp. 2988–2997.
- [26] J. Mehrer, N. Kriegeskorte, and T. Kietzmann, “Beware of the beginnings: intermediate and higherlevel representations in deep neural networks are strongly affected by weight initialization,” in *Conference on Cognitive Computational Neuroscience*, 2018.
- [27] M. Gygli, J. Uijlings, and V. Ferrari, “Towards reusable network components by learning compatible representations,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 9, 2021, pp. 7620–7629.
- [28] A. Csizsárik, P. Kőrösi-Szabó, A. Matszangosz, G. Papp, and D. Varga, “Similarity and matching of neural network representations,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 5656–5668, 2021.
- [29] Y. Bansal, P. Nakkiran, and B. Barak, “Revisiting model stitching to compare neural representations,” *Advances in neural information processing systems*, vol. 34, pp. 225–236, 2021.
- [30] X. Yang, D. Zhou, S. Liu, J. Ye, and X. Wang, “Deep model reassembly,” *Advances in neural information processing systems*, vol. 35, pp. 25 739–25 753, 2022.
- [31] Z. Pan, J. Cai, and B. Zhuang, “Stitchable neural networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 16 102–16 112.