



HAL
open science

WASH: Train your Ensemble with Communication-Efficient Weight Shuffling, then Average

Louis Fournier, Adel Nabli, Masih Aminbeidokhti, Marco Pedersoli, Eugene
Belilovsky, Edouard Oyallon

► **To cite this version:**

Louis Fournier, Adel Nabli, Masih Aminbeidokhti, Marco Pedersoli, Eugene Belilovsky, et al.. WASH: Train your Ensemble with Communication-Efficient Weight Shuffling, then Average. 2024. hal-04588075

HAL Id: hal-04588075

<https://hal.science/hal-04588075>

Preprint submitted on 25 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WASH: Train your Ensemble with Communication-Efficient Weight Shuffling, then Average

Louis Fournier

ISIR - Sorbonne Université
Paris - France
louis.fournier@isir.upmc.fr

Adel Nabli

Mila, Concordia University
Sorbonne Université
adel.nabli@sorbonne-universite.fr

Masih Aminbeidokhti

École de technologie supérieure
Montréal - Québec

Marco Pedersoli

École de technologie supérieure
Montréal - Québec

Eugene Belilovsky

Concordia University
Mila – Quebec AI Institute
Montréal - Québec

Edouard Oyallon

Center for Computational Mathematics
Flatiron Institute
New York, USA

Abstract

The performance of deep neural networks is enhanced by ensemble methods, which average the output of several models. However, this comes at an increased cost at inference. Weight averaging methods aim at balancing the generalization of ensembling and the inference speed of a single model by averaging the parameters of an ensemble of models. Yet, naive averaging results in poor performance as models converge to different loss basins, and aligning the models to improve the performance of the average is challenging. Alternatively, inspired by distributed training, methods like DART and PAPA have been proposed to train several models in parallel such that they will end up in the same basin, resulting in good averaging accuracy. However, these methods either compromise ensembling accuracy or demand significant communication between models during training. In this paper, we introduce WASH, a novel distributed method for training model ensembles for weight averaging that achieves state-of-the-art image classification accuracy. WASH maintains models within the same basin by randomly shuffling a small percentage of weights during training, resulting in diverse models and lower communication costs compared to standard parameter averaging methods.

1 Introduction

In order to enhance the accuracy of a given class of models, aggregating the answers of multiple instances trained in parallel can be done via model *ensembling*. This can lead to significant improvements in modern deep learning models (11), increasing generalization ability. However, this comes at the cost of evaluating multiple instances of a given model at inference. This increases both the required memory and computations, resources which can be critical for on-device inference (31). To resolve this issue, the population of models can be fused into a single model to obtain both the generalization improvements of ensembling and the inference cost of a single model. Since independent models can be linearly connectable (13), a simple technique is to average the weights of the different models to obtain a fused model (51).

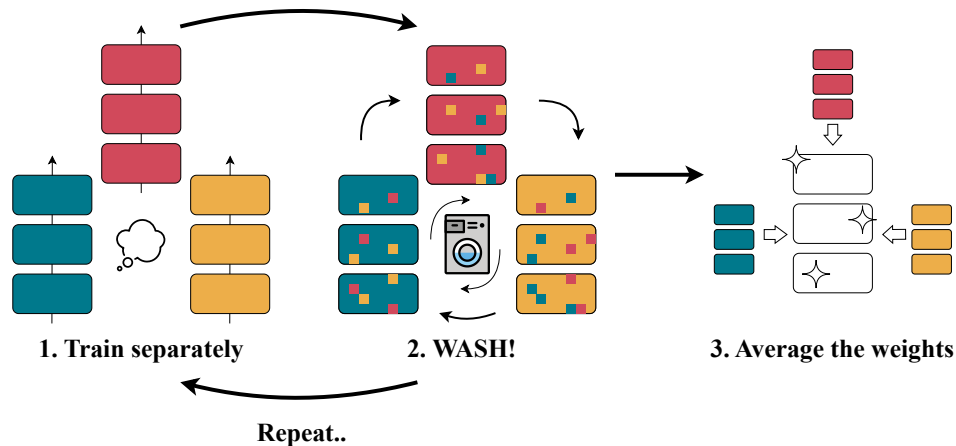


Figure 1: **Representation of training with WASH.** A population of models is being trained separately. **(1)** After each training step, **(2)** a small percentage of parameters are permuted between models. **(3)** At the end of the training, the model weights are averaged, resulting in a high-performance model.

There are however limits to this method. For models that are too dissimilar, the performance of the averaged model may not be better than chance (18). To mitigate this, the ensemble can either use a pre-trained network as a starting point (33) or ensure that models share part of their optimization path (13). Nevertheless, diminishing too much the ensemble diversity comes at the cost of its performance (see Fig. 6 of (11)), revealing a tradeoff between model diversity and weight averageability. Inspired by distributed training, techniques like DART (19) and PAPA (20) have been proposed to train a population of models in parallel on heterogeneous data while communicating to balance this tradeoff. DART, similarly to LocalSGD (43), averages all the models regularly to avoid models diverging. PAPA controls the diversity of the models more finely, by pushing them towards the averaged parameters using an Exponential Moving Average (EMA) like EASGD (54), achieving better performances. Notably, they show that training a population in such a way results in models that generalize better than a single model trained with the same compute as the entire population, demonstrating the potential of these distributed approaches. However, existing methods require a regular computation of the average model using an all-reduce operation, either to remove periodically any diversity in the population (19) or in the case of PAPA, to compute an EMA of the average. This results in a high communication cost during the parallel training of the population of models (35), hindering the scalability of these approaches as the population size increases (34).

We propose in this paper a novel distributed method to train a population of models in parallel while keeping their weights within the same basin. It requires a fraction of the communication cost of PAPA but displays greater model diversity during training, increasing the final averaging accuracy. Our main idea is to shuffle parameters between models during training, forcing them to learn using the others' parameters. We refer to 'parameter shuffling' as the following idea. A permutation is randomly chosen, and models will communicate peer-to-peer their parameters following the permutation. The use of a permutation is distinct from the notion of weight permutation of (1) which is inside one model. We denote our method, which achieves **Weight Averaging** using parameter **SHuffling**, as **WASH**, and represent it schematically in Fig. 1.

Contributions. Our work makes the following contributions: **(1)** We propose a novel method for the training of a population of models that can be weight averaged, which we refer to as WASH (**Weight Averaging** using parameter **SHuffling**). By shuffling a small number of parameters between models during training, the resulting population can be weight-averaged into a high-performance model for a fraction of the communication volume of methods such as PAPA. **(2)** We find that WASH provides state-of-the-art results on image classification tasks, resulting in models with performances at the level of ensembling methods while only requiring a single network at inference time. **(3)** We provide experiments to better understand the improvement provided by WASH, notably on how WASH reduces the distance between models in the population implicitly while maintaining diversity. **(4)** We perform different ablations on our method, showing the impact of the shuffling. **(5)** Our code is made available at github.com/fournierlouis/WASH.

2 Related work

Ensemble and weight averaging. By combining predictions from multiple models, ensemble methods significantly improve a predictive system’s ability to make accurate generalizations (7; 25), while reducing the variance of estimator (4). This variance reduction is especially effective when errors are uncorrelated and models show diversity, meaning they do not fail simultaneously on the same instances (16; 11). However, ensembles require additional passes through each model for inference, leading to increased computational costs. This cost can become prohibitive with a large number of models. As a remedy under some constraints, models can be averaged together to remove the computational burden during inference. Averaging the weights of models was first explored in simple linear (26) and convex scenarios (37; 3). In deep learning, (18) establish that weight averaging is a first-order approximation of the ensemble when models are close in the weight space. Notably, a simple averaging of multiple points along the SGD trajectory leads to better generalization. Following mode connectivity (15; 13) and the observation that many optima of independent models are connectable, (2; 50) propose learning simplexes in parameter space with a regularization penalty to encourage diversity in weight space, and (52; 38) propose to train multiple model branches with different last-layer initialization and hyperparameters concurrently. These models are later averaged to enhance generalization and reduce the inference cost. However, for these models to be amenable to weight averaging, they need to begin with the same pre-trained initialization (33) which can diminish diversity among the models. To alleviate this issue, neuron alignment techniques (41; 1; 36; 17) match the units of multiple networks to make them amenable to weight averaging, however, they rarely work in practical scenarios (21), often obtaining performance below the individual models. DART (19) and Branch-Train-Merge (BTM) (27) propose a three-phase training pipeline. The process begins with an initial shared training phase, followed by the parallelized training of multiple models, each diversified through different data domains or different data augmentations. Finally, these models are merged into a single model. They find that iterative refinement of the last 2 stages enhances the overall optimization trajectory and improves generalization. To enhance diversity among models, PAPA (20) proposes to rather gradually adjust model weights towards the population average throughout the training process, beginning from random initialization. However, these approaches can result in substantial communication costs during training. Conversely, WASH tackles high communication costs by permuting only a small fraction of parameters among models during training, while ensuring the branches remain amenable to weight averaging at the end.

Distributed and federated learning. In distributed training of deep learning models, the trade-off between communications and model performance is a core concern (34; 22), and finding methods efficiently alleviating some of its communication cost is a recurrent theme in different areas of research (45; 12). For instance, communication overhead being a key concern of decentralized optimization, it has been shown in this literature that to train models in a data-parallel setting with a limited communication budget, a key metric to observe is the average distance to the consensus (23; 39; 44; 48; 32). The techniques discussed earlier to train a population of models for weight averaging are similar to methods in the LocalSGD (43; 29) and Federated Learning (30; 22; 28) literature. The training in DART and BTM is similar to the LocalSGD training, where models are averaged regularly after several steps of computations. PAPA, which uses an EMA of the averaged model to gradually move the models towards consensus, is similar to methods like EASGD (54) or SlowMo (47). Only averaging a population at the end of training like in BTM was also proposed for LocalSGD (42), and cross-gradient aggregation (10) can be seen as a way to shuffle gradients locally. Federated learning also uses techniques discussed previously for model merging (46; 53; 5). Finally, our method can be thought of as training a global model where each local model picks randomly from a subset of parameters when shuffled. This can be linked to Bayesian learning (14) in particular for federated learning (49; 8), or to federated subnetwork training (9; 40).

3 Parameter shuffling in an ensemble for weight averaging

Motivation of our training procedure. We aim to balance the benefits of model ensembling with the computational efficiency of using a single model at inference via weight averaging. In other words, our objective is to produce a single model resulting from the ensembling. A set of N model parameters $\{\theta_n\}_{n \leq N} \subset \mathbb{R}^d$ are trained in parallel on the same dataset, with a different data order and

Algorithm 1 Training with WASH

```
1: Input: Datasets  $D_i$ , number of models  $N$ , initial parameters  $\theta_0$ , training steps  $T$ , number of
   layers  $L$ , base probability  $p$ 
2: Initialize parameters  $(\theta_n)_n \leftarrow \theta_0$  and optimizers  $\text{OPT}_i$ 
3: for  $t = 1$  to  $T$  do
4:   # Training step
5:   for  $n = 1$  to  $N$ , in parallel do
6:      $(x_n, y_n) \leftarrow D_n$            # Sample data
7:      $\theta_n \leftarrow \text{OPT}_n(x_n, y_n, \theta_n)$  # Update the model  $n$ 
8:   # Shuffling step
9:   for layer  $l = 0$  to  $L - 1$  do
10:    for parameter  $\theta^i$  in layer  $l$  do
11:      With probability  $p(1 - \frac{l}{L-1})$ ,
12:         $\pi_i \leftarrow \text{Random permutation}$ 
13:         $(\theta_n^i)_n \leftarrow (\theta_{\pi_i(n)}^i)_n$  # Send and permute the parameter
14: Output: the averaged model  $\frac{1}{N} \sum_{n=1}^N \theta_n$ 
```

possibly different data augmentations and regularizations. To avoid divergence among the models, PAPA applies an EMA every T training steps and produces the following update

$$\tilde{\theta}_n \leftarrow \alpha \theta_n + (1 - \alpha) \bar{\theta}, \quad (1)$$

where $\bar{\theta} \triangleq \frac{1}{N} \sum_{n=1}^N \theta_n$ represents the average of the model weights, also referred to as the *consensus*, and $\alpha \in]0, 1[$ is weighted depending on the learning rate. Despite its benefits, this method has drawbacks, including the need for synchronized global communication across all models, which can be inefficient, and the potential reduction in model diversity due to the consensus constraint, which may reduce model expressivity. Indeed, we observe that after each update

$$\sum_n \|\tilde{\theta}_n - \bar{\theta}\|^2 = \alpha^2 \sum_n \|\theta_n - \bar{\theta}\|^2 < \sum_n \|\theta_n - \bar{\theta}\|^2, \quad (2)$$

which shows that the EMA step of methods like PAPA directly reduces the distance of the models to the consensus, hindering their diversity.

Proposed method: WASH. To address these challenges, we propose the following stochastic parameter shuffling step instead of the EMA, defined for any individual parameter $\theta_n^i \in \mathbb{R}$ of a model $\theta_n = [\theta_n^i]_{i=1}^d$ by

$$\hat{\theta}_n^i \leftarrow \begin{cases} \theta_{\pi_i(n)}^i & \text{with probability } p, \\ \theta_n^i & \text{otherwise,} \end{cases} \quad (3)$$

where π_i denotes a random permutation of the indices $\{1, \dots, N\}$, chosen uniformly at each iteration for each parameter index $i \in \{1, \dots, d\}$, and independently from the Bernoulli variable of Eq. (3). Notably, this parameter shuffling reduces in expectation to

$$\mathbb{E}[\hat{\theta}_n] = (1 - p)\theta_n + p\bar{\theta}. \quad (4)$$

Thus, WASH aligns, in expectation, with the EMA of Eq. (1) for $p = (1 - \alpha)$. The expected number of parameters communicated by each model at each step is thus $p \times d$ while for PAPA, each model communicating all of its parameters every T steps, this amounts to $\frac{d}{T}$. Thus, $p \ll \frac{1}{T}$ results in a significantly reduced communication overhead favorable to WASH. However, the model diversity is higher, as WASH preserves the consensus distance, as shown by

$$\sum_n \|\hat{\theta}_n - \bar{\theta}\|^2 = \sum_n \sum_i (\hat{\theta}_n^i - \bar{\theta}^i)^2 = \sum_i \sum_n (\theta_n^i - \bar{\theta}^i)^2 = \sum_n \|\theta_n - \bar{\theta}\|^2. \quad (5)$$

Table 1: **Communication volume and inference costs** of four training techniques. The baseline Ensemble is trained separately but requires a linearly increasing inference cost. In our experiments, we fix the base probability of WASH and WASH+Opt to be equal to 0.001 or 0.05 when training on CIFAR-10/100 or ImageNet, resulting in a reduction of communication volume over PAPA.

Technique	Communication volume		Inference cost
	CIFAR-10/100	ImageNet	
Ensemble	0	0	N
PAPA	1	1	1
WASH	$1/200$	$1/4$	1
WASH+Opt	$1/100$	$1/2$	1

Layer-wise adaptation via WASH. Recognizing that different network layers may require varying levels of adaptation due to their roles and dynamics, we introduce a layer-specific probability adjustment. Assuming L layers in the network, we set for each layer l (where $0 \leq l < L$)

$$p_l = p \left(1 - \frac{l}{L-1} \right), \quad (6)$$

where p is a base probability. In other words, the parameters of the first layer have a shuffling probability of p , while the final layer’s parameters are never shuffled. This adaptation ensures that deeper layers, which are typically slower to train and more sensitive to the input features, undergo fewer permutations than the more generalizable early layers. This strategy not only preserves the specificity required by the initial layers but also further halves the overall communication overhead.

Full procedure. Alg. 1 presents the training of a population of N models using WASH. Starting from the same initialization, our training procedure alternates between local gradient computations and shuffling communications. At inference, we simply average the weights of the models, obtaining a single model with parameters $\bar{\theta}$. Note that techniques like REPAIR (21) or activation alignment (1) could be incorporated to improve the alignment of the models, but we found them to be unnecessary to obtain high accuracy and kept our evaluation framework minimal for the sake of simplicity.

4 Experiments

Training methods. We showcase the capacities of WASH for training a population of neural networks on standard image classification tasks. As a Baseline, we consider a population trained separately, with each model working on a different dataset order and different data augmentations and regularization (if they are used). This is the same baseline as (20), only starting from the same initialization, but we found that this change brought no significant impact on performance. We also compare WASH to PAPA (20) on the same tasks (with PAPA however using models with a different initialization), to show our improvement despite requiring a fraction of the communication cost. We do not provide comparisons to DART (19) or the variants of PAPA as their performances are generally inferior (20). We also propose a variant of WASH named WASH+Opt that also permutes the optimizer state associated with the parameter shuffled (in our case, the momentum of SGD), doubling the volume of communications. We do not permute or recompute the running statistics of BatchNorm layers for simplicity.

Communication cost. Training with PAPA requires computing an all-reduce operation on all of the models’ parameters every $T = 10$ training steps. In comparison, WASH requires, in expectation, a shuffling of $p/2$ of the parameters of the models at every training step. Thus, by keeping a base probability $p \leq 0.2$, WASH results in a more communication-efficient training. In practice, p will be equal in our experiments to 0.001 or 0.05, ensuring a communication volume reduction of 200 or 4.

Evaluation strategy. After training, the population of models obtained can be evaluated in three separate ways. As a baseline, the performance of the population can be evaluated as an Ensemble, averaging the predictions of the models. The parameters of the models can be averaged to obtain a

Table 2: **Ensemble and Averaged Model accuracy for a heterogeneous population of models; trained with varying data augmentations and regularizations.** We compare models trained separately (Baseline), with PAPA, or our method WASH and its variant WASH+Opt. We also report the GreedySoup accuracy for the Baseline models. The best Ensemble (black) and Averaged (blue) accuracy are reported in bold. Except on CIFAR-10, WASH and in particular WASH+Opt provide the best performance for the final Averaged Model, with performances comparable to the Ensemble of models for a fraction of the inference cost.

Method Config	#N	Baseline (trained separately)			PAPA		WASH (ours)		WASH+Opt (ours)	
		Ensemble	Averaged	GreedySoup	Ensemble	Averaged	Ensemble	Averaged	Ensemble	Averaged
CIFAR-10										
VGG-16	3	95.98±.42	10.00±.00	95.26±.05	96.12±.34	96.13±.24	95.89±.23	95.97±.24	95.91±.36	95.85±.27
	5	96.28±.40	10.00±.00	95.42±.10	96.24±.17	96.21±.13	96.15±.10	96.20±.10	96.00±.21	96.04±.14
	10	96.47±.07	10.00±.00	95.39±.24	96.32±.13	96.31±.13	96.27±.10	96.18±.13	96.14±.08	96.20±.05
ResNet18	3	97.15±.28	10.17±.29	96.62±.38	97.33±.05	97.24±.05	97.21±.19	97.19±.17	97.22±.07	97.25±.14
	5	97.33±.08	10.09±.16	96.61±.03	97.35±.12	97.31±.06	97.21±.10	97.25±.12	97.18±.09	97.16±.07
	10	97.59±.01	9.26±1.28	96.79±.14	97.39±.13	97.34±.06	97.30±.10	97.28±.04	97.20±.13	97.16±.13
CIFAR-100										
VGG-16	3	80.36±.15	1.00±.00	77.92±.22	78.89±.10	78.77±.16	79.10±.88	79.05±.68	79.15±.61	79.15±.41
	5	81.32±.56	1.00±.00	77.81±.25	79.51±.38	79.24±.43	79.65±.27	79.39±.21	79.75±.21	79.71±.20
	10	82.24±.15	1.00±.00	77.83±.65	79.95±.11	79.64±.13	80.05±.18	79.70±.25	80.03±.11	79.76±.13
ResNet18	3	82.84±.48	1.00±.01	80.06±1.5	81.58±.12	81.53±.13	81.91±.34	81.90±.36	81.99±.06	82.08±.09
	5	83.72±.49	1.00±.00	80.72±.52	82.09±.30	82.01±.34	82.16±.42	81.97±.28	82.35±.17	82.17±.15
	10	84.18±.20	1.00±.00	80.61±.43	82.32±.09	82.15±.14	82.43±.32	82.31±.38	82.42±.31	82.18±.22
ImageNet										
ResNet50	3	76.16±.28	0.10±.00	74.15±.11	75.62±.15	*	74.39±.14	74.34±.18	74.30±.22	74.18±.26
	5	76.68±.06	0.10±.00	74.47±.06	75.80±.21	*	74.63±.11	74.59±.07	74.44±.21	74.39±.21

single model, which we refer to as Averaged. This is equivalent to UniformSoup in (51) or AvgSoup in (20) for example. More elaborate averaging methods have been proposed, such as GreedySoup (51), which averages an increasing number of models (in order of validation accuracy) until the averaging does not improve accuracy. We report the accuracy of the Ensemble and Averaged model for all training techniques, as well as the GreedySoup accuracy of the Baseline. Like (20), we find that the GreedySoup accuracy corresponds to the accuracy of a single model for the Baseline and that the Averaged model accuracy outperforms the GreedySoup model for the other techniques, and thus chose not to report it. We summarize in Tab. 1 the communication volume and inference cost necessary for training a separate Ensemble of models, or training with PAPA, WASH, or WASH+Opt.

4.1 Main experiments

Experimental setup. We showcase the performance of WASH for training neural networks on image classification tasks on the CIFAR-10, CIFAR-100 (24), and ImageNet (6) datasets. We use the same training framework as (20) for a fair comparison. We train a population of N models for $N \in \{3, 5, 10\}$, on the ResNet-18, 50 and VGG-16 architectures. 2% of the training data is kept as validation for computing the GreedySoup. Like (20), we consider one framework with heterogeneous models, learning with different data augmentations and regularizations, and one homogeneous setting with no data augmentations, where the only difference between the models’ training is the dataset shuffling. Details are presented in the Appendix. Models are trained with SGD with momentum, a weight decay of 10^{-4} , and a cosine annealing scheduler with starting and minimum learning rates 0.1 and 10^{-4} . For CIFAR-10/100, we train over 300 epochs with a batch size of 64, and 90 epochs with a batch size of 256 for ImageNet. For WASH and WASH-Opt we initialize the models with the same parameters and choose p with cross-validation to be equal to 0.001 or 0.05 when training on CIFAR-10/100 or ImageNet. We do not require any alignment technique such as REPAIR.

Main results. Tab. 2 and Tab. 3 correspond respectively to the heterogeneous and homogeneous settings. We report the test accuracies as the average of 3 runs for the Ensemble of models, the Averaged model, and the GreedySoup for the Baseline (equivalent to the best model). Consistent with the findings of (20), we find that networks trained separately have a high Ensemble accuracy, but perform as random when averaged. On CIFAR-10/100, methods like PAPA and WASH result in lower Ensemble accuracy but almost no difference between the Ensemble and Averaged accuracies. In general, WASH and WASH+Opt outperform PAPA, despite requiring a lower communication volume.

Table 3: **Ensemble and Averaged Model accuracy for a homogeneous population of models.** We compare models trained separately (Baseline), with PAPA, or our methods WASH and WASH+Opt. The best Ensemble (black) and Averaged (blue) accuracy are reported in bold. We observe the same results in this setting, with WASH in particular reaching close to the Ensemble performance.

Method Config	#N	Baseline (trained separately)			PAPA		WASH (ours)		WASH+Opt (ours)	
		Ensemble	Averaged	GreedySoup	Ensemble	Averaged	Ensemble	Averaged	Ensemble	Averaged
CIFAR-10										
VGG-16	3	94.93±.06	10.00±.00	93.60±.41	94.38±.14	94.34±.18	94.41±.23	94.58±.17	94.45±.05	94.47±.02
	5	95.29±.05	10.00±.00	93.82±.30	94.55±.12	94.58±.12	94.72±.08	94.70±.17	94.63±.11	94.68±.14
	10	95.23±.06	10.00±.00	93.82±.06	94.79±.18	94.78±.20	94.66±.03	94.54±.07	94.71±.07	94.61±.13
ResNet18	3	96.14±.10	10.00±.00	95.42±.27	95.89±.04	95.89±.06	95.77±.12	95.77±.17	95.85±.04	95.87±.10
	5	96.19±.16	10.00±.00	95.31±.09	95.99±.08	95.99±.08	95.96±.08	95.98±.05	95.94±.12	95.98±.12
	10	96.34±.02	10.00±.00	95.26±.11	96.10±.25	96.11±.24	96.08±.07	96.12±.09	96.07±.07	96.08±.14
CIFAR-100										
VGG-16	3	77.63±.24	1.00±.00	73.76±.35	75.10±.11	75.09±.16	76.30±.37	76.04±.58	76.04±.03	75.96±.18
	5	78.52±.10	1.00±.00	73.76±.18	75.56±.16	75.55±.14	76.63±.27	76.48±.23	76.64±.15	76.13±.18
	10	79.26±.06	1.00±.00	73.99±.26	76.24±.44	76.26±.43	77.06±.12	76.43±.18	76.72±.15	75.94±.26
ResNet18	3	79.54±.17	1.00±.00	76.84±.54	77.83±.26	77.86±.30	78.90±.17	78.76±.25	78.66±.08	78.56±.21
	5	80.11±.23	1.00±.00	76.83±.45	77.94±.16	77.92±.19	79.24±.32	79.09±.43	79.32±.19	79.19±.15
	10	80.55±.13	1.00±.00	76.80±.41	78.40±.15	78.44±.22	79.65±.17	79.43±.16	79.34±.34	79.19±.45

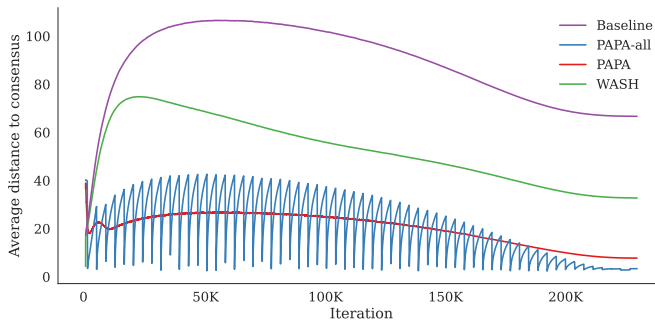


Figure 2: **Average distance to the consensus (i.e. the averaged model)** during training for a heterogeneous population of 5 models trained on CIFAR-100, either separately, with PAPA, PAPA-all, or our method WASH. Starting at consensus, models initially diverge from each other before converging back during convergence, mainly due to weight decay. Models trained with WASH have a smaller distance to consensus than ones trained separately; allowing them to be averaged with no performance loss. By training with PAPA-all (i.e. averaging to a single model every few epochs), models are not able to reach the same diversity as WASH between these averaging steps. Finally, the EMA of PAPA has a strong pulling effect towards consensus, resulting in a similar distance as PAPA-all. The jitter in the curve is due to the immediate distance reduction caused by the EMA steps.

On ImageNet, our parallelization procedure resulted in a slightly lower Baseline accuracy and we were not able to reproduce PAPA’s baseline due to our distributed constraint. The WASH Averaged model reaches a high accuracy, like previously. Both of our methods reduce the gap with accuracies of the baseline Ensemble, indicating that WASH hinders less the diversity of the population of models while maintaining weight averagability. However, a gap still remains, which may be inherent to models being in the same basin. WASH and WASH+Opt have very similar results, with the simpler WASH being better on the homogeneous case and WASH+Opt being better on the heterogeneous case.

4.2 Why do shuffling parameters help?

In this section, we propose to explain the improvement provided by our parameter shuffling over previous mechanisms such as BTM, DART, or PAPA, that focus on the averaging of parameters. We first show that models trained with WASH have a smaller distance to consensus than models trained separately. Then, we argue that despite this, WASH is a weak perturbation on the training of the models and that it incites diversity in the models.

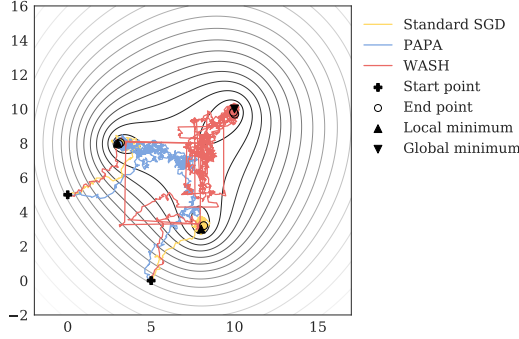


Figure 3: **2D optimization example.** We train 2 points with SGD on a simple loss function with 2 local and 1 global minima (upwards and downwards triangle). The two models are trained from two different starting points (plus signs). If the points are trained separately (yellow), they converge to their closest local minimum (yellow circles). By training with PAPA (blue), the points reach a consensus but then converge to one of the local minima (blue circles). By training with WASH (red), the shuffling (seen by the horizontal and vertical lines in the trajectory) allows more diversity in the optimization path, and the points both reach the global minimum (red circles).

Reducing distance to consensus. To better analyze the diversity of the models trained with WASH, we propose to report the distance of the models to the consensus (the averaged model) during training, as a proxy of the diversity metric. (18; 52) showed that the difference between the Ensemble and the Averaged models depends on the distance between models. We present in Fig. 2 the average distance of the models to the consensus, for models trained separately, with PAPA, PAPA-all, or with WASH. PAPA-all is a variant of PAPA functionally identical to DART. The idea is to average the weights every few epochs before letting the models diversify again. We observe that WASH results in a consistently lower distance to consensus than the baseline, despite explicitly leaving the distance to consensus unchanged during the shuffling step, and shuffling only a small number of parameters. Thus, the smaller distance at the end of the training explains why the averaging of the parameters does not cause a decrease in performance. By comparison, PAPA-all (i.e. DART) results in alternating phases where models diversify before being averaged, and we observe that the models are not able to reach the diversity of WASH. Similarly, the EMA of PAPA has a strong pulling effect and results in average in a similar diversity as PAPA-all. Thus, we find that models trained with WASH have a higher diversity than models trained with PAPA or PAPA-all, while being close enough that averaging them does not cause a loss of performance. More generally, we show in Fig. 6 of the Appendix that various interpolations of models trained by WASH result in a similar performance, showcasing that they all reside in the same loss basin.

Encouraging diversity. WASH can be viewed as a weak perturbation on the models: the parameter shuffling affects more weakly the models than parameter averaging or the EMA of PAPA as only a few parameters are affected at a time, and the consensus distance is unaffected. Furthermore, the shuffling of parameters increases the diversity of the trajectories seen by the models. We showcase this with a toy example, by training jointly two points with SGD on a 2D loss function with 2 local minima and 1 global minimum, either training them separately, with PAPA, or with WASH. We represent the trajectories corresponding to each method in Fig. 3. Training separately the two points makes them converge into a separate local minimum (i.e. a different basin). Training with PAPA allows the two points to reach a consensus, however they converge together into a local minimum. In contrast, by training with WASH, we show that both points reach the global minimum, as the shuffling allows for a greater diversity of points to optimize with. We provide more details in the Appendix.

4.3 Ablations

We present in this section ablations to better understand the effect of the parameter shuffling, varying the layer-wise probability adaptation, the base probability value, and the shuffling period.

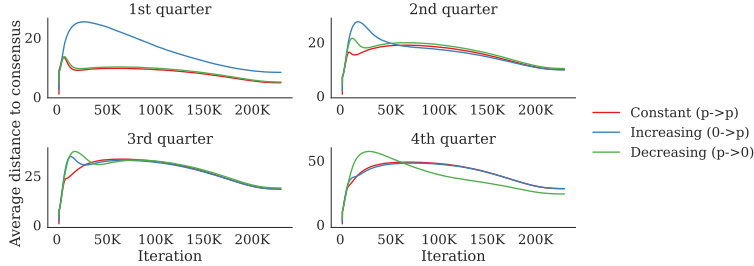
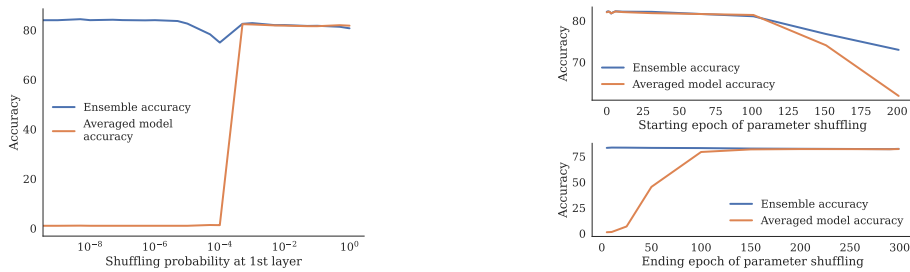


Figure 4: **Average distance to the consensus for different layer-wise adaptations of WASH**, for different slices of the model’s parameters. Keeping the probability constant across layers ensures the lowest distance to consensus for the first quarters. Surprisingly, in the last quarter of parameters, despite initially starting with a higher distance to consensus, the ‘decreasing probability’ shows a lower distance to consensus later in training; despite shuffling being less frequent than the other schedules. The ‘increasing probability’ showcases how early layers are sensible to shuffling.



(a) **Ensemble and Averaged accuracy for varying base probability values.** We observe a phase transition as the base probability augments between a phase where permuting does not improve the averaged model accuracy and a phase where the ensemble accuracy is equal to the averaged model accuracy. Between the phases, the ensemble accuracy diminishes.

(b) **Ensemble and Averaged accuracy depending on the starting or ending epoch of the shuffling.** The parameter shuffling is beneficial both at the start and the end of training. Note that ending early, at epoch 150 out of 300, is less impactful on performances than starting permuting at epoch 150, showing that WASH is more important early in training.

Figure 5: **Ablations of WASH**

Layer-wise adaptation variations. For WASH, we found that a probability decrease with depth provided the best results. We show in Tab. 4 of the Appendix the performances for alternatives, with probability either staying constant or increasing with depth. We find lower performances for both alternatives. In Fig. 4, we report the models’ distances to consensus for all three schedules. More precisely, we provide the distances for different slices of the models’ parameters, indicating the effect of the shuffling depending on the depth. Shuffling equally all layers results in the lowest distance to the consensus as predicted, except in the last quarter of parameters. Here, surprisingly, our base ‘decreasing’ schedule shows a lower distance to consensus despite shuffling less frequently. We also observe a particularly strong effect of the shuffling for early layers, as the distance is more emphasized in the first quarter between the ‘increasing’ curve and the others.

Base probability variation. We present in Fig. 5a the Ensemble and Averaged for different values of p , the base shuffling probability of the first layer. Rather than a smooth increase of the Averaged model accuracy, we observe a phase transition between a phase where the Averaged model accuracy is not improved by the shuffling and a sudden increase in the accuracy where it reaches the accuracy of the Ensemble. Just before the transition, the Ensemble accuracy is decreased, before increasing again back to its previous performance. The accuracy decreases only slightly even when increasing the shuffling probability to 1, indicating the resilience of the models even to heavy shuffling.

Shuffling is beneficial at every step. Finally, we propose to show the impact of the parameter shuffling at different steps of the training by varying the epoch at which the shuffling either starts or stops. In Fig. 5b, we show that there is no improvement by having a warmup or slowdown period in

parameter shuffling, indicating that all phases of the training are improved by WASH. Furthermore, stopping parameter shuffling early results in a much smaller loss of Averaged accuracy compared to starting late. In other words, shuffling at the start of training before models converge is more impactful as models may still reside in different loss basins.

Conclusion

We proposed a novel distributed training method, WASH, aimed at training a population of models in parallel. These models are averaged at the end of training to obtain a highly performing model with accuracies close to the ensemble accuracy for a fraction of the inference cost. Our method requires a fraction of the communication cost of similarly performing techniques while obtaining state-of-the-art results for our weight-averaged models. We show that our novel parameter shuffling does not explicitly reduce the distance between models while increasing the diversity of the optimization paths seen by the population. Still, we observe that the distance between our models is smaller than training them separately, allowing them to be averaged at the end of training.

Acknowledgements

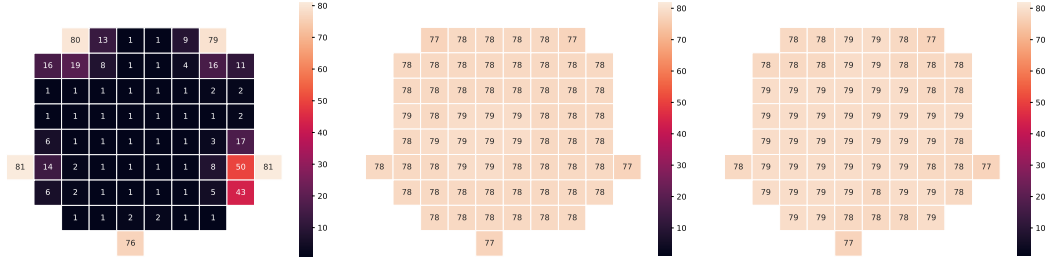
This work was supported by Project ANR-21-CE23-0030 ADONIS, EMERG-ADONIS from Alliance SU, and Sorbonne Center for Artificial Intelligence (SCAI) of Sorbonne University (IDEX SUPER 11-IDEX-0004). This work was granted access to the AI resources of IDRIS under the allocations 2023-A0151014526 made by GENCI. Eugene Belilovsky also acknowledges the FRQNT New Scholar grant.

References

- [1] S. K. Ainsworth, J. Hayase, and S. Srinivasa. Git re-basin: Merging models modulo permutation symmetries, 2022.
- [2] G. Benton, W. Maddox, S. Lotfi, and A. G. G. Wilson. Loss surface simplexes for mode connecting volumes and fast ensembling. In *International Conference on Machine Learning*, pages 769–779. PMLR, 2021.
- [3] L. Bottou, F. E. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. *SIAM review*, 60(2):223–311, 2018.
- [4] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, Aug 1996.
- [5] M. Chen, M. Jiang, Q. Dou, Z. Wang, and X. Li. Fedsoup: Improving generalization and personalization in federated learning via selective model interpolation, 2023.
- [6] 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*, pages 248–255. Ieee, 2009.
- [7] T. G. Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.
- [8] C. Dun, M. Hipolito, C. Jermaine, D. Dimitriadis, and A. Kyrillidis. Efficient and light-weight federated learning via asynchronous distributed dropout, 2022.
- [9] C. Dun, C. R. Wolfe, C. M. Jermaine, and A. Kyrillidis. Resist: Layer-wise decomposition of resnets for distributed training. In J. Cussens and K. Zhang, editors, *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pages 610–620. PMLR, 01–05 Aug 2022.
- [10] Y. Esfandiari, S. Y. Tan, Z. Jiang, A. Balu, E. Herron, C. Hegde, and S. Sarkar. Cross-gradient aggregation for decentralized learning from non-iid data. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 3036–3046. PMLR, 18–24 Jul 2021.

- [11] S. Fort, H. Hu, and B. Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019.
- [12] L. Fournier and E. Oyallon. Cyclic data parallelism for efficient parallelism of deep neural networks, 2024.
- [13] J. Frankle, G. K. Dziugaite, D. M. Roy, and M. Carbin. Linear mode connectivity and the lottery ticket hypothesis, 2020.
- [14] Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning, 2015.
- [15] T. Garipov, P. Izmailov, D. Podoprikin, D. P. Vetrov, and A. G. Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [16] R. Gontijo-Lopes, Y. Dauphin, and E. D. Cubuk. No one representation to rule them all: Overlapping features of training methods. *arXiv preprint arXiv:2110.12899*, 2021.
- [17] S. Horoi, A. M. O. Camacho, E. Belilovsky, and G. Wolf. Harmony in diversity: Merging neural networks with canonical correlation analysis. In *International Conference on Machine Learning (ICML)*, 2024.
- [18] P. Izmailov, D. Podoprikin, T. Garipov, D. Vetrov, and A. G. Wilson. Averaging weights leads to wider optima and better generalization, 2019.
- [19] S. Jain, S. Addepalli, P. Sahu, P. Dey, and R. V. Babu. Dart: Diversify-aggregate-repeat training improves generalization of neural networks, 2023.
- [20] A. Jolicoeur-Martineau, E. Gervais, K. Fatras, Y. Zhang, and S. Lacoste-Julien. Population parameter averaging (papa), 2023.
- [21] K. Jordan, H. Sedghi, O. Saukh, R. Entezari, and B. Neyshabur. Repair: Renormalizing permuted activations for interpolation repair, 2023.
- [22] J. Konecný, H. B. McMahan, D. Ramage, and P. Richtárik. Federated optimization: Distributed machine learning for on-device intelligence. *ArXiv*, abs/1610.02527, 2016.
- [23] L. Kong, T. Lin, A. Koloskova, M. Jaggi, and S. Stich. Consensus control for decentralized deep learning. In *International Conference on Machine Learning*, pages 5686–5696. PMLR, 2021.
- [24] A. Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- [25] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- [26] C. Lakshminarayanan and C. Szepesvari. Linear stochastic approximation: How far does constant step-size and iterate averaging go? In *International Conference on Artificial Intelligence and Statistics*, pages 1347–1355. PMLR, 2018.
- [27] M. Li, S. Gururangan, T. Dettmers, M. Lewis, T. Althoff, N. A. Smith, and L. Zettlemoyer. Branch-train-merge: Embarrassingly parallel training of expert language models, 2022.
- [28] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith. Federated optimization for heterogeneous networks. In *ICML Workshop on Adaptive & Multitask Learning: Algorithms & Systems*, 2019.
- [29] T. Lin, S. U. Stich, K. K. Patel, and M. Jaggi. Don’t use large mini-batches, use local sgd. In *International Conference on Learning Representations*, 2020.
- [30] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In A. Singh and J. Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282. PMLR, 20–22 Apr 2017.

- [31] G. Menghani. Efficient deep learning: A survey on making deep learning models smaller, faster, and better. *ACM Computing Surveys*, 55(12):1–37, 2023.
- [32] A. Nabli, E. Belilovsky, and E. Oyallon. A2cid2: Accelerating asynchronous communication in decentralized deep learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [33] B. Neyshabur, H. Sedghi, and C. Zhang. What is being transferred in transfer learning? In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 512–523. Curran Associates, Inc., 2020.
- [34] J. J. G. Ortiz, J. Frankle, M. Rabbat, A. Morcos, and N. Ballas. Trade-offs of local sgd at scale: An empirical study, 2021.
- [35] S. Pati, S. Aga, M. Islam, N. Jayasena, and M. D. Sinclair. Computation vs. communication scaling for future transformers on future hardware, 2023.
- [36] F. A. G. Peña, H. R. Medeiros, T. Dubail, M. Aminbeidokhti, E. Granger, and M. Pedersoli. Re-basin via implicit sinkhorn differentiation, 2022.
- [37] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.
- [38] A. Ramé, M. Kirchmeyer, T. Rahier, A. Rakotomamonjy, P. Gallinari, and M. Cord. Diverse weight averaging for out-of-distribution generalization, 2022.
- [39] W. Shi, Q. Ling, G. Wu, and W. Yin. Extra: An exact first-order algorithm for decentralized consensus optimization, 2014.
- [40] E. Shulgin and P. Richtárik. Towards a better theoretical understanding of independent subnetwork training, 2023.
- [41] S. P. Singh and M. Jaggi. Model fusion via optimal transport, 2023.
- [42] A. Spiridonoff, A. Olshevsky, and I. C. Paschalidis. Communication-efficient sgd: From local sgd to one-shot averaging, 2021.
- [43] S. U. Stich. Local sgd converges fast and communicates little, 2019.
- [44] H. Tang, X. Lian, M. Yan, C. Zhang, and J. Liu. D²: Decentralized training over decentralized data, 2018.
- [45] G. Wang, H. Qin, S. A. Jacobs, C. Holmes, S. Rajbhandari, O. Ruwase, F. Yan, L. Yang, and Y. He. Zero++: Extremely efficient collective communication for giant model training, 2023.
- [46] H. Wang, M. Yurochkin, Y. Sun, D. Papailiopoulos, and Y. Khazaeni. Federated learning with matched averaging, 2020.
- [47] J. Wang, V. Tantia, N. Ballas, and M. Rabbat. Slowmo: Improving communication-efficient distributed sgd with slow momentum. In *International Conference on Learning Representations*, 2020.
- [48] Z. Wang, J. Zhang, T.-H. Chang, J. Li, and Z.-Q. Luo. Distributed stochastic consensus optimization with momentum for nonconvex nonsmooth problems. *IEEE Transactions on Signal Processing*, 69:4486–4501, 2021.
- [49] D. Wen, K.-J. Jeon, and K. Huang. Federated dropout – a simple approach for enabling federated learning on resource constrained devices, 2021.
- [50] M. Wortsman, M. Horton, C. Guestrin, A. Farhadi, and M. Rastegari. Learning neural network subspaces, 2021.



(a) Accuracy heatmap of the Baseline. The interpolated models' performance is equal to random ones. (b) Accuracy heatmap of WASH. The accuracy is very similar for various interpolations. (c) Accuracy heatmap of WASH+Opt. The results are similar to WASH.

Figure 6: Accuracy heatmap for different weight interpolations, for models trained separately, with WASH or WASH+Opt.

[51] M. Wortsman, G. Ilharco, S. Y. Gadre, R. Roelofs, R. Gontijo-Lopes, A. S. Morcos, H. Namkoong, A. Farhadi, Y. Carmon, S. Kornblith, and L. Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 23965–23998. PMLR, 17–23 Jul 2022.

[52] M. Wortsman, G. Ilharco, J. W. Kim, M. Li, S. Kornblith, R. Roelofs, R. Gontijo-Lopes, H. Hajishirzi, A. Farhadi, H. Namkoong, and L. Schmidt. Robust fine-tuning of zero-shot models, 2022.

[53] M. Yurochkin, M. Agarwal, S. Ghosh, K. Greenewald, N. Hoang, and Y. Khazaeni. Bayesian nonparametric federated learning of neural networks. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 7252–7261. PMLR, 09–15 Jun 2019.

[54] S. Zhang, A. Choromanska, and Y. LeCun. Deep learning with elastic averaging sgd. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, page 685–693, Cambridge, MA, USA, 2015. MIT Press.

5 Appendix

2D optimization example The loss function we consider is a heavily simplified version of the Ackley function. With a minima in (x_m, y_m) defined by

$$g(x, y, x_m, y_m, \lambda) = \exp(-\lambda \sqrt{0.5((x - x_m)^2 + (y - y_m)^2)}), \quad (7)$$

the function we consider in our example is

$$f(x, y) = -10g(x, y, 10, 10, 0.1) - 5g(x, y, 8, 3, 0.3) - 5g(x, y, 3, 8, 0.3). \quad (8)$$

This function has a 2 local minima in $(3, 8)$ and $(8, 3)$ and a global minimum in $(10, 10)$. In all three cases, the starting points are $(0, 5)$ and $(5, 0)$. We compute SGD by first computing the exact gradient of the function and then adding Gaussian noise to the gradient. The learning rate is 0.1 and we optimize for 1000 steps. For PAPA, we consider $\alpha = 0.99$. For WASH, the shuffling probability is equal for both coordinates and equal to 0.01.

Interpolation heatmap Here, we propose to display a heatmap showing the accuracy of more varied interpolations between 5 models trained separately, with WASH, or WASH+Opt. We observe how WASH and WASH+Opt trained models converge to the same loss basin, and that a large number of possible interpolations result in a high accuracy. The heatmaps are presented in Fig. 6.

Layer-wise adaptation variants performance We showcase in Tab.4 the performance of the three variants of layer-wise adaptations of WASH.

Table 4: **Test accuracies of WASH with variants of the shuffling probability per depth.** Trained with a population of 5 models on CIFAR-100 with a ResNet-18. The results show that permuting the first layers is more important than the later layers. Still, a constant probability across layers does not decrease WASH’s performance much.

Proba. at layer		L-1	Technique			Best model	Worst model
0	to		Ensemble	Averaged	GreedySoup		
10^{-3}	\searrow	0	82.22±.38	82.15±.22	81.94± 0.25	80.89±.03	78.80±.77
10^{-3}	\rightarrow	10^{-3}	82.04±.19	81.94±.15	81.69±.23	80.60±.16	78.67±.89
0	\nearrow	10^{-3}	81.75±.35	81.37±.10	81.14±.20	80.08±.40	78.55±.70

Augmentations and regularization used We follow the same data augmentations and regularizations used in (20) for a fair comparison. We use Mixup (random draw from {0, 0.5, 1.0} for CIFAR-10/100 or from {0, 0.2} for ImageNet), Label smoothing (random draw from {0, 0.05, 0.1} for CIFAR-10/100 or from {0, 0.1} for ImageNet), CutMix (random draw from {0, 0.5, 1.0} for CIFAR-10/100 or from {0, 1.0} for ImageNet) and Random Erasing (random draw from {0, 0.15, 0.35} for CIFAR-10/100 or from {0, 0.35} for ImageNet).

For our experiments, we required a single A100 GPU for up to 14 hours to train up to a population of 10 models, and up to 40 hours for a population of 20 models. Similarly, we required 16 A100 GPUs to train in parallel a population of 5 models on ImageNet.