



HAL
open science

Towards efficient feature sharing in MIMO architectures

Rémy Sun, Alexandre Ramé, Clément Masson, Nicolas Thome, Matthieu Cord

► To cite this version:

Rémy Sun, Alexandre Ramé, Clément Masson, Nicolas Thome, Matthieu Cord. Towards efficient feature sharing in MIMO architectures. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), Jun 2022, New Orleans, United States. pp.2696-2700, 10.1109/CVPRW56347.2022.00303 . hal-03951737

HAL Id: hal-03951737

<https://hal.science/hal-03951737>

Submitted on 23 Jan 2023

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.

Towards efficient feature sharing in MIMO architectures

Rémy Sun^{1,3,*} Alexandre Ramé¹ Clément Masson³ Nicolas Thome² Matthieu Cord^{1,4}

¹MLIA, ISIR, Sorbonne Université

²Vertigo, CEDRIC, Conservatoire National des Arts et Métiers

³Thales Land and Air Systems, Elancourt, France

⁴Valeo.ai

Abstract


Multi-input multi-output architectures propose to train multiple subnetworks within one base network and then average the subnetwork predictions to benefit from ensembling for free. Despite some relative success, these architectures are wasteful in their use of parameters. Indeed, we highlight in this paper that the learned subnetwork fail to share even generic features which limits their applicability on smaller mobile and AR/VR devices. We posit this behavior stems from an ill-posed part of the multi-input multi-output framework. To solve this issue, we propose a novel unmixing step in MIMO architectures that allows subnetworks to properly share features. Preliminary experiments on CIFAR 100 show our adjustments allow feature sharing and improve model performance for small architectures.

1. Introduction


The last decade has seen large deep architectures take over many machine learning domains [3, 5] previously solved by more traditional algorithms. As such, deep learning has become ever more present in practical applications. It is therefore now especially important to find ways to maximize model performance [12, 16].

A well known way to obtain better performances given already trained models is to ensemble the predictions given by multiple models [6]. Indeed, predictions from independently trained models have been shown to complement each other such that the aggregated predictions largely outperform the individual model predictions on a test set.

Unfortunately, this increase in performance comes at the cost of dramatically increased overhead : to ensemble multiple models, one must have access to multiple trained models [1, 6]. This is an untenable cost in many real world applications where networks must fit on tiny embedded chips in



(a) Previous MIMO models train subnetworks that fail to share features.



(b) MixShare introduces mechanisms that allow feature sharing.

Figure 1. The lack of feature sharing in MIMO models is a missed opportunity for parameter efficiency on smaller networks. Mixshare introduces unmixing to allow feature sharing among subnetworks in multi-input multi-output models, improving them on smaller networks previous methods struggle on.

mobile and AR/VR devices. Significant emphasis has therefore been put in the ensembling literature on finding ways to minimize the inherent cost of ensembling, typically through some degree of parameter sharing between models [7, 13].

Multi-input multi-output (MIMO) strategies [2, 11] provide an interesting solution to this conundrum by ensembling for virtually free. Through their multiple inputs and outputs, MIMO frameworks train independent subnetworks within a base network. Thanks to the sparse nature of large neural networks [8], the resulting subnetworks yield strong and diverse predictions that can be ensembled. As shown

*Corresponding author

mailto:remy.sun@isir.upmc.fr

on Fig. 1a with $M = 2$, the M inputs are embedded by M subnetworks with no structural differences. Thus, we have M (inputs, labels) pairs in training: $\{(x_i, y_i)\}_{0 \leq i < M}$. More precisely, M **images are fed to the network at once**. The M inputs are encoded by M distinct convolutional layers $\{c_i\}_{0 \leq i < M}$ into a shared latent space before being aggregated (either through summation [2] or more complex mixing operations [11]). This representation is then processed by the core network into a single feature vector, which is classified by M dense layers $\{d_i\}_{0 \leq i < M}$. Diverse subnetworks appear as d_i learns to classify y_i from input x_i . At inference, we can ensemble M predictions by feeding the same image M times to the model.

MixMo [11] has however recently highlighted significant limitations of such architectures: multi-input multi-output architectures require large base models and struggle to fit more than 2 subnetworks. Indeed, [11] shows a significant drop in performance on CIFAR 100 when going from 2 subnetworks to 4 subnetworks.

This scaling issue is explained by analyzing the features inside the network, as we show at the beginning of this paper by extending [2, 11]’s study of subnetwork behavior. Our analysis shows that the aforementioned scaling issues stem from how subnetworks share no features in the base network (see Fig. 1a): each channel or feature is almost exclusively used by one subnetwork. As such, we can explain the scaling issue since each additional subnetwork significantly reduces the effective size of the individual subnetworks. Beyond causing issues on smaller architectures or harder datasets, this leads to very wasteful use of network parameters. This is especially unfortunate as the subnetworks could at the very least share generic features in the first layers. We see this as a missed opportunity, one that can significantly improve multi-input multi-output models’ applicability to real world settings like mobile devices.

In this paper, we first carefully study in Sec. 2 how subnetworks use the base network’s features. After showing the lack of feature sharing, we discuss the impact of this on parameter efficiency and model performance. Secondly, we propose Mixshare in Sec. 3 to address the issues preventing feature sharing (see Fig. 1b). In particular, we introduce a novel unmixing mechanism (Sec. 3.1) to allow sharing and discuss in Sec. 3.2 how proper network initialization is necessary to improve model performance.

2. MIMO Subnetworks do not share features

In this section, we strive to pinpoint the cause of multi-input multi-output architectures’ scaling issues. To this end, we consider the following question: how do subnetworks behave in multi-input multi-output architectures?

Following [11], we check how the inputs are organized in the C feature maps of the mixing space by considering the L_1 norm of the C encoder kernels for each subnetwork (see

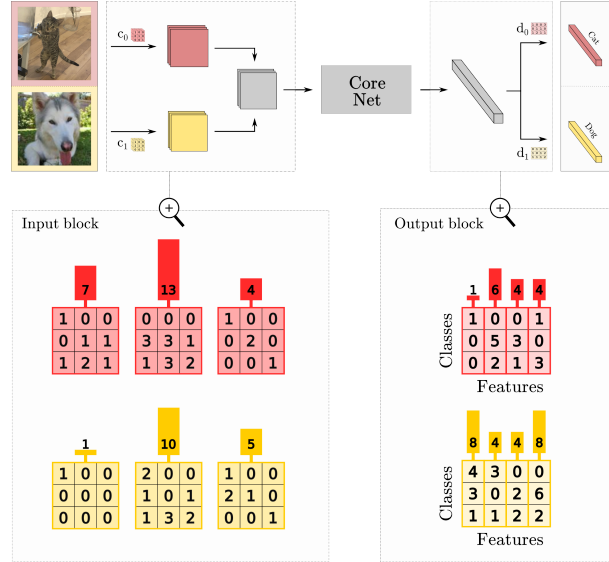


Figure 2. We study the influence of features in the input and output block on the subnetworks (L_1 norm figured by bars). For the input block, we consider the L_1 norm of the feature kernels for the relevant encoders. For the output block, we look at the L_1 norm of columns in the classifier weights matrices. On the figure, this shows us the first input block kernel is mostly used by the first (red) subnetwork (7 vs. 1). Similarly, the first feature of the output block proves important only to the first (red) classifier.

Fig. 2). This tells us whether a feature map contains more information about one input, and we can visualize which maps are used by which subnetwork through histograms h_0 and h_1 of feature influence for each subnetwork. Quantitatively, we can approximate the feature sharing rate through the ratio of $\frac{\min(h_0, h_1)}{\max(h_0, h_1)}$. In the same spirit, we consider the L_1 norm of the columns of classifier weight matrices to quantify the importance of each feature to each classifier.

We conduct our study on a WideResNet-28-2 [15] using the more realistic batch repetition 2 setting from [11] on the CIFAR 100 dataset [4] (see Appendix). We choose to consider this situation as it perfectly showcases the issues encountered by MIMO methods on smaller architectures. To complement this, we also show results on the slightly larger WideResNet-28-5 later on in the paper.

Fig. 3 shows the subnetworks are fully independent in the core network: each channel in the input block encodes information about only one input, as the corresponding kernel of the other encoder’s L_1 is very low. A similar behavior is observed in the output block, and further analysis of input influence on intermediary feature maps shows this behavior remains consistent within the network (See Appendix).

Multi-input multi-output architectures’ scaling issues become much easier to understand in light of this: the amount of weights available to each of the underlying subnetworks decreases quadratically with the number of subnetwork. In-




Figure 3. Features are used by one subnetwork or the other, never both at the same time: the overlap (orange) is very low.

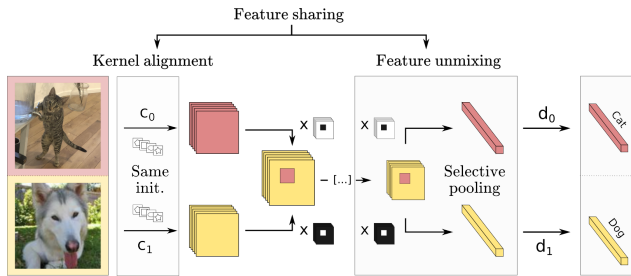


Figure 4. Two steps are necessary to allow feature sharing in networks: 1) Ensure the subnetworks share a “common language” by initializing the convolutional encoders to be close to each other. 2) Extract descriptions of each input from model features.

deed, since feature maps of different subnetworks cannot communicate, only $\frac{1}{M}$ weights can be non-zero. This fraction of non-zero weights must then be distributed between the M subnetworks. Furthermore, the subnetworks likely extract similar generic features, at least in the first layers. Since the subnetworks share no features, this means those features are unnecessarily replicated for each subnetwork.

This is not wholly surprising or undesirable behavior as MIMO strives to train M independent subnetworks to obtain diverse ensembles. By avoiding overlap between subnetworks, the subnetworks act as a standard ensemble of smaller models, with the base model size acting as hard cap on the number of the parameters used by the ensemble.

While it is true not sharing any features ensures subnetworks’ independence, it seems unnecessary. Indeed, the subnetworks are highly unlikely to extract completely different features. As such, subnetworks should benefit from sharing features at least in the early layers even if the classifier still consider fairly different features.

At first blush, nothing in the MIMO training protocol explicitly requires the subnetworks not share any features. Why do the subnetworks avoid sharing features? How could we encourage them to share some parameters?

3. How can subnetworks share features?

We discuss here the obstacles preventing feature sharing in multi-input multi-output architectures, and propose solutions to correct this behavior.




Figure 5. Applying unmixing leads to features being used by both subnetworks: the overlap (orange) is very high.

3.1. Unmixing: extracting features for each input

We build upon an intuition put forth in MixMo [11]: the lack of feature sharing is caused by the need for individual classifier at the end of the network to extract class information for one input specifically. Indeed, the M classifiers have access to the exact same set of extracted features. If two classifiers use the same feature, that feature needs to describe the state of two different inputs. This is an issue when one accounts for the fact inputs are in fact drawn independently and there can therefore be no meaningful feature describing the state of two inputs simultaneously.

Let us now consider how the classifiers should ideally behave on shared features. Since each classifier is paired to one of the input pathways, they should be able to extract two different interpretations of the shared features that still encode the same functional information (see Fig. 4). For instance, the shared feature should encode for the presence of flowers but each classifier should be able to infer from the feature whether its personal input contains flowers.

While this is not the case in traditional CNNs, MixMo [11] introduces a modification to the seminal MIMO architecture that causes feature maps to encode information about the different inputs separately. Indeed, since MixMo mixes inputs according to some binary mixing augmentation scheme (typically CutMix [14]), each pixel on the final feature maps encodes information about one of the inputs.

This is fortunate as it provides us with a fairly natural solution: unmixing. **Unmixing** (illustrated in Fig. 4) recycles the binary masks generated for input mixing in order to filter the feature maps so that only information relevant to a specific input is contained in the unmixed version. This way, a single feature map can describe each of the inputs.

Fig. 5 shows that applying unmixing causes the subnetworks to share features, both in the input and output block. In fact, every feature in the unmixed model is used by all subnetworks which proves unmixing indeed solves the core obstacle to feature sharing in MIMO networks.

Introducing unmixing however leads to unstable and generally worse performance as seen in Tab. 1. Crucially, even individual subnetwork accuracy suffers from unmixing which suggests an underlying issue.

Method	28-2			28-5		
	Acc. ens.	Acc. Ind.	Classifier share rate	Acc. ens.	Acc. Ind.	Classifier share rate
MixMo	74.8 ± 0.2	71.6 ± 0.2	4.9 ± 0.4	81.9 ± 0.1	79.4 ± 0.2	8.2 ± 0.1
MixMo + Unmix	69.7 ± 15.4	69.7 ± 15.4	99.1 ± 0.4	79.4 ± 2.7	79.4 ± 2.7	98.8 ± 0.7
MixMo + Unmix + kernel init.	79.0 ± 0.1	79.0 ± 0.1	99.4 ± 0.1	82.1 ± 0.2	82.1 ± 0.2	99.3 ± 0.1
MixShare (partial on 25% features)	73.3 ± 0.5	71.5 ± 1.0	60.6 ± 6.5	79.9 ± 0.3	78.8 ± 0.4	60.0 ± 2.3
MixShare (fadeout to 100 epochs)	79.0 ± 0.1	76.7 ± 1.3	64.4 ± 6.0	82.4 ± 0.3	81.6 ± 0.5	62.7 ± 8.3

Table 1. Overall ensemble accuracy (%), average subnetwork accuracy (%) and classifier (output block) sharing rate (see Sec. 2) for MixMo variants, *mean ± std* reported over 3 runs. Mixshare with fadeout unmixing yields both strong individual models and ensemble gains.

3.2. Aligning encoder kernels to allow efficient feature sharing

Intuitively, feature sharing should at the very least lead to higher individual subnetwork accuracy as the subnetworks use more parameters. As such, we now investigate why unmixing degrades performance so dramatically.

By extracting multiple possible interpretations of a single feature, unmixing introduces a new problem in the model. Indeed, we need our interpretations of the same feature to encode the same functional characteristics (e.g. flower detection). The issue is that a randomly initialized multi-input multi-output network typically leads to having multiple interpretations of the same feature.

Indeed, the encoders computing the mixed representations are very different. For an input feature, the mixed feature map could contain information about horizontal borders on input 1 and vertical borders on input 2. As such, there is no consistent interpretation for our mixed features.

We can unify the interpretation of unmixed features at the start by simply **aligning the kernels of the encoders**. Indeed, as long as each feature encodes the same sort of information for each encoder, there should be no ambiguity introduced by the unmixing process.

Tab. 1 shows that fixing the initialization scheme of the encoders to the same value does indeed lead the model to outperform normal mixmo models.

3.3. Towards partial feature sharing

While proper unmixing does allow feature sharing in multi-input multi-output networks, Tab. 1 and Fig. 5 show it leads to subnetworks sharing all features: the subnetworks are identical. This is even less desirable than fully separated subnetworks as it makes ensembling pointless [9, 10].

Ideally, subnetworks would share some parameters but still remain distinct functionally. This way, we would be able to strike a compromise between fully separated and fully shared subnetworks. The issue with this however, is that removing obstacles to feature sharing makes it unnecessary for subnetworks to separate in any way.

In this preliminary work, we discuss two solutions: **partial unmixing** and **fadeout unmixing**. **Partial unmixing** is

a straightforward solution where we only apply unmixing to a fixed subset of the final feature maps (e.g. 25%). In **Fadeout unmixing** we start training the network with proper unmixing but progressively reduce the strength of unmixing so that there is no unmixing towards the end of the procedure. For instance, we use the unmixing mask $M + r(1 - M)$ (instead of M) with $r = \min(1, epoch/100)$ if we want to stop unmixing by epoch 100. As such, fadeout unmixing initializes the network in a shared state and progressively pushes the subnetworks to develop independent features.

We now propose the full MixShare framework by combining proper kernel initialization and partial/fadeout unmixing along with slight adjustments to standard MIMO procedures like input repetition [2] and loss balancing [11] (see Appendix). Tab. 1 shows that both MixShare variants succeed in causing partial feature sharing. **Partial** fails to train strong individual subnetworks, but still showcases ensemble benefits. **Fadeout** on the other hand leads to strong performances and retains significant ensembling benefits on medium sized networks like a WideResNet 28-5.

4. Conclusion

We have shown multi-input multi-output models induce fully separated subnetworks because of a difficulty in matching outputs to inputs for the neural network. We have proposed an unmixing mechanism and encoder initialization for MixMo [11] architectures and demonstrated it allows multi-input multi-output architectures to share features. Our preliminary experiments show this corrected architecture outperforms standard multi-input multi-output architectures on smaller networks with a proper unmixing scheme. We hope that by highlighting the main issue at the crux of these architectures’ inefficiency, our work will lead to further research on MIMO architectures that will lead to their deployment smaller mobile and AR/VR devices.

Acknowledgments This work was conducted using HPC resources from GENCI–IDRIS (Grant 2021-AD011013208), and under a CIFRE grant between Thales Land and Air Systems and Sorbonne University.

References

- [1] Lars Kai Hansen and Peter Salamon. Neural network ensembles. *IEEE transactions on pattern analysis and machine intelligence*, 1990. [1](#)
- [2] Marton Havasi, Rodolphe Jenatton, Stanislav Fort, Jeremiah Zhe Liu, Jasper Snoek, Balaji Lakshminarayanan, Andrew Mingbo Dai, and Dustin Tran. Training independent subnetworks for robust prediction. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [4](#), [6](#)
- [3] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, 2016. [1](#)
- [4] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. [2](#)
- [5] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012. [1](#)
- [6] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in neural information processing systems*, 2017. [1](#)
- [7] Stefan Lee, Senthil Purushwalkam, Michael Cogswell, David J. Crandall, and Dhruv Batra. Why M heads are better than one: Training a diverse ensemble of deep networks. *Arxiv preprint*, 2015. [1](#)
- [8] Eran Malach, Gilad Yehudai, Shai Shalev-Schwartz, and Ohad Shamir. Proving the lottery ticket hypothesis: Pruning is all you need. In *International Conference on Machine Learning*, 2020. [1](#)
- [9] Tianyu Pang, Kun Xu, Chao Du, Ning Chen, and Jun Zhu. Improving adversarial robustness via promoting ensemble diversity. In *International Conference on Machine Learning*, 2019. [4](#)
- [10] Alexandre Rame and Matthieu Cord. Dice: Diversity in deep ensembles via conditional redundancy adversarial estimation. In *International Conference on Learning Representations*, 2021. [4](#)
- [11] Alexandre Rame, Remy Sun, and Matthieu Cord. Mixmo: Mixing multiple inputs for multiple outputs via deep subnetworks. In *International Conference on Computer Vision*, 2021. [1](#), [2](#), [3](#), [4](#), [6](#)
- [12] Divya Shanmugam, Davis Blalock, Guha Balakrishnan, and John Guttag. Better aggregation in test-time augmentation. In *International Conference on Computer Vision*, 2021. [1](#)
- [13] Yeming Wen, Dustin Tran, and Jimmy Ba. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. In *International Conference on Learning Representations*, 2019. [1](#)
- [14] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *International Conference on Computer Vision*, 2019. [3](#), [6](#)
- [15] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *British Machine Vision Conference*, 2016. [2](#)
- [16] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018. [1](#)