

RepMode: Learning to Re-parameterize Diverse Experts for Subcellular Structure Prediction

Donghao Zhou^{1,2,3} Chunbin Gu³ Junde Xu^{1,2,3} Furui Liu⁴ Qiong Wang¹
Guangyong Chen^{4*} Pheng-Ann Heng^{1,3,4}

¹ Guangdong Provincial Key Laboratory of Computer Vision and Virtual Reality Technology, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences

² University of Chinese Academy of Sciences ³ The Chinese University of Hong Kong

⁴ Zhejiang Lab

<https://correr-zhou.github.io/RepMode>

Abstract

In biological research, fluorescence staining is a key technique to reveal the locations and morphology of subcellular structures. However, it is slow, expensive, and harmful to cells. In this paper, we model it as a deep learning task termed subcellular structure prediction (SSP), aiming to predict the 3D fluorescent images of multiple subcellular structures from a 3D transmitted-light image. Unfortunately, due to the limitations of current biotechnology, each image is partially labeled in SSP. Besides, naturally, subcellular structures vary considerably in size, which causes the multi-scale issue of SSP. To overcome these challenges, we propose Re-parameterizing Mixture-of-Diverse-Experts (RepMode), a network that dynamically organizes its parameters with task-aware priors to handle specified single-label prediction tasks. In RepMode, the Mixture-of-Diverse-Experts (MoDE) block is designed to learn the generalized parameters for all tasks, and gating re-parameterization (GatRep) is performed to generate the specialized parameters for each task, by which RepMode can maintain a compact practical topology exactly like a plain network, and meanwhile achieves a powerful theoretical topology. Comprehensive experiments show that RepMode can achieve state-of-the-art overall performance in SSP.

1. Introduction

Recent years have witnessed great progress in biological research at the subcellular level [1,2,7,21,23,59,61], which plays a pivotal role in deeply studying cell functions and behaviors. To address the difficulty of observing subcellu-

*Corresponding Author: gychen@zhejianglab.com

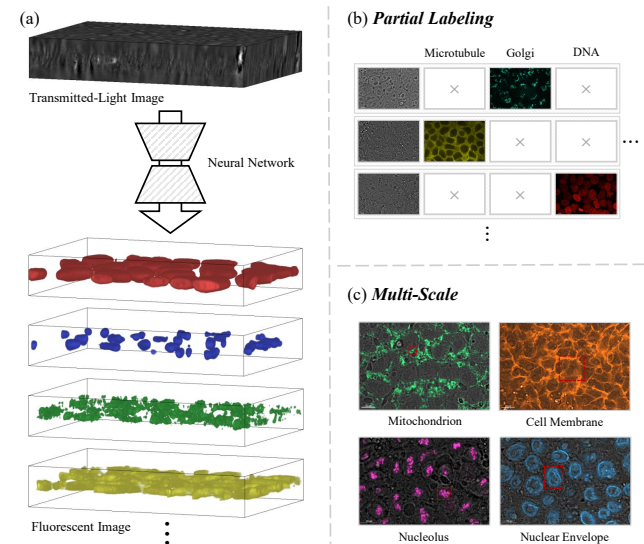


Figure 1. (a) Illustration of subcellular structure prediction (SSP), which aims to predict the 3D fluorescent images of multiple subcellular structures from a 3D transmitted-light image. This task faces two challenges, *i.e.* (b) partial labeling and (c) multi-scale.

lar structures, fluorescence staining was invented and has become a mainstay technology for revealing the locations and morphology of subcellular structures [26]. Specifically, biologists use the antibodies coupled to different fluorescent dyes to “stain” cells, after which the subcellular structures of interest can be visualized by capturing distinct fluorescent signals [64]. Unfortunately, fluorescence staining is expensive and time-consuming due to the need for advanced instrumentation and material preparation [29]. Besides, phototoxicity during fluorescent imaging is detrimental to living cells [28]. In this paper, we model fluorescence

staining as a deep learning task, termed *subcellular structure prediction (SSP)*, which aims to directly predict the 3D fluorescent images of multiple subcellular structures from a 3D transmitted-light image (see Fig. 1(a)). The adoption of SSP can significantly reduce the expenditure on subcellular research and free biologists from this demanding workflow.

Such an under-explored and challenging bioimage problem deserves the attention of the computer vision community due to its high potential in biology. Specifically, SSP is a dense regression task where the fluorescent intensities of multiple subcellular structures need to be predicted for each transmitted-light voxel. However, due to the limitations of current biotechnology, each image can only obtain *partial labels*. For instance, some images may only have the annotations of nucleoli, and others may only have the annotations of microtubules (see Fig. 1(b)). Moreover, different subcellular structures would be presented at *multiple scales* under the microscope, which also needs to be taken into account. For example, the mitochondrion is a small structure inside a cell, while obviously the cell membrane is a larger one since it surrounds a cell (see Fig. 1(c)).

Generally, there are two mainstream solutions: 1) *Multi-Net* [5, 32, 33, 47]: divide SSP into several individual prediction tasks and employs multiple networks; 2) *Multi-Head* [6, 9, 46]: design a partially-shared network composed of a shared feature extractor and multiple task-specific heads (see Fig. 2(a)). However, these traditional approaches organize network parameters in an *inefficient* and *inflexible* manner, which leads to two major issues. First, they fail to make full use of partially labeled data in SSP, resulting in *label-inefficiency*. In Multi-Net, only the images containing corresponding labels would be selected as the training set for each network and thus the other images are wasted, leading to an unsatisfactory generalization ability. As for Multi-Head, although all images are adopted for training, only partial heads are updated when a partially labeled image is input and the other heads do not get involved in training. Second, to deal with the multi-scale nature of SSP, they require exhausting pre-design of the network architecture, and the resultant one may not be suitable for all subcellular structures, which leads to *scale-inflexibility*.

In response to the above issues, herein we propose *Reparameterizing Mixture-of-Diverse-Experts (RepMode)*, an all-shared network that can dynamically organize its parameters with task-aware priors to perform specified single-label prediction tasks of SSP (see Fig. 2(b)). Specifically, RepMode is mainly constructed of the proposed *Mixture-of-Diverse-Experts (MoDE) blocks*. The MoDE block contains the expert pairs of various receptive fields, where these *task-agnostic* experts with diverse configurations are designed to *learn the generalized parameters for all tasks*. Moreover, *gating re-parameterization (GatRep)* is proposed to conduct the *task-specific* combinations of experts to achieve efficient

expert utilization, which aims to *generate the specialized parameters for each task*. With such a parameter organizing manner (see Fig. 2(c)), RepMode can maintain a practical topology exactly like a plain network, and meanwhile achieves a theoretical topology with a better representational capacity. Compared to the above solutions, RepMode can fully learn from all training data, since the experts are shared with all tasks and thus participate in the training of each partially labeled image. Besides, RepMode can adaptively learn the preference of each task for the experts with different receptive fields, thus no manual intervention is required to handle the multi-scale issue. Moreover, by fine-tuning few newly-introduced parameters, RepMode can be easily extended to an unseen task without any degradation of the performance on the previous tasks. Our main contributions are summarized as follows:

- We propose a stronger baseline for SSP, named RepMode, which can switch different “modes” to predict multiple subcellular structures and also shows its potential in task-incremental learning.
- The MoDE block is designed to enrich the generalized parameters and GatRep is adopted to yield the specialized parameters, by which RepMode achieves dynamic parameter organizing in a task-specific manner.
- Comprehensive experiments show that RepMode can achieve state-of-the-art (SOTA) performance in SSP. Moreover, detailed ablation studies and further analysis verify the effectiveness of RepMode.

2. Related Works

Partially labeled dense prediction. In addition to SSP, many other dense prediction tasks could also face the challenge of partial labeling. In general, the previous methods can be divided into two groups. The first one seeks for an effective training scheme by adopting knowledge distillation [17, 72], learning cross-task consistency [36], designing jointly-optimized losses [54], *etc.* The second one aims to improve the network architecture with a dynamic segmentation head [70], task-guided attention modules [62], conditional tensor incorporation [15], *etc.* However, these methods are primarily developed for large-scale datasets. Compared to these well-explored tasks, SSP only has relatively small datasets due to the laborious procedure of fluorescence staining. Thus, the training data of SSP should be utilized in a more efficient way. In light of that, we adopt a task-conditioning strategy in RepMode, where all parameters are shared and thus can be directly updated using the supervision signal of each label. Unlike other task-conditional networks [15, 55, 62, 70], our RepMode is more flexible and capable of maintaining a compact topology.

Multi-scale feature learning. Multi-scale is a fundamental problem of computer vision, caused by the variety

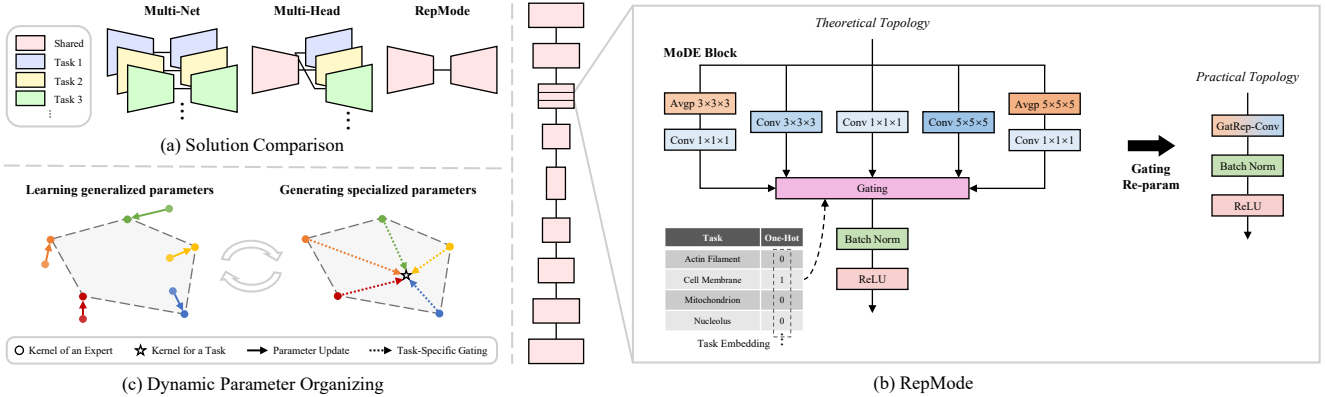


Figure 2. Overview of the proposed method. (a) Comparison of two mainstream solutions (*i.e.* Multi-Net and Multi-Head) and the proposed method (*i.e.* RepMode) for SSP. (b) Illustration of our RepMode which includes two key components, *i.e.* the proposed MoDE block and GatRep. (c) Diagram of how RepMode dynamically organizes its parameters in a MoDE block. Note that the gray region denotes the convex hull decided by the expert kernels, and the convex hull is the area where the task-specific kernels would be situated.

in the size of the objects of interest. The common solutions are adopting multi-resolution input [16, 19, 66, 71], designing parallel branches [3, 25, 37, 39, 57, 74], fusing cross-layer features [42, 44, 50, 56], performing hierarchical predictions [40, 41, 43], *etc.* These methods often adopt a pre-defined architecture for all objects to extract multi-scale features in a unified fashion. In contrast, RepMode learns the dynamic combinations of the experts with different receptive fields for each subcellular structure, and thus is capable of learning multi-scale features in a task-specific manner.

Mixture-of-Experts. Mixture-of-Experts (MoE) typically consists of a gating module and multiple independent learners (*i.e.* experts) [67]. For an input sample, MoE would adaptively assemble the corresponding output of all experts [10, 45, 48, 58] or only route it to a few specific experts [24, 31, 51, 53], which depends on the gating strategy. Benefiting from its divide-and-conquer principle, MoE is widely adopted in computer vision [10, 20, 48, 51, 63], natural language processing [8, 22, 53], and recommendation systems [24, 45, 49, 58]. Our RepMode is established based on the idea of MoE, but is further explored from the following aspects: 1) Instead of performing input-aware gating, RepMode only uses the task embedding for gating, aiming to adjust its behavior for a specified task; 2) The experts of RepMode can be combined together, which can efficiently utilize multiple experts in an MoE-inspired architecture.

Structural re-parameterization. Different from other re-parameterization (re-param) methods [18, 35, 52, 68], structural re-param [13, 14] is a recent technique of equivalently converting multi-branch network structures. With this technique, multi-branch blocks [12–14, 60] are introduced to plain networks for enhancing their performance. However, these methods only achieve inference-time converting, resulting in non-negligible training costs. There are previous works [11, 27] accomplishing training-time con-

verting, but they require model-specific optimizer modification [11] or extra parameters [27] and only explore its potential on one single task. In this work, we elegantly incorporate task-specific gating into structural re-param to achieve both training- and inference-time converting for handling multiple tasks, which is more cost-friendly and with better applicability. Besides, dynamic convolutions [4, 38, 65, 73] also can be roughly considered as re-param methods, which aim to assemble convolutions with the same shape in an input-dependent way. In contrast, using task-dependent gating, our RepMode can combine experts with diverse configurations to generate composite convolutional kernels, and thus is with higher flexibility to model more situations.

3. Methodology

3.1. Problem definition

We start by giving a formal definition of SSP. Following [47], we assume that each image has *only one* fluorescent label, which greatly relaxes the annotation requirement and makes the setting of this task more general and challenging. Let $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n, l_n)\}_{n=1}^N$ denotes a SSP dataset with N samples. The n -th image $\mathbf{x}_n \in \mathcal{I}_n$ is associated with the label $\mathbf{y}_n \in \mathcal{L}_n$, where $\mathcal{I}_n = \mathbb{R}^{D_n \times H_n \times W_n}$ denotes the image space and $D_n \times H_n \times W_n$ is the image size. The label indicator $l_n \in \mathcal{L} = \{1, 2, \dots, S\}$ represents that \mathbf{y}_n is the label of the l_n -th subcellular structure, where S is the total number of subcellular structure categories. In this work, our goal is to learn a network $F : \mathcal{I} \times \mathcal{L} \rightarrow \mathcal{I}$ with the parameters θ from \mathcal{D} . SSP can be considered as a collection of S single-label prediction tasks, each of which corresponds to one category of subcellular structures. To solve SSP, Multi-Net and Multi-Head divide task-specific parameters from θ for each task. In contrast, RepMode aims to share θ with all tasks and dynamically organize θ to handle specified tasks.

3.2. Network architecture

The backbone of RepMode is a 3D U-shape encoder-decoder architecture mainly constructed of the downsampling and upsampling blocks. Specifically, each downsampling block contains two successive MoDE blocks to extract task-specific feature maps and double their channel number, followed by a downsampling layer adopting a convolution with a kernel size of $2 \times 2 \times 2$ and a stride of 2 to halve their resolution. Note that batch normalization (BN) and ReLU activation are performed after each convolutional layer. In each upsampling block, an upsampling layer adopts a transposed convolution with a kernel size of $2 \times 2 \times 2$ and a stride of 2 to upsample feature maps and halve their channel number. Then, the upsampled feature maps are concatenated with the corresponding feature maps passed from the encoder, and the resultant feature maps are further refined by two successive MoDE blocks. Finally, a MoDE block without BN and ReLU is employed to reduce the channel number to 1, aiming to produce the final prediction. We adopt such a common architecture to highlight the applicability of RepMode and more details are provided in Appendix A. Notably, MoDE blocks are employed in both the encoder and decoder, which can facilitate task-specific feature learning and thus helps to achieve superior performance.

3.3. Mixture-of-Diverse-Experts block

To handle various prediction tasks of SSP, the representational capacity of the network should be strengthened to guarantee its generalization ability. Thus, we propose the MoDE block, a powerful alternative to the vanilla convolutional layer, to serve as the basic network component of RepMode. In the MoDE block, diverse experts are designed to explore a unique convolution collocation, and the gating module is designed to utilize the task-aware prior to produce gating weights for dynamic parameter organizing. We delve into the details of these two parts in the following.

Diverse expert design. In the MoDE block, we aim to achieve two types of expert diversity: 1) *Shape diversity*: To tackle the multi-scale issue, the experts need to be equipped with various receptive fields; 2) *Kernel diversity*: Instead of irregularly arranging convolutions, it is a better choice to explore a simple and effective pattern to further enrich kernel combinations. Given these guidelines, we propose to construct *expert pairs* to constitute the multi-branch topology of the MoDE block. The components of an expert pair are 3D convolutions (Conv) and 3D average poolings (Avgp). Specifically, an expert pair contains a Conv $K \times K \times K$ expert and an Avgp $K \times K \times K$ - Conv $1 \times 1 \times 1$ expert, and we utilize a stride of 1 and same-padding to maintain the resolution of feature maps. Overall, the MoDE block is composed of expert pairs with three receptive fields to attain shape diversity (see Fig. 3(a)). When $K = 1$, since these two experts are equal, only one is preserved for simplicity.




Figure 3. (a) Expert pairs with a receptive field size of $K = 5, 3, 1$. Note that each branch denotes an expert. (b) Examples of two types of Conv kernels in an expert pair, including A-Conv and normal Conv. Here we present the kernels of a 2D version with a receptive field size of 3 for simplicity.

Notably, the Avgp $K \times K \times K$ - Conv $1 \times 1 \times 1$ expert is essentially a special form of the Conv $K \times K \times K$ expert. To be specific, merging the serial Avgp $K \times K \times K$ kernel and Conv $1 \times 1 \times 1$ kernel would result in a Conv kernel with limited degrees of freedom (named as A-Conv). Compared to normal Conv, A-Conv has only one learnable parameter and thus acts like a learnable average pooling, which enriches kernel diversity in the same shape (see Fig. 3(b)). The combination of Conv and Avgp is also widely adopted in previous works [13, 27], but we further explore such a characteristic from the perspective of serial merging.

Gating module design. In order to perform a specified single-label prediction task, the task-aware prior needs to be encoded into the network, so that it can be aware of which task is being handled and adjust its behavior to focus on the desired task. Instead of embedding the task-aware prior by a hash function [15] or a complicated learnable module [55], we choose the most simple way, *i.e.* embed the task-aware prior of each input image \mathbf{x}_n with the label indicator l_n into a S -dimensional one-hot vector \mathbf{p}_n , which is expressed as

$$p_{ns} = \begin{cases} 1, & \text{if } s = l_n, \\ 0, & \text{otherwise,} \end{cases} \quad s = 1, 2, \dots, S, \quad (1)$$

where p_{ns} indicates the s -th entry of \mathbf{p}_n . Then, the task embedding \mathbf{p}_n is fed into the gating module and the gating weights \mathbf{G} are generated by a single-layer fully connected network (FCN) $\phi(\cdot)$, shown as $\mathbf{G} = \phi(\mathbf{p}_n) = \{\mathbf{g}_t\}_{t=1}^T$ where $T = 5$. Note that we omit n in \mathbf{G} for brevity. Here $\mathbf{g}_t \in \mathbb{R}^{C_0}$ represents the gating weights for the t -th experts, which is split from \mathbf{G} , and C_0 is the channel number of the output feature maps. Finally, \mathbf{G} would be further activated as $\hat{\mathbf{G}} = \{\hat{\mathbf{g}}_t\}_{t=1}^T$ by Softmax for the balance of the intensity of different experts, which can be formulated as

$$\hat{g}_{ti} = \frac{\exp(g_{ti})}{\sum_{j=1}^T \exp(g_{ji})}, \quad i = 1, 2, \dots, C_0, \quad (2)$$

where g_{ti} (resp. \hat{g}_{ti}) is the i -th entry of \mathbf{g}_t (resp. $\hat{\mathbf{g}}_t$). With the resultant gating weights $\hat{\mathbf{G}}$, RepMode can perform dynamic parameter organizing for these task-agnostic experts conditioned on the task-aware prior.

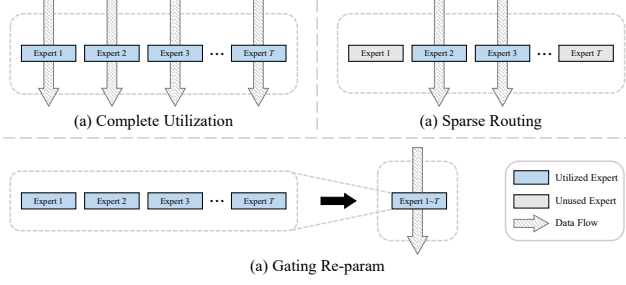


Figure 4. Diagram of expert utilization manners of MoE, including (a) complete utilization, (b) sparse routing, and (c) our GatRep.

3.4. Gating re-parameterization

In addition to studying expert configurations, how to efficiently utilize multiple experts is also worth further exploration. The traditional manner is to completely utilize all experts to process the input feature maps [10,45,48,58] (see Fig. 4(a)). However, the output of all experts needs to be calculated and stored, which would slow down training and inference and increase the GPU memory utilization [14,69]. The advanced one is to sparsely route the input feature maps to specific experts [24,31,51,53] (see Fig. 4(b)). However, only a few experts are utilized and the others remain unused, which would inevitably reduce the representational capacity of MoE. To avoid these undesired drawbacks and meanwhile preserve the benefits of MoE, we elegantly introduce task-specific gating to structural re-param, and thus propose GatRep to adaptively fuse the kernels of experts in the MoE block, through which only one convolution operation is explicitly required (see Fig. 4(c)).

Preliminary. GatRep is implemented based on the homogeneity and additivity of Conv and Avgp, which are recognized in [13,14]. The kernels of a Conv with C_1 input channels, C_O output channels, and $K \times K \times K$ kernel size is a fifth-order tensor $\mathbf{W} \in \mathcal{Z}(K) = \mathbb{R}^{C_O \times C_1 \times K \times K \times K}$, where $\mathcal{Z}(K)$ denotes this kernel space. Besides, the kernels of an Avgp $\mathbf{W}^a \in \mathbb{R}^{C_1 \times C_1 \times K \times K \times K}$ can be constructed by

$$\mathbf{W}_{c_1, c'_1, :, :, :}^a = \begin{cases} \frac{1}{K^3}, & \text{if } c_1 = c'_1, \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

where $c_1, c'_1, :, :, :$ is the indexes of the tensor and $c_1, c'_1 = 1, \dots, C_1$. Note that the Avgp kernel is fixed and thus unlearnable. Moreover, we omit the biases here as a common practice. Specifically, GatRep can be divided into two steps, *i.e.* serial merging and parallel merging (see Fig. 5).

Step 1: serial merging. The first step of GatRep is to merge Avgp and Conv into an integrated kernel. For brevity, here we take an Avgp - Conv expert as an example. Let \mathbf{M}^1 denote the input feature maps. The process of producing the output feature maps \mathbf{M}^0 can be formulated as

$$\mathbf{M}^0 = \mathbf{W} \circledast (\mathbf{W}^a \circledast \mathbf{M}^1), \quad (4)$$

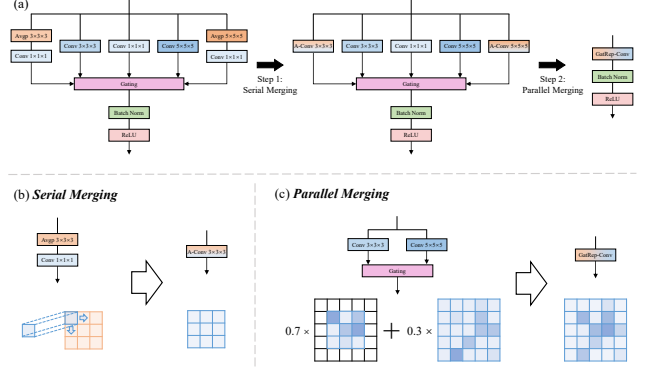


Figure 5. (a) Process of GatRep which includes two steps, *i.e.* (b) serial merging and (c) parallel merging.

where \circledast denotes the convolution operation. According to the associative law, we can perform an equivalent transformation for Eq. (4) by first combining \mathbf{W}^a and \mathbf{W} . Such a transformation can be expressed as

$$\mathbf{M}^0 = \underbrace{(\mathbf{W} \circledast \mathbf{W}^a)}_{\mathbf{W}^c} \circledast \mathbf{M}^1, \quad (5)$$

which means that we first adopt \mathbf{W} to perform a convolution operation on \mathbf{W}^a , and then use the resultant kernel \mathbf{W}^c to process the input feature maps. With this transformation, the kernels of Avgp and Conv can be merged as an integrated one for the subsequent step.

Step 2: parallel merging. The second step of GatRep is to merge all experts in a task-specific manner. We define $\text{Pad}(\cdot, K')$ as a mapping function equivalently transferring a kernel to the kernel space $\mathcal{Z}(K')$ by zero-padding, and set $K' = 5$ which is the biggest receptive field size of these experts. Let $\hat{\mathbf{M}}^0$ denote the final task-specific feature maps. This transformation can be formulated as

$$\hat{\mathbf{M}}^0 = \underbrace{\left(\sum_{t=1}^T \hat{\mathbf{g}}_t \odot \text{Pad}(\mathbf{W}_t^c, K') \right)}_{\mathbf{W}^c} \circledast \mathbf{M}^1, \quad (6)$$

where \odot denotes the channel-wise multiplication and \mathbf{W}_t^c is the kernel of the t -th expert. Note that \mathbf{W}_t^c is an integrated kernel (resp. Conv kernel) for an Avgp - Conv expert (resp. Conv expert). To be specific, the detailed pixel-level form is provided in Appendix B. Finally, \mathbf{W}^c is the resultant task-specific kernel dynamically generated by GatRep.

4. Experiments

4.1. Experimental setup

We conduct the experiments based on the following experimental setup unless otherwise specified. Due to space limitations, more details are included in Appendix C.

Methods	Actin Filament			Actom. Bundle			Cell Membrane			Desmosome			DNA			Endop. Reticulum			Golgi Apparatus		
	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2
Multi-Net [47]	.4241	.4716	.5695	.7247	.4443	.2606	.5940	.4351	.3930	.8393	.5640	.0162	.5806	.5033	.3822	.4635	.4914	.5262	.8023	.5732	.0801
Multi-Head (Dec.)	.4278	.4803	.5657	.7052	.4363	.2804	.5785	.4625	.4089	.8431	.5677	.0118	.5312	.4764	.4346	.4454	.4832	.5448	.7925	.5768	.0910
Multi-Head (Las.)	.4648	.4978	.5281	.6697	.4222	.3168	.5568	.4441	.4310	.8402	.5637	.0148	.5088	.4824	.4581	.4372	.4697	.5531	.7918	.5807	.0921
CondNet [15]	.4246	.4719	.5688	.6873	.4286	.2988	.5635	.4157	.4242	.8422	.5655	.0126	.4967	.4707	.4712	.4290	.4697	.5615	.7996	.5823	.0831
TSNs [55]	.4279	.4779	.5656	.6691	.4111	.3174	.5309	.4346	.4575	.8392	.5630	.0160	.4974	.4682	.4702	.4362	.4785	.5543	.7892	.5777	.0949
PIPO-FAN [16]	.4063	.4603	.5873	.6815	.4306	.3046	.5440	.4389	.4441	.8417	.5674	.0131	.4868	.4626	.4813	.4433	.4832	.5470	.7968	.5861	.0861
DoDNet [70]	.4215	.4706	.5721	.6989	.4204	.2870	.5459	.4390	.4422	.8415	.5633	.0133	.5280	.4810	.4382	.4414	.4844	.5490	.7927	.5774	.0909
TGNet [62]	.3917	.4535	.6023	.6843	.4213	.3018	.5856	.4227	.4015	.8392	.5654	.0160	.5011	.4746	.4666	.4441	.4806	.5460	.7870	.5774	.0973
RepMode	.3936	.4558	.6004	.6572	.4103	.3295	.5443	.4136	.4437	.8358	.5619	.0199	.4852	.4598	.4831	.4046	.4445	.5865	.7792	.5694	.1064

Methods	Microtubule			Mitochondria			Nuclear Envelope			Nucleolus			Tight Junction			All			Δ_{Imp} (%)		
	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2	MSE	MAE	R^2
Multi-Net [47]	.3682	.4348	.6296	.4684	.3921	.5172	.3014	.3006	.6954	.2164	.1789	.7826	.6474	.3369	.3370	.5341	.4269	.4337	0.000	0.000	0.000
Multi-Head (Dec.)	.3932	.4594	.6044	.4545	.3888	.5315	.2687	.2895	.7284	.2114	.1762	.7877	.6396	.3252	.3451	.5226	.4258	.4456	2.149	0.272	2.752
Multi-Head (Las.)	.3781	.4465	.6196	.4649	.3991	.5208	.2909	.3057	.7059	.2213	.1870	.7778	.6547	.3367	.3298	.5223	.4275	.4461	2.218	-0.12	2.868
CondNet [15]	.3868	.4523	.6108	.4673	.4067	.5184	.2876	.3014	.7094	.2203	.1865	.7787	.6569	.3241	.3274	.5206	.4232	.4478	2.534	0.884	3.249
TSNs [55]	.3407	.4235	.6572	.4625	.3956	.5233	.2904	.2991	.7064	.2116	.1751	.7874	.6479	.3320	.3367	.5113	.4192	.4572	4.263	1.804	5.437
PIPO-FAN [16]	.3604	.4365	.6373	.4750	.4171	.5105	.2904	.3003	.7065	.2097	.1782	.7894	.6437	.3282	.3410	.5141	.4237	.4543	3.747	0.766	4.764
DoDNet [70]	.3972	.4606	.6004	.4772	.4119	.5081	.2976	.3164	.6992	.2250	.1934	.7740	.6703	.3336	.3137	.5276	.4291	.4406	1.227	-0.49	1.607
TGNet [62]	.3569	.4310	.6410	.4585	.3971	.5274	.2748	.2940	.7222	.2093	.1799	.7897	.6232	.3238	.3619	.5108	.4183	.4578	4.363	2.022	5.566
RepMode	.3389	.4171	.6590	.4459	.3885	.5404	.2631	.2820	.7340	.1995	.1682	.7997	.6168	.3245	.3685	.4956	.4078	.4735	7.209	4.482	9.176

Table 1. Experimental results of the proposed RepMode and the comparing methods on twelve prediction tasks of SSP. The best performance (lowest MSE and MAE, highest R^2) is marked in bold. Note that ‘‘All’’ indicates the overall performance.

Datasets. For a comprehensive comparison, the dataset is constructed from a dataset collection [47] containing twelve partially labeled datasets, each of which corresponds to one category of subcellular structures (*i.e.* one single-label prediction task). All images are resized to make each voxel correspond to $0.29 \times 0.29 \times 0.29 \mu\text{m}^3$. Moreover, we perform per-image z-scored normalization for voxels to eliminate systematic differences in illumination intensity. For each dataset, we randomly select 25% samples for evaluation and then withhold 10% of the rest for validation.

Implementation details. Mean Squared Error (MSE) is adopted as the loss function, which is commonly used to train a regression model. Besides, Adam [34] is employed as the optimizer with a learning rate of 0.0001. Each model is trained for 1000 epochs from scratch and validation is performed every 20 epochs. Finally, the validated model that attains the lowest MSE is selected for evaluation on the test set. In a training epoch, we randomly crop a patch with a size of $32 \times 128 \times 128$ from each training image as the input with a batch size of 8, and random flip is performed for data augmentation. In the inference stage, we adopt the Gaussian sliding window strategy [30] to aggregate patch-based output for a full prediction. To ensure fairness, the same backbone architecture, training configuration, and inference strategy are applied to all comparing models.

Evaluation metrics. In addition to MSE, Mean Absolute Error (MAE) and Coefficient of Determination (R^2) are also used as the evaluation metrics. MAE measures absolute differences and thus is less sensitive to outliers than MSE.

R^2 measures correlations by calculating the proportion of variance in a label that can be explained by its prediction. For a clear comparison, we also present the relative overall performance improvement over Multi-Net (*i.e.* Δ_{Imp}).

4.2. Comparing to state-of-the-art methods

We compared our RepMode to the following methods: 1) Multi-Net: [47]; 2) Multi-Head: include two variants, *i.e.* multiple task-specific decoders (denoted by Dec.) or last layers (denoted by Las.); 3) CondNet [15] and TSNs [55]: two SOTA task-conditional networks for multi-task learning; 4) PIPO-FAN [16], DoDNet [70], and TGNet [62]: three SOTA methods of a similar task, *i.e.* partially labeled multi-organ and tumor segmentation (note that DoDNet and TGNet also adopt task-conditioning strategies).

The experimental results on twelve tasks of SSP are reported in Tab. 1. As recognized in [47], the performance of Multi-Net is sufficient to assist with biological research in some cases, thus it can be a reference of reliable metric values for real-life use. Furthermore, two Multi-Head variants can achieve better performance, which verifies the importance of learning from the complete dataset. Notably, PIPO-FAN is an improved Multi-Head variant that additionally constructs a pyramid architecture to handle the multi-scale issue. The results show that such an architecture can further improve performance but still can not address this issue well. Moreover, the competitive performance of CondNet and TSNs demonstrates that adopting an appropriate task-conditioning strategy is beneficial. However, these

Ablation	Methods	MSE	MAE	R^2
Scope	only in Dec.	.5097	.4139	.4590
	only in Enc.	.5079	.4184	.4607
Expert	w/o $1 \times 1 \times 1$ expert pair	.5027	.4106	.4662
	w/o $3 \times 3 \times 3$ expert pair	.5080	.4141	.4605
	w/o $5 \times 5 \times 5$ expert pair	.5017	.4108	.4672
	w/o Conv expert	.5631	.4346	.4042
	w/o Avgp - Conv expert	.5037	.4101	.4651
Average Pooling	w/o Avgp	.4999	.4112	.4691
	all use Avgp $3 \times 3 \times 3$.4974	.4072	.4716
	all use Avgp $5 \times 5 \times 5$.4964	.4091	.4725
Gating	use Gauss. task embedding	.5071	.4155	.4616
	use two-layer FCN	.4980	.4060	.4710
	use Sigmoid activation	.4992	.4094	.4698
	Input-dep. gating	.7958	.5527	.1619
Original	RepMode	.4956	.4078	.4735

Table 2. Ablation studies from four aspects. Note that ‘‘Methods’’ denotes the variants of the proposed RepMode.

networks remain Multi-Head variants since multiple task-specific heads are still required. As an advanced version of DoDNet, TGNet additionally modulates the feature maps in the encoder and skip connections, leading to more competitive performance. It can be observed that, SSP is an extremely tough task since it is hard to attain a huge performance leap in SSP, even for those powerful SOTA methods of related tasks. However, the proposed RepMode, which aims to learn the task-specific combinations of diverse task-agnostic experts, outperforms the existing methods on ten of twelve tasks of SSP and achieves SOTA overall performance. Notably, RepMode can even achieve 7.209% (resp. 4.482%, 9.176%) Δ_{Imp} on MSE (resp. MAE, R^2), which is near twice the second best method (*i.e.* TGNet).

4.3. Ablation studies

To verify the effectiveness of the proposed RepMode, we conduct comprehensive ablation studies totally including four aspects, where the results are reported in Tab. 2.

Scope of use of the MoDE block. As we mentioned in Sec. 3.2, we employ the MoDE block in both the encoder and decoder of the network. Therefore, we change its scope of use to explore its influence on performance. The results show that employing it only in the encoder can achieve better performance than only in the decoder, since the encoder can extract task-specific features and pass them to the decoder through skip connections. Moreover, employing it in both the encoder and decoder is superior since the whole network can perform dynamic parameter organizing.

Effectiveness of the expert design. The MoDE block is composed of three expert pairs, each of which contains a Conv expert and an Avgp - Conv expert. It can be observed that removing experts (especially the Conv experts) from the MoDE block could cause a performance drop due to the

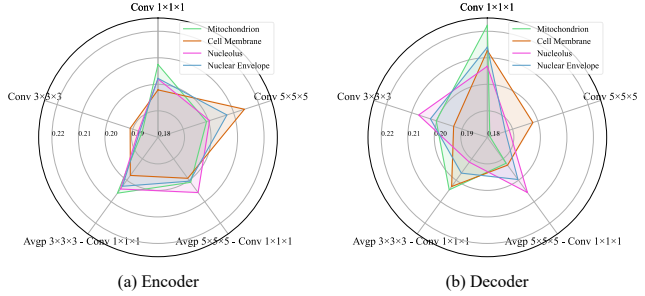


Figure 6. Visualization of channel-wisely averaged gating weights of two MoDE blocks, which are randomly selected from the encoder and decoder of a well-trained RepMode, respectively. The results of subcellular structures shown in Fig. 1(c) are presented.

degradation of the representational capacity. Moreover, it is also an interesting finding that the expert pair with the commonly used $3 \times 3 \times 3$ receptive field is most critical.

Average poolings matter. Avgp is one of the basic components of the MoDE block. The results show that removing such an unlearnable component could also reduce performance, which further verifies the effectiveness of the expert design. Besides, we can observe that setting the receptive fields of all Avgp to the same one could also cause a performance drop. This is because being equipped with different receptive fields could facilitate expert diversity.

Different gating strategies. For task-specific gating, the one-hot task embedding is fed into the single-layer FCN followed by Softmax activation. Accordingly, we conduct the following modifications: 1) Use the task embedding with each entry sampled from $\mathcal{N}(0, 1)$; 2) Use the two-layer FCN with the hidden unit number set to 6; 3) Use Sigmoid activation; 4) Input-dependent gating: input feature maps are first processed by a global average pooling and then fed into the gating module. The superior performance and simplicity of the original gating approach demonstrate the applicability of RepMode. Notably, input-dependent gating underperforms since an all-shared network can not be aware of the desired task of input without access to any priors.

4.4. Further analysis

In this subsection, we perform further analysis of RepMode to further reveal its capability. Additional analysis and discussion are provided in Appendix D.

Gating weights visualization. In the MoDE block, the gating weights are produced for dynamic parameter organizing, through which the preference of each task for diverse experts can be learned. As shown in Fig. 6, the cell membrane relatively prefers the Conv $5 \times 5 \times 5$ expert while the mitochondrion relatively prefers the Conv $1 \times 1 \times 1$ one as we expect. Besides, the preference of the mid-scale structures (*i.e.* nucleolus and nuclear envelope) is more variable. Notably, the Avg - Conv experts also be assigned sufficient weights, which verifies the effectiveness of the expert pairs.




Figure 7. Examples of the prediction results on the test set, including (a) microtubule, (b) actin filament, and (c) DNA. We compare the predictions of our RepMode with the ones of TGNets [62] which is a competitive method. Note that the dotted boxes indicate the major prediction difference. More examples are provided in Appendix E.




Figure 8. Diagram of the MoDE block for task-incremental learning. Note that here we employ an extra Conv $3 \times 3 \times 3$ expert.

It can also be observed that the network pays more attention to small-scale features in the decoder, which could be due to the need for producing a detailed prediction.

Qualitative results. Subcellular structures are hard to be distinguished in transmitted-light images (see Fig. 7). The second best method (*i.e.* TGNets) suffers from incomplete (see Fig. 7(a)) and redundant (see Fig. 7(b)) predictions, and even yields inexistent patterns (see Fig. 7(c)). But relatively, RepMode can produce more precise predictions for various subcellular structures at multiple scales even though there are some hard cases (see Fig. 7(a)&(b)). Such a practical advance is crucial in biological research, since inaccurate predictions at some key locations may mislead biologists into making incorrect judgments.

RepMode as a better task-incremental learner. For a well-trained RepMode, we fine-tuned a newly-introduced expert and gating module for each MoDE block with the other experts frozen, aiming to extend it to an unseen task (see Appendix D.1 for more details). As we expect, RepMode can preserve and transfer its domain knowledge through the pretrained experts, which helps to achieve better performance compared to the plain networks (see Tab. 3). As long as the previous gating weights have been stored, such a task-incremental learning manner would not cause any degradation of the performance on the previous tasks.

Methods	Strategies	Nucleolus			Cell Membrane		
		MSE	MAE	R^2	MSE	MAE	R^2
Multi-Net [47]	Individual Training	.2164	.1789	.7826	.5940	.4351	.3930
Multi-Head (Dec.)	All Fine-Tuning	.2121	.1811	.7870	.5339	.4097	.4543
RepMode	Experts Frozen	.2052	.1774	.7939	.5260	.4077	.4625

Table 3. Experimental results of task-incremental learning for two basic subcellular structures. Note that the strategies of the experimental methods for extending to a new task are presented.

5. Conclusions

In this paper, we focus on an under-explored and challenging bioimage problem termed SSP, which faces two main challenges, *i.e.* partial labeling and multi-scale. Instead of constructing a network in a traditional manner, we choose to dynamically organize network parameters with task-aware priors and thus propose RepMode. Experiments show that RepMode can achieve SOTA performance in SSP. We believe that RepMode can serve as a stronger baseline for SSP and help to motivate more advances in both the biological and computer vision community.

Acknowledgements

We would like to thank Danruo Deng, Bowen Wang, and Jiancheng Huang for their valuable discussion and suggestions. This work is supported by the National Key R&D Program of China (2022YFE0200700), the National Natural Science Foundation of China (Project No. 62006219 and No. 62072452), the Natural Science Foundation of Guangdong Province (2022A1515011579), the Regional Joint Fund of Guangdong under Grant 2021B1515120011, and the Hong Kong Innovation and Technology Fund (Project No. ITS/170/20 and ITS/241/21).

References

- [1] Florian J Bock and Stephen WG Tait. Mitochondria as multifaceted regulators of cell death. *Nature reviews Molecular cell biology*, 21(2):85–100, 2020. **1**
- [2] Jeremy G Carlton, Hannah Jones, and Ulrike S Eggert. Membrane and organelle dynamics during cell division. *Nature Reviews Molecular Cell Biology*, 21(3):151–166, 2020. **1**
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. **3**
- [4] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic convolution: Attention over convolution kernels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11030–11039, 2020. **3**
- [5] Shiyi Cheng, Sipei Fu, Yumi Mun Kim, Weiye Song, Yunzhe Li, Yujia Xue, Ji Yi, and Lei Tian. Single-cell cytometry via multiplexed fluorescence prediction by label-free reflectance microscopy. *Science advances*, 7(3):eabe0431, 2021. **2**
- [6] Eric M Christiansen, Samuel J Yang, D Michael Ando, Ashkan Javaherian, Gaia Skibinski, Scott Lipnick, Elliot Mount, Alison O’neil, Kevan Shah, Alicia K Lee, et al. In silico labeling: predicting fluorescent labels in unlabeled images. *Cell*, 173(3):792–803, 2018. **2**
- [7] Josie A Christopher, Charlotte Stadler, Claire E Martin, Marcel Morgenstern, Yanbo Pan, Cora N Betsinger, David G Rattray, Diana Mahdessian, Anne-Claude Gingras, Bettina Warscheid, et al. Subcellular proteomics. *Nature Reviews Methods Primers*, 1(1):1–24, 2021. **1**
- [8] Aidan Clark, Diego de Las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake Hechtman, Trevor Cai, Sebastian Borgeaud, et al. Unified scaling laws for routed language models. In *International Conference on Machine Learning*, pages 4057–4086. PMLR, 2022. **3**
- [9] Jan Oscar Cross-Zamirski, Elizabeth Mouchet, Guy Williams, Carola-Bibiane Schönlieb, Riku Turkki, and Yin-hai Wang. Label-free prediction of cell painting from bright-field images. *Scientific reports*, 12(1):1–13, 2022. **2**
- [10] Yongxing Dai, Xiaotong Li, Jun Liu, Zekun Tong, and Ling-Yu Duan. Generalizable person re-identification with relevance-aware mixture of experts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16145–16154, 2021. **3, 5**
- [11] Xiaohan Ding, Honghao Chen, Xiangyu Zhang, Kaiqi Huang, Jungong Han, and Guiguang Ding. Reparameterizing your optimizers rather than architectures. *arXiv preprint arXiv:2205.15242*, 2022. **3**
- [12] Xiaohan Ding, Yuchen Guo, Guiguang Ding, and Jungong Han. Acnet: Strengthening the kernel skeletons for powerful cnn via asymmetric convolution blocks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1911–1920, 2019. **3, 14**
- [13] Xiaohan Ding, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Diverse branch block: Building a convolution as an inception-like unit. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10886–10895, 2021. **3, 4, 5, 14**
- [14] Xiaohan Ding, Xiangyu Zhang, Ningning Ma, Jungong Han, Guiguang Ding, and Jian Sun. Repvgg: Making vgg-style convnets great again. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13733–13742, 2021. **3, 5, 14**
- [15] Konstantin Dmitriev and Arie E Kaufman. Learning multi-class segmentations from single-class datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9501–9511, 2019. **2, 4, 6, 13**
- [16] Xi Fang and Pingkun Yan. Multi-organ segmentation over partially labeled datasets with multi-scale feature abstraction. *IEEE Transactions on Medical Imaging*, 39(11):3619–3629, 2020. **3, 6, 13**
- [17] Shixiang Feng, Yuhang Zhou, Xiaoman Zhang, Ya Zhang, and Yanfeng Wang. Ms-kd: Multi-organ segmentation with multiple binary-labeled datasets. *arXiv preprint arXiv:2108.02559*, 2021. **2**
- [18] Mikhail Figurnov, Shakir Mohamed, and Andriy Mnih. Implicit reparameterization gradients. *Advances in neural information processing systems*, 31, 2018. **3**
- [19] Jianlong Fu, Heliang Zheng, and Tao Mei. Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4438–4446, 2017. **3**
- [20] Sam Gross, Marc’Aurelio Ranzato, and Arthur Szlam. Hard mixtures of experts for large scale weakly supervised vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6865–6873, 2017. **3**
- [21] Yuting Guo, Di Li, Siwei Zhang, Yanrui Yang, Jia-Jia Liu, Xinyu Wang, Chong Liu, Daniel E Milkie, Regan P Moore, U Serdar Tulu, et al. Visualizing intracellular organelle and cytoskeletal interactions at nanoscale resolution on millisecond timescales. *Cell*, 175(5):1430–1442, 2018. **1**
- [22] Suchin Gururangan, Mike Lewis, Ari Holtzman, Noah A Smith, and Luke Zettlemoyer. Demix layers: Disentangling domains for modular language modeling. *arXiv preprint arXiv:2108.05036*, 2021. **3**
- [23] Gabriele Gut, Markus D Herrmann, and Lucas Pelkmans. Multiplexed protein maps link subcellular organization to cellular states. *Science*, 361(6401):eaar7042, 2018. **1**
- [24] Hussein Hazimeh, Zhe Zhao, Aakanksha Chowdhery, Maheswaran Sathiamoorthy, Yihua Chen, Rahul Mazumder, Lichan Hong, and Ed Chi. Dselect-k: Differentiable selection in the mixture of experts with applications to multi-task learning. *Advances in Neural Information Processing Systems*, 34:29335–29347, 2021. **3, 5**
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1904–1916, 2015. **3**
- [26] Brian Herman. *Fluorescence microscopy*. Garland Science, 2020. **1**

- [27] Mu Hu, Junyi Feng, Jiashen Hua, Baisheng Lai, Jianqiang Huang, Xiaojin Gong, and Xian-Sheng Hua. Online convolutional re-parameterization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 568–577, 2022. 3, 4
- [28] Jaroslav Icha, Michael Weber, Jennifer C Waters, and Caren Norden. Phototoxicity in live fluorescence microscopy, and how to avoid it. *BioEssays*, 39(8):1700003, 2017. 1
- [29] Kyuseok Im, Sergey Mareninov, M Diaz, and William H Yong. An introduction to performing immunofluorescence staining. *Biobanking*, pages 299–311, 2019. 1
- [30] Fabian Isensee, Jens Petersen, Andre Klein, David Zimmerer, Paul F Jaeger, Simon Kohl, Jakob Wasserthal, Gregor Koehler, Tobias Norajitra, Sebastian Wirkert, et al. nnu-net: Self-adapting framework for u-net-based medical image segmentation. *arXiv preprint arXiv:1809.10486*, 2018. 6, 13
- [31] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991. 3, 5
- [32] YoungJu Jo, Hyungjoo Cho, Wei Sun Park, Geon Kim, DongHun Ryu, Young Seo Kim, Moosung Lee, Sangwoo Park, Mahn Jae Lee, Hosung Joo, et al. Label-free multiplexed microtomography of endogenous subcellular dynamics using generalizable deep learning. *Nature Cell Biology*, 23(12):1329–1337, 2021. 2
- [33] Mikhail E Kandel, Yuchen R He, Young Jae Lee, Taylor Hsuan-Yu Chen, Kathryn Michele Sullivan, Onur Aydin, M Taher A Saif, Hyunjoon Kong, Nahil Sobh, and Gabriel Popescu. Phase imaging with computational specificity (pics) for measuring dry mass changes in sub-cellular compartments. *Nature communications*, 11(1):1–10, 2020. 2
- [34] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [35] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 3
- [36] Wei-Hong Li, Xialei Liu, and Hakan Bilen. Learning multiple dense prediction tasks from partially annotated data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18879–18889, 2022. 2
- [37] Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 510–519, 2019. 3
- [38] Yunsheng Li, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Ye Yu, Lu Yuan, Zicheng Liu, Mei Chen, and Nuno Vasconcelos. Revisiting dynamic convolution via matrix decomposition. *arXiv preprint arXiv:2103.08756*, 2021. 3
- [39] Yanghao Li, Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang. Scale-aware trident networks for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6054–6063, 2019. 3
- [40] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 3
- [41] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 3
- [42] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8759–8768, 2018. 3
- [43] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 3
- [44] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 3
- [45] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1930–1939, 2018. 3, 5
- [46] Bryce Manifold, Shuaiqian Men, Ruoqian Hu, and Dan Fu. A versatile deep learning architecture for classification and label-free prediction of hyperspectral images. *Nature machine intelligence*, 3(4):306–315, 2021. 2
- [47] Chawin Ounkomol, Sharmishta Seshamani, Mary M Maleckar, Forrest Collman, and Gregory R Johnson. Label-free prediction of three-dimensional fluorescence images from transmitted-light microscopy. *Nature methods*, 15(11):917–920, 2018. 2, 3, 6, 8, 12, 13
- [48] Svetlana Pavlitskaya, Christian Hubschneider, Michael Weber, Ruby Moritz, Fabian Huger, Peter Schlicht, and Marius Zollner. Using mixture of expert models to gain insights into semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 342–343, 2020. 3, 5
- [49] Zhen Qin, Yicheng Cheng, Zhe Zhao, Zhe Chen, Donald Metzler, and Jingzheng Qin. Multitask mixture of sequential experts for user activity streams. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3083–3091, 2020. 3
- [50] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 3
- [51] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34:8583–8595, 2021. 3, 5
- [52] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems*, 29, 2016. 3
- [53] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017. 3, 5

- [54] Gonglei Shi, Li Xiao, Yang Chen, and S Kevin Zhou. Marginal loss and exclusion loss for partially supervised multi-organ segmentation. *Medical Image Analysis*, 70:101979, 2021. [2](#)
- [55] Guolei Sun, Thomas Probst, Danda Pani Paudel, Nikola Popović, Menelaos Kanakis, Jagruti Patel, Dengxin Dai, and Luc Van Gool. Task switching network for multi-task learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8291–8300, 2021. [2](#), [4](#), [6](#), [13](#), [15](#)
- [56] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. [3](#)
- [57] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. [3](#)
- [58] Hongyan Tang, Junning Liu, Ming Zhao, and Xudong Gong. Progressive layered extraction (ple): A novel multi-task learning (mtl) model for personalized recommendations. In *Fourteenth ACM Conference on Recommender Systems*, pages 269–278, 2020. [3](#), [5](#)
- [59] Peter J Thul, Lovisa Åkesson, Mikaela Wiking, Diana Mahdessian, Aikaterini Geladaki, Hammou Ait Blal, Tove Alm, Anna Asplund, Lars Björk, Lisa M Breckels, et al. A subcellular map of the human proteome. *Science*, 356(6340):eaal3321, 2017. [1](#)
- [60] Xintao Wang, Chao Dong, and Ying Shan. Reprsr: Training efficient vgg-style super-resolution networks with structural re-parameterization and batch normalization. *arXiv preprint arXiv:2205.05671*, 2022. [3](#)
- [61] Georg Wolff, Ronald WAL Limpens, Jessika C Zevenhoven-Dobbe, Ulrike Laugks, Shawn Zheng, Anja WM de Jong, Roman I Koning, David A Agard, Kay Grünewald, Abraham J Koster, et al. A molecular pore spans the double membrane of the coronavirus replication organelle. *Science*, 369(6509):1395–1398, 2020. [1](#)
- [62] Hao Wu, Shuchao Pang, and Arcot Sowmya. Tgnet: A task-guided network architecture for multi-organ and tumour segmentation from partially labelled datasets. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2022. [2](#), [6](#), [8](#), [12](#), [13](#), [15](#)
- [63] Lemeng Wu, Mengchen Liu, Yinpeng Chen, Dongdong Chen, Xiyang Dai, and Lu Yuan. Residual mixture of experts. *arXiv preprint arXiv:2204.09636*, 2022. [3](#)
- [64] Junde Xu, Donghao Zhou, Danruo Deng, Jingpeng Li, Cheng Chen, Xiangyun Liao, Guangyong Chen, and Pheng Ann Heng. Deep learning in cell image analysis. *Intelligent Computing*, 2022, 2022. [1](#)
- [65] Brandon Yang, Gabriel Bender, Quoc V Le, and Jiquan Ngiam. Condconv: Conditionally parameterized convolutions for efficient inference. *Advances in Neural Information Processing Systems*, 32, 2019. [3](#)
- [66] Taojiannan Yang, Sijie Zhu, Chen Chen, Shen Yan, Mi Zhang, and Andrew Willis. Mutualnet: Adaptive convnet via mutual learning from network width and resolution. In *European conference on computer vision*, pages 299–315. Springer, 2020. [3](#)
- [67] Seniha Esen Yuksel, Joseph N Wilson, and Paul D Gader. Twenty years of mixture of experts. *IEEE transactions on neural networks and learning systems*, 23(8):1177–1193, 2012. [3](#)
- [68] Sergey Zagoruyko and Nikos Komodakis. Diracnets: Training very deep neural networks without skip-connections. *arXiv preprint arXiv:1706.00388*, 2017. [3](#)
- [69] Guodong Zhang, Aleksandar Botev, and James Martens. Deep learning without shortcuts: Shaping the kernel with tailored rectifiers. *arXiv preprint arXiv:2203.08120*, 2022. [5](#)
- [70] Jianpeng Zhang, Yutong Xie, Yong Xia, and Chunhua Shen. Dodnet: Learning to segment multi-organ and tumors from multiple partially labeled datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1195–1204, 2021. [2](#), [6](#), [12](#), [13](#)
- [71] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE signal processing letters*, 23(10):1499–1503, 2016. [3](#)
- [72] Lefei Zhang, Shixiang Feng, Yu Wang, Yanfeng Wang, Ya Zhang, Xin Chen, and Qi Tian. Unsupervised ensemble distillation for multi-organ segmentation. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2022. [2](#)
- [73] Yikang Zhang, Jian Zhang, Qiang Wang, and Zhao Zhong. Dynet: Dynamic convolution for accelerating convolutional neural networks. *arXiv preprint arXiv:2004.10694*, 2020. [3](#)
- [74] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [3](#)