

# DenseNets Reloaded: Paradigm Shift Beyond ResNets and ViTs

Donghyun Kim<sup>1\*</sup>, Byeongho Heo<sup>2</sup>, and Dongyoon Han<sup>2\*</sup>

<sup>1</sup>NAVER Cloud AI, <sup>2</sup>NAVER AI Lab

**Abstract.** This paper revives Densely Connected Convolutional Networks (DenseNets) and reveals the underrated effectiveness over predominant ResNet-style architectures. We believe DenseNets’ potential was overlooked due to untouched training methods and traditional design elements not fully revealing their capabilities. Our pilot study shows dense connections through concatenation are strong, demonstrating that DenseNets can be revitalized to compete with modern architectures. We methodically refine suboptimal components - architectural adjustments, block redesign, and improved training recipes towards widening DenseNets and boosting memory efficiency while keeping concatenation shortcuts. Our models, employing simple architectural elements, ultimately surpass Swin Transformer, ConvNeXt, and DeiT-III — key architectures in the residual learning lineage. Furthermore, our models exhibit near state-of-the-art performance on ImageNet-1K, competing with the very recent models and downstream tasks, ADE20k semantic segmentation, and COCO object detection/instance segmentation. Finally, we provide empirical analyses that uncover the merits of the concatenation over additive shortcuts, steering a renewed preference towards DenseNet-style designs. Our code is available at <https://github.com/naver-ai/rdnet>.

## 1 Introduction

The “ImageNet moment” was sparked by the emergence of Convolutional Neural Networks (ConvNets), starting with the milestone AlexNet [31]. Subsequently, VGG [52] and GoogleNet [54] further highlighted the benefits of stacking multiple convolutional layers in ConvNets. In the same era, a monumental architecture ResNet [23] and its family [24, 73] stands out for introducing a groundbreaking concept - additive skip connections (also known as additive shortcuts or identity mapping [24]), which allowed for the stacking of up to 1,000 layers. The introduction of residual learning with it was a game-changer, diminishing the gradient vanishing problem by ensuring the input gradient always remained at one from the derivative of the identity mapping. This innovation sparked a series of successors, including the milestone ConvNets - EfficientNet [56] and ConvNeXt [41]; it paved the way for the next leap, such as Transformers [63], Vision Transformers (ViTs) [15], and Hierarchical ViTs [40], which accentuates the lasting influence of additive shortcuts.

---

\* Equal contribution. Correspondence to Dongyoon Han.

In the early stage of this period governed by residual learning, **Densely Connected Convolutional Networks** (DenseNets [26]) introduced a novel approach: maintaining shortcut connections through feature concatenation instead of using additive shortcuts. This led to the concept of feature reuse [26], allowing more compact models and reducing overfitting through explicit supervision propagation to the early layers. DenseNets showcased efficiency and superior performance in tasks like semantic segmentation [29]. The evolution of architectural designs post-DenseNet appeared to challenge the dominance of ResNets but saw a decline in their popularity, shaded by the advantages of additive shortcuts.

Successors of DenseNets [33,64,66] revisited DenseNets to advance its design spirit but struggled against more dominant architectural trends again. We argue the potential of DenseNets still remained underexplored due to low accessibility, being gradually hindered by outdated training methods and the limitations of low-capacity components; they struggled to keep pace with the advancements in modern architectures that benefited from years of evolutionary refinements. Furthermore, we presume DenseNets requires an overhaul due to its limited applicability and memory challenges caused by increasing feature dimensionality. While the authors addressed memory concerns [26,45], these issues continue to restrict the expansion of the architecture, particularly for width scaling. Despite the drawbacks, we conjecture the core design concept is still highly potent.

Bearing this in mind, this paper revitalizes DenseNets by highlighting the undervalued efficacy of concatenations. Through a comprehensive pilot study training with over 10k random networks across varied setups, we validate our claim that concatenation can surpass the additive shortcut. Afterward, we modernize DenseNet with a more memory-efficient design to widen it, abandoning ineffective components and enhancing architectural and block designs, while preserving the essence of dense connectivity via concatenation. We employ contemporary strategies that synergize with DenseNets as well. Our methodology eventually exceeds strong modern architectures [17, 21, 35, 38, 48, 83] and some milestones like Swin Transformer [40], ConvNeXt [41], and DeiT-III [60] in performance trade-offs on ImageNet-1K [49]. Our models demonstrate competitive performance on downstream tasks such as ADE20K semantic segmentation and COCO object detection/instance segmentation. Remarkably, our models do not exhibit slowdown or degradation as the input size increases. Ultimately, our empirical analyses shed light on the unique benefits of concatenation.

## 2 Related Work

**Densely Connected Neural Networks** (DenseNets) [26] pioneered dense connections within Convolutional Neural Networks (ConvNets) beyond additive shortcuts, highlighted by parameter efficiency and enhanced precisions. Building on this framework, several variants were proposed. PeleeNet [66] successfully proposed modifications to achieve real-time inference capabilities upon DenseNet. VovNet [33] departed from DenseNets’ dense feature reuse in favor of a sparser one-shot aggregation aimed at real-time object detection. CSPNet [64], by omit-

ting features and later concatenating them at cross-stage layers, reduces computational demands, barely affecting precision. DenseNets were further highlighted with the effectiveness on dense prediction tasks; for example, Jegou *et al.* [29] showed the effectiveness of DenseNets on semantic segmentation. MDU-Net [80] exploited the dense connectivity for enhanced biomedical image segmentation. DCCT [44] integrated dense connections into a Transformer architecture [63] to facilitate image dehazing. For video snapshot compressive imaging, EfficientSCI [65] also leveraged the benefits of dense connectivity. Wang *et al.* [69] utilized dense connections to improve the detection of small objects.

We believe these references demonstrate DenseNet-based designs’ potential, to our knowledge, but none recently challenged the ImageNet benchmarks using the principle of dense connections.

**Modern architectures.** DeiT [58] and AugReg [53] exhibited modernized training recipes [3, 56, 70] could replace massive training data for ViT [15] training. Descendant hierarchical ViTs [14, 40, 67, 68, 74, 75], which got closer to ConvNets, showed locality offers efficacy along with computational efficiency. Hybrid architectures [10, 21, 36, 62, 83] then explicitly equip convolutions for the locality. Ironically, incomers have become closer to ConvNets, aiming not to forsake the proven effectiveness of simple convolution, albeit using Transformers.

ConvNets [4, 23, 47, 56] initially predominated due to strong capability along with efficiency. Interestingly, advancements from the ViT side have also contributed to modernizing ConvNets; many recent architectures [19, 59, 61, 76] were inspired by ViT’s designs but armed with locality, demonstrating the continued high competitiveness of convolutions. Successors like RepLKNet [13] and SLaK [39] employed large-scale kernel convolutions built upon the predecessor’s legacy to emulate the globality of attention [15], offering to learn enhanced global representations. RevCol [5] introduced a new concept to mix multi-level features repeatedly through multiple columns. InceptionNeXt [77] adopted the inception module [54] inside ConvNeXt to show improved performance. HorNet [48] and MogaNet [35] both have presented remarkable performance by employing multiple gated convolution and multi-level features, respectively, which also took advantage of multi-scale features for globality.

Those architectures surpassed ViTs on ImageNet and dense prediction tasks as well, but similarities like using additive shortcuts and architectural complexity continue to restrict architectural diversity and innovation. Furthermore, network modernization methods [3, 41, 66, 70] have successfully revisited existing architectures but did not handle beyond baselines using additive shortcuts. This work follows a general direction but ensures our starts from a distinct baseline, acknowledging uncertainties about the effectiveness of existing roadmaps.

### 3 Methodology: Revitalizing DenseNets

This section starts with our conjecture that DenseNet may not fall behind modern architectures and proves it by substantiating revised DenseNet architectures. Based on our conjecture with our pilot experiments, we propose our methodology, which encompasses some modernized materials to revive DenseNet.

### 3.1 Preliminary

**Motivation.** ResNets [23] have renowned due to the simple formulation at  $l$ -th layer:  $\mathbf{X}_{l+1} = \mathbf{X}_l + f(\mathbf{X}_l \mathbf{W})$ , where the input  $\mathbf{X}_l$  and the weight  $\mathbf{W}$ . A pivotal element is the *residual connection* (*i.e.*, additive shortcut, +), which facilitates modularized architectural designs as evidenced in Swin [40], ConvNeXt [41], and ViT [15]. While DenseNets [26] follow the formulation:  $\mathbf{X}_{l+1} = [\mathbf{X}_l, f(\mathbf{X}_l \mathbf{W})]$  based on the *dense connection* through concatenation having explicit parameter efficiency. However, this formulation should constrain feature dimensionality due to memory concerns, making it challenging to scale in width.

DenseNets [26] initially outperformed ResNets [23] but failed to realize a complete paradigm shift, losing initial momentum due to their applicability. In particular, despite the efforts for memory [33,45], width scaling remains problematic for DenseNets, with wider models like DenseNet-161/-233 consuming more memory [26]. Nonetheless, inspired by the prior works [3, 41, 70] and motivated by successes in dense prediction tasks such as semantic segmentation [80], we believe DenseNets would outperform popular architectures and warrant further exploration of their potential: 1) feature concatenation merits strong capability; 2) the above concerns in DenseNets can be mitigated through a strategic design.

**Our conjecture.** *Concatenation shortcut is an effective way of increasing rank.* Consider the layer output  $f(\mathbf{X}\mathbf{W})$  with the weight  $\mathbf{W} \in \mathbb{R}^{d_{in} \times d_{out}}$  and the input  $\mathbf{X} \in \mathbb{R}^{N \times d_{in}}$  with a nonlinearity  $f$ , where we assume the number of instances  $N \gg d_{in}$ . As focusing on the matrix rank of  $f$ ,  $\text{rank}(f(\mathbf{X}\mathbf{W}))$  generally gets closer to  $d_{out}$  due to the nonlinearity when  $d_{in}$  is not that small [20]. Literature [7, 18, 20] manifested the layer  $\mathbf{W}$  with  $d_{out} > d_{in}$  offers increased representational capacity. Intriguingly, DenseNets enjoy a similar aspect because we can decompose  $\mathbf{W} = [\mathbf{W}_P, \mathbf{I}]$ , where  $\mathbf{W}_P$  and  $\mathbf{I}$  denote the weights in the building block and concatenation. We further argue that increasing rank like this frequently would be more beneficial. The output dimension of  $\mathbf{W}_P$  is called *growth rate*.

*A strategic design mitigates memory concerns.* Consider the output of stacked layers  $\mathbf{X}\mathbf{W}_1 f(\mathbf{W}_2)$ , where the weights  $\mathbf{W}_1 \in \mathbb{R}^{d_{in} \times d_r}$ ,  $\mathbf{W}_2 \in \mathbb{R}^{d_r \times d_{out}}$ ,  $d_r < d_{in}, d_{out}$ , and a nonlinearity  $f$  after  $\mathbf{W}_2$ . Likewise, the rank is likely preserved as  $d_r$  is not that small. This suggests that using intermediate dimension reducers like  $\mathbf{W}_1$  (*i.e.*, *transition layer*) may not impact the rank significantly. We argue a frequent application would effectively address memory concerns.

**Pilot study.** We conduct a pilot study to verify our conjecture by sampling over 15k networks on Tiny-ImageNet [32], where their shortcuts are either additive like ResNets [23] or concatenation in DenseNets [26]. We carefully control experiments regarding computational costs and involve diverse training setups to ensure a balanced and comprehensive comparison. Intriguingly, concatenation shortcuts all outperform additive ones with the averaged Tiny-ImageNet accuracy - **56.8±3.9 (concat)** vs. **55.9±4.1 (add)**, thereby empirically supporting our claim. Detailed results and setups are provided in §5.1.

**Table 1: ImageNet-1K performance progressions.** Beginning from the baseline - DenseNet-201 [26], we report every performance change throughout progressions. We uphold DenseNet’s principle of feature reuse through concatenation as the core of the model progression.  $+\alpha$  denotes a new element  $\alpha$  was added to each prior model. Both enhancements in efficiency or accuracy are colored in **red**, while degradations are marked in **blue**.; GR denotes the growth rate, the amount of feature concatenation [26].

Elements	Top-1 Acc (%)	Param (M)	FLOPs (G)	Lat (ms) (b1, Infer)	Lat (ms) (b128, Infer)	Lat (ms) (Train)	Mem (GB) (Train)
(a) DenseNet-201 [26]	79.7	20.0	4.3	38.4	190	131	3.9
(b) (a) + Wider & shallower	79.5 (-0.2)	21.8 (+1.8)	11.1 (+6.8)	8.5 (-29.9)	170 (-20)	85 (-46)	3.2 (-0.7)
(c) (b) + Modernized blocks	80.4 (+0.9)	12.9 (-8.9)	4.8 (-6.3)	10.4 (+ 2.9)	230 (+60)	112 (+27)	3.4 (+0.2)
(d) (c) + Channel dim $\uparrow$ (GR $\downarrow$ )	80.8 (+0.4)	19.9 (+7.0)	4.7 (-0.1)	11.8 (+ 1.4)	184 (-46)	88 (-24)	3.1 (-0.3)
(e) (d) + Trans. layers $\uparrow$ (GR $\uparrow$ )	82.3 (+1.5)	21.2 (+1.3)	5.0 (+0.3)	11.0 (+ 0.8)	183 (- 1)	90 (+ 2)	3.4 (+0.3)
(f) (e) + Patchification stem	82.4 (+0.1)	21.2 (-0.0)	4.9 (-0.1)	11.0 (- 0.0)	179 (- 6)	88 (- 2)	3.2 (-0.2)
(g) (f) + Refined Trans. layers	82.6 (+0.2)	22.4 (+1.2)	4.9 (+0.0)	13.6 (+ 2.6)	170 (- 9)	97 (+ 9)	3.1 (-0.1)
(h) (g) + Channel re-scaling	82.8 (+0.2)	23.9 (+1.5)	5.0 (+0.1)	14.0 (+ 0.4)	175 (+ 5)	99 (+ 2)	3.1 (+0.0)

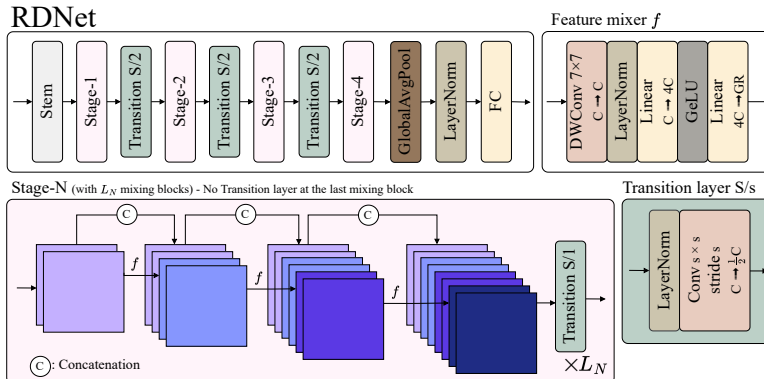
### 3.2 Revitalizing DenseNets

We revisit DenseNets while maintaining its core principle via concatenation. Our strategy explores ways to widen DenseNets and identify effective elements. Elements that contribute to the performance improvements are detailed in Table 1.

**Baseline.** As the series of revisiting ResNets [3, 70] showed, refined training recipes bring significant improvements. Likewise, we train DenseNet-201 with a modern training setup, establishing it as our baseline. Following the well-explored setups [40, 41, 57, 58, 70], we include Label Smoothing [55], RandAugment [9], Random Erasing [12, 81], Mixup [79], Cutmix [78], and Stochastic Depth [27]; we use AdamW [43] with the cosine learning rate schedule [42] and linear warmup [16] with a popular large epochs training setup (300).

**Going wider and shallower.** DenseNets originally proposed exceedingly deep architectures (*e.g.*, DenseNet-265 [26]), which effectively showed the scalability. We argue that enhancing feature dimension through a high growth rate (GR) and increasing depth is hardly achieved simultaneously under resource constraints. Prior works [20, 33, 41, 47, 66] designed shallower networks to achieve efficiency, particularly latency. Inspired by this, we modify DenseNet to a favorable baseline accordingly; widening the network by augmenting GR while diminishing its depth. Specifically, we vastly increase GR - from 32 to 120 here - to achieve it; we adjust the number of blocks per stage, being reduced from (6, 12, 48, 32) to a much smaller (3, 3, 12, 3) for a depth adjustment. We do not shrink the depth as much to maintain minimal nonlinearity. Table 1(b) shows this strategic modification has led to notable latencies and memory efficiency - around 35% and 18% decreases in training speed and memory, respectively. The marked increase in GFLOPs to 11.1 will be adjusted through the later elements. Further study supports our decision - prioritizing width while balancing depth (see Table 8a).

**Improved feature mixers.** We employ the base block [41] for our feature mixer block, which has been extensively studied to reveal its effectiveness. Before using it, we should reevaluate the studies for our case because 1) DenseNets did not use



**Fig. 1: Schematic illustration of RDNet.** RDNet features a unique design distinguishing it from ResNet-style architectures, primarily due to the use of feature concatenation. We design four stages in RDNet across all scales, where each stage- $N$  comprises  $L_N$  mixing blocks consisting of three feature mixers and one transition layer (the last mixing block does not employ the transition layer). Feature mixer  $f$  denotes our building block combines previously concatenated features to compress them into GR-dimensional features for concatenation. The growth rate (GR) adjusts the amount of concatenated features and is predetermined for each stage. Transition layers for downsampling are positioned after each stage as before. S and C denote stride and channel size. This figure illustratively sets GR to two.

additive shortcuts, and 2) the building block was originally designed to reduce dimensions successively. We find using the following setups still holds: using 1) Layer Normalization (LN) [2] instead of Batch Normalization (BN) [28]; 2) post-activation; 3) depthwise convolution [25] 4) fewer normalizations and activations; 5) a kernel size of 7. A unique aspect of our block is that the output channel (GR) is smaller than the input channel (C); mixed features are eventually more compressed features. As can be seen in Table 1(c), our design improves accuracy by a large margin (+0.9%p) while slightly increasing computational costs. We supplement factor analyses for our study here (see Table 8b).

**Larger intermediate channel dimensions.** A large input dimension for the depthwise convolution is crucial [50]. By adeptly modulating expansion ratio (ER) for inverted bottlenecks in the previous works [20, 41, 50, 56, 57] successfully achieved significant performance, by enlarging intermediate tensor size within the block beyond input dimensions (*e.g.*, ER was tuned to 6).

DenseNets similarly employed the ER concept; however, they distinctively applied it to the growth rate (GR) (*e.g.*, ER=4×GR) rather than to the input dimension to reduce both input and output dimensions. We argue that this harms the capability of encoded features through the nonlinearity [20]. Thus, we reengineer the approach by directing ER proportional to the input dimension (*i.e.*, decoupling ER from GR). This change results in increased computational demands from a larger intermediate dimension; thus, halving GR (*e.g.*, from 120 to 60) manages these demands without compromising accuracy. Namely, we enrich the features before applying nonlinearity and further compress the

channels to control computational costs. Thereafter, we achieve both a faster training speed of 21% and 0.4%p improvement in accuracy shown in Table 1. Additionally, we conduct a factor analysis to ascertain whether reducing ER and increasing GR is preferable, or conversely, elevating ER and decreasing GR; Table 8c displays employing GR of 4 ultimately yields the optimal results.

**More transition layers.** The transition layers [26] between stages are intended to reduce the number of channels. Due to the dense connections in every block, the intensified accumulation of features does not allow a high growth rate (GR). This gets worse as multiple blocks are stacked within a single stage, such as in the third stage, where numerous blocks accumulate in a single stage with low GRs. We introduce a novel aspect using more transition layers to address it. To be specific, we propose to use a transition layer in a stage, not solely after each stage, but after every three blocks with a stride of 1. These transition layers focus on dimension reduction rather than downsampling. This modification evidently reduces the computational costs substantially; therefore, we successfully increase overall GRs thanks to it<sup>1</sup>. This is further supported by the results in Table 8e, which reveals using transition layers frequently often improves accuracy.

Additionally, we note that the models exhibit low parameter counts compared to their FLOPs. We remedy this by introducing variable GR at different stages (*e.g.*, 64, 104, 128, 192) instead of a uniform GR. Our further study in Table 8d suggests that a uniform growth rate (GR) compromises both accuracy and efficiency. Finally, Table 1(e) shows our design achieves significant accuracy improvements without greatly affecting computational costs.

**Patchification stem** Recent advancements revealed the effectiveness of using image patches as inputs within a stem [41, 48, 59]. We use the identical setup of a patch size 4 with a stride 4 as the patchification (LN [2] follows). Our empirical findings suggest that employing the patchification yields a notable acceleration in computational speed without loss of precision (see Table 1(f)).

**Refined transition layers** Another role of the transition layers was downsampling, and extra average poolings to downsample were adopted. We refine the transition layers, removing the average pooling and replacing the convolution by adjusting the kernel size and stride with the stride (LN replaces BN as well). Therefore, our transition layers play two additional roles: 1) dimension reduction, as aforementioned; 2) downsampling. Placing the transition layer after each stage exhibits +0.2%p gain, barely hurting efficiency (see Table 1(g)). For the dimension reduction ratio, we reexamine the impact, previously explored in [26]; Table 8f reconfirms 0.5 is optimal; higher transition ratios degrade precision.

**Channel re-scaling.** We investigate if channel re-scaling is required due to the diverse variance of concatenated features. We examine our proposed re-scaling approach, which has a similar formulation by merging the channel layer-scale [58] and an effective squeeze-excitation network [34]. Table 1(h) indicates it achieves a slight +0.2%p improvement, albeit with very minor inefficiency.

<sup>1</sup> Increase in GR aims to address the overall low GR in the baseline at an *architecture level*, whereas the abovementioned GR decrease was to boost ER on a *block level*.

### 3.3 Revitalized DenseNet (RDNet)

We finally introduce Revitalized DenseNet (dubbed RDNet), illustrated in Fig. 1. Our final model achieves both enhanced precision and efficiency, particularly enjoying significantly faster speed (see Table 1(h) vs. Table 1(a)). RDNet model family aligns with the widely-adopted scales [23,40,41]. Our models distinctively include the Growth Rate,  $GR=(GR_1, GR_2, GR_3, GR_4)$ , and the number of the feature mixers in each stage,  $B=(B_1, B_2, B_3, B_4)$ , where we assign the number of the feature mixers per each stage being a multiple of 3 (*i.e.*,  $B_N=3L_N$ ), where  $L_N$  is the number of the mixing blocks. We summarize the configurations below:

- RDNet-T:  $GR = (64, 104, 128, 224), B = (3, 3, 12, 3)$
- RDNet-S:  $GR = (64, 128, 128, 240), B = (3, 3, 21, 6)$
- RDNet-B:  $GR = (96, 128, 168, 336), B = (3, 3, 21, 6)$
- RDNet-L:  $GR = (128, 192, 256, 360), B = (3, 3, 24, 6)$

## 4 Experiment

### 4.1 Image Classification

We evaluate our model family on ImageNet-1K [49]. Our models are trained following the training setups in Swin Transformer [40] and ConvNeXt [41] to ensure a fair comparison and not aimed to finetune the setups. The models are trained using AdamW [43] with a batch size of 512 and an initial learning rate of  $1e-4$  for 300 epochs. As aforementioned in our baseline in §3.2, we employed identical data augmentations/regularization techniques to ConvNeXt’s; EMA is not used for our training. The comprehensive details of the training recipe are detailed in Appendix. We follow the standard evaluation protocols [23,40,41].

Our superiority is first underscored as compared with those of the current top-performing architectures [17,21,38,48,83]. We visualize the trade-off plots in Fig. 2 and detail the accuracies with diverse computational costs in Table 2. Ours show very competitive results compared with state-of-the-art models. Table 2 exhibits that while our models slightly fall behind in accuracy, they significantly make up with speed metrics. For example, RDNet-S can match with other lighter models such as SMT-S or MogaNet-S. Notably, ours do not require large memory usage as we aimed but achieve further efficiency.

We further exhibit a comparison with the popular models in Table 4. Ours surpass competitors by high precision, with decent memory usage and faster speeds. We further visualize trade-offs in Fig. 3, where RDNet demonstrates competitive performance even when juxtaposed with the milestone architectures.


### 4.2 Zero-shot Image Classification

We evaluate RDNet on ImageNet-1k zero-shot performance by training CLIP [46] to verify the applicability under a different training scheme. We follow the training protocol in ConvNeXt-OpenCLIP [8] using

**Table 3: ImageNet-1K zero-shot classification results.** Ours outperforms ConvNeXt further in efficiency.

Models	Param	Top-1	Top-5
ConvNeXt-B	152	51.2	79.3
<b>RDNet-B</b>	<b>150</b>	<b>54.1</b>	<b>82.1</b>





**Fig. 2: ImageNet-1K performance trade-off among state-of-the-arts.** We provide comparative visualizations among *state-of-the-art models*, which were known for top-performing models. It turns out that RDNet is highly competitive in practice in terms of model speed and memory consumption.

**Table 2: ImageNet-1K comparison with the latest models.** Fig. 2 visualized this table. We thoroughly compare our models against the latest architectures in practical latency and memory usage to demonstrate superiority. *bn* denotes latency (ms), measured with a batch size of *n*. *Mem* denotes the memory occupation (GB) measured with a batch size of 16. Interestingly, while our models slightly lag in accuracy, they significantly compensate with superior speed metrics.


Model	Date	Param	FLOPs	Top-1	b1	b8	b16	b32	b64	b128	Mem
RDNet-T	Ours	24	5.0	82.8	7.4	13.4	24.6	45.7	88.9	175.2	4.1
HorNet-T <sub>7×7</sub> [48]	NeurIPS'2022	22	4.0	82.8	21.2	23.2	27.0	50.7	96.1	183.7	4.1
VAN-B2 [17]	CVMJ'2023	27	5.0	82.8	24.2	28.0	39.0	75.4	144.4	274.8	5.2
BiFormer-S [83]	CVPR'2023	26	4.5	83.8	51.6	50.5	50.9	86.8	167.5	197.2	8.5
NAT-T [21]	CVPR'2023	28	4.3	83.2	26.5	28.0	33.2	53.0	102.2	335.3	3.8
SMT-S [38]	ICCV'2023	21	4.7	83.7	46.9	48.2	55.8	96.0	176.5	335.3	5.3
MogaNet-S [35]	ICLR'2024	25	5.0	83.4	20.0	22.4	40.9	77.2	147.4	288.1	6.4
RDNet-S	Ours	50	8.7	83.7	11.9	21.5	39.8	74.0	144.2	289.0	5.4
HorNet-S <sub>7×7</sub> [48]	NeurIPS'2022	50	8.8	84.0	23.2	25.7	46.0	88.3	171.4	328.9	5.7
VAN-B3 [17]	CVMJ'2023	45	9.0	83.9	45.1	49.4	64.1	123.1	237.2	446.9	7.4
BiFormer-B [83]	CVPR'2023	57	9.8	84.3	60.4	67.8	85.8	161.2	311.9	584.2	12.2
NAT-S [21]	CVPR'2023	51	7.8	83.7	28.1	28.2	43.4	82.7	159.9	310.4	5.2
SMT-B [38]	ICCV'2023	32	7.7	84.3	69.3	70.9	87.1	149.6	272.5	518.7	7.8
MogaNet-B [35]	ICLR'2024	44	9.9	84.3	37.9	43.8	80.9	152.9	294.0	576.6	11.1
RDNet-B	Ours	87	15.4	84.4	11.7	32.2	61.4	116.6	233.7	471.6	6.9
HorNet-B <sub>7×7</sub> [48]	NeurIPS'2022	87	15.6	84.3	22.9	37.9	71.5	134.5	259.6	500.0	7.7
VAN-B4 [17]	CVMJ'2023	60	12.2	84.2	60.4	67.8	85.8	161.2	311.9	584.2	9.0
NAT-B [21]	CVPR'2023	90	13.7	84.3	28.3	33.5	43.4	82.7	159.9	310.4	5.2
MogaNet-L [35]	ICLR'2024	83	15.9	84.7	60.8	64.9	118.5	224.3	429.2	838.6	14.9
RDNet-L	Ours	186	34.7	84.8	15.7	63.2	121.0	233.3	460.7	933.7	10.9
MogaNet-XL [35]	ICLR'2024	181	34.5	85.1	66.3	112.3	207.5	394.0	771.9	1512.5	24.1

1.28B seen images from the aggregated set of CC3M [51], CC12M [6], and RedCaps [11]. We use the OpenCLIP codebase<sup>2</sup>.

### 4.3 Semantic Segmentation

We employ ImageNet-1K pre-trained weights to perform semantic segmentation on the ADE20K [82] dataset using UperNet [72]. We use a learning rate of 8e-5 with a weight decay of 0.03, and utilize stochastic depth rate 0.1, 0.2, and 0.3 for the RDNet-T, -S, and -B, respectively. The remainder of the training

<sup>2</sup> [https://github.com/mlfoundations/open\\_clip](https://github.com/mlfoundations/open_clip)



**Fig. 3: ImageNet-1K performance trade-off among previous milestones.** We provide comparative visualizations between previous architectures and our models. Notice that we also include speed comparisons to highlight actual differences in practice. Our models outperform the competing modern architectures revealing the potential of feature concatenation in designing networks.

**Table 4: ImageNet-1K performance comparison with milestones.** We report top-1 accuracy (%), parameter count (M), FLOPs (G), inference time (ms) for 128 images, and memory usage (GB) with a batch size of 16. All models are (pre-)trained on ImageNet-1k from scratch.

Model	Res	Param	FLOPs	Lat	Mem	Top-1	Model	Res	Param	FLOPs	Lat	Mem	Top-1
RSB-ResNet50 [70]	224 <sup>2</sup>	26	4.1	115	2.1	80.4	RSB-ResNet152 [70]	224 <sup>2</sup>	60	11.6	270	4.7	82.0
RegNetY-4GF [47]	224 <sup>2</sup>	21	4.0	128	2.7	81.5	RegNetY-16GF [47]	224 <sup>2</sup>	84	15.9	389	5.4	82.2
DeiT-S [58]	224 <sup>2</sup>	22	4.6	128	1.9	79.8	DeiT-B [58]	224 <sup>2</sup>	87	17.5	418	3.8	81.8
CoaT-Lite-S [36]	224 <sup>2</sup>	22	4.0	211	3.3	81.9	Swin-B [40]	224 <sup>2</sup>	89	15.4	445	5.4	83.5
Swin-T [40]	224 <sup>2</sup>	28	4.5	173	2.6	81.3	PVTv2-B5 [68]	224 <sup>2</sup>	82	11.8	414	7.0	83.8
PVTv2-B2-Li [68]	224 <sup>2</sup>	23	3.9	173	4.4	82.1	ConvNeXt-B [41]	224 <sup>2</sup>	89	15.4	417	5.4	83.8
FocalNet-T [74]	224 <sup>2</sup>	29	4.5	181	4.0	82.3	CSWin-B [14]	224 <sup>2</sup>	78	15.0	543	6.5	84.2
ConvNeXt-T [41]	224 <sup>2</sup>	29	4.5	150	2.7	82.1	RepLkNet-31B [13]	224 <sup>2</sup>	79	15.3	461	2.7	83.5
CSWin-T [14]	224 <sup>2</sup>	23	4.3	194	2.7	82.8	DeiT III-B [60]	224 <sup>2</sup>	87	17.5	422	4.0	83.8
DeiT III-S [60]	224 <sup>2</sup>	22	4.6	128	2.0	81.4	FocalNet-B [74]	224 <sup>2</sup>	89	15.4	476	6.1	83.9
RevCol-T [5]	224 <sup>2</sup>	30	4.5	189	2.0	82.2	RevCol-B [5]	224 <sup>2</sup>	138	16.6	653	3.5	84.1
SLaK-T [39]	224 <sup>2</sup>	30	5.0	238	3.3	82.5	SLaK-B [39]	224 <sup>2</sup>	95	17.1	558	6.9	84.0
InceptionNeXt-T [77]	224 <sup>2</sup>	28	4.2	132	3.3	82.3	InceptionNeXt-B [77]	224 <sup>2</sup>	87	14.9	405	6.1	84.0
RDNet-T	224 <sup>2</sup>	24	5.0	175	4.1	82.8	RDNet-B	224 <sup>2</sup>	87	15.4	472	6.9	84.4
RSB-ResNet101 [70]	224 <sup>2</sup>	45	7.9	190	3.9	81.5	RegNetY-32GF [47]	224 <sup>2</sup>	145	32.3	638	7.3	82.5
RegNetY-8GF [47]	224 <sup>2</sup>	39	8.0	238	4.0	82.2	NFNet-F1 [4]	320 <sup>2</sup>	133	35.5	421	5.9	84.7
NFNet-F0 [4]	224 <sup>2</sup>	71	12.4	235	3.5	83.6	DeiT III-L [60]	224 <sup>2</sup>	304	61.6	1375	10.5	84.9
CoaT-Lite-M [36]	224 <sup>2</sup>	45	9.8	396	5.5	83.6	DeiT III-L [60]	384 <sup>2</sup>	304	191.2	4586	28.1	85.8
Swin-S [40]	224 <sup>2</sup>	50	8.7	293	3.9	83.0	ConvNeXt-L [41]	224 <sup>2</sup>	198	34.4	857	8.6	84.3
PVTv2-B4 [68]	224 <sup>2</sup>	63	10.1	370	7.2	83.6	ConvNeXt-L [41]	384 <sup>2</sup>	198	101.1	2550	19.0	85.5
ConvNeXt-S [41]	224 <sup>2</sup>	50	8.7	266	4.0	83.1	RDNet-L	224 <sup>2</sup>	186	34.7	934	10.9	84.8
CSWin-S [14]	224 <sup>2</sup>	35	6.9	313	4.0	83.6	RDNet-L	384 <sup>2</sup>	186	101.9	2714	24.3	85.8
FocalNet-S [74]	224 <sup>2</sup>	50	8.7	313	4.6	83.5							
RevCol-S [5]	224 <sup>2</sup>	60	9.0	377	2.4	83.5							
SLaK-S [39]	224 <sup>2</sup>	55	9.8	372	5.0	83.8							
InceptionNeXt-S [77]	224 <sup>2</sup>	49	8.4	245	3.2	83.5							
RDNet-S	224 <sup>2</sup>	50	8.7	289	5.4	83.7							

settings follows ConvNeXt [41]. As demonstrated in Table 5, RDNet exhibits strong performance, which reveals the effectiveness on dense prediction tasks.

#### 4.4 Object Detection

We evaluate object detection performance on COCO [37] using Mask-RCNN [22]. We use a learning rate of 3e-5 with a stochastic depth rate of 0.2. The remainder of the training settings follows ConvNeXt [41] again. As demonstrated in Table 6, RDNet exhibits competitive performance.

**Table 5: ADE20K semantic segmentation results.** All trained with the unified head UperNet (160K) on ADE20K. FLOPs (G) are measured at  $512 \times 2048$  resolutions.

Architecture	Crop	Param	FLOPs	mIoU <sup>ss</sup>	mIoU <sup>ms</sup>
Swin-T [40]	512 <sup>2</sup>	60	945	44.5	46.1
ConvNeXt-T [41]	512 <sup>2</sup>	60	939	46.0	46.7
RevCol-T [5]	512 <sup>2</sup>	60	937	47.4	47.8
NAT-T [21]	512 <sup>2</sup>	58	934	47.1	48.4
RDNet-T	512 <sup>2</sup>	58	961	<b>47.6</b>	<b>48.6</b>
Swin-S [40]	512 <sup>2</sup>	81	1038	47.6	49.5
ConvNeXt-S [41]	512 <sup>2</sup>	82	1027	<b>48.7</b>	49.6
RevCol-S [5]	512 <sup>2</sup>	90	1031	47.9	49.0
NAT-S [21]	512 <sup>2</sup>	82	1010	48.0	49.5
RDNet-S	512 <sup>2</sup>	86	1040	<b>48.7</b>	<b>49.8</b>
Swin-B [40]	512 <sup>2</sup>	121	1188	48.1	49.7
ConvNeXt-B [41]	512 <sup>2</sup>	122	1170	49.1	49.9
DeiT III-B [60]	512 <sup>2</sup>	128	1283	49.3	50.2
RevCol-B [5]	512 <sup>2</sup>	122	1169	49.0	50.1
NAT-B [21]	512 <sup>2</sup>	123	1137	48.5	49.7
RDNet-B	512 <sup>2</sup>	127	1187	<b>49.6</b>	<b>50.5</b>

## 5 Discussions

### 5.1 Pilot Study - Random Network Experiments

This study aims to reveal the effectiveness of dense connections over residual connections. We train tons of random networks across various scenarios, which include 1) multiple network scales; 2) multiple types of building blocks; 3) a range of network architectural elements; and 4) different training setups.

**Parameter spaces and cost constraints.** Table 7 (left) shows our parameter spaces for three individual scales, where  $\text{RandNet}_{A,B,C}$  are trained. We diversify the search space with respect to the budgets, such as parameter count, FLOPs, and memory consumption. We expand space from  $\mathcal{C}$  to  $\mathcal{D}$  by incorporating data augmentation and further to  $\mathcal{E}$  with both data augmentation and a different optimizer [43]. Only randomly generated networks that meet the pre-defined budget are trained. We use the 90-epochs training setup [23] trained on Tiny-ImageNet [71]. For  $\mathcal{C}, \mathcal{E}$  spaces using data augmentation [9, 27, 55, 78, 79, 81], training is done for 180 epochs. Overall, the cumulative number of trained networks reach over 15k.

**RandNet architecture.** Based on [23], we stack random building blocks within the first stage. We generate random networks in the parameter space containing

**Table 6: COCO object detection and segmentation results.** We utilize MaskRCNN with 3x schedule. FLOPs (G) are calculated with image size (1280, 800). The result of Swin-T is from the official repository [1].

Backbone	Param	FLOPs	AP <sup>box</sup>	AP <sub>50</sub> <sup>box</sup>	AP <sub>75</sub> <sup>box</sup>	AP <sup>mask</sup>	AP <sub>50</sub> <sup>mask</sup>	AP <sub>75</sub> <sup>mask</sup>
PVT-S [67]	44M	304G	43.0	65.3	46.9	39.9	62.5	42.8
Swin-T [40]	48M	267G	46.0	68.1	50.3	41.6	65.1	44.9
ConvNeXt-T [41]	48M	262G	46.2	67.9	50.8	41.7	65.0	44.9
RDNet-T	43M	278G	<b>47.5</b>	68.5	52.1	<b>42.4</b>	65.6	45.7

**Table 7: Random network experiments.** We present our experimental setups (left) and results (right). Five parameter spaces guide random network generations for two distinct shortcuts. We sample random networks within each parameter space, ensuring similar computational costs. Each parameter space varies in 1) architectural elements - channel sizes, activations, normalizations, and convolution kernel sizes in  $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$ ; 2) data augmentations in  $\mathcal{D}, \mathcal{E}$ ; 3) optimizers in  $\mathcal{E}$ .  $\mathcal{D}, \mathcal{E}$  is based on the architectural space  $\mathcal{C}$ .  $[a_1, \dots, a_n]$  and  $(a, b, c)$  denote a closed interval: a list of  $n$  elements and a range of elements from  $a$  to  $b$  with a step of  $c$ , respectively. All results are averaged.

Parameter space	$\mathcal{A}$	$\mathcal{B}$	$\mathcal{C}$	Model	Skip	FLOPs	Param	Mem	Top-1 (%)
Param ( $xM$ )	$2 < x < 2.5$	$4 < x < 5$	$9 < x < 10$	RandNet $_{\mathcal{A}}$	add	$2.25 \pm 0.13$	$2.21 \pm 0.13$	$0.65 \pm 0.03$	$45.8 \pm 2.0$
FLOPs ( $xG$ )	$2 < x < 2.5$	$4 < x < 5$	$9 < x < 10$	RandNet $_{\mathcal{A}}$	concat	$2.24 \pm 0.14$	$2.24 \pm 0.13$	$0.75 \pm 0.08$	<b><math>47.5 \pm 2.1</math></b>
Depth	(3, 6, 1)	(3, 8, 1)	(3, 12, 1)	RandNet $_{\mathcal{B}}$	add	$4.53 \pm 0.27$	$4.44 \pm 0.26$	$0.78 \pm 0.05$	$50.1 \pm 2.1$
Inter. channel dim	(32, 96, 8)	(64, 128, 8)	(64, 192, 8)	RandNet $_{\mathcal{B}}$	concat	$4.53 \pm 0.29$	$4.51 \pm 0.28$	$0.90 \pm 0.11$	<b><math>51.2 \pm 2.1</math></b>
Output channel dim	(32, 96, 8)	(64, 128, 8)	(64, 192, 8)	RandNet $_{\mathcal{C}}$	add	$9.61 \pm 0.23$	$9.41 \pm 0.23$	$1.01 \pm 0.09$	$53.2 \pm 2.2$
Activations	[ReLU, SiLU, Mish, GELU, LeakyReLU]			RandNet $_{\mathcal{C}}$	concat	$9.53 \pm 0.26$	$9.44 \pm 0.26$	$1.24 \pm 0.17$	<b><math>54.3 \pm 2.0</math></b>
Normalization layers	[BatchNorm, LayerNorm]			RandNet $_{\mathcal{D}}$	add	$9.60 \pm 0.23$	$9.40 \pm 0.23$	$1.02 \pm 0.09$	$57.4 \pm 1.5$
Kernel sizes		[3, 5, 7, 9]		RandNet $_{\mathcal{D}}$	concat	$9.54 \pm 0.26$	$9.44 \pm 0.26$	$1.24 \pm 0.16$	<b><math>58.1 \pm 1.4</math></b>
Parameter space		$\mathcal{D}$	$\mathcal{E}$	RandNet $_{\mathcal{E}}$	add	$9.59 \pm 0.23$	$9.38 \pm 0.23$	$1.02 \pm 0.09$	$58.2 \pm 1.5$
Base space		$\mathcal{C}$	$\mathcal{C}$	RandNet $_{\mathcal{E}}$	concat	$9.54 \pm 0.26$	$9.44 \pm 0.26$	$1.25 \pm 0.17$	<b><math>58.9 \pm 1.6</math></b>
Optimizer		-	AdamW						
Data augmentation		✓	✓						

**Fig. 4: Cumulative probability vs. error** of trained models in Table 7 is visualized here following Radosavovic *et al.* [47]. Across all scales and settings, we observe concatenation-based models outperform those employing additive shortcuts.


diverse depths, widths, activations, normalizations, and kernel sizes to provide flexibility under constrained costs (see Table 7). Additionally, we diversify building blocks across all search spaces to conduct more extensive experiments. Three distinct architectural blocks - dubbed PreNorm, PostNorm, and PostNorm (w/o act) - are differentiated by the use of pre-activation and shortcut positions. PreNorm block adopts the pre-normalization [24, 26] precedes a skip connection. In contrast, two PostNorms enjoy post-normalization [23, 41]. PostNorm varies from PostNorm (w/o act) based on the activation function post-skip connection.

**Result interpretation.** Table 7 (right) exhibits that concatenation consistently outperforms additive shortcuts across all configurations. Furthermore, Fig. 4 demonstrates the superior capability of concatenation-based architectures.


## 5.2 Impact of Input Size on Performance

We provide compelling findings regarding versus input size. First, Fig 5 (left) shows RDNet enjoys strong adaptability to input size variations. Intriguingly, DenseNet161, even trained without strong data augmentations, still enjoys adaptability, surpassing DeiT-S trained with strong data augmentations. We attribute this to the effectiveness of dense connections.

Our finding further shows that, unlike width-oriented networks that slow with larger input sizes (due to the large intermediate tensors), our model’s



**Fig. 5: Accuracy/latency/memory vs. resolution.** RDNet enjoys resolution-robustness against various input image sizes to maintain accuracy. Furthermore, RDNet exhibits a similar latency/memory trend to ConvNeXt and Swin Transformer, maintaining minimal increase with larger images compared to DeiT-S and DenseNet161.



**Fig. 6: CKA analysis.** We compute CKA using features before passing through short-cuts (either concatenation or addition). This result suggests RDNet generates diverse features across all layers in comparison to ConvNeXt. The third column presents a direct comparison between RDNet and ConvNeXt. Overall, RDNet learns distinct features from what ConvNeXt does and more varied features.


optimized width avoids latency/memory loss. Fig 5 (middle, right) illustrates that RDNet compete with ConvNeXt and Swin Transformer, diverging from DenseNet161 [26] that gets slower and consumes more memory as image size grows. We note that larger scales (*e.g.*, -S, -B, and -L) all follow the same trend.

### 5.3 CKA analysis

We analyze the layer-specific features of RDNet compared to ConvNeXt using Centered Kernel Alignment (CKA) [30]. Fig. 6 displays RDNet learns distinct features at every layer, showcasing different patterns compared to ConvNeXt. In the third column, ConvNeXt and RDNet astonishingly learn different features when compared, highlighting the unique learning dynamics of each model.

### 5.4 Revisiting Stochastic Depth

Notably, DenseNets primitively did not employ Stochastic Depth [27] for model training due to sharing the similarity in connectivity patterns of networks. We posit that Stochastic Depth should not be overlooked; our results demonstrate a noticeable improvement when it is incorporated into our model, as illustrated in Fig. 7. We also observe that a small stochastic depth ratio affects profoundly (see Table 9).



**Fig. 7: Stochastic depth proves effective with dense connections.** It still acts as a regularizer.

**Table 8: Ablation study results** are reported here with each ImageNet-1K accuracy (%) with parameter count (M) and FLOPS (G). The best options are marked in gray.

(a) Depth/width scaling				(b) Block configuration				(c) Expansion ratio (ER)			
Depth	Param	FLOPs	Top-1	Block conf	Param	FLOPs	Top-1	ER	Param	FLOPs	Top-1
3, 3, 9, 3	23.3	5.0	82.5	(a) RDNet-T	23.9	5.0	82.8	1.0	24.4	5.0	82.1
3, 3, 12, 3	23.9	5.0	82.8	(a) + more act	23.9	5.0	81.9	2.0	23.8	5.0	82.6
3, 3, 15, 3	23.5	5.0	82.8	(a) $\leftrightarrow$ 3 $\times$ 3 dwconv	23.5	4.9	82.4	3.0	24.2	5.0	82.7
3, 3, 18, 6	50.3	8.7	83.5	(a) $\leftrightarrow$ 5 $\times$ 5 dwconv	23.7	5.0	82.5	4.0	23.9	5.0	82.8
3, 3, 21, 6	50.4	8.7	83.7	(a) $\leftrightarrow$ 9 $\times$ 9 dwconv	23.9	5.1	82.8	6.0	24.3	5.0	82.6
3, 3, 24, 6	49.9	8.7	83.6	(a) $\leftrightarrow$ 11 $\times$ 11 dwconv	24.4	5.2	82.7				
3, 3, 18, 6	89.2	15.4	84.2	(a) $\leftrightarrow$ 13 $\times$ 13 dwconv	24.8	5.3	82.7				
3, 3, 21, 6	86.2	15.4	84.4	(b) (a) $\leftrightarrow$ dwconv at last	23.6	5.0	82.3				
3, 3, 24, 6	87.4	15.4	84.2	(b) + more act/norm	23.6	5.0	81.1				
				(b) LN $\leftrightarrow$ BN	23.6	5.0	82.2				

(d) Growth rate (GR)				(e) Transition layer intervals				(f) Transition ratio			
GR	Param	FLOPs	Top-1	Interval	Param	FLOPs	Top-1	Ratio	Param	FLOPs	Top-1
90, 90, 90, 90	13.2	5.0	81.6	2	24.3	5.0	82.7	0.3	24.4	5.0	82.6
120, 120, 120, 120	23.9	8.9	83.0	3	23.9	5.0	82.8	0.4	23.9	5.0	82.6
64, 104, 128, 224	23.9	5.0	82.8	4	24.1	5.0	82.6	0.5	23.9	5.0	82.8
				6	23.7	5.0	82.1	0.6	23.8	5.0	82.5
								0.7	23.6	5.0	82.3

## 5.5 Ablation Studies

We gather all ablation studies in Table 8. Each table contains several models that are meticulously adjusted for almost equivalent computational costs with others to ensure a fair comparison of our specific focuses. Our methodology, in Section 3.2, methodically referenced each study.

**Table 9: Stochastic depth is compatible with dense connections.**

Ratio	Param	FLOPs	Top-1
0	23.9	5.0	81.6
0.05	23.9	5.0	82.5
0.10	23.9	5.0	82.6
0.15	23.9	5.0	82.8
0.20	23.9	5.0	82.6

## 6 Conclusion

In this paper, we have revisited the past success of DenseNet, which once outperformed ResNet in this era dominated by models using addition-based shortcuts, such as ResNet, ConvNeXt, and ViT. We have first rediscovered the potential of DenseNet, focusing on the underappreciated fact that DenseNet’s concatenation shortcuts surpass the expressivity of the convention of ResNet-style addition-based shortcuts through our pilot study. We then highlight the outdated training setups and classical macro-block designs that diminish DenseNet’s effectiveness against modernized architectures. By achieving our goal to widen DenseNet with modernized elements, we have proven that DenseNet’s foundational principles are competitive in achieving robust modeling performance on their own. Our models exhibit strong performance competitive to the latest modern architectures; the employment of diverse concatenated features has significantly enhanced performance in dense prediction tasks, showcasing an advantage overlooked in models utilizing addition shortcuts. We hope that our work sheds light on the advantages of using concatenations in network design, advocating for the consideration of DenseNet-style architectures alongside ResNet-style ones.

**Limitations.** Our models have been scaled to a ‘large’ level, but resource limitation prevents more extensions to upper scales such as ViT-G.

## References

1. Github repository: Swin transformer for object detection, <https://github.com/SwinTransformer/Swin-Transformer-Object-Detection>
2. Ba, J., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
3. Bello, I., Fedus, W., Du, X., Cubuk, E.D., Srinivas, A., Lin, T.Y., Shlens, J., Zoph, B.: Revisiting resnets: Improved training and scaling strategies. Conference on Neural Information Processing Systems (NeurIPS) **34**, 22614–22627 (2021)
4. Brock, A., De, S., Smith, S.L., Simonyan, K.: High-performance large-scale image recognition without normalization. In: International Conference on Machine Learning (ICML). pp. 1059–1071 (2021)
5. Cai, Y., Zhou, Y., Han, Q., Sun, J., Kong, X., Li, J., Zhang, X.: Reversible column networks. In: International Conference on Learning Representations (ICLR) (2023)
6. Changpinyo, S., Sharma, P., Ding, N., Soricut, R.: Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 3558–3568 (2021)
7. Chen, Y., Li, J., Xiao, H., Jin, X., Yan, S., Feng, J.: Dual path networks. Conference on Neural Information Processing Systems (NIPS) **30** (2017)
8. Cherti, M., Beaumont, R., Wightman, R., Wortsman, M., Ilharco, G., Gordon, C., Schuhmann, C., Schmidt, L., Jitsev, J.: Reproducible scaling laws for contrastive language-image learning. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 2818–2829 (2023)
9. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: IEEE Transactions on Computer Vision and Pattern Recognition Workshop (CVPRW). pp. 702–703 (2020)
10. Dai, Z., Liu, H., Le, Q.V., Tan, M.: Coatnet: Marrying convolution and attention for all data sizes. In: Conference on Neural Information Processing Systems (NeurIPS). pp. 3965–3977 (2021)
11. Desai, K., Kaul, G., Aysola, Z., Johnson, J.: Redcaps: Web-curated image-text data created by the people, for the people. arXiv preprint arXiv:2111.11431 (2021)
12. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. arXiv preprint arXiv:1708.04552 (2017)
13. Ding, X., Zhang, X., Zhou, Y., Han, J., Ding, G., Sun, J.: Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR) (2022)
14. Dong, X., Bao, J., Chen, D., Zhang, W., Yu, N., Yuan, L., Chen, D., Guo, B.: Cswin transformer: A general vision transformer backbone with cross-shaped windows. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR) (2022)
15. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (ICLR) (2020)
16. Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch sgd: Training imagenet in 1 hour. arXiv preprint arXiv:1706.02677 (2017)
17. Guo, M.H., Lu, C.Z., Liu, Z.N., Cheng, M.M., Hu, S.M.: Visual attention network. Computational Visual Media (CVMJ) (2023)

18. Han, D., Kim, J., Kim, J.: Deep pyramidal residual networks. In: *IEEE Transactions on Computer Vision and Pattern Recognition (CVPR)*. pp. 5927–5935 (2017)
19. Han, D., Yoo, Y., Kim, B., Heo, B.: Learning features with parameter-free layers. *arXiv preprint arXiv:2202.02777* (2022)
20. Han, D., Yun, S., Heo, B., Yoo, Y.: Rethinking channel dimensions for efficient model design. In: *IEEE Transactions on Computer Vision and Pattern Recognition (CVPR)*. pp. 732–741 (2021)
21. Hassani, A., Walton, S., Li, J., Li, S., Shi, H.: Neighborhood attention transformer. In: *IEEE Transactions on Computer Vision and Pattern Recognition (CVPR)*. pp. 6185–6194 (2023)
22. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *International Conference on Computer Vision (ICCV)* (2017)
23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Transactions on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)
24. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: *ECCV*. pp. 630–645. Springer (2016)
25. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017)
26. Huang, G., Liu, Z., Pleiss, G., Maaten, L.v.d., Weinberger, K.Q.: Convolutional networks with dense connectivity. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **44**(12), 8704–8716 (2022)
27. Huang, G., Sun, Y., Liu, Z., Sedra, D., Weinberger, K.Q.: Deep networks with stochastic depth. In: *ECCV*. pp. 646–661. Springer (2016)
28. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *International Conference on Machine Learning (ICML)*. pp. 448–456. PMLR (2015)
29. Jégou, S., Drozdal, M., Vazquez, D., Romero, A., Bengio, Y.: The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation. In: *IEEE Transactions on Computer Vision and Pattern Recognition Workshop (CVPRW)*. pp. 11–19 (2017)
30. Kornblith, S., Norouzi, M., Lee, H., Hinton, G.: Similarity of neural network representations revisited. In: *International Conference on Machine Learning (ICML)*. pp. 3519–3529 (2019)
31. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Communications of the ACM* **60**, 84 – 90 (2012)
32. Le, Y., Yang, X.S.: Tiny imagenet visual recognition challenge (2015), <https://api.semanticscholar.org/CorpusID:16664790>
33. Lee, Y., Hwang, J.w., Lee, S., Bae, Y., Park, J.: An energy and gpu-computation efficient backbone network for real-time object detection. In: *IEEE Transactions on Computer Vision and Pattern Recognition Workshop (CVPRW)*. pp. 752–760 (2019)
34. Lee, Y., Park, J.: Centermask: Real-time anchor-free instance segmentation. In: *IEEE Transactions on Computer Vision and Pattern Recognition (CVPR)* (2020)
35. Li, S., Wang, Z., Liu, Z., Tan, C., Lin, H., Wu, D., Chen, Z., Zheng, J., Li, S.Z.: Moganet: Multi-order gated aggregation network. In: *International Conference on Learning Representations (ICLR)* (2024), <https://openreview.net/forum?id=XhYWgjqCrV>



36. Li, Y., Chen, Y., Dai, X., Chen, D., Liu, M., Yuan, L., Liu, Z., Zhang, L., Vasconcelos, N.: Micronet: Improving image recognition with extremely low flops. In: International Conference on Computer Vision (ICCV). pp. 468–477 (2021)
37. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European Conference on Computer Vision (ECCV). pp. 740–755. Springer (2014)
38. Lin, W., Wu, Z., Chen, J., Huang, J., Jin, L.: Scale-aware modulation meet transformer. In: International Conference on Computer Vision (ICCV). pp. 5992–6003 (10 2023)
39. Liu, S., Chen, T., Chen, X., Chen, X., Xiao, Q., Wu, B., Pechenizkiy, M., Mocanu, D.C., Wang, Z.: More convnets in the 2020s: Scaling up kernels beyond 51x51 using sparsity. In: International Conference on Learning Representations (ICLR) (2023)
40. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: International Conference on Computer Vision (ICCV) (2021)
41. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 11976–11986 (2022)
42. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
43. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (ICLR) (2019)
44. Parihar, A.S., Java, A.: Densely connected convolutional transformer for single image dehazing. Journal of Visual Communication and Image Representation p. 103722 (2023)
45. Pleiss, G., Chen, D., Huang, G., Li, T., van der Maaten, L., Weinberger, K.Q.: Memory-efficient implementation of densenets. arXiv preprint arXiv:1707.06990 (2017)
46. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning (ICML). pp. 8748–8763. PMLR (2021)
47. Radosavovic, I., Kosaraju, R.P., Girshick, R.B., He, K., Dollár, P.: Designing network design spaces. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 10425–10433 (2020)
48. Rao, Y., Zhao, W., Tang, Y., Zhou, J., Lim, S.N., Lu, J.: Hornet: Efficient high-order spatial interactions with recursive gated convolutions. In: Conference on Neural Information Processing Systems (NeurIPS) (2022)
49. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. International Journal of Computer Vision (IJCV) **115**(3), 211–252 (2015)
50. Sandler, M., Howard, A.G., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 4510–4520 (2018)
51. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2556–2565 (2018)
52. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)

53. Steiner, A., Kolesnikov, A., Zhai, X., Wightman, R., Uszkoreit, J., Beyer, L.: How to train your vit? data, augmentation, and regularization in vision transformers. arXiv preprint arXiv:2106.10270 (2021)
54. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 1–9 (2015)
55. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. IEEE Transactions on Computer Vision and Pattern Recognition (CVPR) pp. 2818–2826 (2016)
56. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International Conference on Machine Learning (ICML). pp. 6105–6114. PMLR (2019)
57. Tan, M., Le, Q.V.: Efficientnetv2: Smaller models and faster training. In: International Conference on Machine Learning (ICML) (2021)
58. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning (ICML). pp. 10347–10357 (2021)
59. Touvron, H., Cord, M., El-Nouby, A., Bojanowski, P., Joulin, A., Synnaeve, G., Verbeek, J., J’egou, H.: Augmenting convolutional networks with attention-based aggregation. arXiv preprint arXiv:2112.13692 (2021)
60. Touvron, H., Cord, M., J’egou, H.: Deit iii: Revenge of the vit. In: European Conference on Computer Vision (ECCV) (2022)
61. Trockman, A., Kolter, J.Z.: Patches are all you need? arXiv preprint arXiv:2201.09792 (2022)
62. Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A.C., Li, Y.: Maxvit: Multi-axis vision transformer. In: European Conference on Computer Vision (ECCV) (2022)
63. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Conference on Neural Information Processing Systems (NIPS) (2017)
64. Wang, C.Y., Liao, H.y., Wu, Y.H., Chen, P.Y., Hsieh, J.W., Yeh, I.H.: Cspnet: A new backbone that can enhance learning capability of cnn. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 1571–1580 (2020)
65. Wang, L., Cao, M., Yuan, X.: Efficientsci: Densely connected network with space-time factorization for large-scale video snapshot compressive imaging. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 18477–18486 (2023)
66. Wang, R.J., Li, X., Ling, C.X.: Pelee: a real-time object detection system on mobile devices. In: Conference on Neural Information Processing Systems (NeurIPS). p. 1967–1976 (2018)
67. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: International Conference on Computer Vision (ICCV). pp. 548–558 (2021)
68. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pvtv2: Improved baselines with pyramid vision transformer. Computational Visual Media (CVMJ) (2022)
69. Wang, Z., Xie, K., Zhang, X.Y., Chen, H.Q., Wen, C., He, J.: Small-object detection based on yolo and dense block via image super-resolution. IEEE Access **9**, 56416–56429 (2021)

70. Wightman, R., Touvron, H., Jégou, H.: Resnet strikes back: An improved training procedure in timm. <https://github.com/huggingface/pytorch-image-models> (2021)
71. Wu, K., Zhang, J., Peng, H., Liu, M., Xiao, B., Fu, J., Yuan, L.: Tinyvit: Fast pretraining distillation for small vision transformers. In: European Conference on Computer Vision (ECCV) (2022)
72. Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: European Conference on Computer Vision (ECCV). Springer (2018)
73. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 1492–1500 (2017)
74. Yang, J., Li, C., Dai, X., Gao, J.: Focal modulation networks. In: Conference on Neural Information Processing Systems (NeurIPS) (2022)
75. Yang, J., Li, C., Zhang, P., Dai, X., Xiao, B., Yuan, L., Gao, J.: Focal self-attention for local-global interactions in vision transformers. In: Conference on Neural Information Processing Systems (NeurIPS) (2021)
76. Yu, W., Luo, M., Zhou, P., Si, C., Zhou, Y., Wang, X., Feng, J., Yan, S.: Metaformer is actually what you need for vision. In: IEEE Transactions on Computer Vision and Pattern Recognition (CVPR). pp. 10819–10829 (2022)
77. Yu, W., Zhou, P., Yan, S., Wang, X.: Inceptionnext: When inception meets convnext. arXiv preprint arXiv:2303.16900 (2023)
78. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: International Conference on Computer Vision (ICCV). pp. 6023–6032 (2019)
79. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (ICLR) (2018)
80. Zhang, J., Jin, Y., Xu, J., Xu, X., Zhang, Y.: Mdu-net: Multi-scale densely connected u-net for biomedical image segmentation. Health Information Science and Systems (HISS) (2023)
81. Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y.: Random erasing data augmentation. In: AAAI Conference on Artificial Intelligence (AAAI). pp. 13001–13008 (2020)
82. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. International Journal of Computer Vision (IJCV) **127**, 302–321 (2018)
83. Zhu, L., Wang, X., Ke, Z., Zhang, W., Lau, R.: Biformer: Vision transformer with bi-level routing attention. IEEE Transactions on Computer Vision and Pattern Recognition (CVPR) (2023)