



# 딥러닝 모델 엑기스 추출 Knowledge Distillation

Machine Learning & Visual Computing Lab  
김유민




# CONTEXT

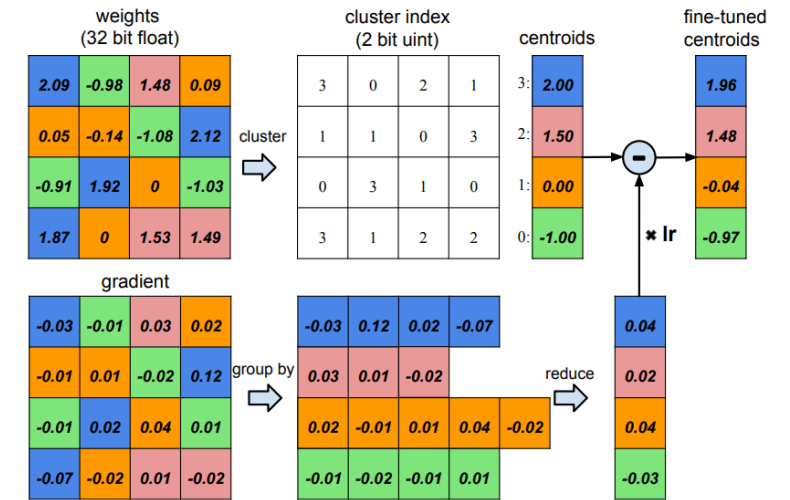
- Compression Methods
- Distillation
- Papers
- Summary

# COMPRESSION METHODS


- Pruning
- Quantization
- Efficient Design
- Distillation
- Etc...




<Weight Pruning<sub>[1]</sub>>



<Weight Quantization<sub>[2]</sub>>



<Depthwise separable convolution<sub>[3]</sub>>



<Big-little net architecture<sub>[4]</sub>>

[1] Song Han, Jeff Pool, John Tran and Willan J. Dally. Learning both Weights and Connections for Efficient Neural Networks. In *NIPS*, 2015.

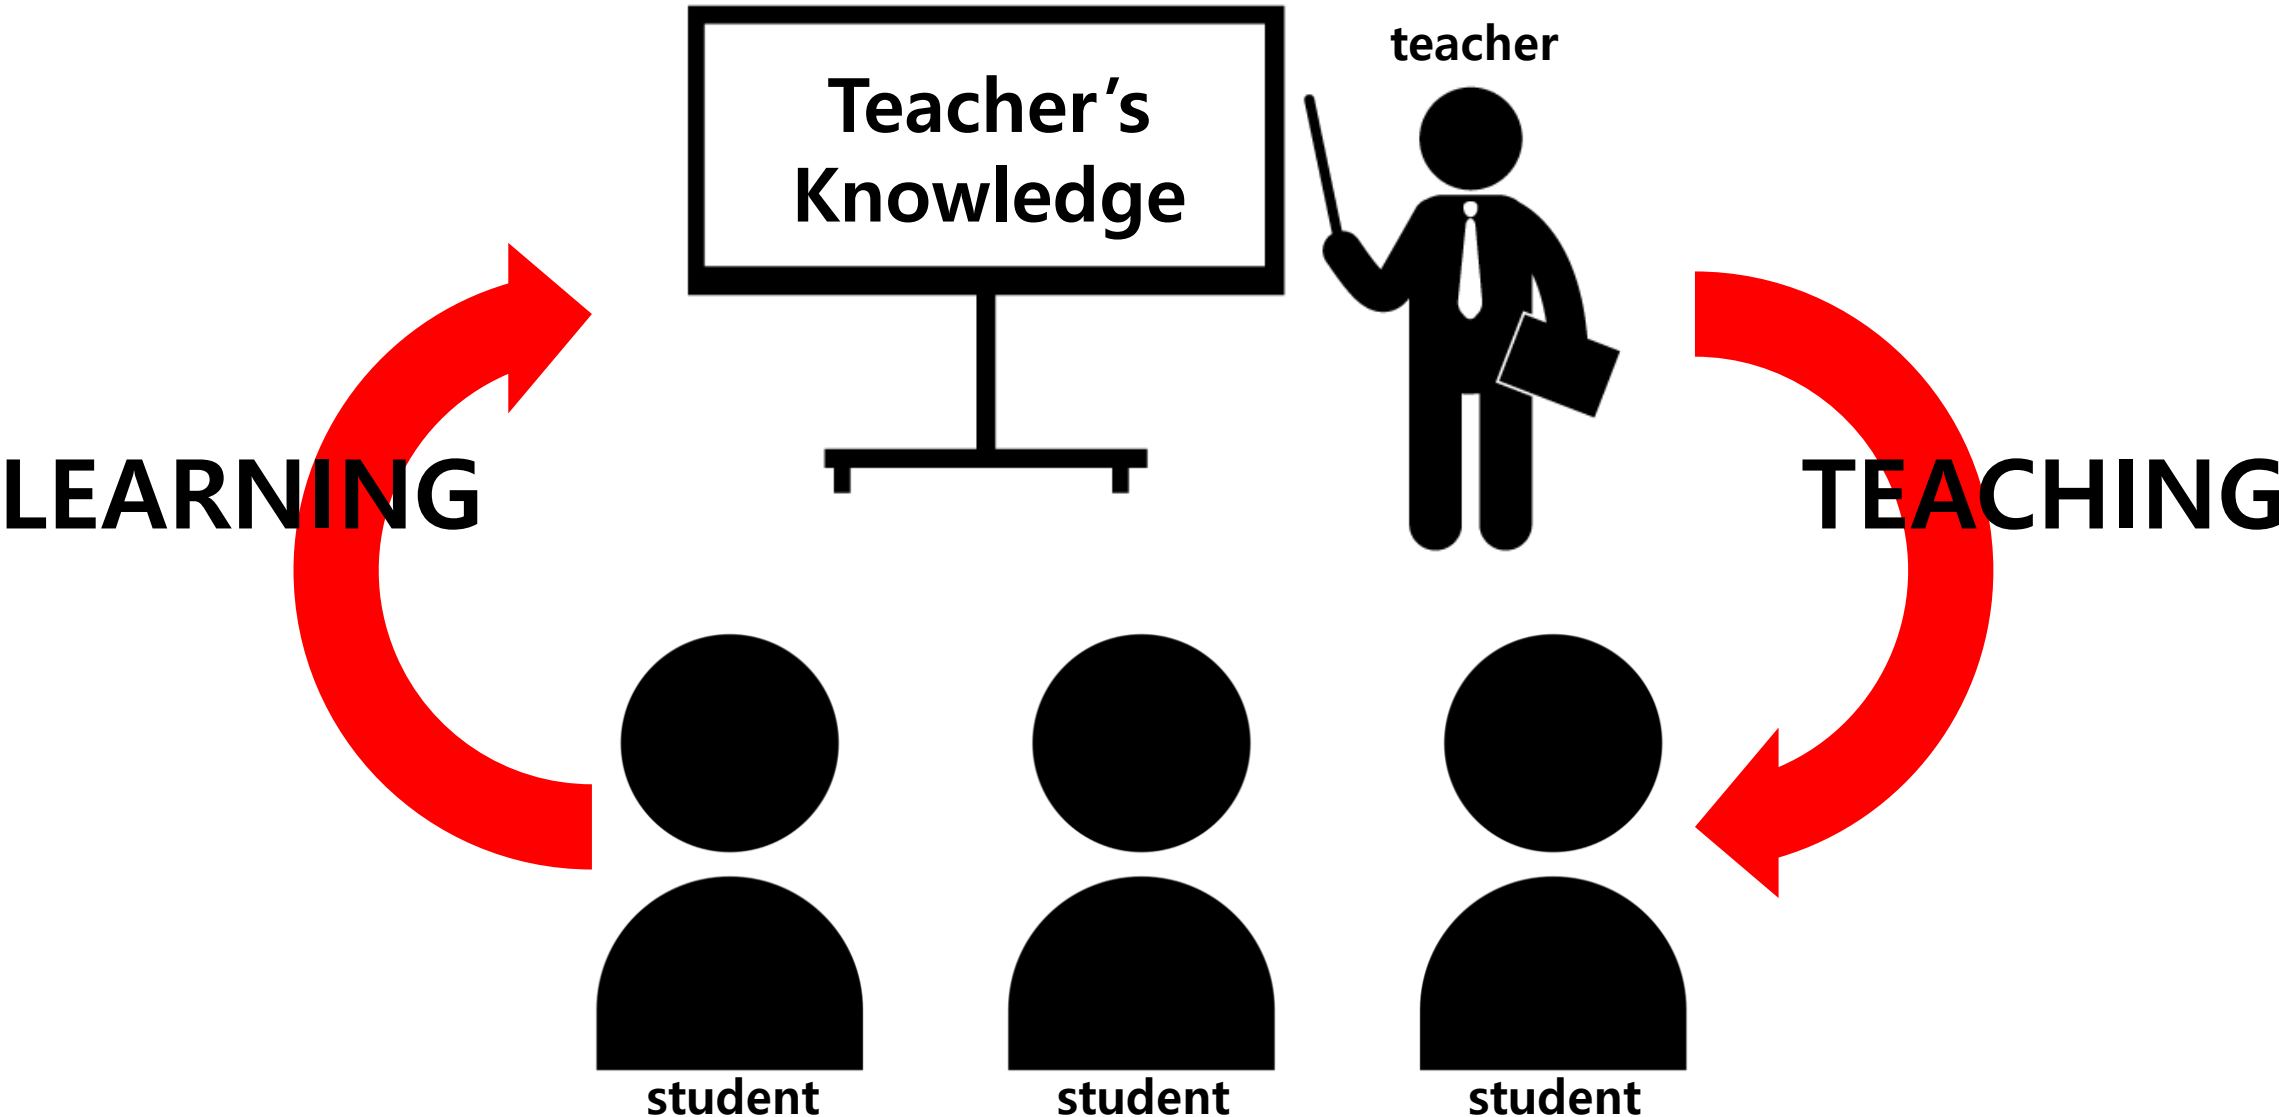
[2] Song Han, Huizi Mao and William J. Dally. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. In *ICLR*, 2016.

[3] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto and Hartwig Adam. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. In *arXiv*, 2017.


[4] Chun-Fu(Richard) Chen et al. Big-Little Net: an Efficient Multi-Scale Feature Representation for visual and speech recognition. In *ICLR*, 2019.

# **DISTILLATION**

# Teacher-Student Relation

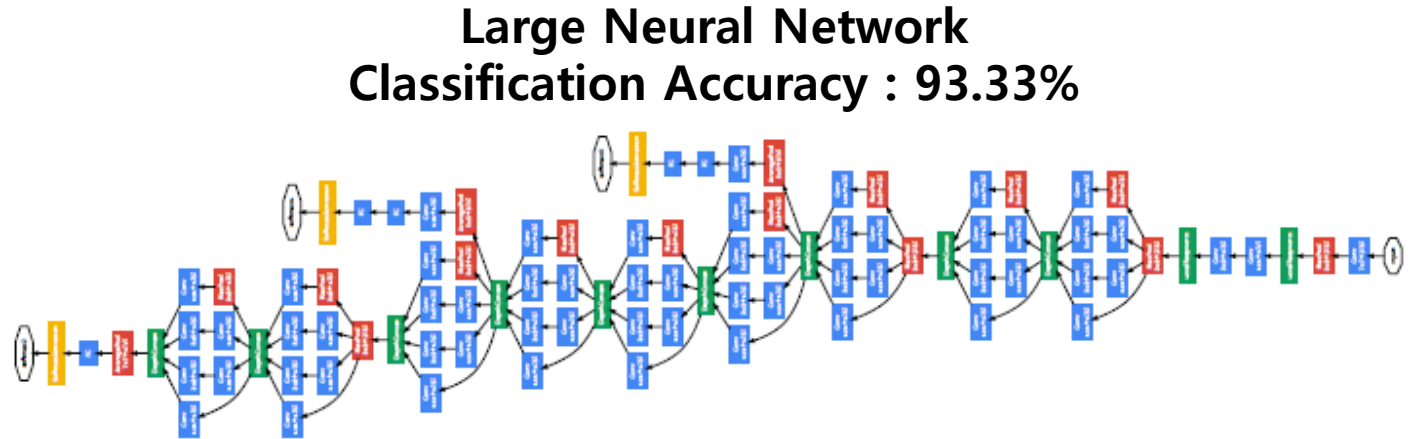


# Teacher-Student Relation in Deep Neural Network




teacher

=




**Large Neural Network**  
**Classification Accuracy : 93.33%**

<Teacher Network>



student

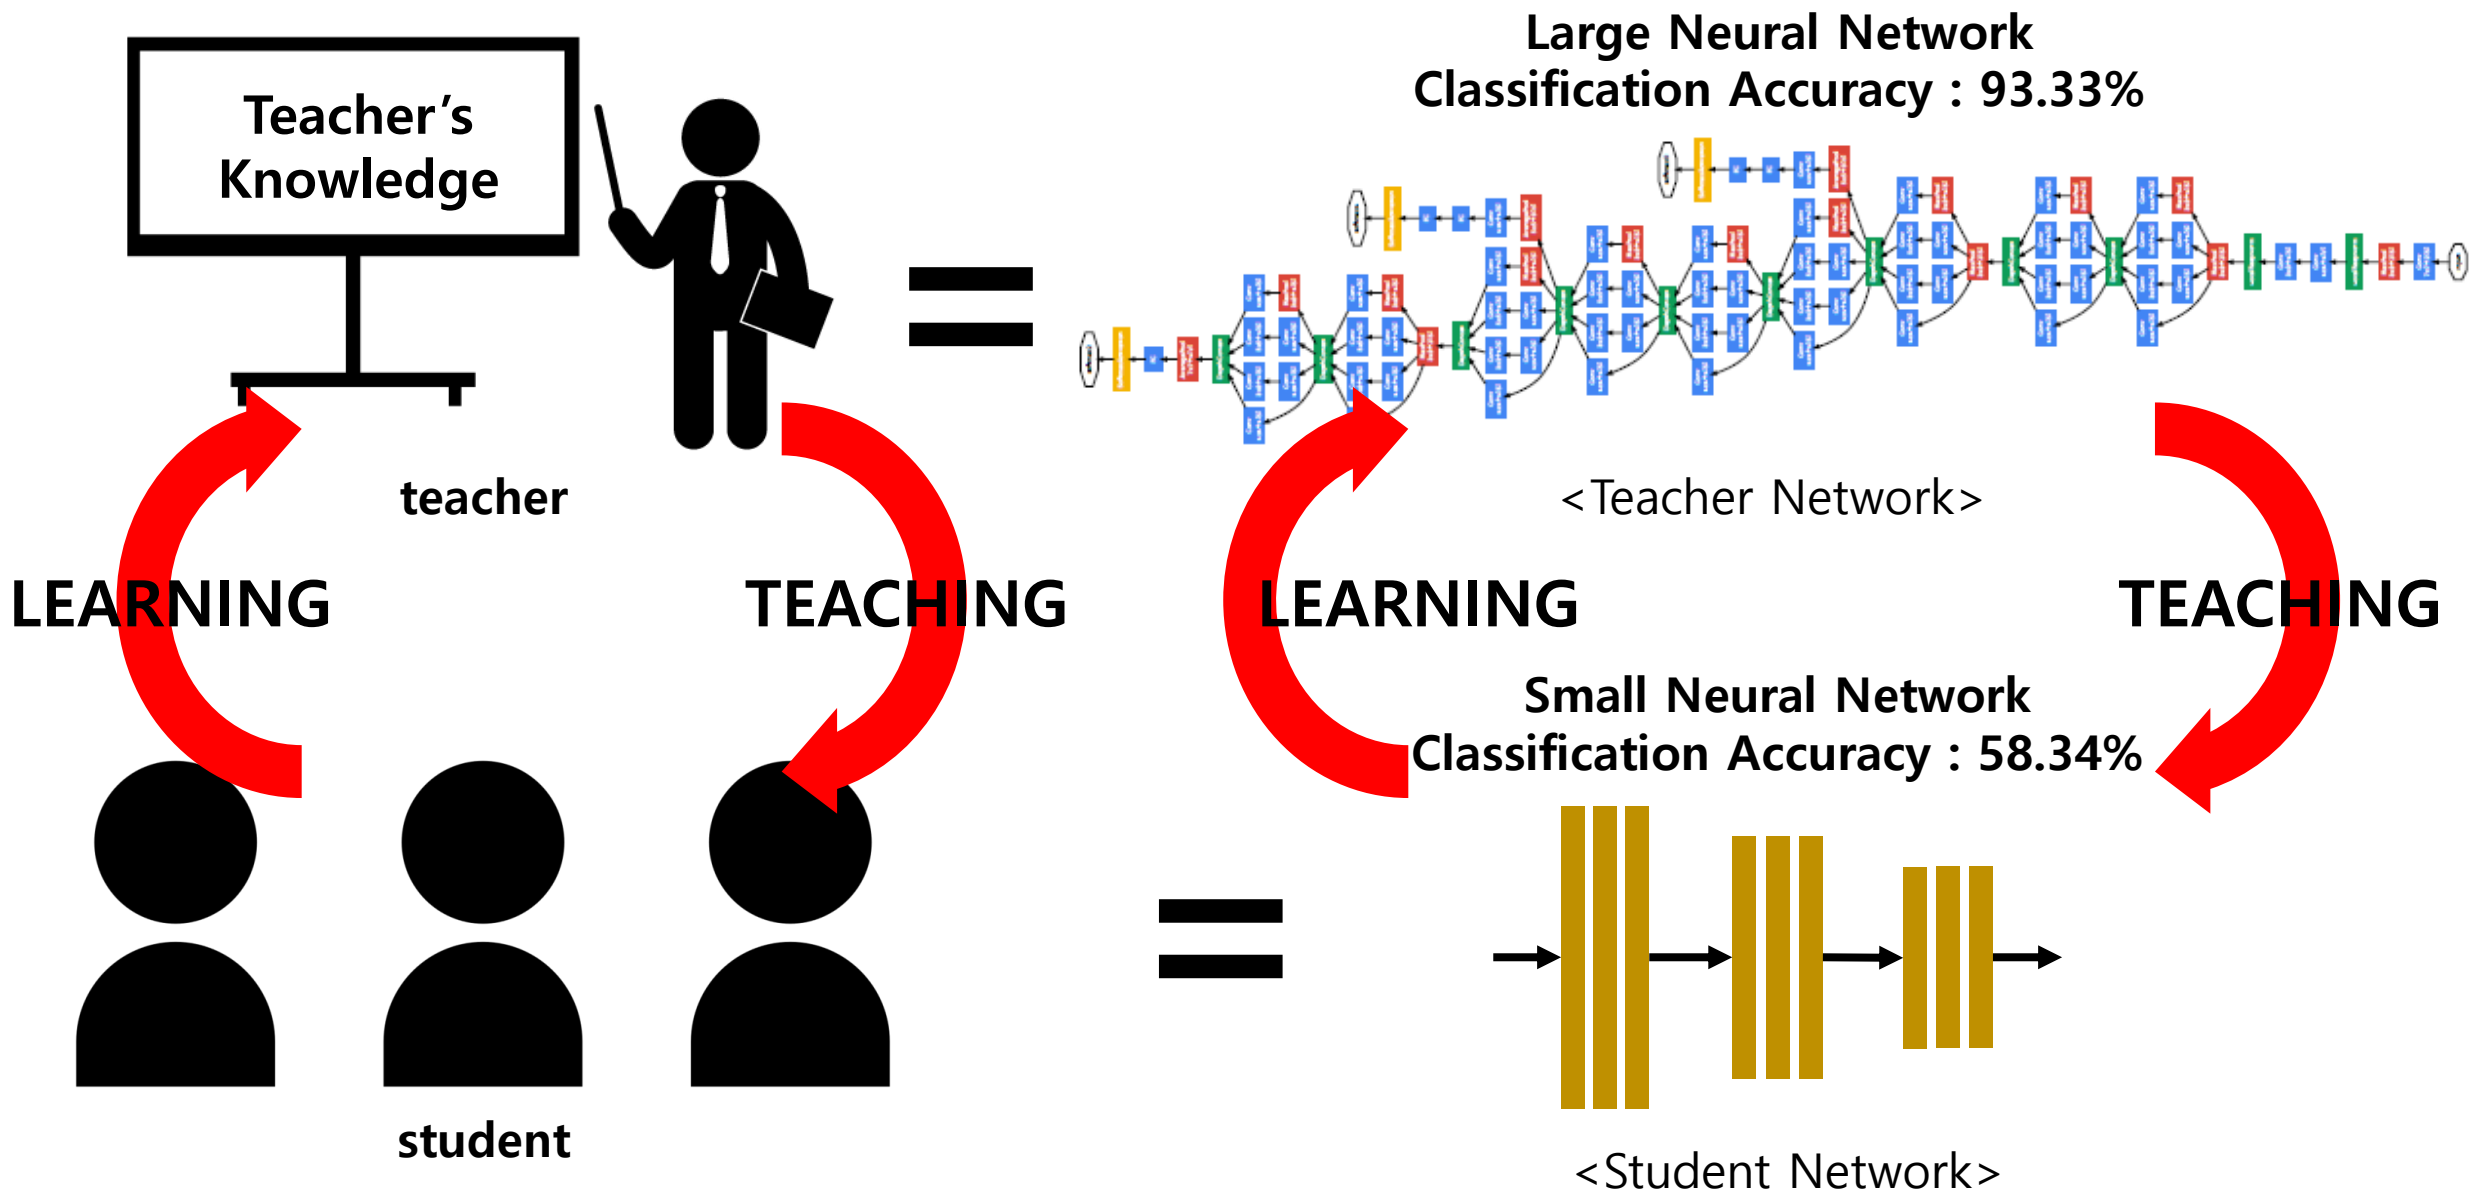
=



**Small Neural Network**  
**Classification Accuracy : 58.34%**

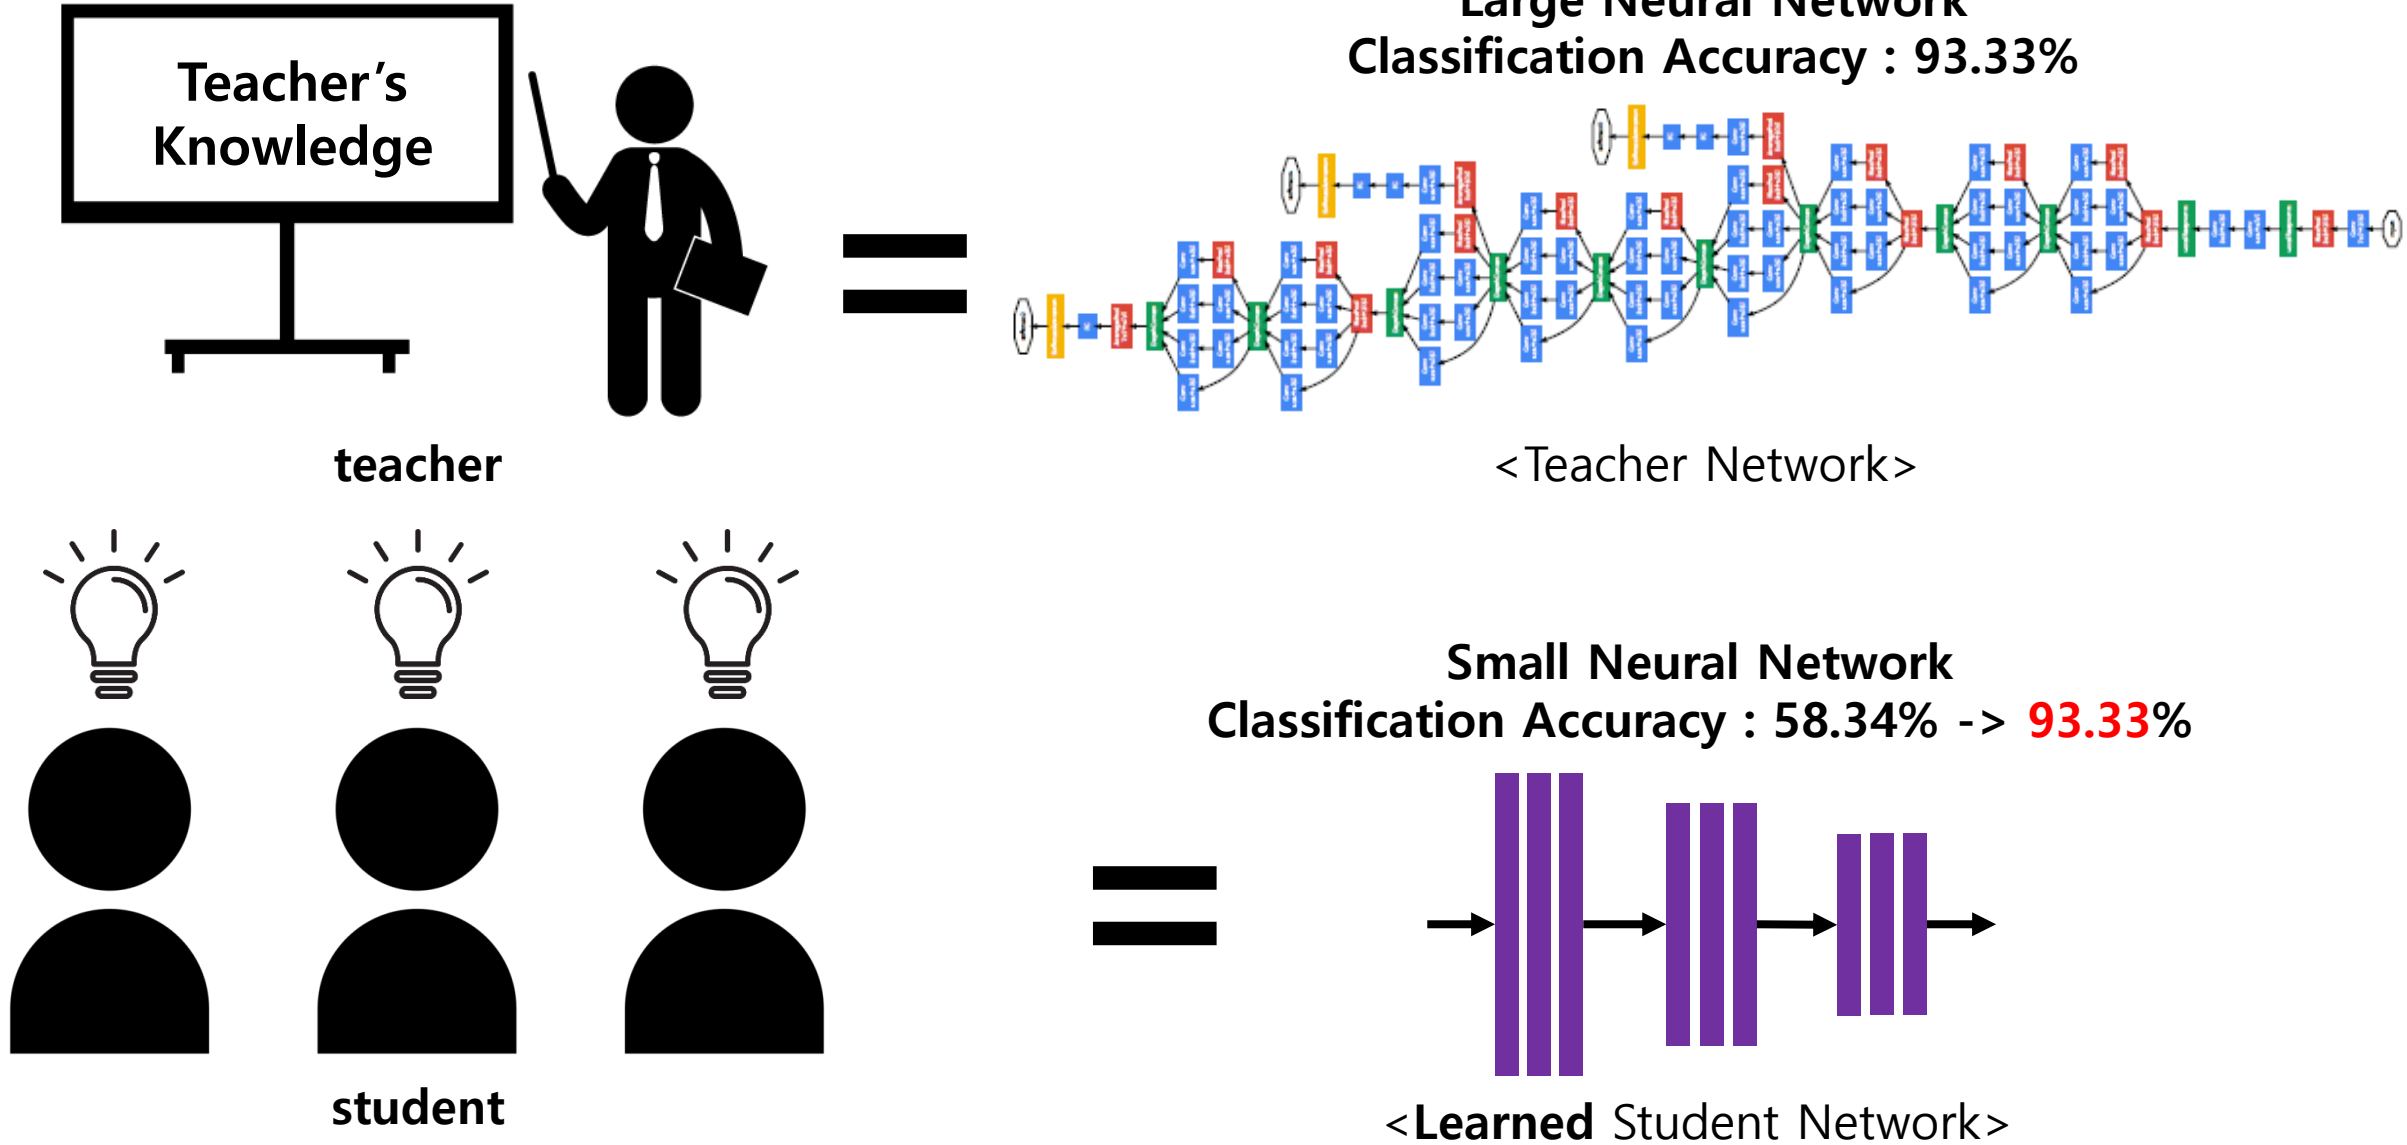
<Student Network>

# Teacher-Student Relation in Deep Neural Network





# Teacher-Student Relation in Deep Neural Network




# Do Deep Nets Really Need to be Deep?


# Do Deep Nets Really Need to be Deep?. In *NIPS*, 2014.

- Lei Jimmy Ba and Rich Caruana.

$n - 1_{th}$  Layer →




$n_{th}$  Layer(=last layer) →




$$\text{Softmax}(\hat{Y}_i = \frac{e^{z_i}}{\sum_{i=1}^n e^{z_i}})$$

Final output softmax distribution →



Cross-entropy Loss :  $L(Y, \hat{Y}) = -\sum_i Y_i \log \hat{Y}_i$

True label →



Propagation Direction


<Last section of DNN>

# Do Deep Nets Really Need to be Deep?. In *NIPS*, 2014.

- Lei Jimmy Ba and Rich Caruana.

$n - 1_{th}$  Layer →

$n_{th}$  Layer(=last layer)  
=> Logits




Propagation Direction




$$\text{Softmax}(\hat{Y}_i = \frac{e^{z_i}}{\sum_{i=1}^n e^{z_i}})$$

Final output softmax distribution



Cross-entropy Loss :  $L(Y, \hat{Y}) = -\sum_i Y_i \log \hat{Y}_i$

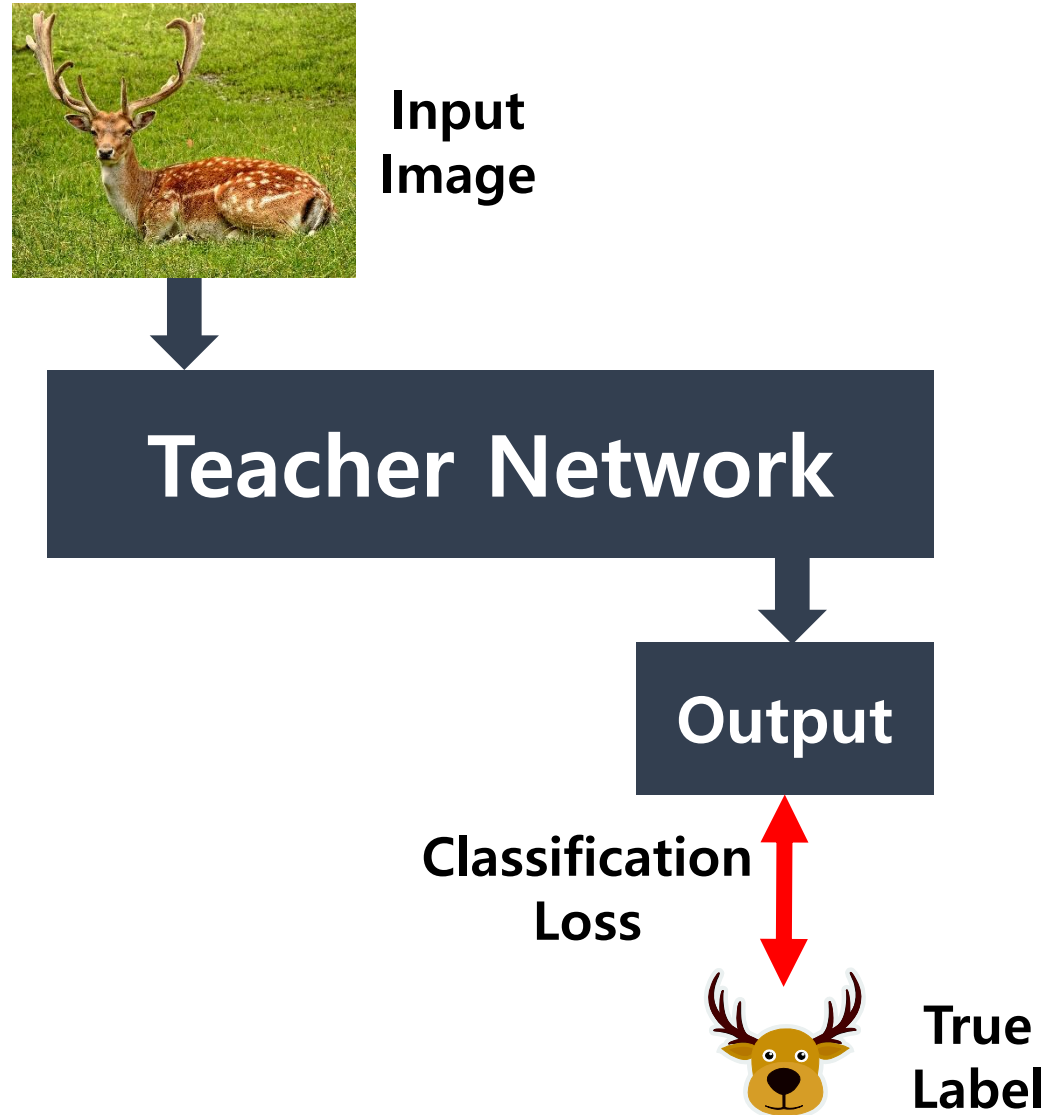
True label



<Last section of DNN>

# Do Deep Nets Really Need to be Deep?. In *NIPS*, 2014.

- Lei Jimmy Ba and Rich Caruana.



<1<sup>st</sup> step. Training Teacher Network>


# Do Deep Nets Really Need to be Deep?. In *NIPS*, 2014.

- Lei Jimmy Ba and Rich Caruana.

**Freezing all weights(Pre-trained)**



**L2 Loss with each Logits**



**Classification Loss**



Cat: 32  
Tiger: 24  
Dog: 15  
Lion: 12  
:  
:  
<Logits>

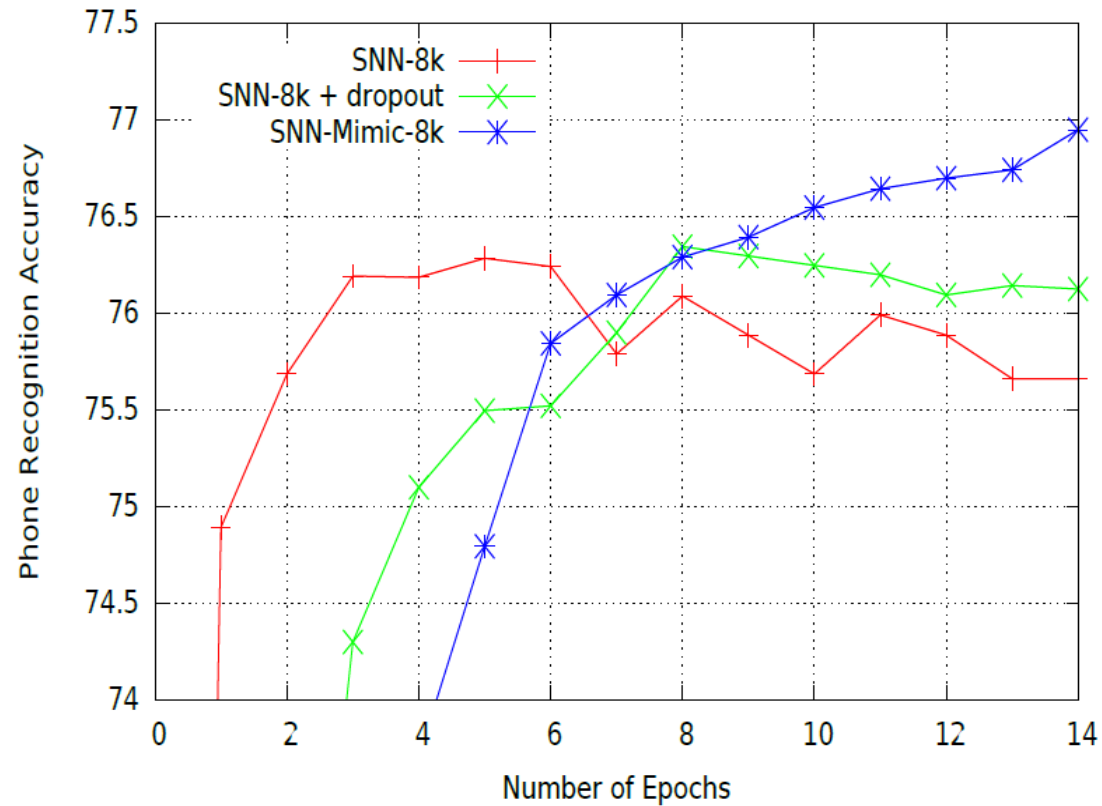


Cat: 84%  
Tiger: 7%  
Dog: 3%  
Lion: 2%  
:  
:  
<softmax probability>

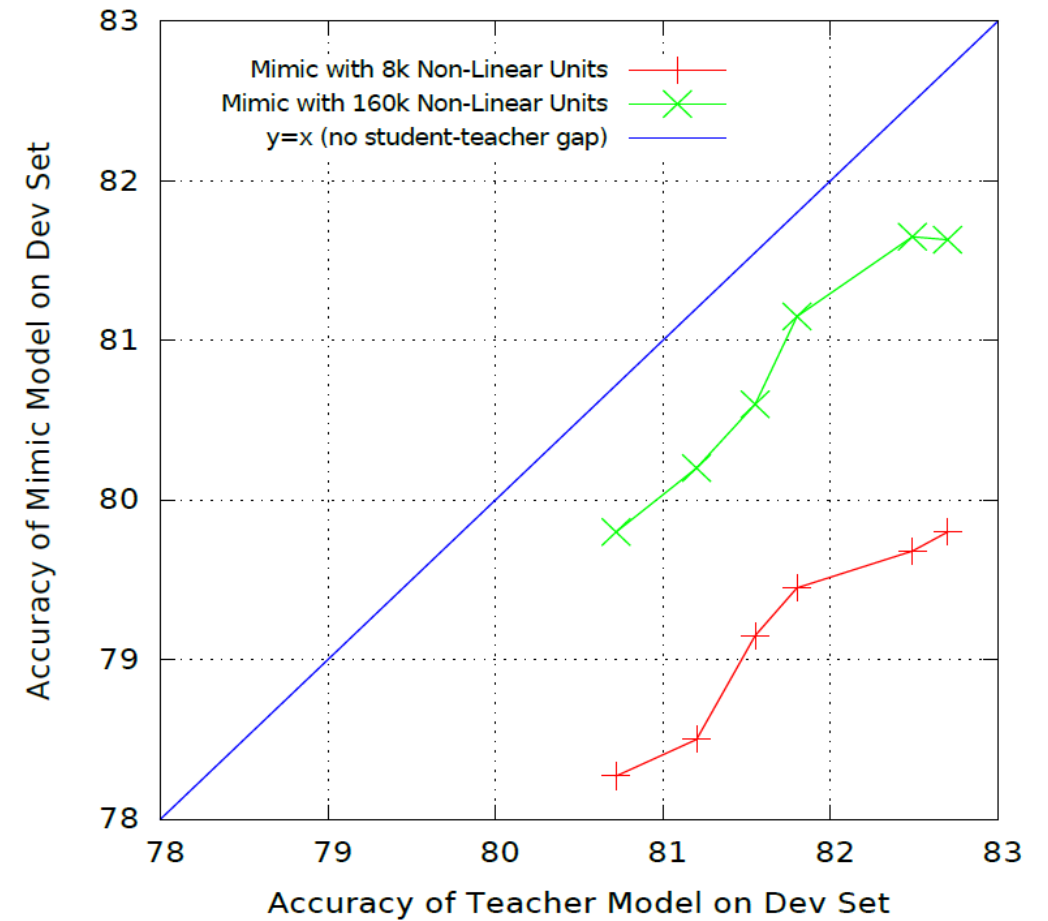
<2<sup>nd</sup> step. Training Student Network>

# Do Deep Nets Really Need to be Deep?. In *NIPS*, 2014.

- Lei Jimmy Ba and Rich Caruana.



<Student doesn't overfit>



<Better teacher, Better student>


# Distilling the Knowledge in a Neural Network



# Distilling the knowledge in a Neural Network. In *NIPS workshop*, 2014.

– Geoffrey Hinton, Oriol Vinyals and Jeff Dean.


$n - 1_{th}$  Layer →



$n_{th}$  Layer(=last layer)  
=> Logits →



$$\text{Softmax}(\hat{Y}_i = \frac{e^{z_i}}{\sum_{i=1}^n e^{z_i}})$$

Final output  
softmax distribution →



Cross-entropy Loss :  $L(Y, \hat{Y}) = -\sum_i Y_i \log \hat{Y}_i$

True label →




<Last section of DNN>

# Distilling the knowledge in a Neural Network. In *NIPS workshop*, 2014.

– Geoffrey Hinton, Oriol Vinyals and Jeff Dean.


$n - 1_{th}$  Layer →

~~$n_{th}$  Layer (=last layer)  
=> Logits~~




$$\text{Softmax}(\hat{Y}_i = \frac{e^{z_i}}{\sum_{i=1}^n e^{z_i}})$$

Final output softmax distribution



Cross-entropy Loss :  $L(Y, \hat{Y}) = -\sum_i Y_i \log \hat{Y}_i$

True label



Propagation Direction

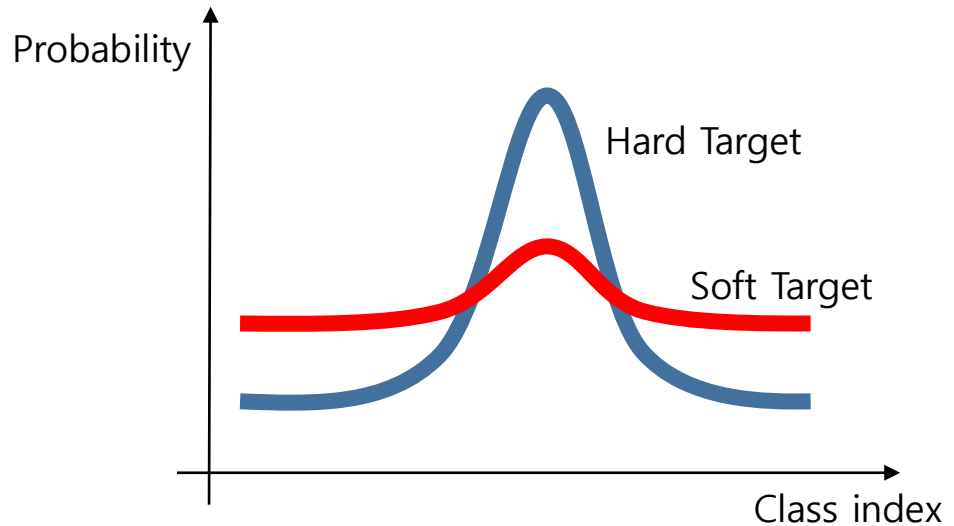
<Last section of DNN>

# Distilling the knowledge in a Neural Network. In *NIPS workshop*, 2014.


– Geoffrey Hinton, Oriol Vinyals and Jeff Dean.

$$\hat{Y}_i = \frac{e^{z_i/T}}{\sum_{i=1}^n e^{z_i/T}}$$

<Soft target function>




<Soften the hard target distribution>




$$\text{Softmax}(\hat{Y}_i = \frac{e^{z_i}}{\sum_{i=1}^n e^{z_i}})$$



**HARD TARGET**



$$\text{Soft Target Function}(\hat{Y}_i = \frac{e^{z_i/T}}{\sum_{i=1}^n e^{z_i/T}})$$

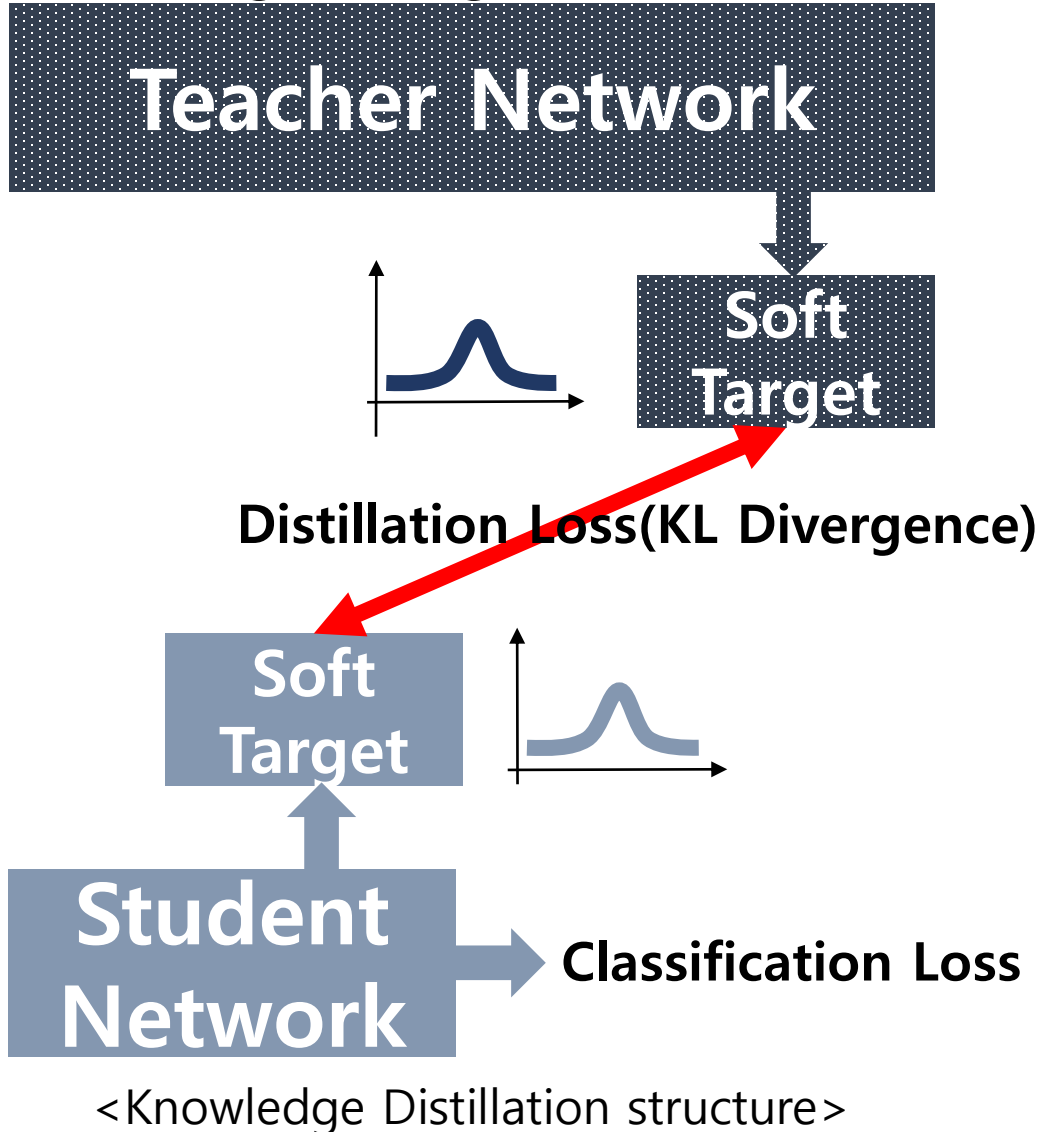


**SOFT TARGET => KNOWLEDGE**

# Distilling the knowledge in a Neural Network. In *NIPS workshop*, 2014.

– Geoffrey Hinton, Oriol Vinyals and Jeff Dean.

## Freezing all weights(Pre-trained)



Cat: 74%  
Tiger: 12%  
Dog: 7%  
Lion: 2%  
⋮



Cat: 40%  
Tiger: 33%  
Dog: 10%  
Lion: 5%  
⋮



Cat: 2%  
Tiger: 84%  
Dog: 5%  
Lion: 3%  
⋮



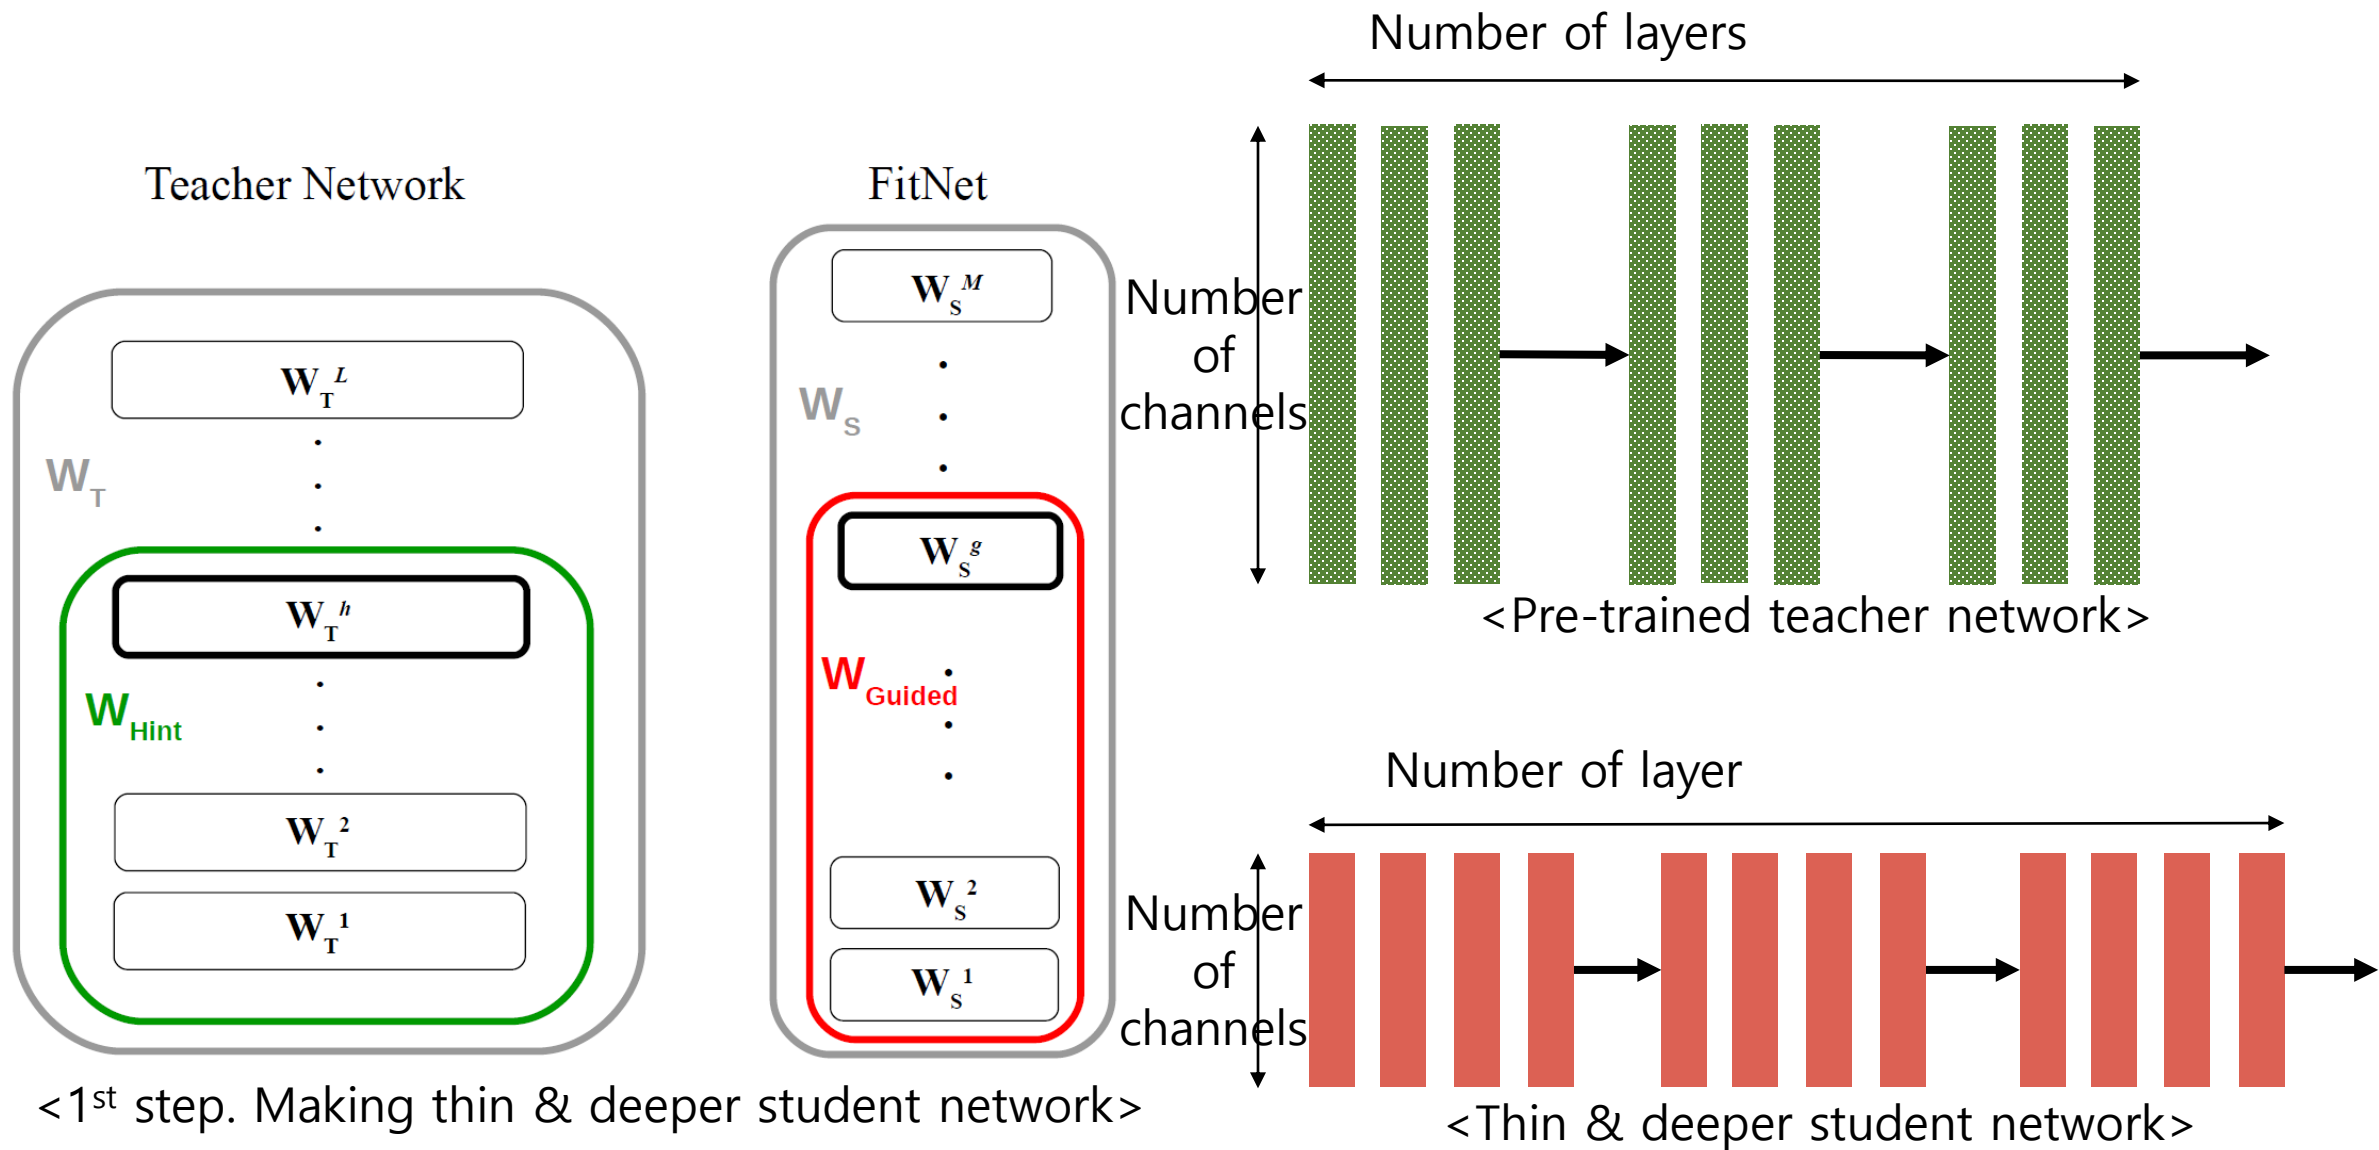
Cat: 7%  
Tiger: 50%  
Dog: 35%  
Lion: 3%  
⋮

<Hard target to soft target distribution>

# **FitNets: Hints for Thin Deep Nets**

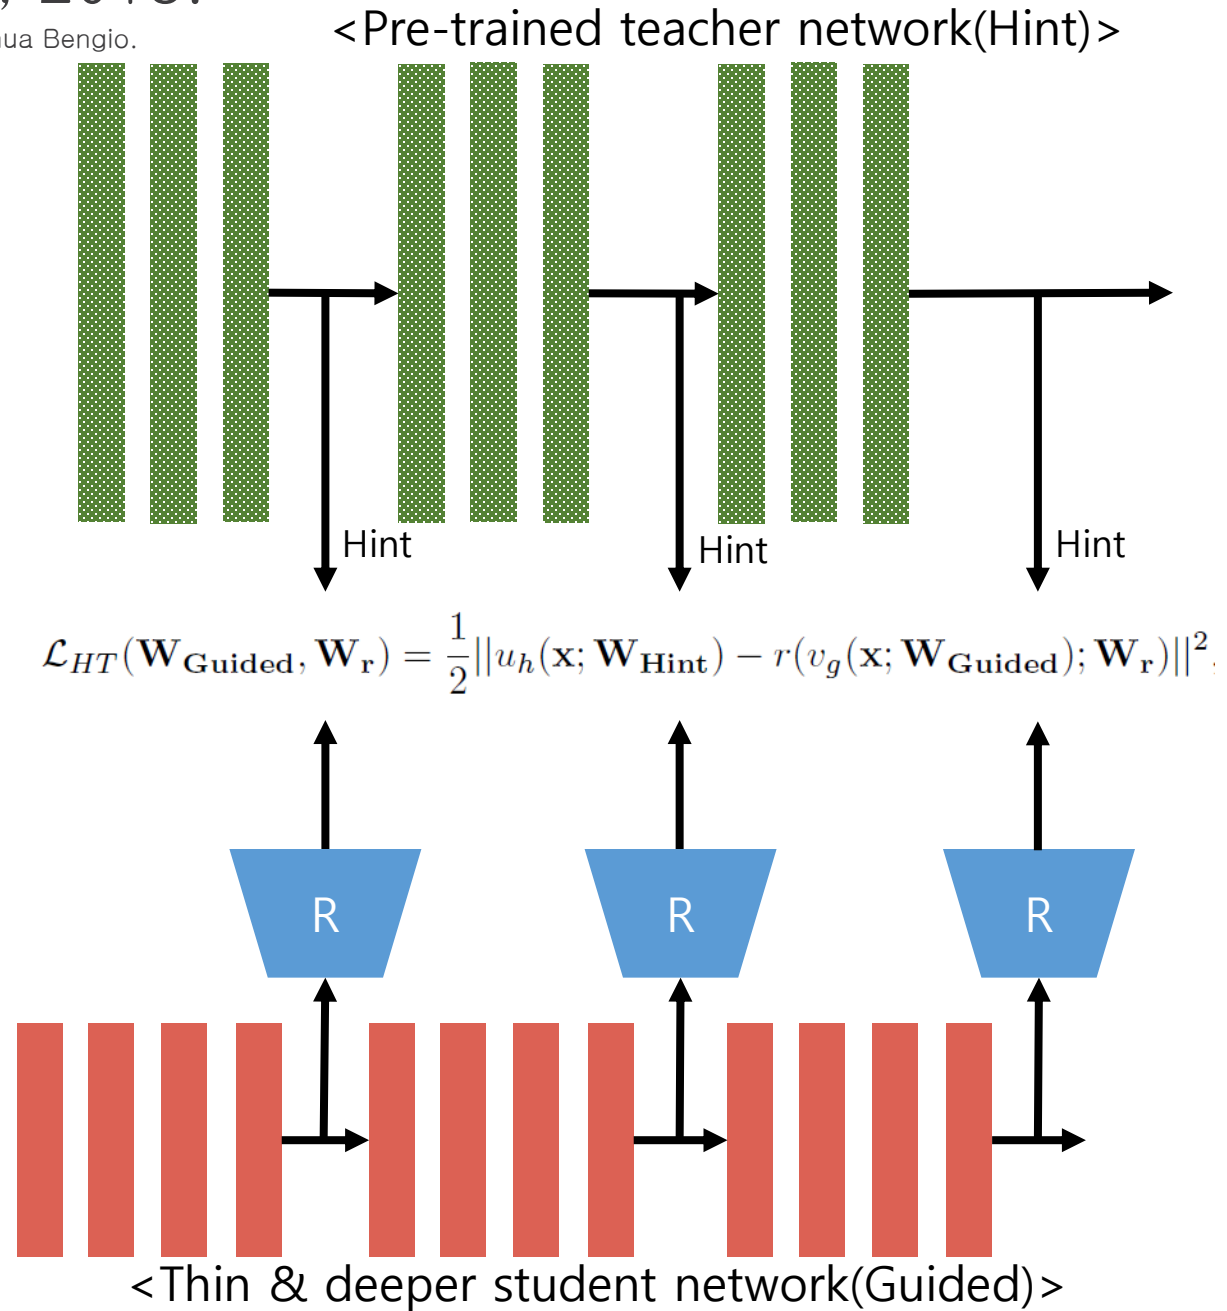
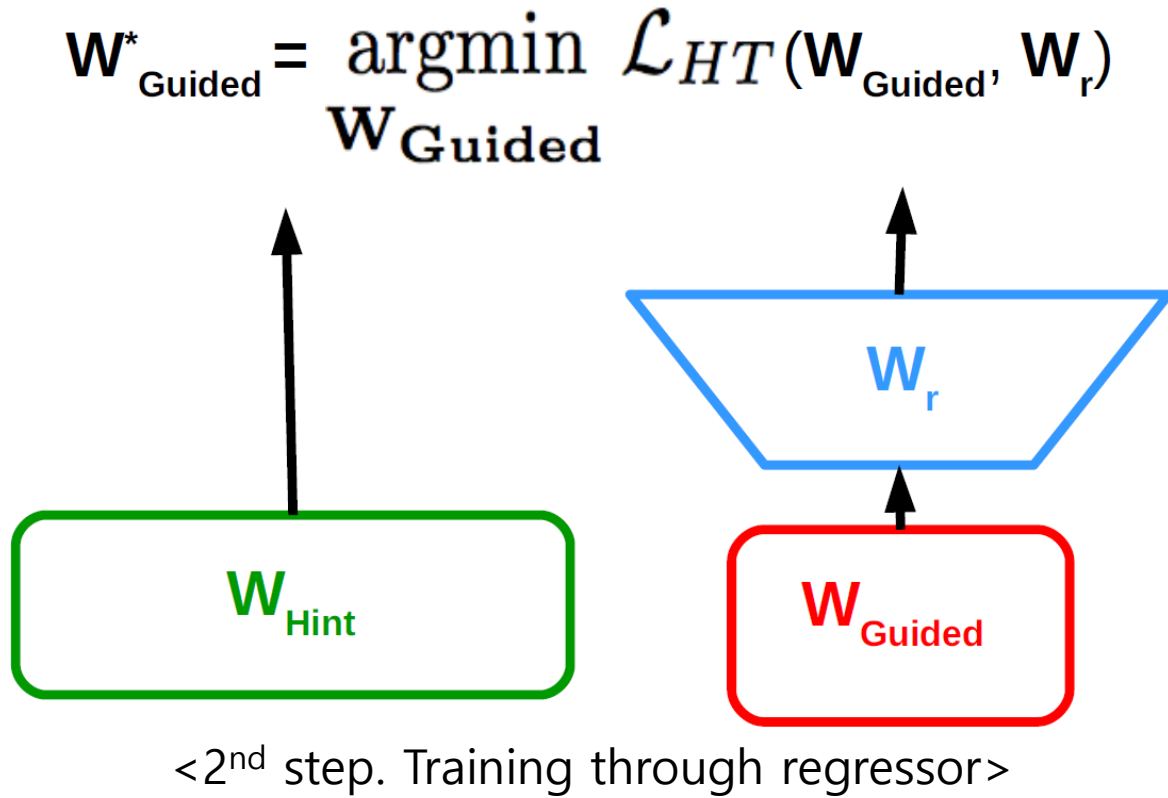
# FitNets: Hints for Thin Deep Nets. In *ICLR*, 2015.

– Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta and Yoshua Bengio.



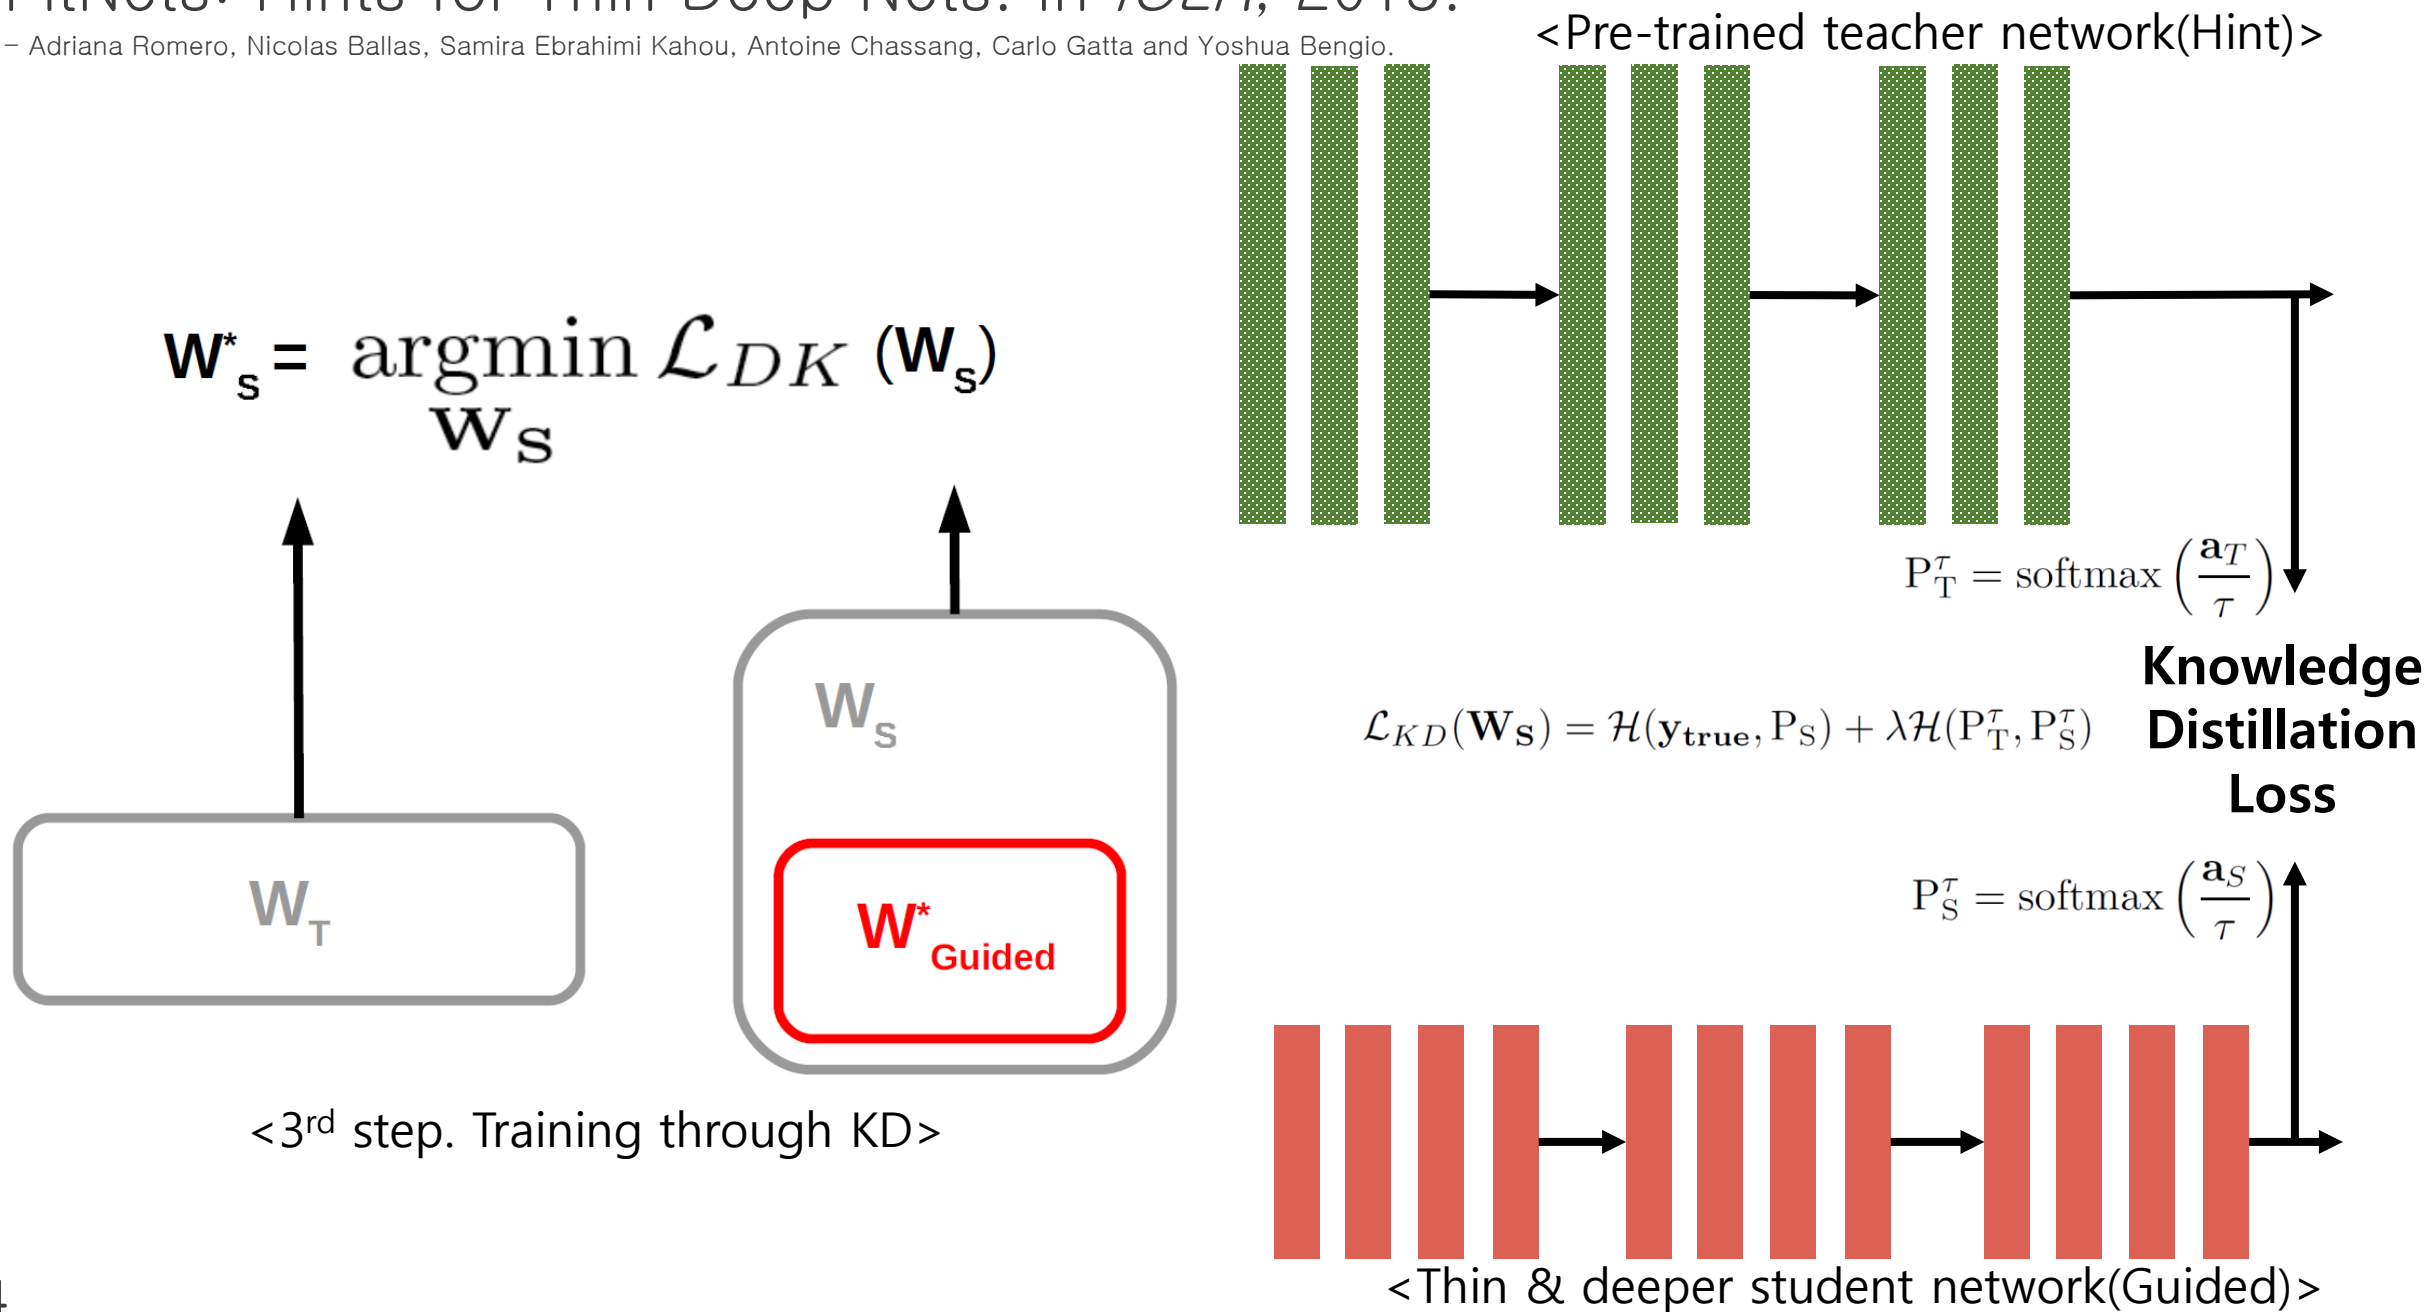
# FitNets: Hints for Thin Deep Nets. In *ICLR*, 2015.

- Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta and Yoshua Bengio.



# FitNets: Hints for Thin Deep Nets. In *ICLR*, 2015.

– Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta and Yoshua Bengio.





# FitNets: Hints for Thin Deep Nets. In *ICLR*, 2015.

– Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta and Yoshua Bengio.

<b>Algorithm</b>	<b># params</b>	<b>Accuracy</b>
<i>Compression</i>		
FitNet	~2.5M	<b>64.96%</b>
Teacher	~9M	63.54%
<i>State-of-the-art methods</i>		
Maxout		61.43%
Network in Network		64.32%
Deeply-Supervised Networks		<b>65.43%</b>

Table 2: Accuracy on CIFAR-100

<b>Network</b>	<b># layers</b>	<b># params</b>	<b># mult</b>	<b>Acc</b>	<b>Speed-up</b>	<b>Compression rate</b>
Teacher	5	~9M	~725M	90.18%	1	1
FitNet 1	11	~250K	~30M	89.01%	<b>13.36</b>	<b>36</b>
FitNet 2	11	~862K	~108M	91.06%	4.64	10.44
FitNet 3	13	~1.6M	~392M	91.10%	1.37	5.62
FitNet 4	19	~2.5M	~382M	<b>91.61%</b>	1.52	3.60

Table 5: Accuracy/Speed Trade-off on CIFAR-10.

# **A Gift from Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning**

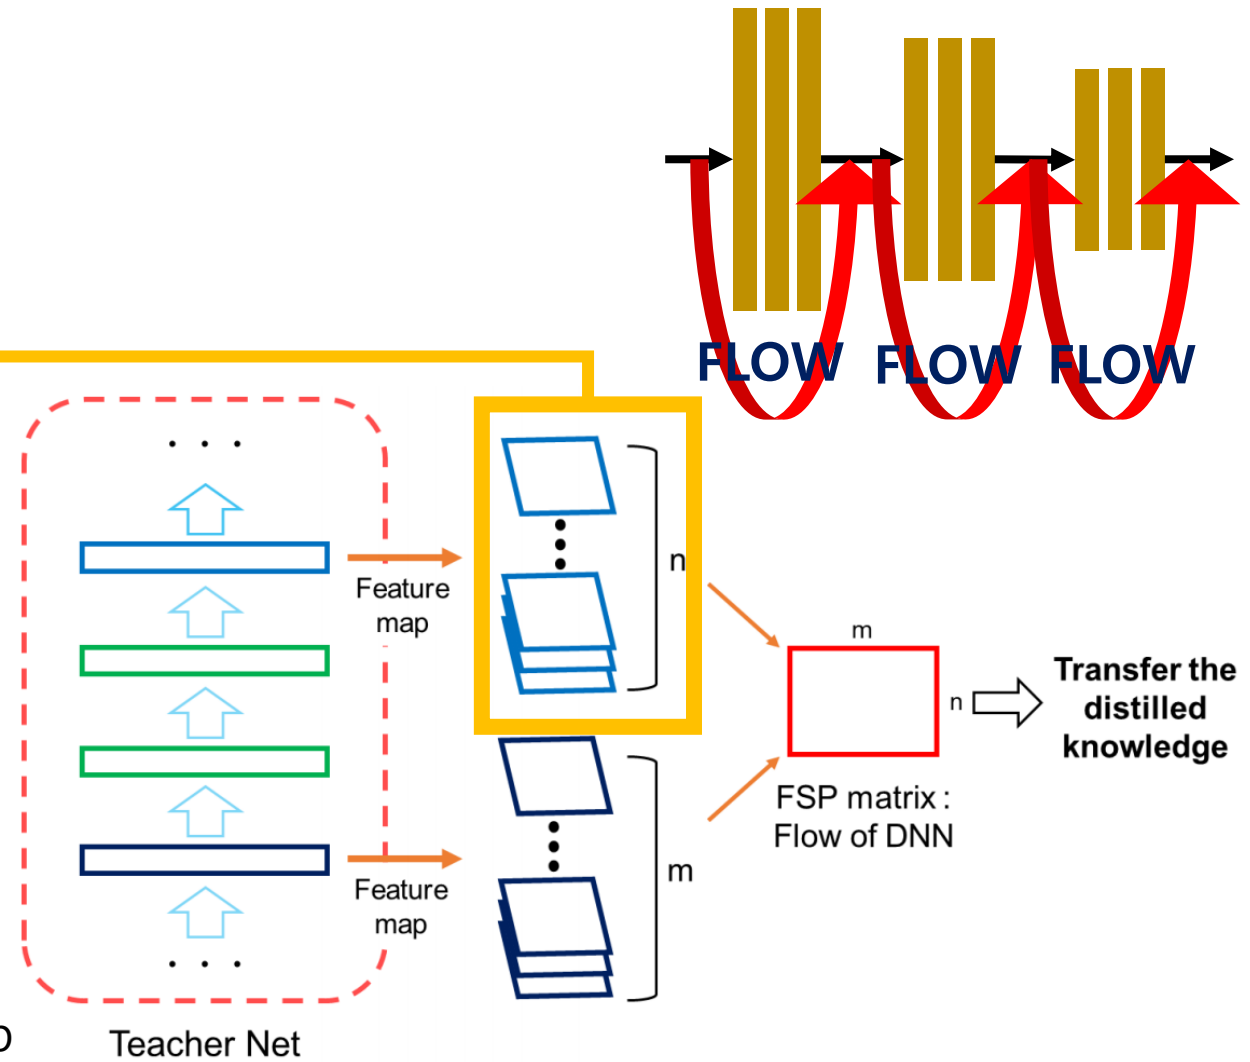
# A Gift from Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning. In *CVPR*, 2017.

- Junho Yim, Donggyu Joo, Jihoon Bae and Junmo Kim.

$$\mathbf{G}(a_1, \dots, a_n) = \begin{pmatrix} (a_1, a_1) & (a_1, a_2) & \dots & (a_1, a_n) \\ (a_2, a_1) & (a_2, a_2) & \dots & (a_2, a_n) \\ \vdots & \vdots & \ddots & \vdots \\ (a_n, a_1) & (a_n, a_2) & \dots & (a_n, a_n) \end{pmatrix}$$

<Gram Matrix definition>

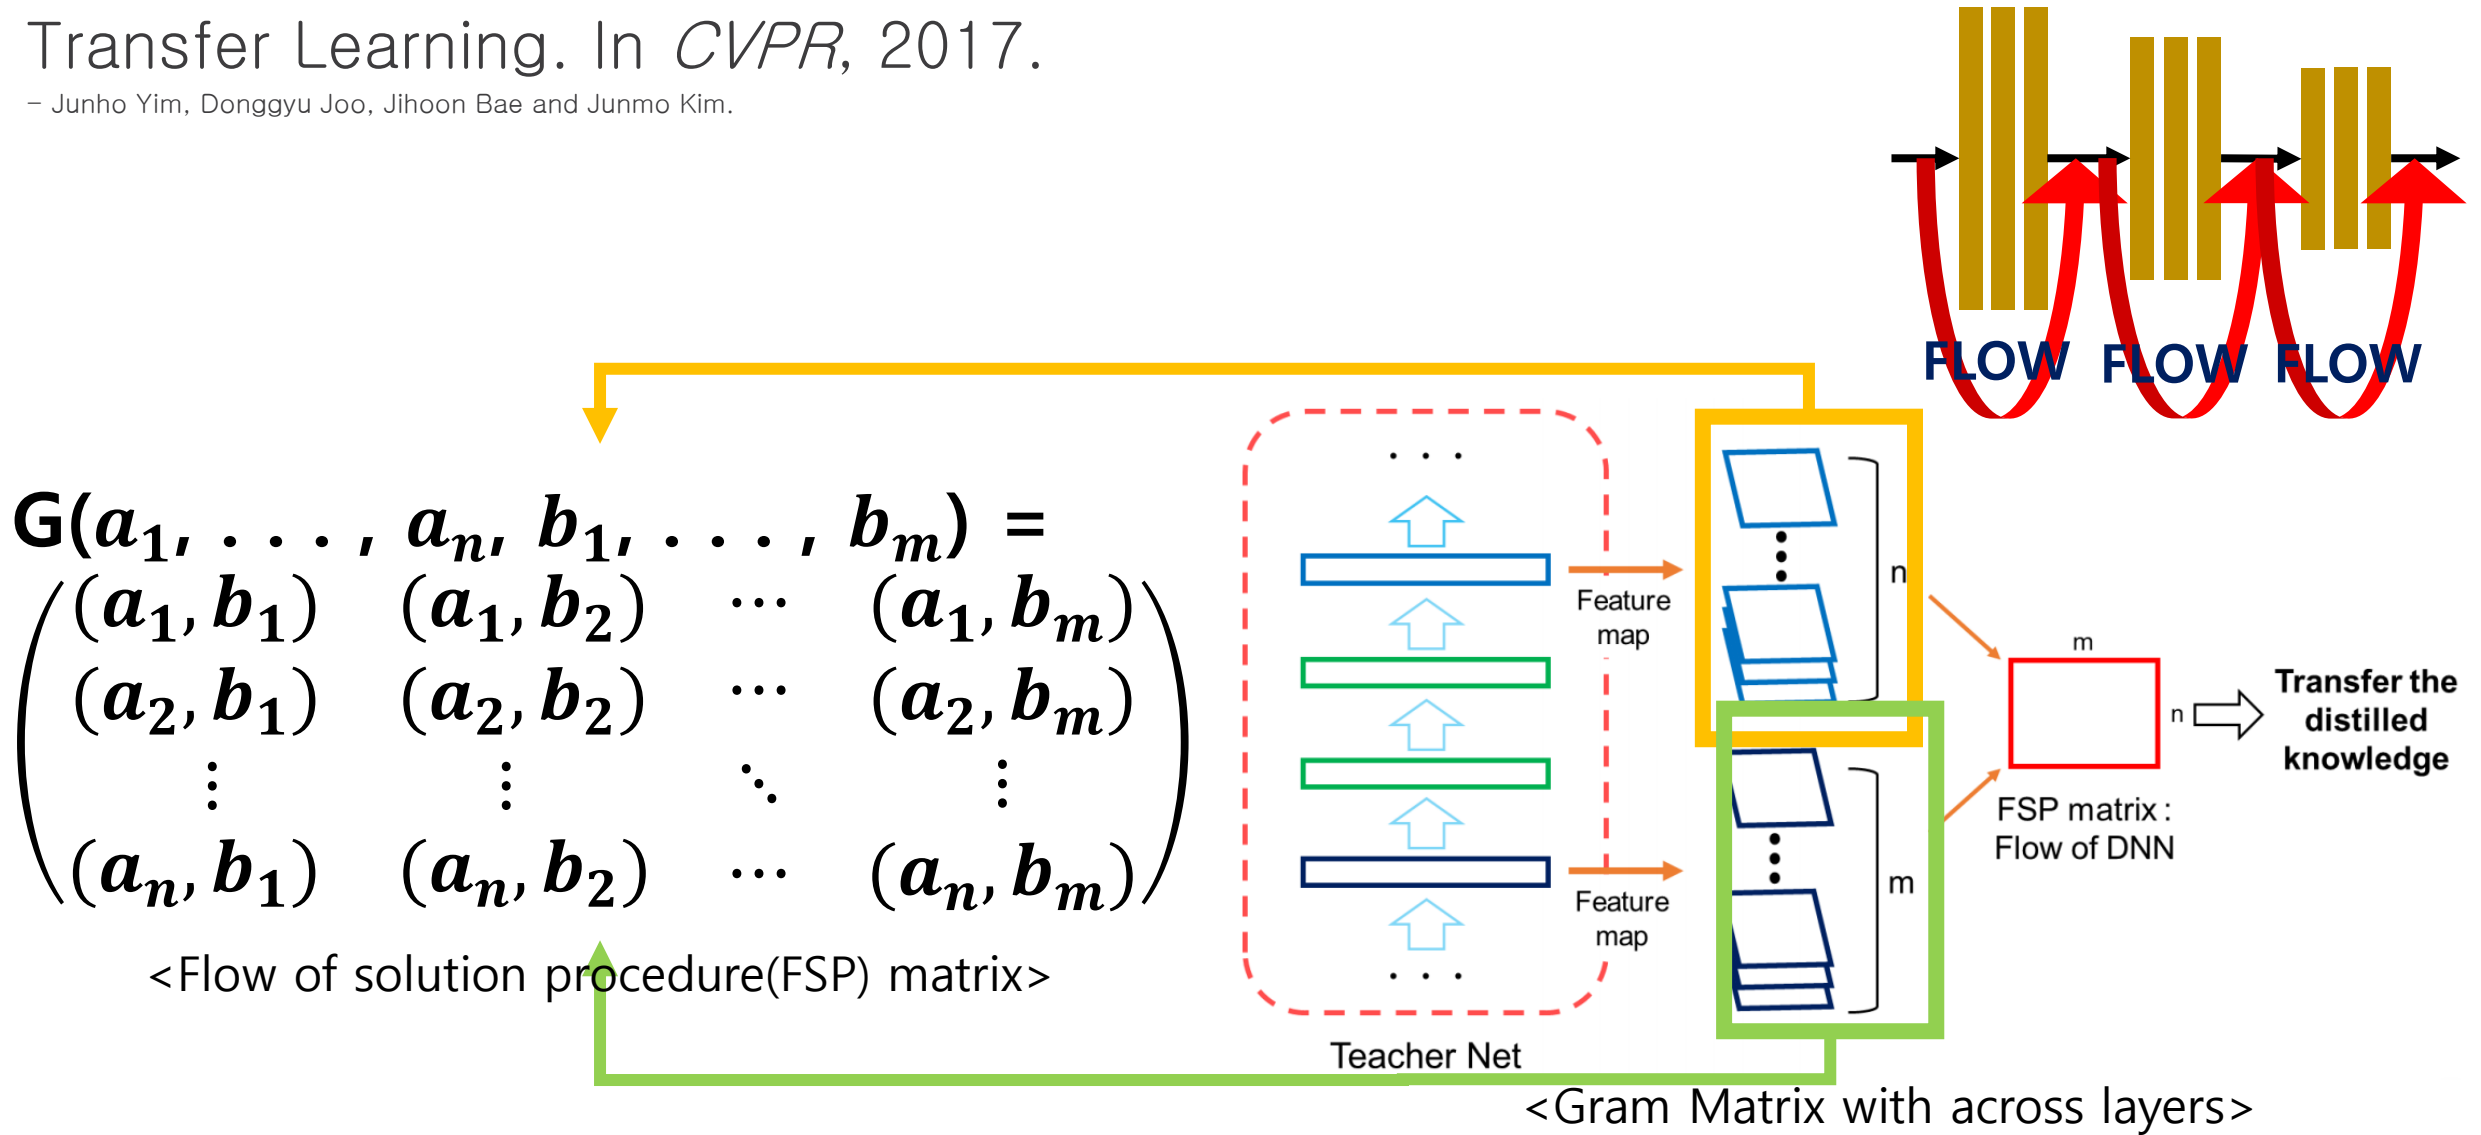
\* (a, b) = Inner product of a and b



<Gram Matrix with across layers>

# A Gift from Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning. In *CVPR*, 2017.

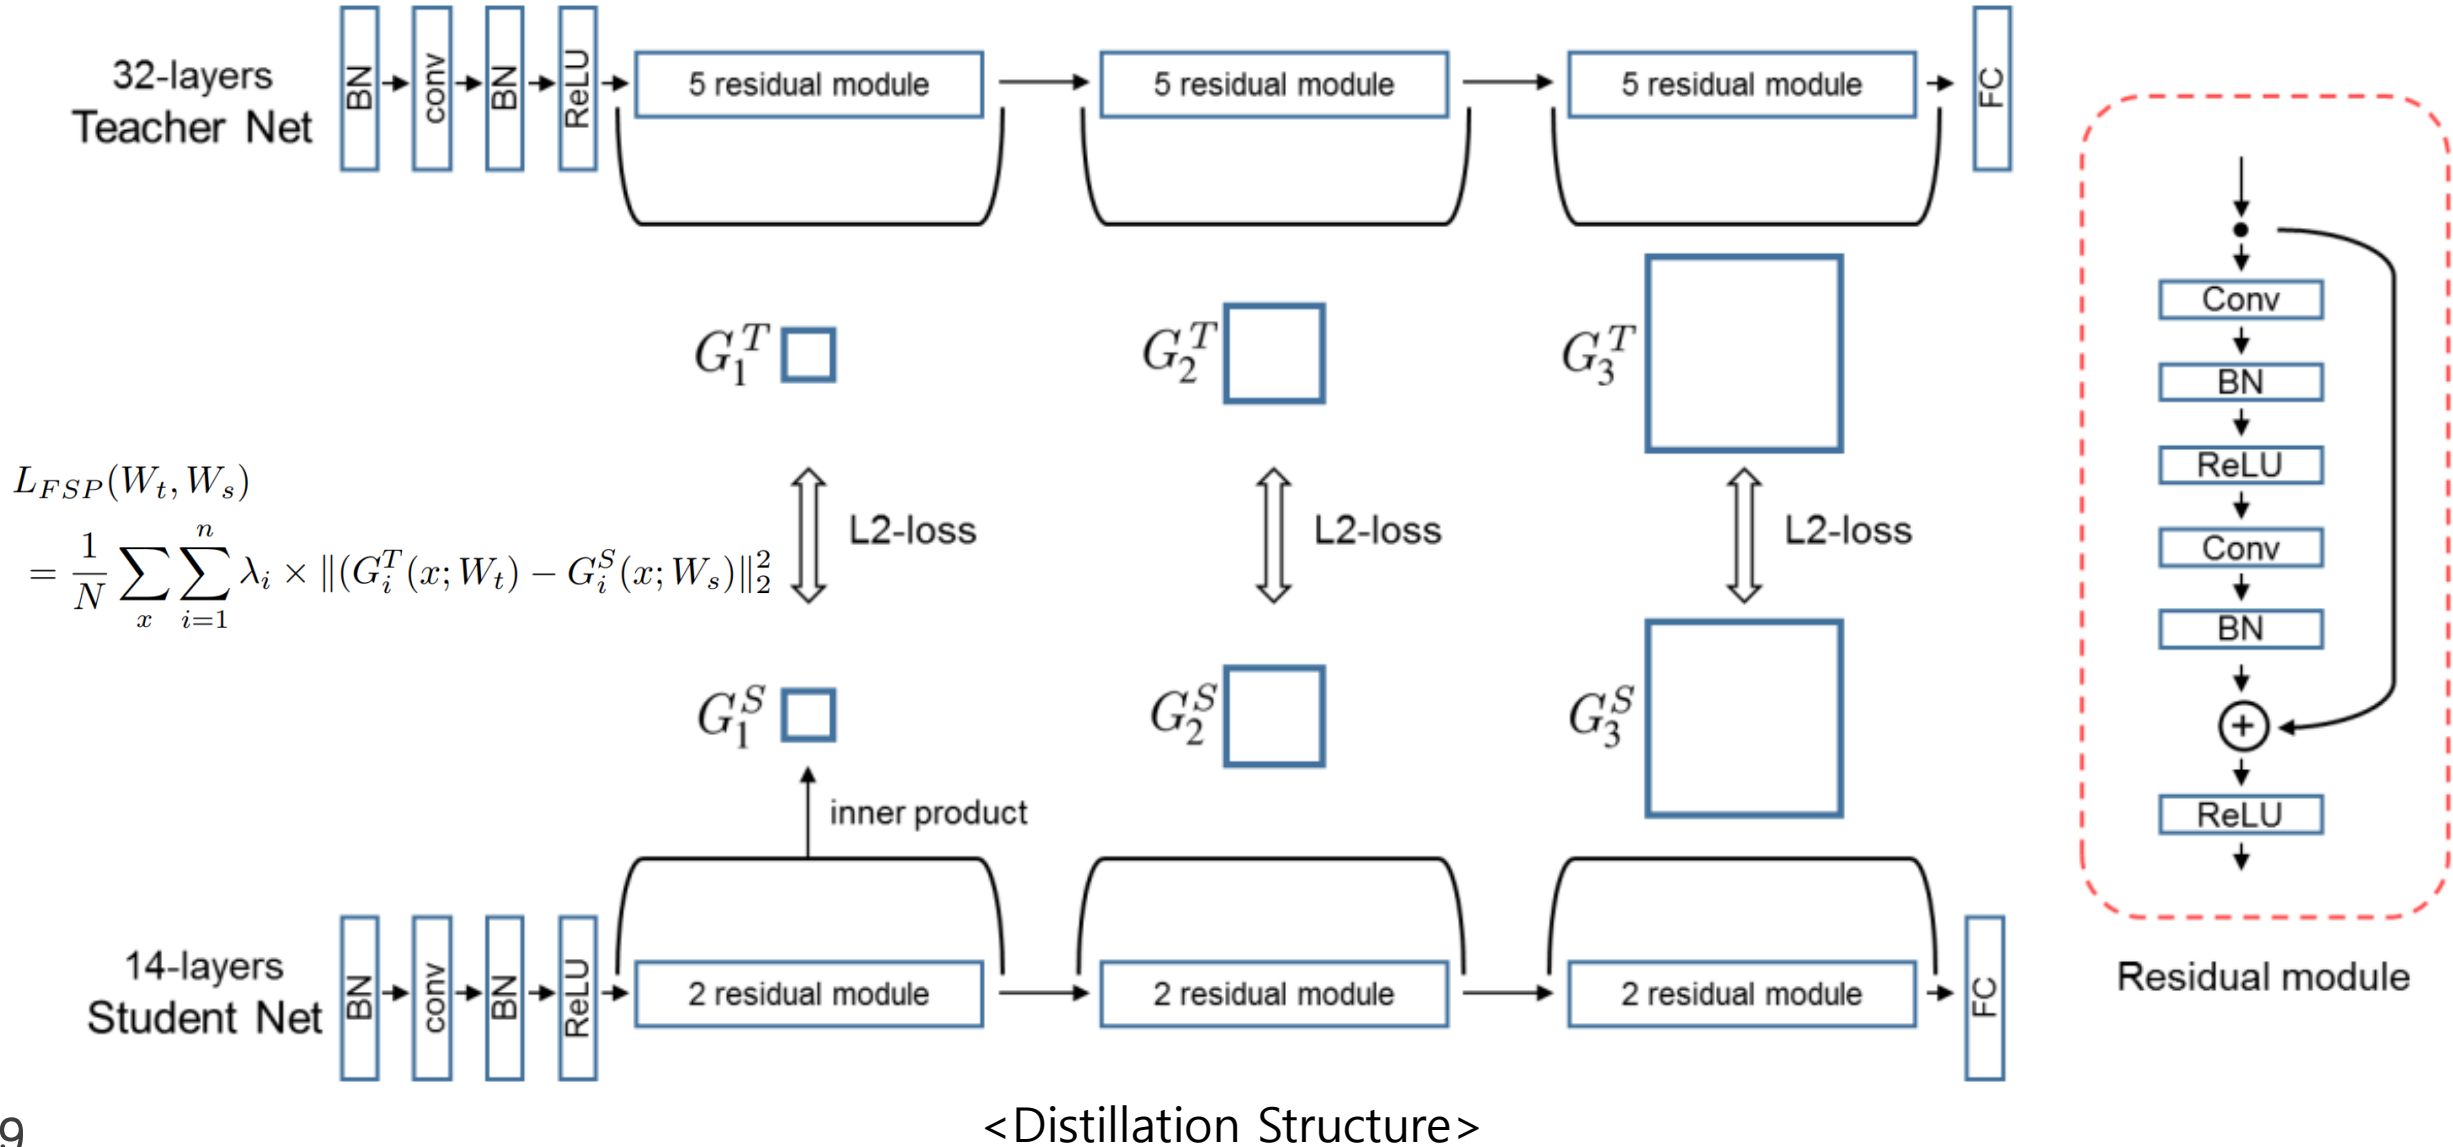
- Junho Yim, Donggyu Joo, Jihoon Bae and Junmo Kim.



\* (a, b) = Inner product of a and b

# A Gift from Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning. In *CVPR*, 2017.

– Junho Yim, Donggyu Joo, Jihoon Bae and Junmo Kim.



$$L_{FSP}(W_t, W_s)$$

$$= \frac{1}{N} \sum_x \sum_{i=1}^n \lambda_i \times \|(G_i^T(x; W_t) - G_i^S(x; W_s))\|_2^2$$

L2-loss

L2-loss

L2-loss

Residual module

# A Gift from Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning. In *CVPR*, 2017.

– Junho Yim, Donggyu Joo, Jihoon Bae and Junmo Kim.

	Accuracy
Teacher-original	91.91
Student-original	87.91
FitNet [20]	88.57
Proposed Method	88.70

Table 3. Recognition rates (%) on CIFAR-10. We used a residual DNN with 8 layers for the student DNN and 26 layers for the teacher DNN.

	Accuracy
Teacher-original	64.06
Student-original	58.65
FitNet [20]	61.28
Proposed Method	63.33

Table 4. Recognition rates (%) on CIFAR-100. We used a residual DNN with 14 layers for the student DNN and 32 layers for the teacher DNN.

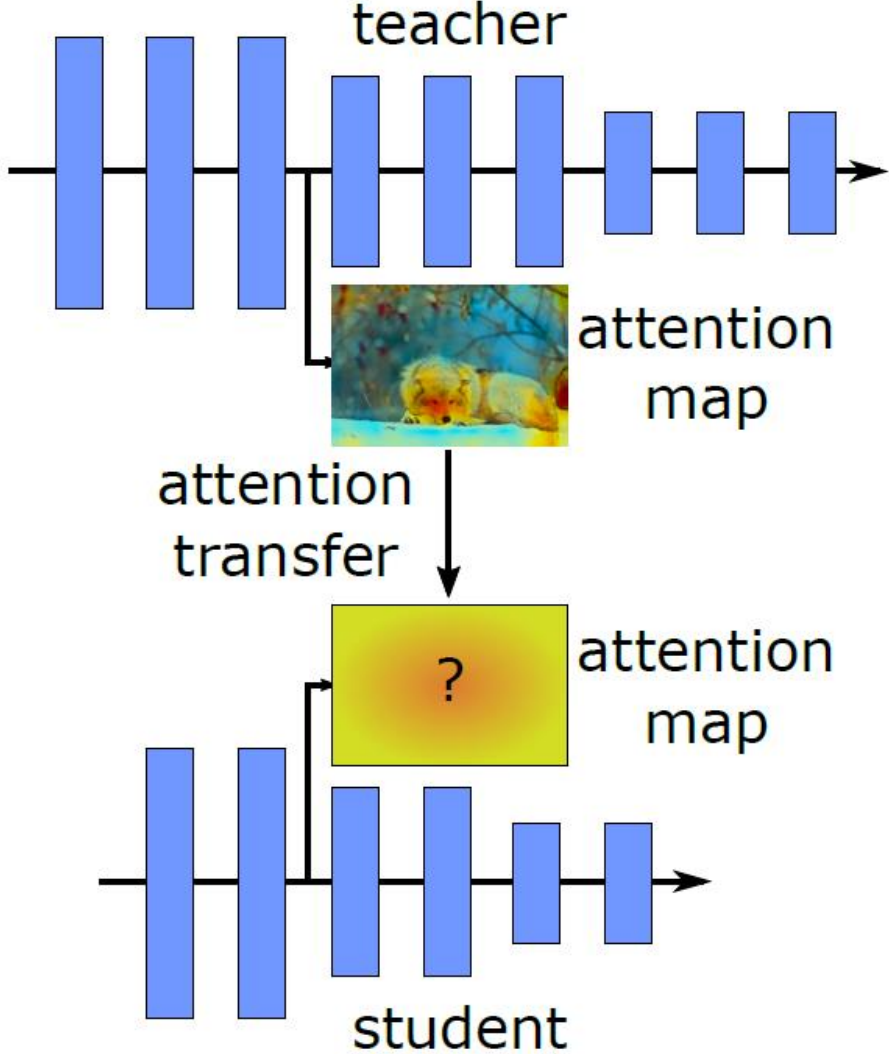
# **Paying More Attention to Attention: Improving the Performance of CNN via Attention Transfer**

# Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR*, 2017.

– Sergey Zagoruyko and Nikos Komodakis.



<Normal image & Attention map>




<Attention Transfer>




# Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR*, 2017.

– Sergey Zagoruyko and Nikos Komodakis.



$$\begin{aligned}
 \mathcal{F} &: R^{C \times H \times W} \rightarrow R^{H \times W} \\
 F_{\text{sum}}(A) &= \sum_{i=1}^C |A_i| \\
 F_{\text{sum}}^p(A) &= \sum_{i=1}^C |A_i|^p \\
 F_{\text{max}}^p(A) &= \max_{i=1, C} |A_i|^p
 \end{aligned}$$

<Attention mapping functions>



63 x 63



32 x 32




8 x 8

<Feature characteristic of each layer>

# Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR*, 2017.

– Sergey Zagoruyko and Nikos Komodakis.



$$\mathcal{L}_{AT} = \mathcal{L}(\mathbf{W}_S, x) + \frac{\beta}{2} \sum_{j \in \mathcal{I}} \left\| \frac{Q_S^j}{\|Q_S^j\|_2} - \frac{Q_T^j}{\|Q_T^j\|_2} \right\|_p$$

<Attention Transfer structure>

# Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR*, 2017.

– Sergey Zagoruyko and Nikos Komodakis.




Figure 4: Activation attention maps for various ImageNet networks: Network-In-Network (62% top-1 val accuracy), ResNet-34 (73% top-1 val accuracy), ResNet-101 (77.3% top-1 val accuracy). Left part: mid-level activations, right part: top-level pre-softmax activations

# Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR*, 2017.

– Sergey Zagoruyko and Nikos Komodakis.

student	teacher	student	AT	F-ActT	KD	AT+KD	teacher
NIN-thin, 0.2M	NIN-wide, 1M	9.38	8.93	9.05	8.55	8.33	7.28
WRN-16-1, 0.2M	WRN-16-2, 0.7M	8.77	7.93	8.51	7.41	7.51	6.31
WRN-16-1, 0.2M	WRN-40-1, 0.6M	8.77	8.25	8.62	8.39	8.01	6.58
WRN-16-2, 0.7M	WRN-40-2, 2.2M	6.31	5.85	6.24	6.08	5.71	5.23

Table 1: Activation-based attention transfer (AT) with various architectures on CIFAR-10. Error is computed as median of 5 runs with different seed. F-ActT means full-activation transfer (see §4.1.2).

Model	top1, top5
ResNet-18	30.4, 10.8
AT	29.3, 10.0
ResNet-34	26.1, 8.3

Table 5: Attention transfer validation error (single crop) on ImageNet. Transfer losses are added on epoch 60/100.

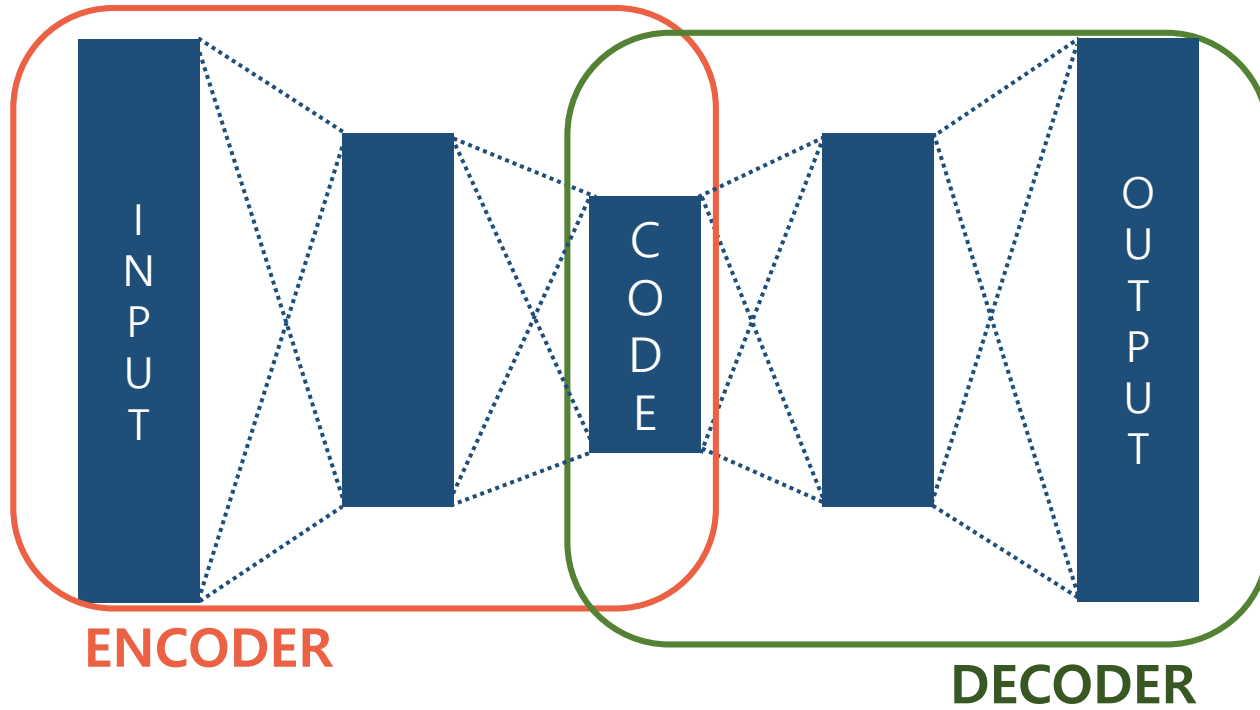
# Paraphrasing Complex Network: Network Compression via Factor Transfer




# Paraphrasing Complex Network: Network Compression via Factor Transfer.

In *NIPS*, 2018.

– Jangho Kim, SeongUk Park and Nojun Kwak.



<Autoencoder structure>

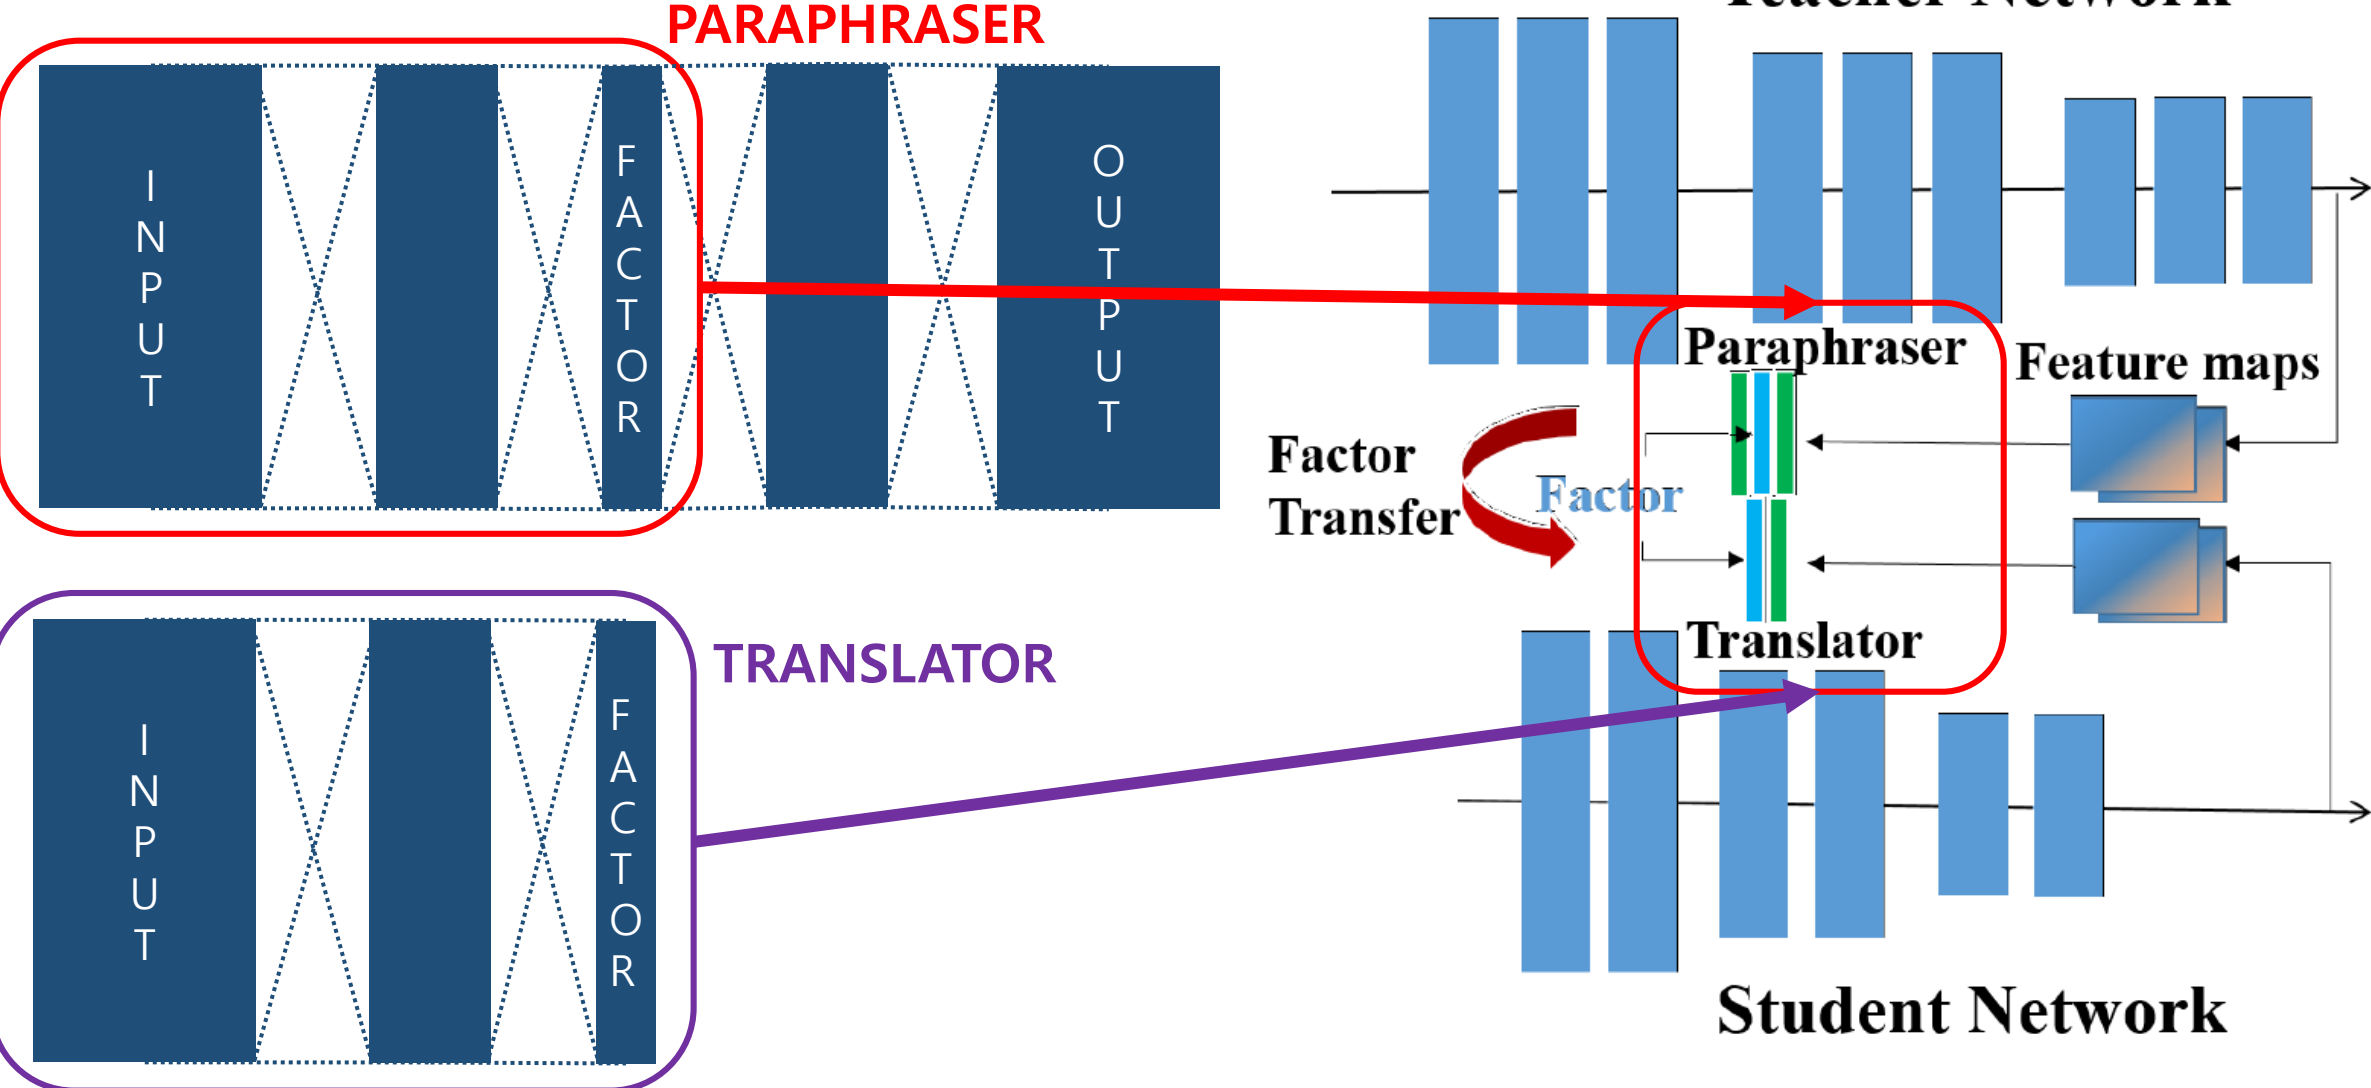


<t-SNE visualization of factor space>

# Paraphrasing Complex Network: Network Compression via Factor Transfer.

In *NIPS*, 2018.

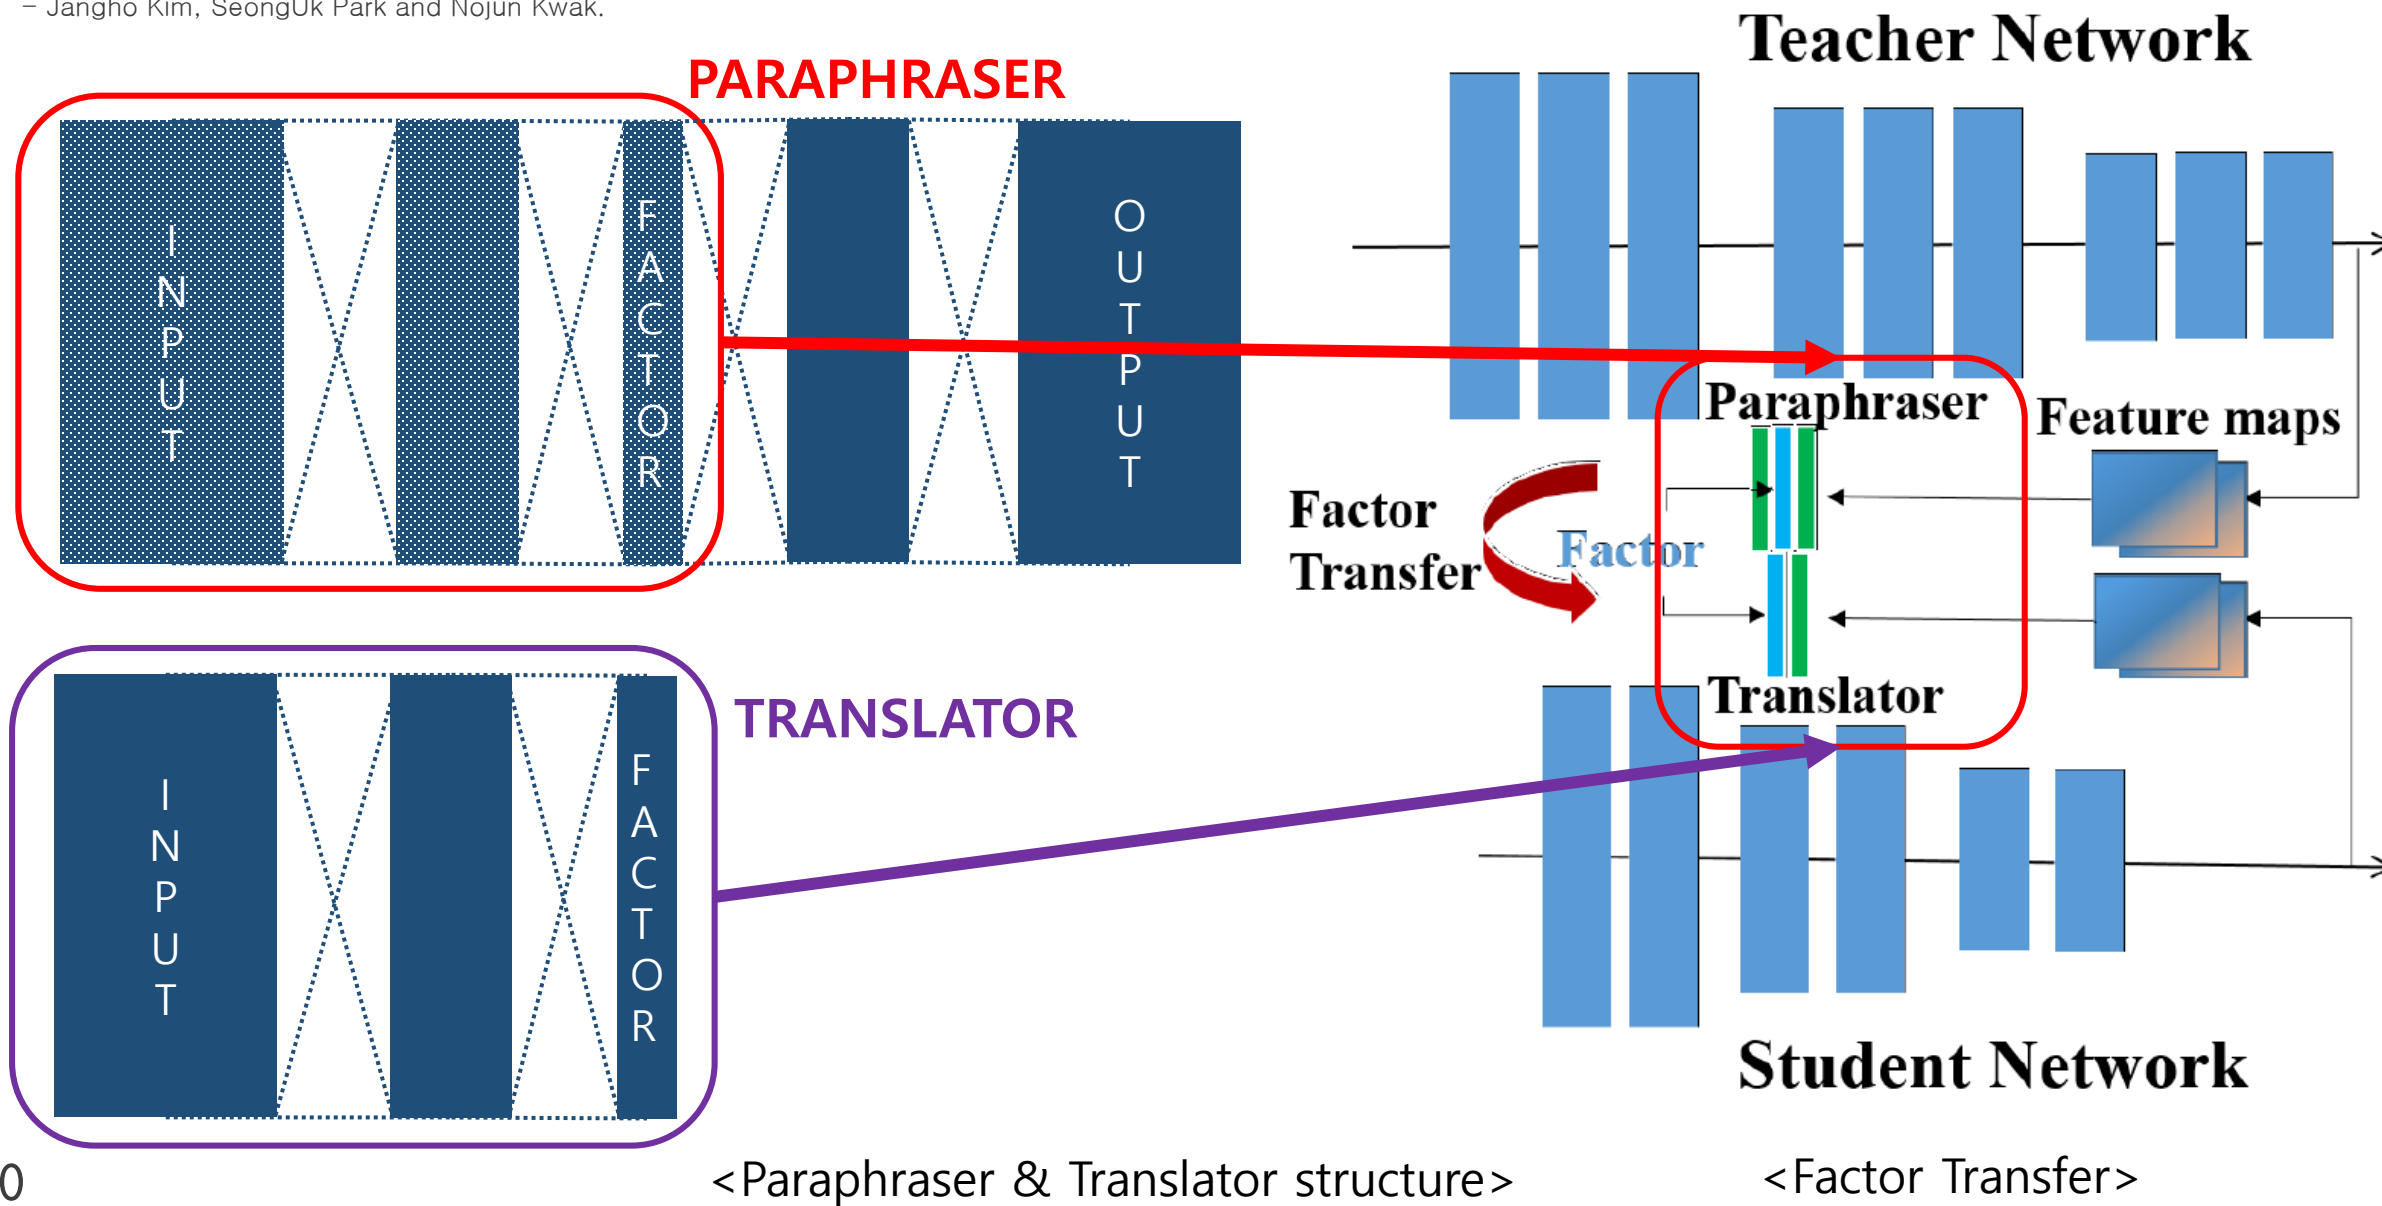
- Jangho Kim, SeongUk Park and Nojun Kwak.



# Paraphrasing Complex Network: Network Compression via Factor Transfer.

In *NIPS*, 2018.

– Jangho Kim, SeongUk Park and Nojun Kwak.

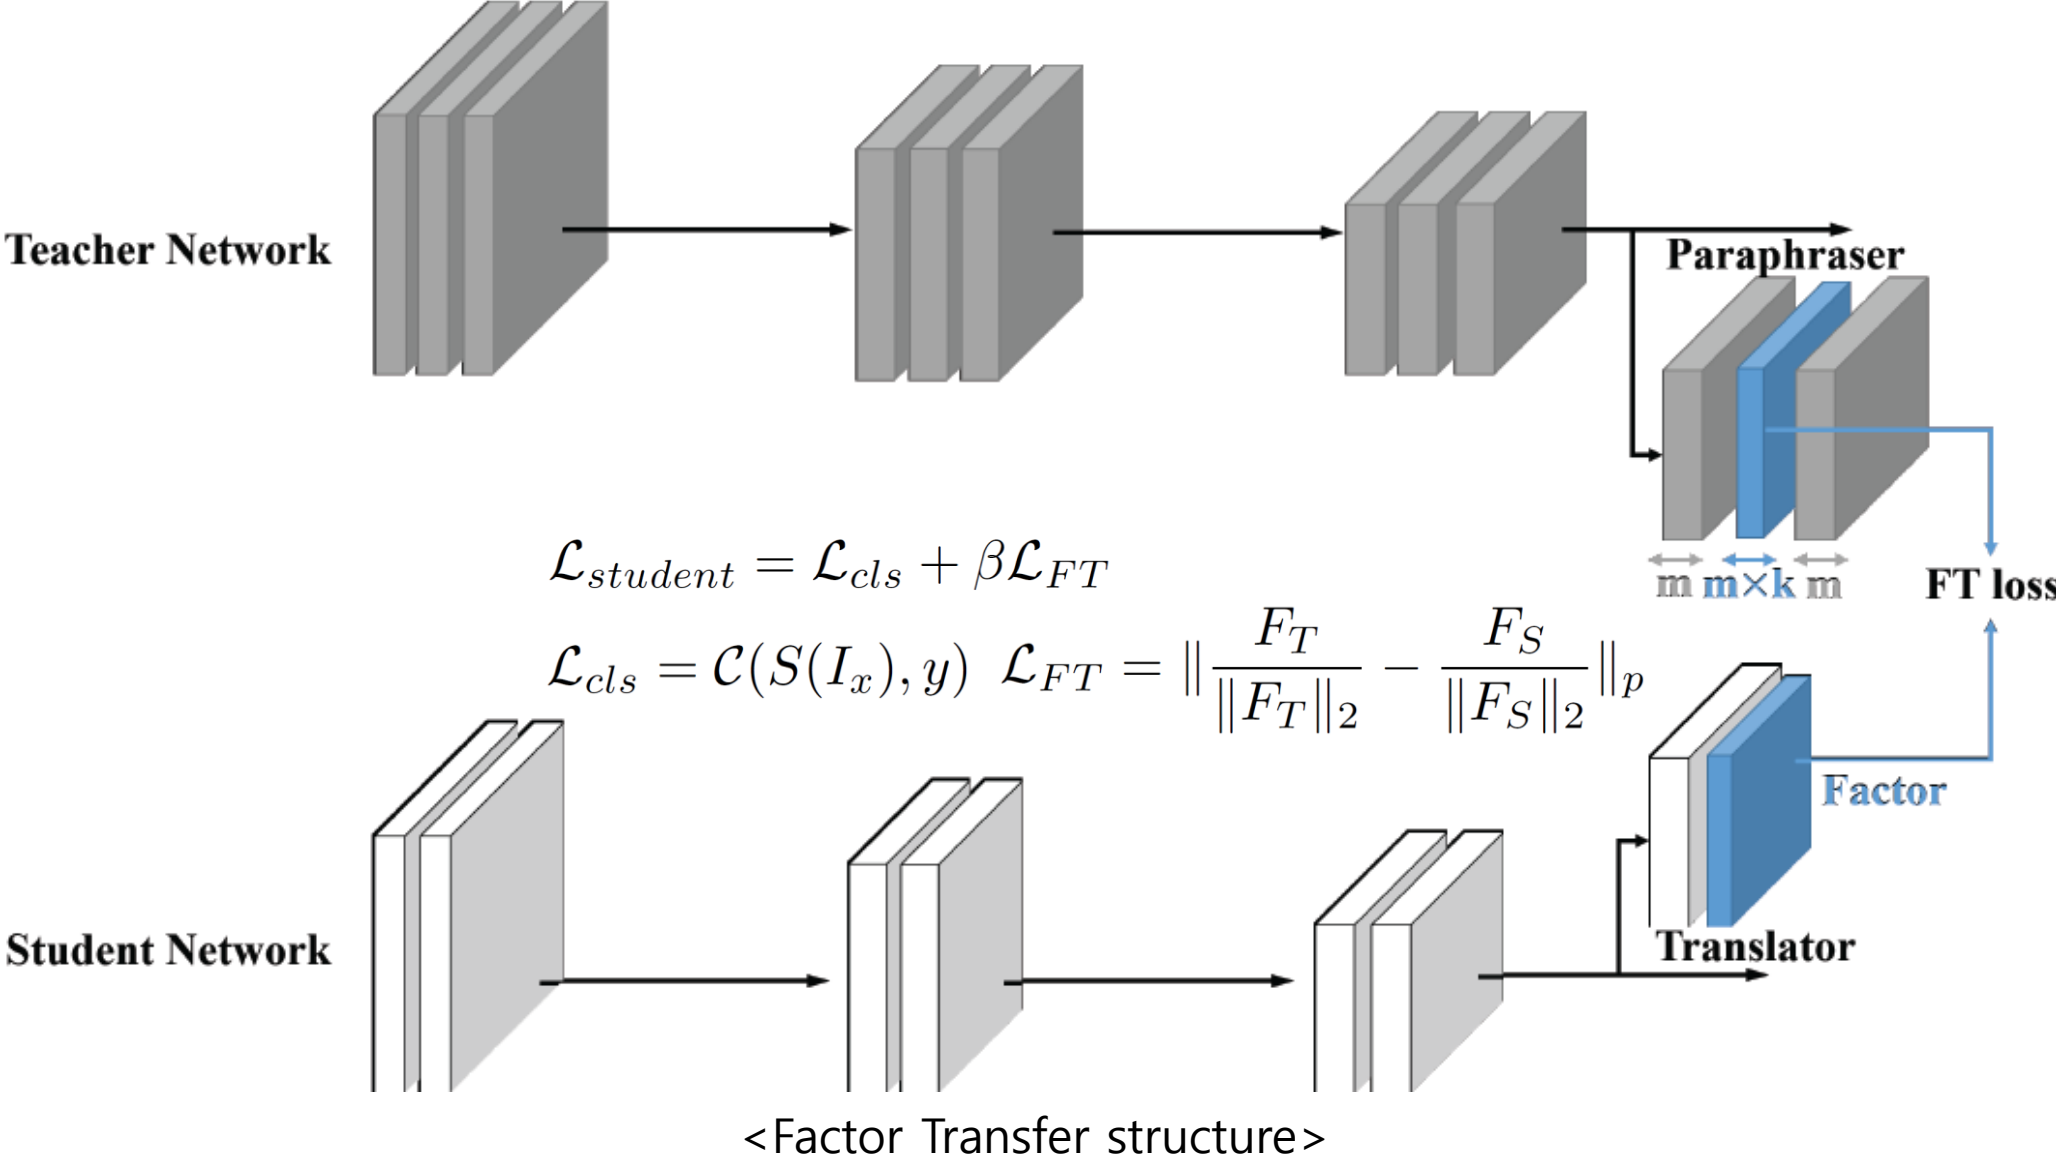




# Paraphrasing Complex Network: Network Compression via Factor Transfer.

In *NIPS*, 2018.

– Jangho Kim, SeongUk Park and Nojun Kwak.



# Paraphrasing Complex Network: Network Compression via Factor Transfer.

In *NIPS*, 2018.

– Jangho Kim, SeongUk Park and Nojun Kwak.

Student	Teacher	Student	AT	KD	FT	AT+KD	FT+KD	Teacher
ResNet-56 (0.85M)	ResNet-110 (1.73M)	28.04	27.28	27.96	<b>25.62</b>	28.01	26.93	26.91
ResNet-20 (0.27M)	ResNet-110 (1.73M)	31.24	31.04	33.14	<b>29.08</b>	34.78	32.19	26.91
Student	Teacher	$k = 0.5$	$k = 0.75$	$k = 1$	$k = 2$	$k = 4$	CAE	RAE
ResNet-56 (0.85M)	ResNet-110 (1.73M)	<b>25.62</b>	25.78	25.85	25.63	25.87	26.41	26.29
ResNet-20 (0.27M)	ResNet-110 (1.73M)	29.20	29.25	29.28	29.19	<b>29.08</b>	29.84	30.11

Table 3: Mean classification error (%) on CIFAR-100 dataset (5 runs). All the numbers are from our implementation.

Paraphraser	Translator	CIFAR-10	CIFAR-100
Yes	No	6.18	27.61
No	Yes	6.12	27.39
Yes	Yes	5.71	26.91
Student (WRN-40-1[0.6M])		7.02	28.81
Teacher (WRN-40-2[2.2M])		4.96	24.10

Table 4: Ablation study with and without the paraphraser ( $k = 0.5$ ) and the Translator. (Mean test error (%) of 5 runs).

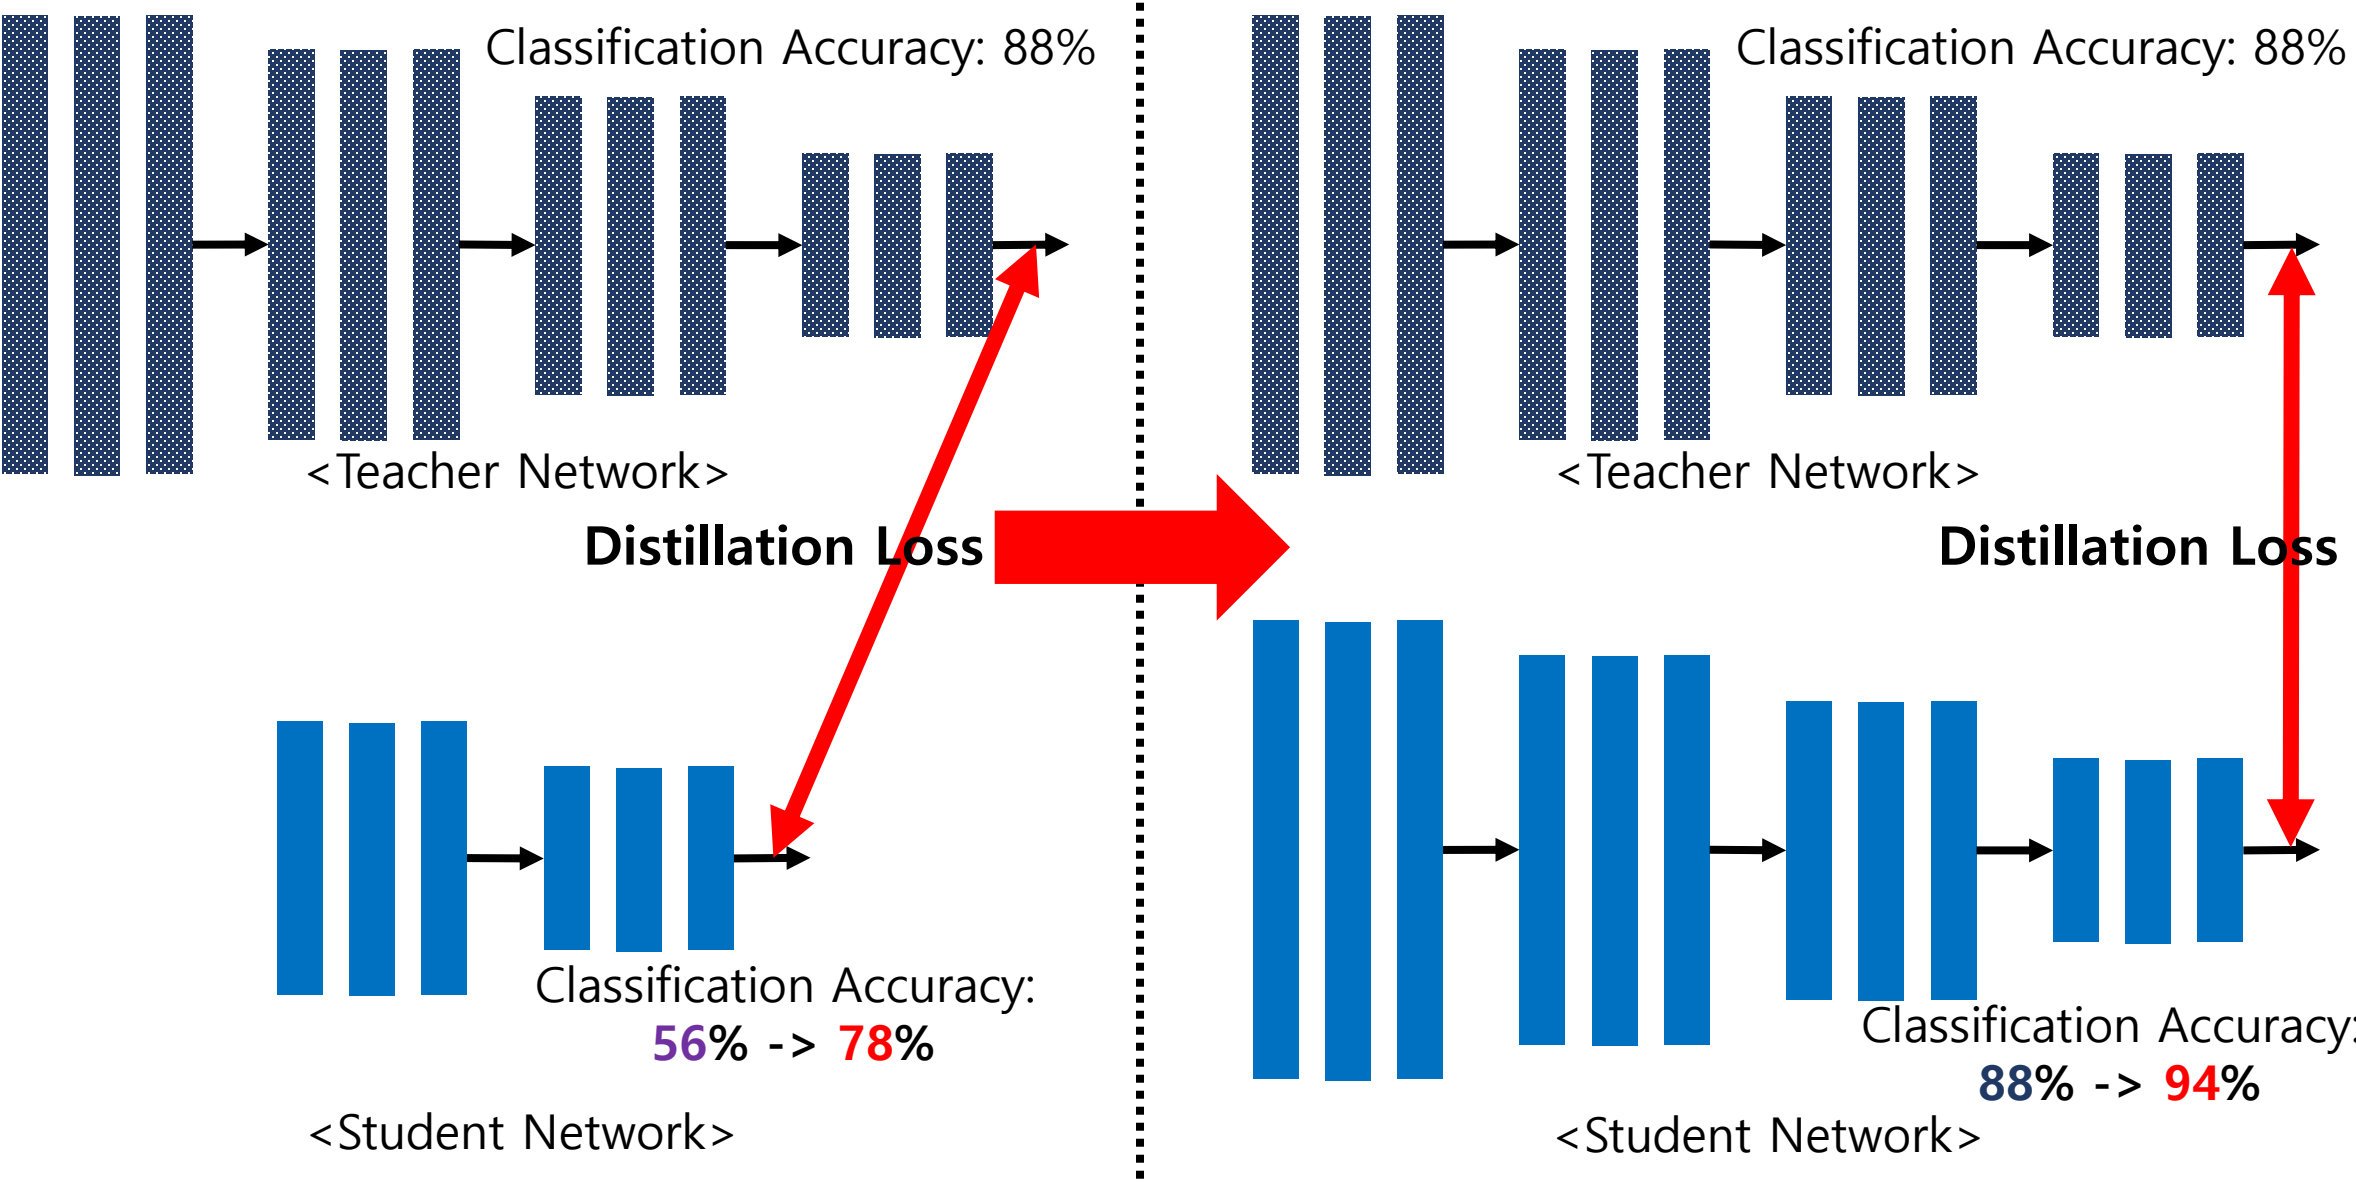
Method	Network	Top-1	Top-5
Student	Resnet-18	<b>29.91</b>	<b>10.68</b>
KD	Resnet-18	33.83	12.55
AT	Resnet-18	29.36	10.23
FT ( $k = 0.5$ )	Resnet-18	<b>28.57</b>	<b>9.71</b>
Teacher	Resnet-34	26.73	8.57

Table 5: Top-1 and Top-5 classification error (%) on ImageNet dataset. All the numbers are from our implementation.

# Born-Again Neural Networks

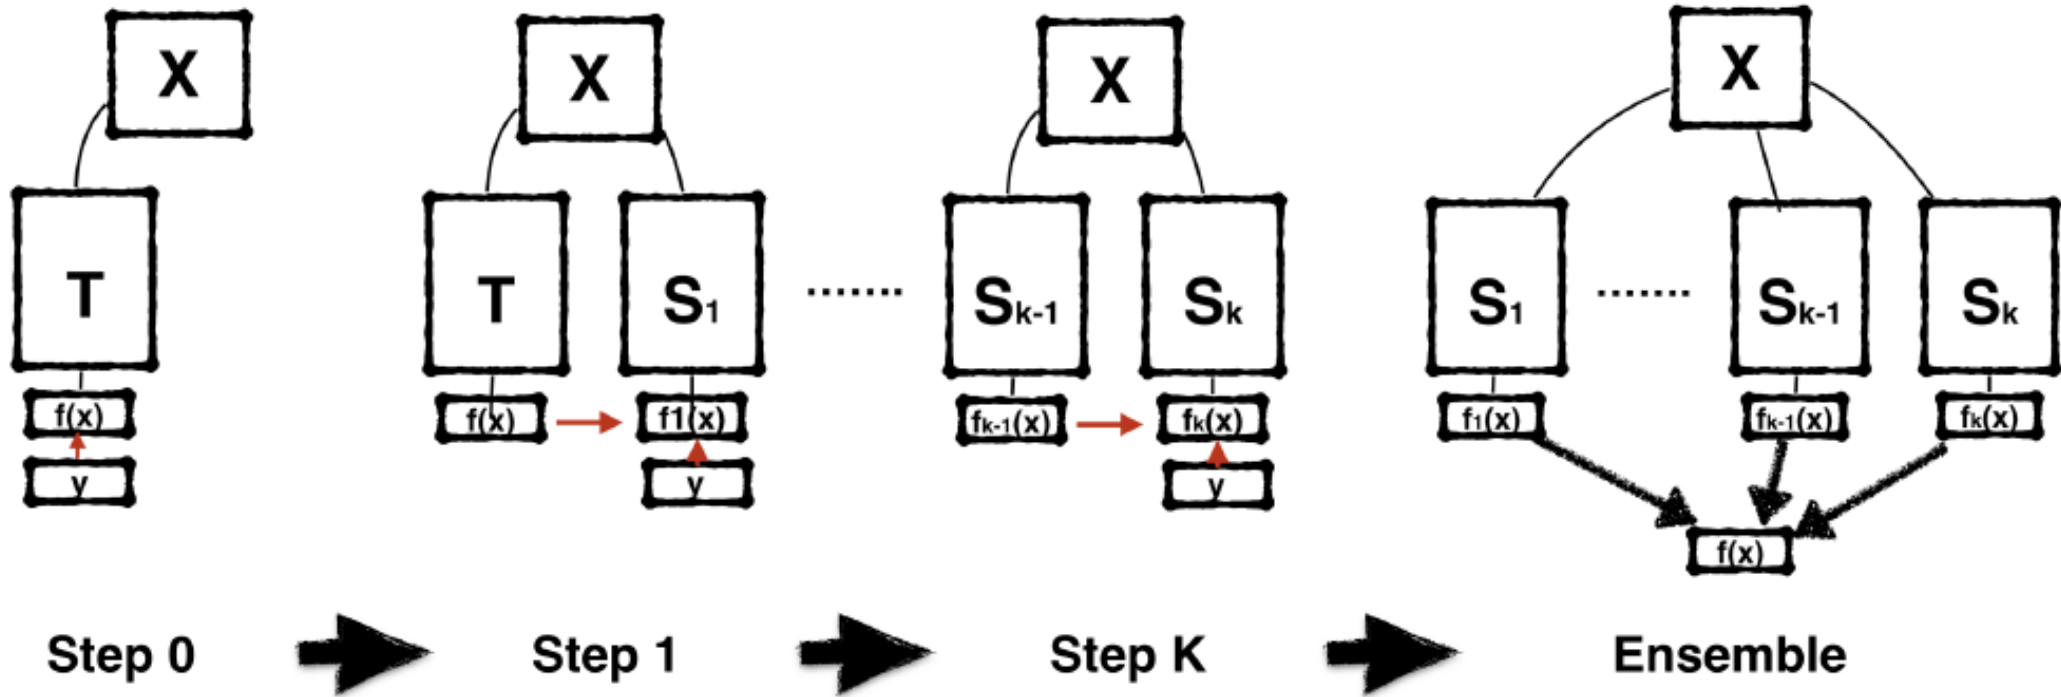
# Born-Again Neural Networks. In *ICML*, 2018.

- Tommaso Furlanello, Zachary C. Lipton, Michael Tschannen, Laurent Itti and Anima Anandkumar.



# Born-Again Neural Networks. In *ICML*, 2018.

– Tommaso Furlanello, Zachary C. Lipton, Michael Tschannen, Laurent Itti and Anima Anandkumar.



**Figure 1. Graphical representation of the BAN training procedure:** during the first step the teacher model  $T$  is trained from the labels  $Y$ . Then, at each consecutive step, a new identical model is initialized from a different random seed and trained from the supervision of the earlier generation. At the end of the procedure, additional gains can be achieved with an ensemble of multiple students generations.

$$\mathcal{L}(f(x, \arg \min_{\theta_{k-1}} \mathcal{L}(f(x, \theta_{k-1}))), f(x, \theta_k)) \quad \hat{f}^k(x) = \sum_{i=1}^k f(x, \theta_i) / k$$

# Born-Again Neural Networks. In *ICML*, 2018.

– Tommaso Furlanello, Zachary C. Lipton, Michael Tschannen, Laurent Itti and Anima Anandkumar.

*Table 1. Test error on CIFAR-10* for Wide-ResNet with different depth and width and DenseNet of different depth and growth factor.

Network	Parameters	Teacher	BAN
Wide-ResNet-28-1	0.38 M	6.69	<b>6.64</b>
Wide-ResNet-28-2	1.48 M	5.06	<b>4.86</b>
Wide-ResNet-28-5	9.16 M	4.13	<b>4.03</b>
Wide-ResNet-28-10	36 M	<b>3.77</b>	3.86
DenseNet-112-33	6.3 M	3.84	<b>3.61</b>
DenseNet-90-60	16.1 M	3.81	<b>3.5</b>
DenseNet-80-80	22.4 M	<b>3.48</b>	3.49
DenseNet-80-120	50.4 M	<b>3.37</b>	3.54

*Table 3. Test error on CIFAR-100* for Wide-ResNet students trained from identical Wide-ResNet teachers and for DenseNet-90-60 students trained from Wide-ResNet teachers

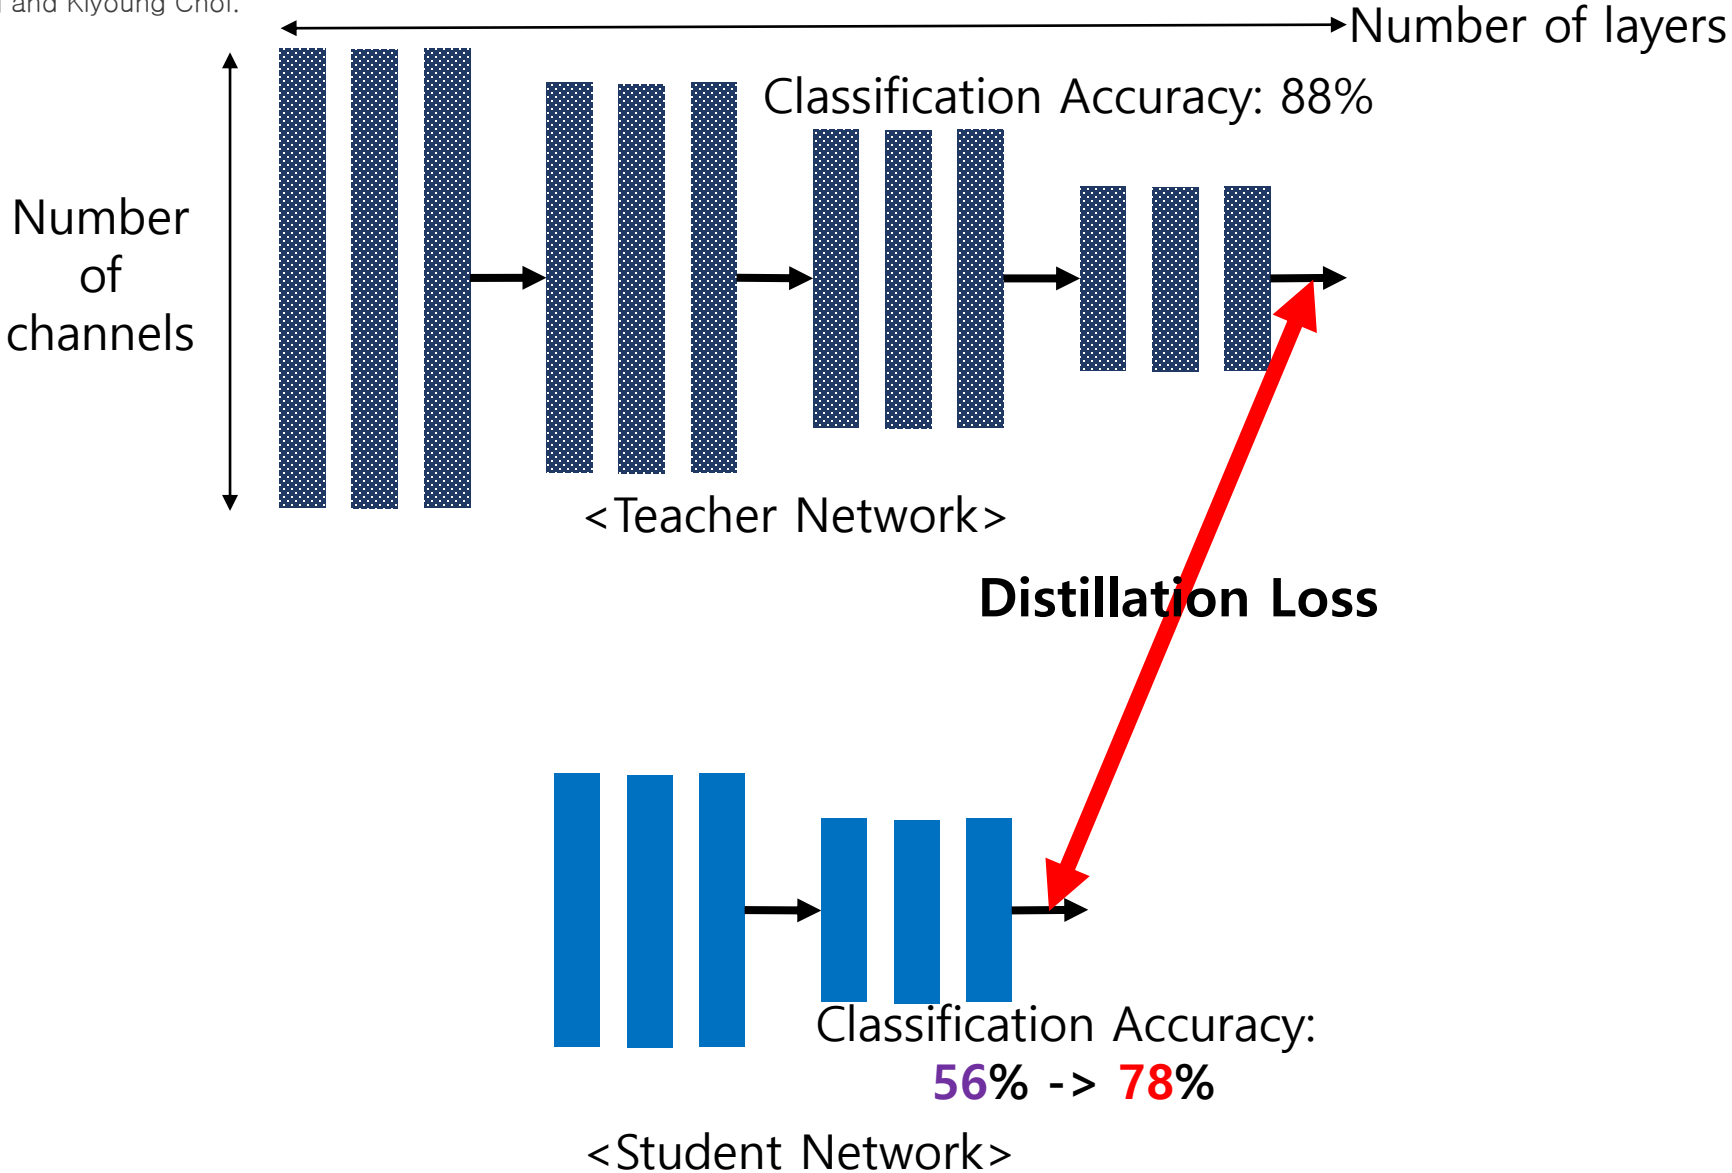
Network	Teacher	BAN	Dense-90-60
Wide-ResNet-28-1	30.05	29.43	24.93
Wide-ResNet-28-2	25.32	24.38	18.49
Wide-ResNet-28-5	20.88	20.93	17.52
Wide-ResNet-28-10	19.08	18.25	16.79

# **Network Recasting: A Universal Method for Network Architecture Transformation**

# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

- Joonsang Yu, Sungbum Kang and Kiyoung Choi.





# Network Recasting: A Universal Method for Network Architecture Transformation. In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.

<Teacher Network>

**MSE Loss**

<Student Network>


$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$



# Network Recasting: A Universal Method for Network Architecture Transformation. In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.

<Teacher Network>



**MSE Loss**



<Student Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

# Network Recasting: A Universal Method for Network Architecture Transformation. In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.


<Teacher Network>



**MSE Loss**



<Student Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

# Network Recasting: A Universal Method for Network Architecture Transformation. In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.


<Teacher Network>



**MSE Loss**



<Student Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.


<Teacher Network>



**MSE Loss**



<Student Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.


<Teacher Network>



**MSE Loss**



<Student Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

# Network Recasting: A Universal Method for Network Architecture Transformation. In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.

<Teacher Network>




**MSE Loss**



$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

<Student Network>




# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.

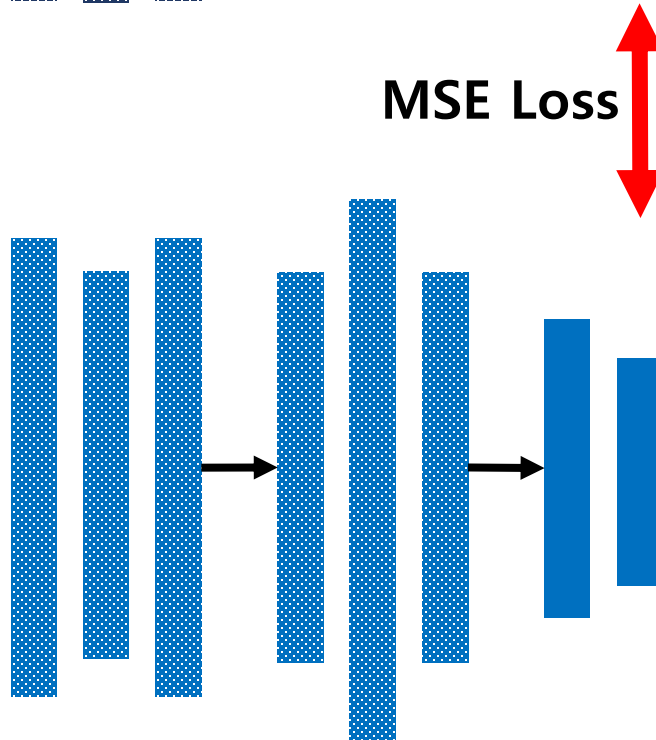
<Teacher Network>



**MSE Loss**

$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

<Student Network>






# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.

<Teacher Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

**MSE Loss**



<Student Network>




# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.

<Teacher Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

**MSE Loss**



<Student Network>




# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.

<Teacher Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

**MSE Loss**



<Student Network>




# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.

<Teacher Network>




$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

**MSE Loss**



<Student Network>



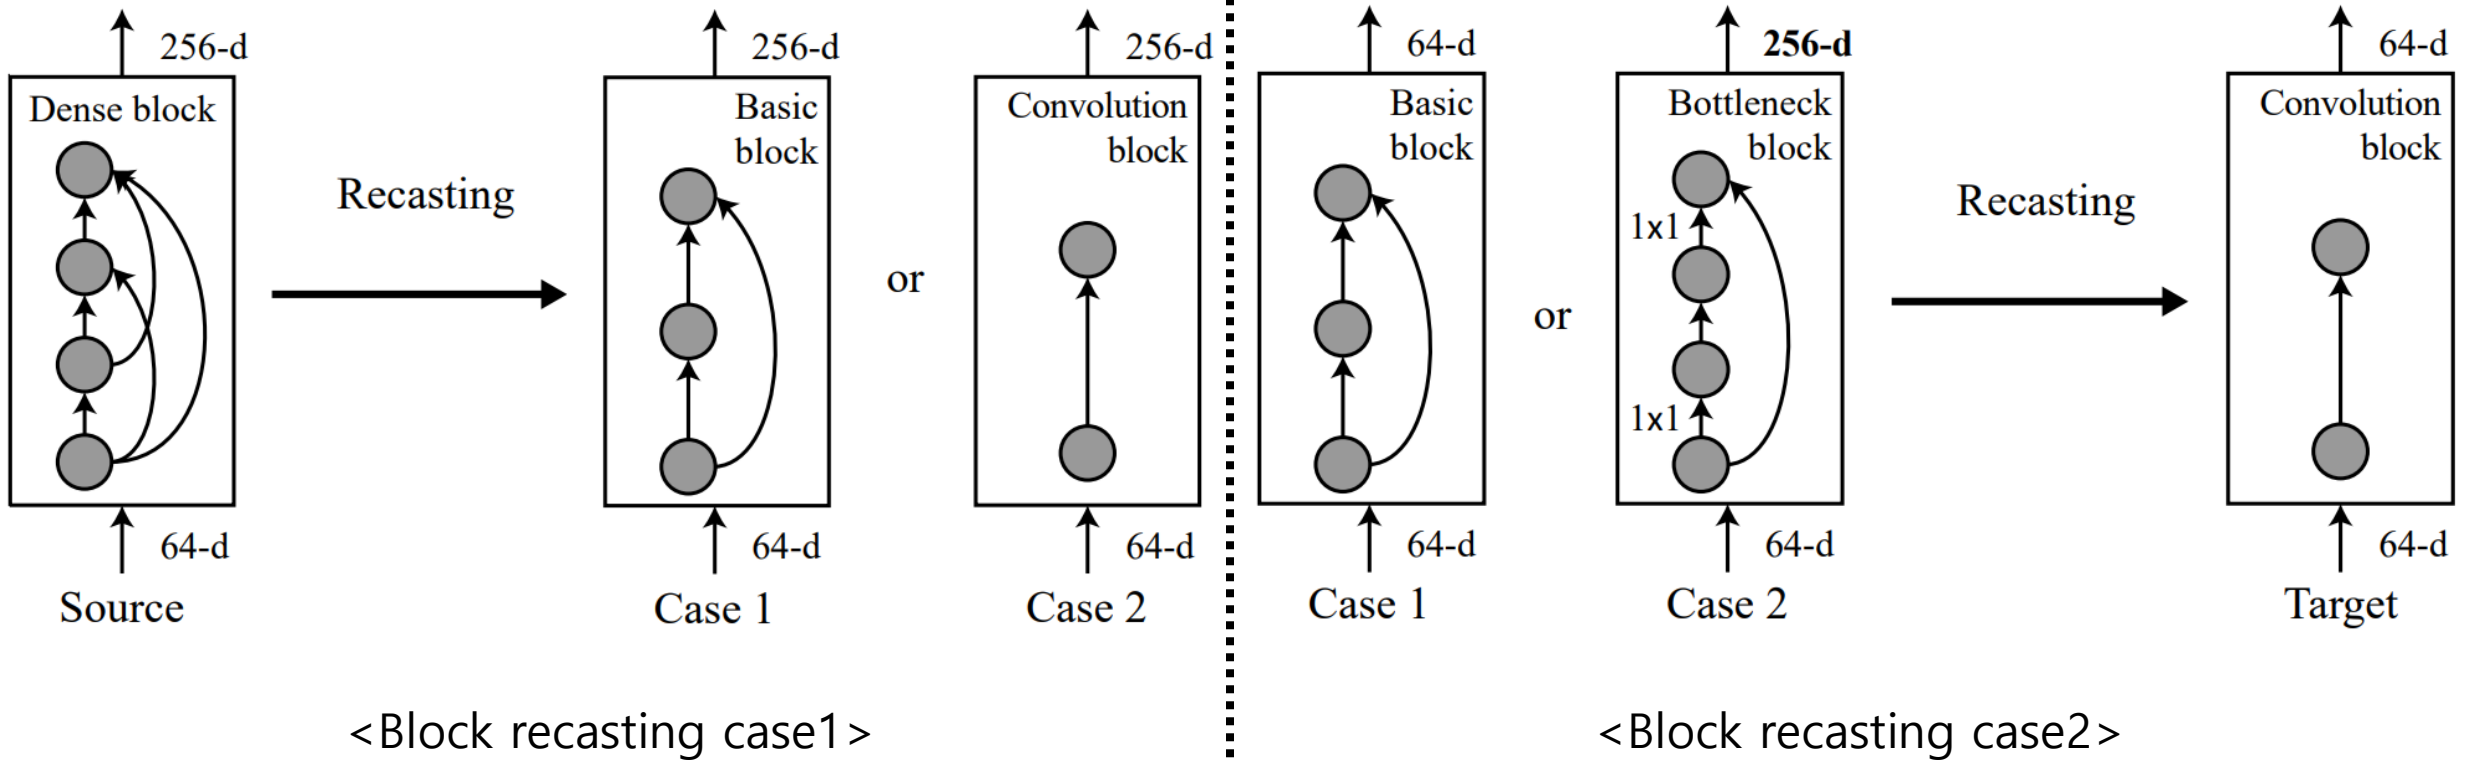
# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

- Joonsang Yu, Sungbum Kang and Kiyoung Choi.

$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

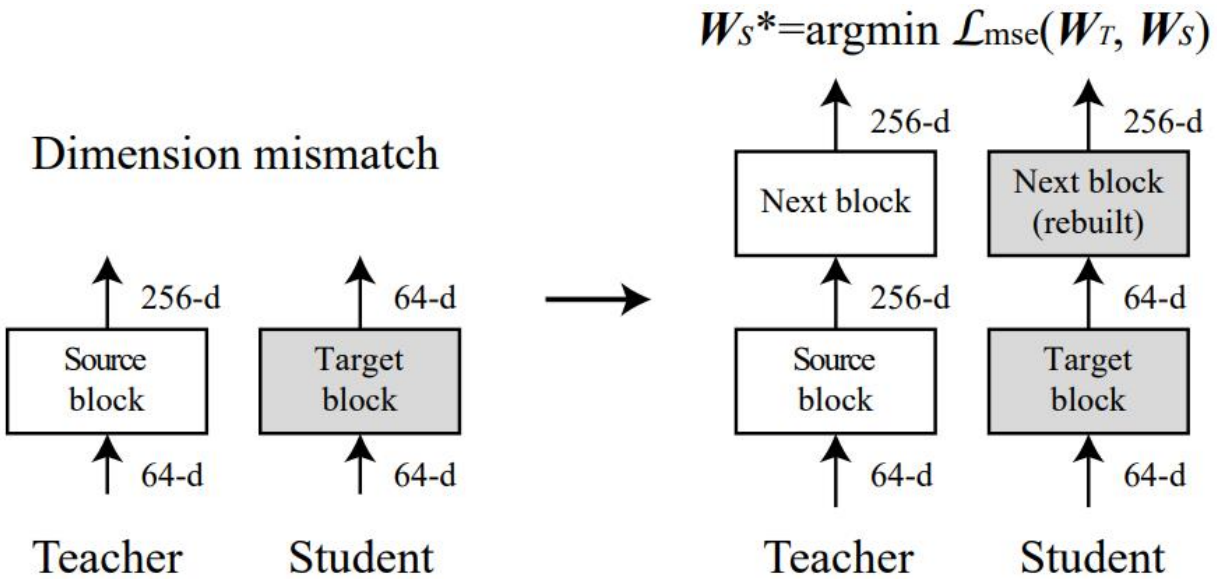
$$\mathcal{L}_{mse}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$



# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.



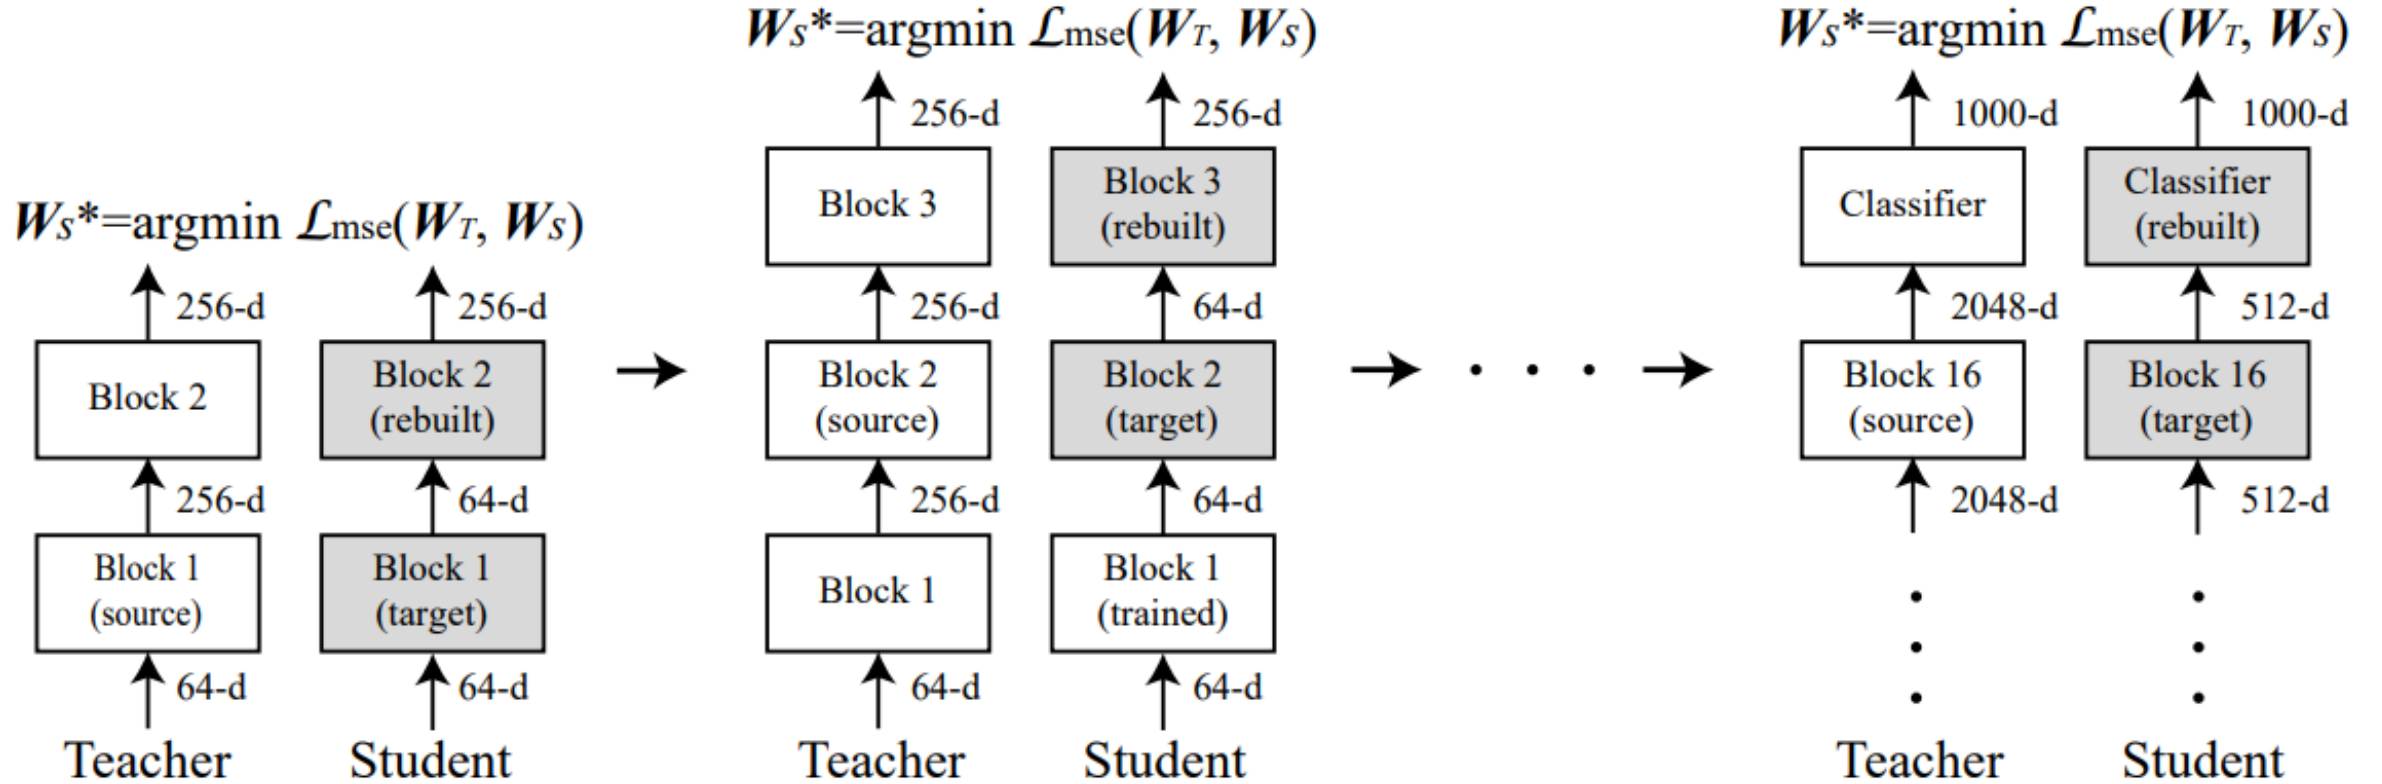
<Method for dimension mismatch>

Recasting Type	Source	Target	Dimension
Transformation	Dense	Basic	Preserved
	Dense	Convolution	Preserved
	Basic	Convolution	Preserved
Compression	Bottleneck	Convolution	Reduced
	Basic	Basic	Reduced
	Convolution	Convolution	Reduced

<Recasting Methods>

# Network Recasting: A Universal Method for Network Architecture Transformation. In AAAI, 2019.

- Joonsang Yu, Sungbum Kang and Kiyoung Choi.



$$\mathcal{L}_{\text{mse}}(W_T, W_S) = \frac{1}{N} \|A(x; W_T) - A(x; W_S)\|_2^2$$

$$\mathcal{L}_{kd}(W_T, W_S) = \mathcal{L}_{\text{mse\_logit}}(W_T, W_S) + \mathcal{L}_{ce}(y_{\text{true}}, W_S)$$

<Network recasting structure>

# Network Recasting: A Universal Method for Network Architecture Transformation. In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyoung Choi.




Figure 6: Visualization of filters in the first layer of AlexNet (*left*) and a student network (*right*). Redundant filters are removed after network recasting.



# Network Recasting: A Universal Method for Network Architecture Transformation.

In *AAAI*, 2019.

– Joonsang Yu, Sungbum Kang and Kiyong Choi.

Table 3: Error rates (%) of compression results on CIFAR datasets. (B/M: billion/million)

Method	Type	C10+	C100+	Params	Mults	Acts/image	Time/image
VGG-16							
Baseline		6.85	28.80	14.71M (1.0×)	313.20M (1.0×)	0.31M (1.0×)	0.37ms
Recasting	Conv	<b>8.31</b>	<b>31.56</b>	2.36M (6.2×)	50.63M (6.2×)	0.13M (2.4×)	0.31ms
KD	Conv	9.24	33.14	2.36M (6.2×)	50.63M (6.2×)	0.13M (2.4×)	0.31ms
Backprop	Conv	8.71	35.13	2.36M (6.2×)	50.63M (6.2×)	0.13M (2.4×)	0.31ms
WRN-28-10							
Baseline		4.06	19.54	36.45M (1.0×)	5.24B (1.0×)	2.52M (1.0×)	0.81ms
Recasting	Basic	<b>5.18</b>	<b>24.13</b>	1.46M (24.9×)	0.21B (24.5×)	0.52M (4.9×)	0.56ms
KD	Basic	5.48	25.28	1.46M (24.9×)	0.21B (24.5×)	0.52M (4.9×)	0.56ms
Backprop	Basic	5.39	25.78	1.46M (24.9×)	0.21B (24.5×)	0.52M (4.9×)	0.56ms


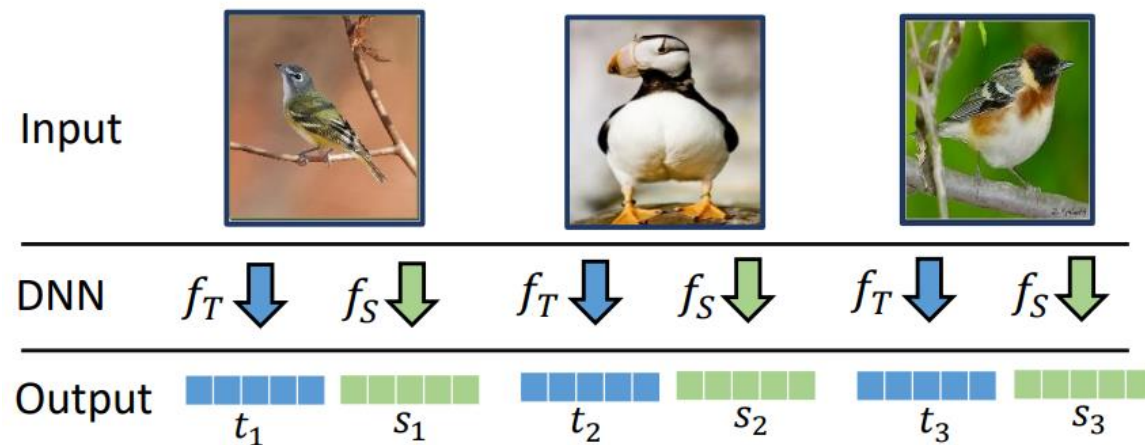
Table 5: Comparison of error rate (%) with previous works on ILSVRC2012. (B/M: billion/million)

Method	Top1	Top5	Params	Mults	Acts/batch	Actual speed-up
ResNet-50						
Recasting(C+R <sub>bt</sub> )	<b>25.00</b>	<b>7.71</b>	21.72M	2.40B	236.16M	<b>2.1</b> ×
ThiNet-30 (Luo, Wu, and Lin 2017)	31.58	11.7	8.66M	1.10B	-	1.3×
AutoPruner ( $r = 0.3$ ) (Luo and Wu 2018)	27.47	8.89	-	1.32B	-	-
VGG-16						
Recasting(C_A)	<b>30.05</b>	<b>10.38</b>	120.61M	3.12B	220.61M	<b>3.2</b> ×
ThiNet-Conv (Luo, Wu, and Lin 2017)	30.20	10.47	131.44M	4.79B	-	2.5×
RNP (3×) (Lin et al. 2017)	-	12.42	-	-	-	2.3×
Channel Pruning (3×) (He, Zhang, and Sun 2017)	-	11.10	-	-	-	2.5×
AutoPruner ( $r = 0.4$ ) (Luo and Wu 2018)	31.57	11.57	-	4.09B	-	-

# Relational Knowledge Distillation


# Relational Knowledge Distillation. In *CVPR*, 2019.

– Wonpyo Park, Dongju Kim, Yan Lu and Minsu Cho.



Point to Point

**Conventional KD**



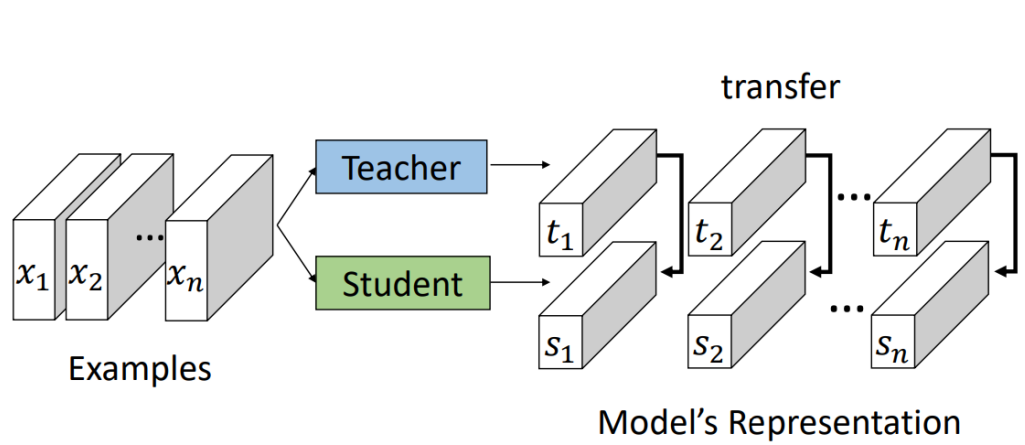
Structure to Structure

**Relational KD**

<Difference between KD and RKD>

# Relational Knowledge Distillation. In *CVPR*, 2019.

– Wonpyo Park, Dongju Kim, Yan Lu and Minsu Cho.



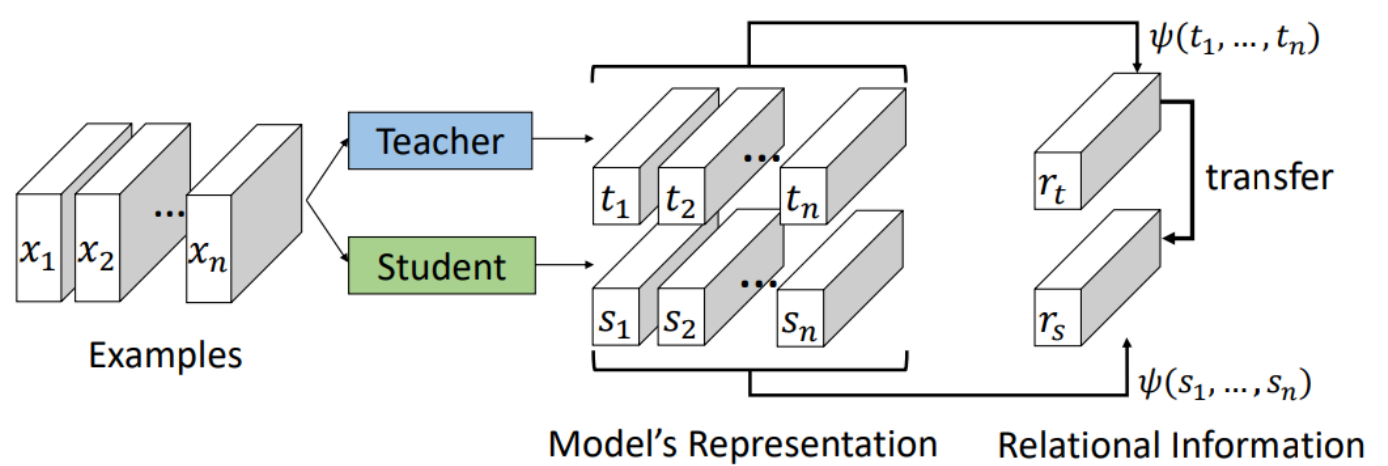
## Individual Knowledge Distillation

$$\psi_D(t_i, t_j) = \frac{1}{\mu} \|t_i - t_j\|_2$$

$$\mu = \frac{1}{|\mathcal{X}^2|} \sum_{(x_i, x_j) \in \mathcal{X}^2} \|t_i - t_j\|_2$$

$$\mathcal{L}_{\text{RKD-D}} = \sum_{(x_i, x_j) \in \mathcal{X}^2} l_\delta(\psi_D(t_i, t_j), \psi_D(s_i, s_j))$$

<Distance-wise distillation loss>



## Relational Knowledge Distillation

$$\psi_A(t_i, t_j, t_k) = \cos \angle t_i t_j t_k = \langle \mathbf{e}^{ij}, \mathbf{e}^{kj} \rangle$$

where  $\mathbf{e}^{ij} = \frac{t_i - t_j}{\|t_i - t_j\|_2}$ ,  $\mathbf{e}^{kj} = \frac{t_k - t_j}{\|t_k - t_j\|_2}$

$$l_\delta(x, y) = \begin{cases} \frac{1}{2}(x - y)^2 & \text{for } |x - y| \leq 1 \\ |x - y| - \frac{1}{2}, & \text{otherwise.} \end{cases}$$

$$\mathcal{L}_{\text{RKD-A}} = \sum_{(x_i, x_j, x_k) \in \mathcal{X}^3} l_\delta(\psi_A(t_i, t_j, t_k), \psi_A(s_i, s_j, s_k))$$

<Angle-wise distillation loss>

# Relational Knowledge Distillation. In *CVPR*, 2019.

– Wonpyo Park, Dongju Kim, Yan Lu and Minsu Cho.

Table 4: Accuracy (%) on CIFAR-100 and Tiny ImageNet.

	CIFAR-100 [15]	Tiny ImageNet [46]
Baseline	71.26	54.45
RKD-D	72.27	54.97
RKD-DA	72.97	56.36
HKD [11]	74.26	57.65
HKD+RKD-DA	<b>74.66</b>	<b>58.15</b>
FitNet [27]	70.81	55.59
FitNet+RKD-DA	72.98	55.54
Attention [47]	72.68	55.51
Attention+RKD-DA	73.53	56.55
Teacher	77.76	61.55

# SUMMARY

- Regularizer
- Teaching & Learning
- Various Fields  
(Model Compression, Transfer Learning, Few shot Learning, Meta Learning...)

# Q&A