

# Laplacian Regularized Few Shot Learning (LaplacianShot)

Imtiaz Masud Ziko, Jose Dolz, Eric Granger and Ismail Ben Ayed



---

ETS Montreal

# Overview

## Few-Shot Learning

- *What and Why ?*
- *Brief discussion on existing approaches.*






## Proposed **LaplacianShot**

- *The context*
- *Proposed formulation*
- *Optimization*
- *Proposed Algorithm*

## Experiments

- *Experimental Setup*
- *SOTA results on 5 different few-shot benchmarks.*

# Few-Shot Learning (An example)

1	2	3	4	5	?	?
						
Support $\mathbb{X}_s$					Query $\mathbb{X}_q$	

# Few-Shot Learning (An example)



Support  $\mathbb{X}_s$



Query  $\mathbb{X}_q$

- Given  $C = \mathbf{5}$  classes
- Each class  $c$  having  $\mathbf{1}$  examples.

**(5-way 1-shot)**

From  
these

**Learn a  
Model**

To classify  
this

# Few-Shot Learning (An example)



Support  $\mathbb{X}_s$



Query  $\mathbb{X}_q$

- Given  $C = \mathbf{5}$  classes
- Each class  $c$  having  $\mathbf{1}$  examples.

**(5-way 1-shot)**

From  
these

**Learn a  
Model**

To classify  
this

# Few-Shot Learning



Support  $\mathbb{X}_s$



Query  $\mathbb{X}_q$

**Humans recognize perfectly**  
with few examples



# Few-Shot Learning



- ❑ Modern ML methods **generalize poorly**
- ❑ Need a better way.

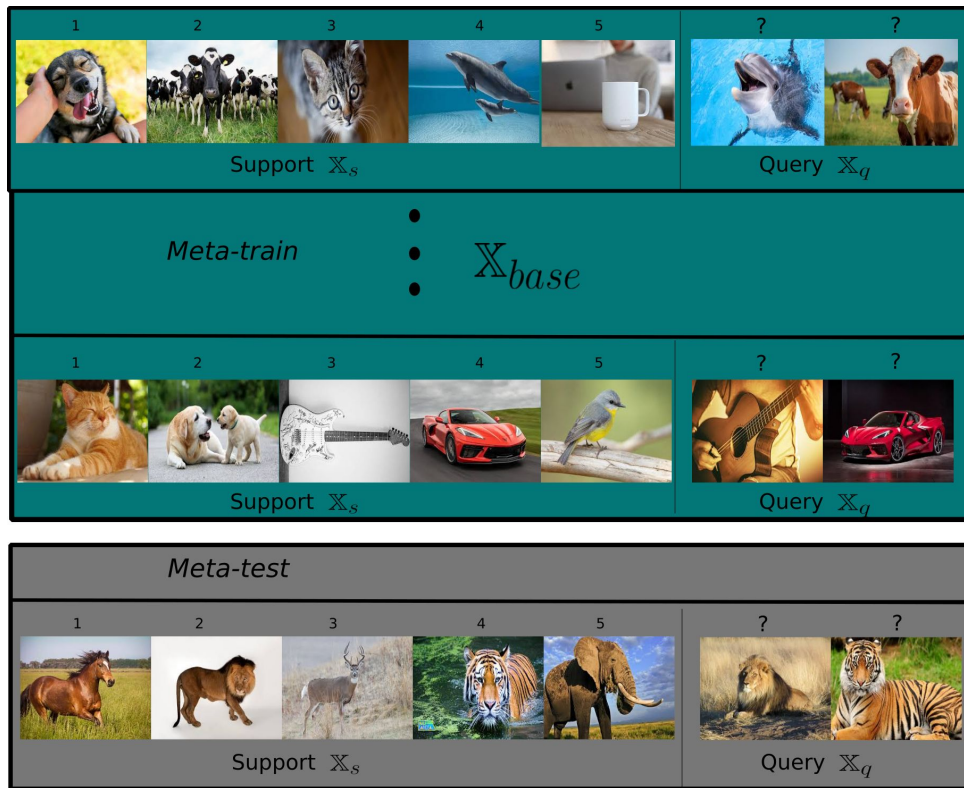
# Few-shot learning

A very large body of recent works, mostly based on:

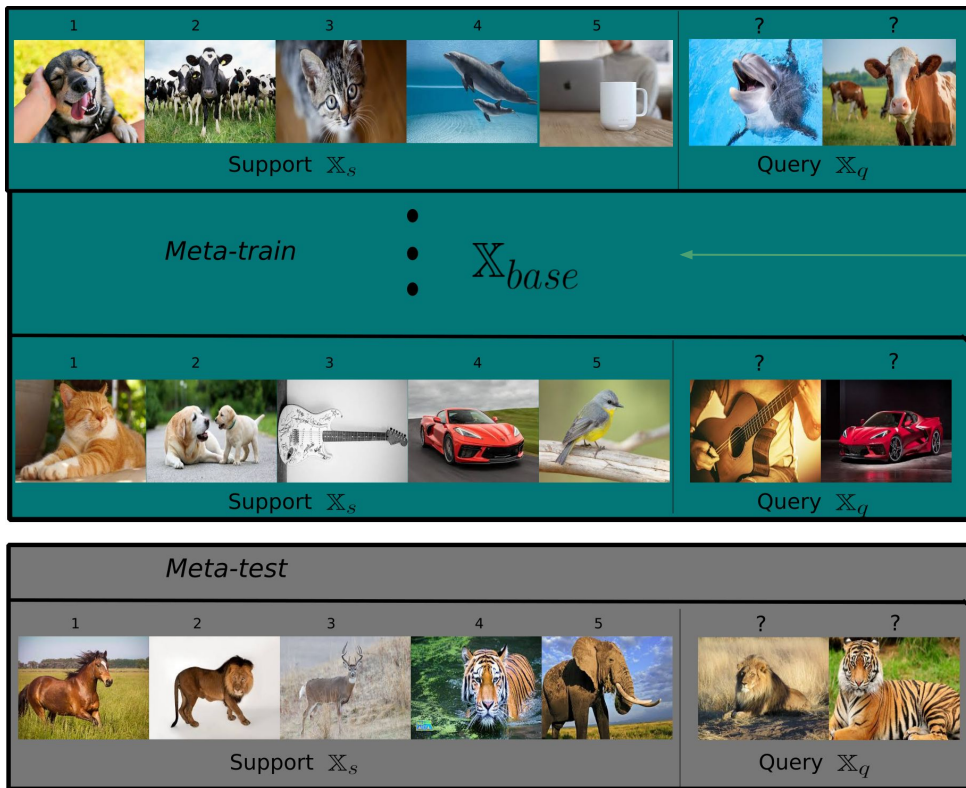
**Meta-learning framework**



# Meta-Learning Framework



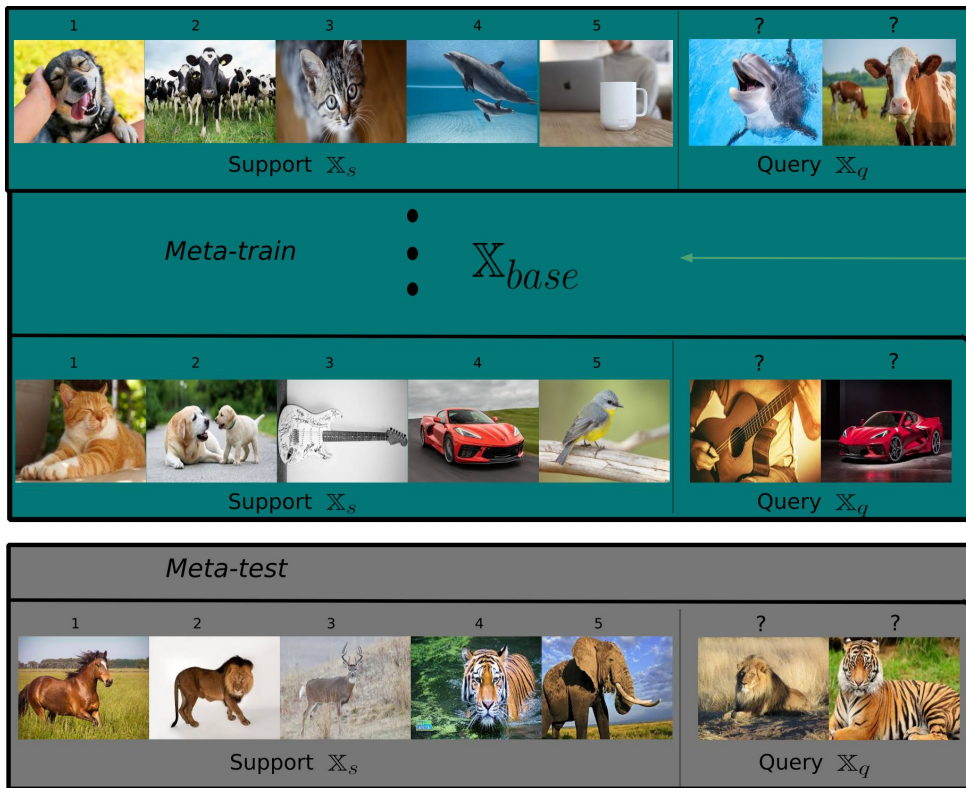
# Meta-Learning Framework



Training set with  
enough labeled data

(base classes ***different***  
***from the*** test classes)

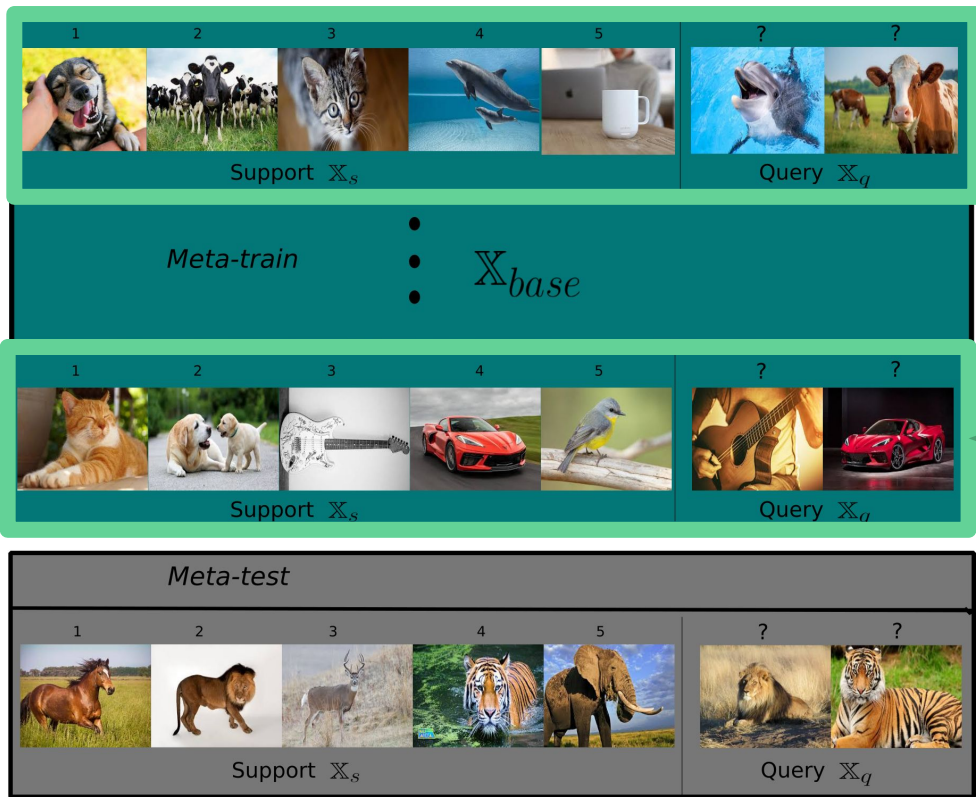
# Meta-Learning Framework



Training set with  
enough labeled data  
**to learn initial model**

$$f_{\theta}$$

# Meta-Learning Framework



Create episodes and do **episodic training** to learn **meta-learner**

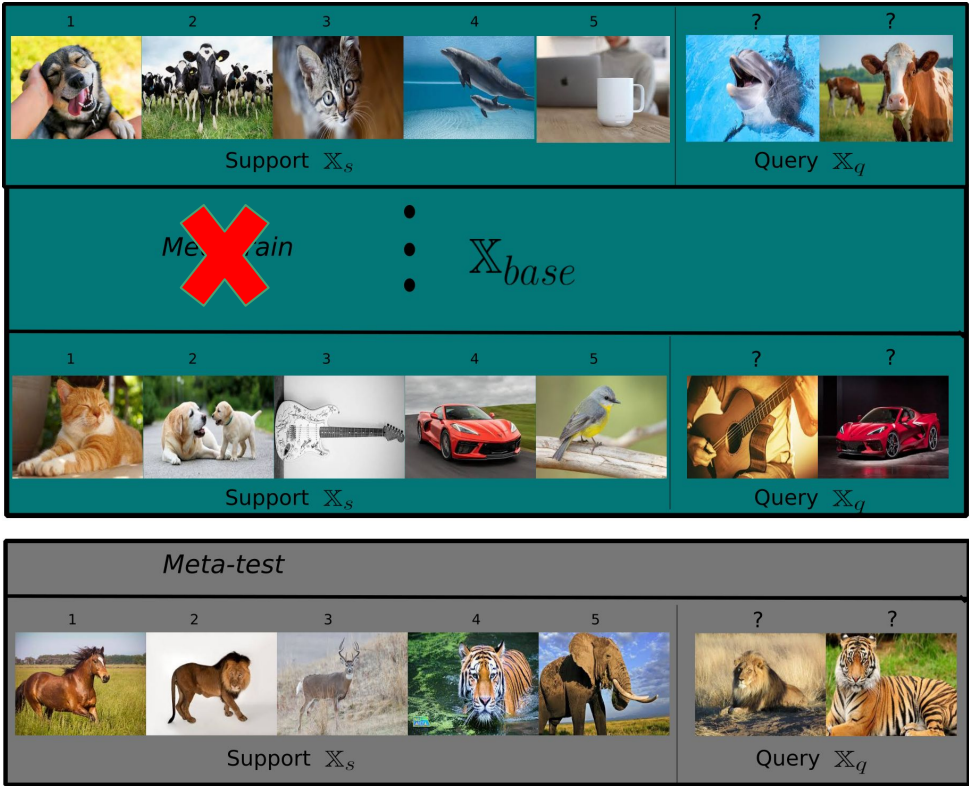
Vinyal et al, (Neurips '16),  
Snell et al, (Neurips '17),  
Sung et al, (CVPR '18),  
Finn et al, (ICML' 17),  
Ravi et al, (ICLR' 17),  
Lee et al, (CVPR' 19),  
Hu et al, (ICLR '20),  
Ye et al, (CVPR '20), . . .

# Taking a few steps backward . .

**Recently** [Chen et al., ICLR'19, Wang et al., '19, Dhillon et al., ICLR'20] :

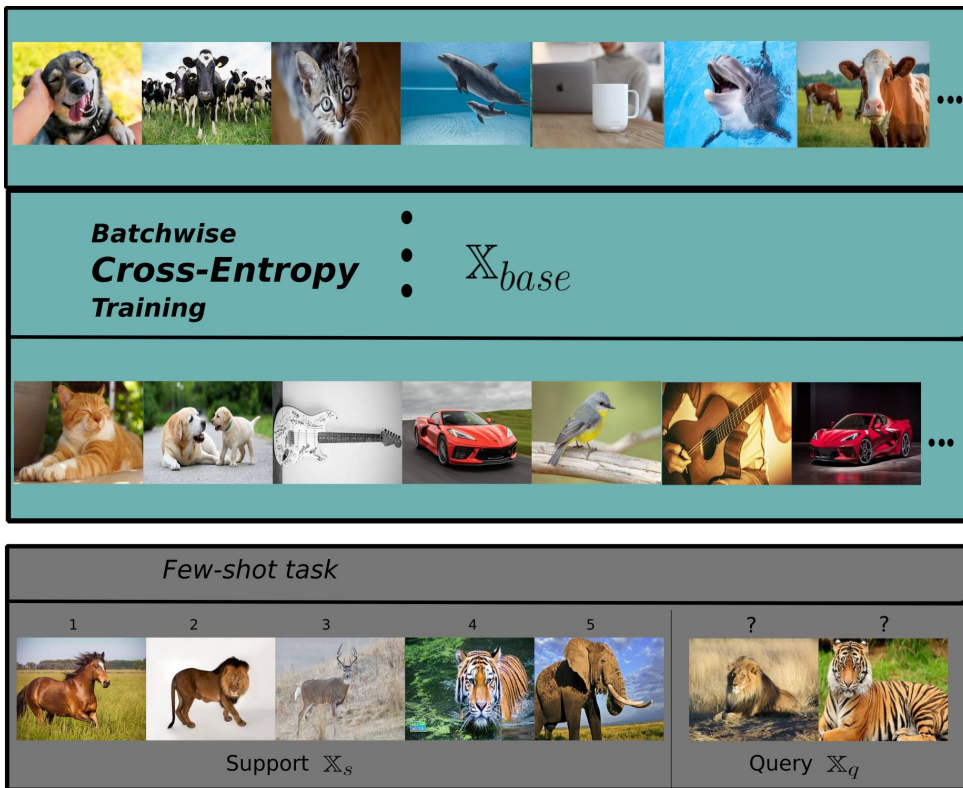
Simple baselines **outperform** the overly convoluted **meta-learning** based approaches.

# Baseline Framework



No need to *meta-train*

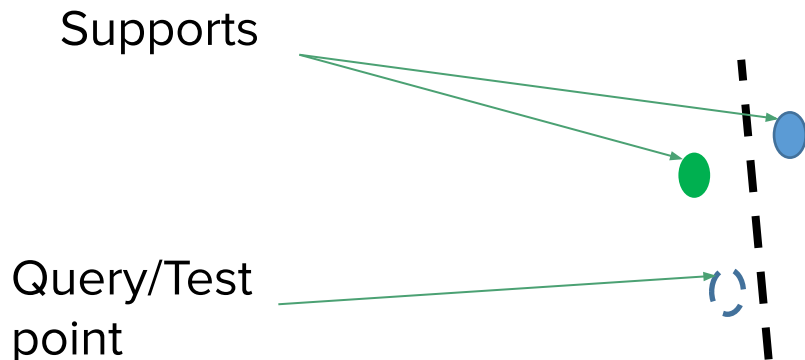
# Baseline Framework



**Simple**  
**conventional**  
cross-entropy  
training

The approaches  
mostly differ  
**during inference**

# Inductive vs Transductive inference



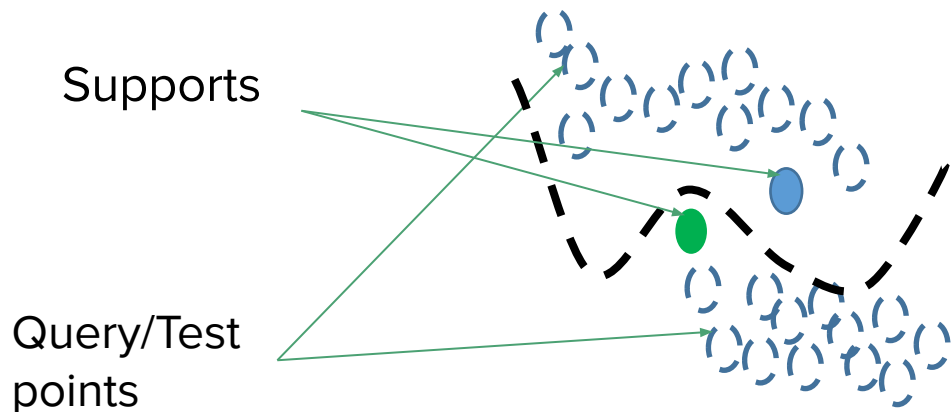
## Examples

Vinayls et al., NEURIPS' 16  
(Attention mechanism)

Snell et al., NEURIPS' 17  
(Nearest Prototype)



# Inductive vs Transductive inference



**Transductive:** Predict for all test points, instead of one at a time

## Examples

Liu et. al., ICLR'19  
(Label propagation)

Dhillon, ICLR'20  
(Transductive  
fine-tuning)

# Proposed LaplacianShot

Laplacian-regularized  
objective:

$$\mathcal{E}(\mathbf{Y}) = \mathcal{N}(\mathbf{Y}) + \frac{\lambda}{2} \mathcal{L}(\mathbf{Y})$$

- Latent Assignment matrix for N query samples:

$$\mathbf{Y} = [\mathbf{y}_q] \in \{0, 1\}^{N \times C}$$

- Label assignment for each query:

$$\mathbf{y}_q = [y_{q,1}, \dots, y_{q,C}]^t \in \{0, 1\}^C$$

- And Simplex Constraints:

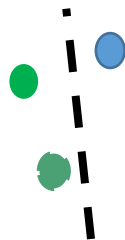
$$\mathbf{1}^t \mathbf{y}_q = 1$$

# Proposed LaplacianShot

Laplacian-regularized objective:

$$\mathcal{E}(\mathbf{Y}) = \boxed{\mathcal{N}(\mathbf{Y})} + \frac{\lambda}{2} \boxed{\mathcal{L}(\mathbf{Y})}$$

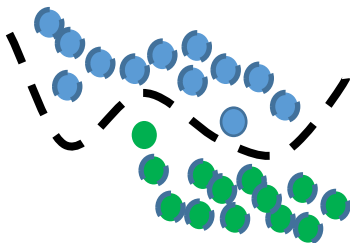
**Nearest Prototype classification**



When  $\lambda = 0$

Similar to ProtoNet (Snell '17) or SimpleShot (Wang '19)

**Laplacian Regularization**



Well known in Graph Laplacian:  
Spectral clustering (Shi '00, Von '07), SLK (Ziko '18)  
SSL (Weston '12, Belkin '06)

# LaplacianShot Takeaways

- ✓ **SOTA results** without bell and whistles.
- ✓ Simple **constrained graph clustering** works very well.
- ✓ **No** network **fine-tuning**, neither **meta-learning**
- ✓ **Model Agnostic**
- ✓ **Fast** transductive inference: almost inductive time

LapLacianShot

**More Details**

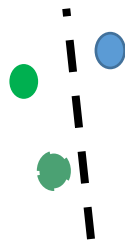
# Proposed LaplacianShot

Laplacian-regularized objective:

$$\mathcal{E}(\mathbf{Y}) = \boxed{\mathcal{N}(\mathbf{Y})} + \frac{\lambda}{2} \mathcal{L}(\mathbf{Y})$$

$$\mathcal{N}(\mathbf{Y}) = \sum_{q=1}^N \sum_{c=1}^C y_{q,c} d(\mathbf{x}_q - \mathbf{m}_c)$$

Nearest Prototype classification



When  $\lambda = 0$

Labeling according to nearest support prototypes

- Feature embedding:  $\mathbf{x}_q = f_{\theta}(x_q)$

- Prototype  $\mathbf{m}_c$  can be :

- **The support example** in 1-shot or
- **Simple mean** from support examples or
- **Weighted mean** from both support and initially predicted query samples

# Proposed LaplacianShot

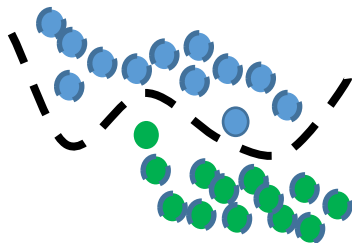
Laplacian-regularized objective:

$$\mathcal{E}(\mathbf{Y}) = \mathcal{N}(\mathbf{Y}) + \frac{\lambda}{2} \mathcal{L}(\mathbf{Y})$$

Pairwise similarity

$$\mathcal{L}(\mathbf{Y}) = \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \|\mathbf{y}_q - \mathbf{y}_p\|^2$$

Laplacian Regularization



Well known in Graph Laplacian:

**Encourages nearby points to have similar assignments**

# Proposed Optimization

Laplacian-regularized  
objective:

$$\mathcal{E}(\mathbf{Y}) = \mathcal{N}(\mathbf{Y}) + \frac{\lambda}{2} \mathcal{L}(\mathbf{Y})$$

$$\mathcal{N}(\mathbf{Y}) = \sum_{q=1}^N \sum_{c=1}^C y_{q,c} d(\mathbf{x}_q - \mathbf{m}_c)$$

$$\mathcal{L}(\mathbf{Y}) = \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \|\mathbf{y}_q - \mathbf{y}_p\|^2$$

Tricky to optimize due to:



# Proposed Optimization

Laplacian-regularized  
objective:

$$\mathcal{E}(\mathbf{Y}) = \mathcal{N}(\mathbf{Y}) + \frac{\lambda}{2} \mathcal{L}(\mathbf{Y})$$

$$\mathcal{N}(\mathbf{Y}) = \sum_{q=1}^N \sum_{c=1}^C y_{q,c} d(\mathbf{x}_q - \mathbf{m}_c)$$

$$\mathcal{L}(\mathbf{Y}) = \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \|\mathbf{y}_q - \mathbf{y}_p\|^2$$

Tricky to optimize due to:

 Simplex/Integer Constraints.

$$\mathbf{Y} = [\mathbf{y}_q] \in \{0, 1\}^{N \times C}$$
$$\mathbf{1}^t \mathbf{y}_q = 1$$

# Proposed Optimization

Laplacian-regularized  
objective:

$$\mathcal{E}(\mathbf{Y}) = \mathcal{N}(\mathbf{Y}) + \frac{\lambda}{2} \mathcal{L}(\mathbf{Y})$$

$$\mathcal{N}(\mathbf{Y}) = \sum_{q=1}^N \sum_{c=1}^C y_{q,c} d(\mathbf{x}_q - \mathbf{m}_c)$$

$$\mathcal{L}(\mathbf{Y}) = \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \|\mathbf{y}_q - \mathbf{y}_p\|^2$$

Tricky to optimize due to:

 Laplacian over discrete variables.

# Proposed Optimization

Laplacian-regularized objective:

$$\mathcal{E}(\mathbf{Y}) = \mathcal{N}(\mathbf{Y}) + \frac{\lambda}{2} \mathcal{L}(\mathbf{Y})$$

$$\mathcal{N}(\mathbf{Y}) = \sum_{q=1}^N \sum_{c=1}^C y_{q,c} d(\mathbf{x}_q - \mathbf{m}_c)$$

$$\mathcal{L}(\mathbf{Y}) = \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \|\mathbf{y}_q - \mathbf{y}_p\|^2$$

Relax integer constraints:

➤ Convex quadratic problem



Require solving for the N×C variables all together



Extra projection steps for the simplex constraints

# Proposed Optimization

Laplacian-regularized  
objective:

$$\mathcal{E}(\mathbf{Y}) = \mathcal{N}(\mathbf{Y}) + \frac{\lambda}{2} \mathcal{L}(\mathbf{Y})$$

$$\mathcal{N}(\mathbf{Y}) = \sum_{q=1}^N \sum_{c=1}^C y_{q,c} d(\mathbf{x}_q - \mathbf{m}_c)$$

$$\mathcal{L}(\mathbf{Y}) = \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \|\mathbf{y}_q - \mathbf{y}_p\|^2$$

We do:

✓ Concave relaxation

- ✓ Independent and closed-form updates for each assignment variable
- ✓ Efficient bound optimization

# Concave Laplacian

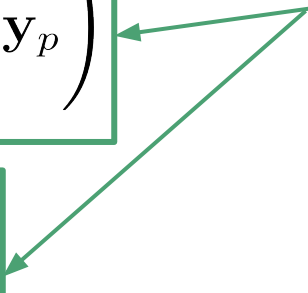
$$\mathcal{L}(\mathbf{Y}) = \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \|\mathbf{y}_q - \mathbf{y}_p\|^2$$

# Concave Laplacian

When  $\mathbf{y}_q \in \{0, 1\}^C$

$$\begin{aligned}\mathcal{L}(\mathbf{Y}) &= 2 \left( \sum_q \mathbf{y}_q^t \mathbf{y}_q D_q - \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \mathbf{y}_q^t \mathbf{y}_p \right) \\ &= 2 \left( \sum_q D_q - \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \mathbf{y}_q^t \mathbf{y}_q \right)\end{aligned}$$

**Equal**



# Concave Laplacian

When  $\mathbf{y}_q \in [0, 1]^C$

$$\mathcal{L}(\mathbf{Y}) = 2 \left( \sum_q \mathbf{y}_q^t \mathbf{y}_q D_q - \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \mathbf{y}_q^t \mathbf{y}_p \right)$$

$$\neq 2 \left( \sum_q D_q - \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \mathbf{y}_q^t \mathbf{y}_q \right)$$

**Not  
Equal**

# Concave Laplacian

When  $\mathbf{y}_q \in [0, 1]^C$

$$\mathcal{L}(\mathbf{Y}) = 2 \left( \sum_q \mathbf{y}_q^t \mathbf{y}_q D_q - \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \mathbf{y}_q^t \mathbf{y}_p \right)$$

**Not  
Equal**

$$\neq 2 \left( \sum_q D_q - \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \mathbf{y}_q^t \mathbf{y}_q \right)$$

$$\text{Degree } D_q = \sum_p w(\mathbf{x}_q, \mathbf{x}_p)$$



# Concave Laplacian

Remove constant terms

$$\mathcal{L}(\mathbf{Y}) = 2 \left( \sum_q D_q - \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \mathbf{y}_q^t \mathbf{y}_p \right)$$

# Concave Laplacian

$$\mathcal{L}(\mathbf{Y}) = -2 \sum_{q,p} w(\mathbf{x}_q, \mathbf{x}_p) \mathbf{y}_q^t \mathbf{y}_p$$

Concave for PSD matrix  $\mathbf{W} = [w(\mathbf{x}_q, \mathbf{x}_p)]$

# Concave-Convex relaxation

Putting it altogether

$$\mathcal{R}(\mathbf{Y}) = \boxed{\mathbf{Y}^t \log \mathbf{Y}} + \mathcal{N}(\mathbf{Y}) + \lambda \mathcal{L}(\mathbf{Y})$$

Convex barrier function:

- Avoids extra dual variables for  $\mathbf{y}_q \geq 0$
- Closed-form update for the simplex constraint dual

# Bound optimization

First-order approximation  
of concave term

$$\mathcal{R}(\mathbf{Y}) = \mathbf{Y}^t \log \mathbf{Y} + \mathcal{N}(\mathbf{Y}) + \lambda \mathcal{L}(\mathbf{Y})$$

Fixed unary

# Bound optimization

We get **Iterative tight upper bound**:

$$B_i(\mathbf{Y}) = \sum_{q=1}^N \mathbf{y}_q^t (\log(\mathbf{y}_q) + \mathbf{a}_q - \lambda \mathbf{b}_q^i)$$

Iteratively optimize:

$$\mathbf{Y}^{i+1} = \arg \min_{\mathbf{Y}} \mathcal{B}_i(\mathbf{Y})$$

**Where:**

$$\mathbf{a}_q = [a_{q,1}, \dots, a_{q,C}]^t; \quad a_{q,c} = d(\mathbf{x}_q, \mathbf{m}_c)$$

$$\mathbf{b}_q^i = [b_{q,1}^i, \dots, b_{q,C}^i]^t; \quad b_{q,c}^i = \sum_p w(\mathbf{x}_q, \mathbf{x}_p) y_{p,c}^i$$

# Bound optimization

Independent **upper bound**:

$$B_i(\mathbf{Y}) = \sum_{q=1}^N \boxed{\mathbf{y}_q^t (\log(\mathbf{y}_q) + \mathbf{a}_q - \lambda \mathbf{b}_q^i)}$$

# Bound optimization

Minimize Independent **upper bound**:

$$\min_{\mathbf{y}_q \in \nabla_C} \mathbf{y}_q^t (\log(\mathbf{y}_q) + \mathbf{a}_q - \lambda \mathbf{b}_q^i), \forall q$$

KKT conditions brings **closed form updates**:

$$\mathbf{y}_q^{i+1} = \frac{\exp(-\mathbf{a}_q^i + \lambda \mathbf{b}_q^i)}{\mathbf{1}^t \exp(-\mathbf{a}_q^i + \lambda \mathbf{b}_q^i)} \quad \forall q$$

# LaplacianShot Algorithm

**Input:**  $\mathbb{X}_s, \mathbb{X}_q, \lambda, f_\theta$

**Output:** Labels  $\in \{1, \dots, C\}^N$  for  $\mathbb{X}_q$

Get prototypes  $\mathbf{m}_c$ .

Compute  $\mathbf{a}_q \forall \mathbf{x}_q \in \mathbb{X}_q$ .

Initialize  $i = 1$ .

Initialize  $\mathbf{y}_q^i = \frac{\exp(-\mathbf{a}_q)}{\mathbf{1}^t \exp(-\mathbf{a}_q)}$ .

**repeat**

Compute  $\mathbf{y}_q^{i+1}$

$\mathbf{y}_q^i \leftarrow \mathbf{y}_q^{i+1}$ .

$\mathbf{Y} = [\mathbf{y}_q^i]; \forall q$ .

$i = i + 1$ .

**until**  $\mathcal{B}_i(\mathbf{Y})$  does not change

$l_q = \arg \max_c \mathbf{y}_q; \forall \mathbf{y}_q \in \mathbf{Y}$ .

Labels =  $\{l_q\}_{q=1}^N$

$$\mathbf{a}_q = [a_{q,1}, \dots, a_{q,C}]^t; a_{q,c} = d(\mathbf{x}_q, \mathbf{m}_c)$$

$$\mathbf{b}_q^i = [b_{q,1}^i, \dots, b_{q,C}^i]^t; b_{q,c}^i = \sum_p w(\mathbf{x}_q, \mathbf{x}_p) y_{p,c}^i$$

$$\mathbf{y}_q^{i+1} = \frac{\exp(-\mathbf{a}_q^i + \lambda \mathbf{b}_q^i)}{\mathbf{1}^t \exp(-\mathbf{a}_q^i + \lambda \mathbf{b}_q^i)} \quad \forall q$$



# Experiments

Datasets:

1. *Mini-ImageNet*
2. *Tierd-ImageNet*
3. CUB 200-2001
4. Inat

## Generic Classification

***miniImageNet* splits:** 64 base, 16 validation and 20 test classes

***tieredImageNet* splits:** 351 base, 97 validation and 160 test classes

## Fine-Grained Classification

**Splits:** 100 base, 50 validation and 50 test classes

# Experiments

## Datasets:

1. *Mini-ImageNet*
2. *Tierd-ImageNet*
3. CUB 200-2001
4. Inat

## Evaluation protocol:

- **5**-way **1**-shot/**5**-shot .
- **15** query samples per class (N=75).
- **Average accuracy** over **10,000** few-shot tasks with 95% confidence interval.

# Experiments

## Datasets:

1. *Mini-ImageNet*
2. *Tierd-ImageNet*
3. CUB 200-2001
4. Inat

- More realistic and challenging
- Recently introduced (Wertheimer& Hariharan, 2019)
- Slight class distinction
- **Imbalanced class distribution** with variable number of supports/query per class

# Experiments

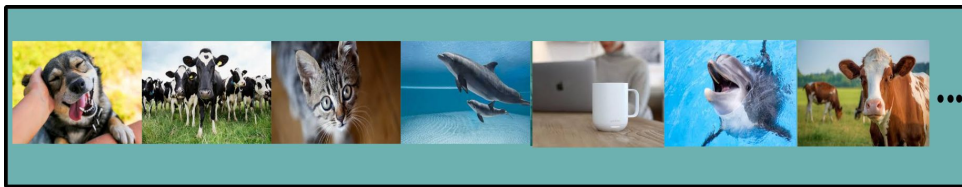
## Datasets:

1. *Mini-ImageNet*
2. *Tierd-ImageNet*
3. CUB 200-2001
4. Inat

## Evaluation protocol:

- **227-way multi-shot** .
- **Top-1 accuracy** averaged over the test images **Per Class**.
- **Top-1 accuracy** averaged over **all the test images (Mean)**

# Experiments



**Batchwise  
Cross-Entropy  
Training**

•  
•  
•  $\mathbb{X}_{base}$



*Few-shot task*



Support  $\mathbb{X}_s$

Query  $\mathbb{X}_q$

We do  
**Cross-entropy  
training with base  
classes**

**LaplacianShot  
during inference**

# Results (Mini-ImageNet)

Methods	Network	1-shot	5-shot
MAML [Finn et al., 2017]	ResNet-18	49.61 $\pm$ 0.92	65.72 $\pm$ 0.77
Chen [Chen et al., 2019]	ResNet-18	51.87 $\pm$ 0.77	75.68 $\pm$ 0.63
RelationNet [Sung et al., 2018]	ResNet-18	52.48 $\pm$ 0.86	69.83 $\pm$ 0.68
MatchingNet [Vinyals et al., 2016]	ResNet-18	52.91 $\pm$ 0.88	68.88 $\pm$ 0.69
ProtoNet [Snell et al., 2017]	ResNet-18	54.16 $\pm$ 0.82	73.68 $\pm$ 0.65
Gidaris [Gidaris and Komodakis, 2018]	ResNet-15	55.45 $\pm$ 0.89	70.13 $\pm$ 0.68
SNAIL[Mishra et al., 2018]	ResNet-15	55.71 $\pm$ 0.99	68.88 $\pm$ 0.92
AdaCNN [Munkhdalai et al., 2018]	ResNet-15	56.88 $\pm$ 0.62	71.94 $\pm$ 0.57
TADAM [Oreshkin et al., 2018]	ResNet-15	58.50 $\pm$ 0.30	76.70 $\pm$ 0.30
CAML [Jiang et al., 2019]	ResNet-12	59.23 $\pm$ 0.99	72.35 $\pm$ 0.71
TPN [Yanbin et al., 2019]	ResNet-12	59.46	75.64
TEAM [Qiao et al., 2019]	ResNet-18	60.07	75.90
MTL [Sun et al., 2019]	ResNet-18	61.20 $\pm$ 1.80	75.50 $\pm$ 0.80
VariationalFSL [Zhang et al., 2019]	ResNet-18	61.23 $\pm$ 0.26	77.69 $\pm$ 0.17
Transductive tuning [Dhillon et al., 2020]	ResNet-12	62.35 $\pm$ 0.66	74.53 $\pm$ 0.54
MetaoptNet[Lee et al., 2019]	ResNet-18	62.64 $\pm$ 0.61	78.63 $\pm$ 0.46
SimpleShot [Wang et al., 2019]	ResNet-18	63.10 $\pm$ 0.20	79.92 $\pm$ 0.14
▼ CAN+T [Hou et al., 2019]	ResNet-12	67.19 $\pm$ 0.55	80.64 $\pm$ 0.35
LaplacianShot (ours)	ResNet-18	<b>72.11 <math>\pm</math> 0.19</b>	<b>82.31 <math>\pm</math> 0.14</b>

# Results (Mini-ImageNet)

Methods	Network	1-shot	5-shot
Qiao [Qiao et al., 2018]	WRN	59.60 $\pm$ 0.41	73.74 $\pm$ 0.19
LEO [Rusu et al., 2019]	WRN	61.76 $\pm$ 0.08	77.59 $\pm$ 0.12
ProtoNet [Snell et al., 2017]	WRN	62.60 $\pm$ 0.20	79.97 $\pm$ 0.14
CC+rot[Gidaris et al., 2019]	WRN	62.93 $\pm$ 0.45	79.87 $\pm$ 0.33
MatchingNet [Vinyals et al., 2016]	WRN	64.03 $\pm$ 0.20	76.32 $\pm$ 0.16
FEAT [Ye et al., 2020]	WRN	65.10 $\pm$ 0.20	81.11 $\pm$ 0.14
Transductive tuning [Dhillon et al., 2020]	WRN	65.73 $\pm$ 0.68	78.40 $\pm$ 0.52
SimpleShot [Wang et al., 2019]	WRN	65.87 $\pm$ 0.20	82.09 $\pm$ 0.14
SIB [Hu et al., 2020]	WRN	70.0 $\pm$ 0.6	79.2 $\pm$ 0.4
BD-CSPN [Liu et al., 2019]	WRN	70.31 $\pm$ 0.93	81.89 $\pm$ 0.60
LaplacianShot (ours)	WRN	<b>74.86 <math>\pm</math> 0.19</b>	<b>84.13 <math>\pm</math> 0.14</b>

# Results (Tiered-ImageNet)

Methods	Network	1-shot	5-shot
MetaoptNet [Lee et al., 2019]	ResNet-18	65.99 $\pm$ 0.72	81.56 $\pm$ 0.53
SimpleShot [Wang et al., 2019]	ResNet-18	69.68 $\pm$ 0.22	84.56 $\pm$ 0.16
CAN+T [Hou et al., 2019]	ResNet-12	73.21 $\pm$ 0.58	84.93 $\pm$ 0.38
LaplacianShot (ours)	ResNet-18	<b>78.98</b> $\pm$ 0.21	<b>86.39</b> $\pm$ 0.16
Meta SGD [Li et al., 2017]	WRN	62.95 $\pm$ 0.03	79.34 $\pm$ 0.06
LEO [Rusu et al., 2019]	WRN	66.33 $\pm$ 0.05	81.44 $\pm$ 0.09
FEAT [Ye et al., 2020]	WRN	70.41 $\pm$ 0.23	84.38 $\pm$ 0.16
CC+rot [Gidaris et al., 2019]	WRN	70.53 $\pm$ 0.51	84.98 $\pm$ 0.36
SimpleShot [Wang et al., 2019]	WRN	70.90 $\pm$ 0.22	85.76 $\pm$ 0.15
Transductive tuning [Dhillon et al., 2020]	WRN	73.34 $\pm$ 0.71	85.50 $\pm$ 0.50
BD-CSPN [Liu et al., 2019]	WRN	78.74 $\pm$ 0.95	86.92 $\pm$ 0.63
LaplacianShot (ours)	WRN	<b>80.18</b> $\pm$ 0.21	<b>87.56</b> $\pm$ 0.15



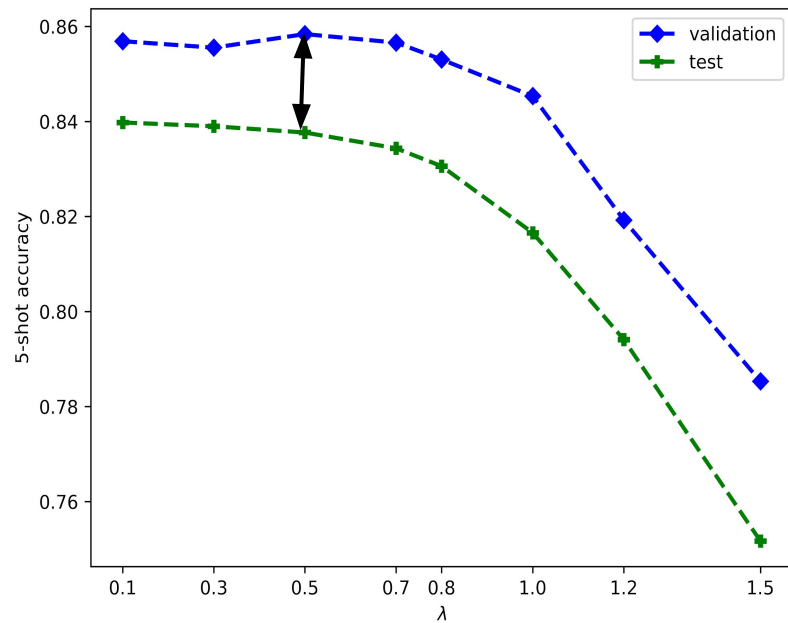
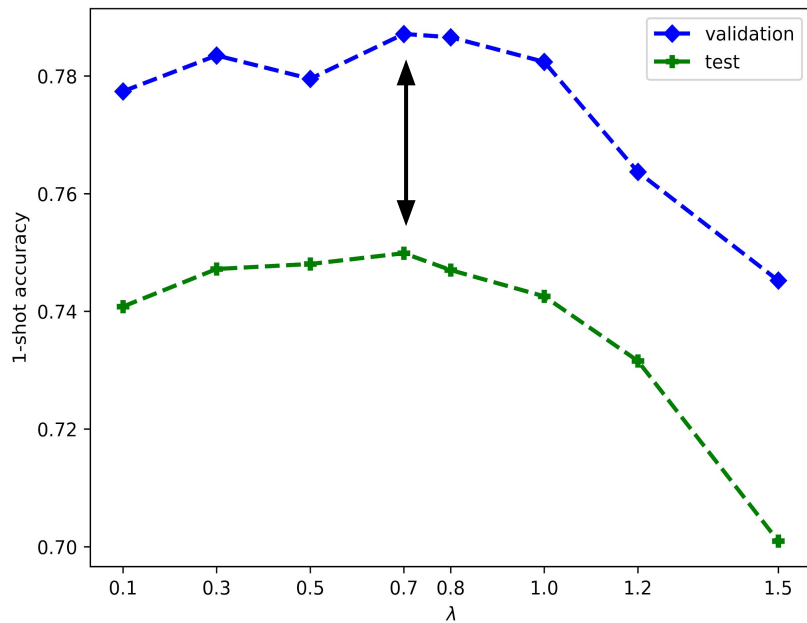
# Results (CUB)

Methods	Network	CUB		<i>miniImagenet</i> → CUB	
		1-shot	5-shot	1-shot	5-shot
MatchingNet [Vinyals et al., 2016]	ResNet-18	73.49	84.45	-	53.07
MAML [Finn et al., 2017]	ResNet-18	68.42	83.47	-	51.34
ProtoNet [Snell et al., 2017]	ResNet-18	72.99	86.64	-	62.02
RelationNet [Sung et al., 2018]	ResNet-18	68.58	84.05	-	57.71
Chen [Chen et al., 2019]	ResNet-18	67.02	83.58	-	65.57
SimpleShot [Wang et al., 2019]	ResNet-18	70.28	86.37	48.56	65.63
LaplacianShot(ours)	ResNet-18	<b>80.96</b>	<b>88.38</b>	<b>55.46</b>	<b>66.33</b>
<b>Cross Domain</b>					

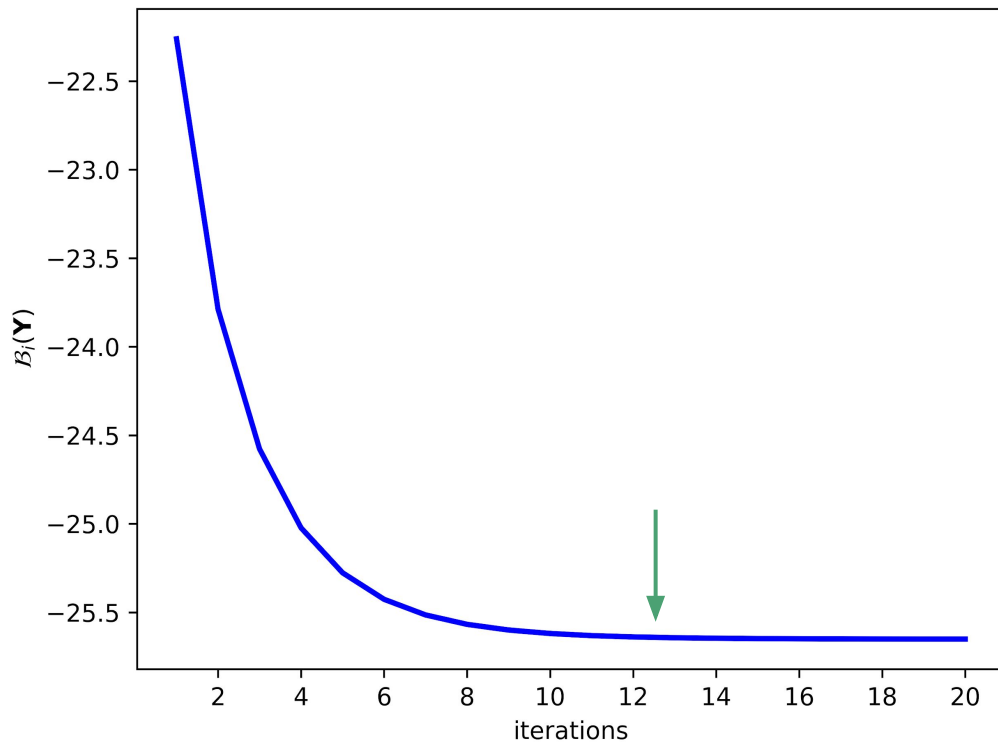
# Results (iNat)

Methods	Network	Per Class	Mean
SimpleShot [Wang et al., 2019]	ResNet-18	55.80	58.56
LaplacianShot	ResNet-18	<b>62.80</b>	<b>66.40</b>
PN+BF+fsL+CP [Wertheimer and Hariharan, 2019]	ResNet-50	46.04	51.25
SimpleShot [Wang et al., 2019]	ResNet-50	58.45	61.07
LaplacianShot	ResNet-50	<b>65.96</b>	<b>69.13</b>
SimpleShot [Wang et al., 2019]	WRN	62.44	65.08
LaplacianShot	WRN	<b>71.55</b>	<b>74.97</b>

# Ablation: Choosing $\lambda$



# Ablation: Convergence



# Ablation: Average Inference time

Methods	Network	inference time (s)
SimpleShot [Wang et al., 2019]	WRN	0.009
Transductive tuning [Dhillon et al., 2020]	WRN	20.7
LaplacianShot	WRN	0.012

Transductive

# LaplacianShot Takeaways

- ✓ **SOTA results** without bell and whistles.
- ✓ Simple **constrained graph clustering** works very well.
- ✓ **No** network **fine-tuning**, neither **meta-learning**
- ✓ **Model Agnostic**: during inference with any training model and gain up to 4/5%!!!
- ✓ **Fast** transductive inference: almost inductive time

# Thank you

Code On:

<https://github.com/imtiaziko/LaplacianShot>