SOSELETO: A Unified Approach to Transfer Learning and Training with Noisy Labels

Or Litany, Daniel Freedman
orlitany@gmail.com, danielfreedman@google.com

Not all train samples are created equal

Key observation: Some source examples are more informative than others for the target classification problem

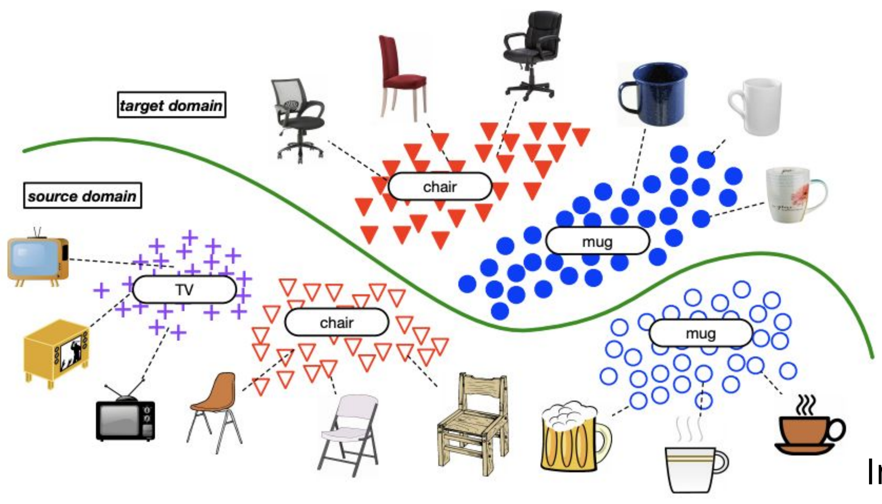

Image credit: Cao et al., 2018

Problem: We do not know *a priori* which source examples will be important

SOSELETO: We let the target decide!

SOurce SElection for Target Optimization

***Source* (many) and *Target* (few) share a representation**

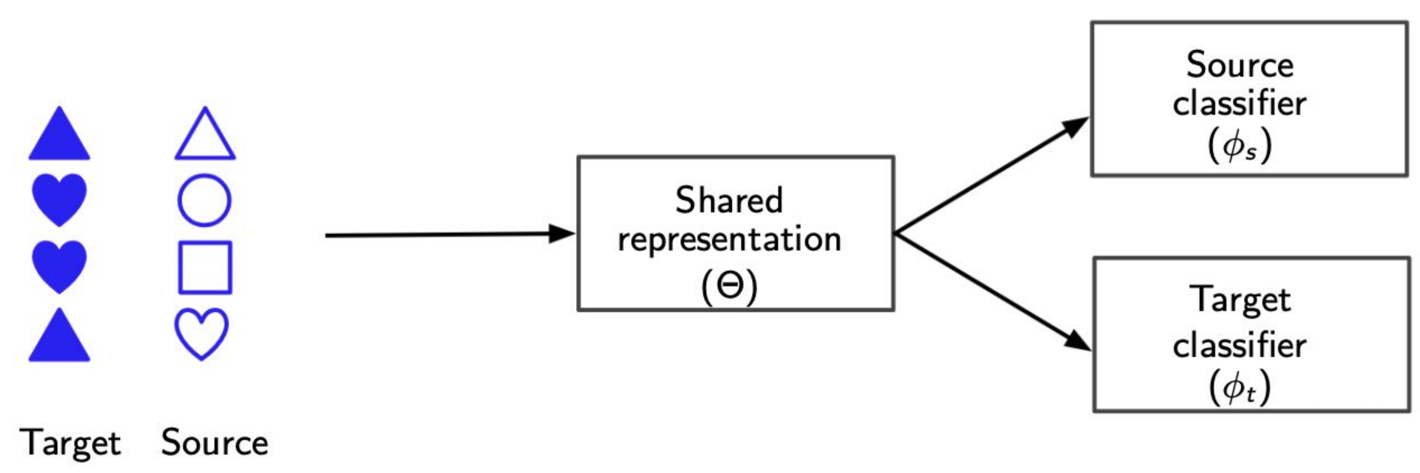

Interior level:
Source samples *directly* modify the representation

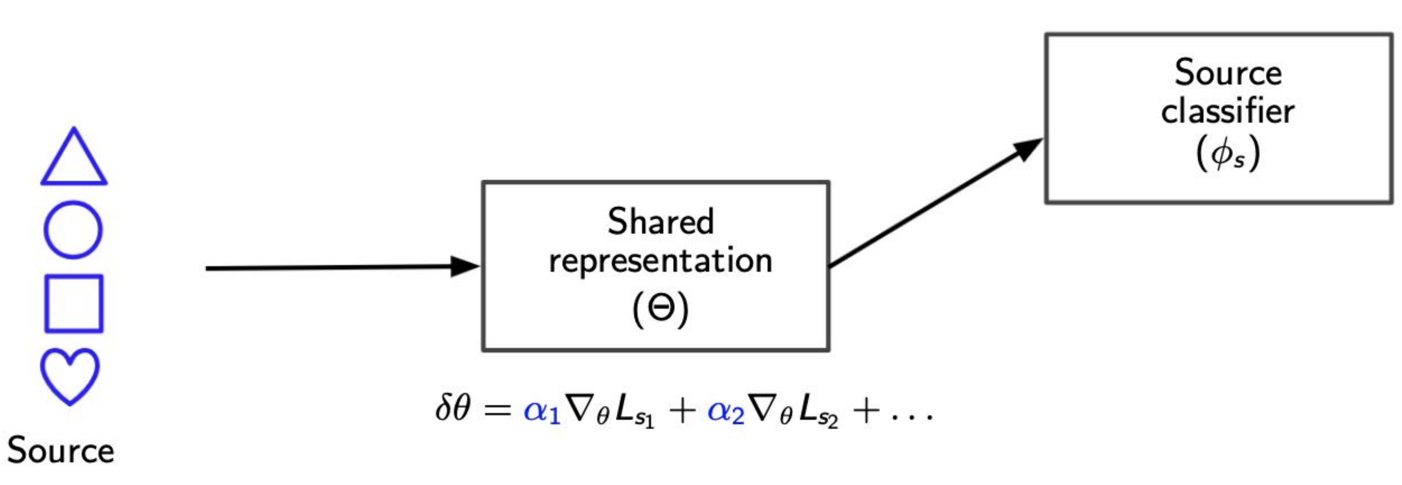

Optimize: $\theta^*(\alpha), \phi^{s*}(\alpha) = \min_{\theta, \phi^s} L_s(\theta, \phi^s, \alpha) \;\; ; \;\; L_s(\theta, \phi^s, \alpha) = \frac{1}{n^s} \sum_{j=1}^{n^s} \alpha_j \ell(y_j^s, F(x_j^s; \theta, \phi^s))$

Exterior level:
Target samples modify representation via source weights

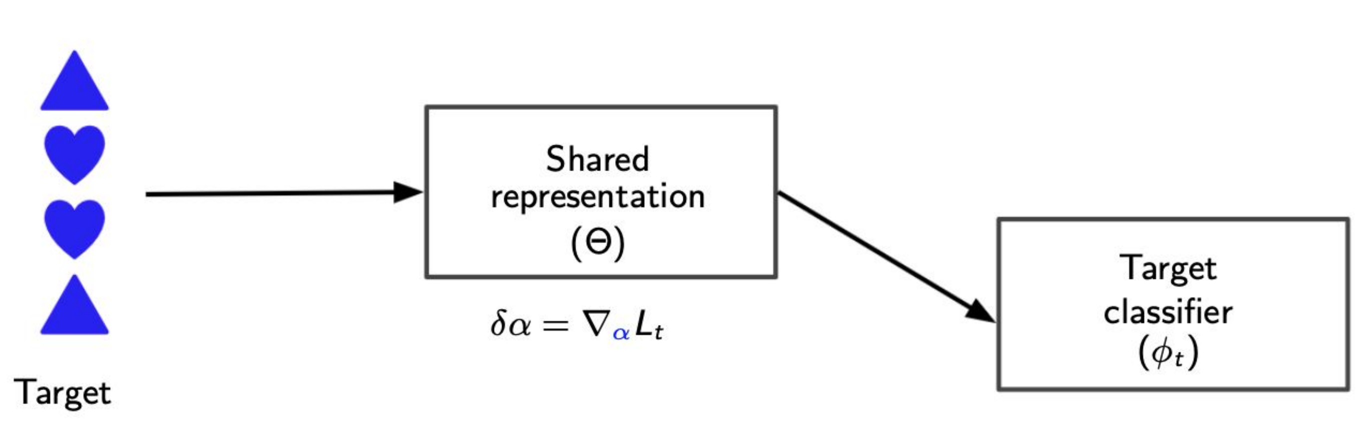

Optimize: $\min_{\alpha, \phi^t} L_t(\theta^*(\alpha), \phi^t) \;\; ; \;\; L_t(\theta, \phi^t) = \frac{1}{n^t} \sum_{i=1}^{n^t} \ell(y_i^t, F(x_i^t; \theta, \phi^t))$

Noisy labels: synthetic experiment

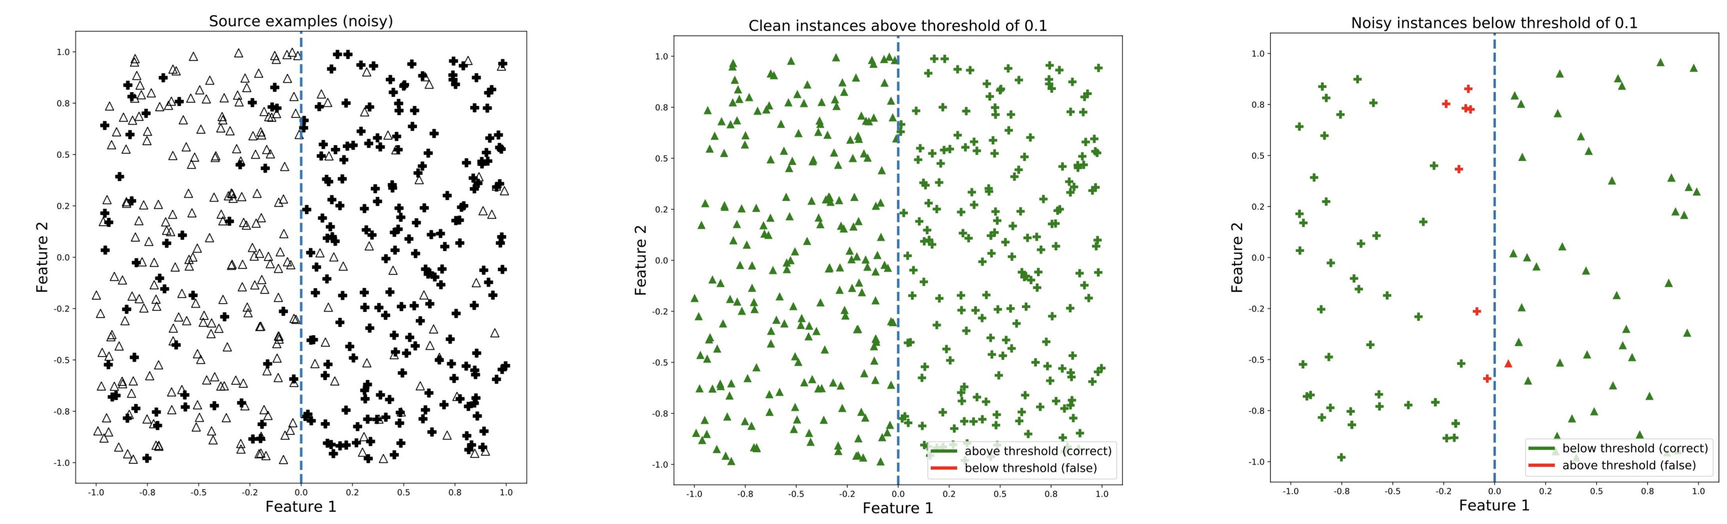

Inputs:
Source: 500 points with 20% noisy labels
Target: 50 points with clean labels
Result:
correctly labeled inputs are assigned large weights
noisy inputs are assigned small weights

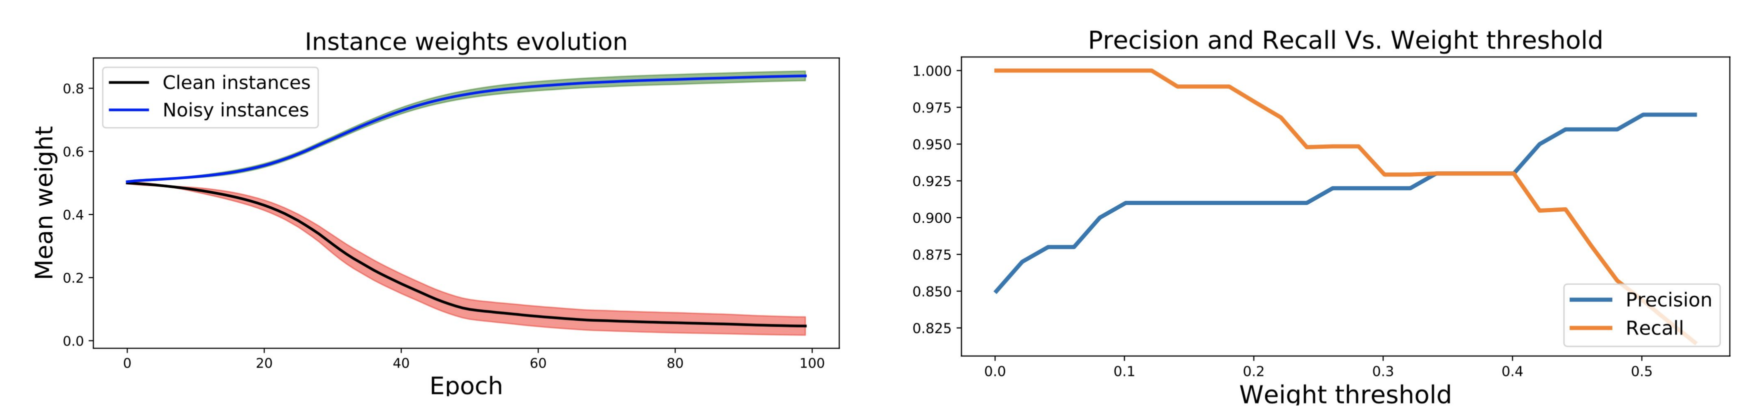

Noisy labels: CIFAR-10

Input: 50K samples, of which 30-50% are noisy.

Noise Level	CIFAR-10 Quick –	Sukhbaatar et al. (2014) 10K clean examples	Xiao et al. (2015) 10K clean examples	Ours 5K clean examples
30%	65.57	69.73	69.81	**72.41**
40%	62.38	66.66	66.76	**69.98**
50%	57.36	63.39	63.00	**66.33**

Data filtering:

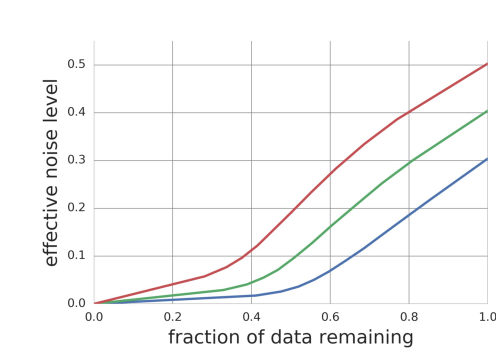

Transfer: SVHN to MNIST

No class overlap: SVHN 0-4 to MNIST 5-9

Uses Unlabelled Data?	Method	$n^t = 20$	$n^t = 25$
No	Target only	80.1	84.0
No	Fine-tuning	80.2	83.0
No	SOSELETO	**83.2**	**87.9**
Yes	Vinyals et al. (2016)	56.6	51.3
Yes	Fine-tuned variant of Vinyals et al. (2016)	79.3	82.7
Yes	Luo et al. (2017)	80.4	83.1
Yes	Label-efficient version of Luo et al. (2017)	**94.2**	**95.0**

Leveraging *unlabeled* MNIST: +4%
Partial overlap (SVHN 0-9) boosts performance to 90.3%
Sample weight corresponds to appearance (and not class)