Adapted Deep Embeddings: A Synthesis of Methods for k-Shot Inductive Transfer Learning

Tyler R. Scott[1,2], Karl Ridgeway[1,2], Michael C. Mozer[1,3]

[1] *University of Colorado, Boulder*
[2] *Sensory Inc.*
[3] *Presently at Google Brain*

Inductive Transfer Learning

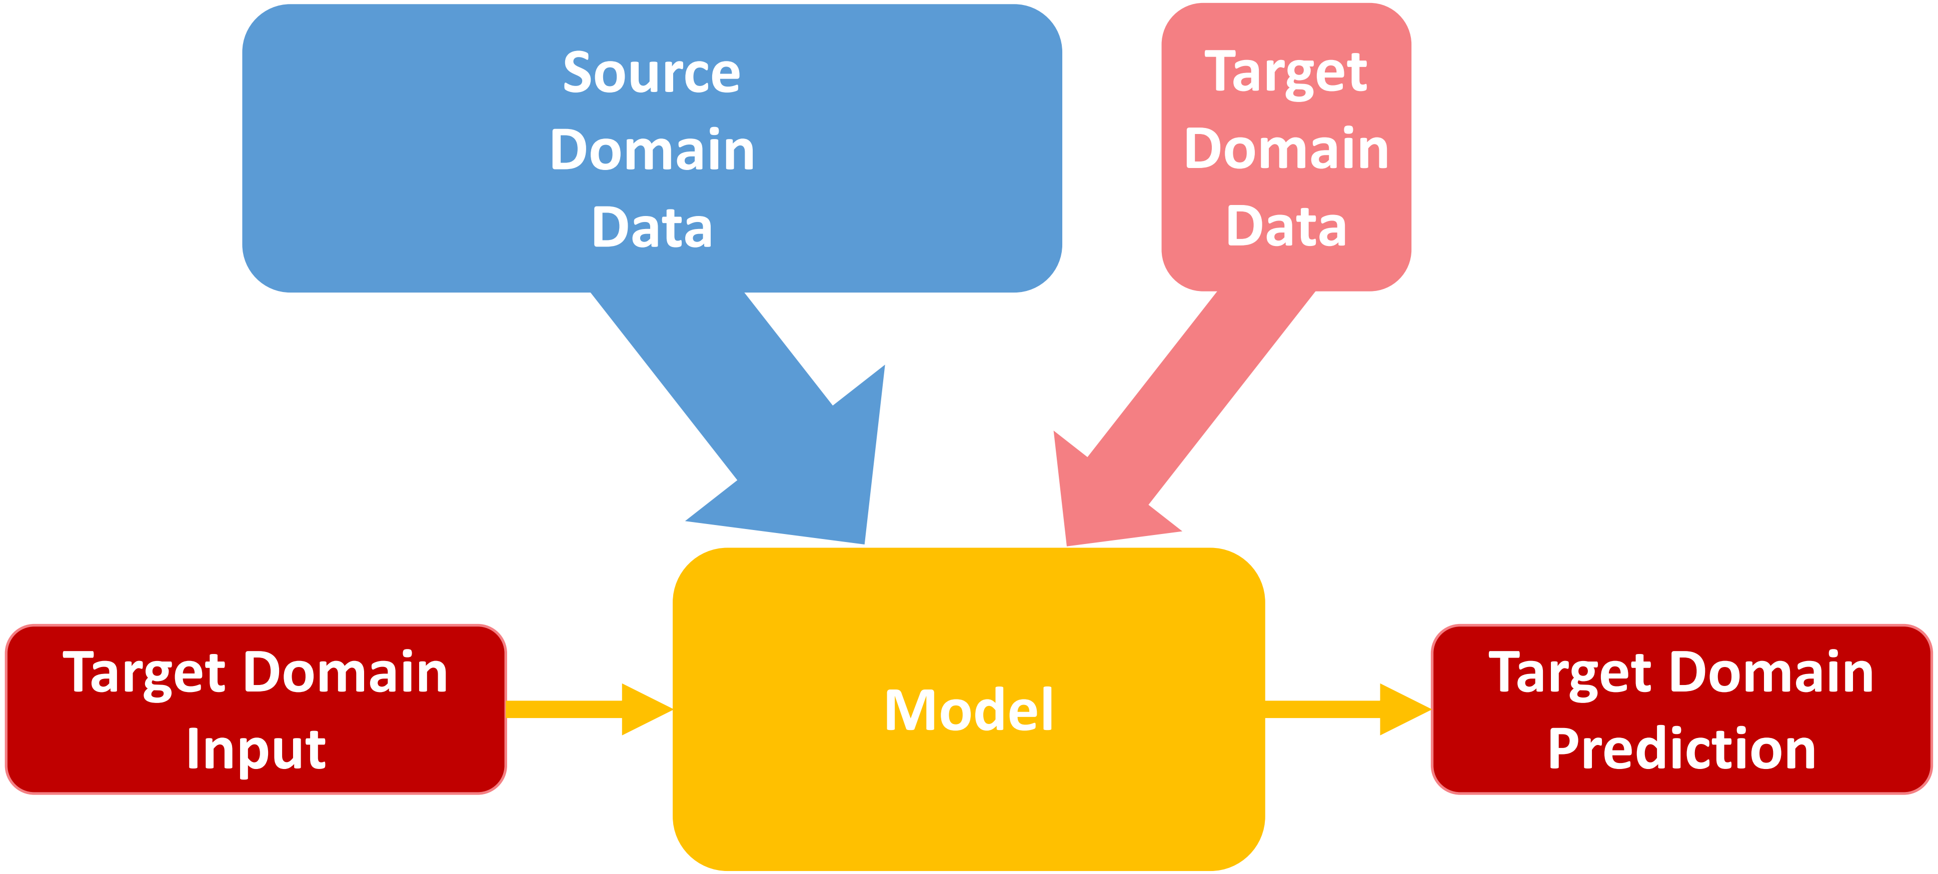

Inductive Transfer Learning

Weight Transfer

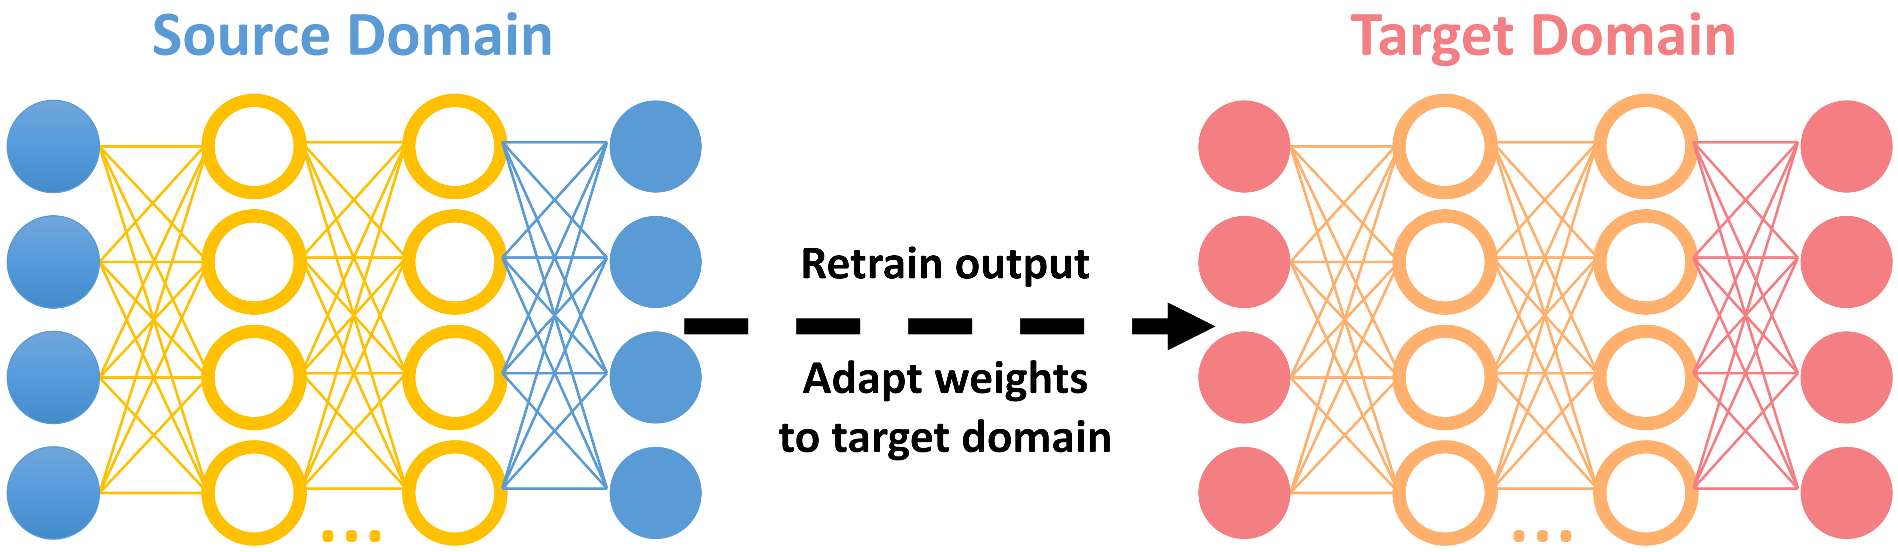

Yosinski et al. (2014)

Inductive Transfer Learning

Weight Transfer

Deep Metric Learning

Source Domain

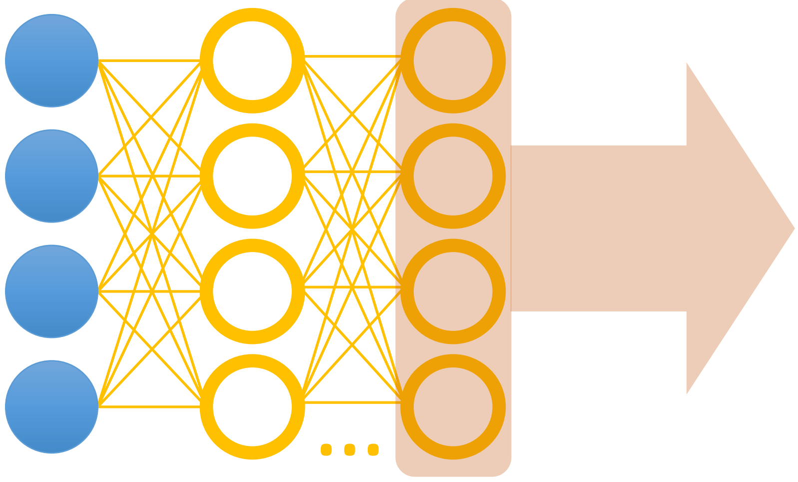

Source & Target
Domain Embedding

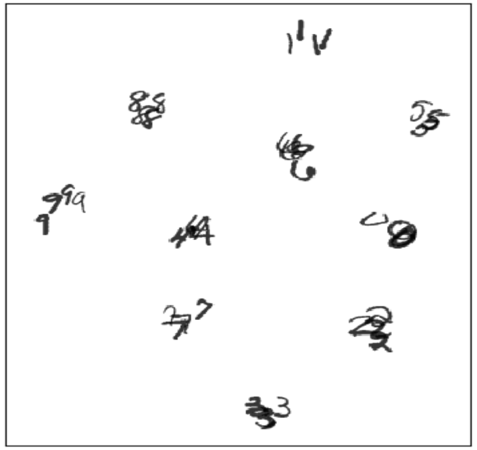

Histogram loss

(Ustinova & Lempitsky, 2016)

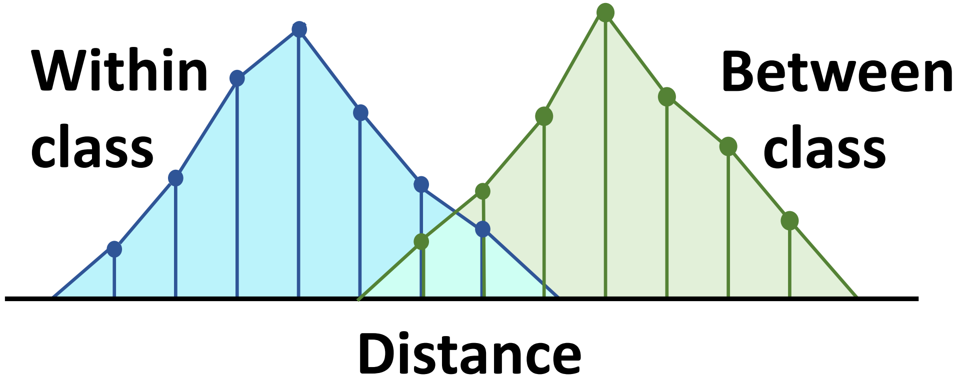

Inductive Transfer Learning

Weight Transfer

Deep Metric Learning

Few-Shot Learning

Source Domain

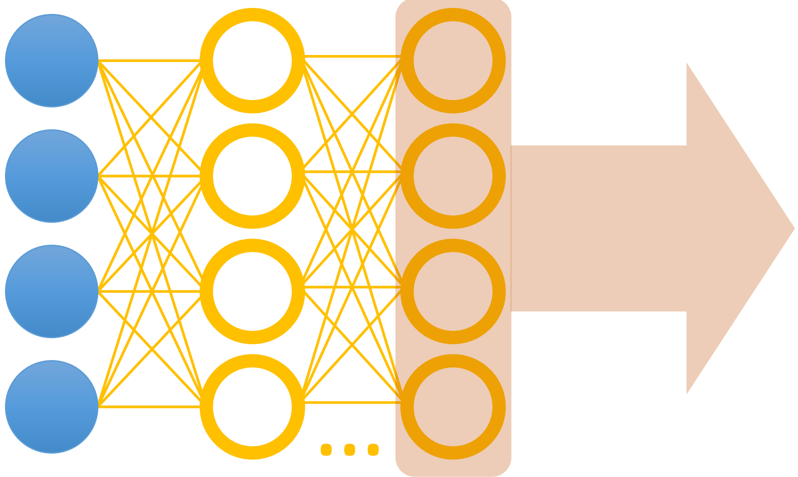

Source & Target Domain Embedding

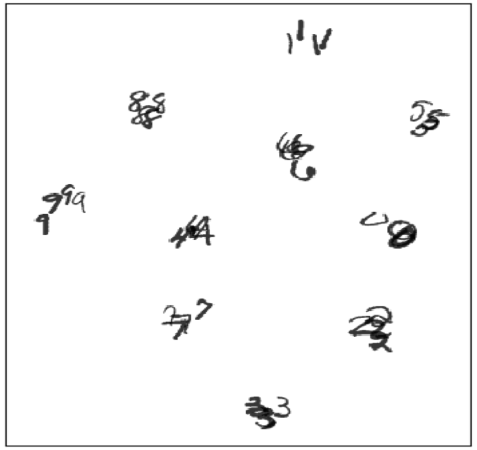

Prototypical nets

(Snell et al., 2017)

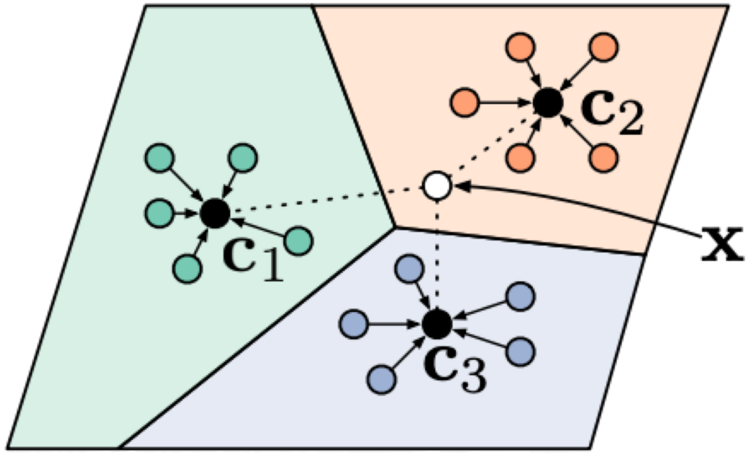

Inductive Transfer Learning

Weight Transfer

Deep Metric Learning

Few-Shot Learning

→ **Adapted Deep Embeddings**

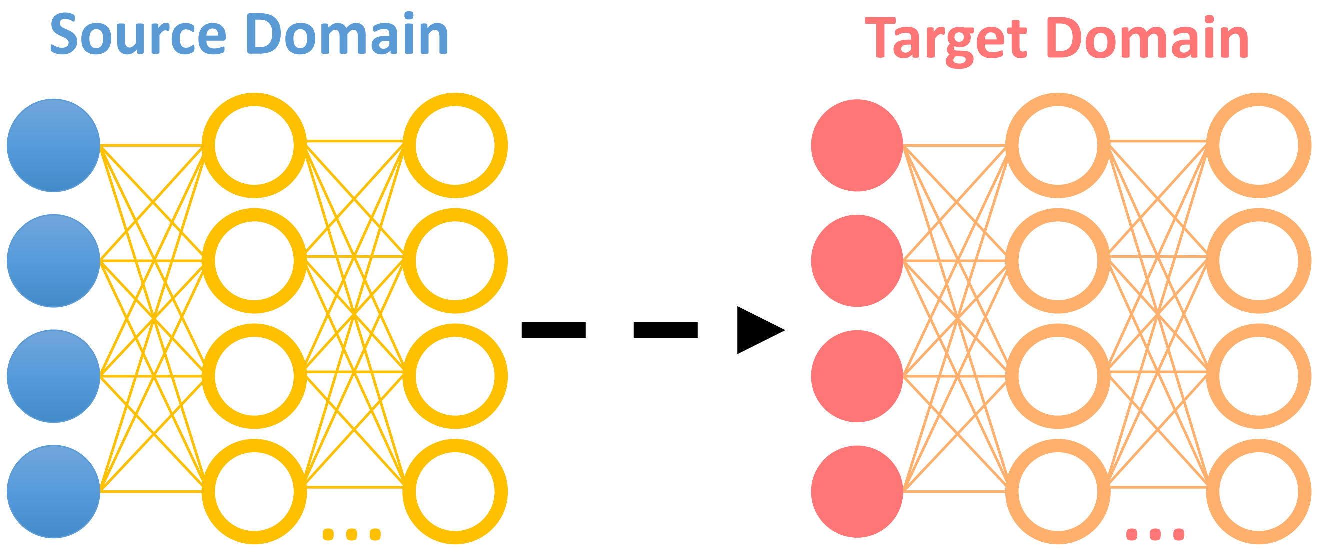

1. Train network using embedding loss
   - Histogram loss, Prototypical nets
2. Adapt weights using limited target-domain data

Inductive Transfer Learning

Why hasn't a comparison been explored?

	# labeled examples per target class (*k*)
Weight Transfer	> 100
Deep Metric Learning	agnostic
Few-Shot Learning	< 20

MNIST

Source Domain

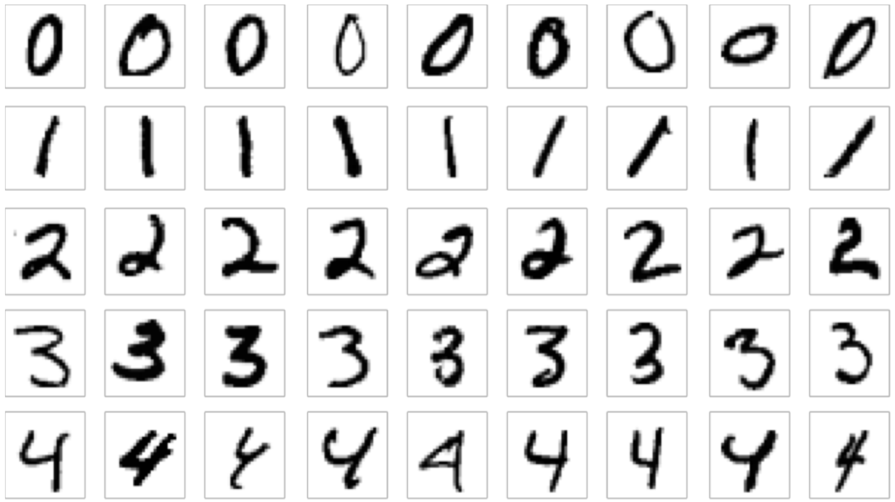

2200 labeled examples per class

Target Domain

k labeled examples per class

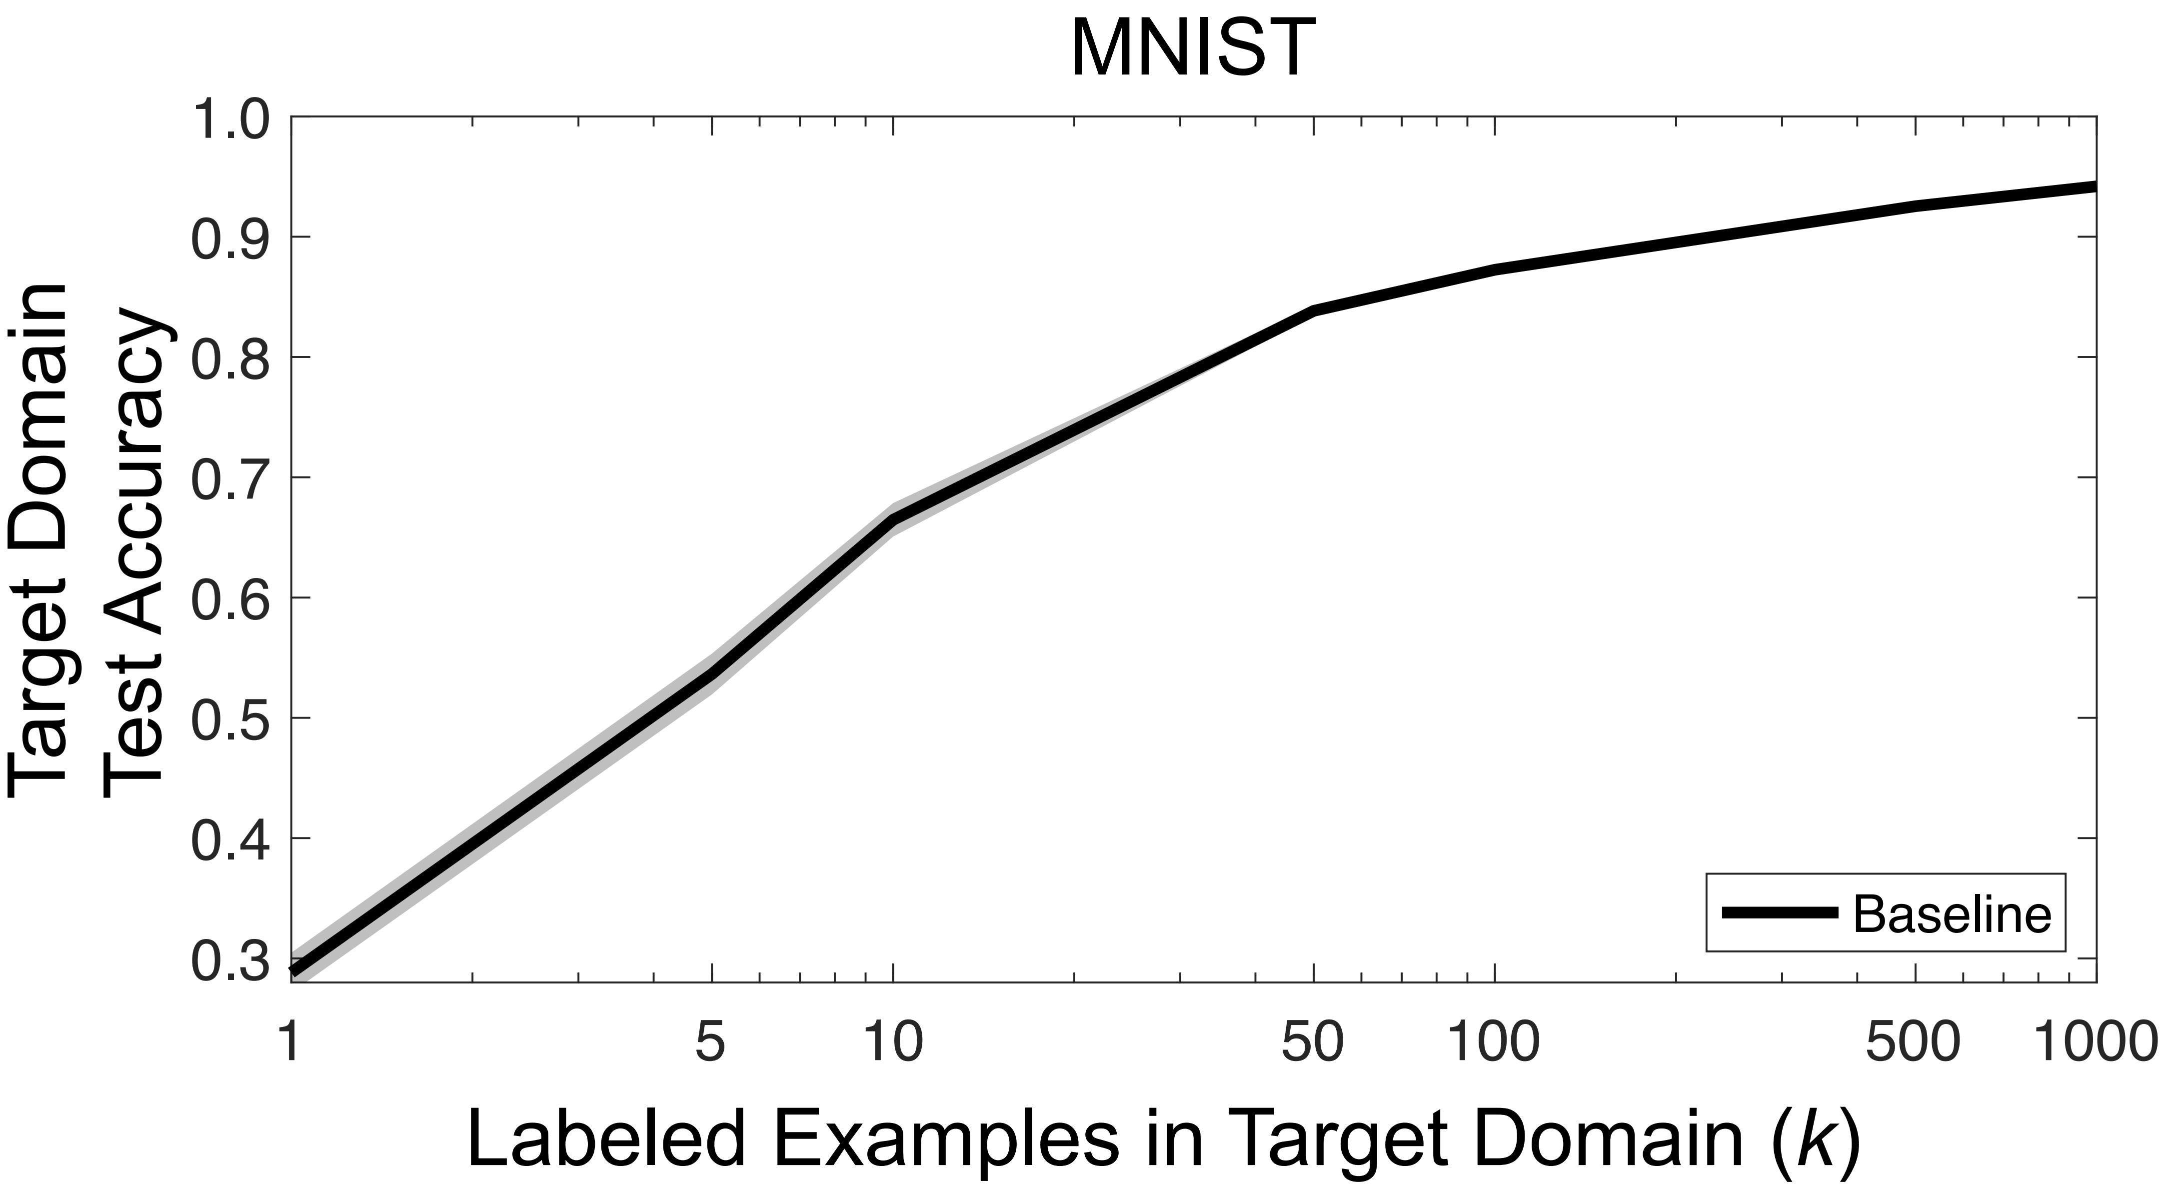
MNIST
Target Domain
Test Accuracy
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
Baseline
1
5
10
50
100
500
1000
Labeled Examples in Target Domain (k)

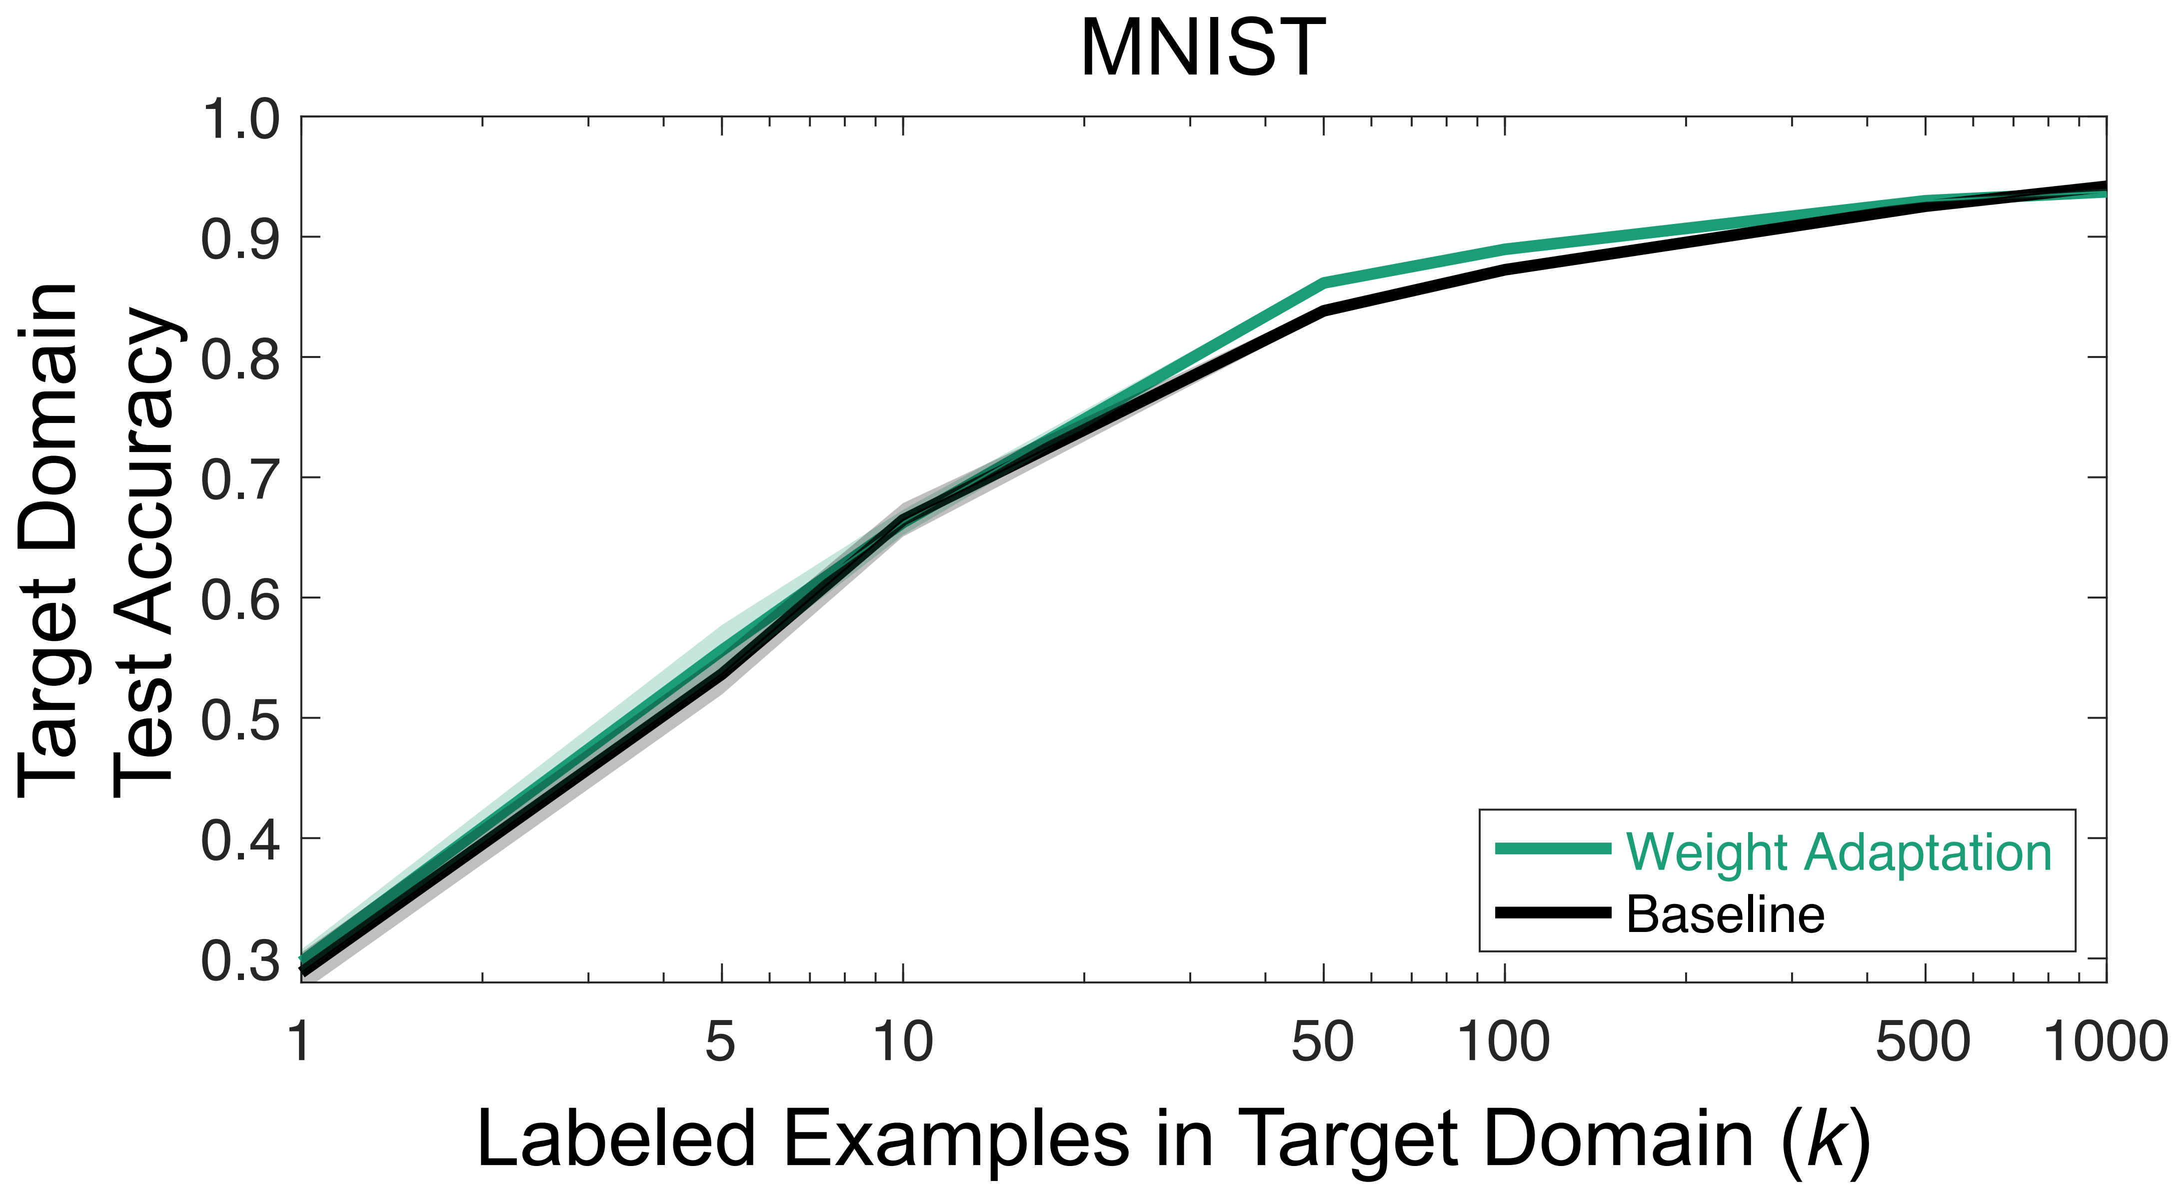

MNIST
Target Domain
Test Accuracy
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
1
5
10
50
100
500
1000
Labeled Examples in Target Domain (k)
Weight Adaptation
Baseline

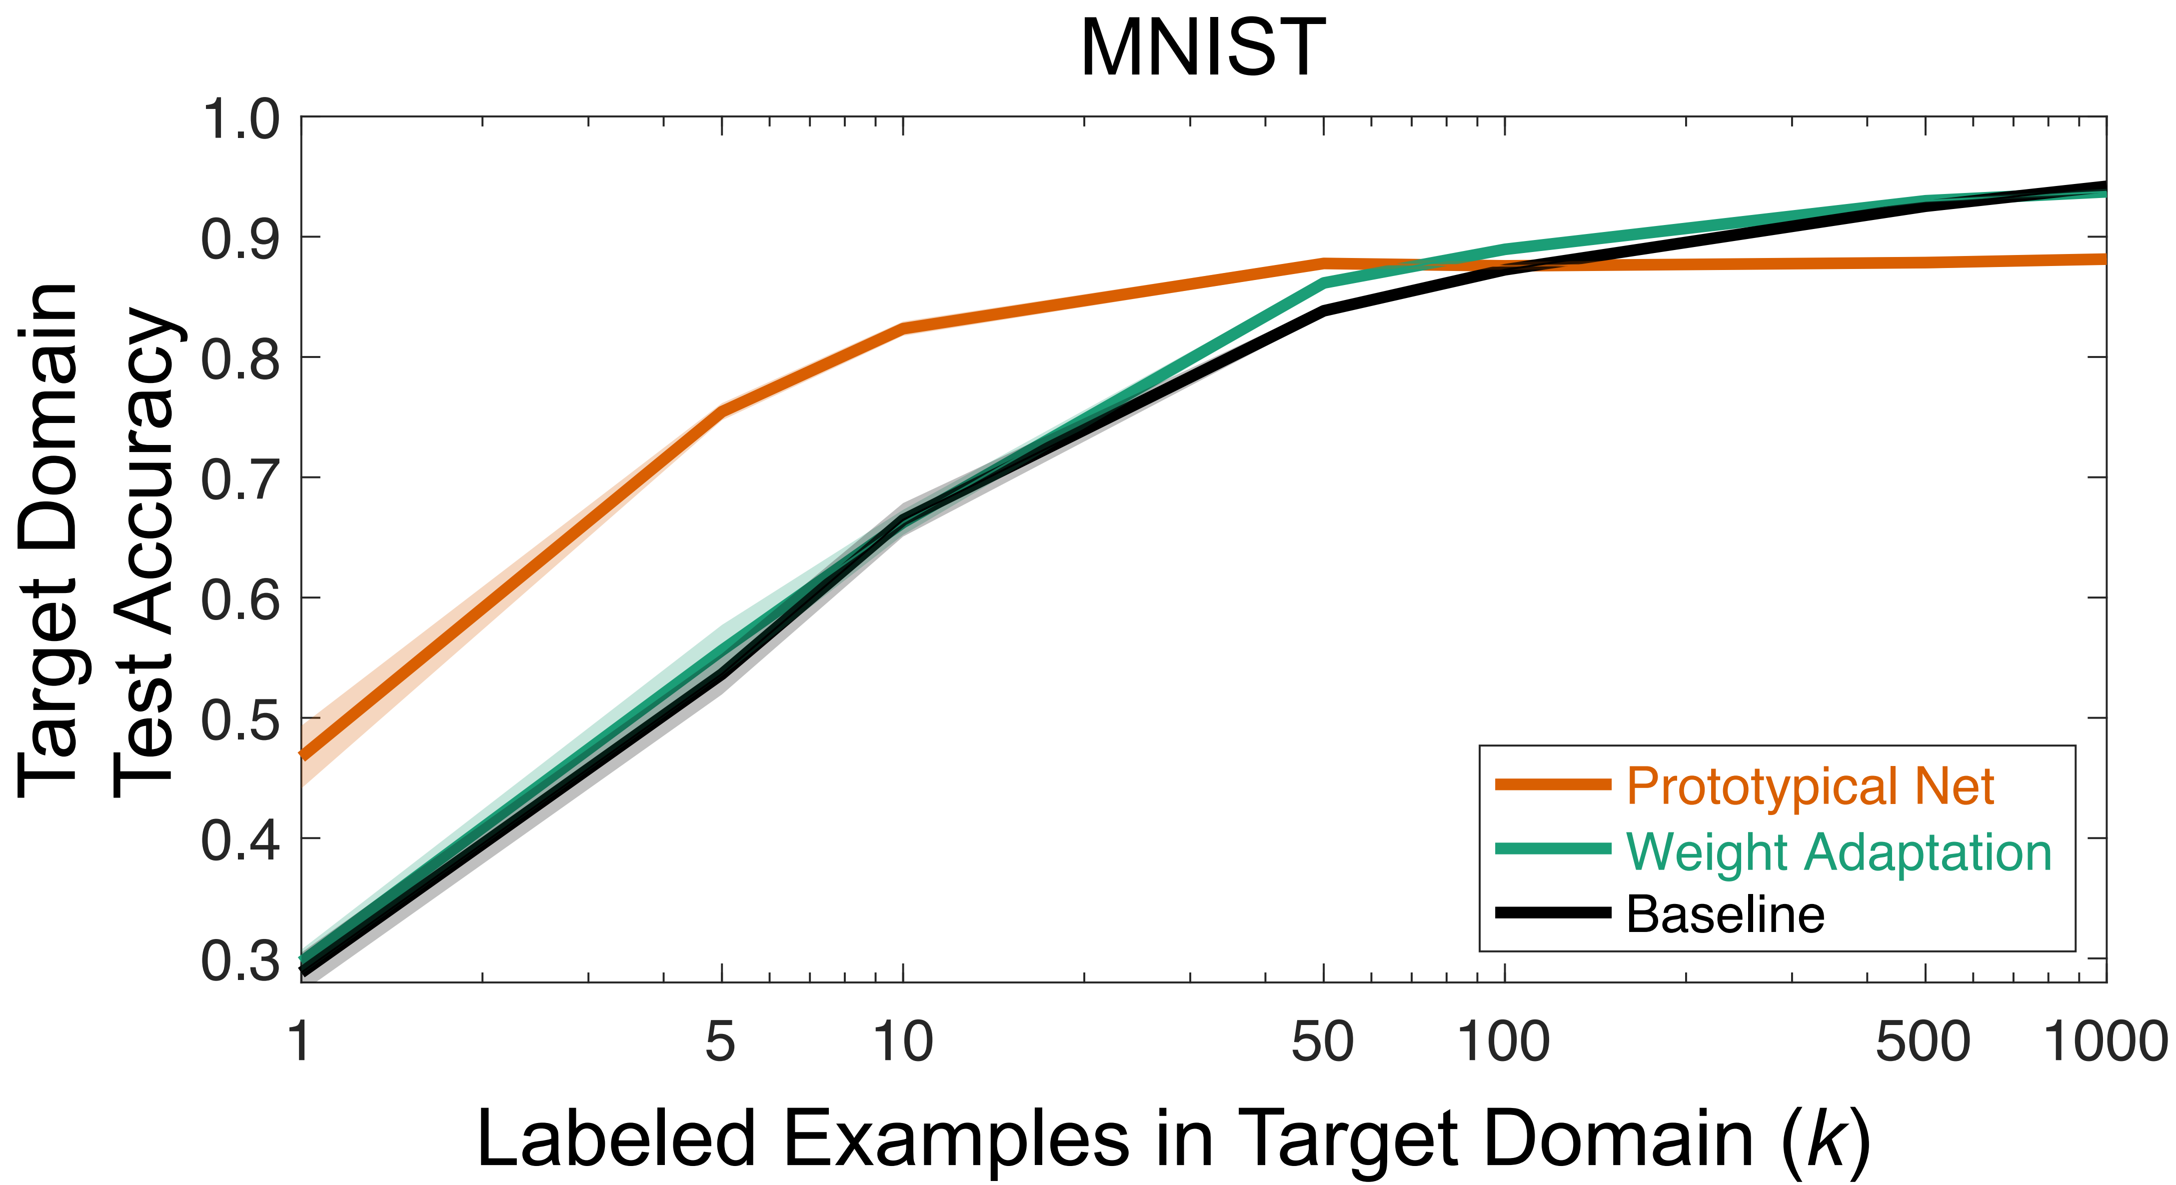
MNIST
Target Domain
Test Accuracy
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
1
5
10
50
100
500
1000
Labeled Examples in Target Domain (k)
Prototypical Net
Weight Adaptation
Baseline

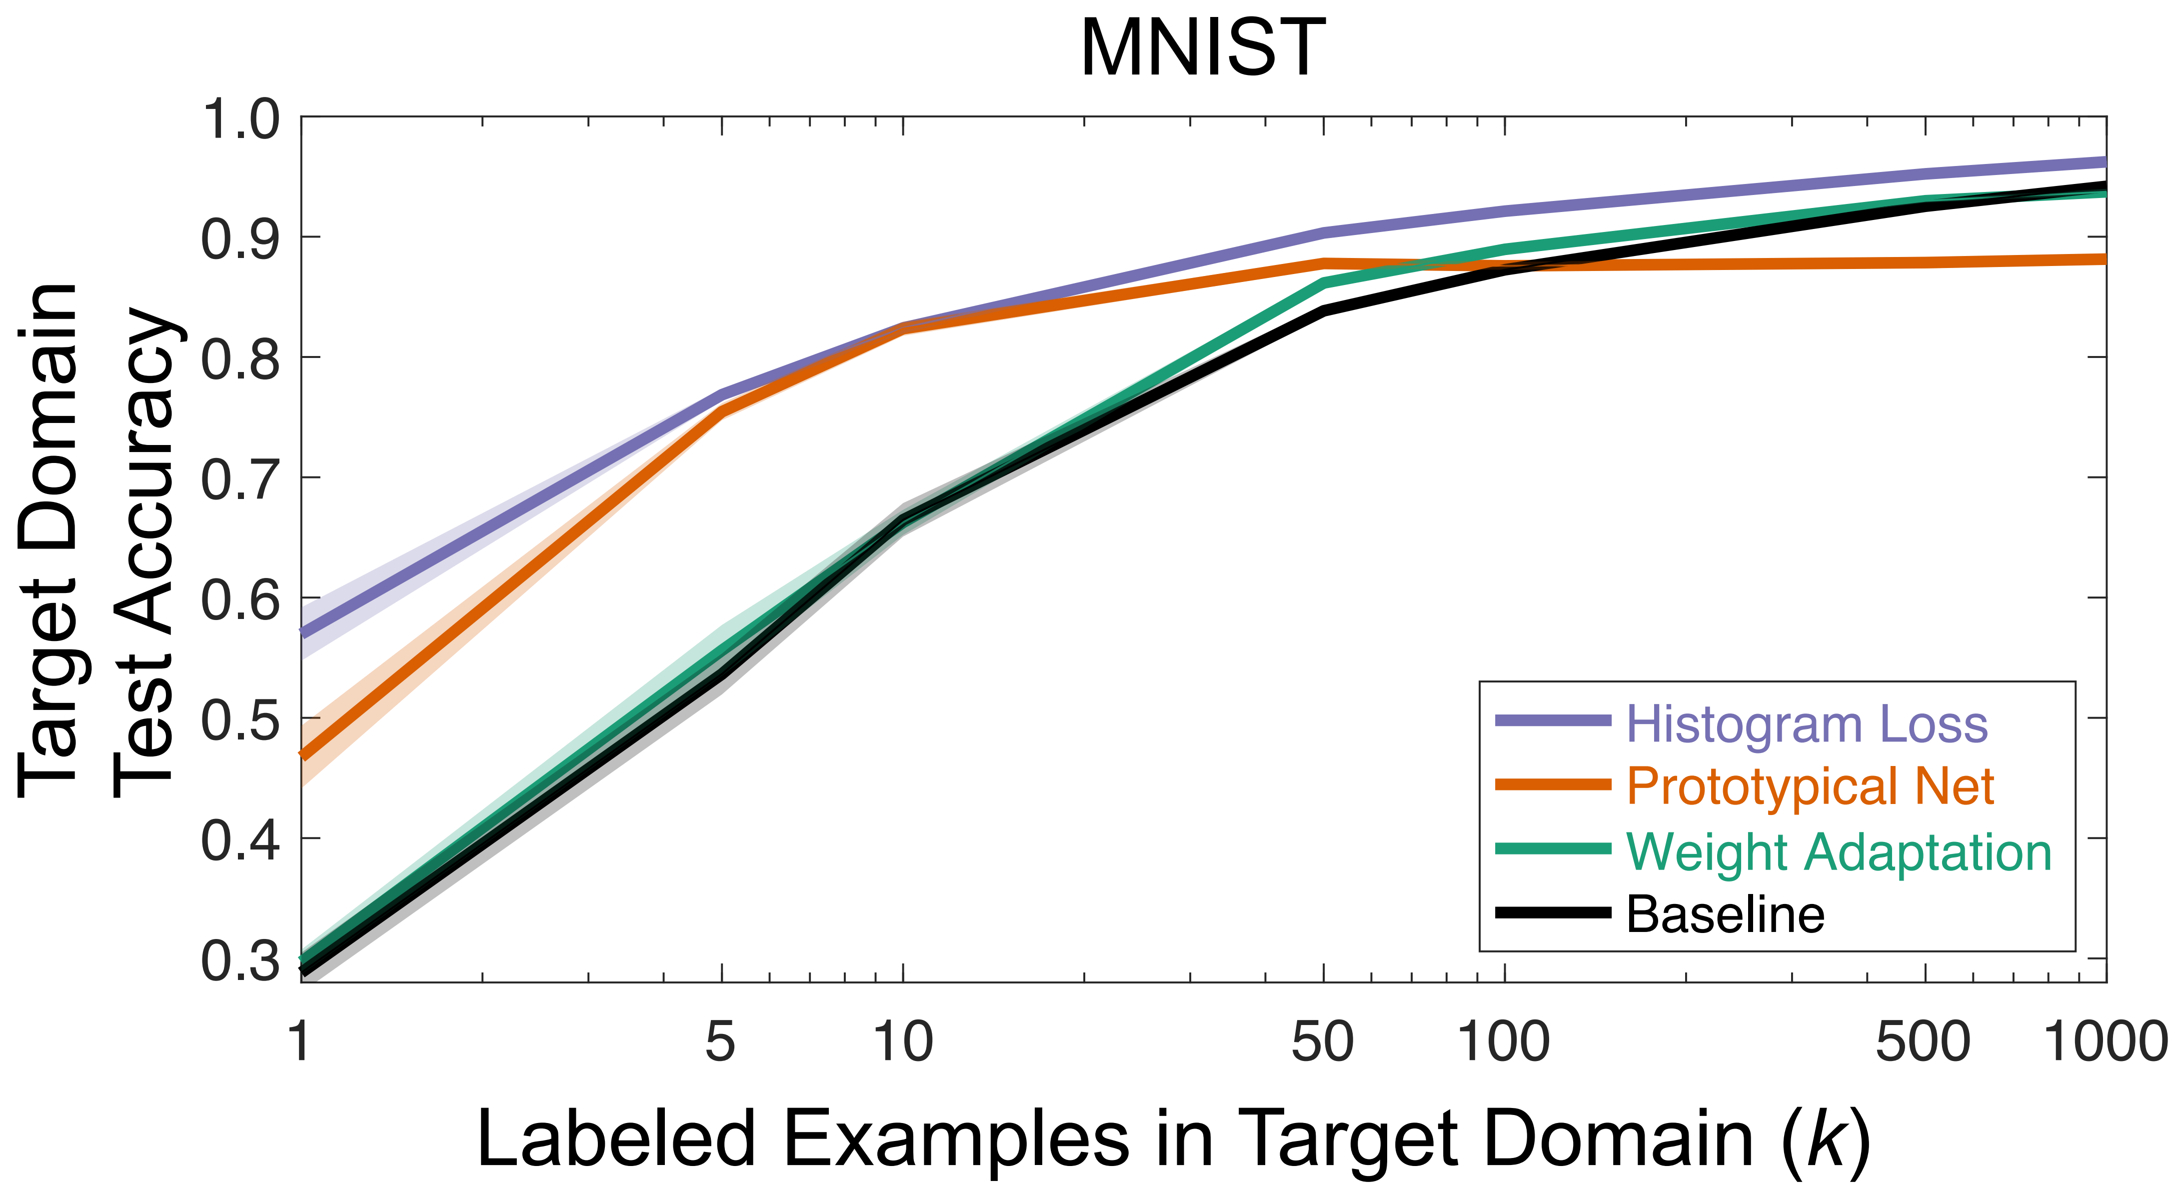
MNIST
Target Domain
Test Accuracy
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
1
5
10
50
100
500
1000
Labeled Examples in Target Domain (k)
Histogram Loss
Prototypical Net
Weight Adaptation
Baseline

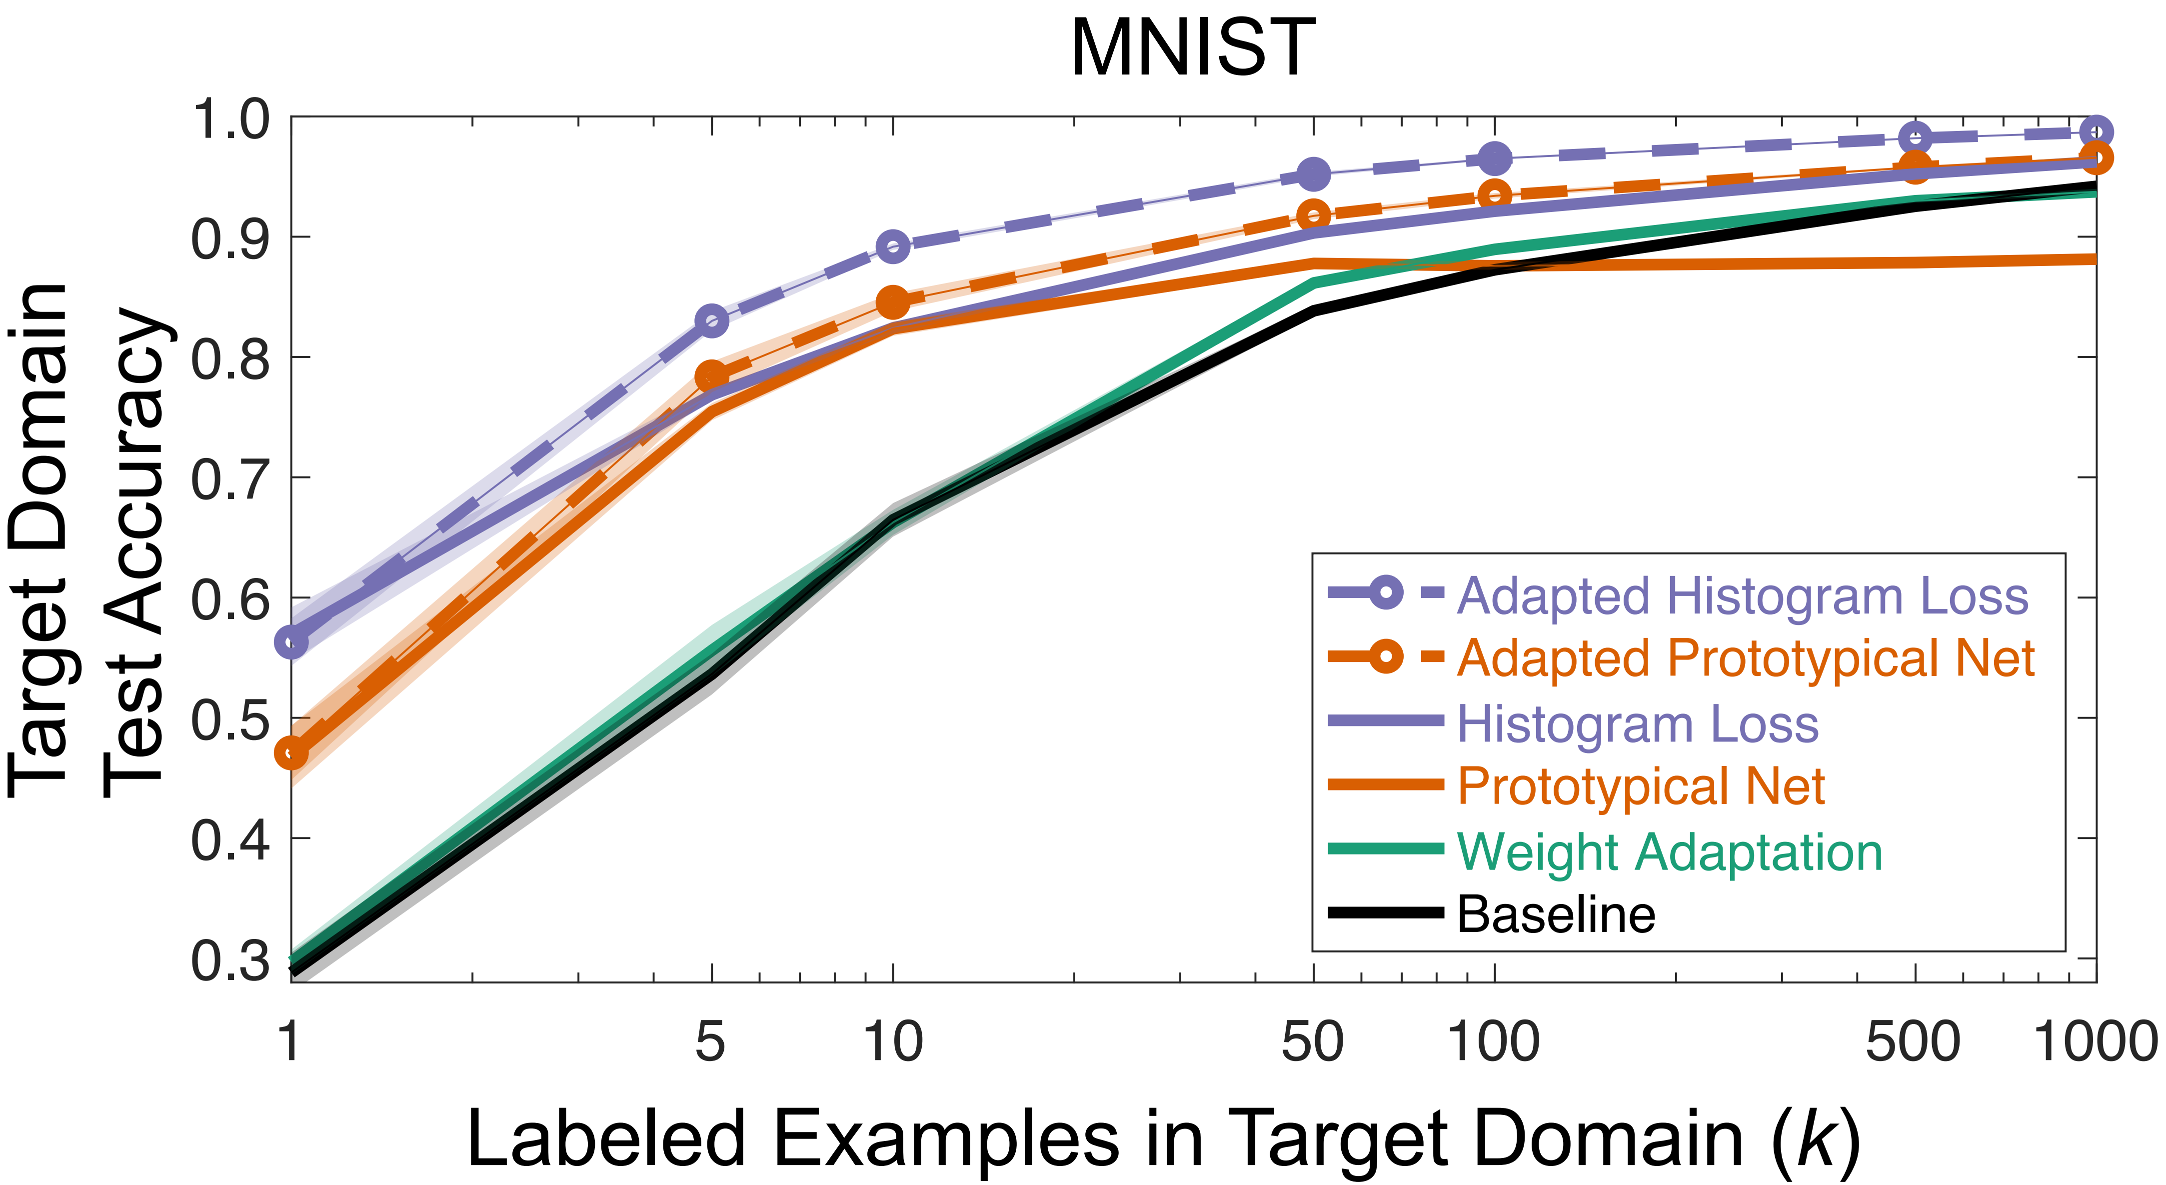
MNIST
Target Domain
Test Accuracy
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
1
5
10
50
100
500
1000
Labeled Examples in Target Domain (k)
Adapted Histogram Loss
Adapted Prototypical Net
Histogram Loss
Prototypical Net
Weight Adaptation
Baseline

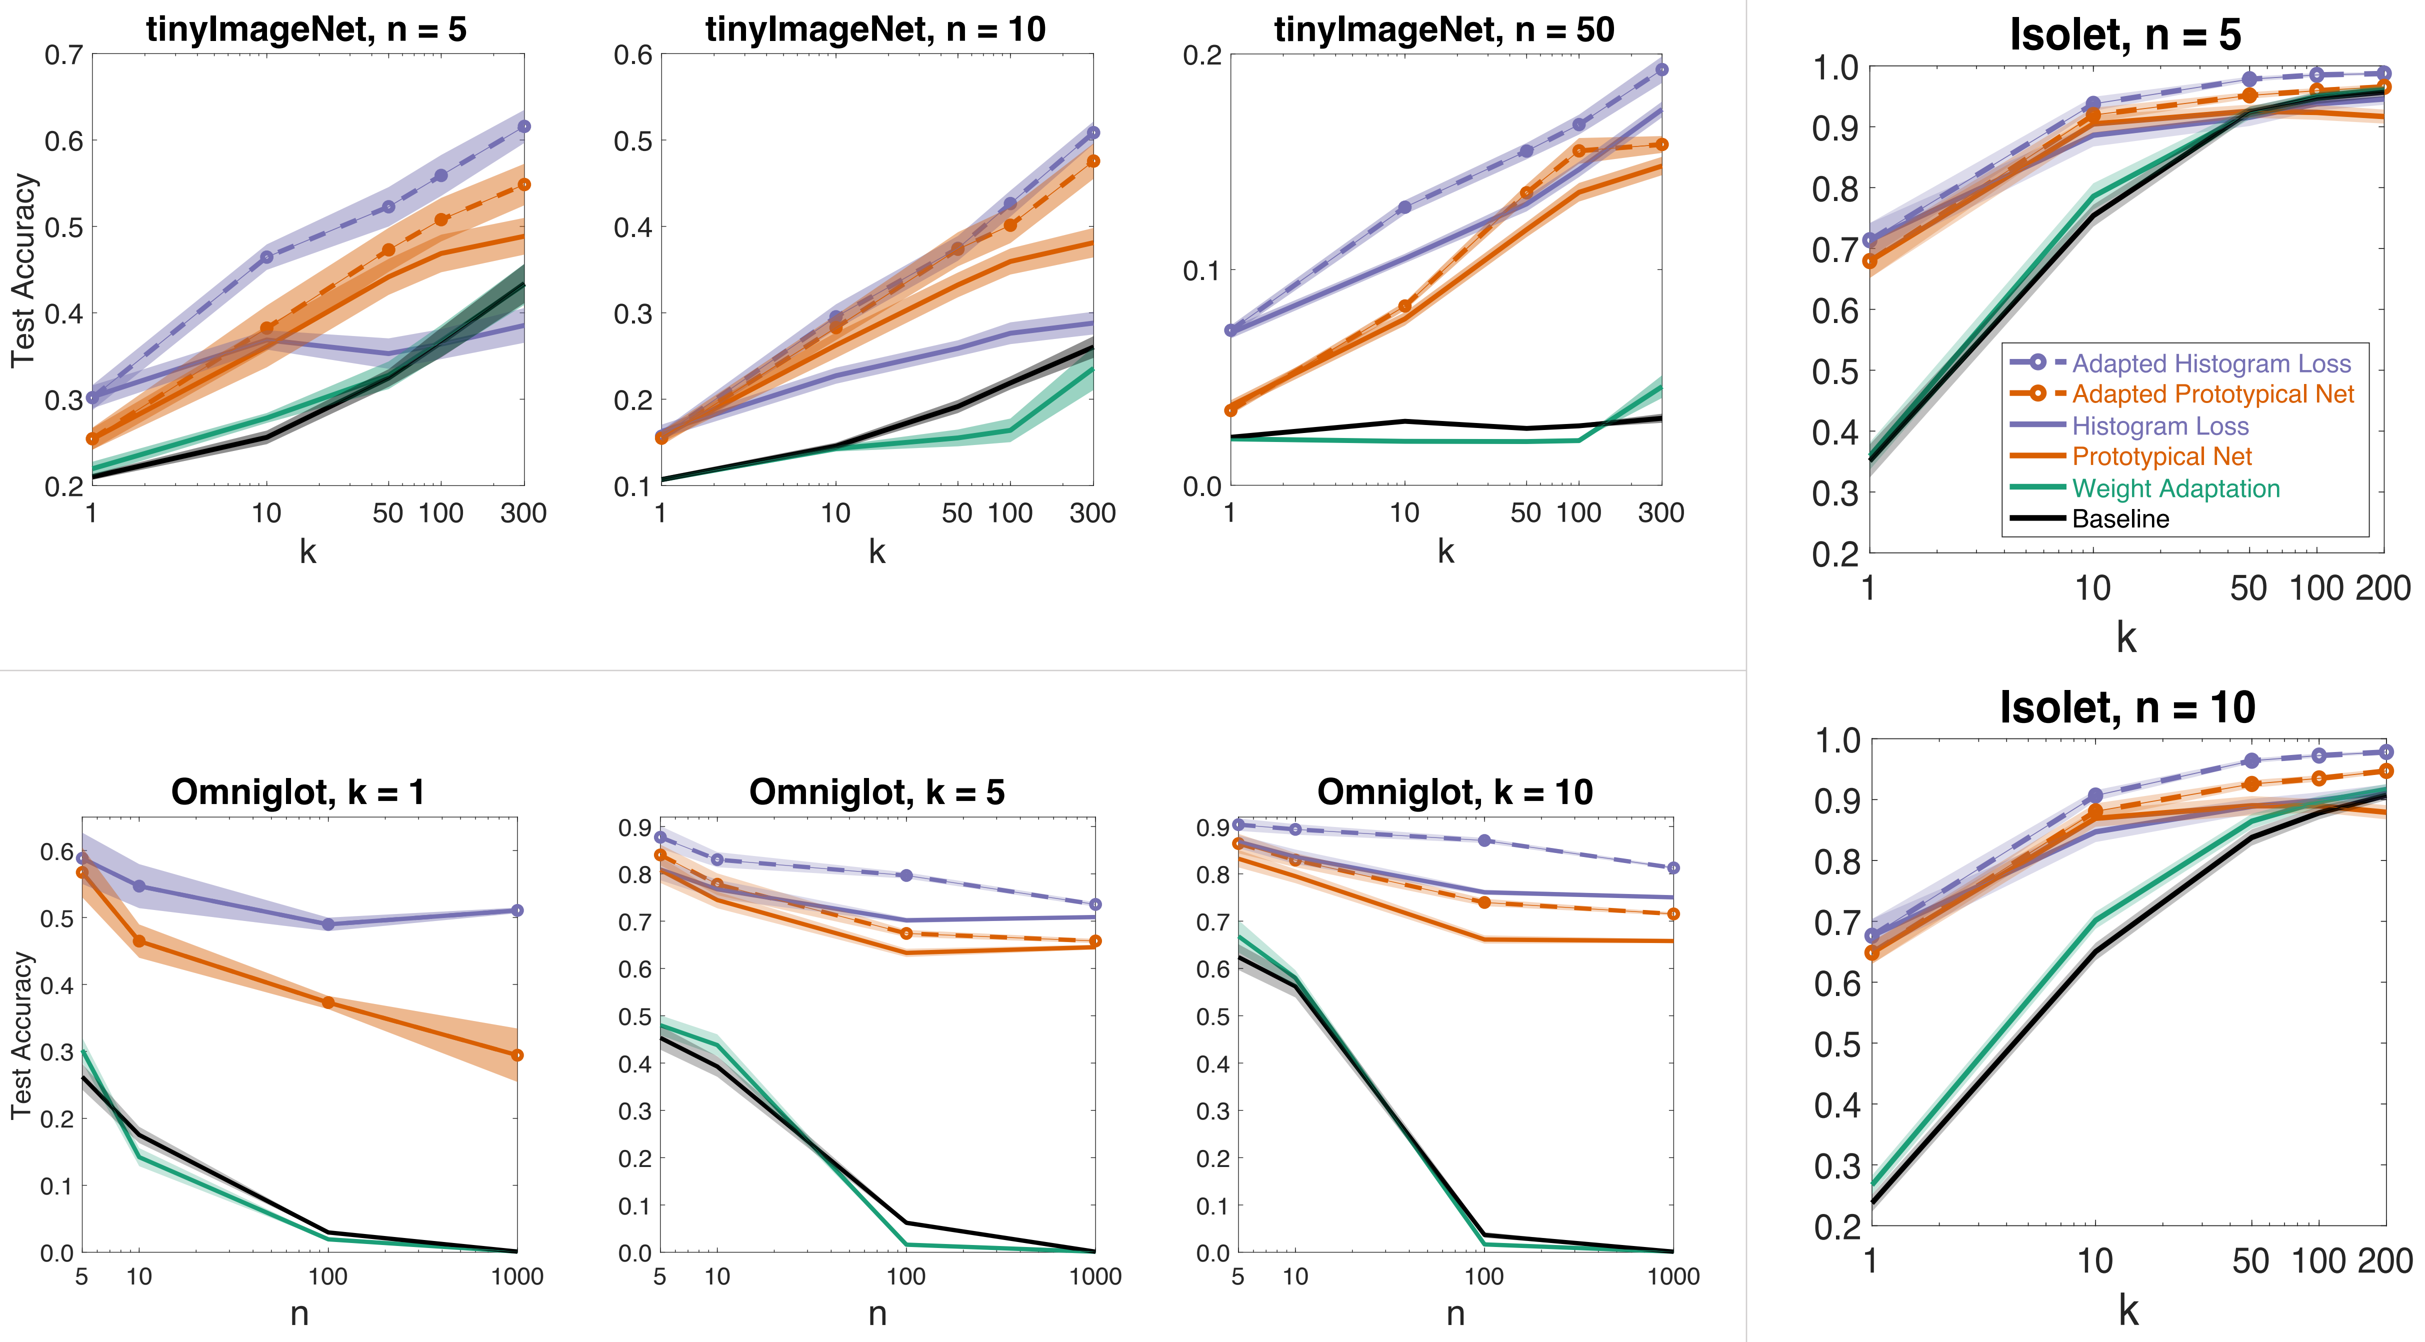

tinyImageNet, n = 5
tinyImageNet, n = 10
tinyImageNet, n = 50
Isolet, n = 5
Omniglot, k = 1
Omniglot, k = 5
Omniglot, k = 10
Isolet, n = 10
Test Accuracy
k
n
Adapted Histogram Loss
Adapted Prototypical Net
Histogram Loss
Prototypical Net
Weight Adaptation
Baseline

Conclusion

- **Weight transfer is the least effective method for inductive transfer learning**

- **Histogram loss is robust regardless of the amount of labeled data in the target domain**

- **Adapted embeddings outperform *every* static embedding method previously proposed**

Poster #167

Room 210 & 230 AB

Today, 5:00 - 7:00 PM