Representation Learning

Lecture slides for Chapter 15 of *Deep Learning*

www.deeplearningbook.org

Ian Goodfellow

2017-10-03

Unsupervised Pretraining Usually Hurts but Sometimes Helps

(Ma et al, 2015)

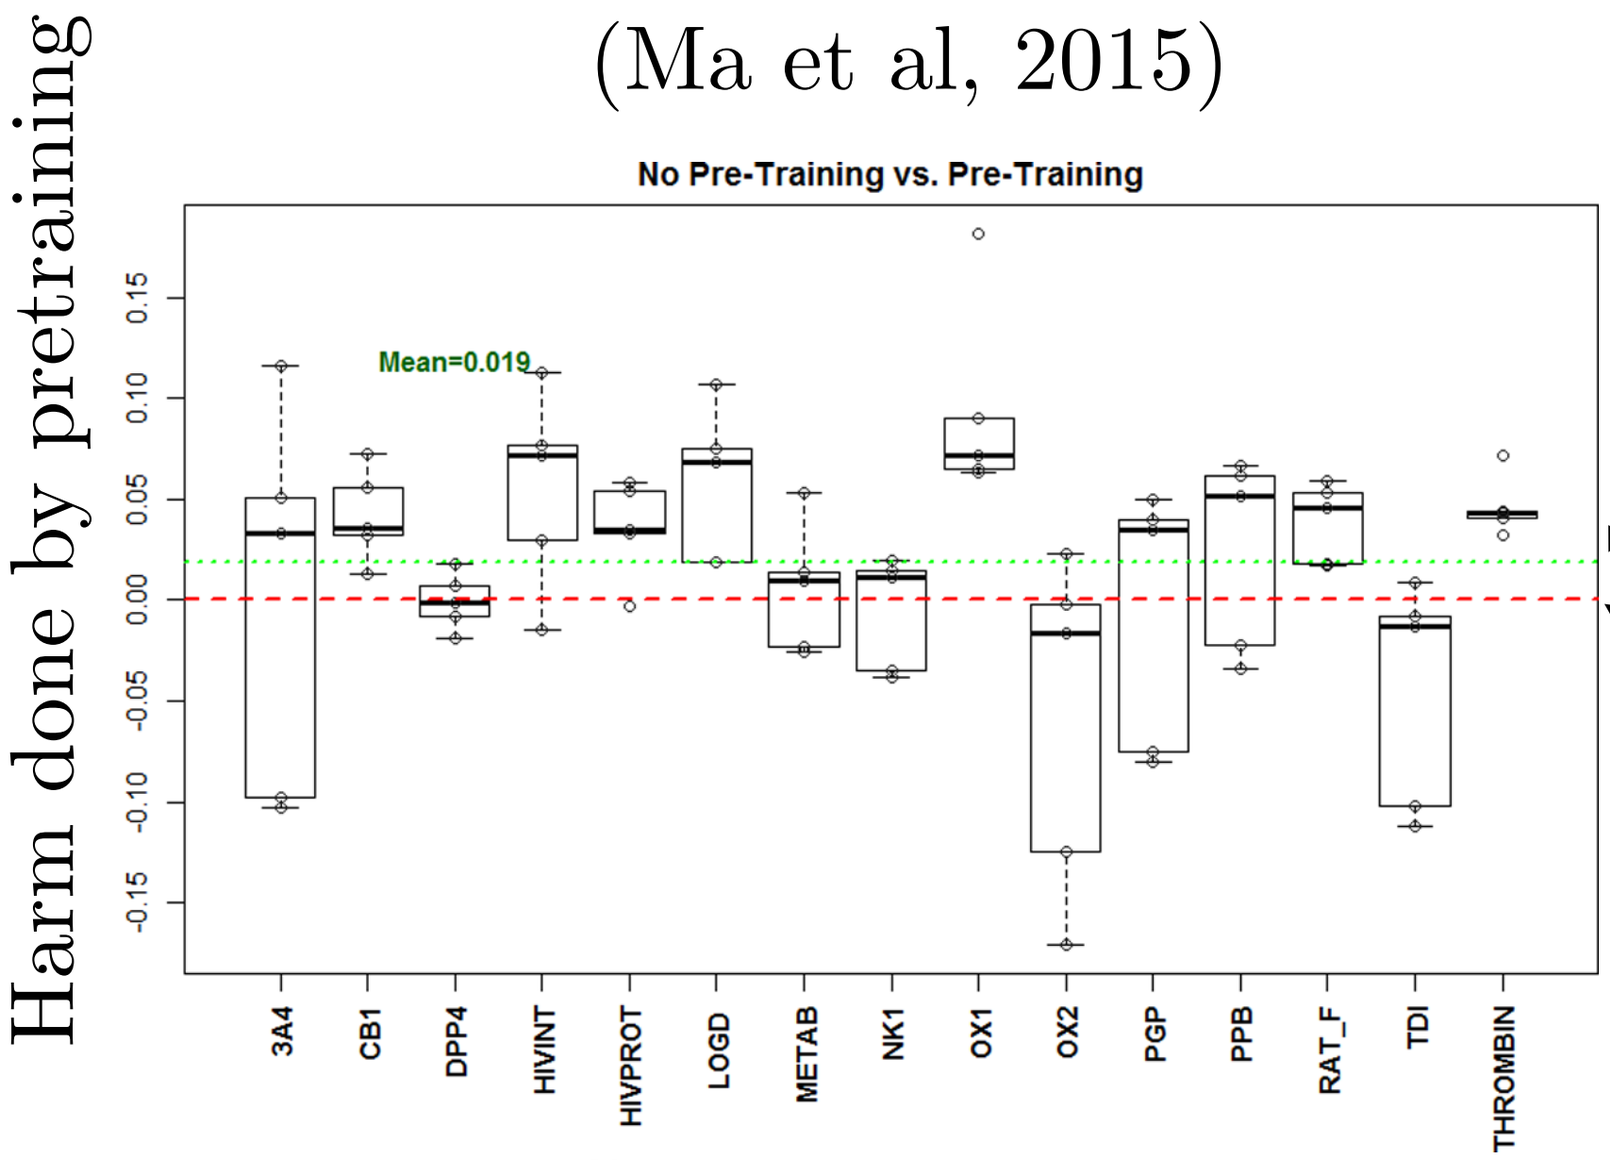

Pretraining Changes Learning Trajectory

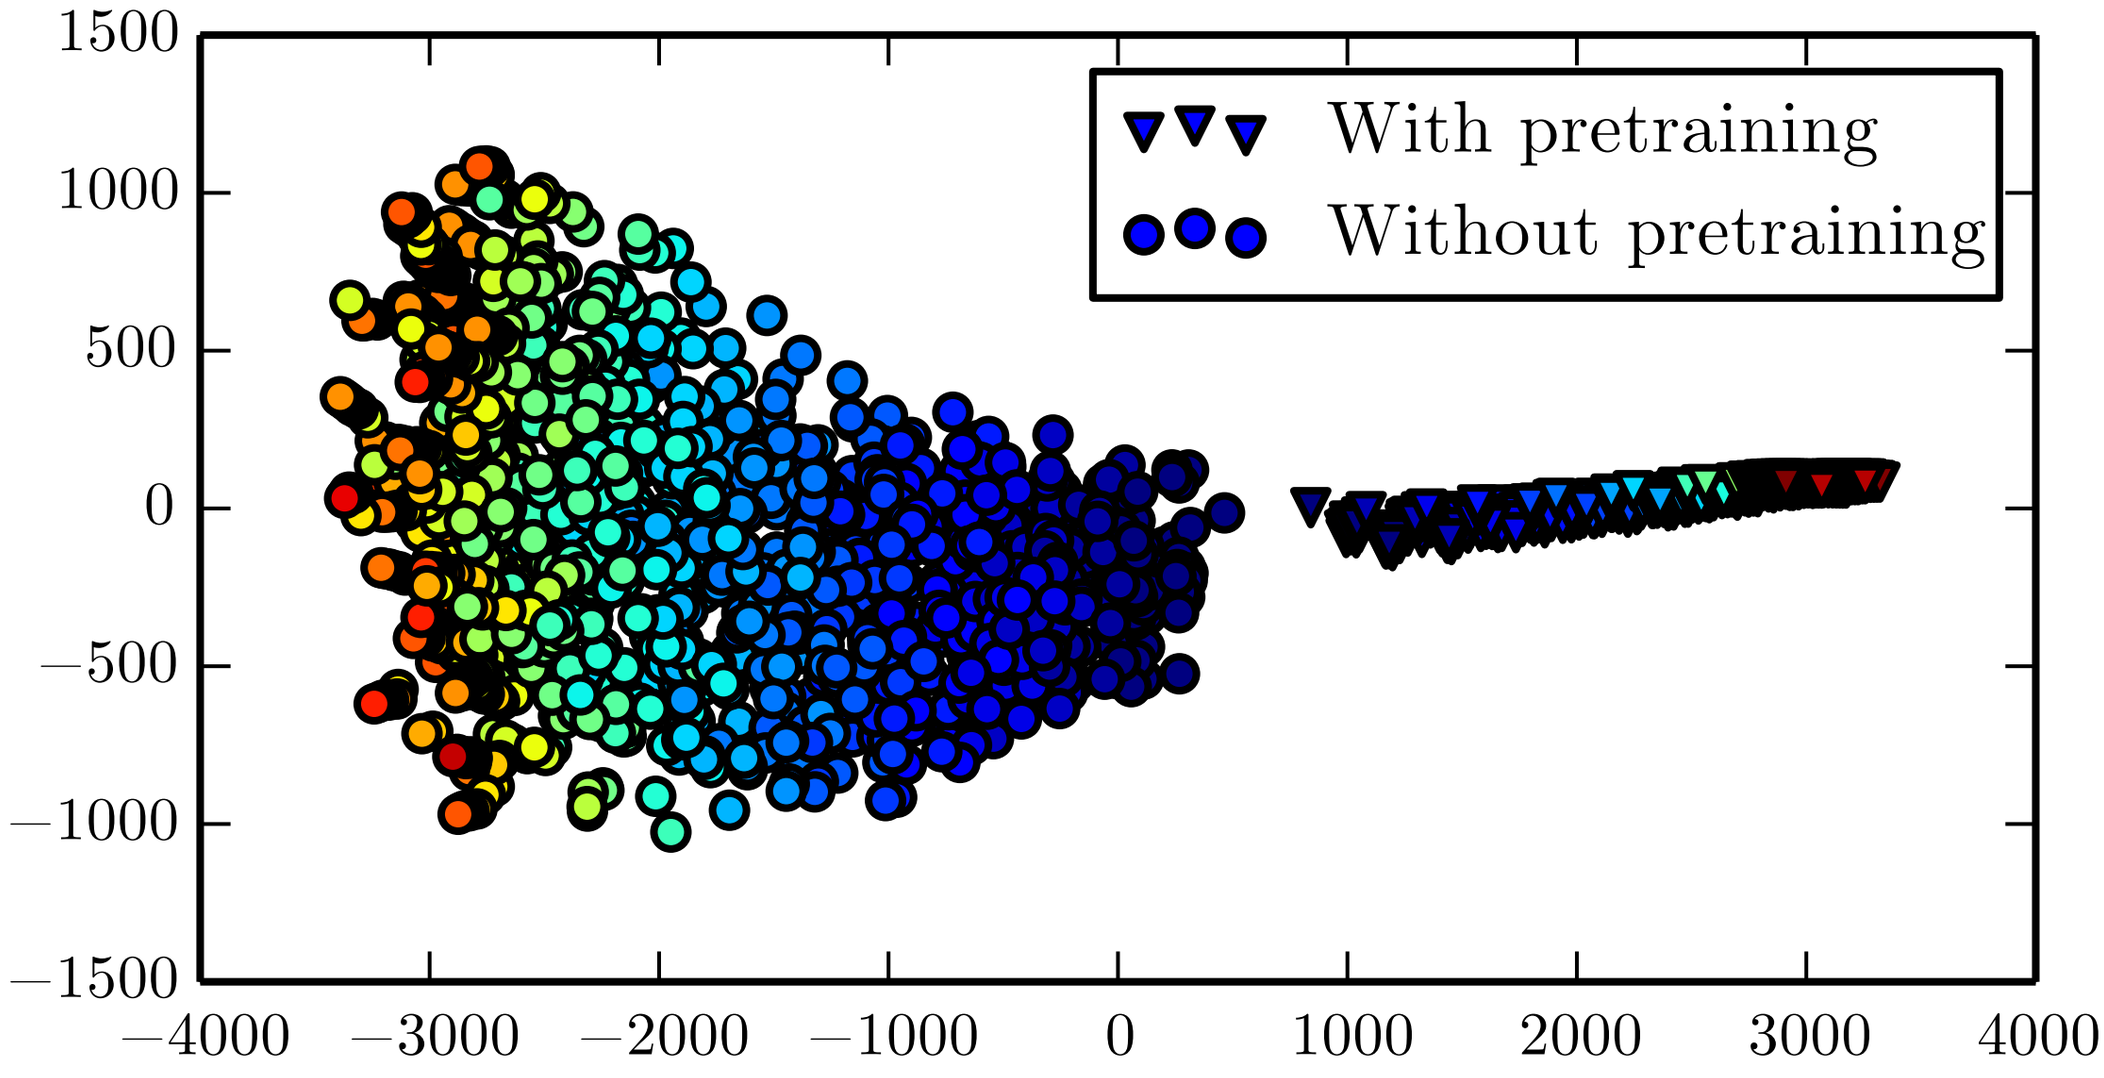

Figure 15.1

Representation Sharing for Multi-Task or Transfer Learning

One representation used for many input formats or many tasks

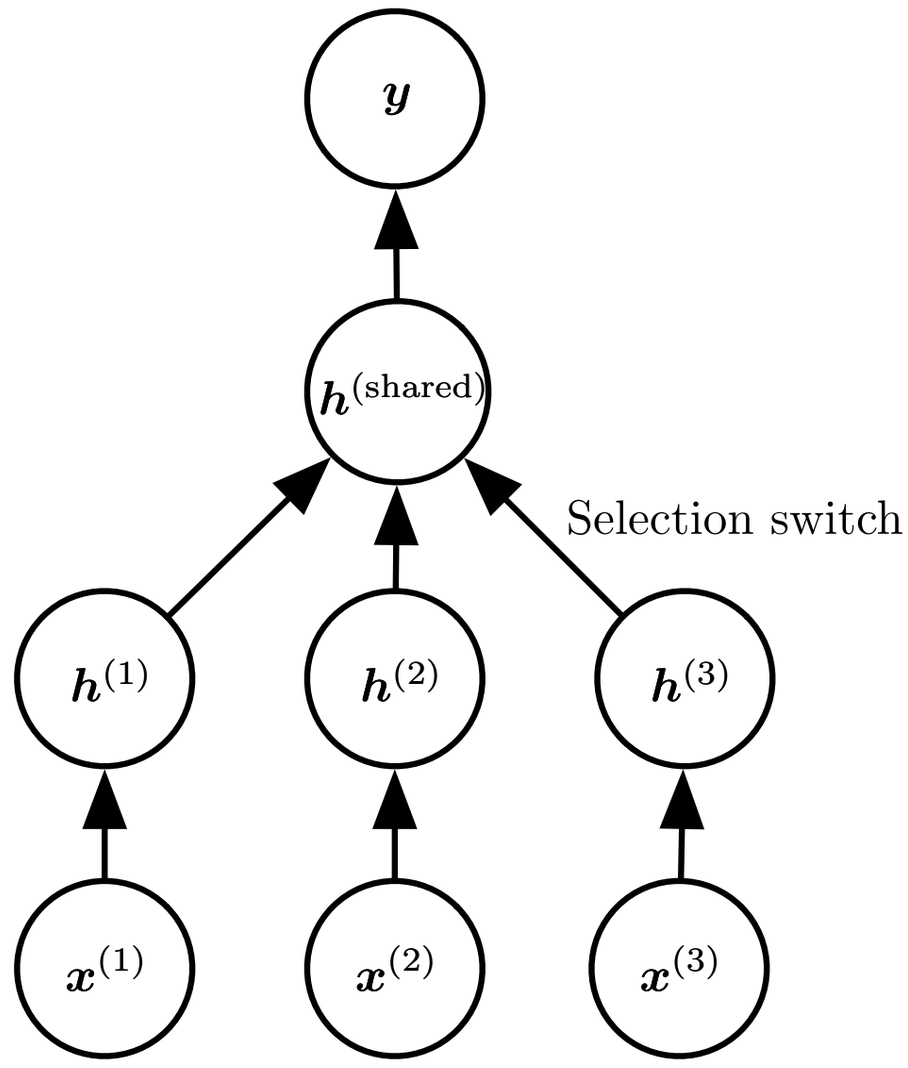

Figure 15.2

Zero Shot Learning

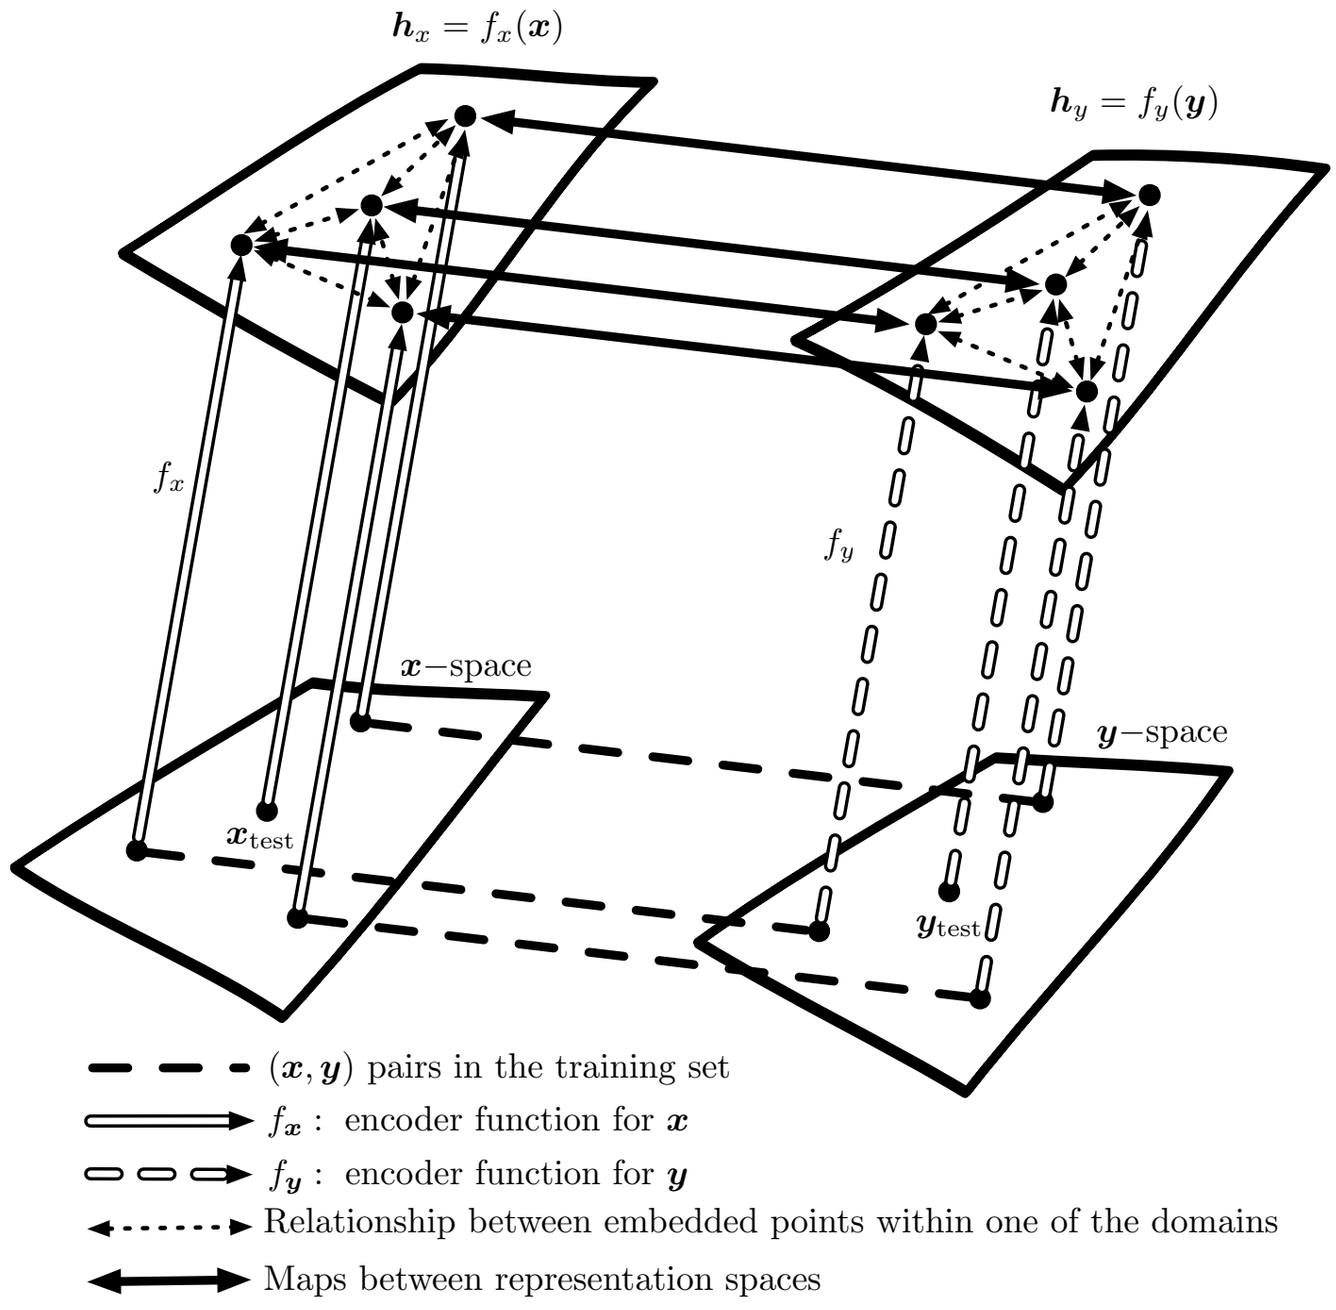

Figure 15.3

Mixture Modeling Discovers Separate Classes

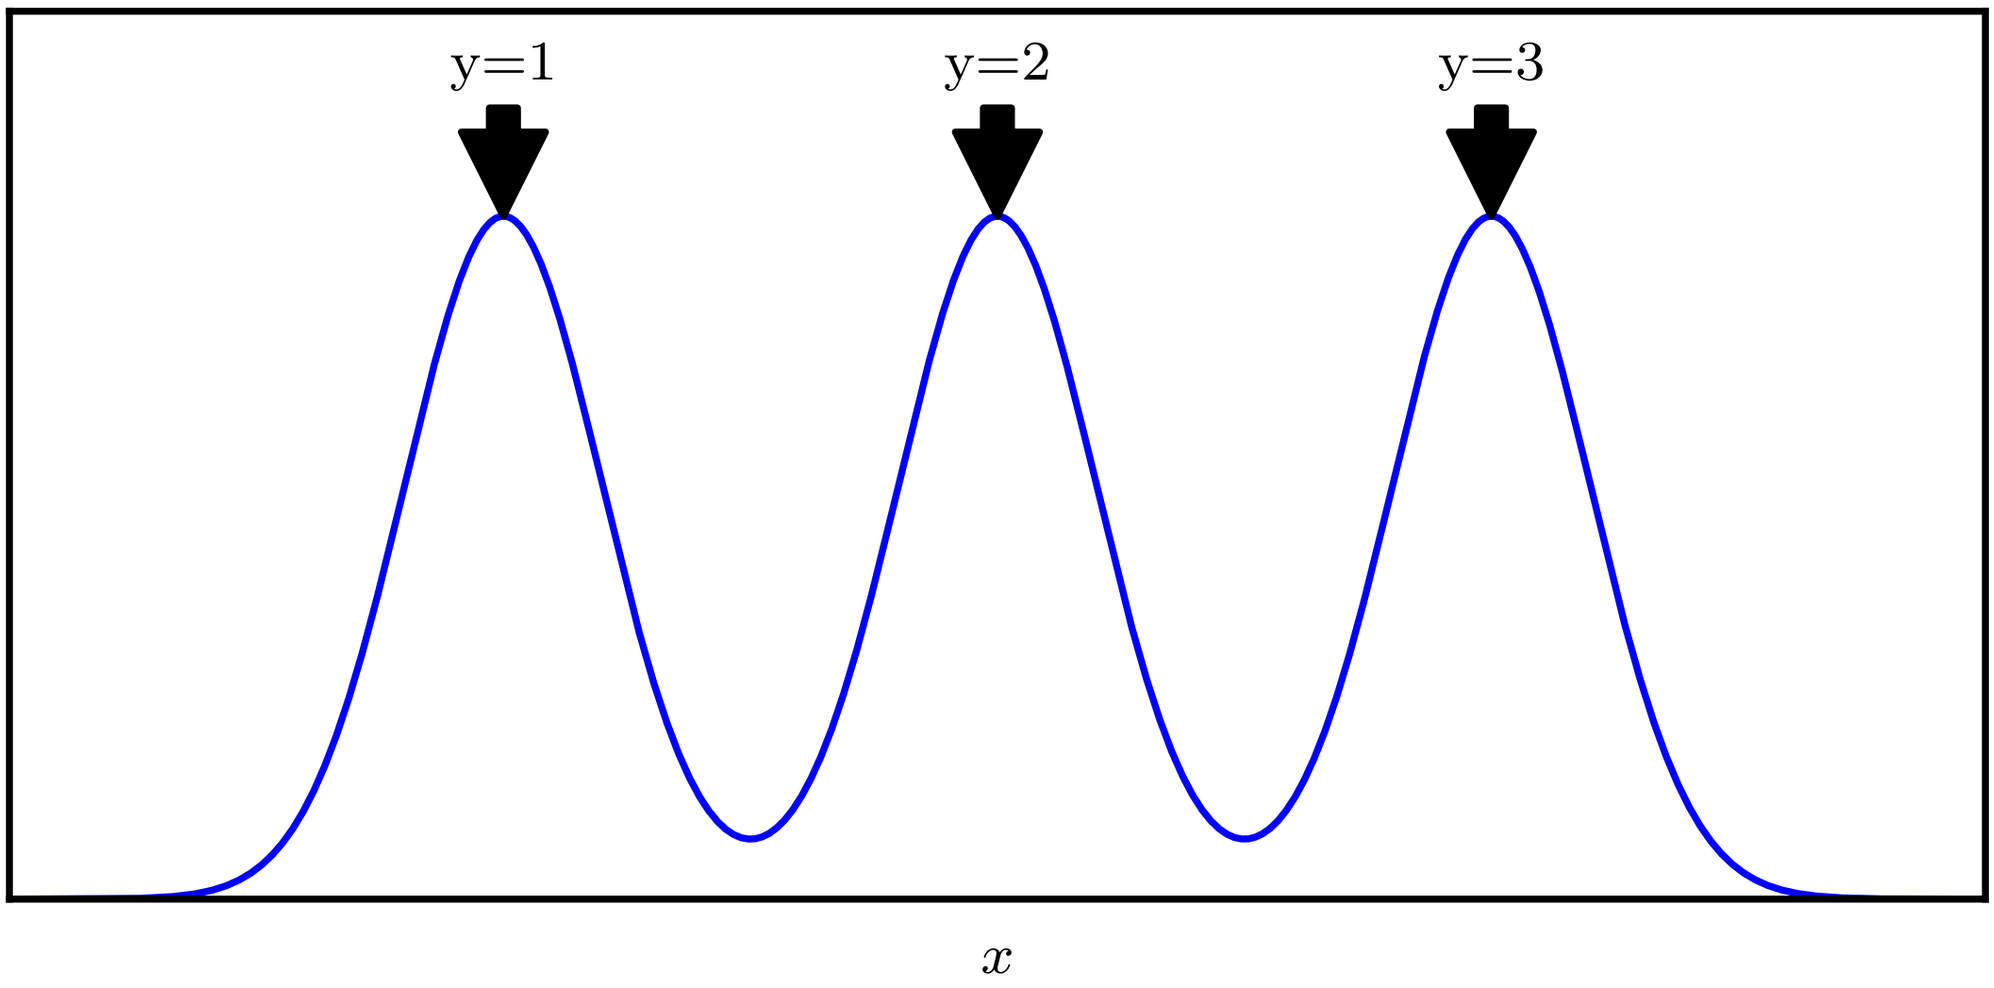

Figure 15.4

Mean Squared Error Can Ignore Small but Task-Relevant Features

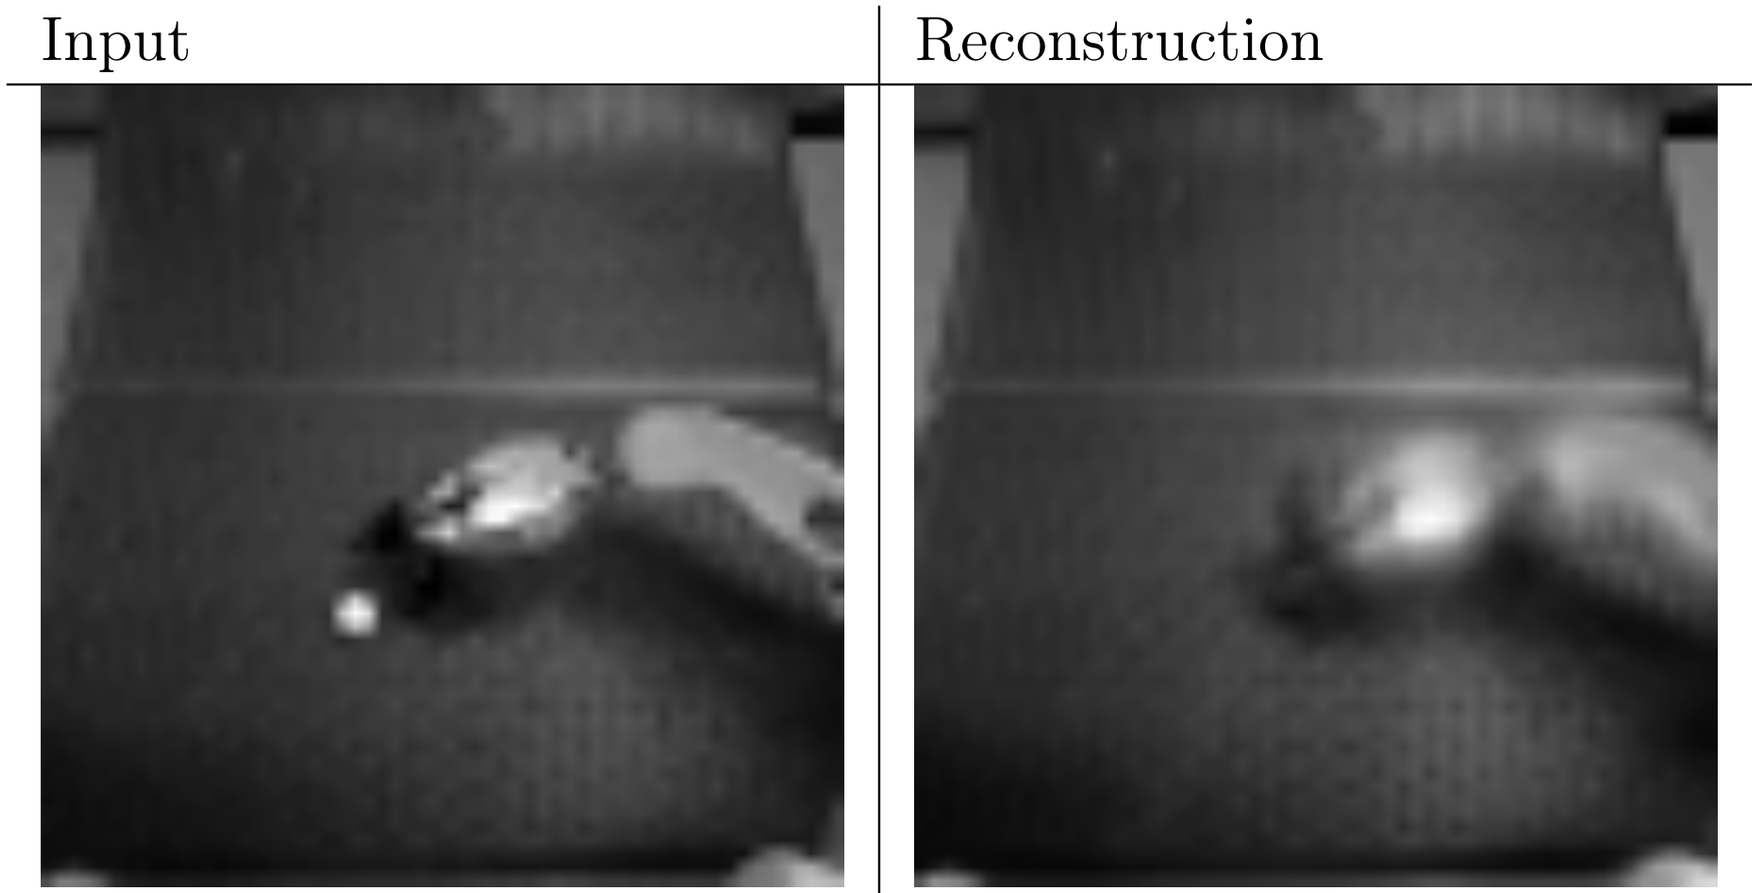

Figure 15.5

The ping pong ball vanishes because it is not large enough to significantly affect the mean squared error

Adversarial Losses Preserve Any Features with Highly Structured Patterns

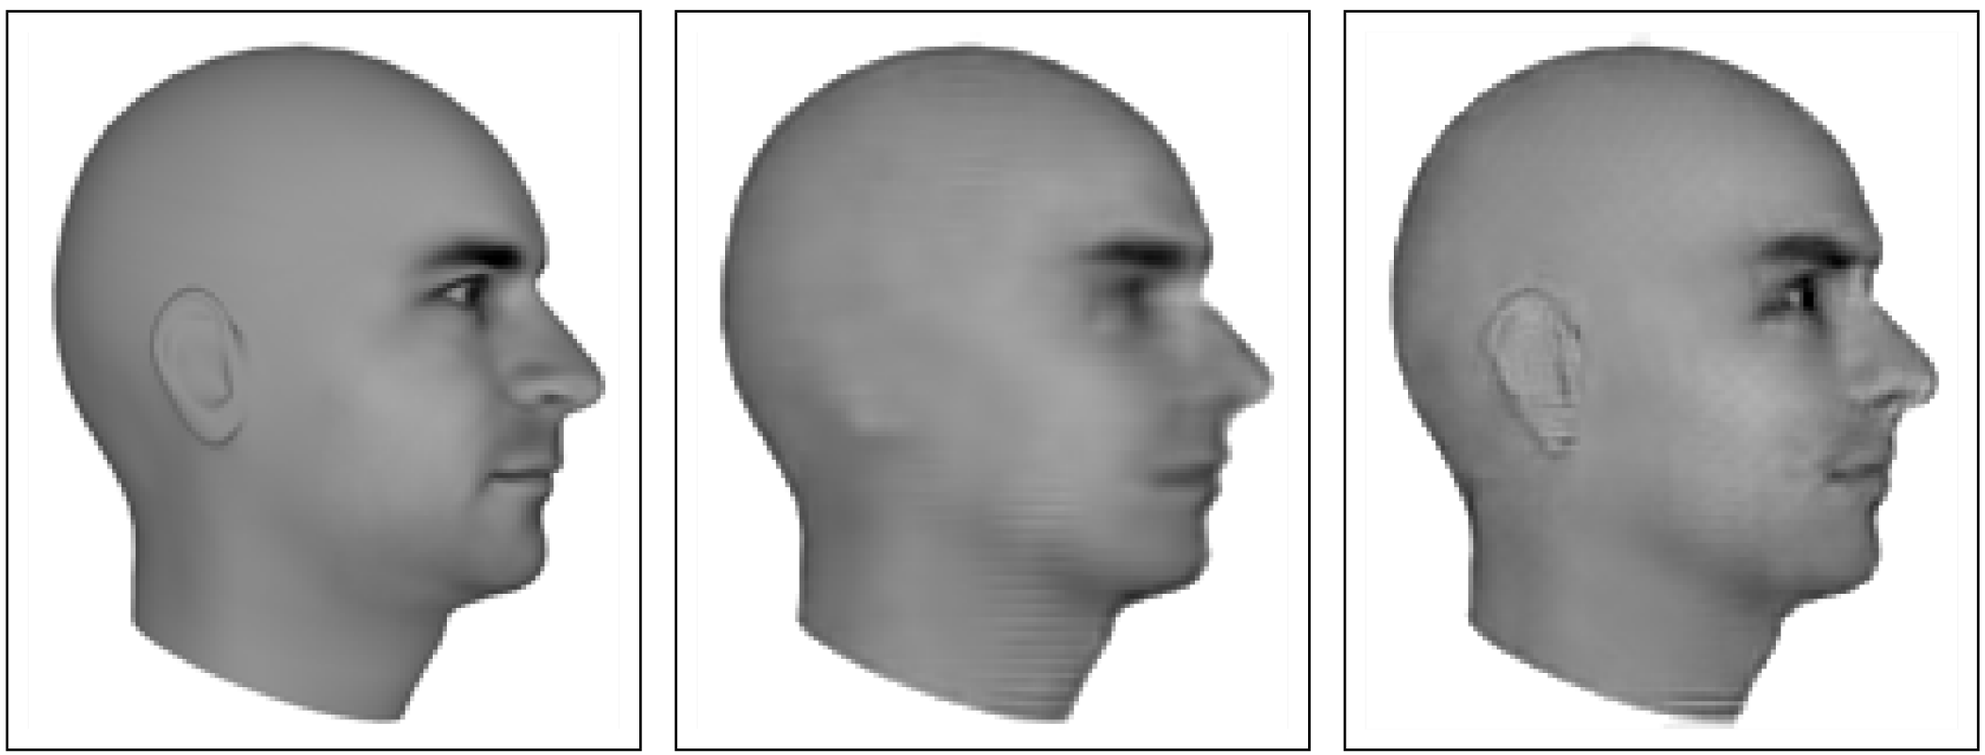

Figure 15.6

Mean squared error loses the ear because it causes a small change in few pixels. Adversarial loss preserves the ear because it is easy to notice its absence.

Binary Distributed Representations Divide Space Into Many Uniquely Identifiable Regions

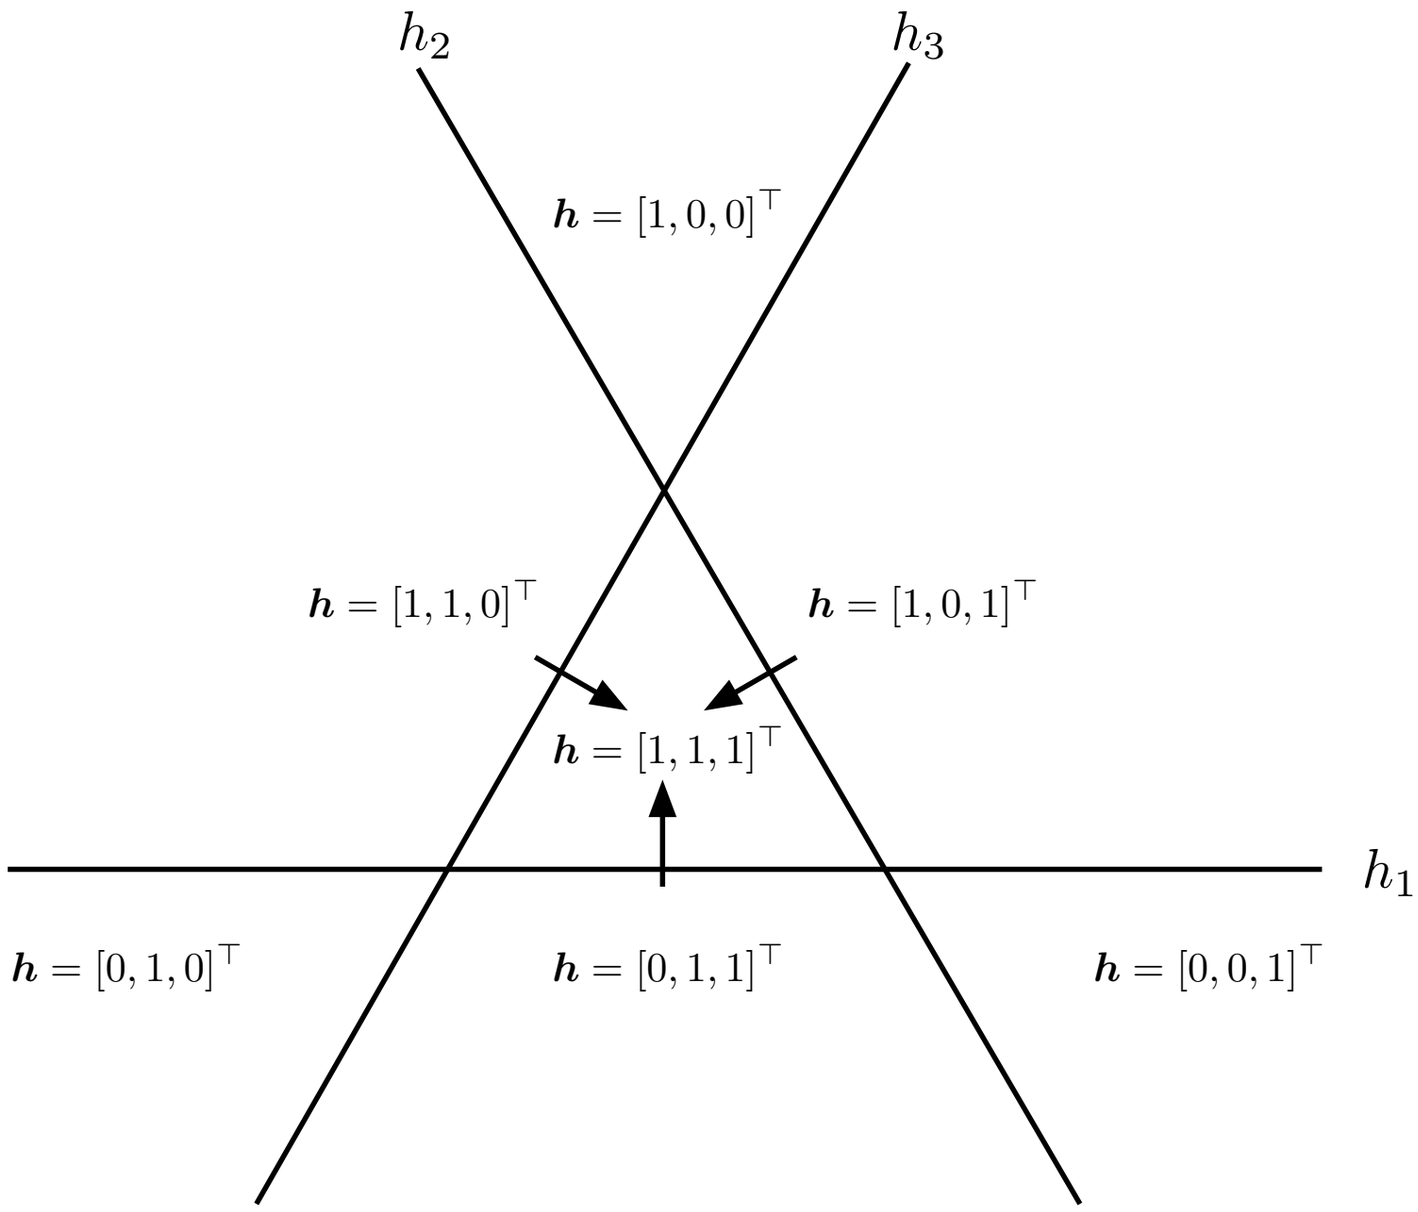

Figure 15.7

Binary Distributed Representations Divide Space Into Many Uniquely Identifiable Regions

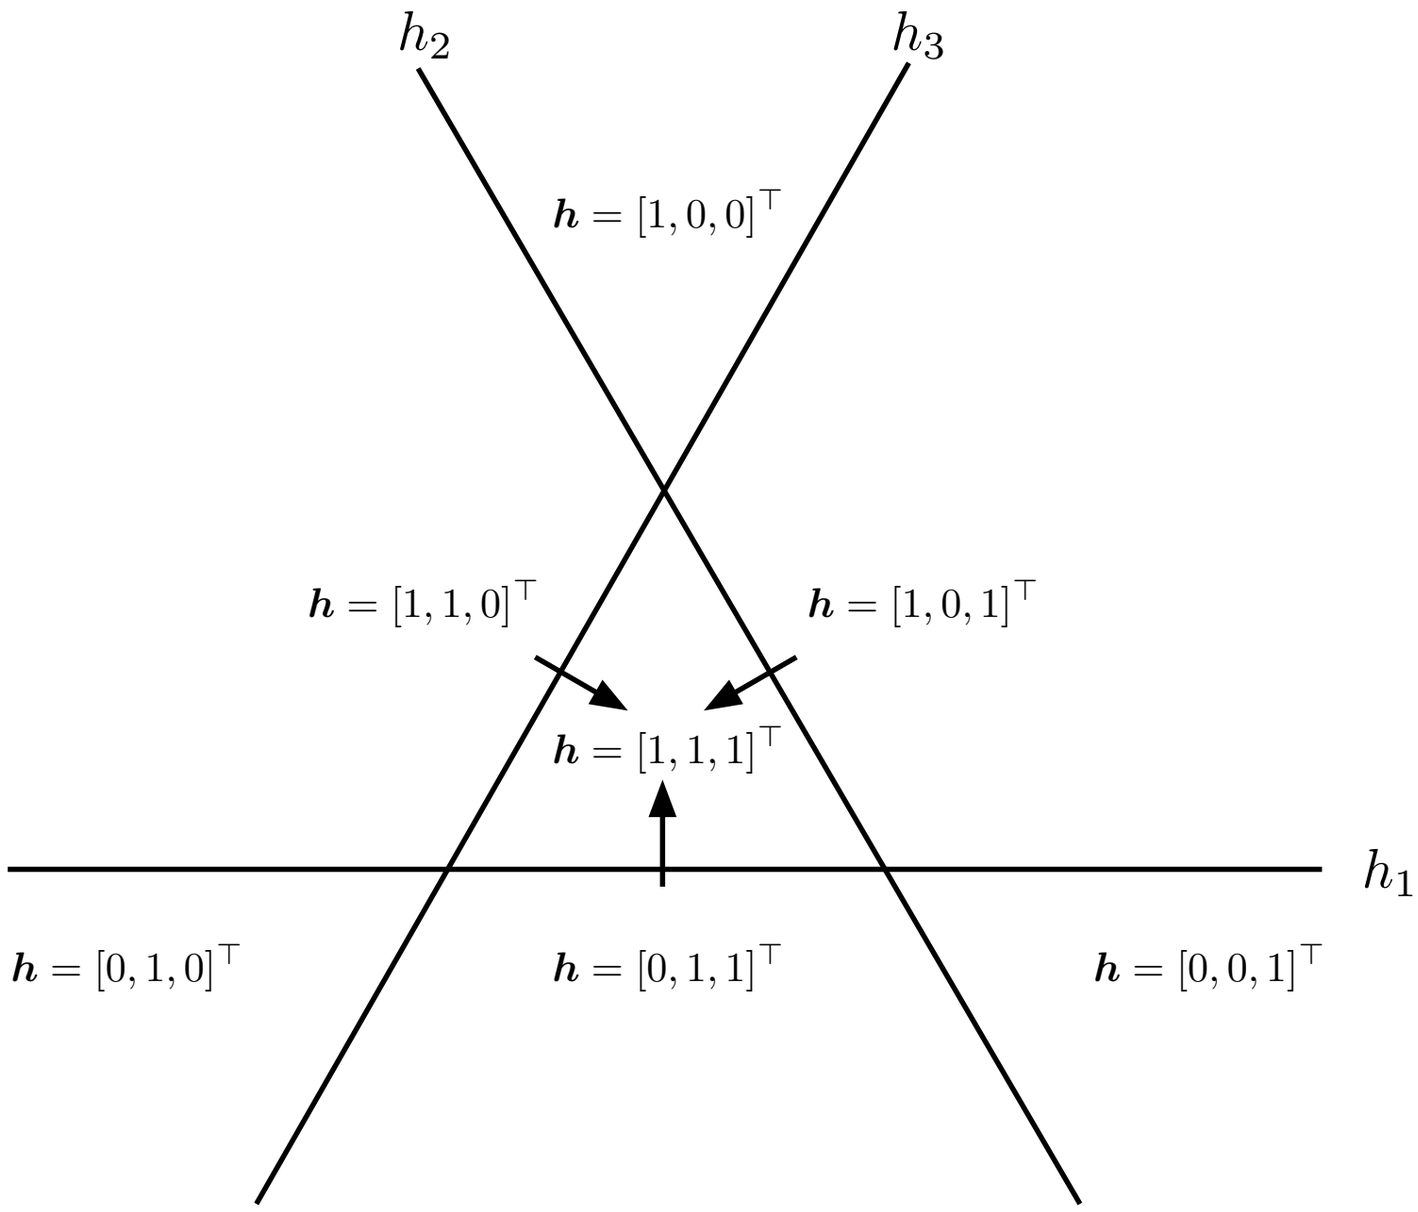

Figure 15.7

Nearest Neighbor Divides Space into one Region Per Centroid

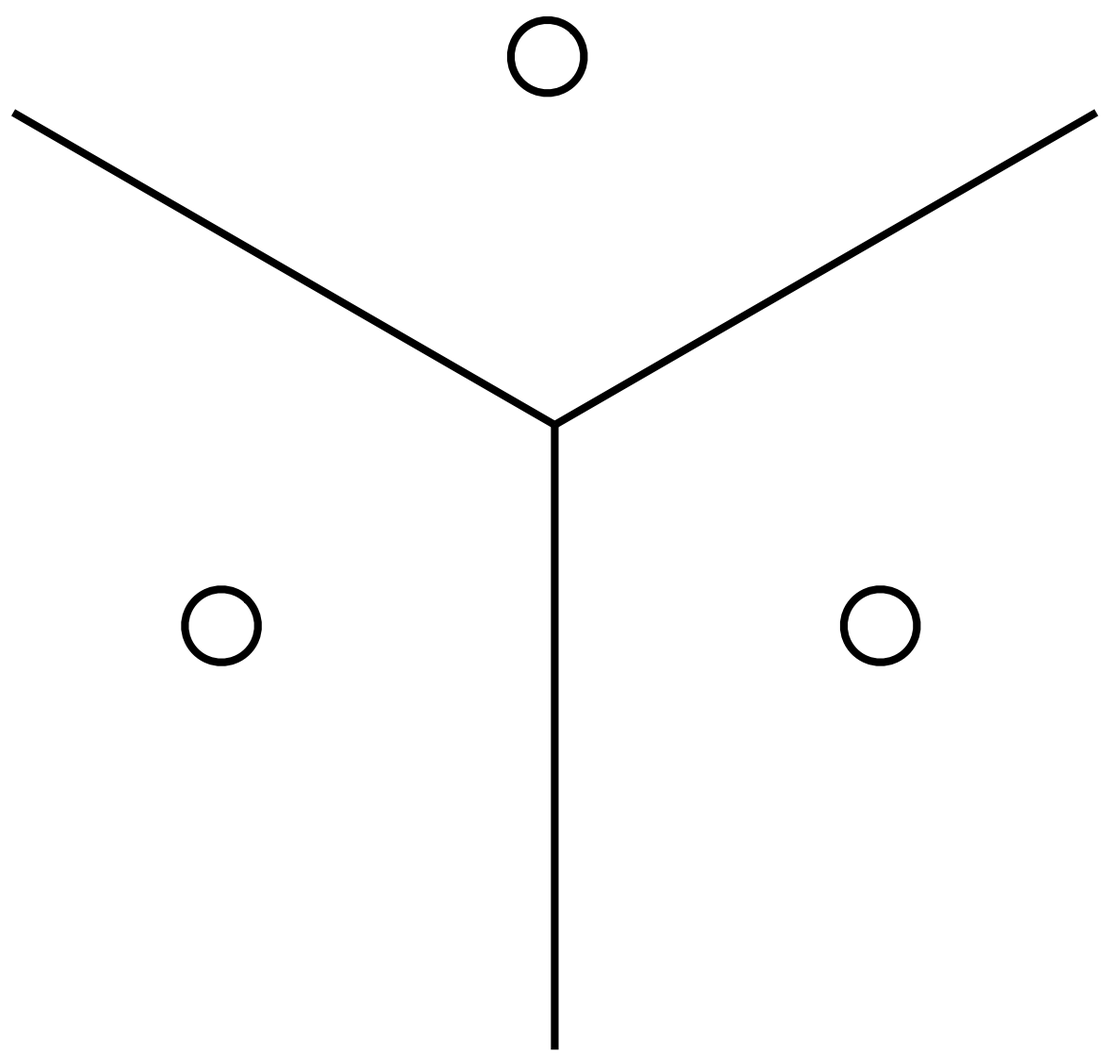

Figure 15.8

GANs learn vector spaces that support semantic arithmetic

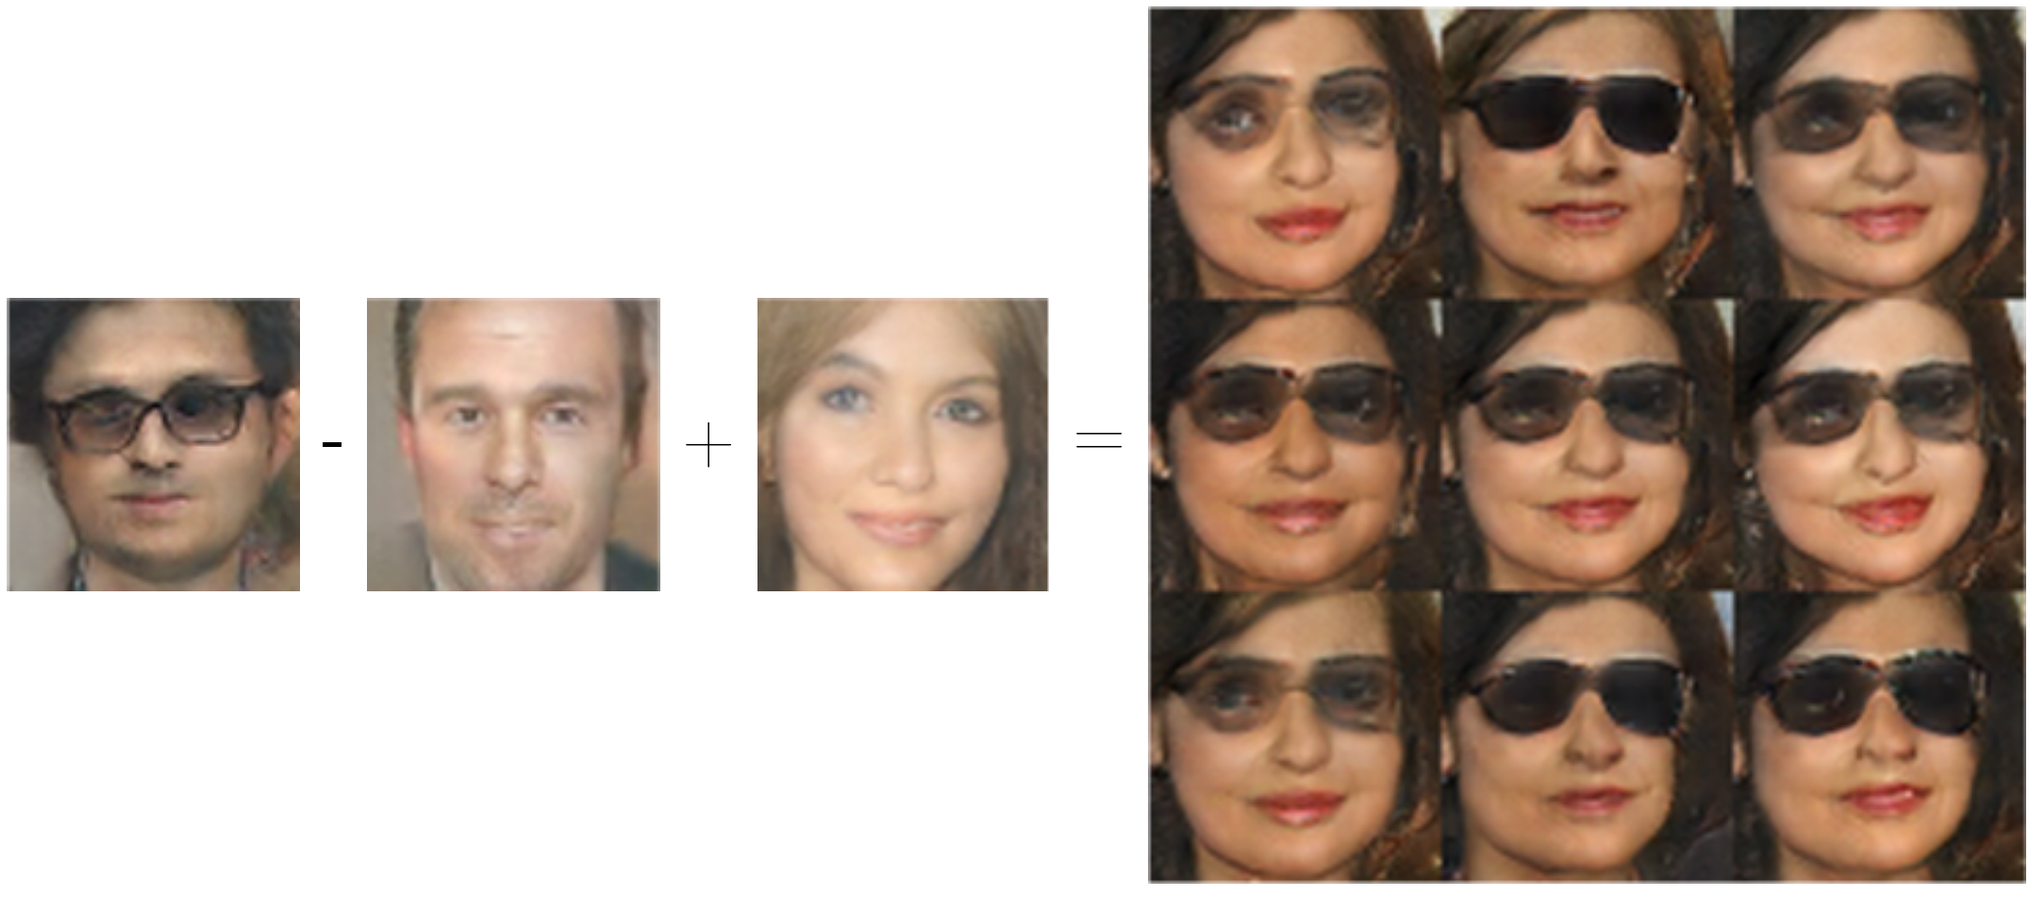

Figure 15.9