60K images → distill → 10 images

+ Fixed init → train

13% accuracy → 94% accuracy

50K images → distill → 100 images

+ Fixed init → train

9% accuracy → 54% accuracy

Dataset distillation on MNIST and CIFAR10

73K images
60K images
→ distill → 100 images encoding domain difference

+ Trained for SVHN → train

52% MNIST accuracy → 85% MNIST accuracy
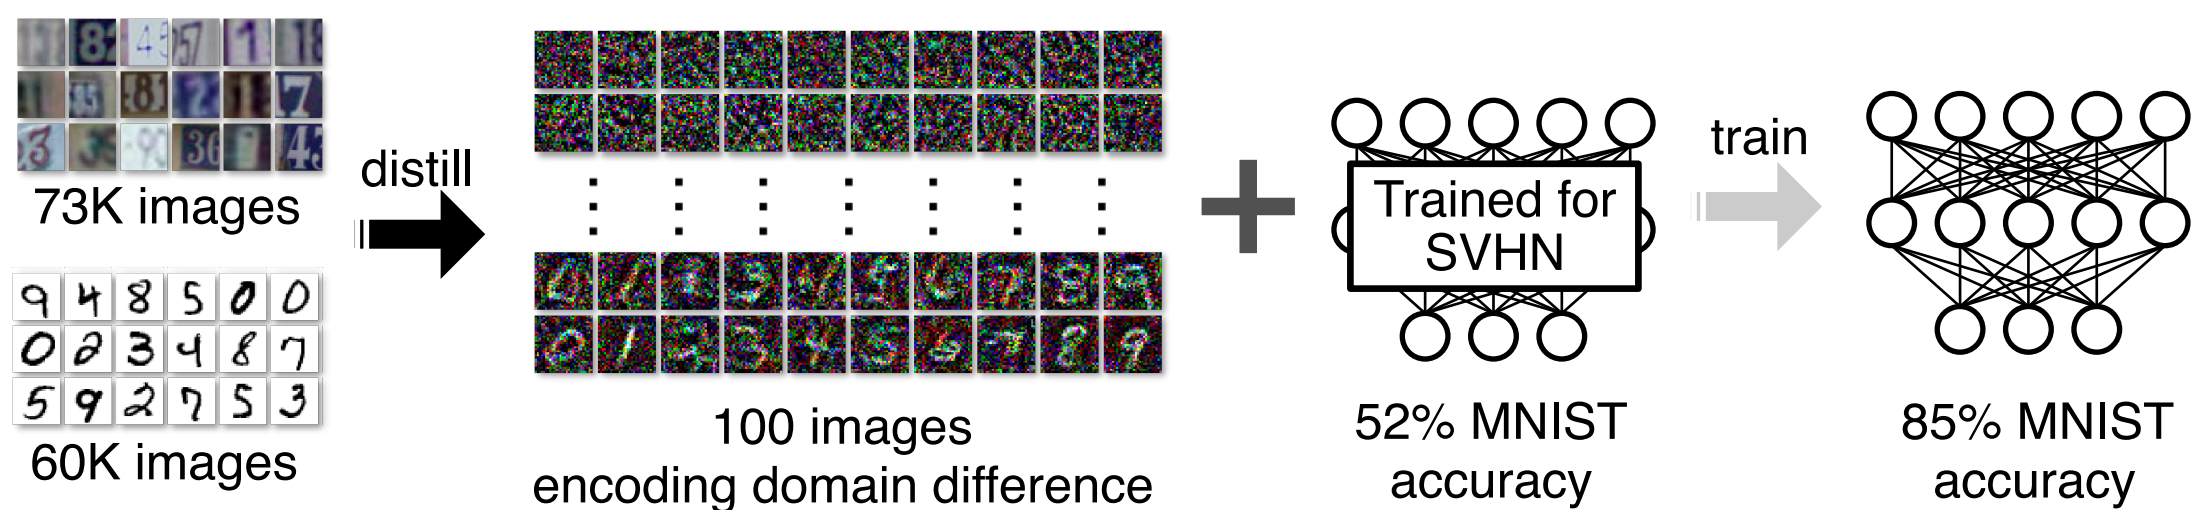
Dataset distillation can quickly fine-tune pre-trained networks on new datasets

50K images → distill → 300 attack images

+ Trained for CIFAR10 → train → Attacked Model

82% accuracy on class "plane" → 7% accuracy on class "plane"
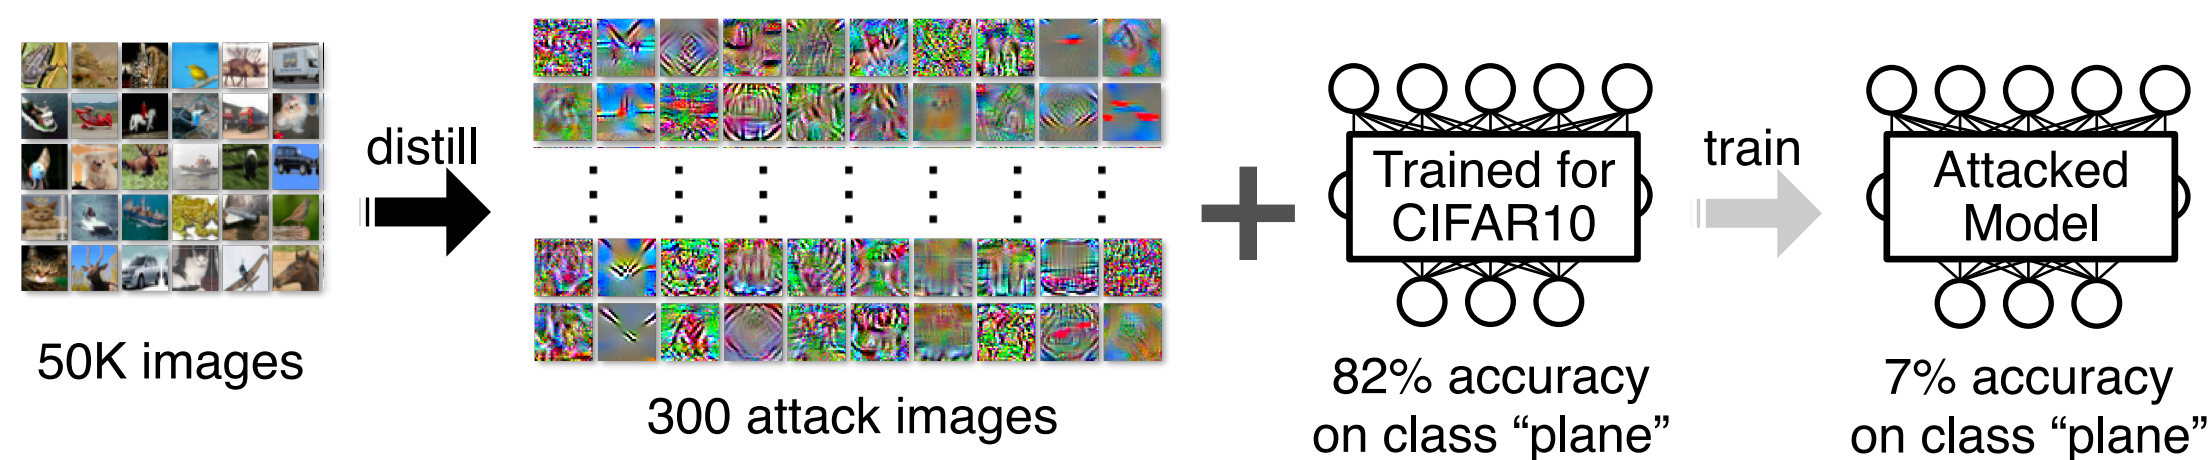
Dataset distillation can maliciously attack classifier networks