

Wearable ImageNet: Synthesizing Tileable Textures via Dataset Distillation

Supplementary Material

George Cazenavette¹ Tongzhou Wang² Antonio Torralba² Alexei Efros³ Jun-Yan Zhu¹

¹Carnegie Mellon University ²Massachusetts Institute of Technology ³UC Berkeley

A. Algorithm

Algorithm 1 Dataset Distillation via Trajectory Matching

Input: $\{\tau_i^*\}$: set of expert parameter trajectories trained on $\mathcal{D}_{\text{real}}$.

Input: M : # of updates between starting and target expert params.

Input: N : # of updates to student network per distillation step.

Input: $\mathcal{A}$: Differentiable augmentation function.

Input: $T^+ < T$: Maximum start epoch.

```
1: Initialize distilled data  $\mathcal{D}_{\text{syn}} \sim \mathcal{N}(0, I)$ 
2: Initialize trainable learning rate  $\alpha := \alpha_0$ 
3: for each distillation step... do
4:    $\triangleright$  Sample expert trajectory:  $\tau^* \sim \{\tau_i^*\}$  with  $\tau^* = \{\theta_t^*\}_0^T$ 
5:    $\triangleright$  Choose random start epoch,  $t \leq T^+$ 
6:    $\triangleright$  Initialize student network with expert params:
7:      $\hat{\theta}_t := \theta_t^*$ 
8:   for  $n = 0 \rightarrow N - 1$  do
9:      $\triangleright$  Sample a mini-batch of distilled images:
10:     $b_{t+n} \sim \mathcal{D}_{\text{syn}}$ 
11:     $\triangleright$  Take random crop with circular padding:
12:     $b_{t+n}^+ = \text{crop\_and\_pad}(b_{t+n})$ 
13:     $\triangleright$  Update student network w.r.t. classification loss:
14:     $\hat{\theta}_{t+n+1} = \hat{\theta}_{t+n} - \alpha \nabla \ell(\mathcal{A}(b_{t+n}^+); \hat{\theta}_{t+n})$ 
15:   end for
16:    $\triangleright$  Compute loss between ending student and expert params:
17:    $\mathcal{L} = \|\hat{\theta}_{t+N} - \theta_{t+M}^*\|_2^2 / \|\theta_t^* - \theta_{t+M}^*\|_2^2$ 
18:    $\triangleright$  Update  $\mathcal{D}_{\text{syn}}$  and  $\alpha$  with respect to  $\mathcal{L}$ 
19: end for
```

Output: distilled data $\mathcal{D}_{\text{syn}}$ and learning rate α

B. Performance as Training Data

Despite not explicitly following the dataset distillation objective, our distilled textures still achieve surprising classification results. By training on only random patches of our ImageSquawk textures, we achieve %39.6 test accuracy. Likewise, we achieve %26.8 test accuracy by only training on random patches of ImageFruit.

C. More Visualizations




Figure 1. ImageFruit Distilled Textures Tiled 3x3