Semi-supervised Synthesis of High-Resolution Editable Textures for 3D Humans Supplementary Material

Bindita Chaudhuri[1], Nikolaos Sarafianos[2], Linda Shapiro[1], Tony Tung[2]
[1]University of Washington, [2]Facebook Reality Labs Research, Sausalito
[1]{bindita,shapiro}@cs.washington.edu, [2]{nsarafianos,tony.tung}@fb.com

Introduction

One of the main advantages of our method is the ability to perform texture synthesis under two different test scenarios: (a) synthesis guided by exemplar images and (b) synthesis with random styles. This provides the user with the flexibility to either copy a unique style from an exemplar image of choice or explore random styles outside the training distribution without the need to search for a corresponding exemplar image. In this supplementary material, we present more results of texture map generation by our method under these test scenarios. We also provide examples of complex patterns generated by our method in Fig. 1, which shows that our network learns complex patterns in addition to uniform textures as garment styles.

Figure 1. Examples of checkered and floral patterns in the textures synthesized by our proposed method.

Synthesis guided by exemplar images

Exemplar-guided texture synthesis can be further divided into three cases:

1. *Reconstruction*, where the styles of all the regions of the output texture are taken from the input texture.
2. *Exemplar-guided style mixing*, where the styles from the input texture for some regions are mixed with the styles from an exemplar texture (different from the input texture) for the remaining regions.
3. *Non-exemplar guided style mixing*, where the styles from the input texture for some regions are mixed with the styles generated by random vectors for the remaining regions.

Fig. 2 shows examples of textures generated with these three cases. We show examples with varying geometry, different garment types and shapes, and a variety of skin color and hair styles.

Synthesis with random styles

Texture synthesis with random styles can be divided into two cases:

1. *Random synthesis for all classes*, where we can randomly sample from the learned per-class style distributions to generate a random texture map.
2. *Style control of selected classes*, where we can fix the styles of some regions and change the styles of the remaining regions by controlling the corresponding style vectors in the style matrix at the generator's input.

We show examples to demonstrate these two test cases in Fig. 3. Our network can successfully disentangle the contextual relationship among the styles of various regions. For example, although our training data contains examples of grey-haired persons wearing mainly suits and pants, our network can generate a grey-haired person wearing a t-shirt and shorts.

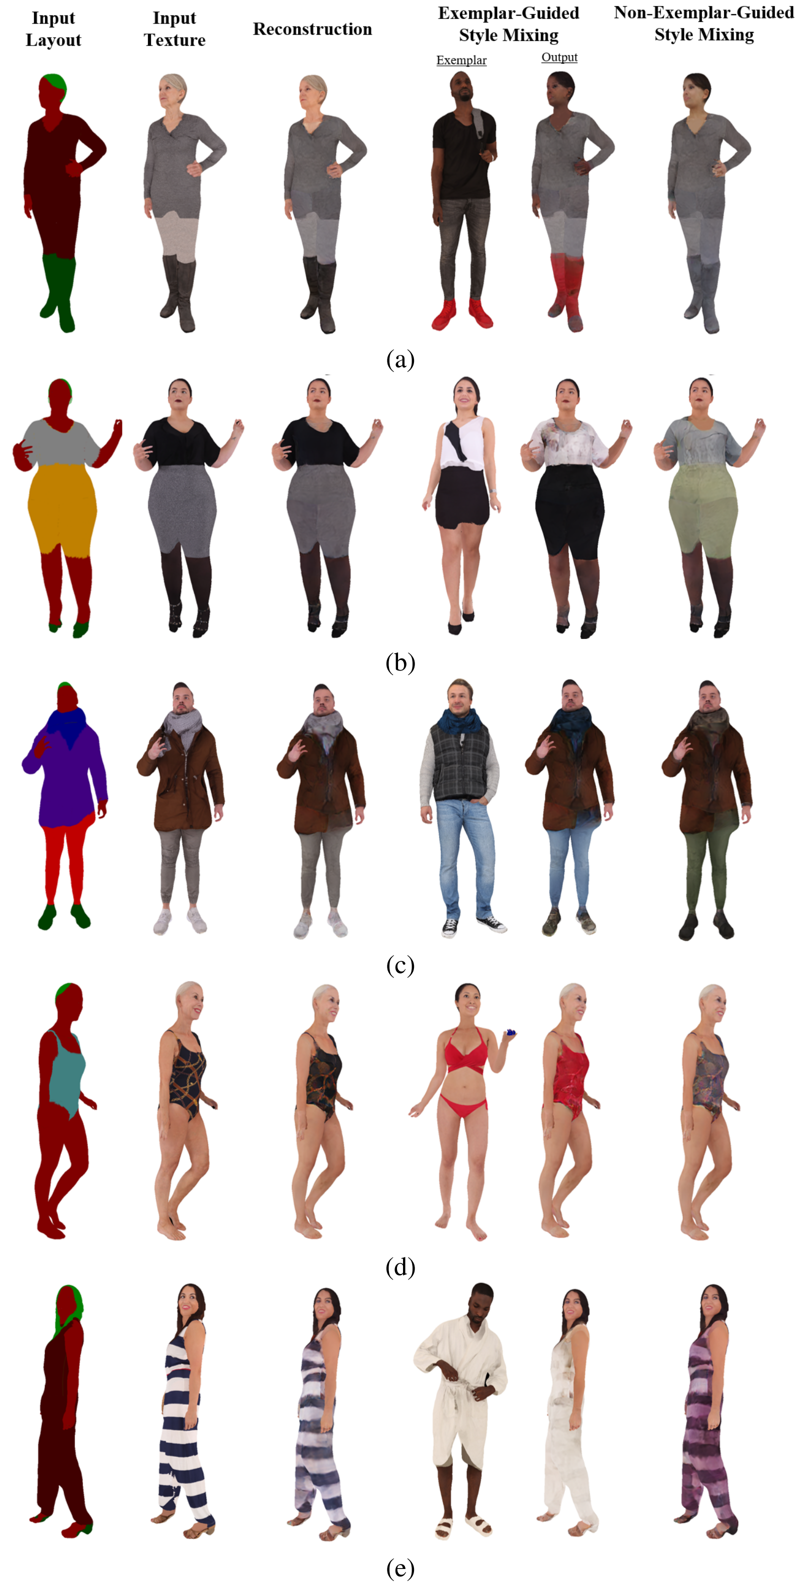

Figure 2. **Synthesis guided by exemplar images**. We change the styles of the following classes: (a) hair, skin and shoes, (b) top and skirt, (c) scarf, jeans and shoes, (d) swimsuit, (e) jumpsuit and shoes.

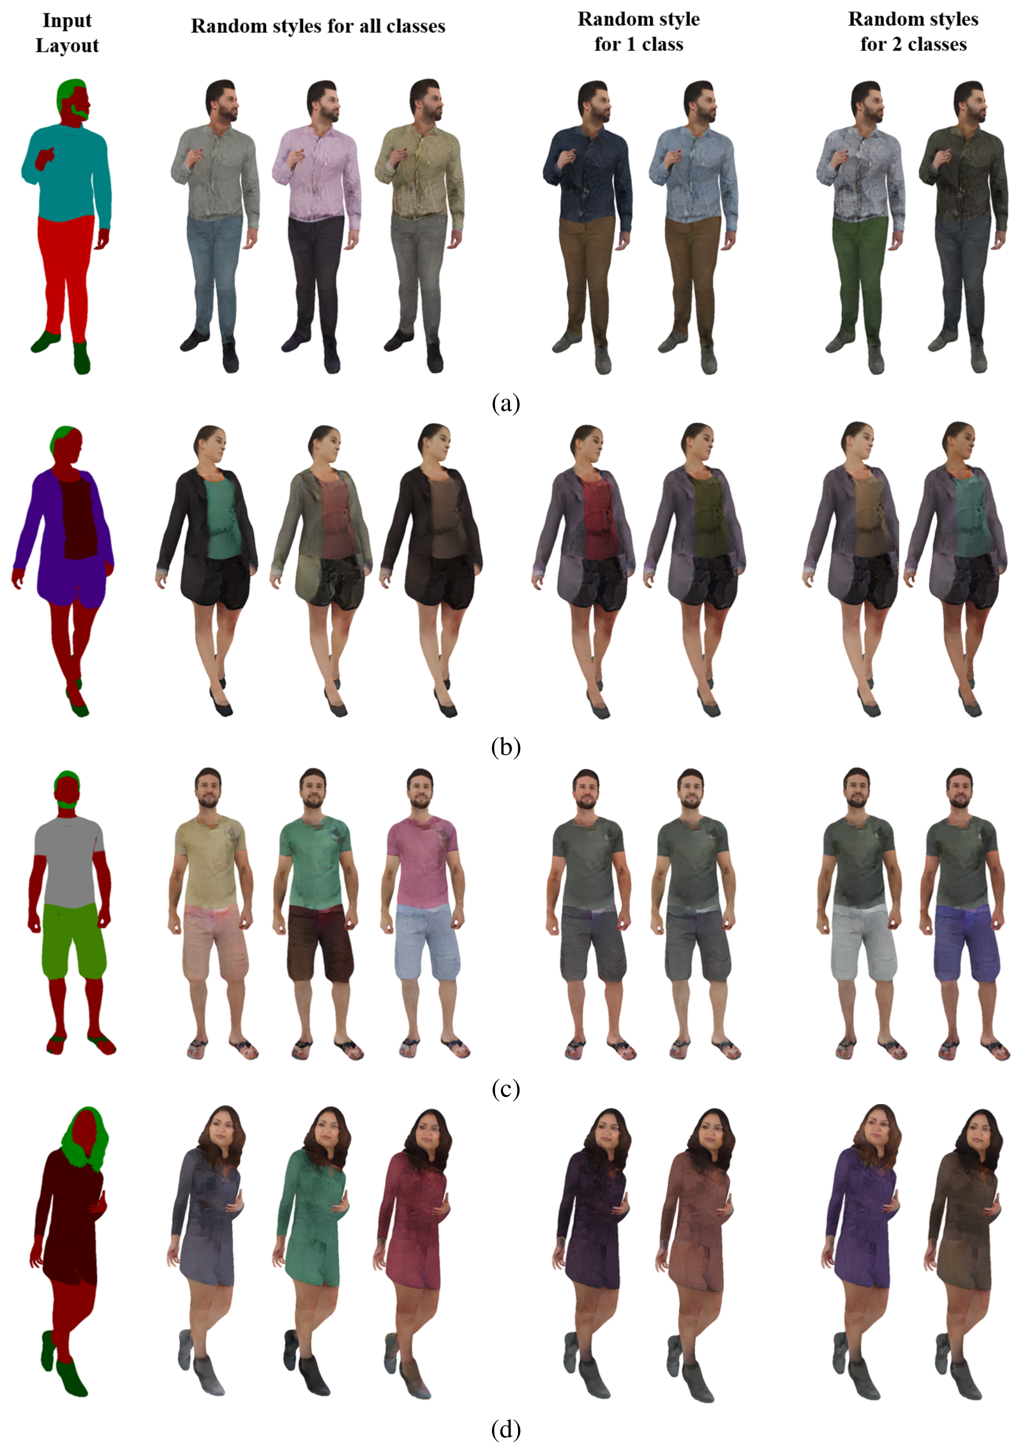

Figure 3. **Synthesis with random styles**. The first 3 textures for each example are generated by sampling random styles for all the classes. The next 2 textures are generated by keeping the styles for all classes fixed except for a single class which is (a) shirt, (b) inner top, (c) skin, and (d) dress. The last 2 textures are generated by keeping the styles for all classes fixed except for two classes which are (a) shirt and pants, (b) inner top and skin, (c) skin and shorts, and (d) dress and hair.