Generative Escher Meshes

NOAM AIGERMAN* , University of Montreal and Adobe Research, Canada
THIBAULT GROUEIX* , Adobe Research, USA

INTRODUCTION

INTRODUCTION

- This paper proposes a fully-automatic, text-guided generative method for producing periodic, repeating, tile-able 2D art.
- our method generates non-square tilings which comprise solely of repeating copies of the same object.
- It achieves this by optimizing both geometry and color of a 2D mesh, in order to generate a non-square tile in the shape and appearance of the desired object.

RELATED WORK

RELATED WORK

TilinGNN: learning to tile with self-supervised graph

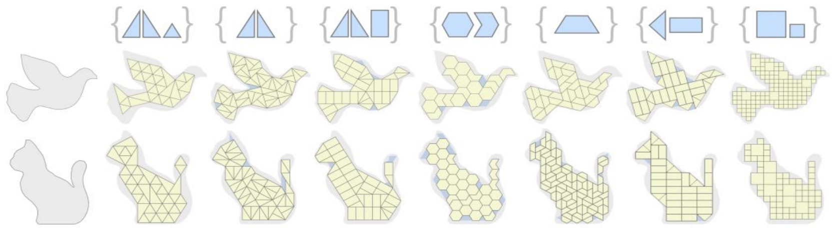

SeamlessGAN: Self-Supervised Synthesis of Tileable Texture Maps

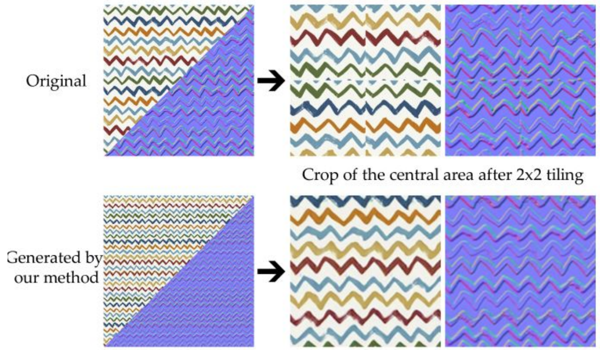

METHOD

METHOD

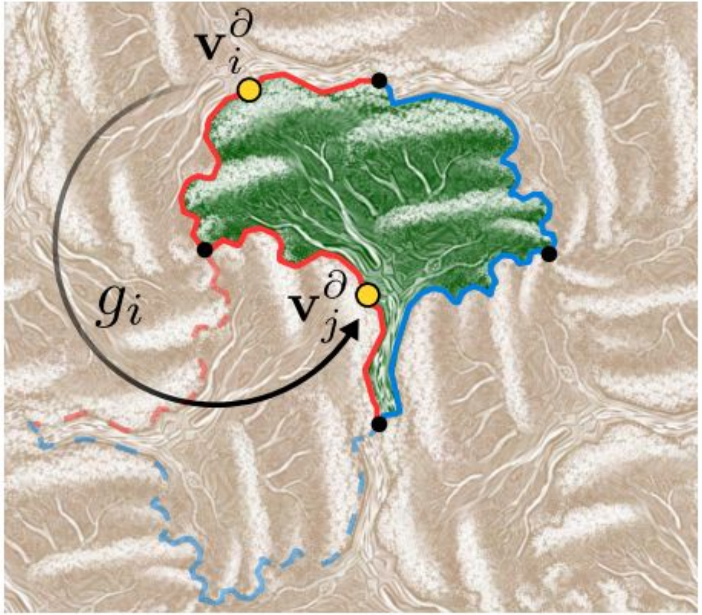

$\mathbf{v}_i^\partial$
g_i
$\mathbf{v}_j^\partial$

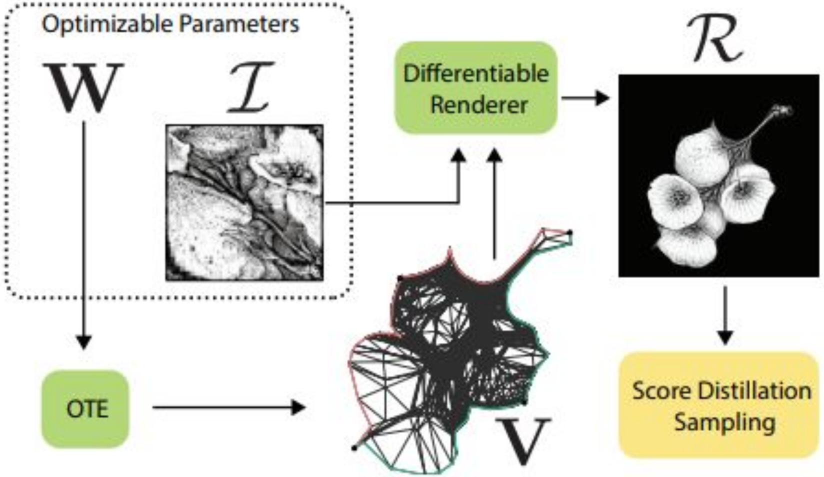

Optimizable Parameters
$\mathbf{W}$
$\mathcal{I}$
Differentiable Renderer
$\mathcal{R}$
OTE
$\mathbf{V}$
Score Distillation Sampling

METHOD

- Necessary and sufficient conditions for tiling
  - (1) All boundary vertices satisfy the boundary conditions Eq. (1).
  - (2) the mesh has no overlaps with itself.
  -
  - $$\mathbf{v}_i^{\partial} = g_i\left(\mathbf{v}_j^{\partial}\right), \qquad (1)$$

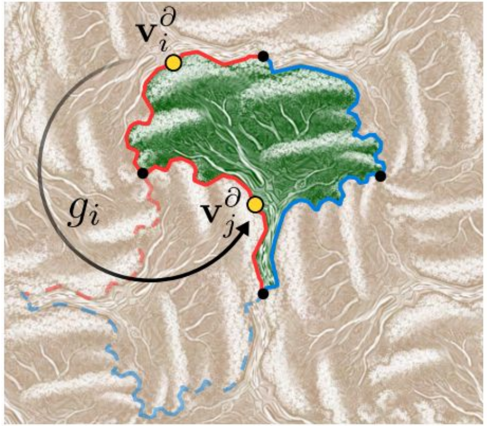

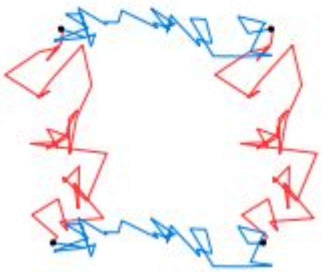

METHOD

- An optimizable, unconstrained representation of tiles
    - Their method works by solving a discrete Laplace equation
      , placing each vertex s.t. it is a weighted average of its neighbors.
    - Their method works by solving a discrete Laplace equation, placing each vertex s.t.
    -

$$\sum_{j \in \mathcal{N}^i} w_{ij}(\mathbf{v}_j - \mathbf{v}_i) = 0, \qquad (2)$$

METHOD

- Text-guided optimization of the tile's appearance
  - Rendering the tile into an image
  - leverages a pre-trained diffusion model to define a loss(DREAMFUSION)
  - Given a rendered image R, SDS "injects" it into the diffusion process, by adding noise to it and then applying a denoising step of the model, conditioned on the desired text prompt, thereby obtaining a modified version of the input image, R˜ , slightly more-correlated to the given text prompt

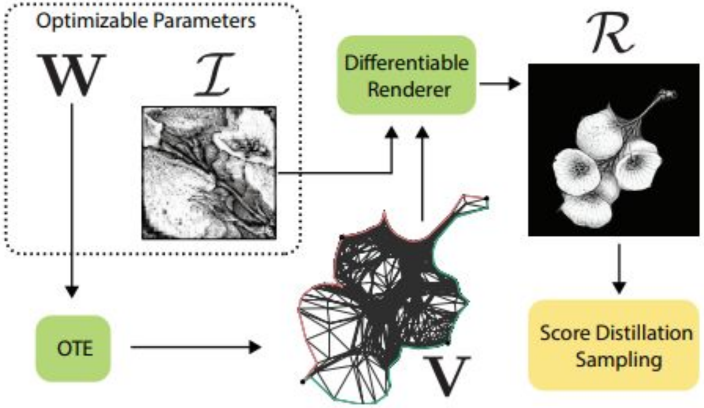

METHOD

- We employ two simple strategies to prevent this:
  - first, we randomizethe background color of the renderings to prevent SDS from using a specific color as a background color the mesh
  - second, we set a higher learning rate for the geometry optimization than for the texture, such that textures emerge slower, giving the shape time to form.

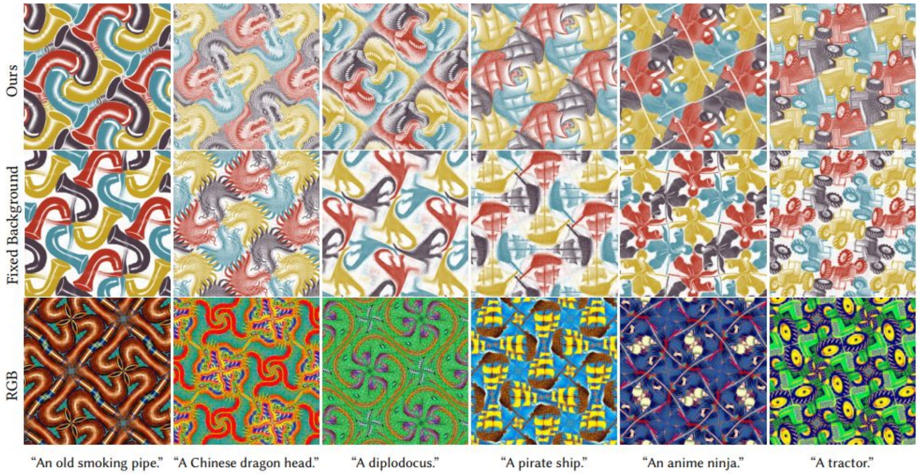

RESULTS

RESULTS

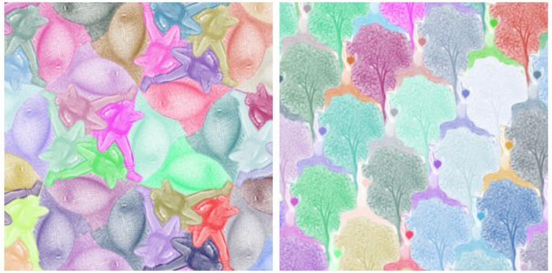

"A scuba-diver."
"A tropical fish."

"Man doing yoga."
"A tree."

"An open hand."

"An espresso."

"An ant."

CONCLUSION

CONCLUSION

- in its current state, our method is heavily-dependent on Score Distillation Sampling and inherits its limitations, namely slow runtime, tendency to produce oversaturated colors, as well as lack of granular control over the resulting imagery. As the boom of diffusion-based generative techniques is at its relative infancy, we believe these problems will be solved quite soon. Since we treat SDS merely as a blackbox for providing our framework with a loss, we could seamlessly swap it out for any new technique which will exhibit better performance.