Supplemental Material
Transport-Based Neural Style Transfer for Smoke Simulations

1 EXTENDED 3D RESULTS

Figure 1 shows additional results computed with our transport-based neural style transfer method.

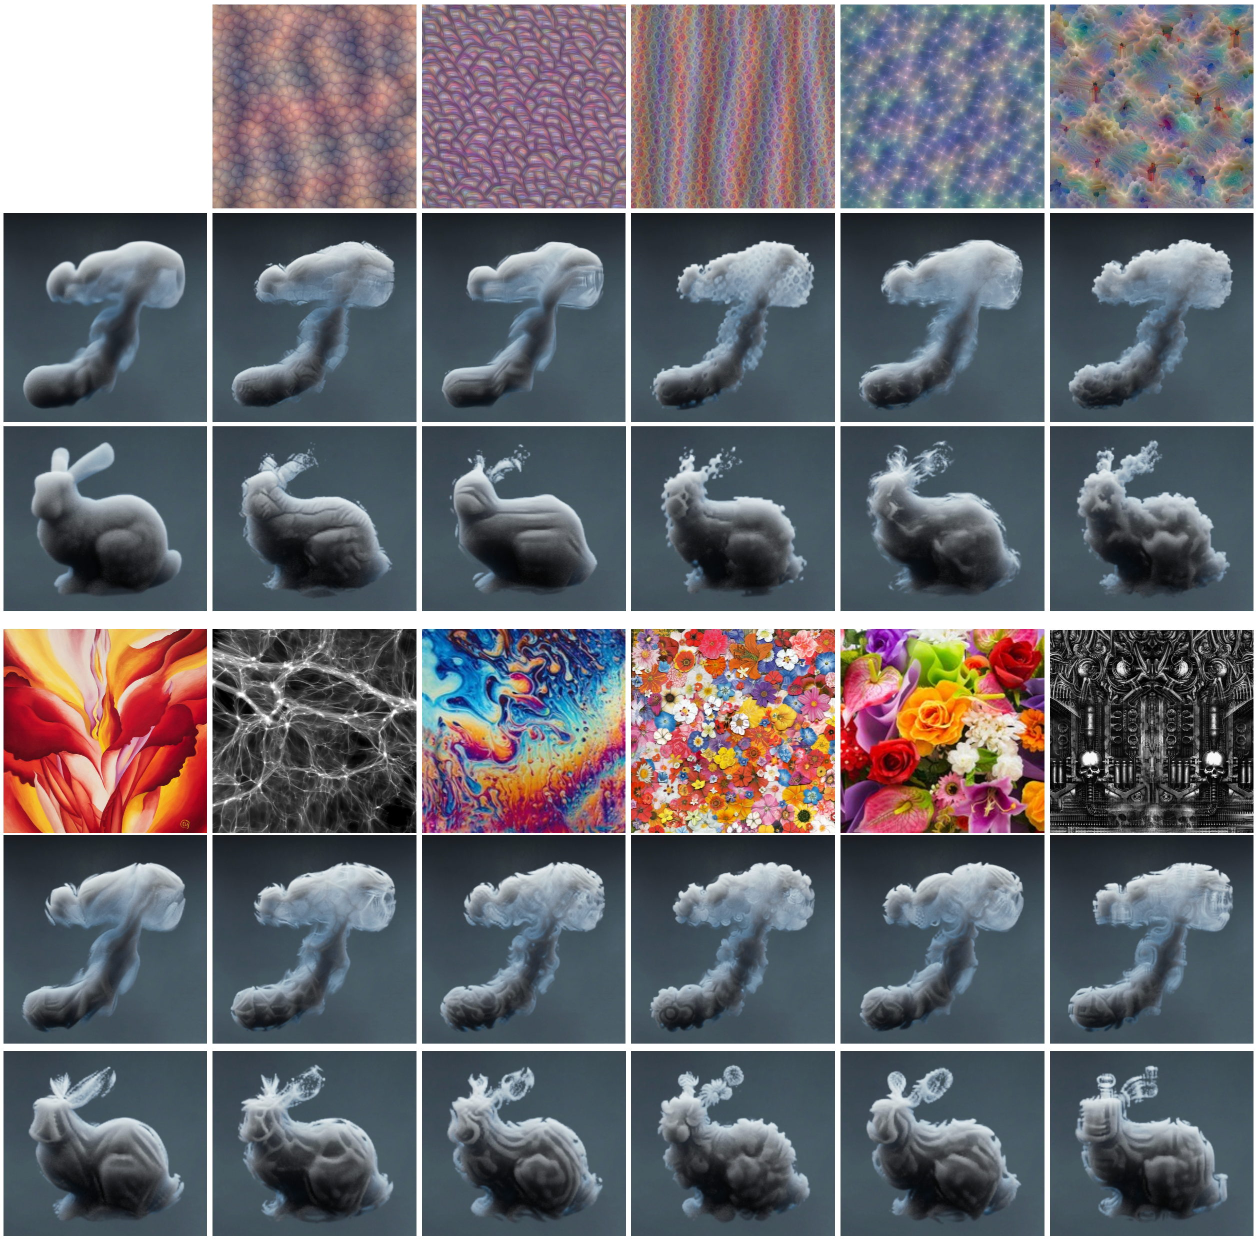

Fig. 1. Additional results showing semantic (top) and style transfer (bottom) applied to the smokejet and bunny smoke simulation (top left).

Examples in the paper apply a soft mask (Figure 2, right) to the stylization velocity field, in order to conform the stylization output to the original smoke silhouette. Figure 2 compares results obtained without employing the soft mask for irrotational, mixed and incompressible examples. Not applying the mask causes the smoke to spread, especially in the case of incompressible velocity fields.

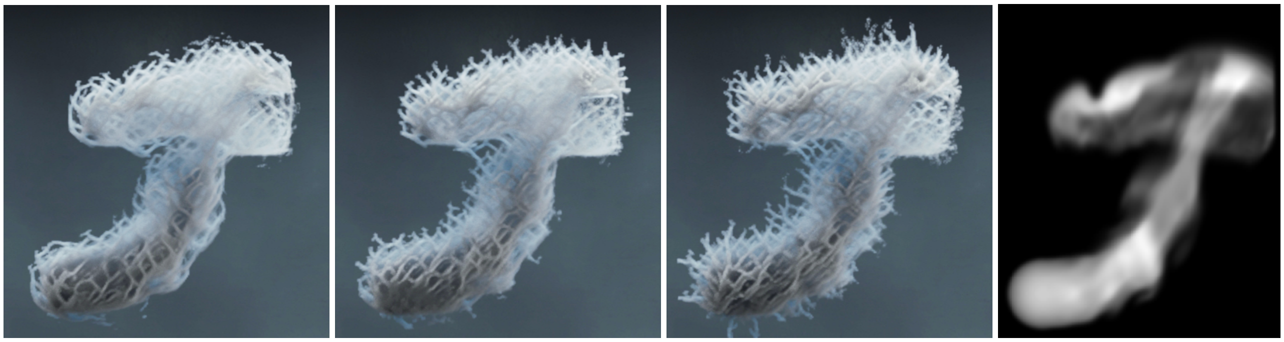

Fig. 2. Results from semantic transfer of a net structure without soft mask (rightmost). Irrotational, mixed and incompressible velocity fields, from left to right.

Figure 3 shows the amount of smoke dissipation over time by the decomposition of the stylization velocity field into its incompressible and irrotational parts. As expected, a streamfunction-based (incompressible) velocity stylization shows less dissipation than those of potential field based (irrotational) velocity stylization. Conservation errors due numerical integration on the advection algorithm, although, are still present. This causes the streamfunction-based velocity field to not exactly conserve the original amount of smoke.

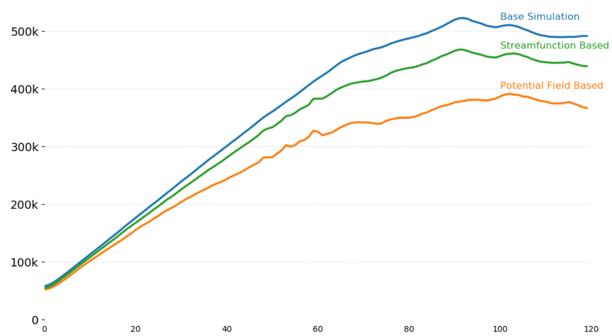

Fig. 3. Density amount plot comparison for incompressible and irrotational velocity fields. Using a streamfunction-based (incompressible) velocity fields reduces the loss of density amount compared to a irrotational approach.

We provide an example of stylization on smokejet simulation with a sphere shape obstacle to see how our method works on the presence of boundaries. Note that the mask is not applied in this case, since the masking would guarantee zero penetration on boundaries. Figure 4 shows density field visualizations of the base simulation, stylized density field, simulation velocity fields and stylized velocity fields, from top to bottom. Note that since the original simulation velocity field is already boundary-respecting, thus our method only yields slight penetrations on the obstacle boundary.

Figure 5 shows the effects of varying the number of iterations and the learning rate size. The learning rate affects structures significantly, higher learning rates result in more pronounced details. This is also visible in Figure 6 where different learning rates for the Starry Night Style transfer were used.

2 ANALYSIS ON DIFFERENTIABLE RENDERING METHODS

We compared a simpler alternative rendering method to the one presented on the original paper. The image is calculated by simply integrating the transmittance along the viewing ray as

$$I_{ij} = 1 - e^{-\alpha \int_0^{rmax} d(\mathbf{r}_{ij})\, dr}. \tag{1}$$

Figure 7 shows the difference between the rendering method proposed on the paper and the one in Equation (1). We vary the transmittance absorption factor (γ) for the stylization of the bunny example (left to right); top image sequence shows the differentiable rendering method, while bottom one shows the alternative rendering computed from Equation (1). While low transmittance values (leftmost) produce similar renderings and stylizations for both approaches, using Equation (1) quickly saturates the image as γ increases. The differentiable rendering method, however, creates structures that are present even in examples with higher transmittance absorption factors (rightmost).

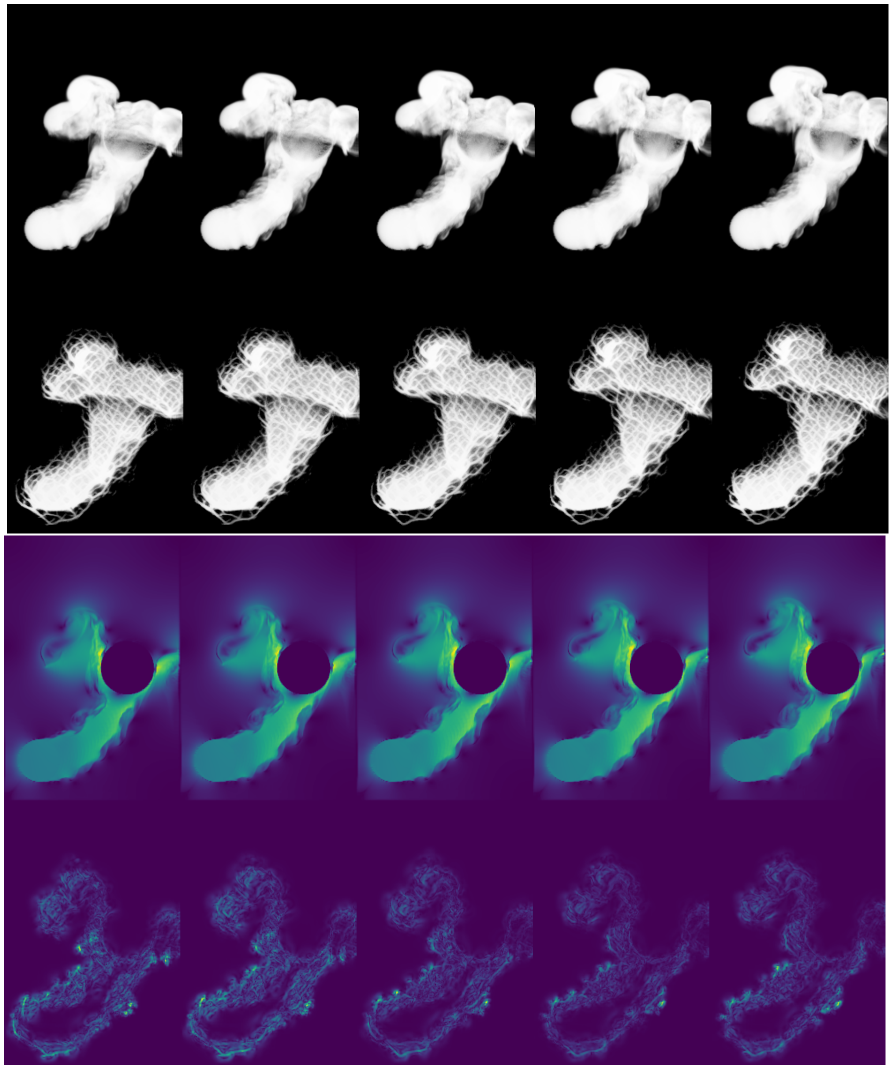

Fig. 4. Stylization on smokejet simulation with a sphere shape obstacle. First row shows base simulation, second row shows stylized density fields, third row represents the middle slice view of magnitude of base simulation velocity fields, while fourth row shows those of stylization velocity fields. Note that no soft mask is used.

3 2D EXAMPLES

We compare the value-based and transport-based density optimization for a 2D smoke simulation in Figure 8. The value-based method shows sharper details, but introduces ghosting artifacts in temporal sequences as spurious density sources and sinks are added. Our transport-based approach leads to temporally smoother stylizations as the total volume is conserved. We also compare the temporal coherency with different window sizes for a 2D smoke simulation in Figure 9. The images are cropped from the top right part of the stylized image sequence. The top row shows the results using a window size of 1, and the bottom row is generated with a window size of 9. The larger the window size, the more overall structure is conserved (highlighted by the color-coded circles).

Figure 10 shows different 2D semantic and style transfer examples. We also provide comparisons about different abstraction levels of style features for 2D simulations (Figure 11). We can control low (left), medium (middle) and high (right) levels of features. Figure 12 shows how changing the resolution of the input fluid simulation affects the stylization. Higher resolution simulations (right) lead to sharper structures than coarse simulations (left).

Fig. 5. Influence of iteration number and learning rate. From left to right: 5, 10, and 20 iterations. From top to bottom: learning rate of 0.001, 0.0005, 0.0001.

Fig. 6. Influence of using different learning rates for the Starry Night style transfer example. A higher learning rate (right, 5× higher) results in more pronounced structures than when using lower learning rates (left), but also more noisy results.

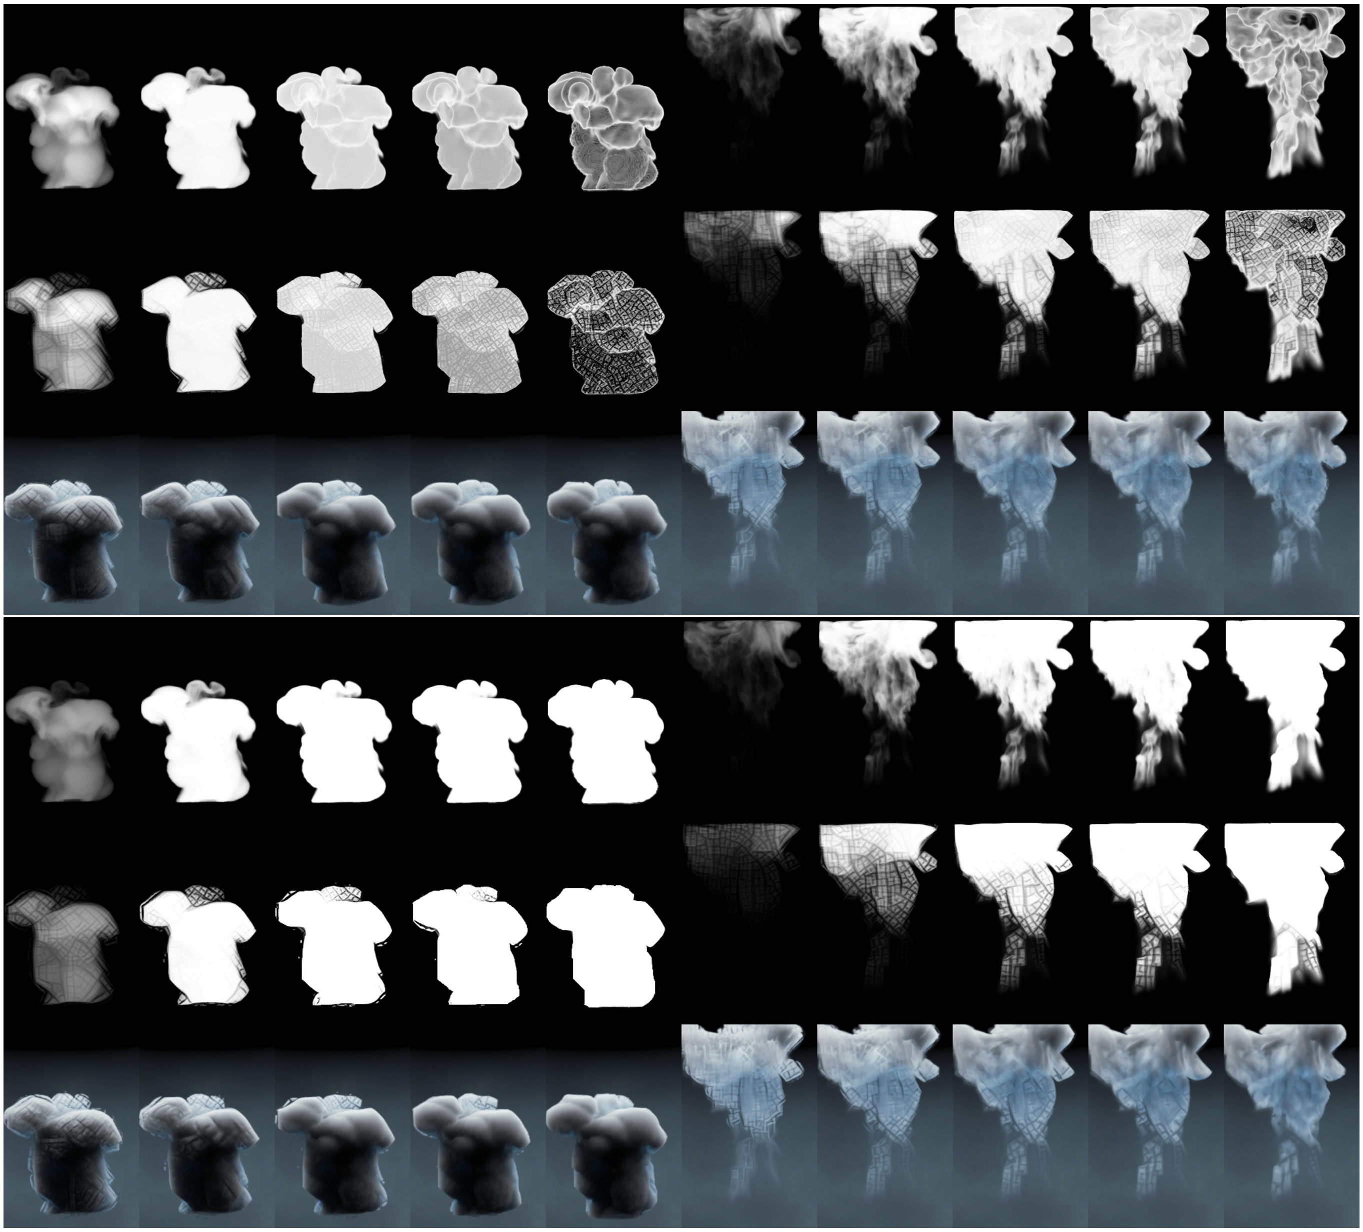

Fig. 7. Our differentiable rendering method (top) versus the one with Equation (1) (bottom).

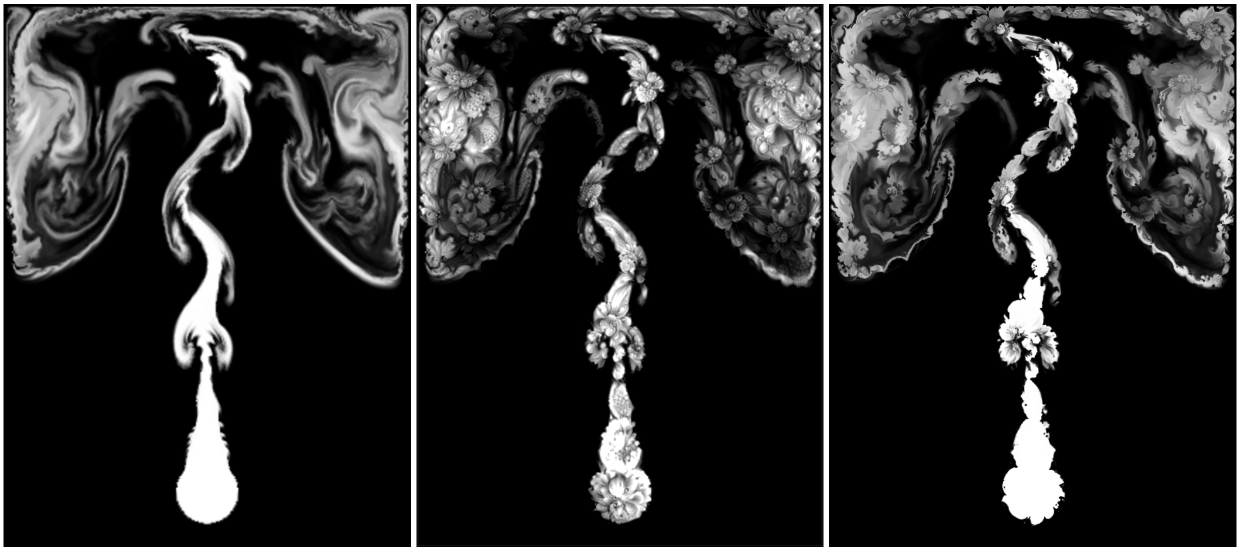

Fig. 8. Value-based density optimization (middle) versus transport-based density optimization (right). The input smoke simulation is shown on the left.

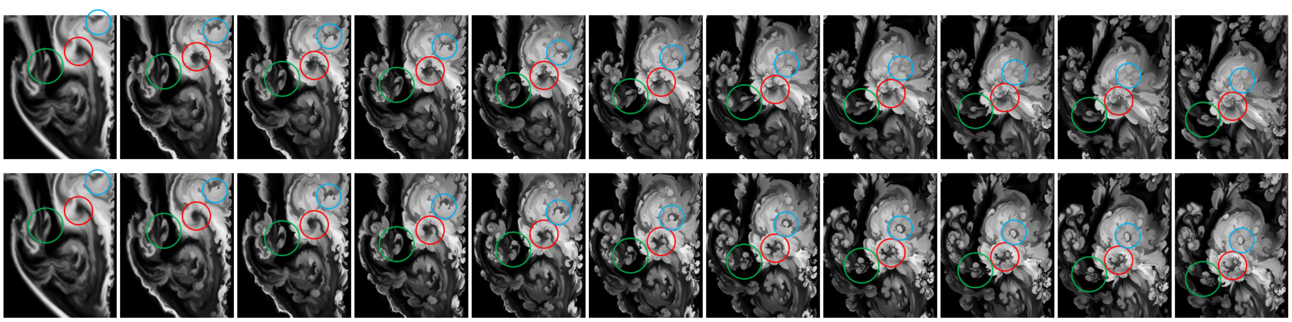

Fig. 9. Comparison of temporal coherency using different window sizes of 1 and 9.

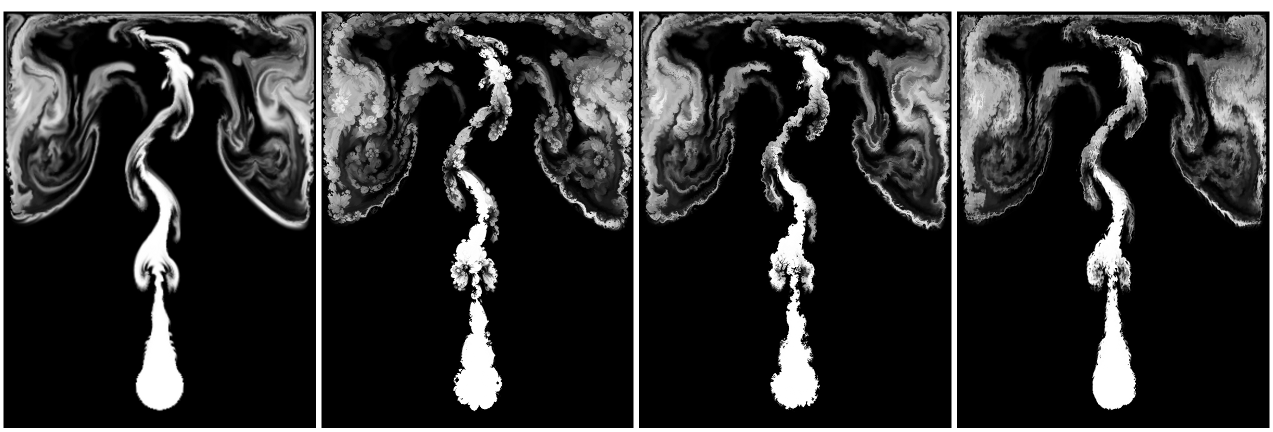

Fig. 10. From left to right: coarse input simulation, flower, volcano, and fire stylizations.

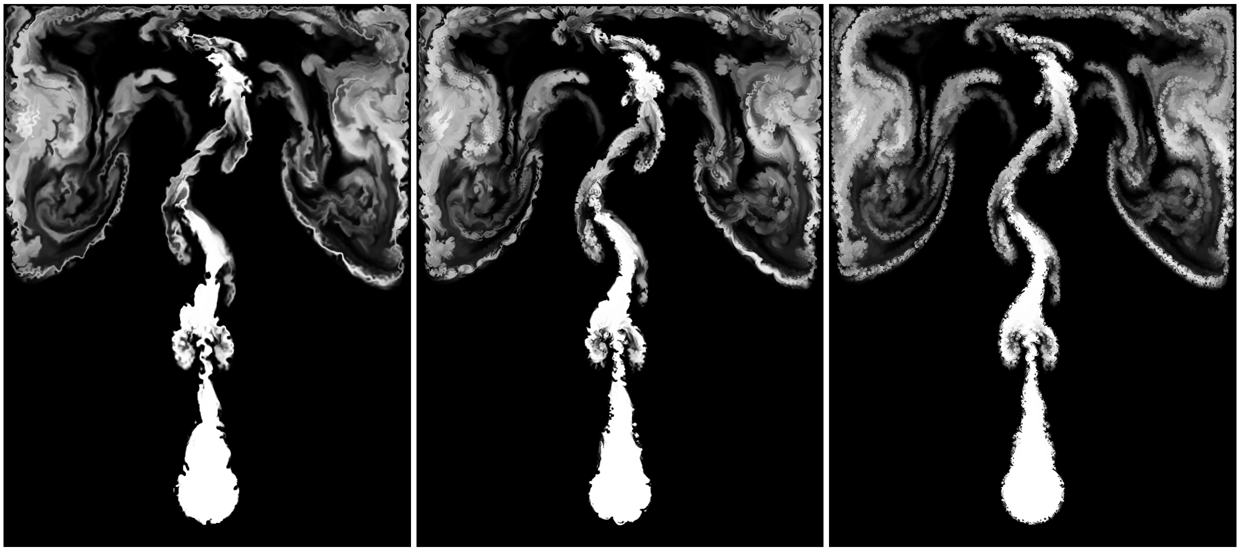

Fig. 11. Low (left), medium (middle) and high (right) levels of style features.

Fig. 12. Stylization applied to a low-resolution (left) and high-resolution (right) fluid simulation. More detailed structures are synthesized with higher resolutions.