SIGGRAPH ASIA 2024 TOKYO 東京

Conference | 3–6 December 2024
Exhibition | 4–6 December 2024
Venue | Tokyo International Forum, Japan

Spatiotemporal Bilateral Gradient Filtering for Inverse Rendering

Sponsored by acm

Organized by koelnmesse

Spatiotemporal Bilateral Gradient Filtering for Inverse Rendering

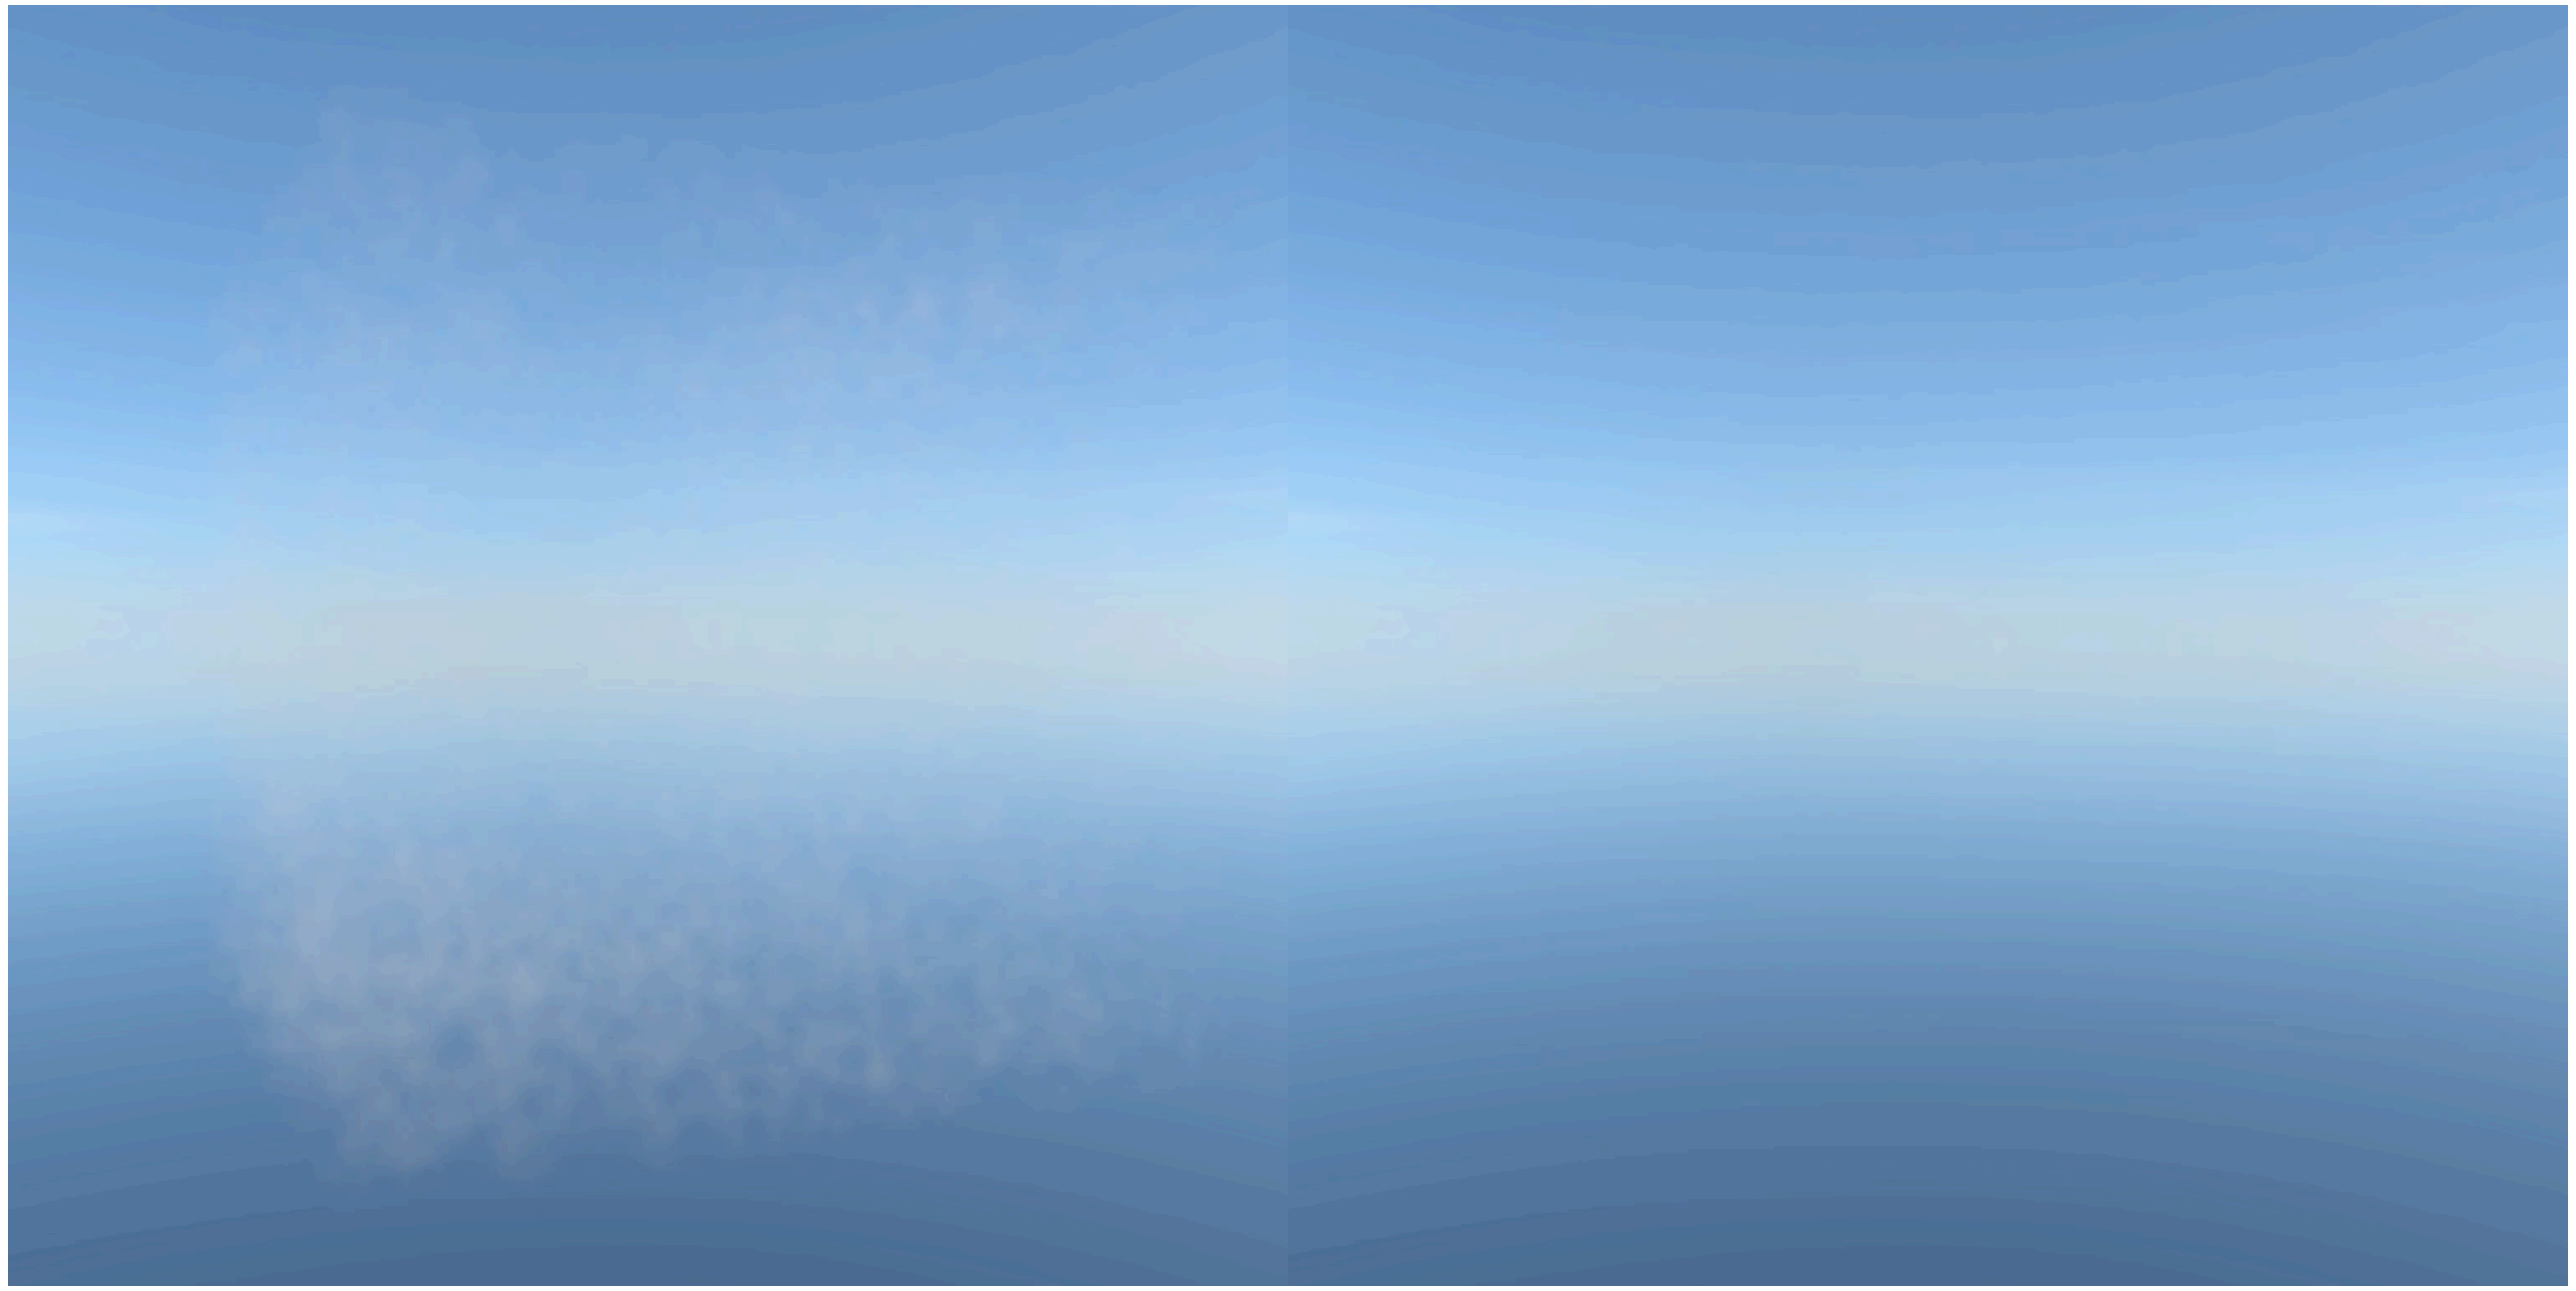

Wesley Chang*, **Xuanda Yang***, Yash Belhe*, Ravi Ramamoorthi,
Tzu-Mao Li - University of California San Diego

Forward and Inverse Rendering

Forward and Inverse Rendering

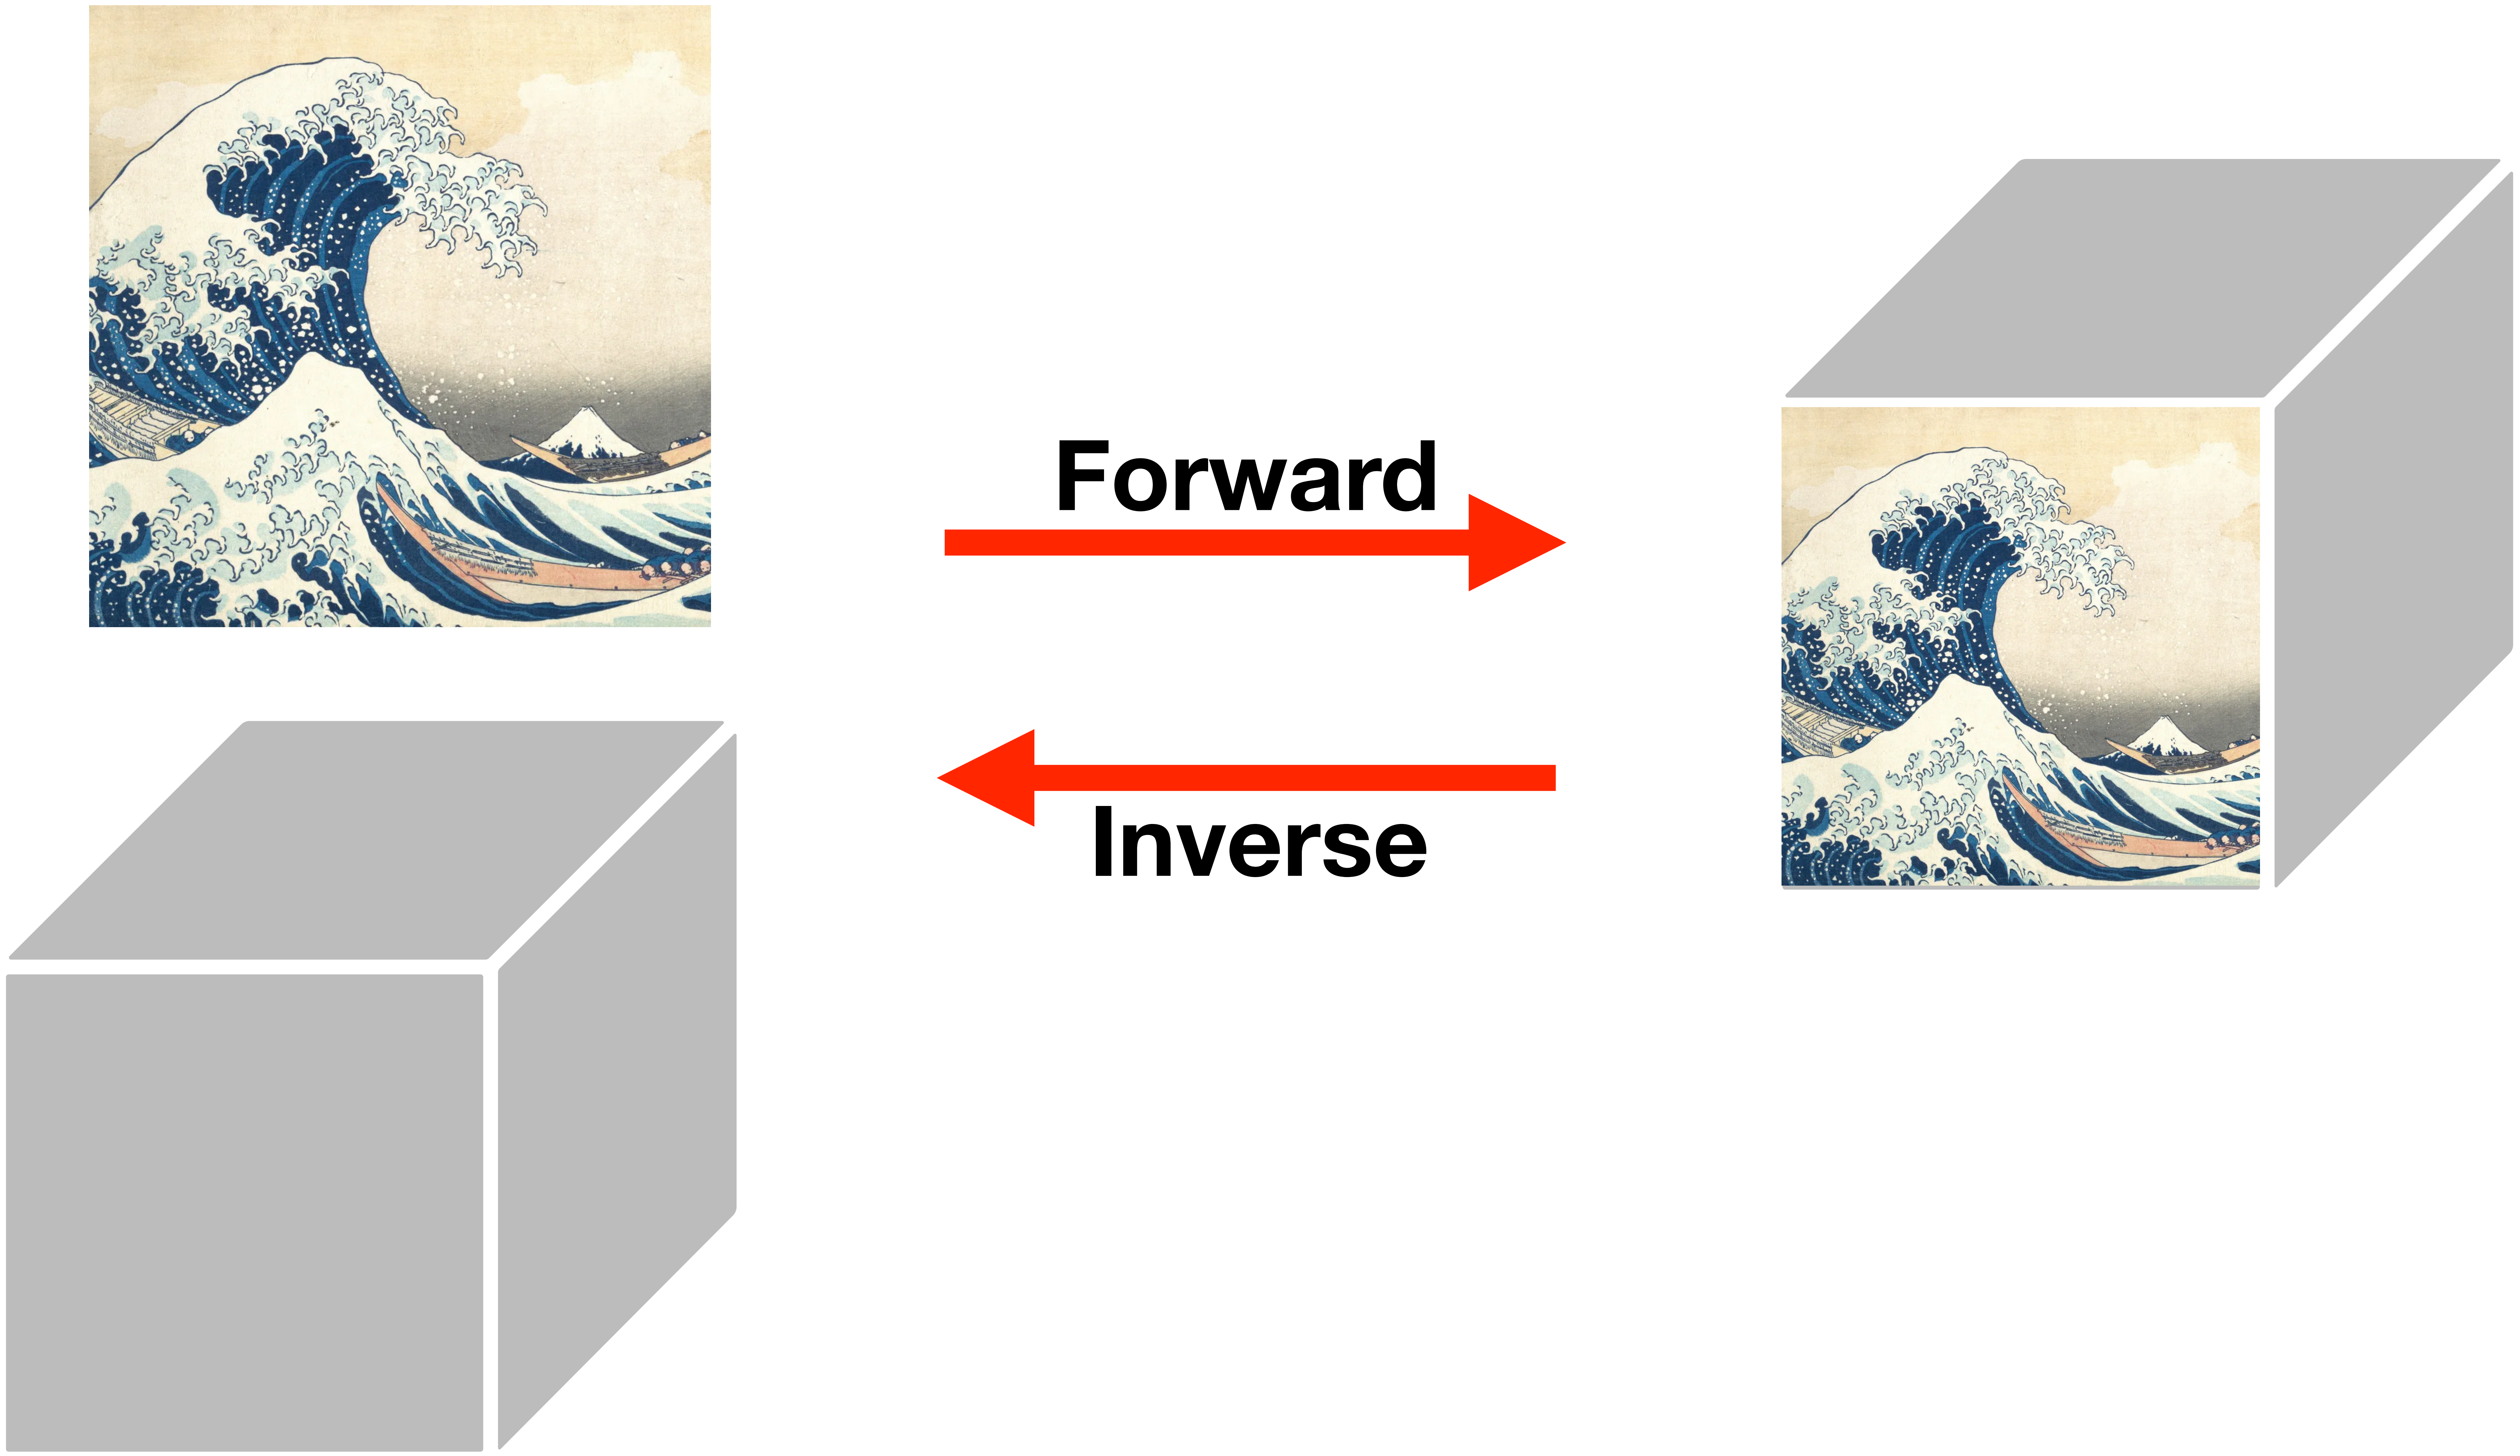

Optimization: Slow and Noisy

Optimization: Slow and Noisy

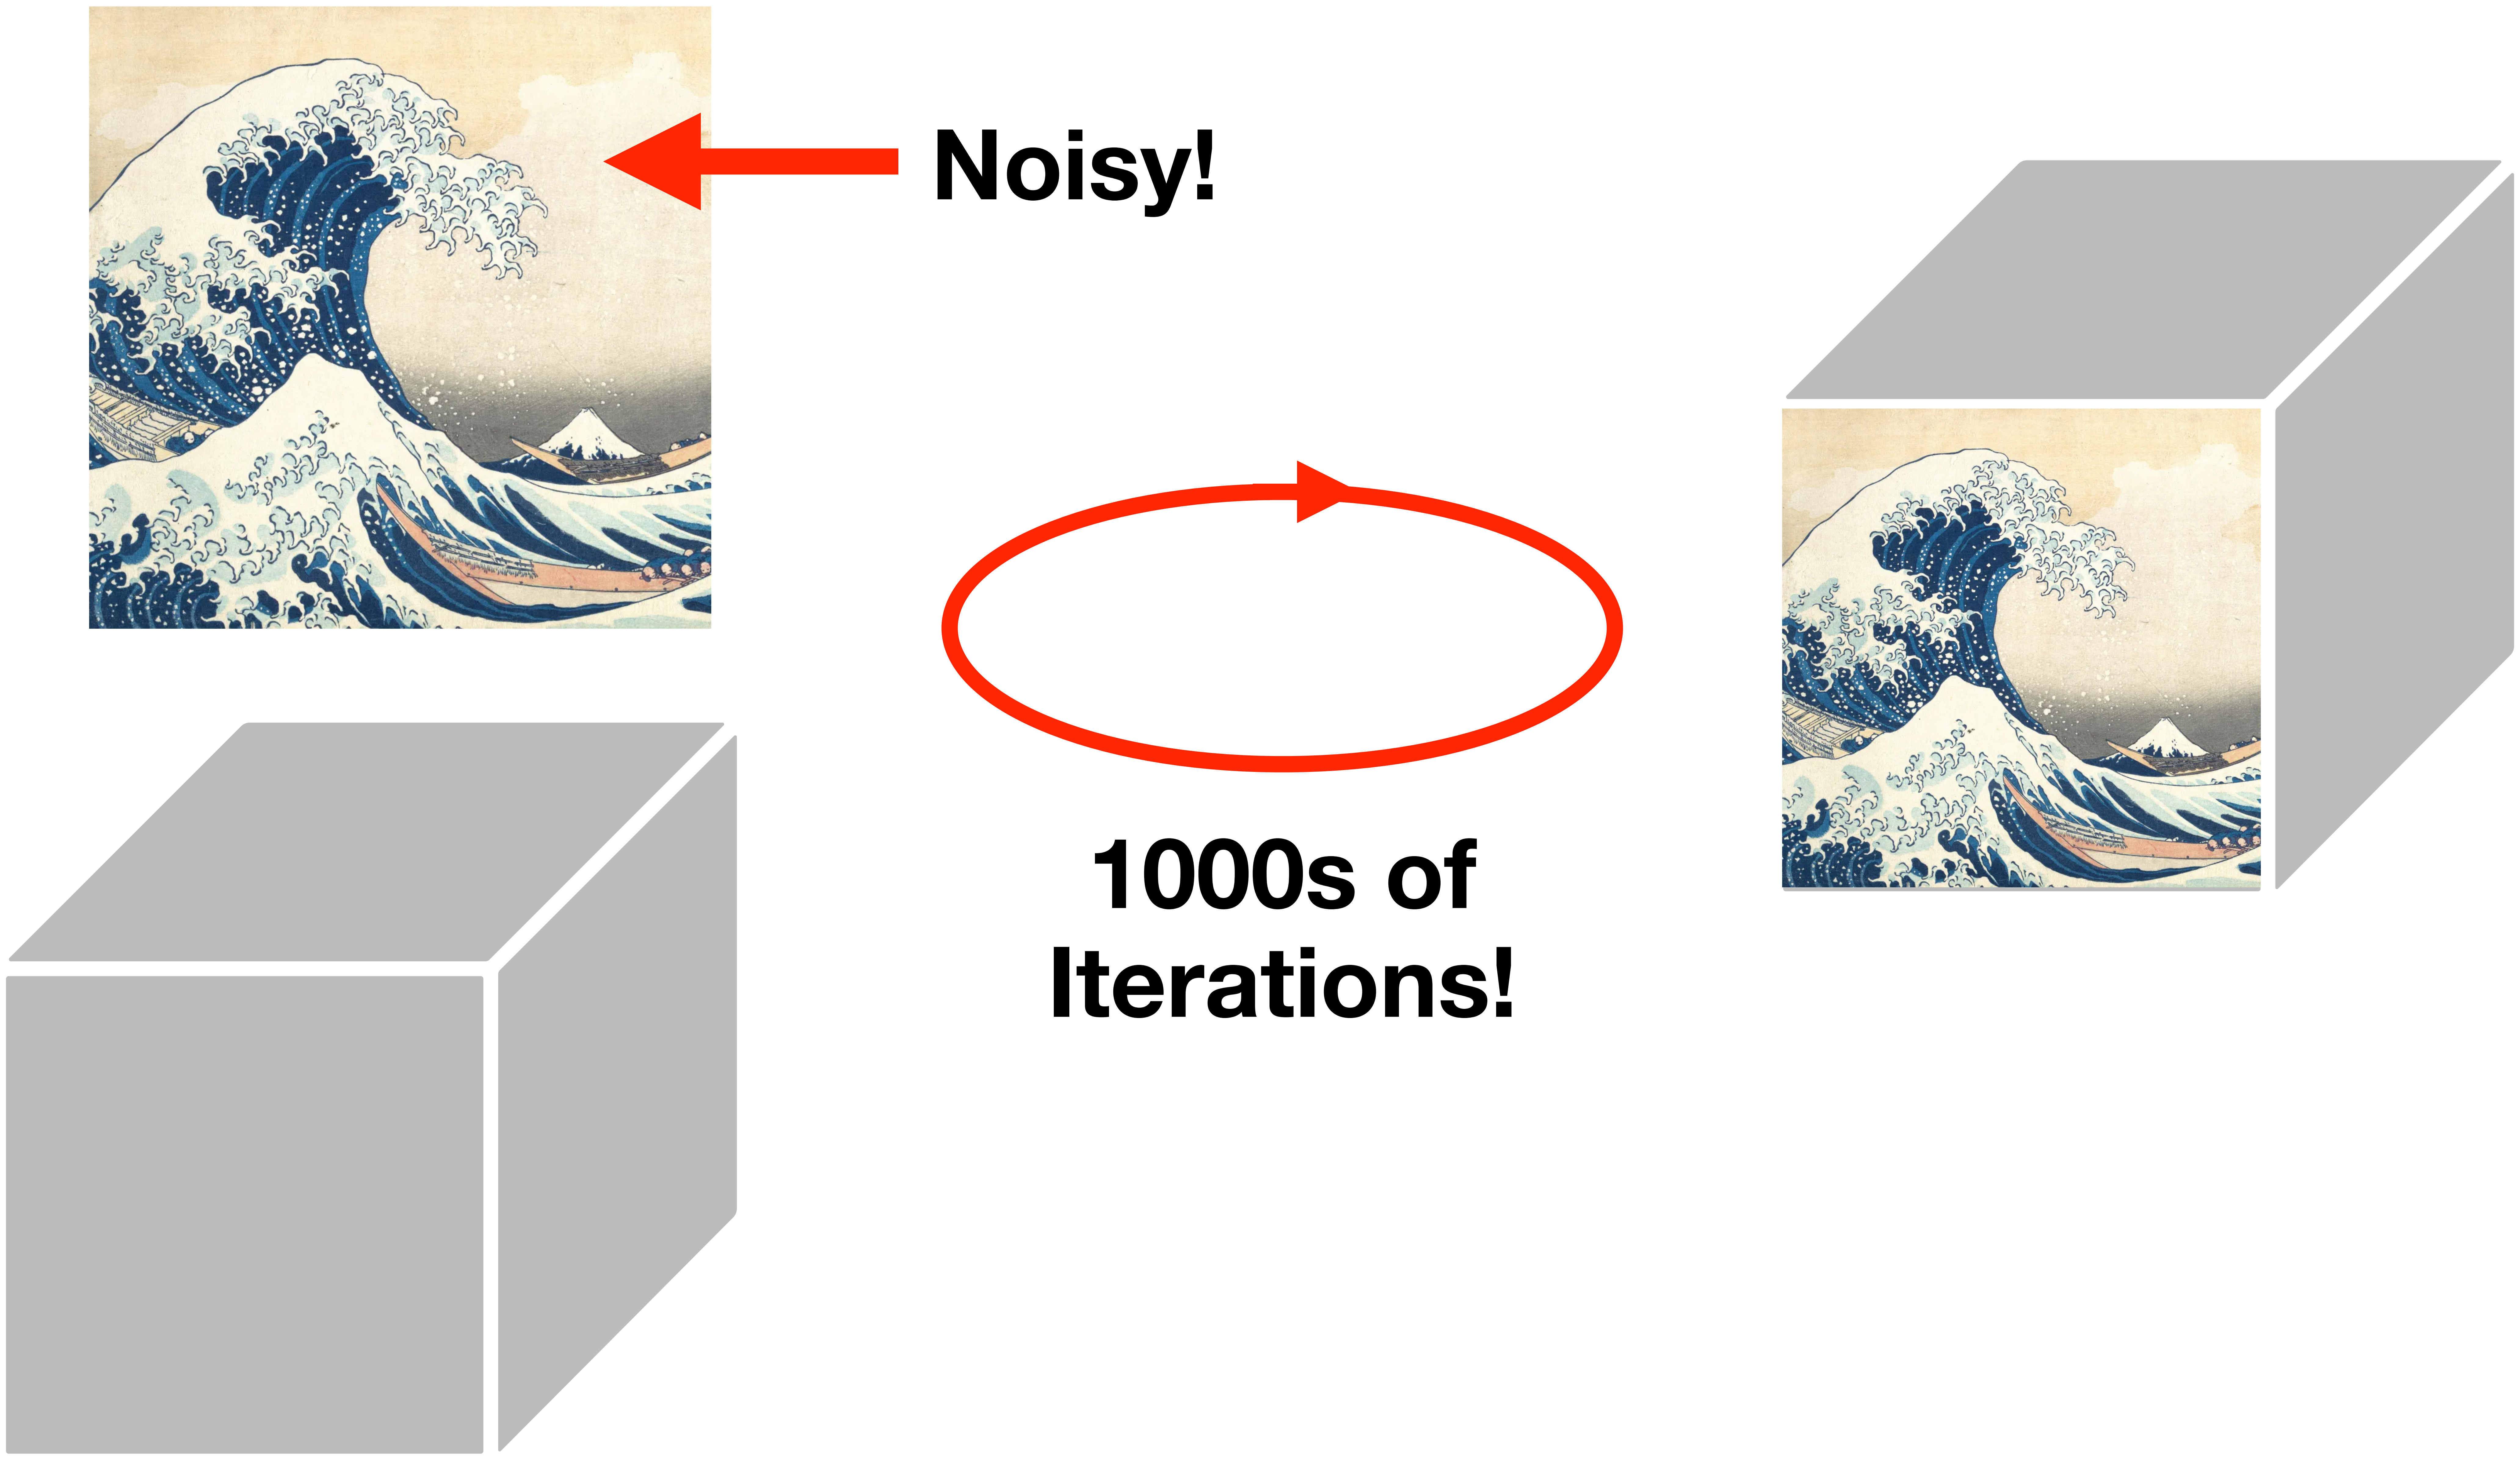

Our Goal

Improve inverse rendering with a new gradient-based optimizer

- **Faster** optimizations: fewer samples, fewer iterations
- **Better** recovery quality: deal with noise

Spatial signals are often piecewise smooth

Spatial signals are often piecewise smooth

smooth
regions

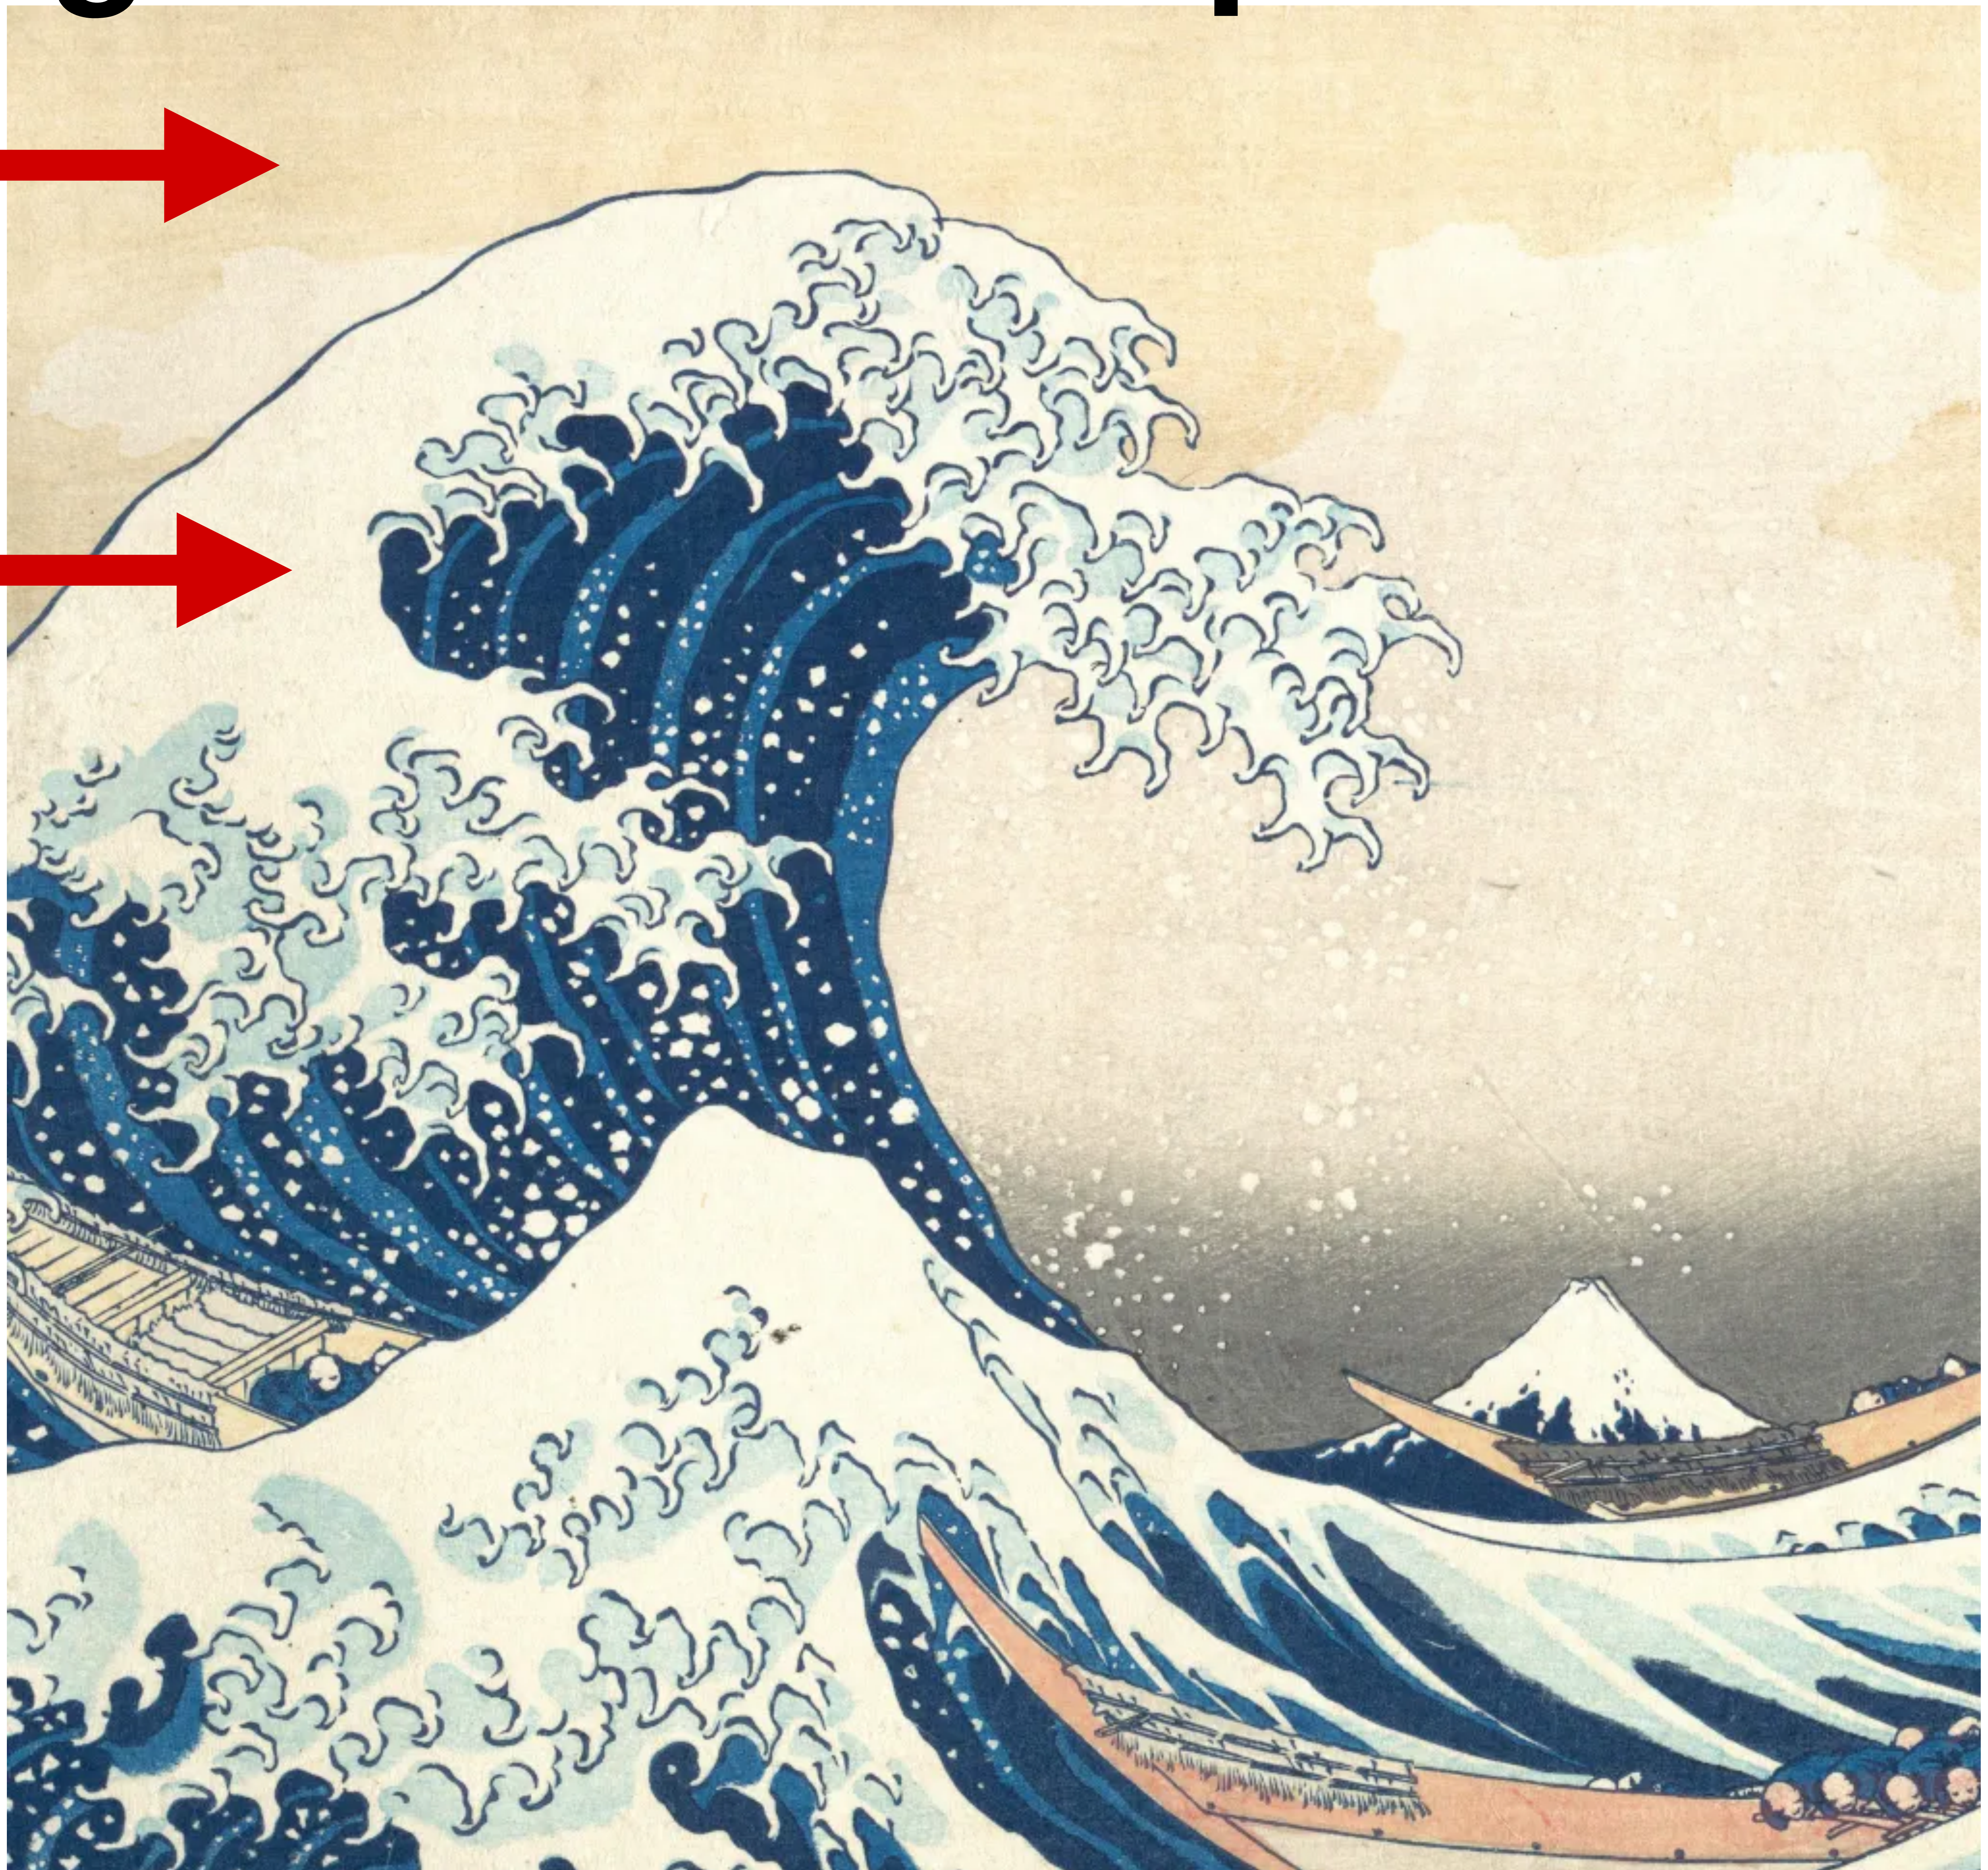

Spatial signals are often piecewise smooth

smooth
regions

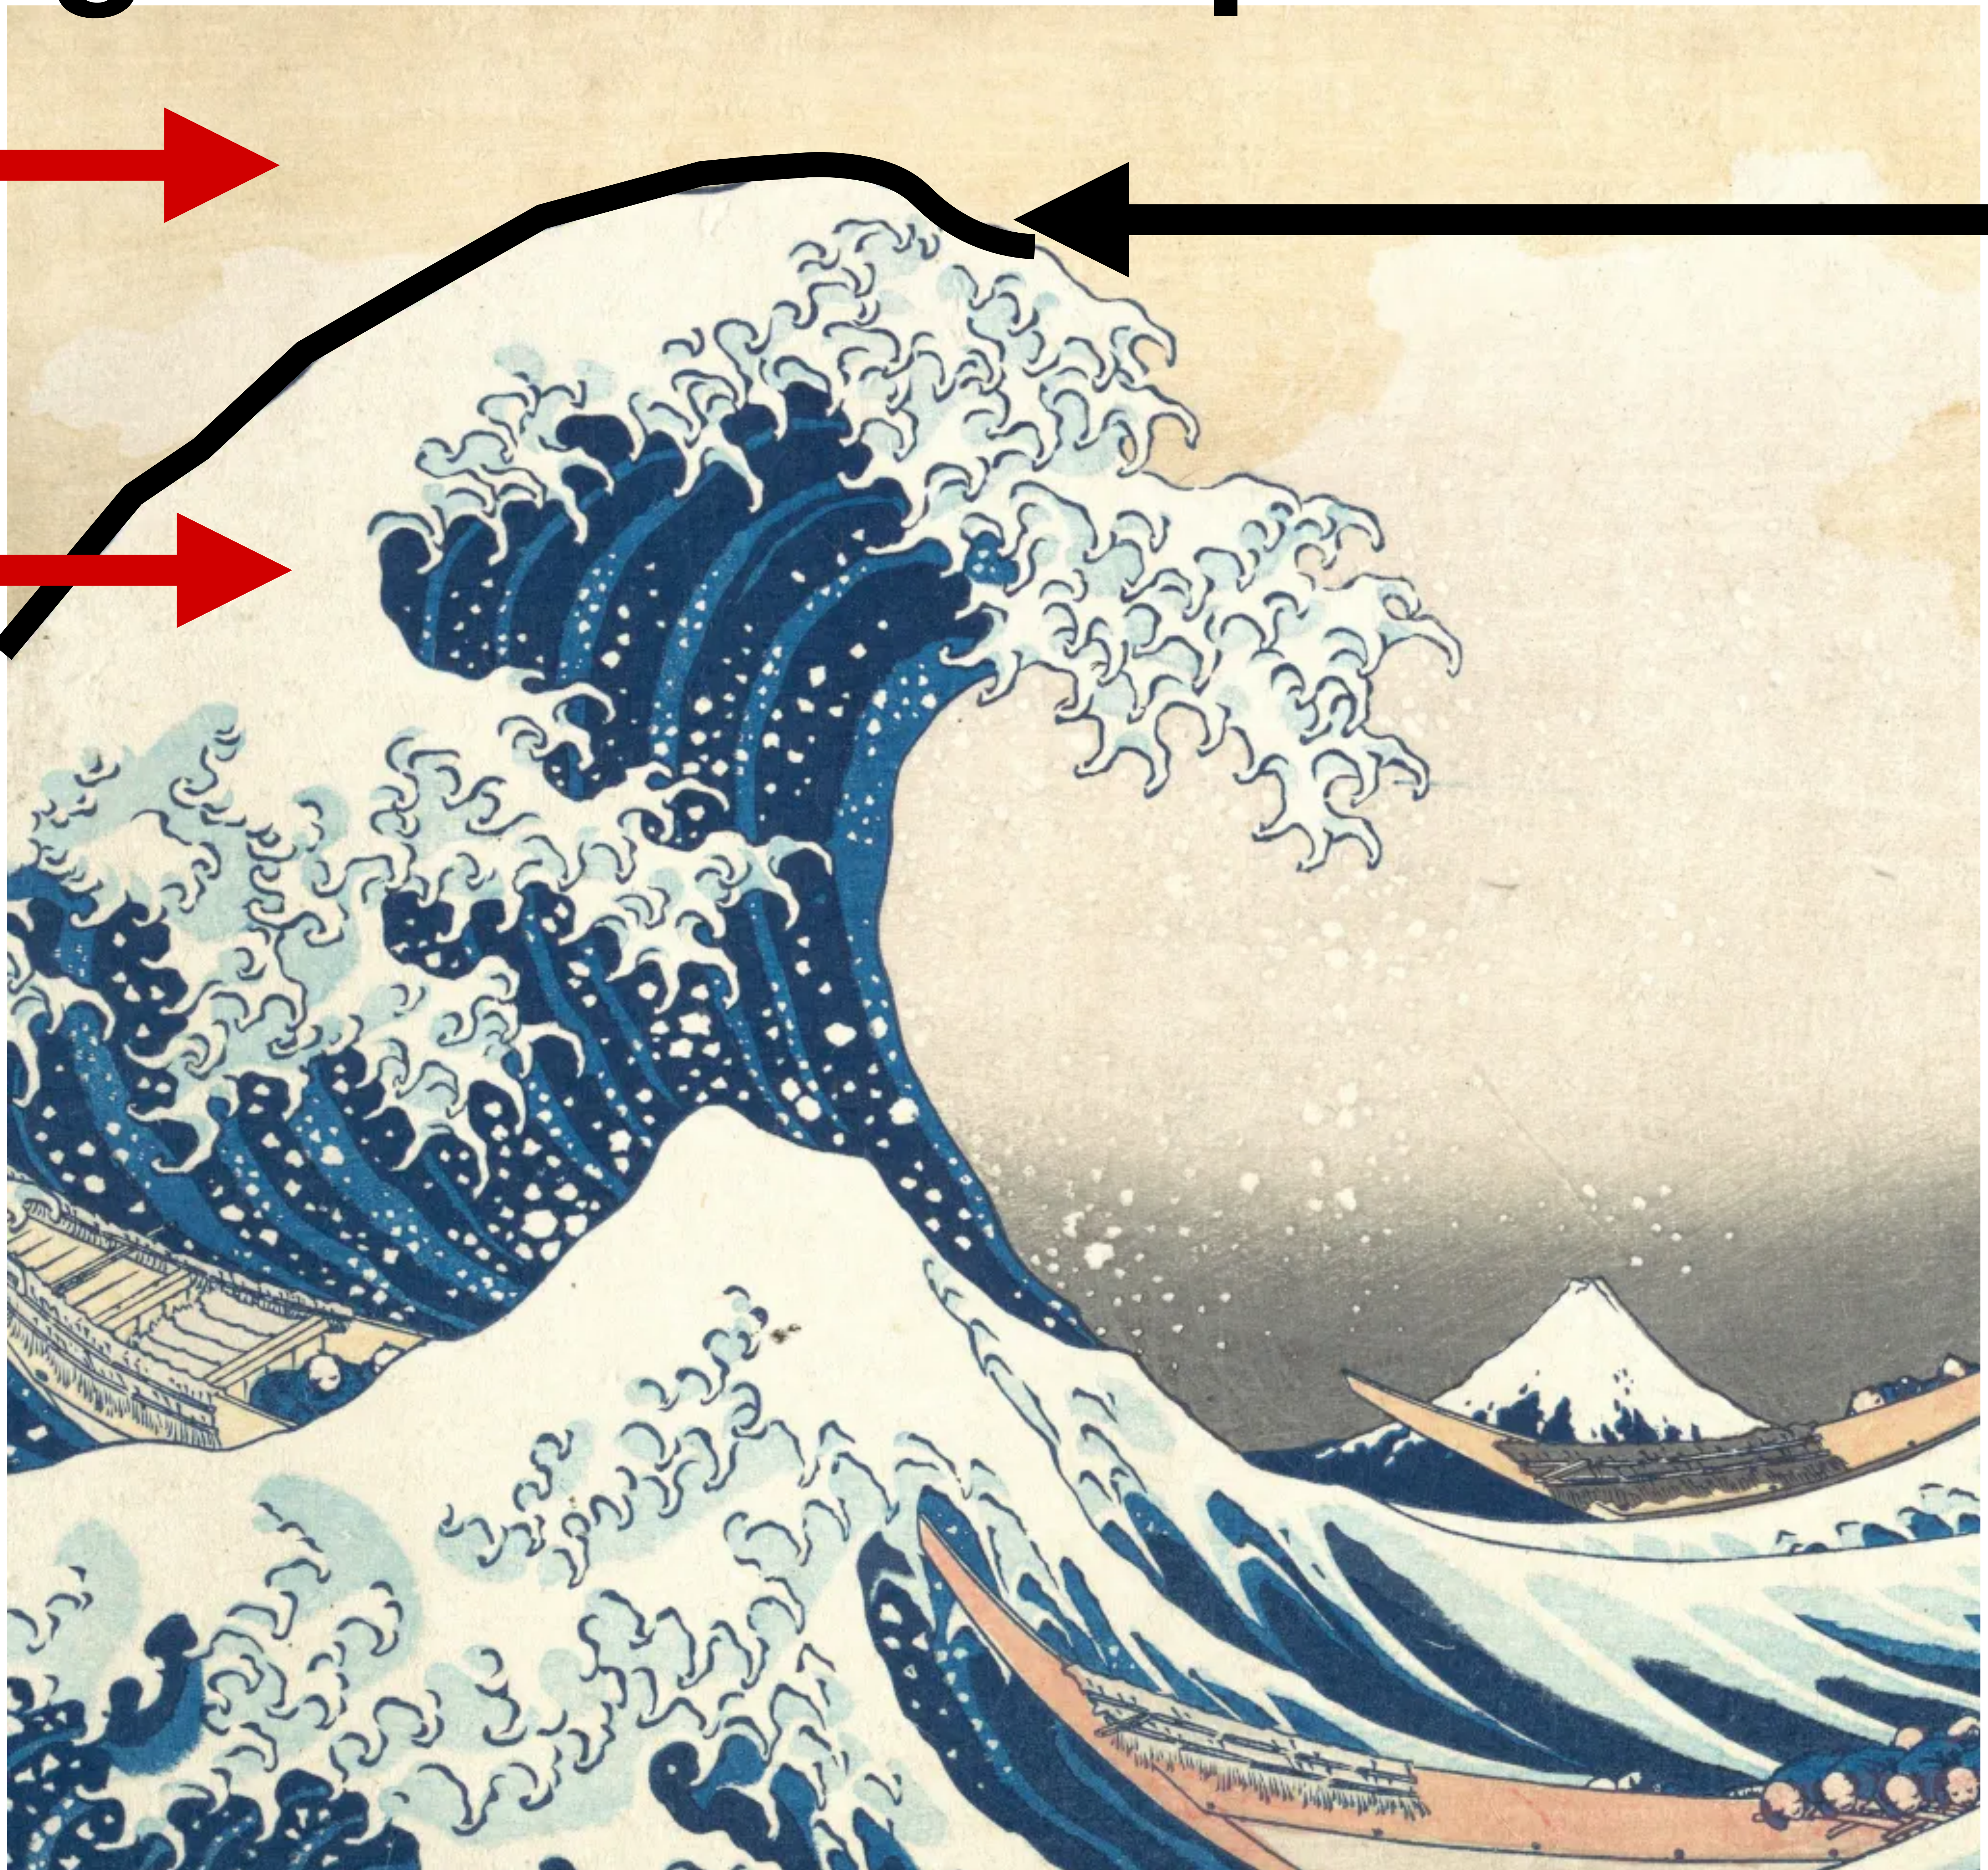

Edge

Spatial signals are often piecewise smooth

Spatial signals are often piecewise smooth

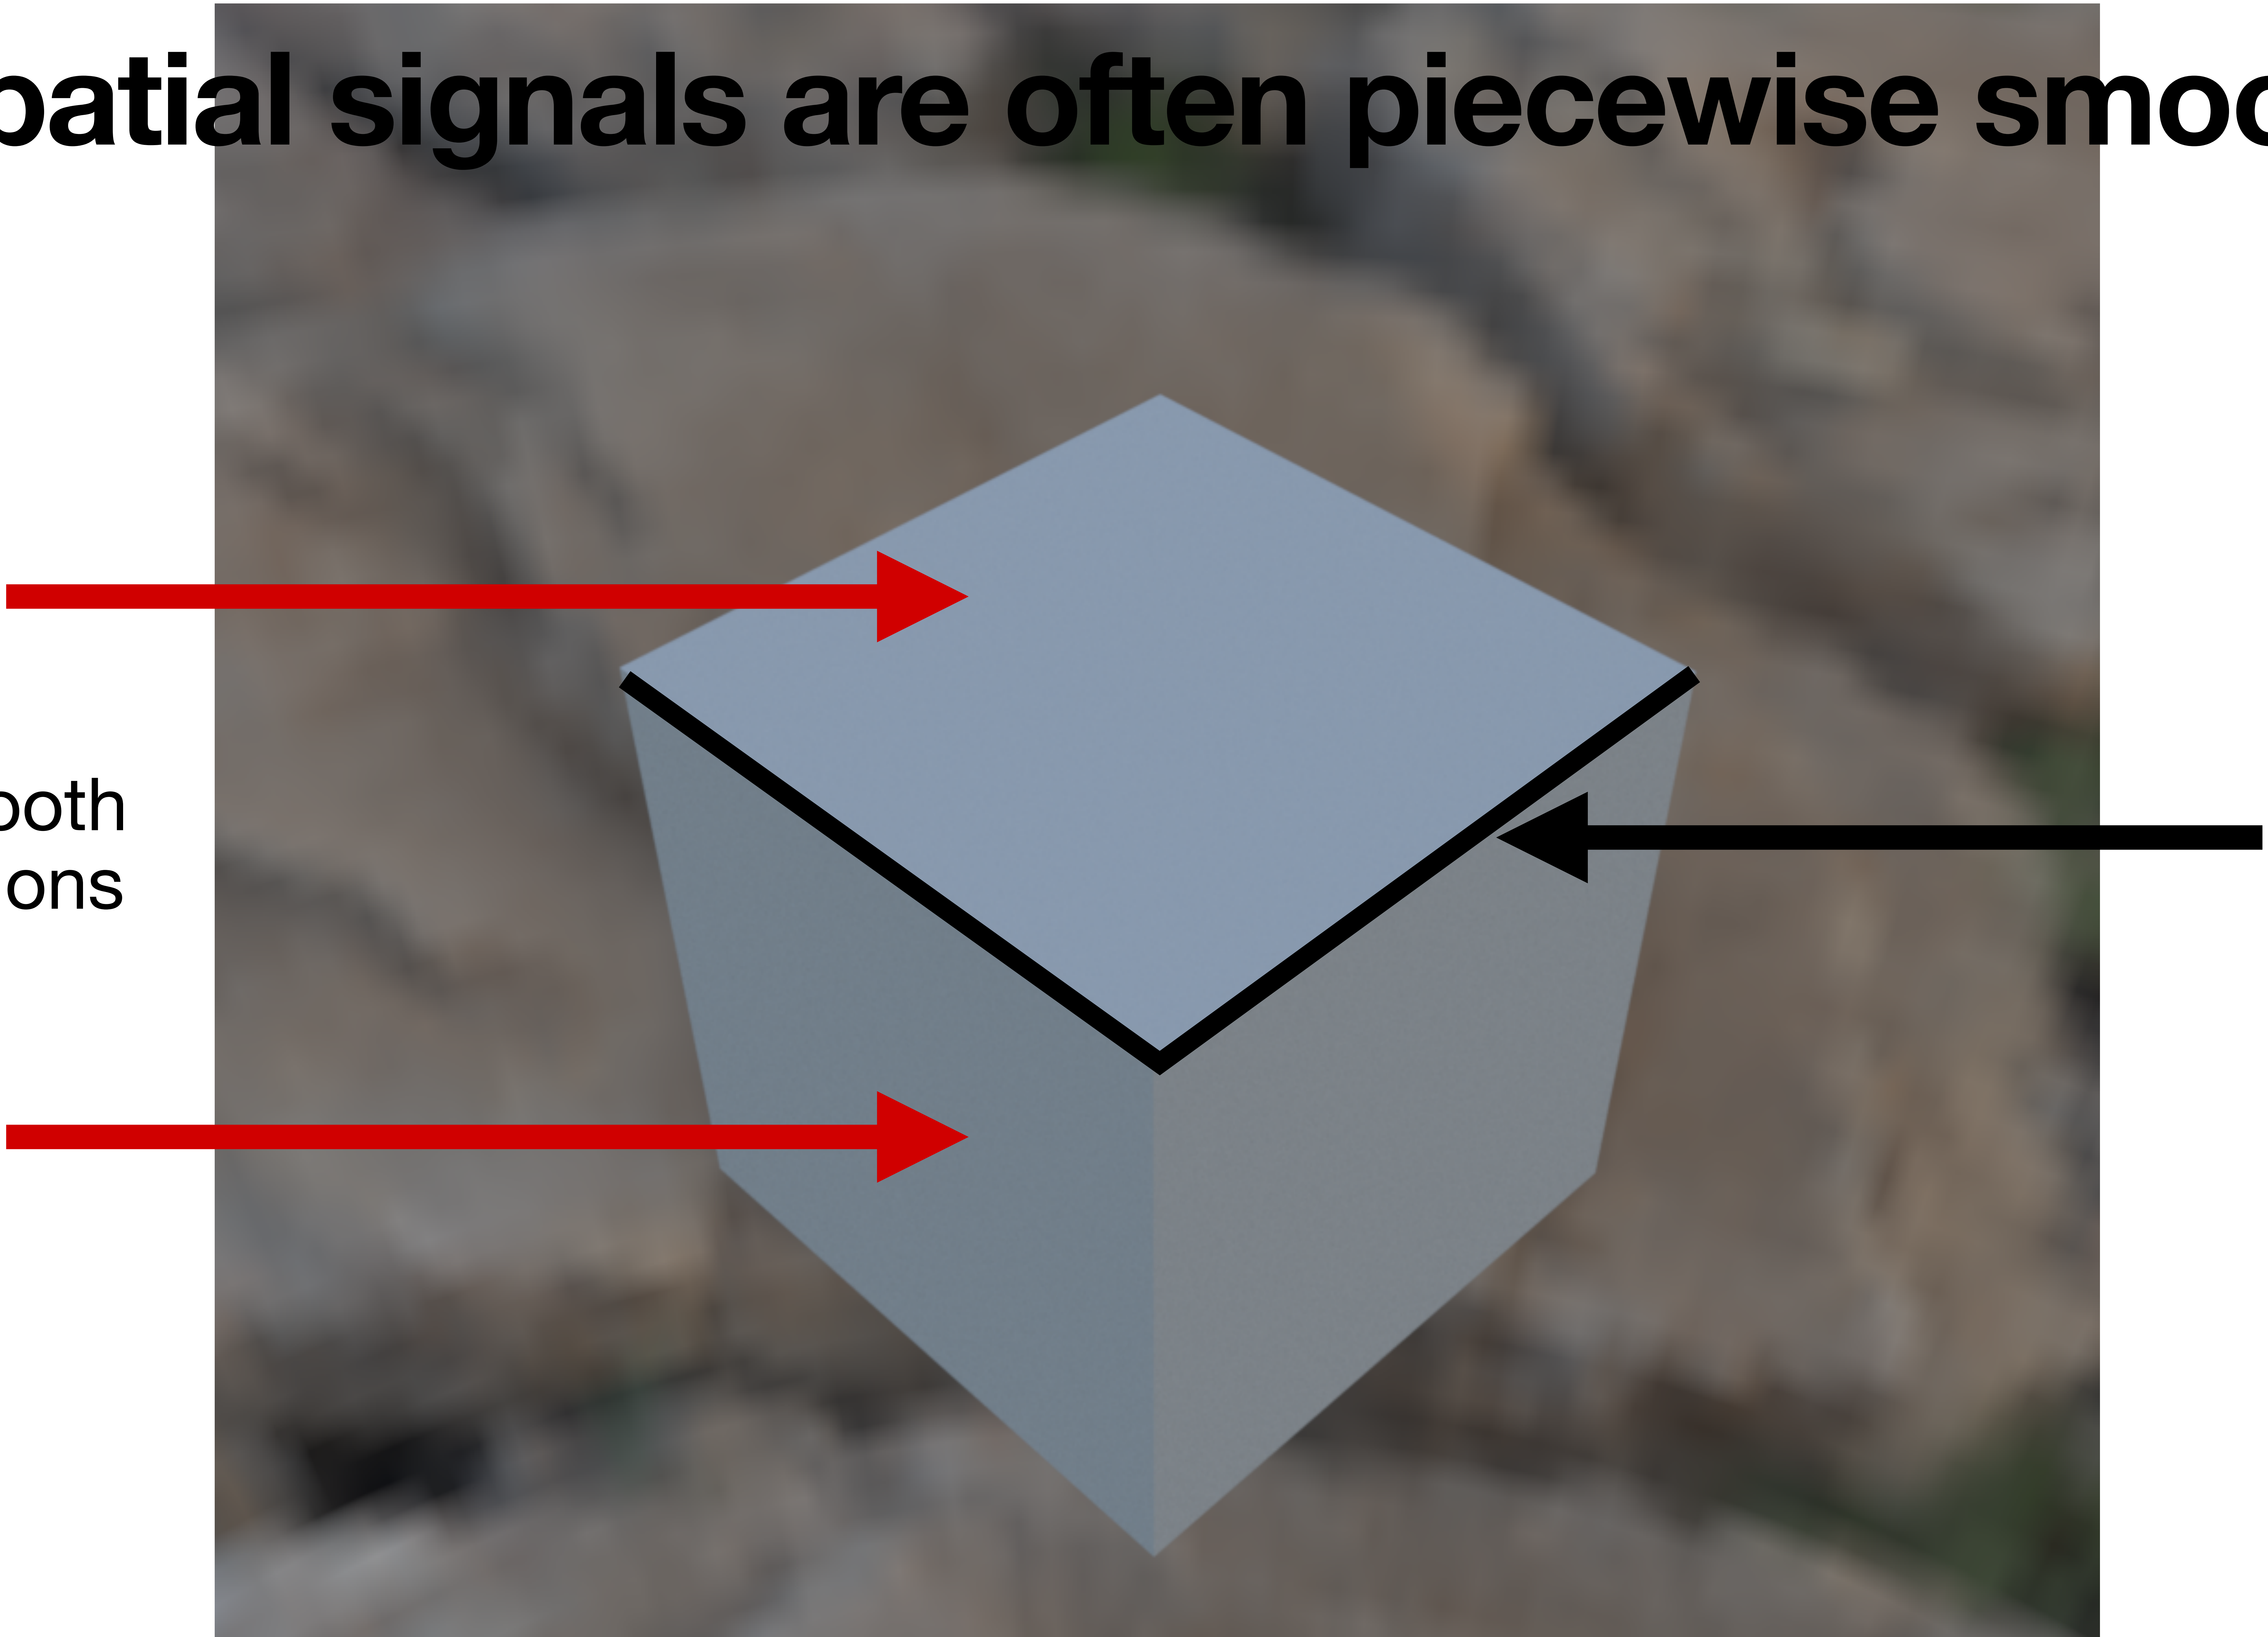

Can we leverage this piecewise smoothness?

Gradient descent struggles with noisy gradients

Let's try to optimize a 1D example:

Initial

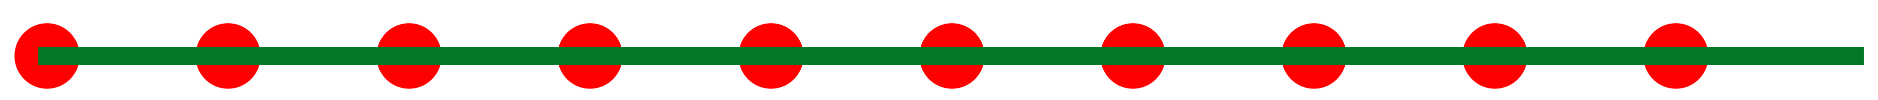

Reference

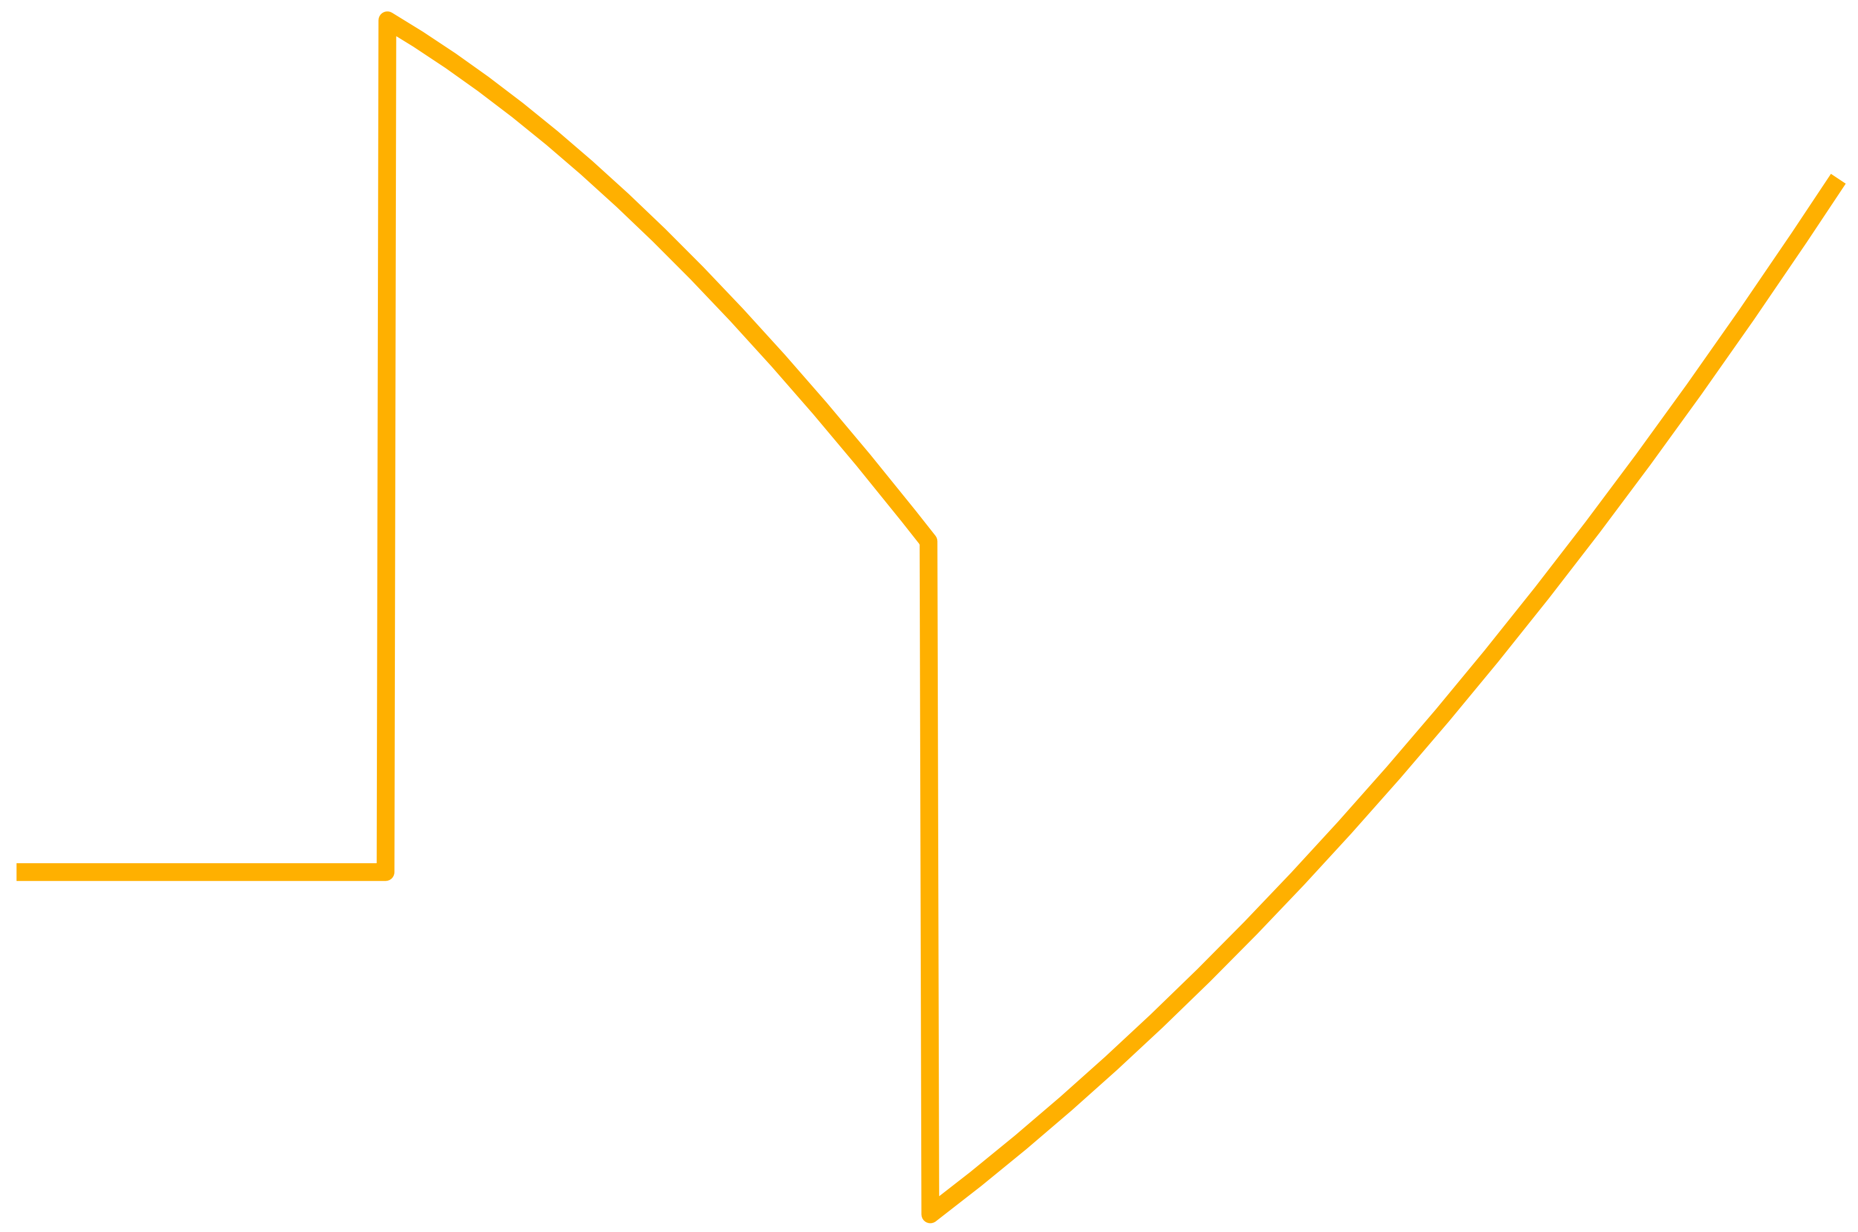

Gradient descent struggles with noisy gradients

Without noise, we get clean and accurate gradients

Noise-Free Gradients

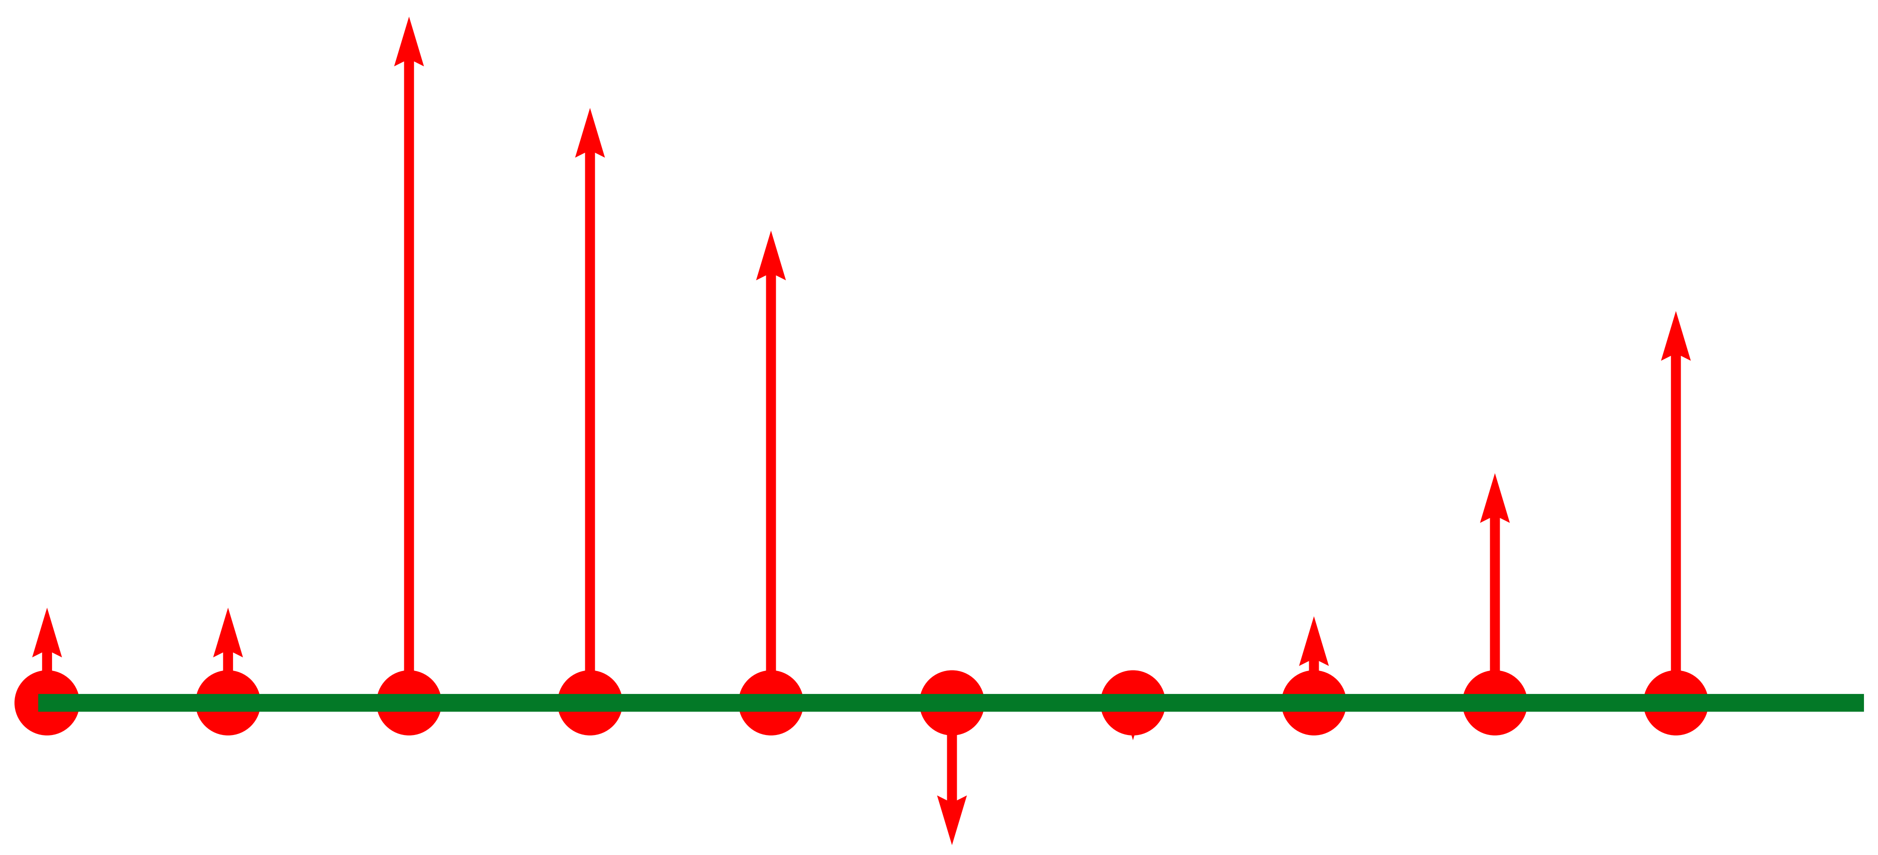

Reference

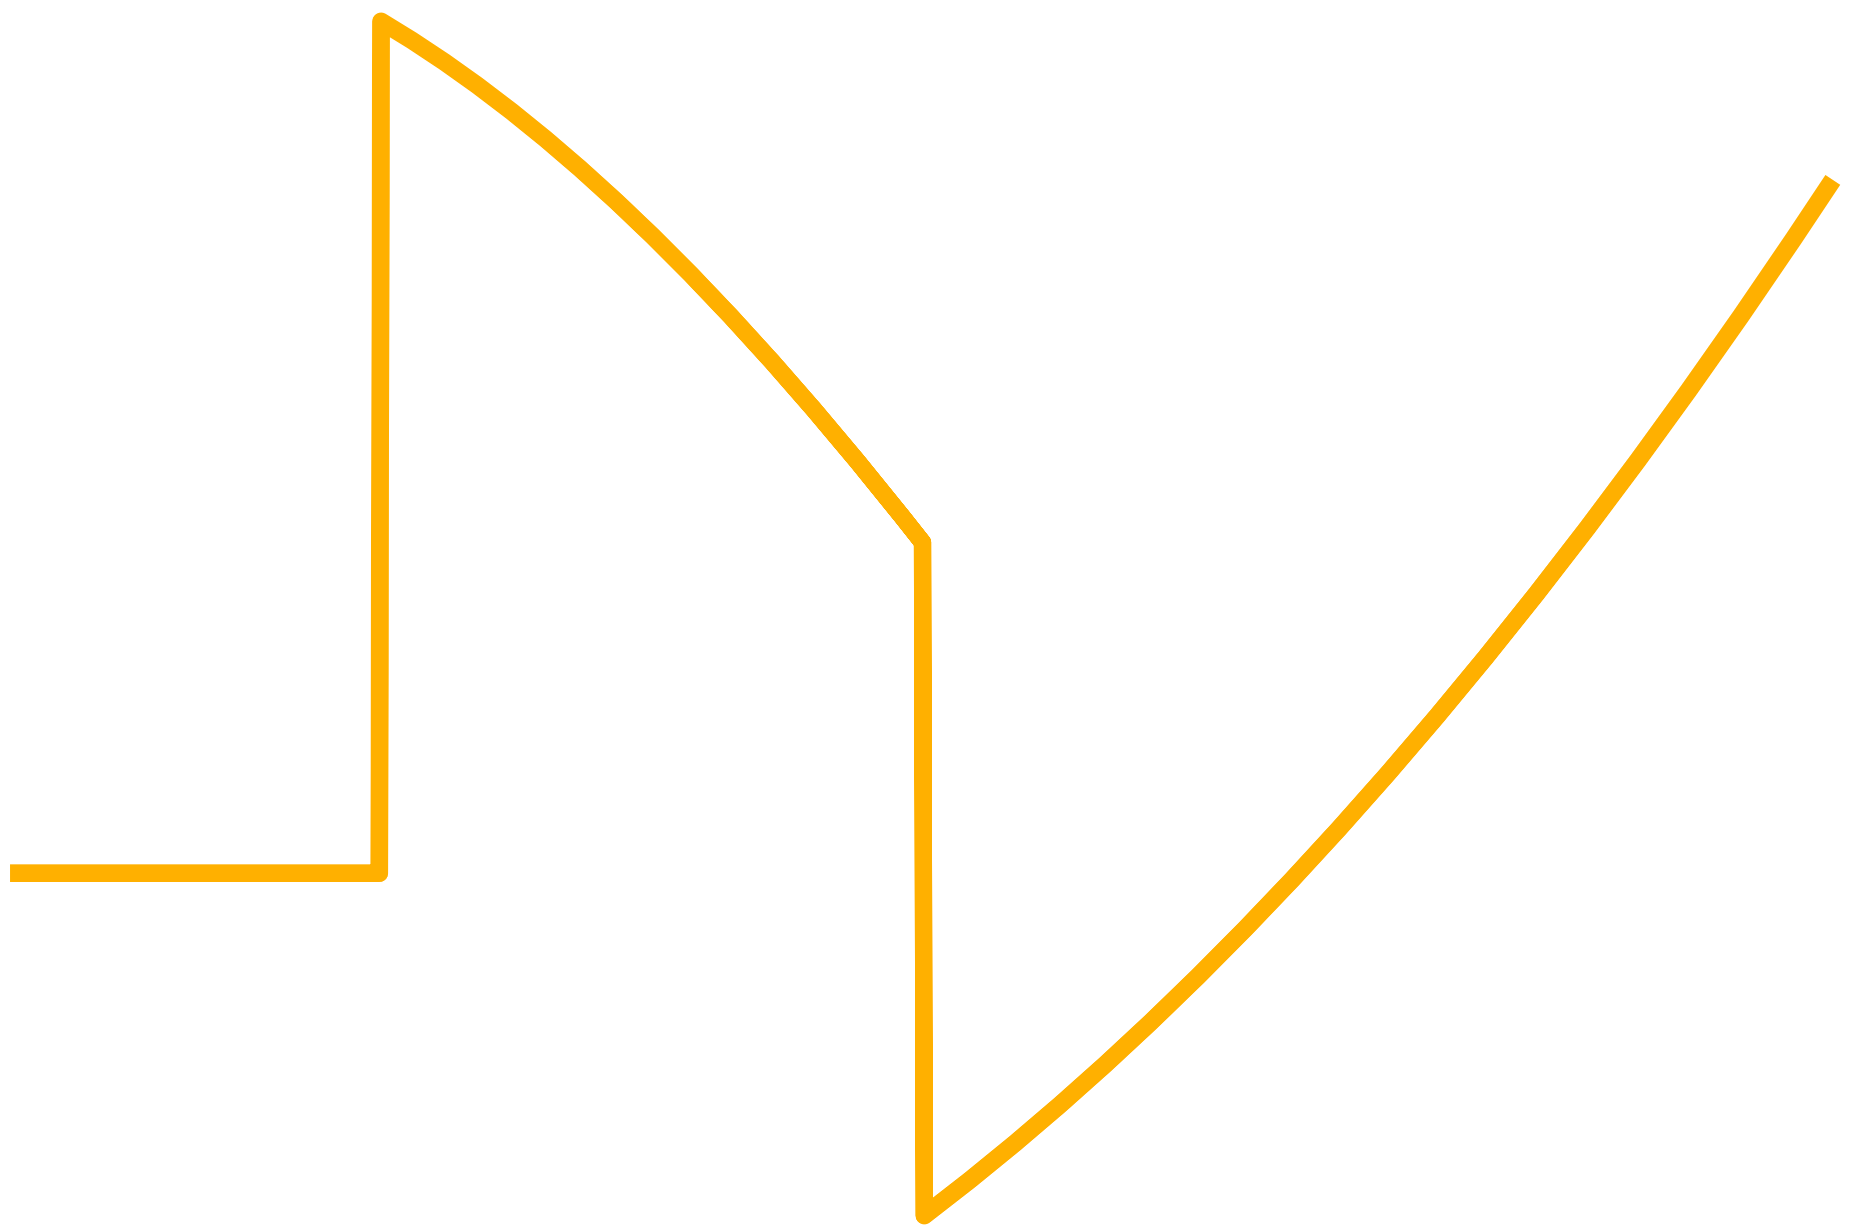

Gradient descent struggles with noisy gradients

Noise-free gradients lead to fast and high-quality optimization

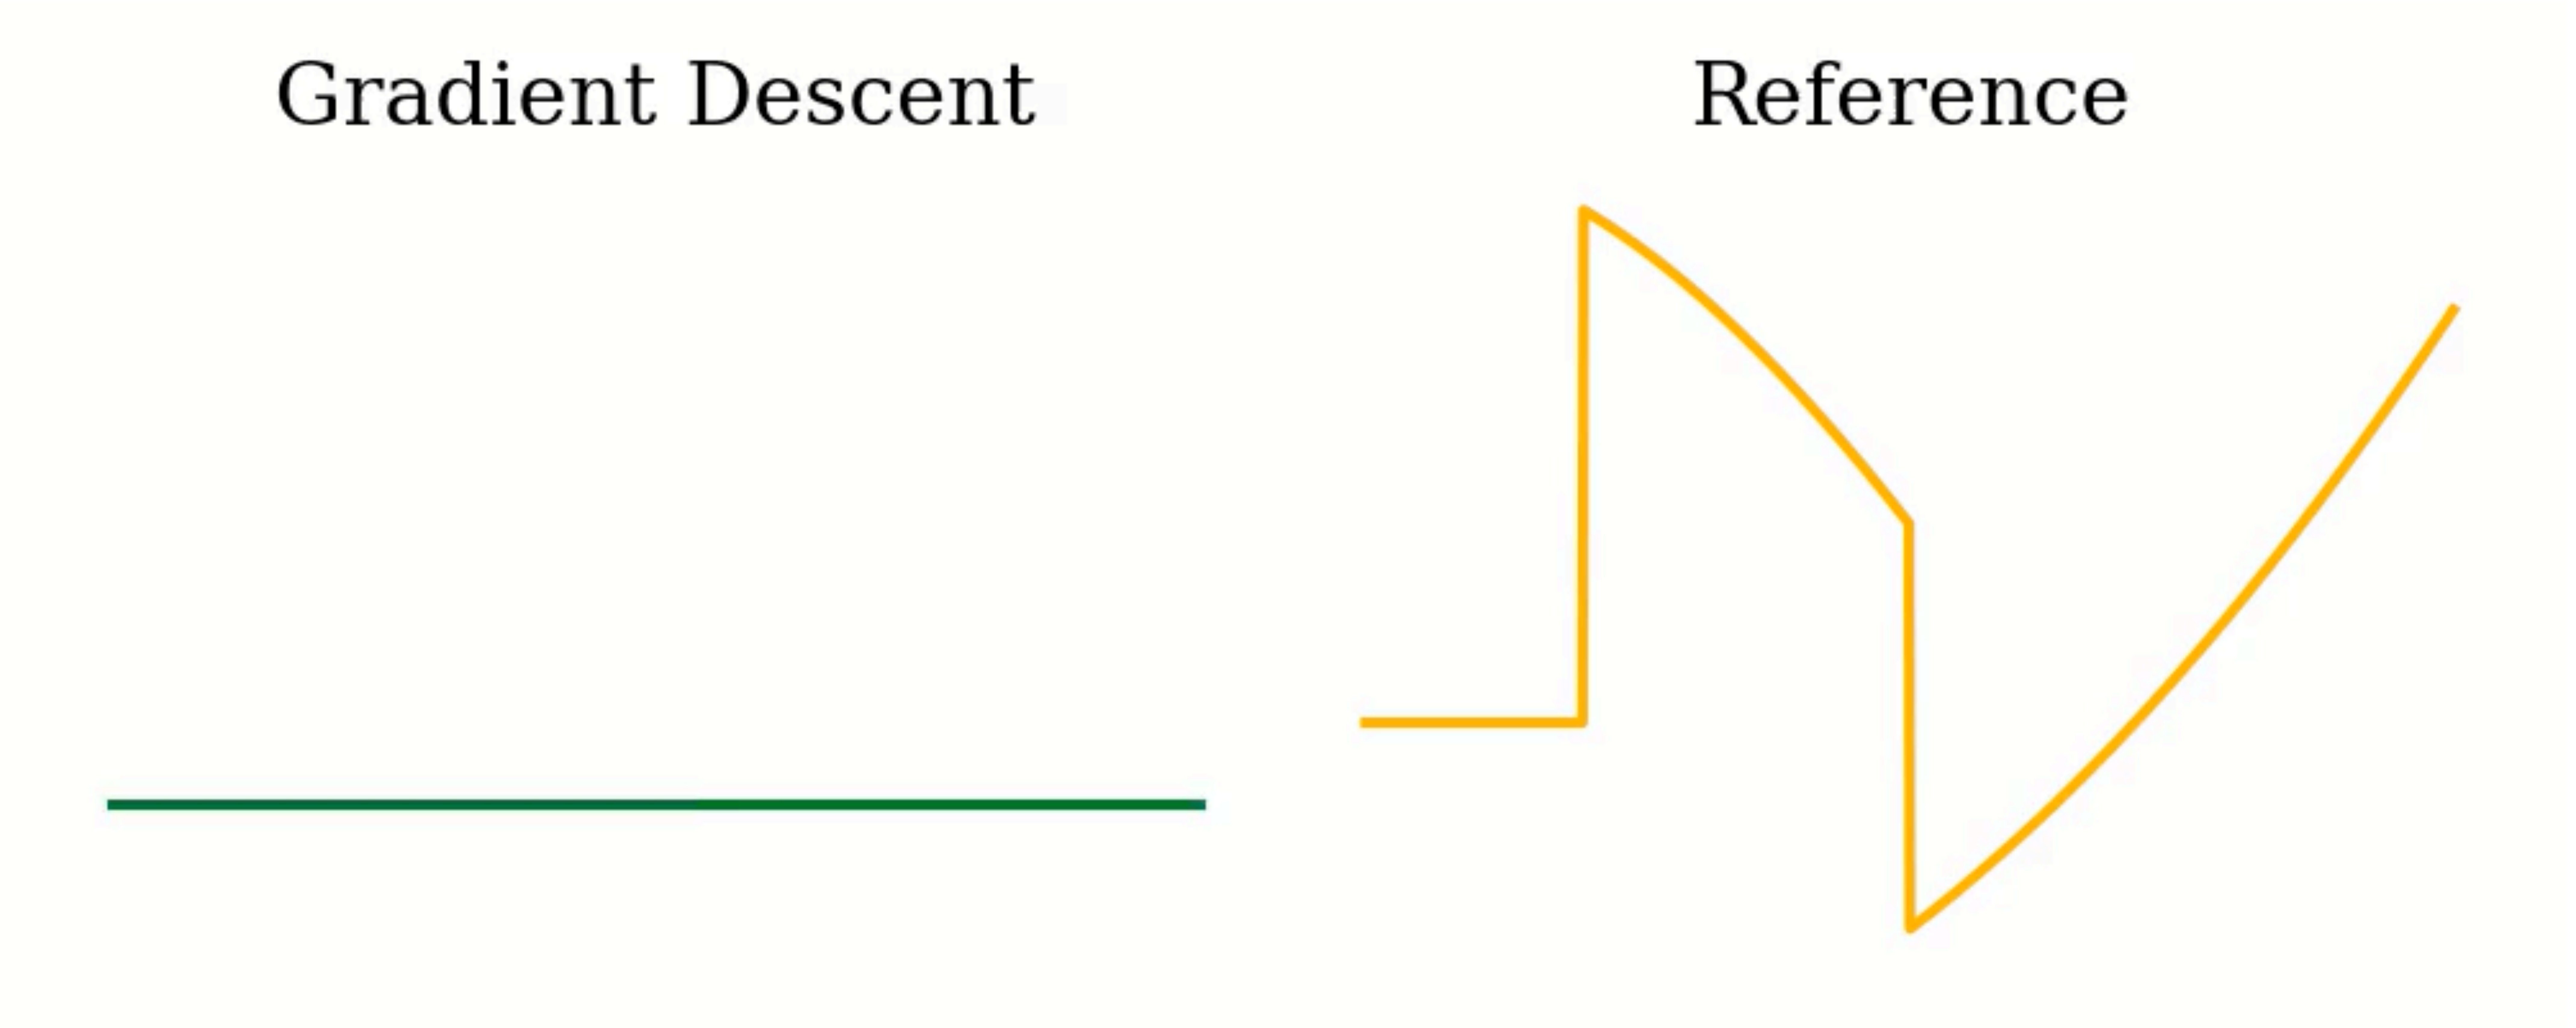

Gradient descent struggles with noisy gradients

Now add some noise to the gradients:

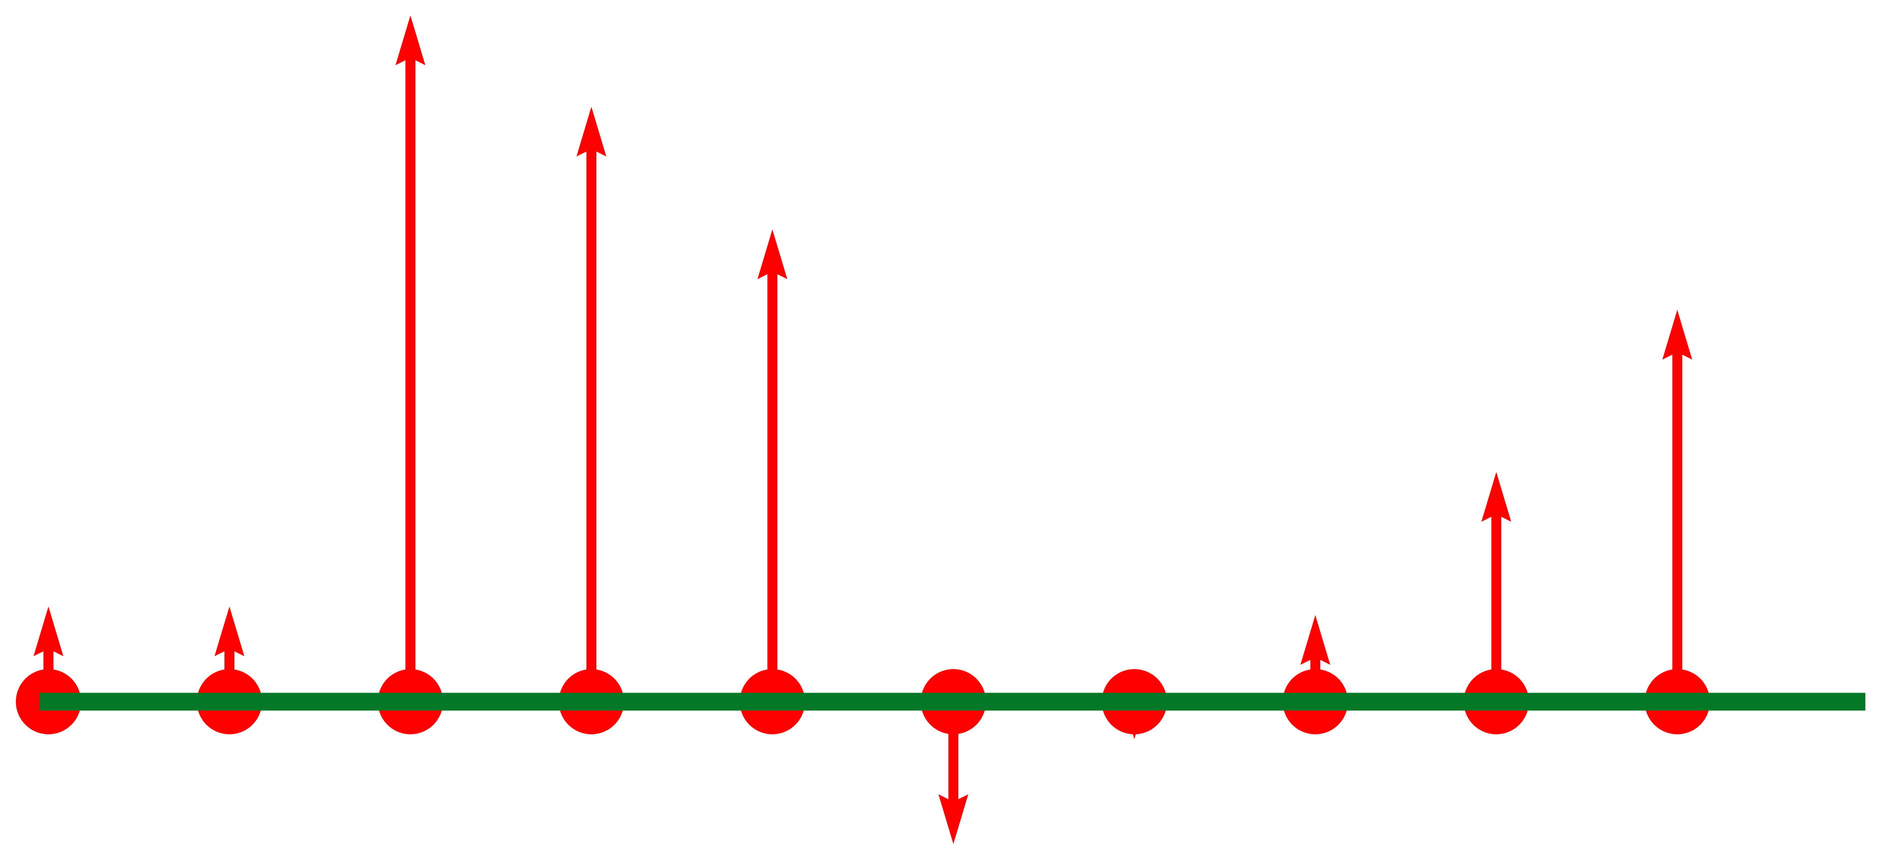

Gradient descent struggles with noisy gradients

Now add some noise to the gradients:

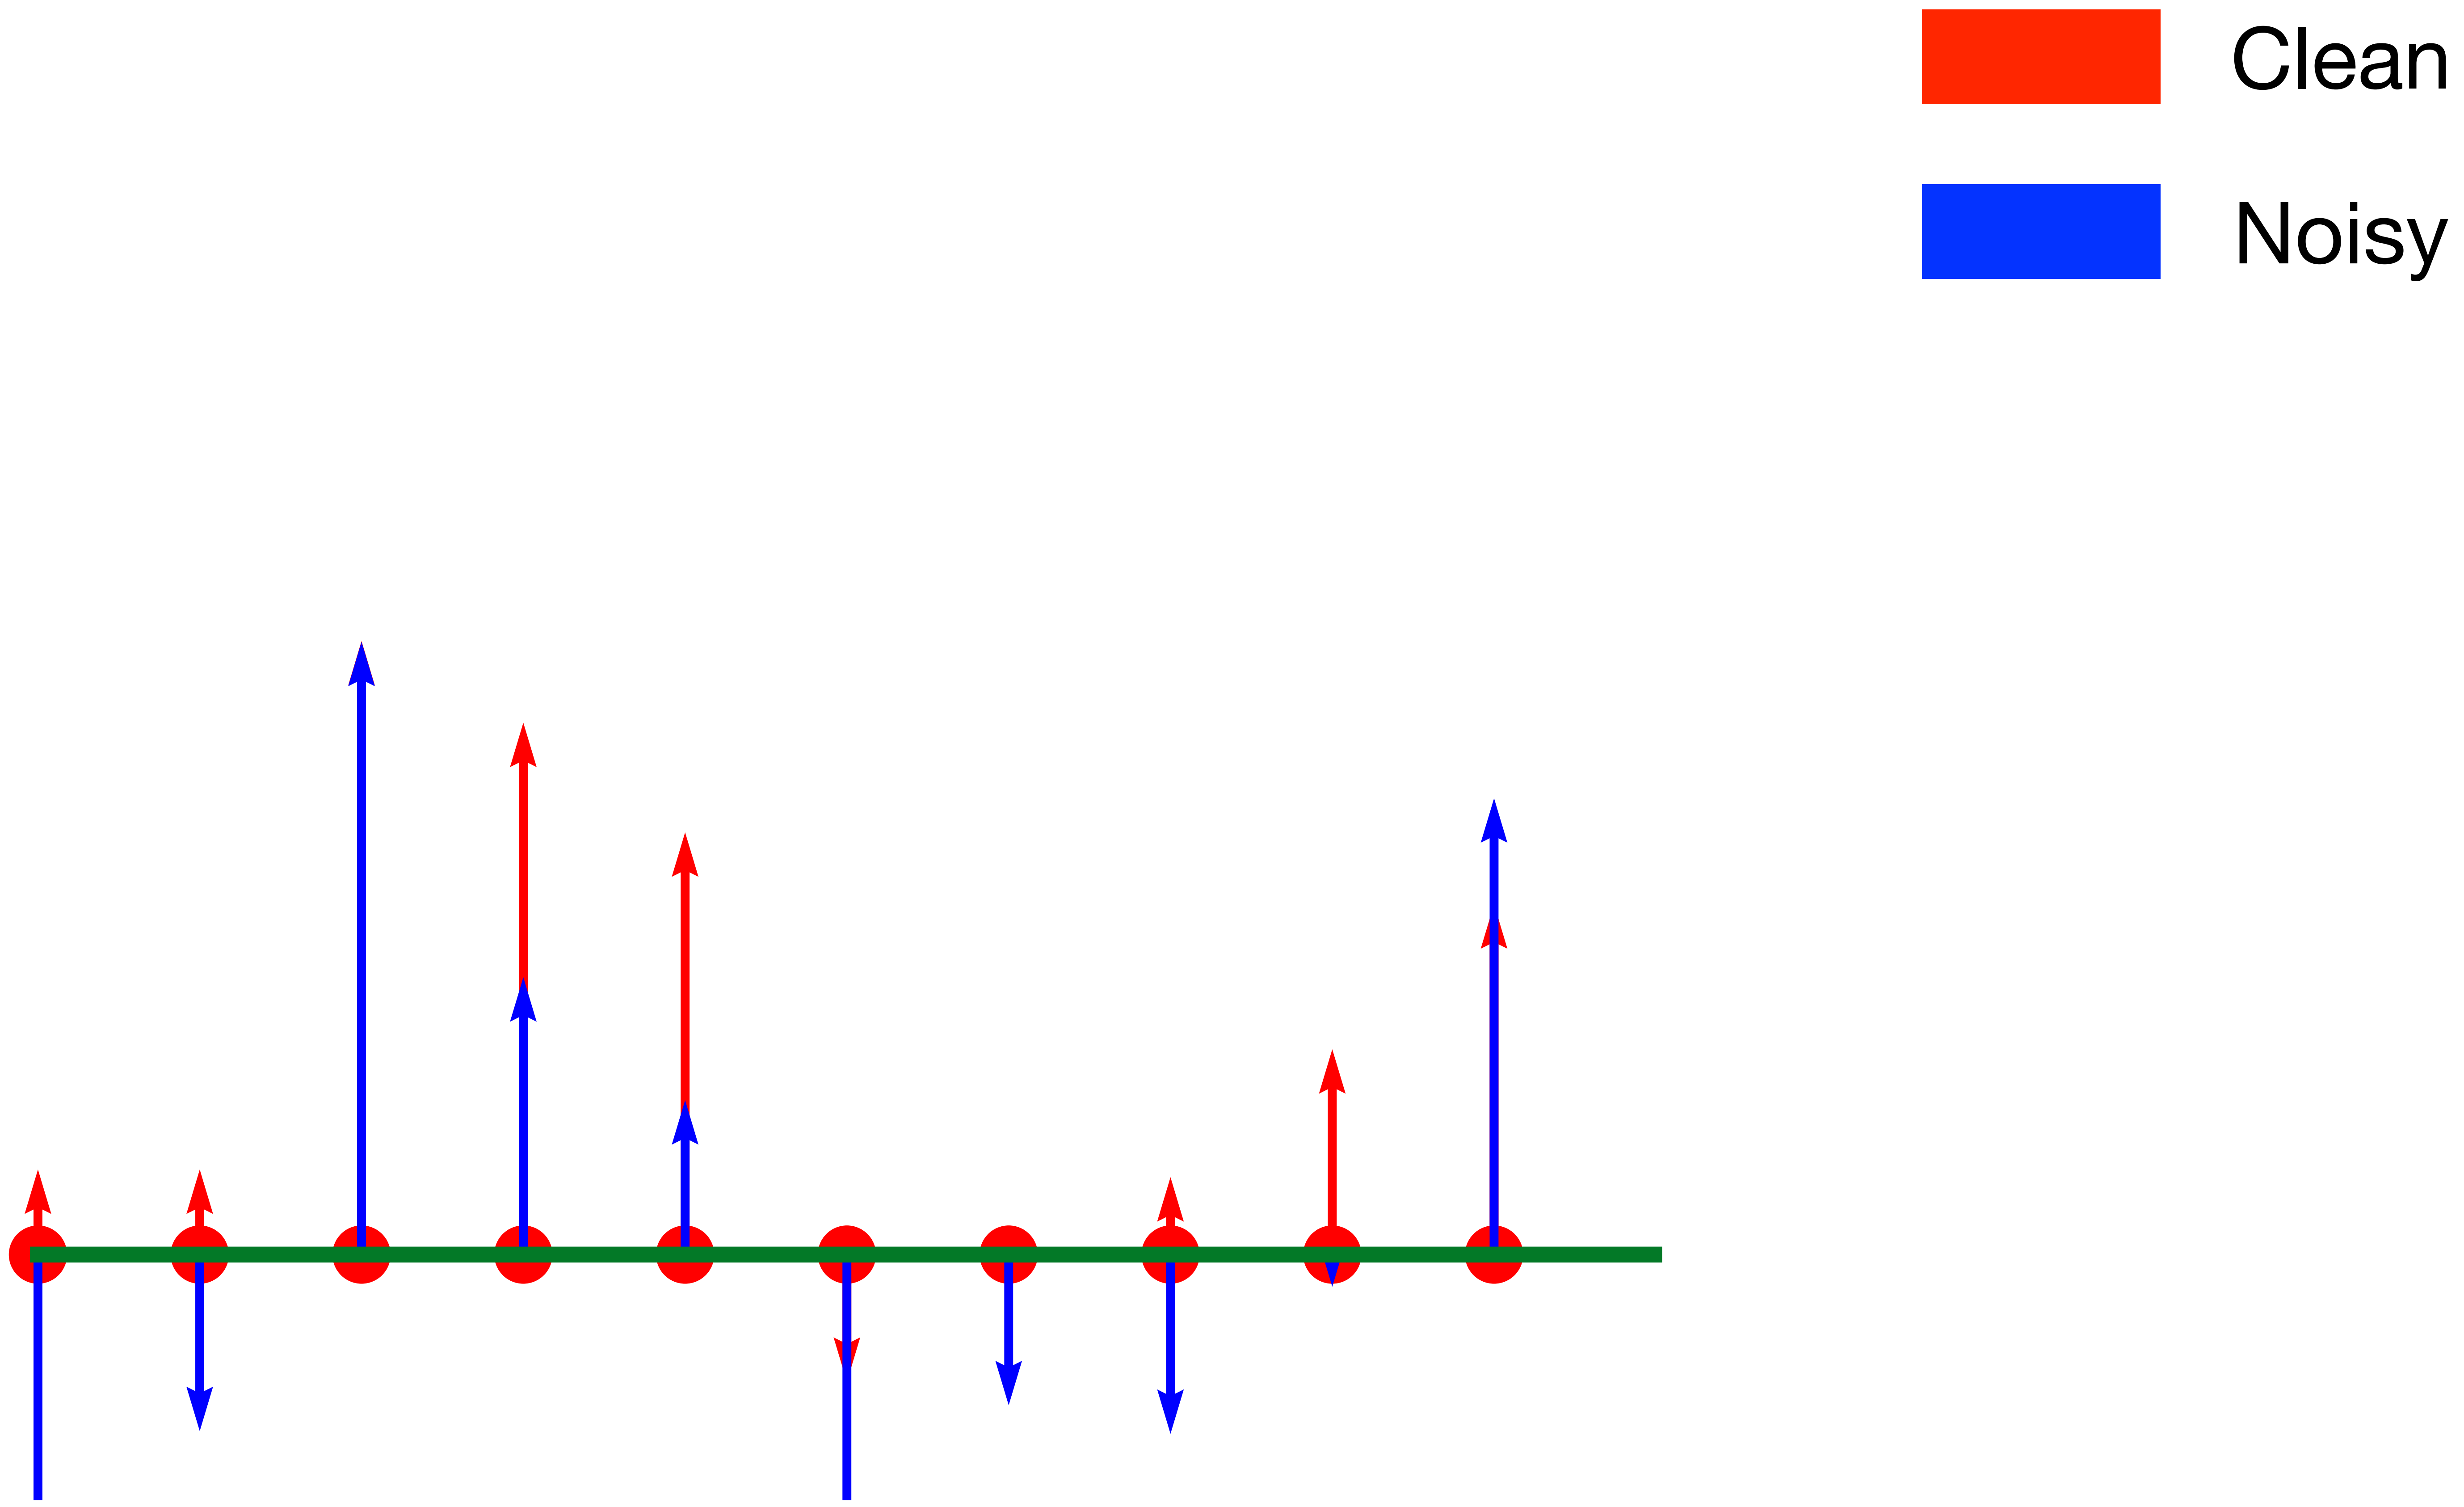

Gradient descent struggles with noisy gradients

Noisy gradients lead to poor convergence

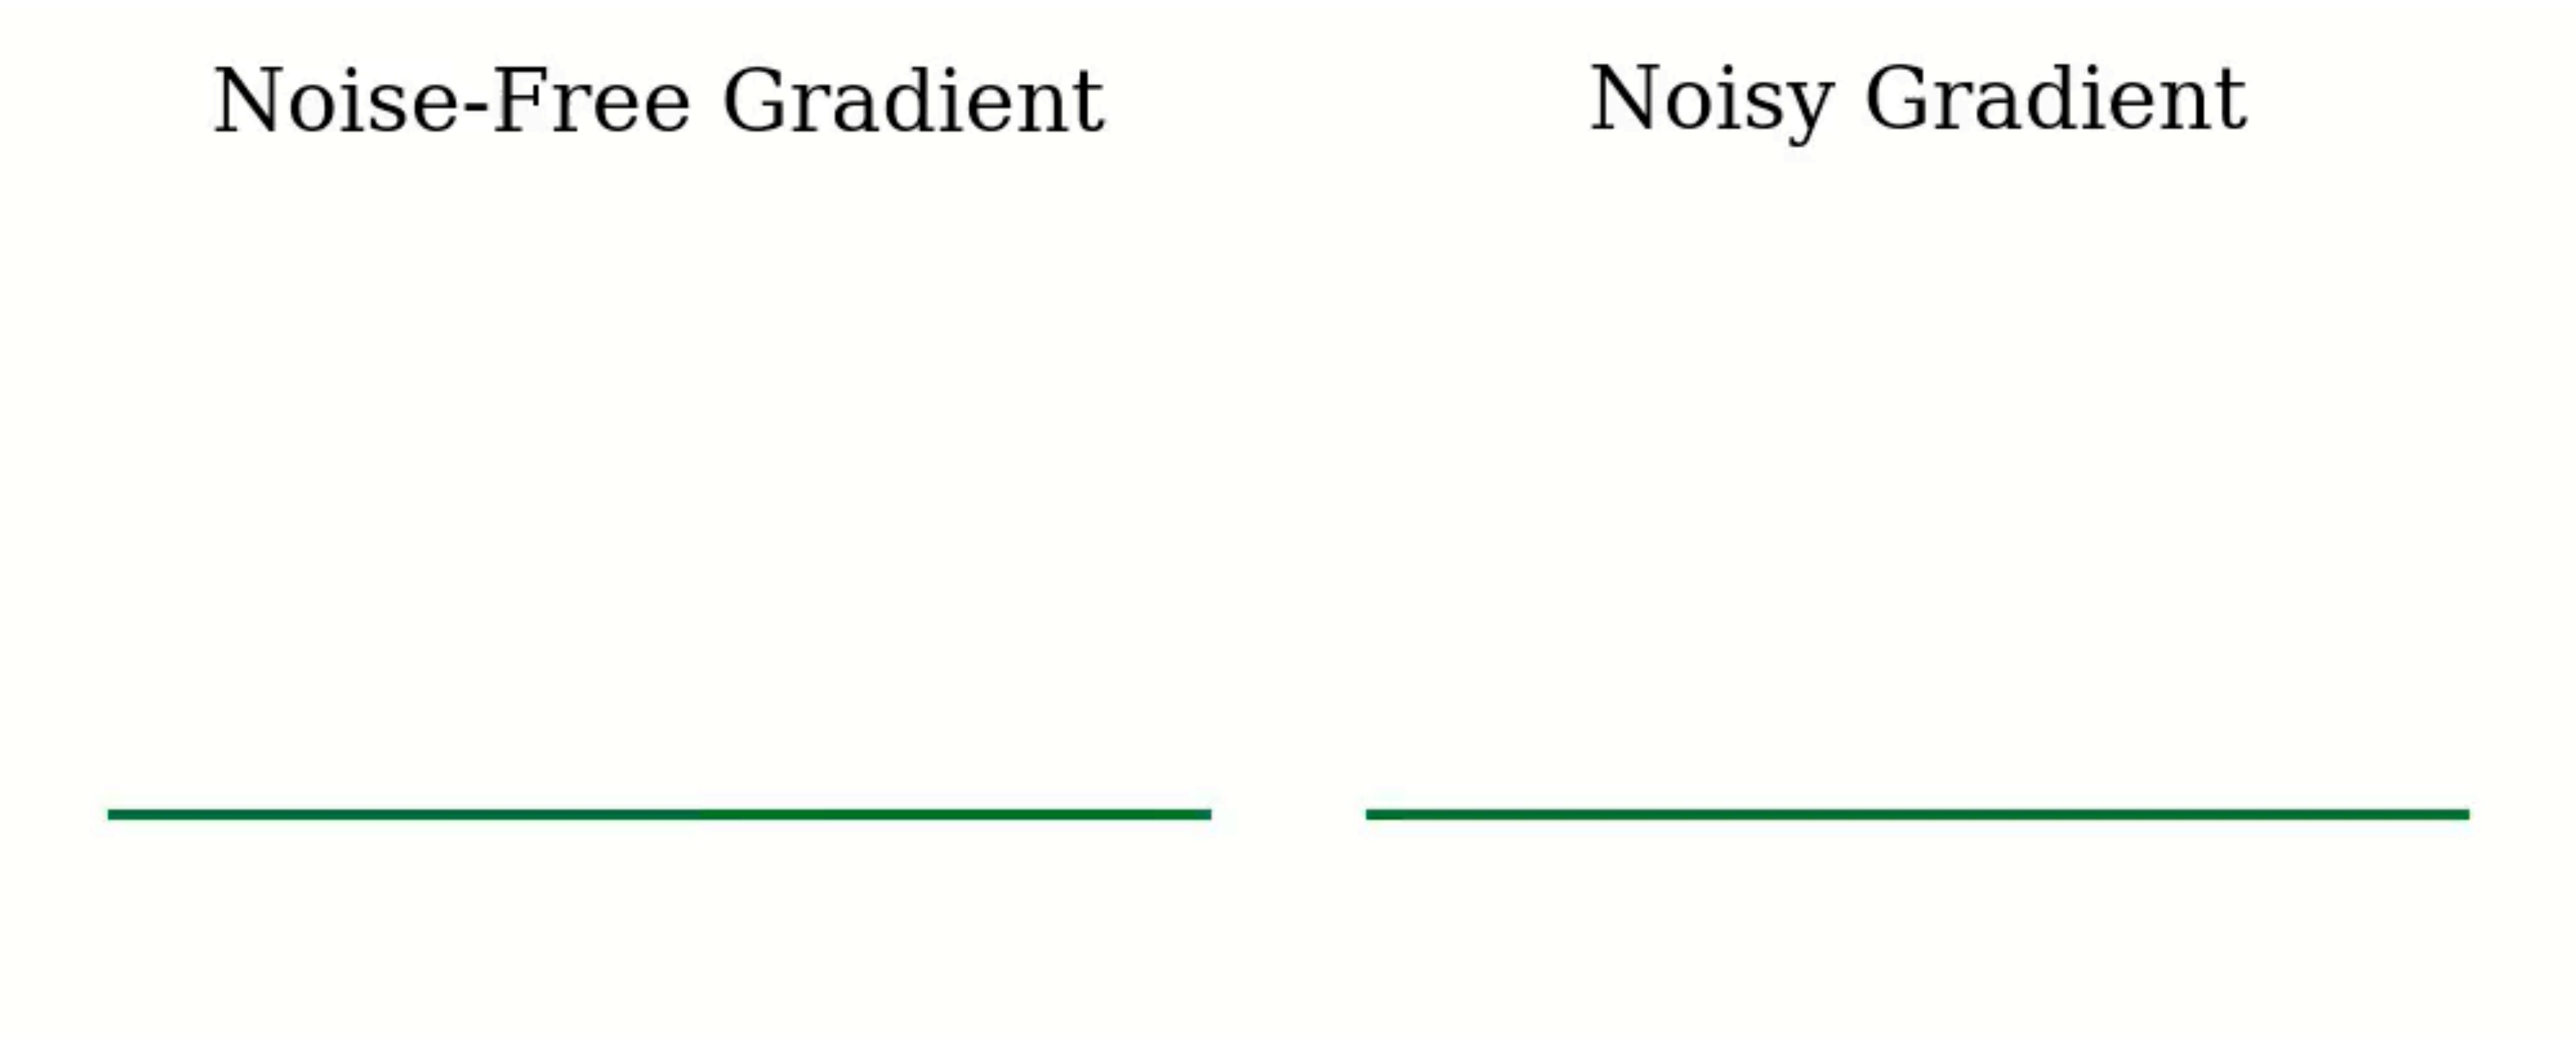

ADAM Optimizer to the rescue?

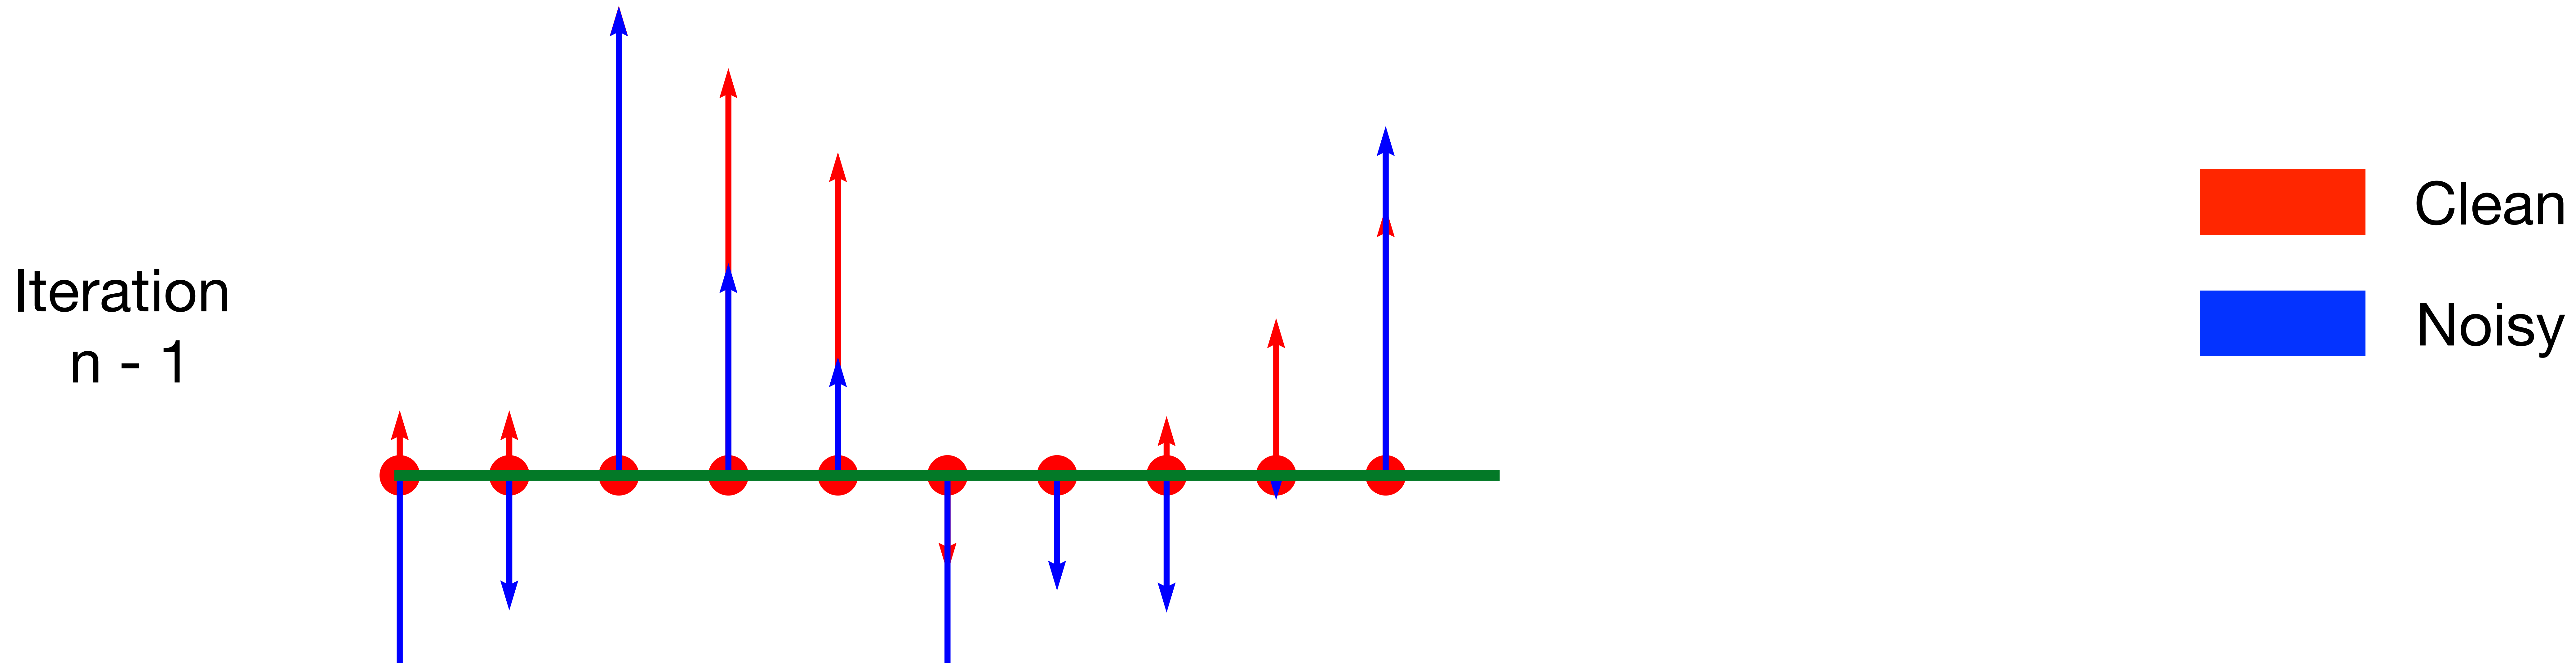

ADAM Optimizer to the rescue?

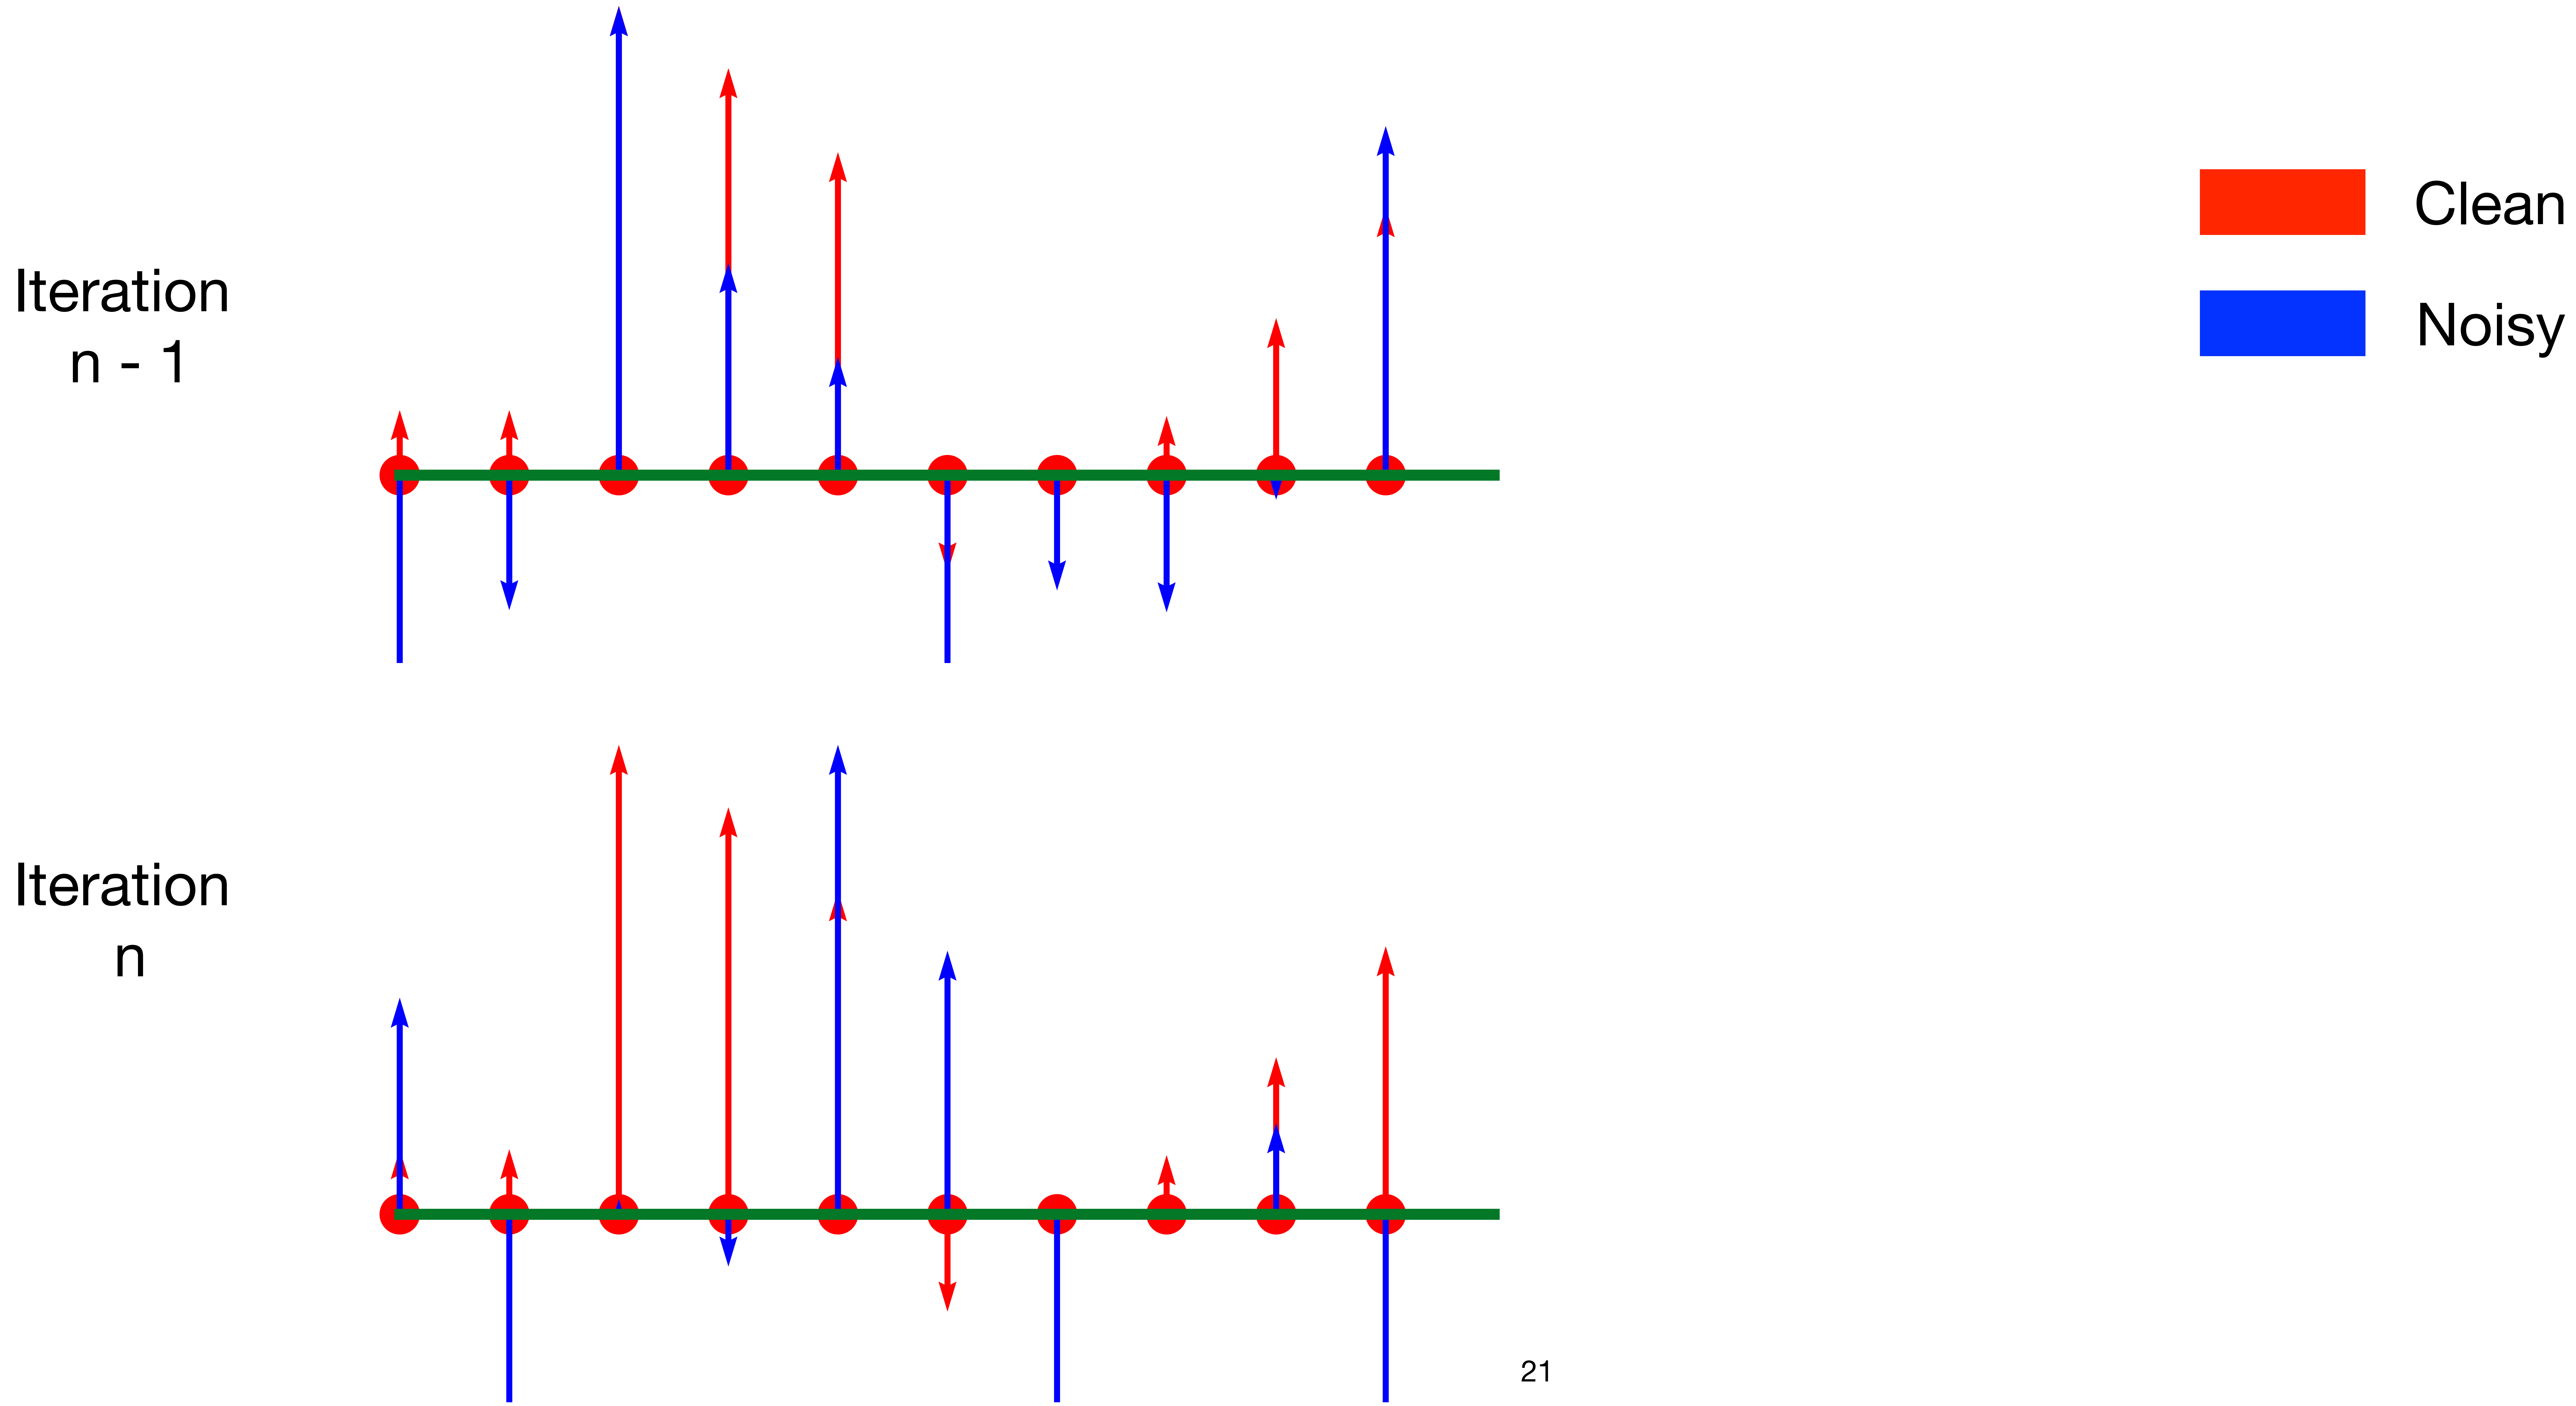

ADAM Optimizer to the rescue?

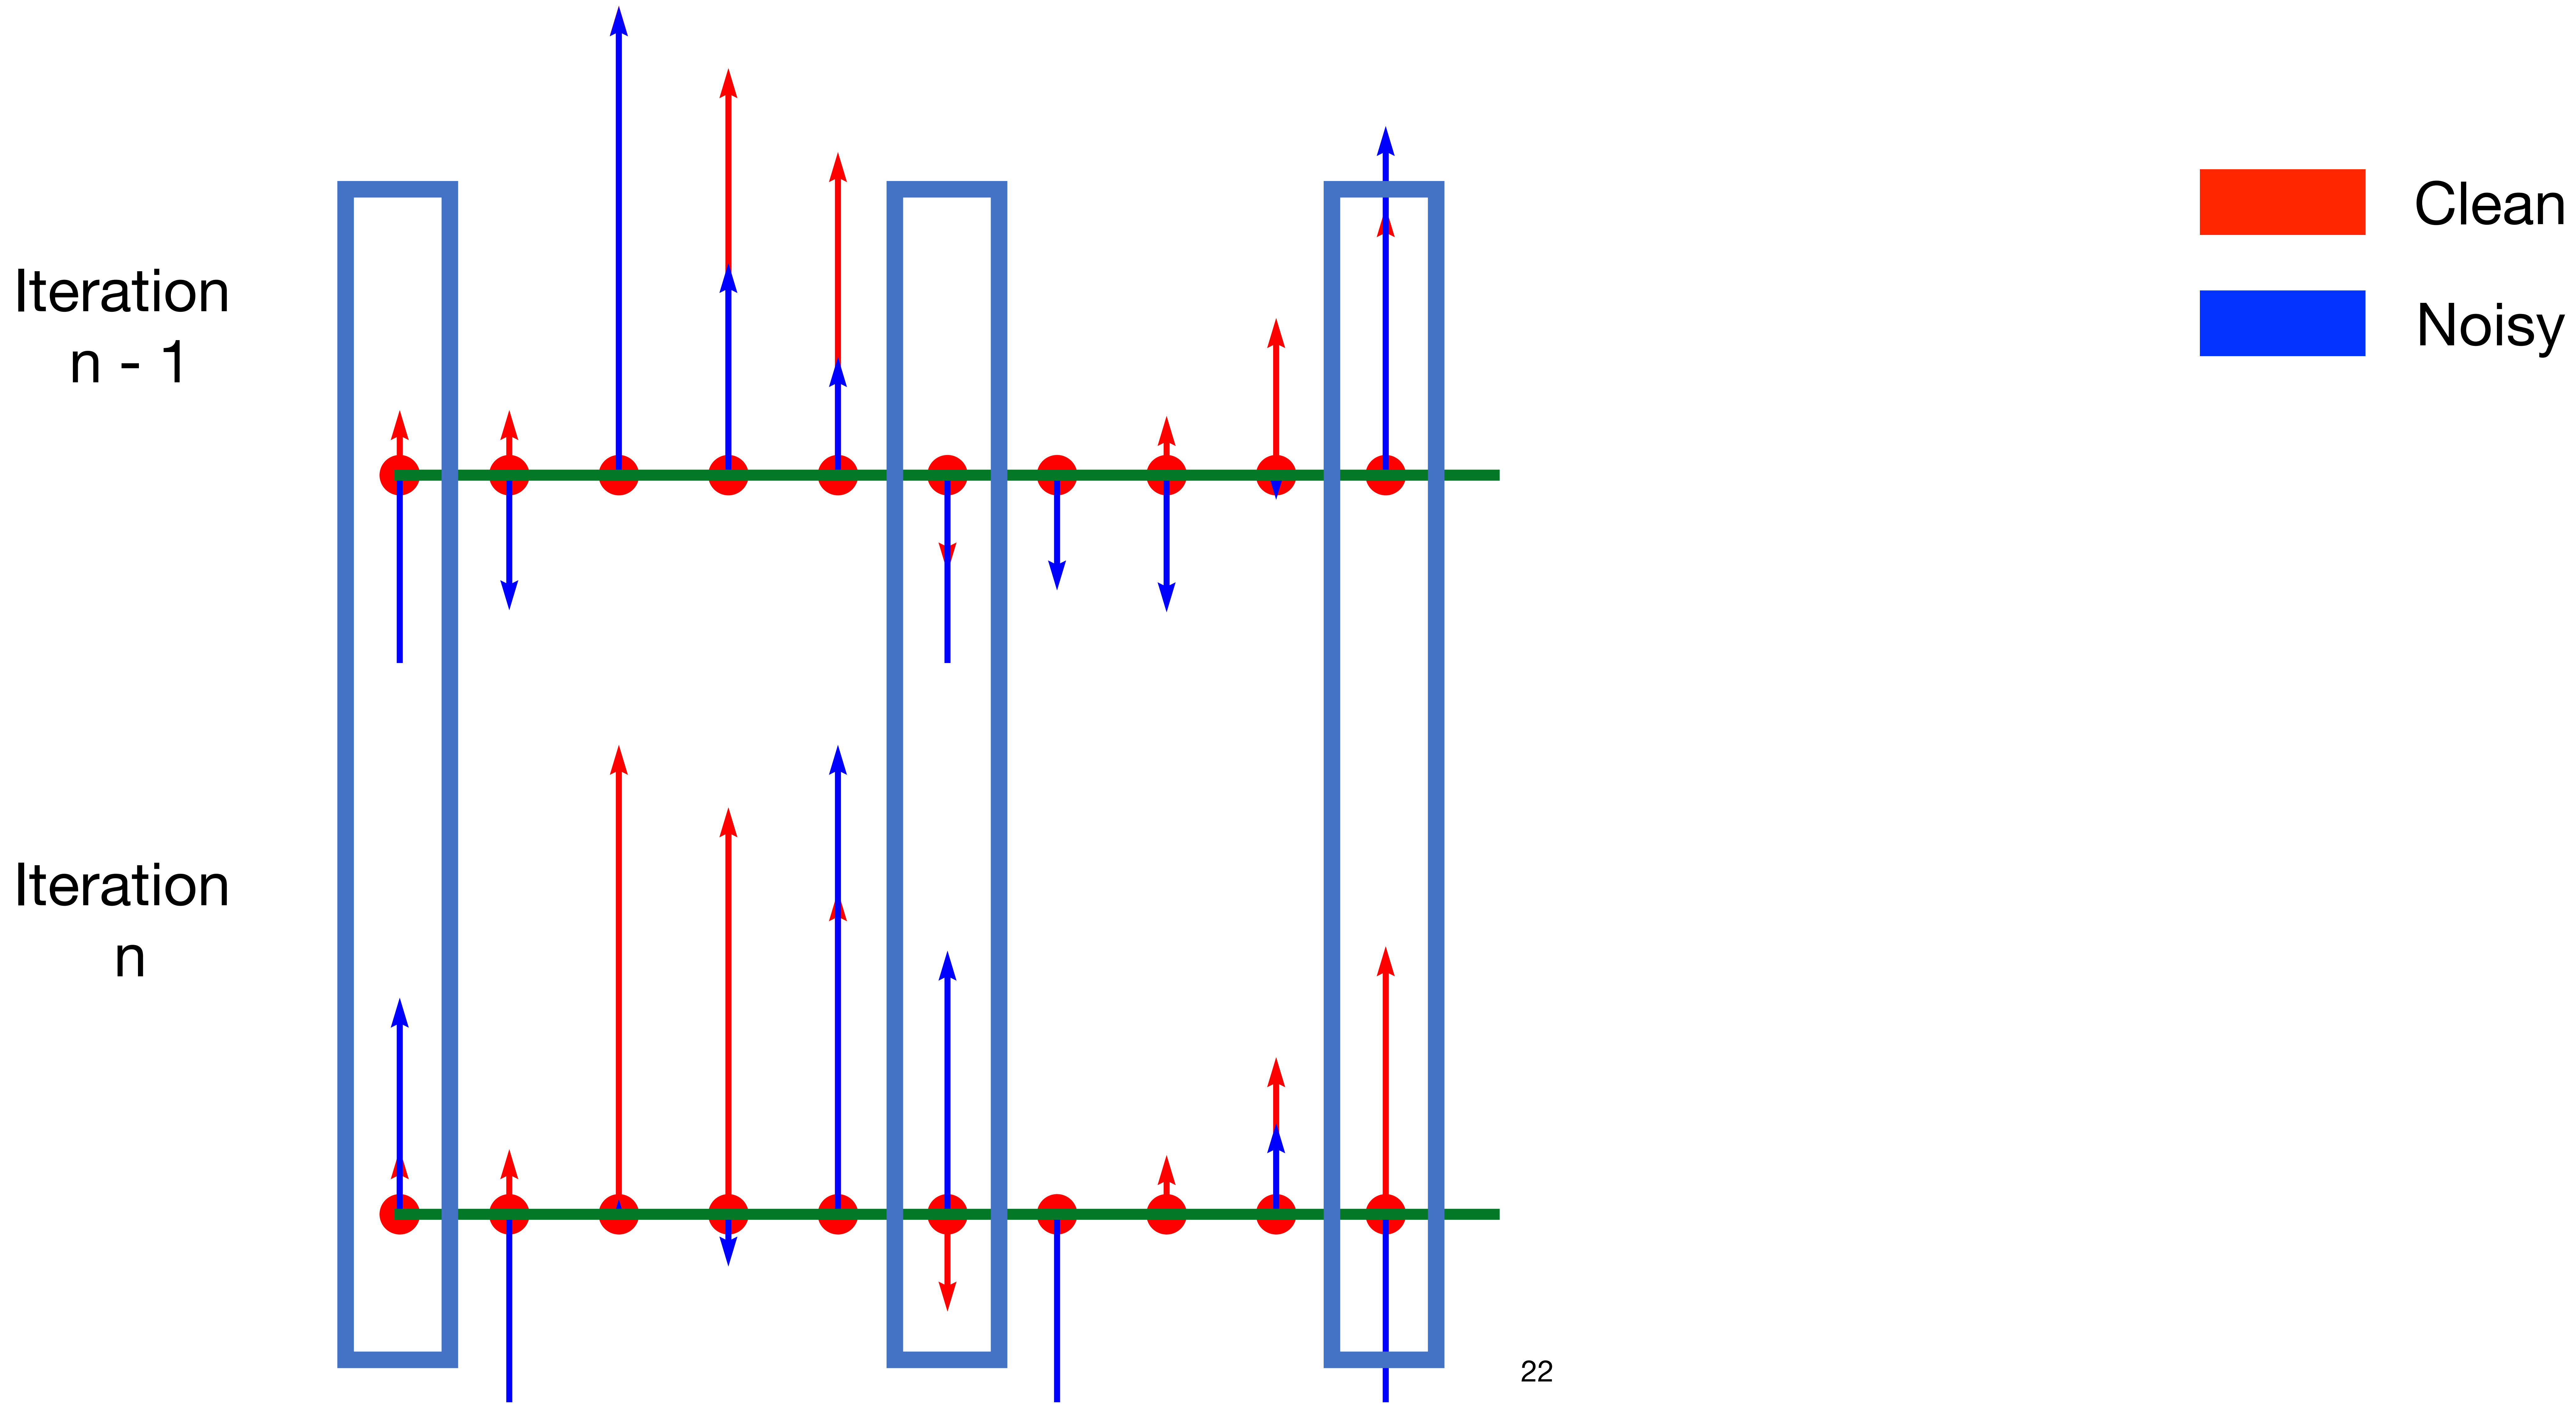

ADAM Optimizer to the rescue?

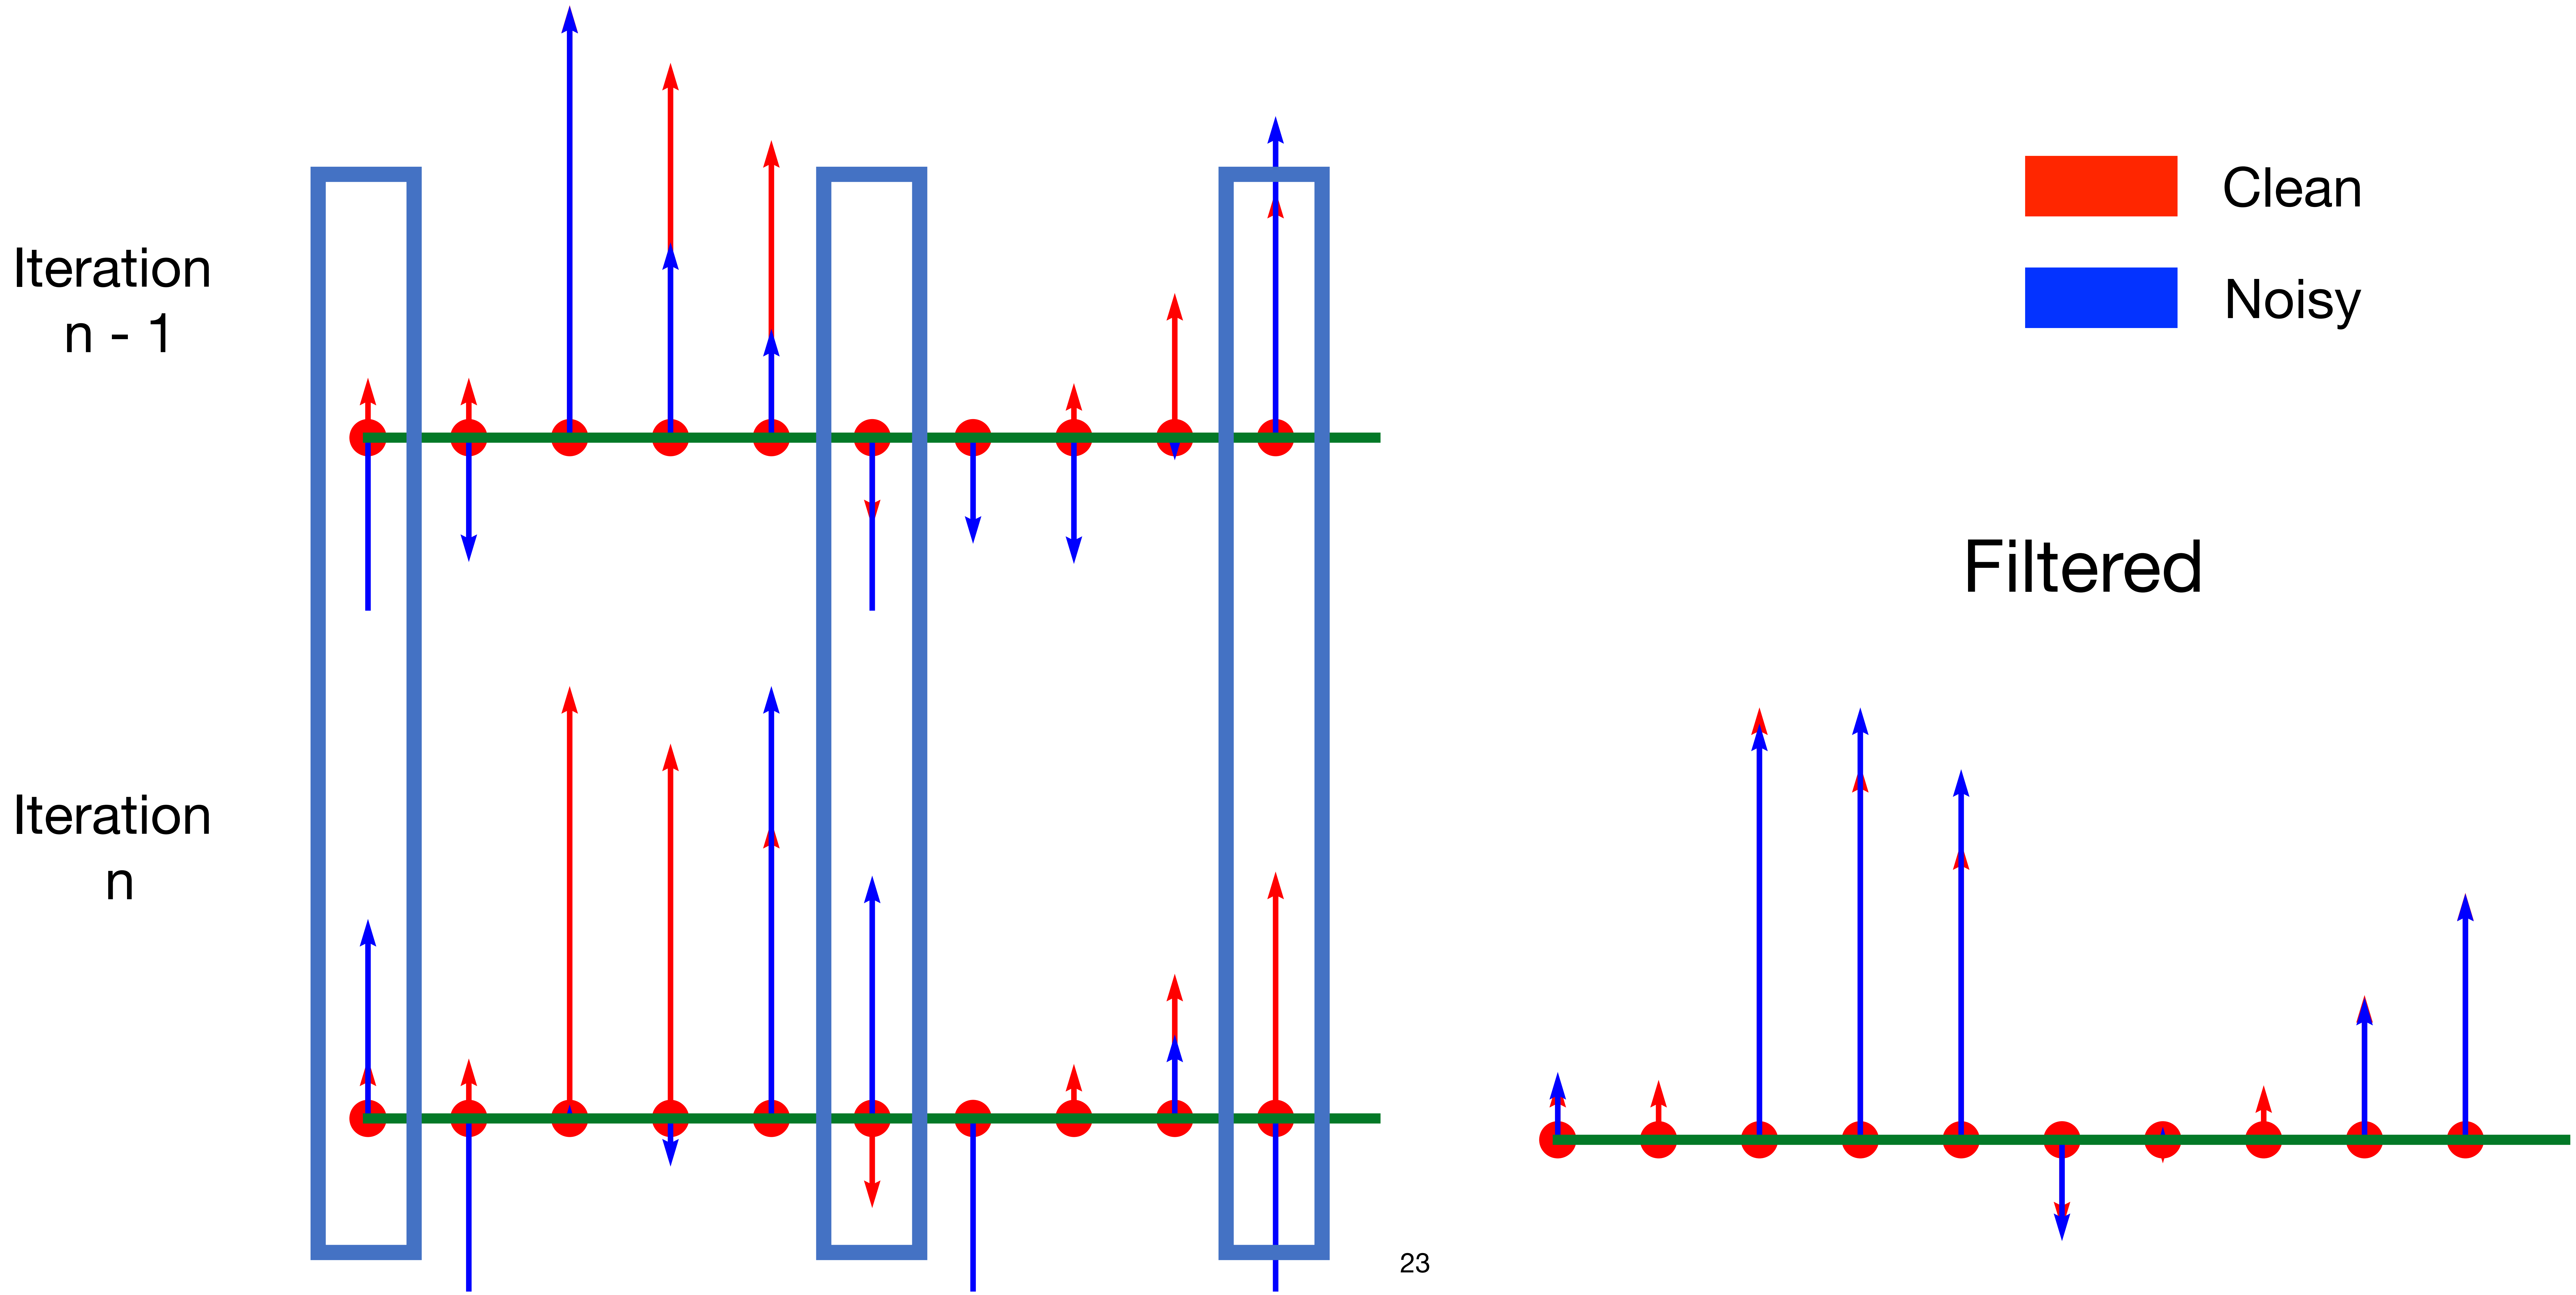

ADAM Optimizer to the rescue?

ADAM optimizer **temporally** filters the gradients

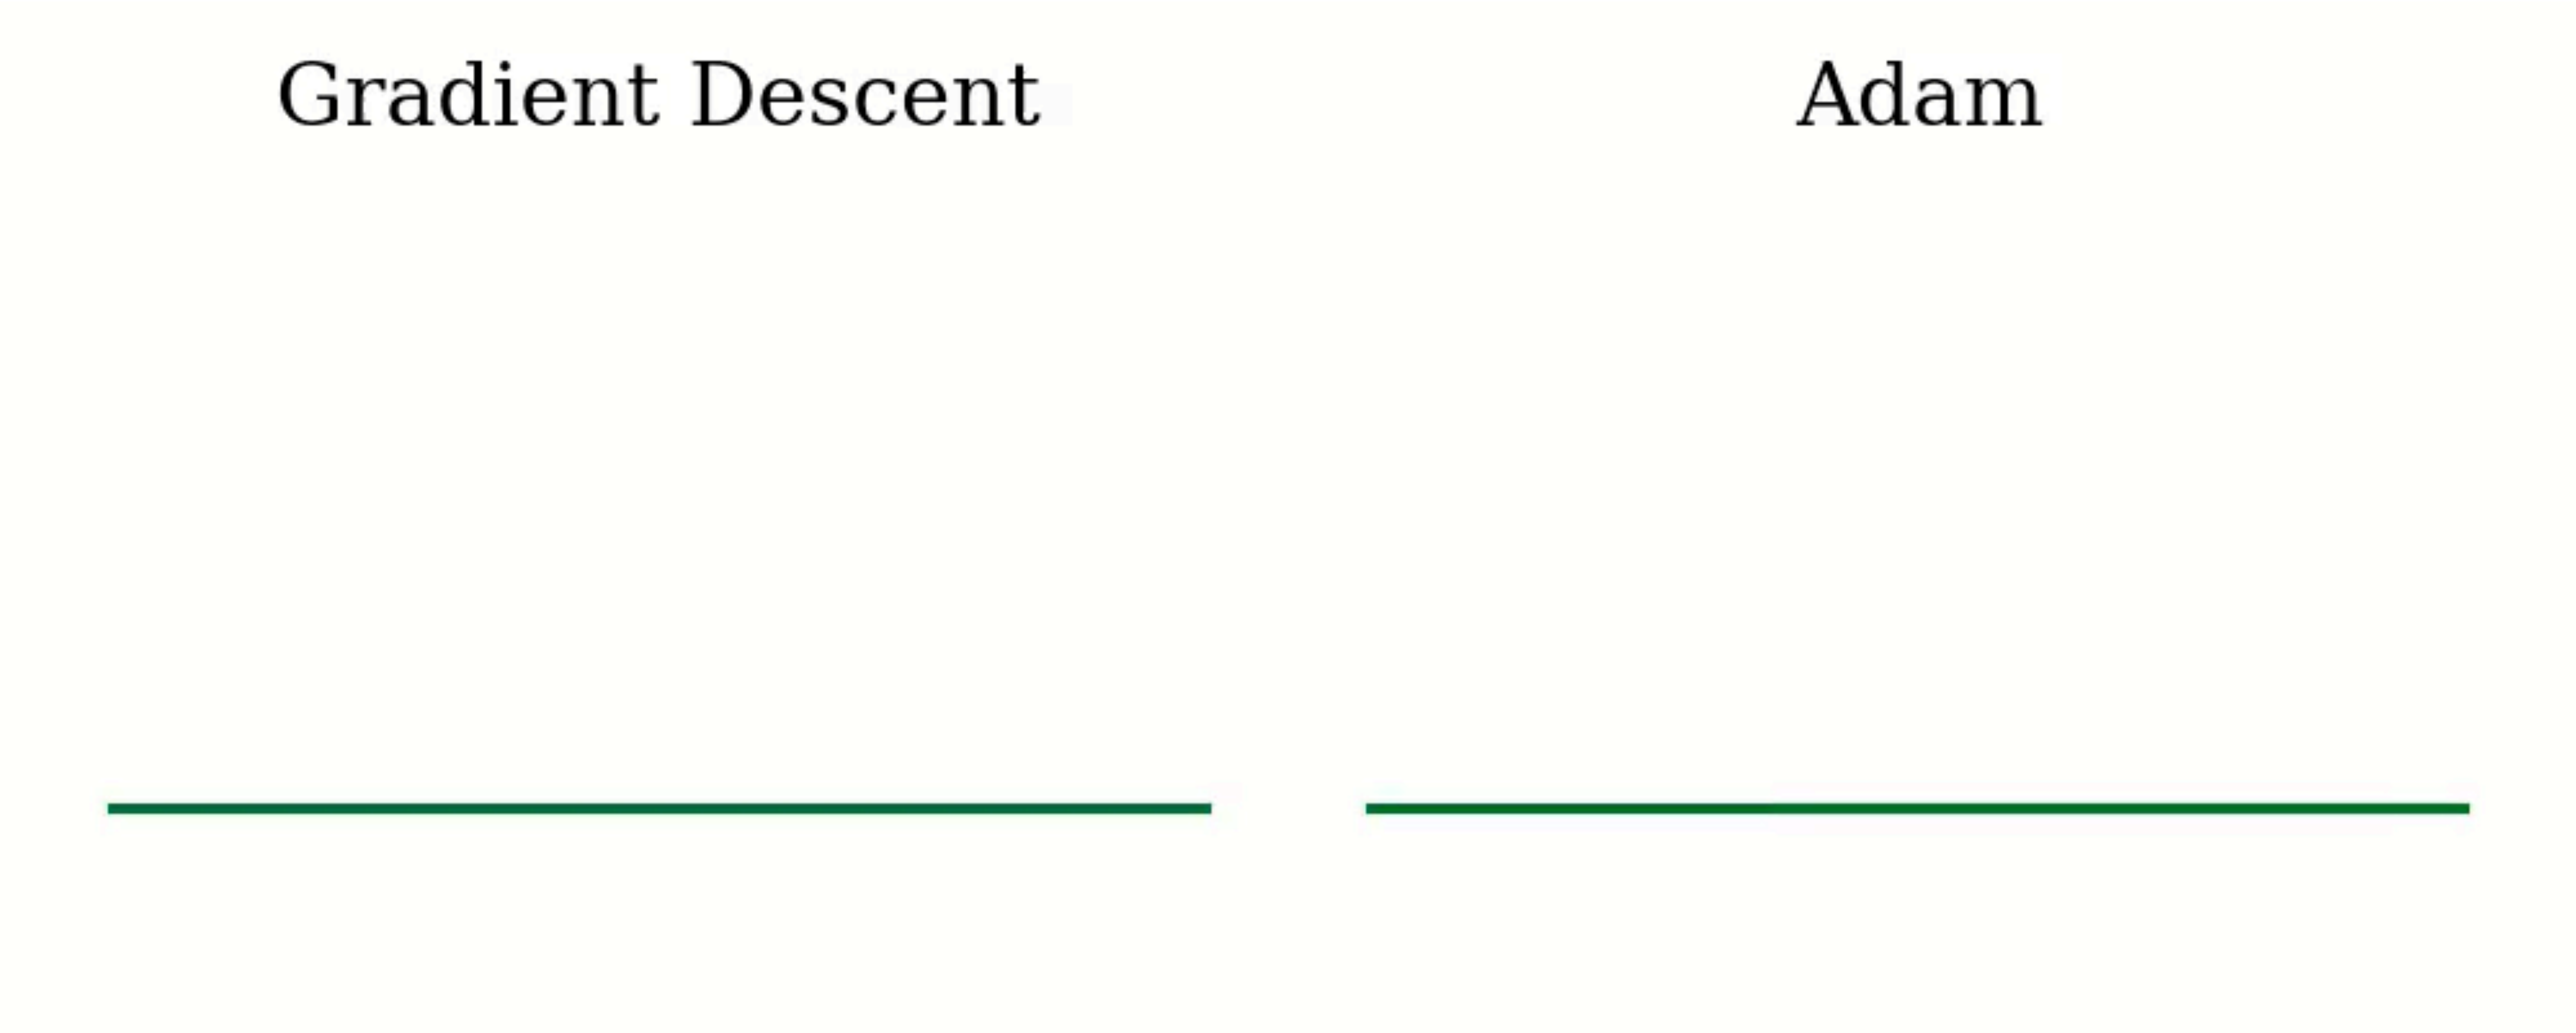

Can we go beyond temporal filtering?

Spatial Filtering: Laplacian Smoothing

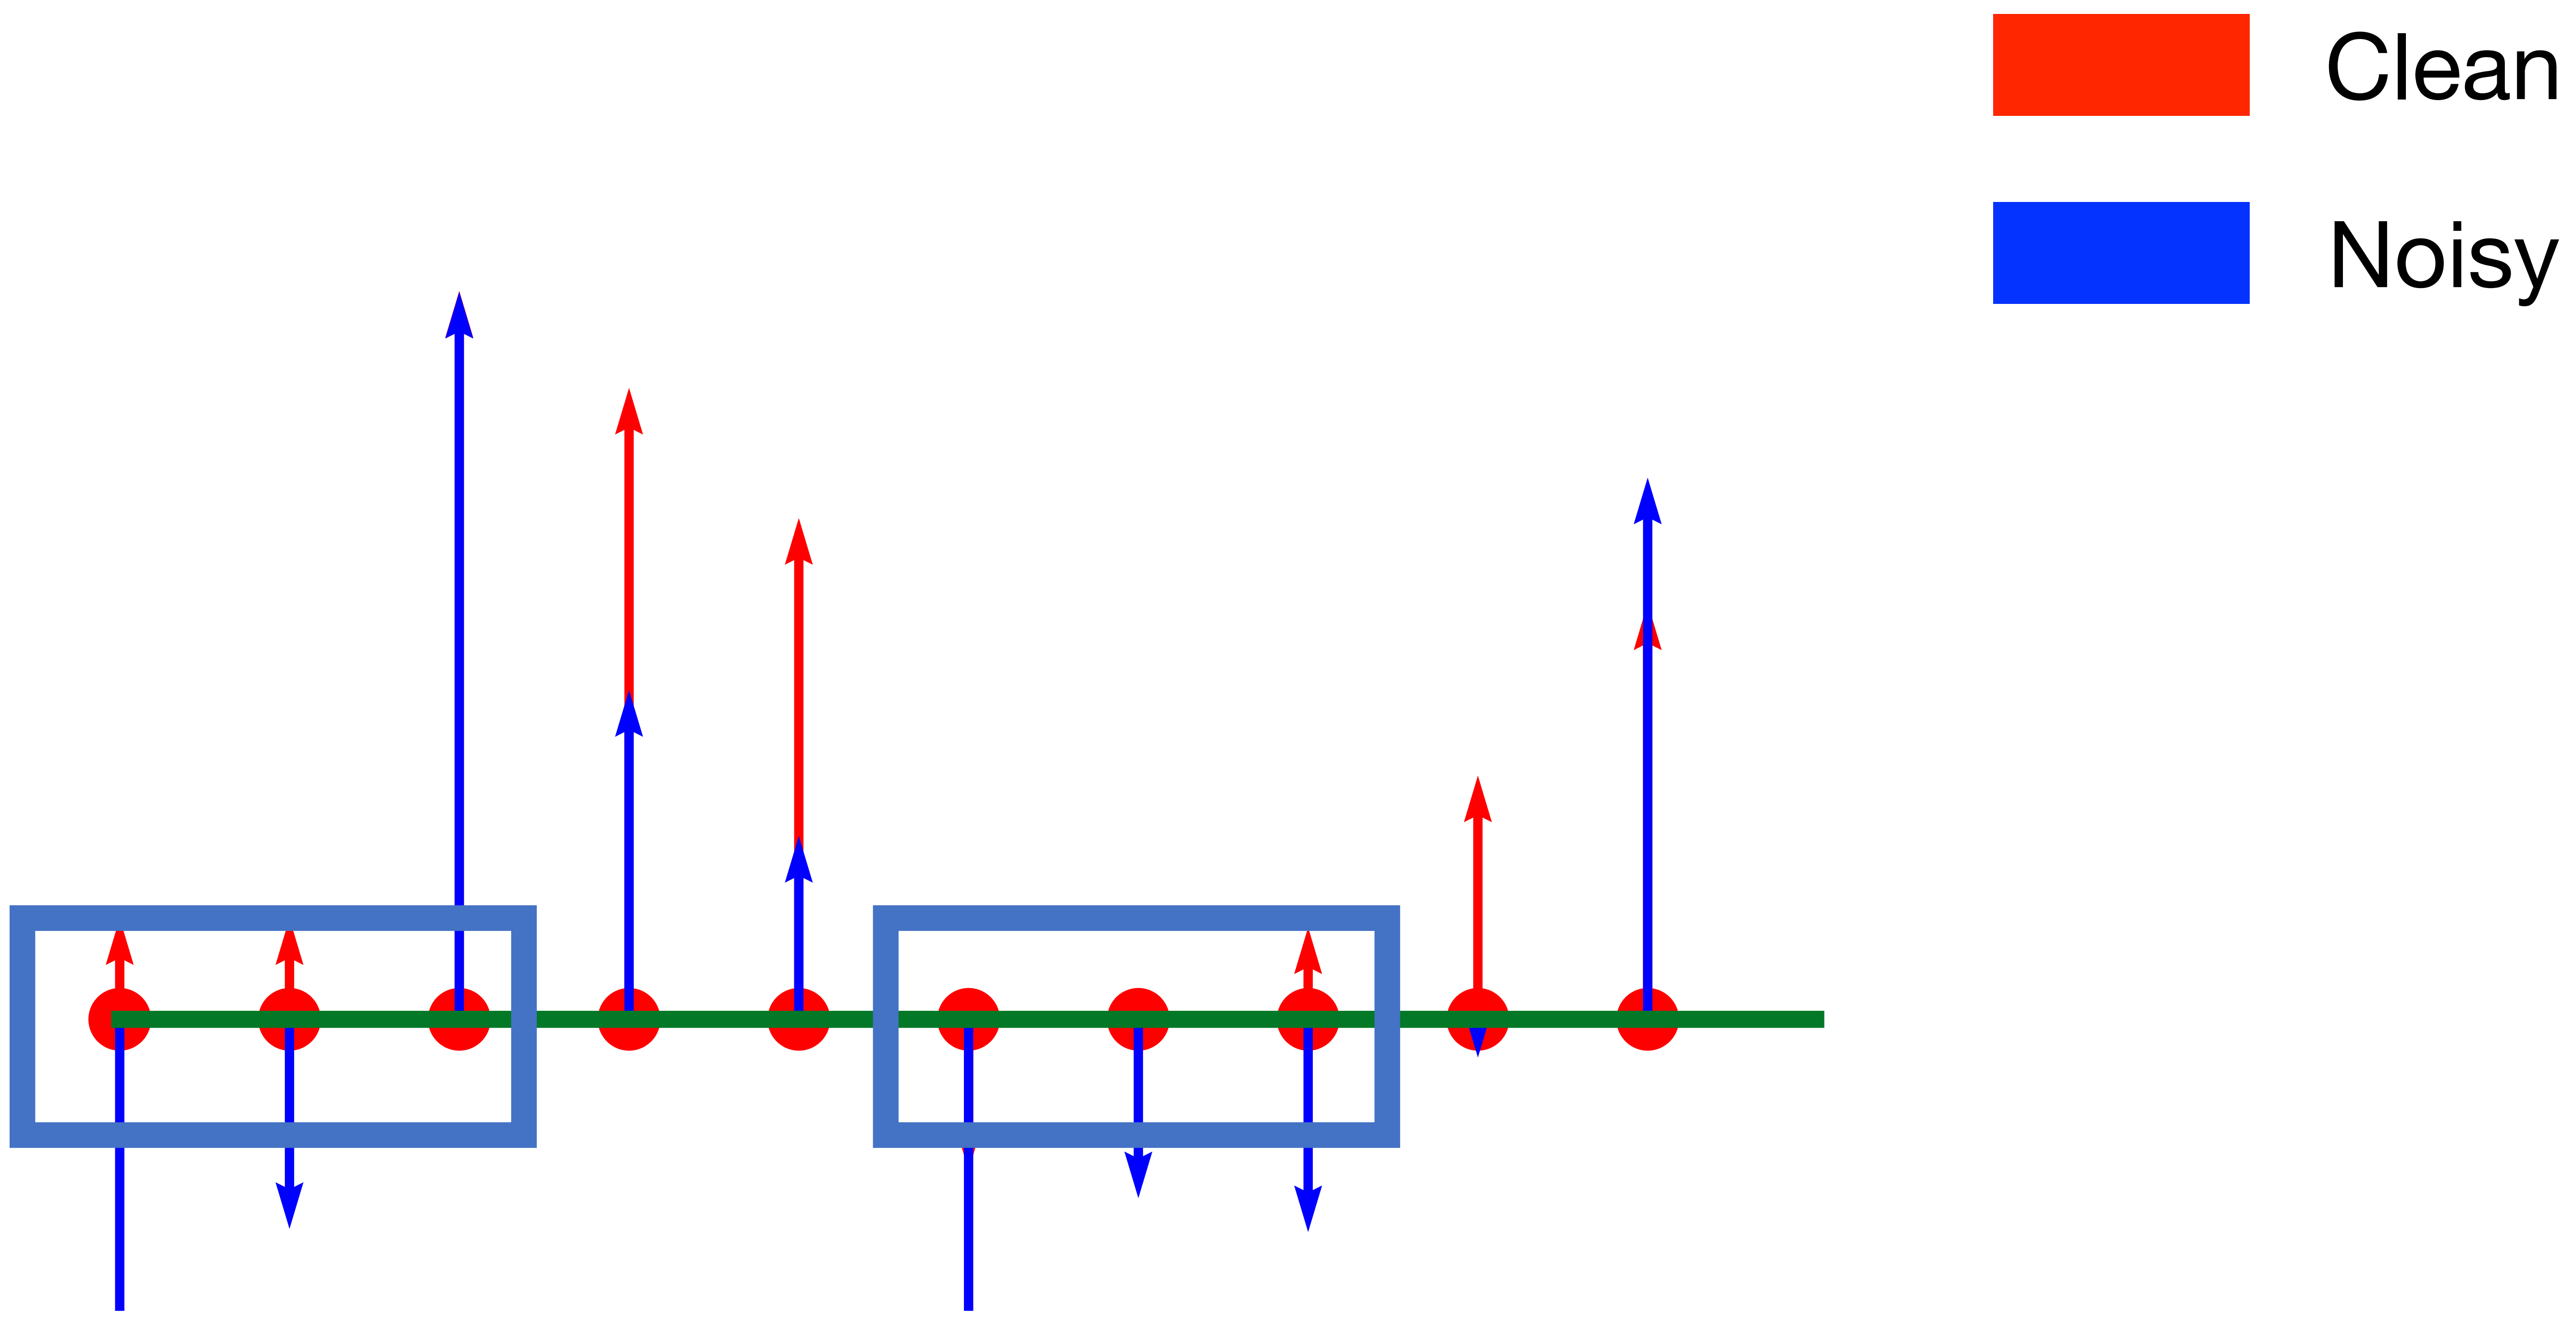

Laplacian smoothing gradient descent [Osher et al. 2018]

Large steps in inverse rendering of geometry [Nicolet et al. 2021]

Spatial Filtering: Laplacian Smoothing

Smoother optimization!

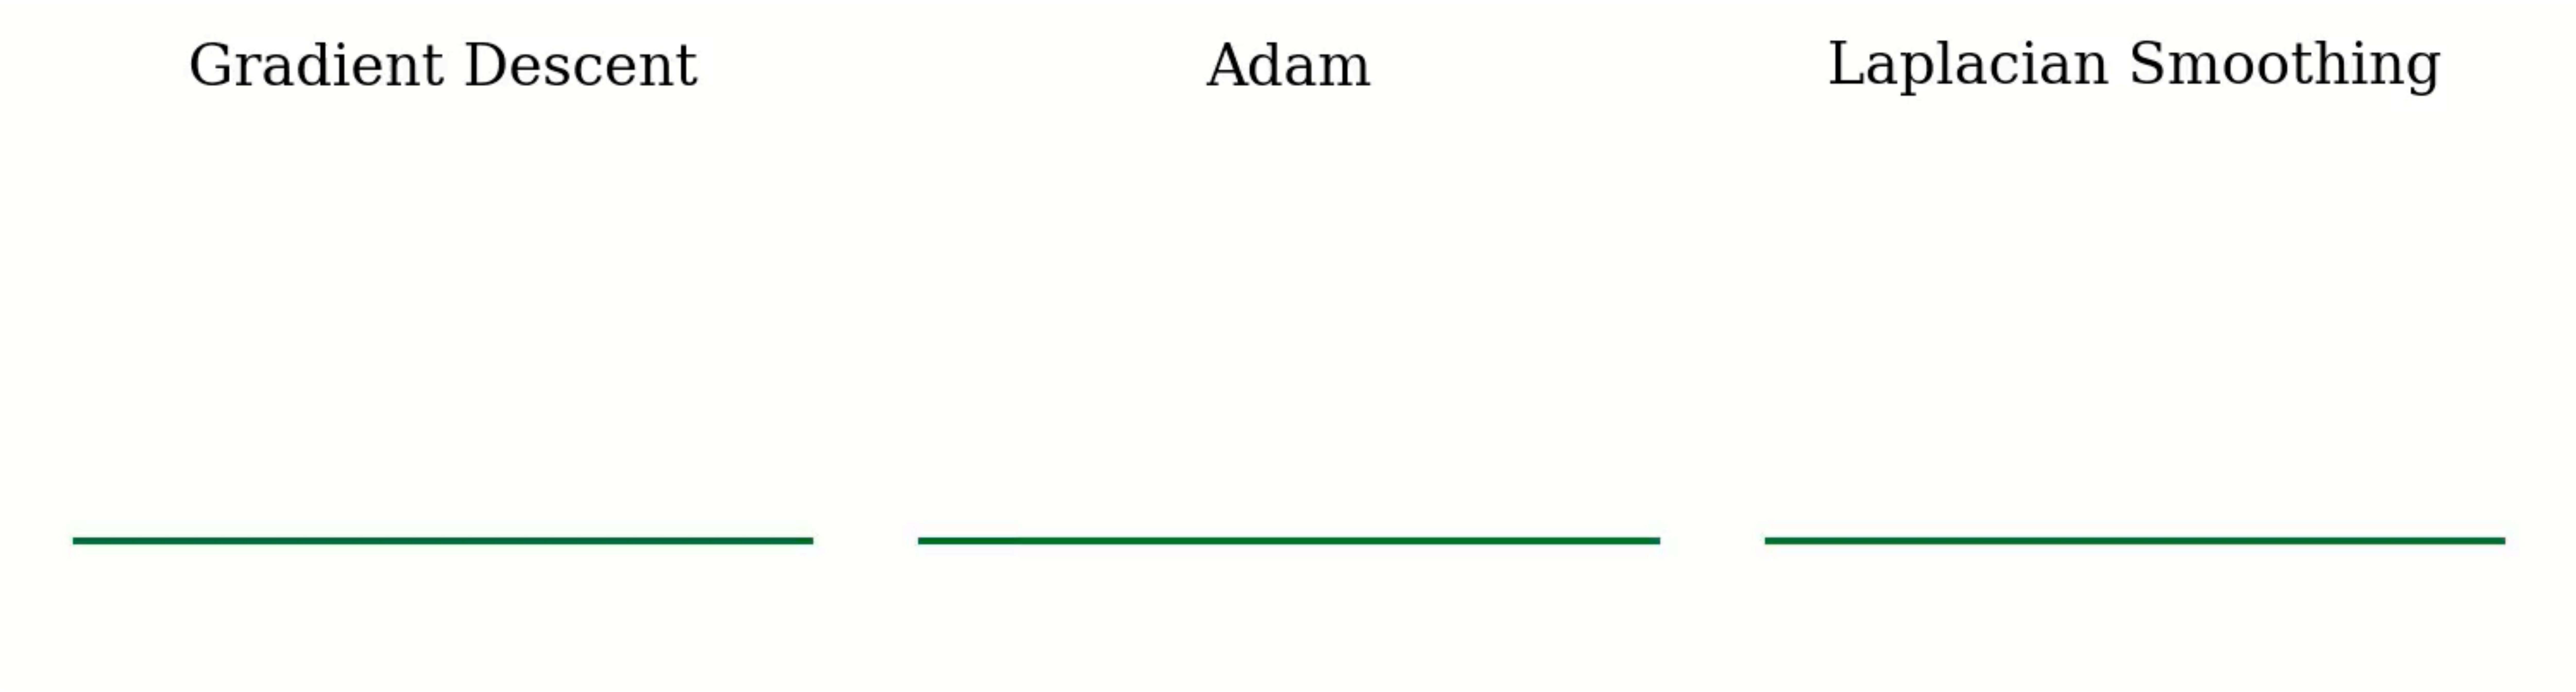

Spatial Filtering: Laplacian Smoothing

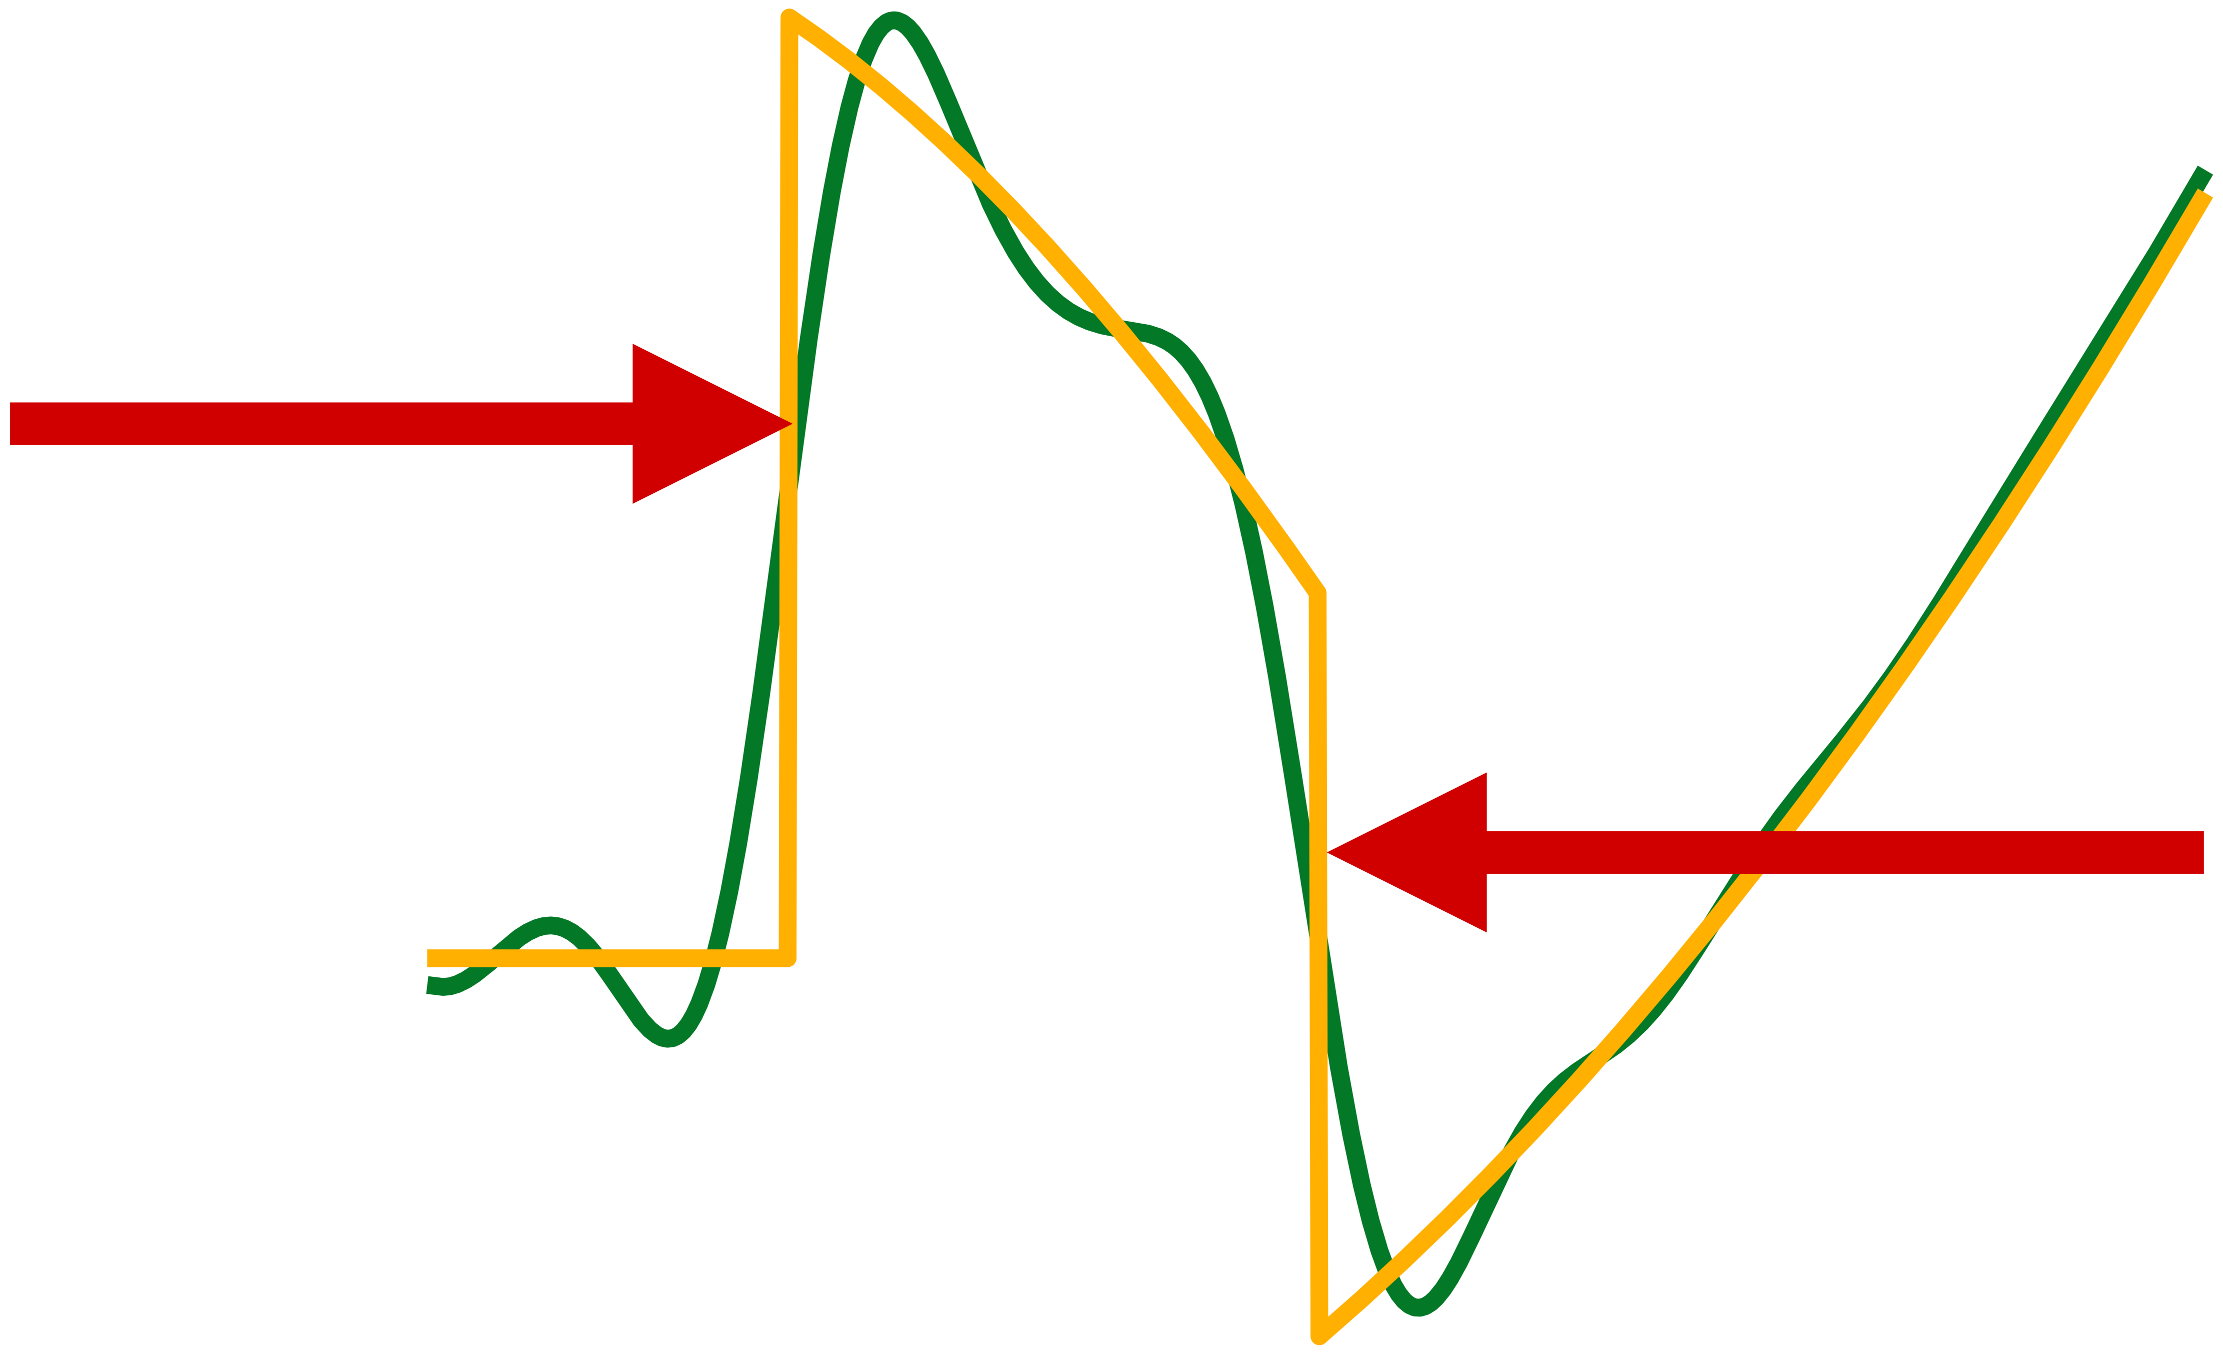

Can we preserve edges?

Can we preserve edges?

Edge-aware filters!

Ours: Edge-preserving Spatiotemporal Filtering

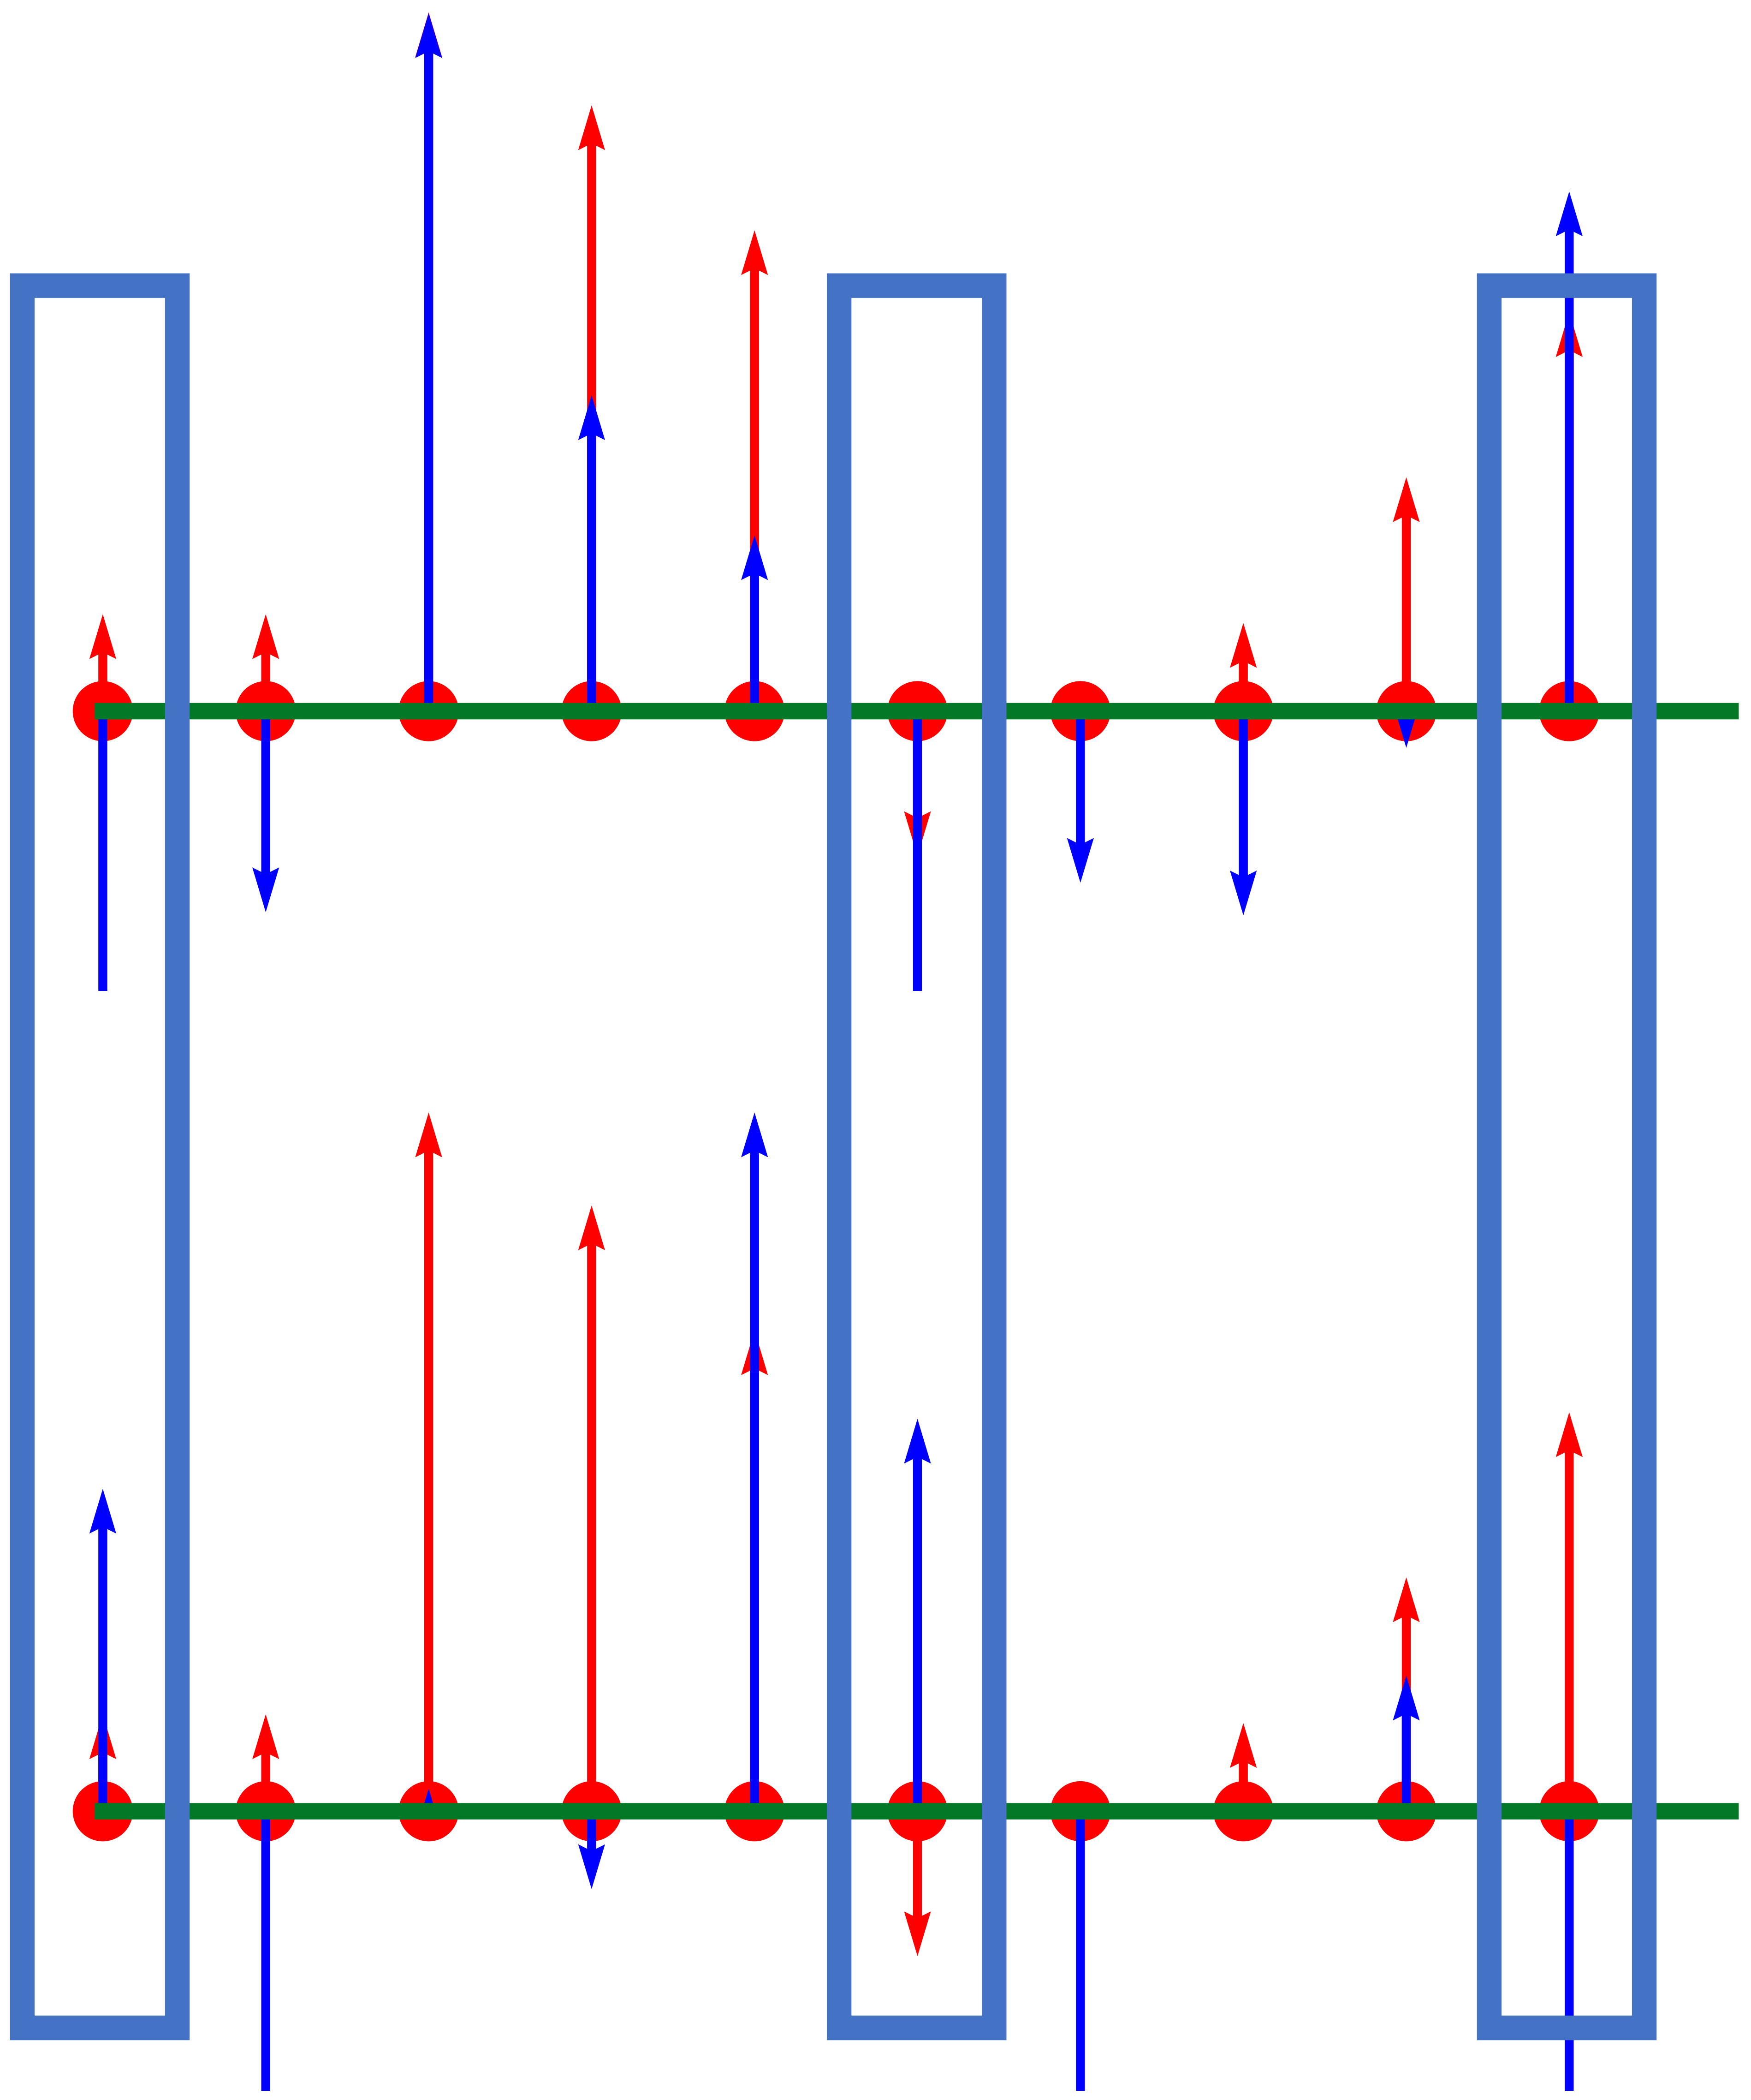

Ours: Edge-preserving Spatiotemporal Filtering

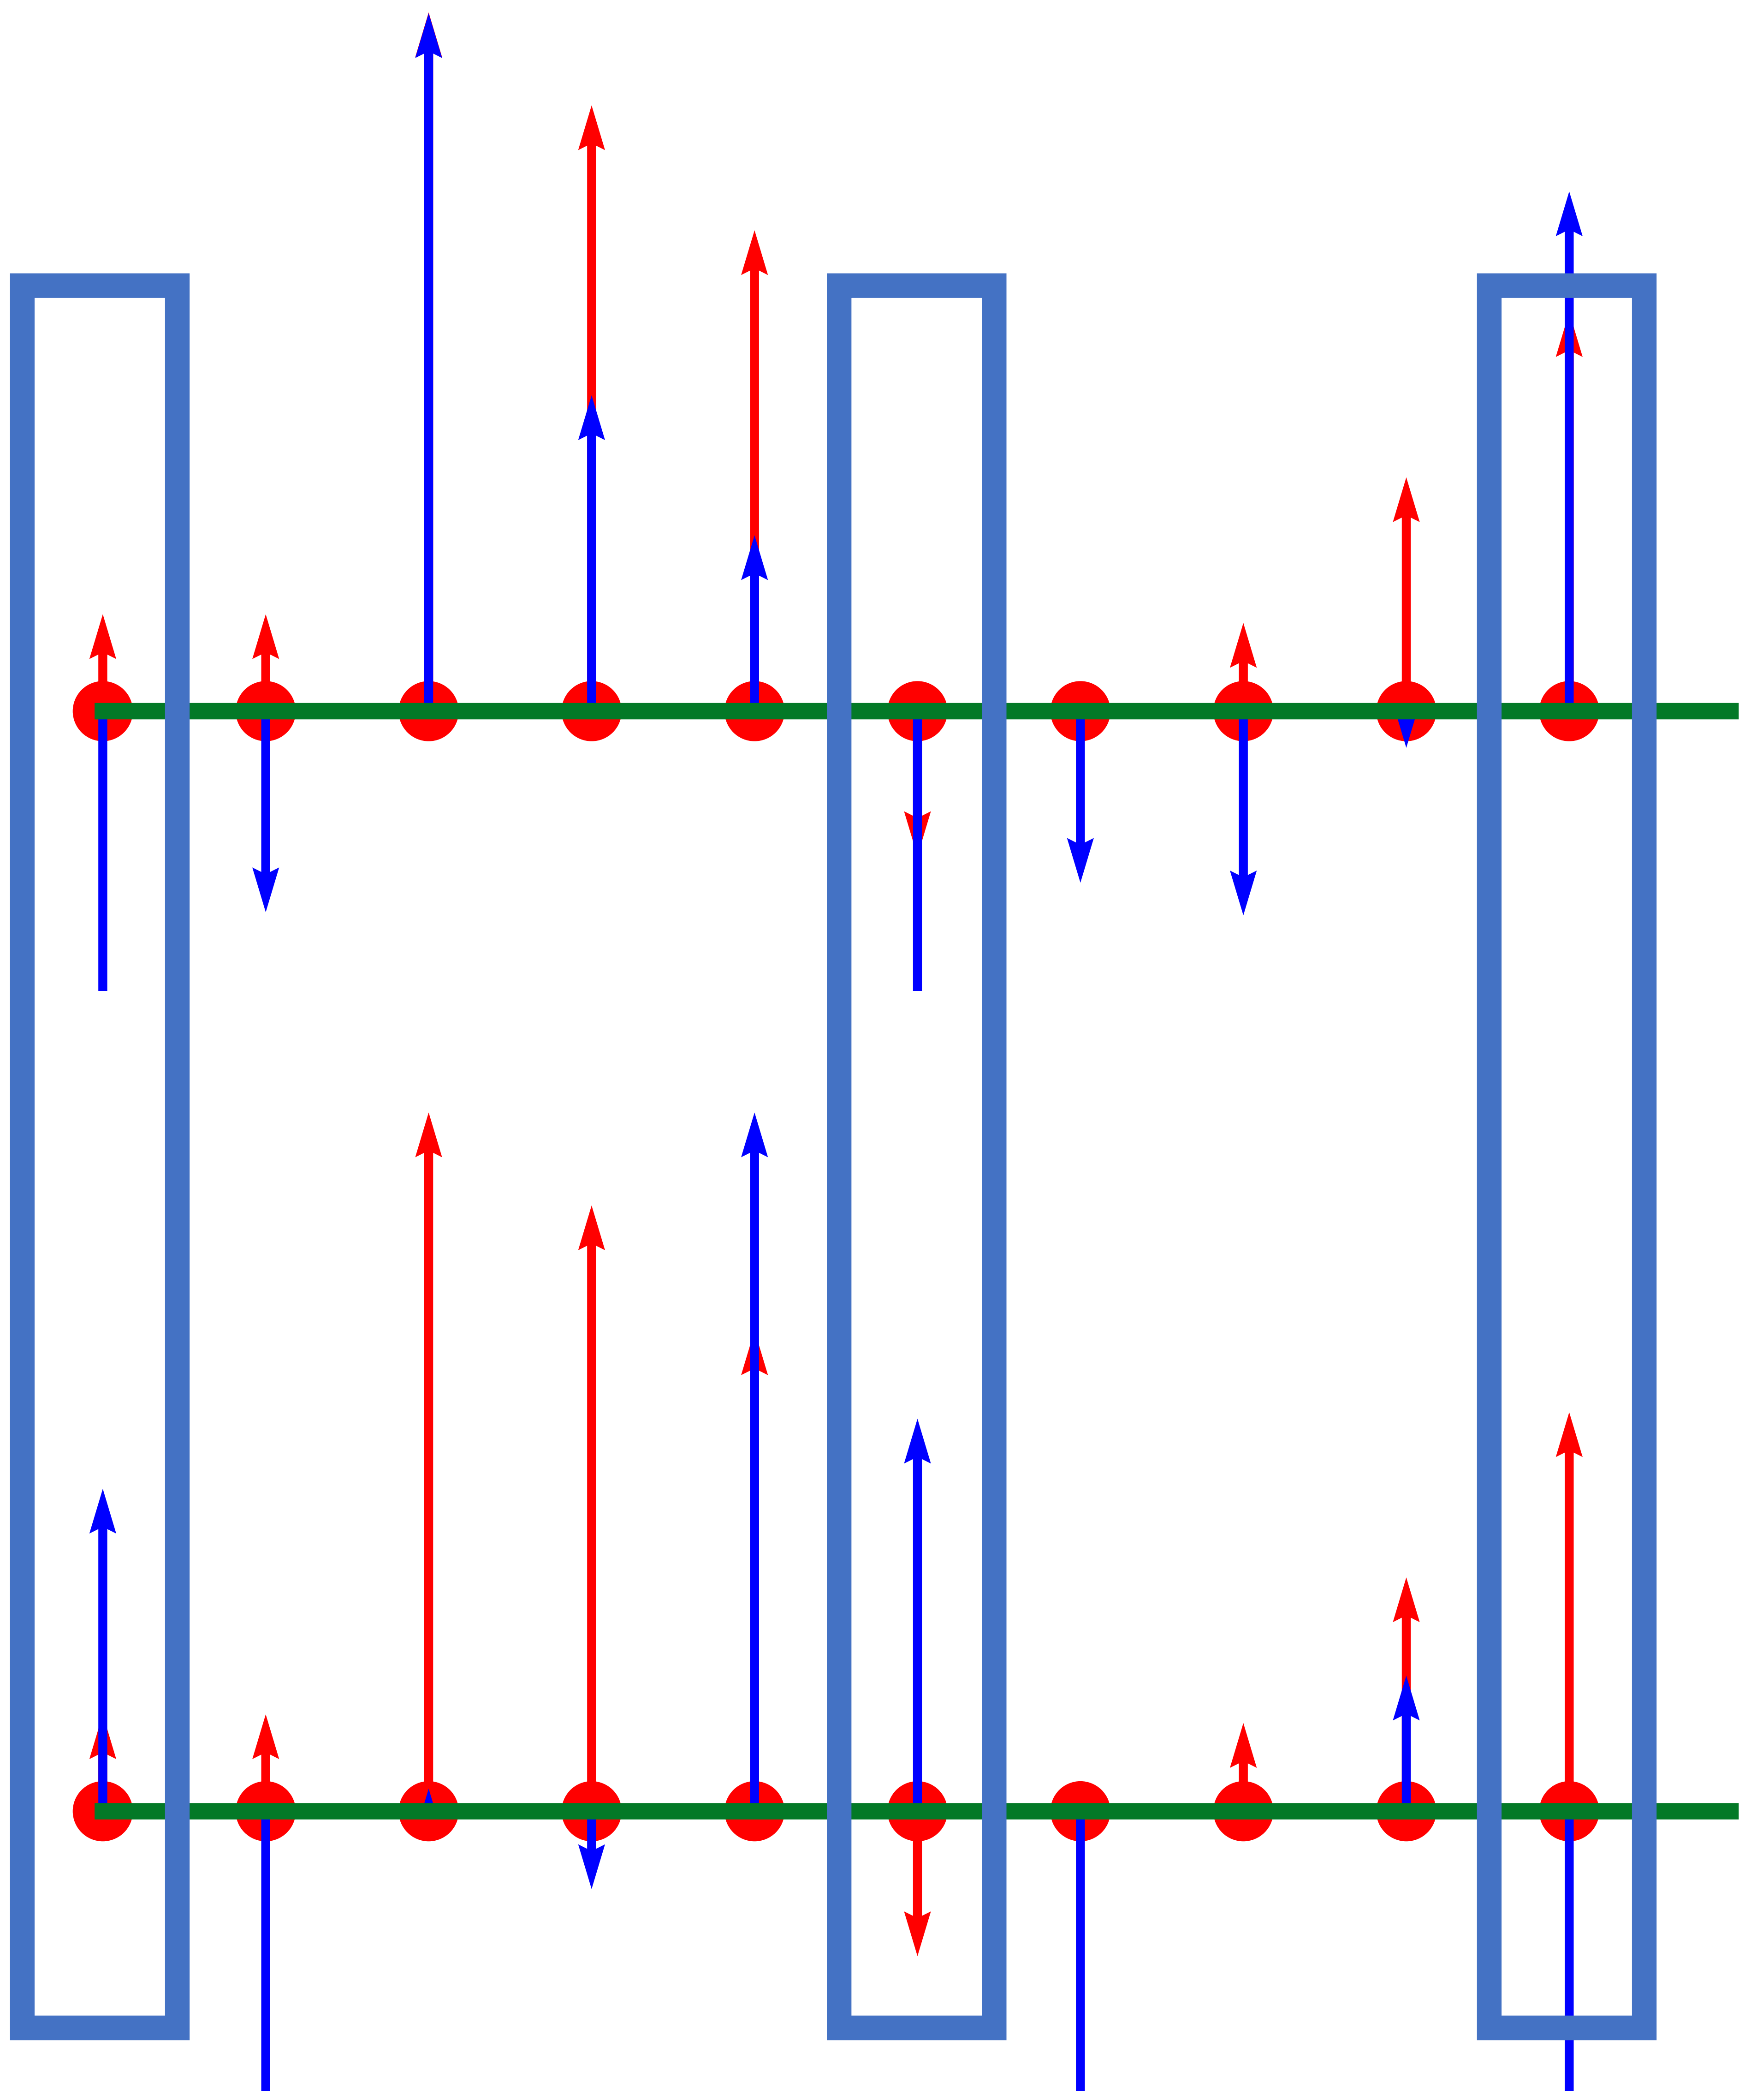

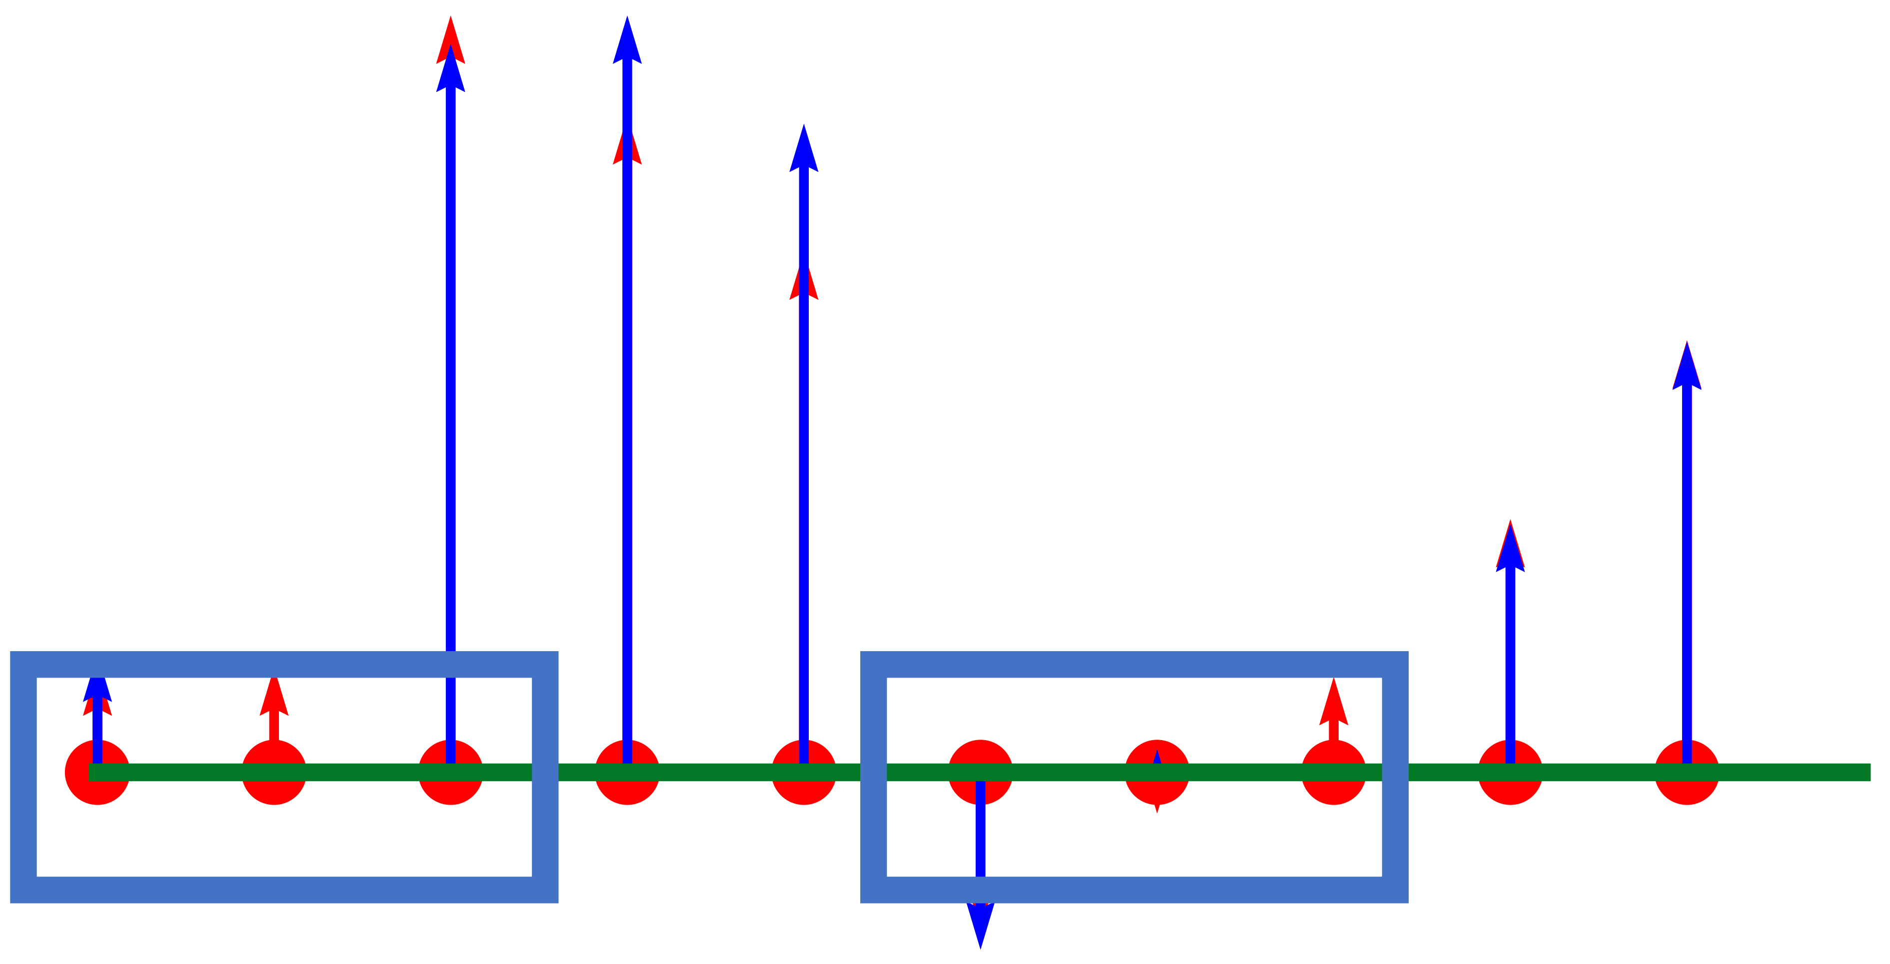

Ours: Edge-preserving Spatiotemporal Filtering

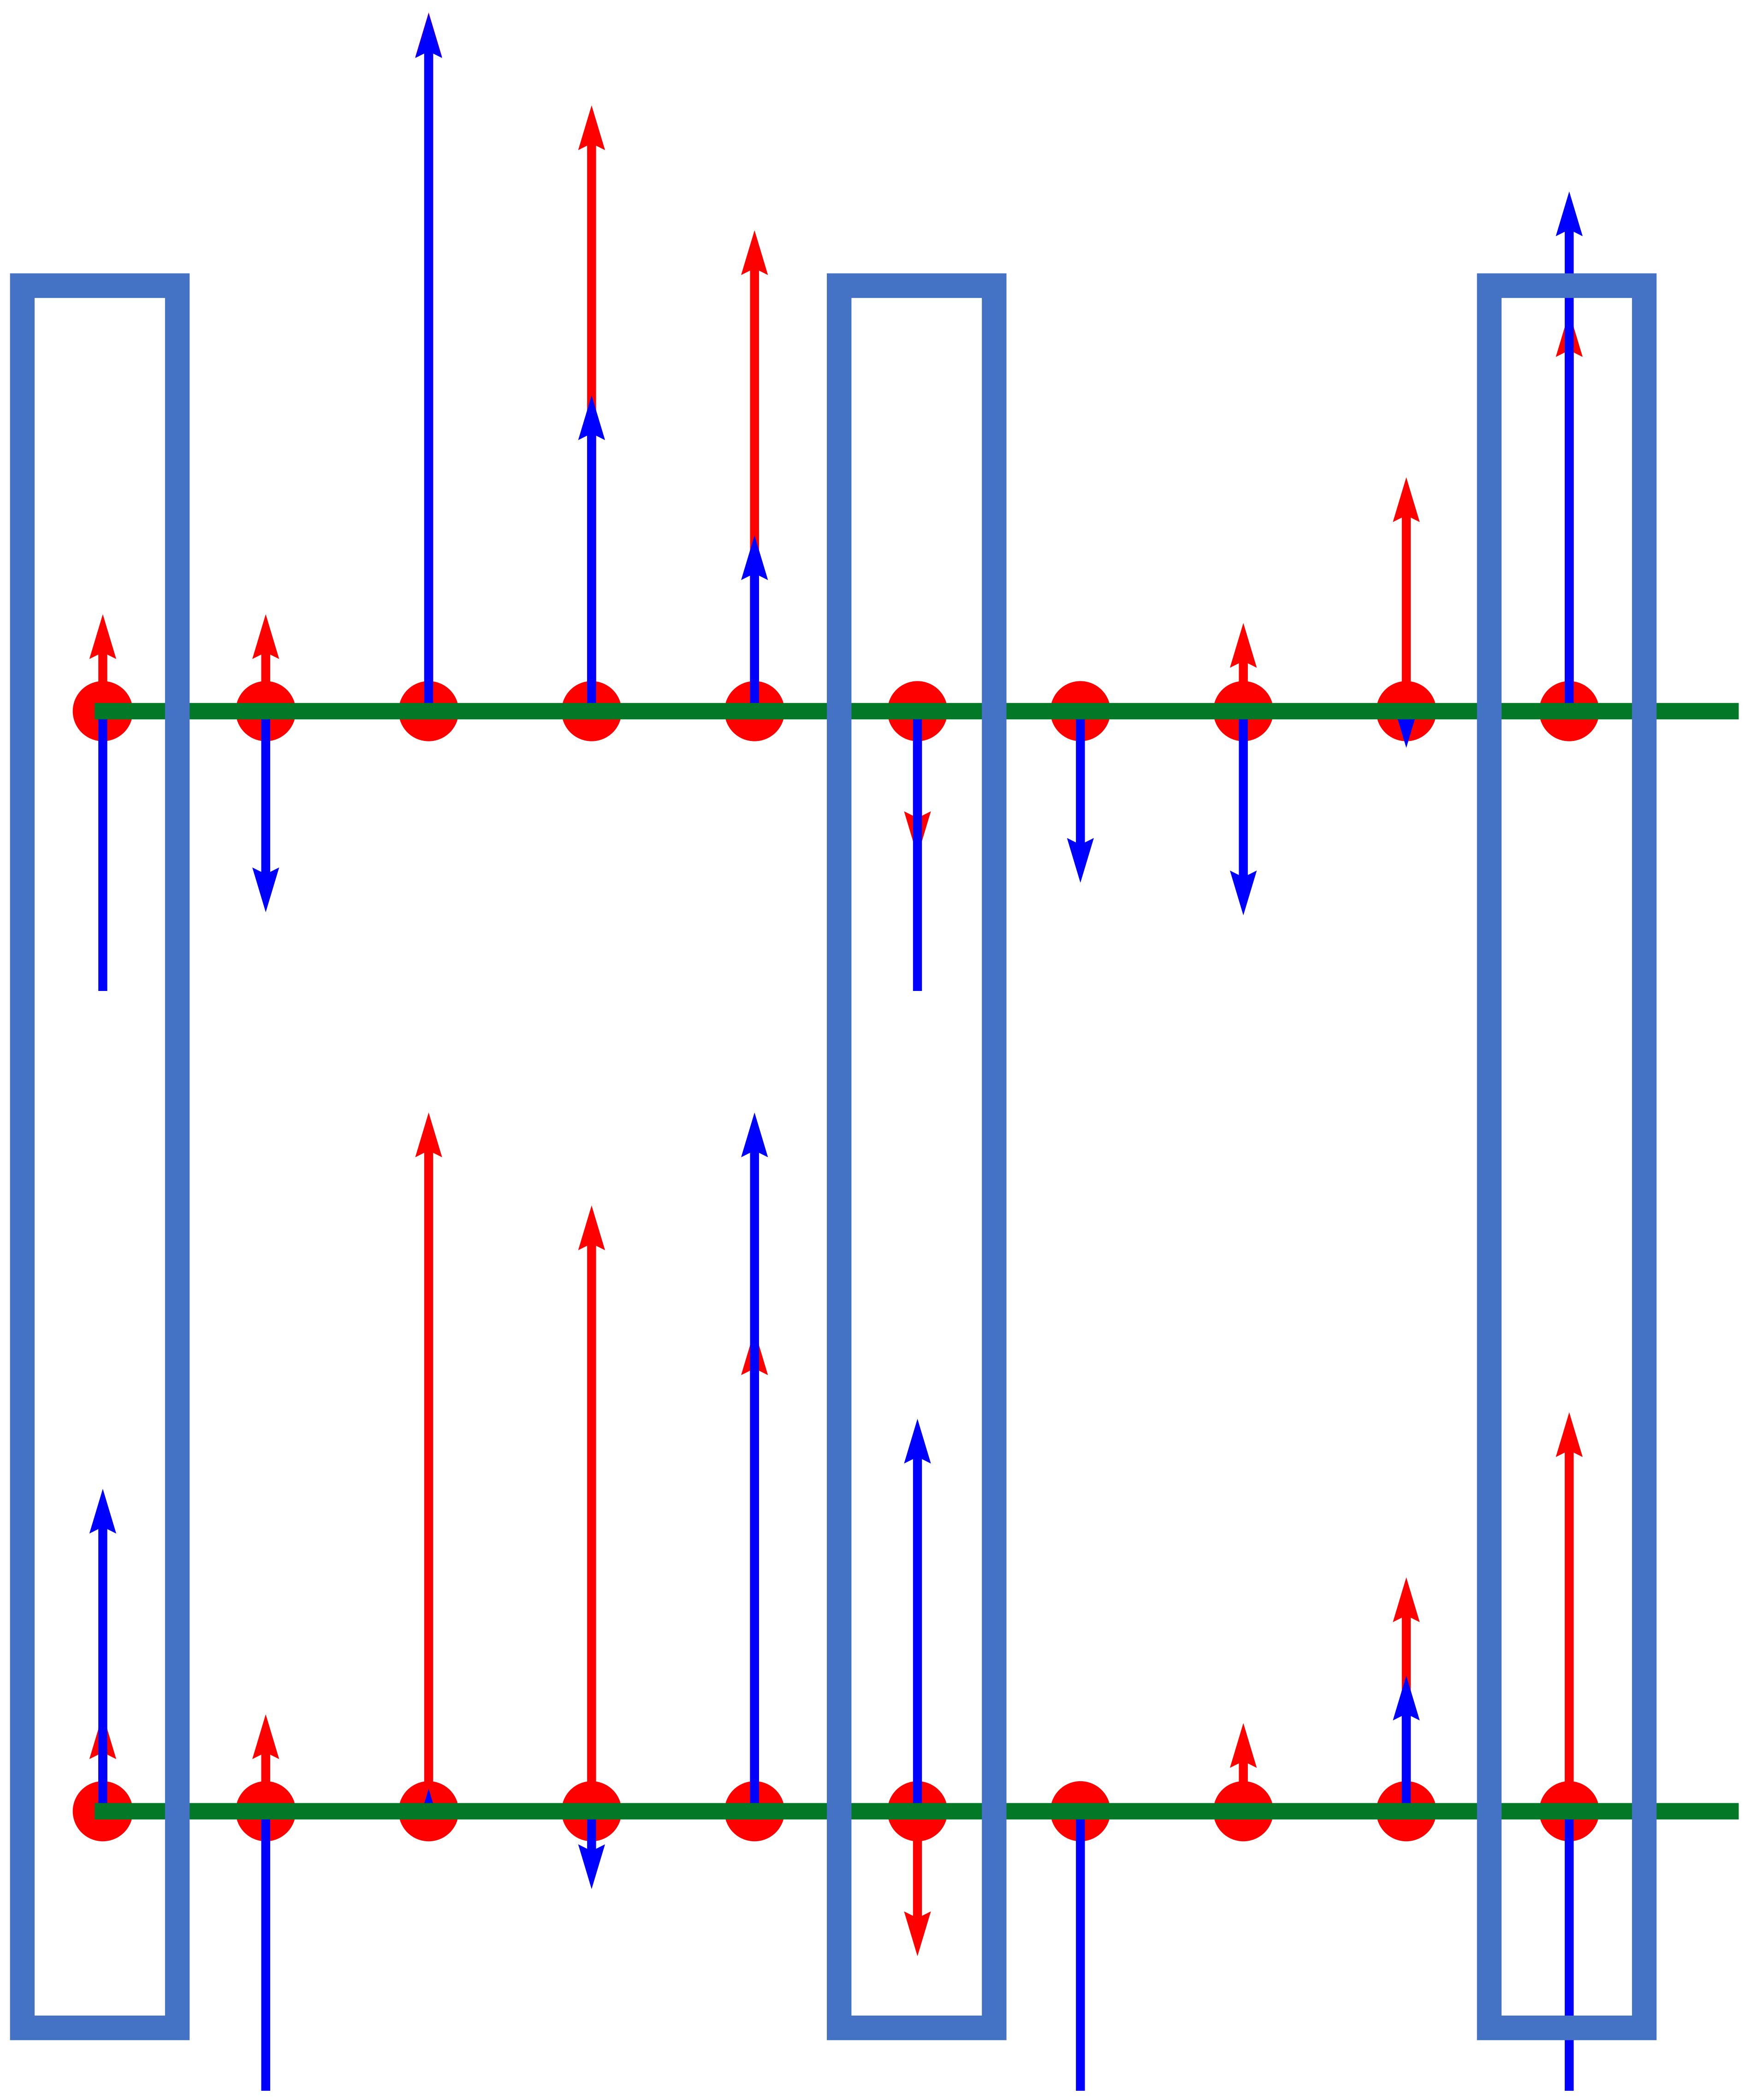

Edge-aware filters!

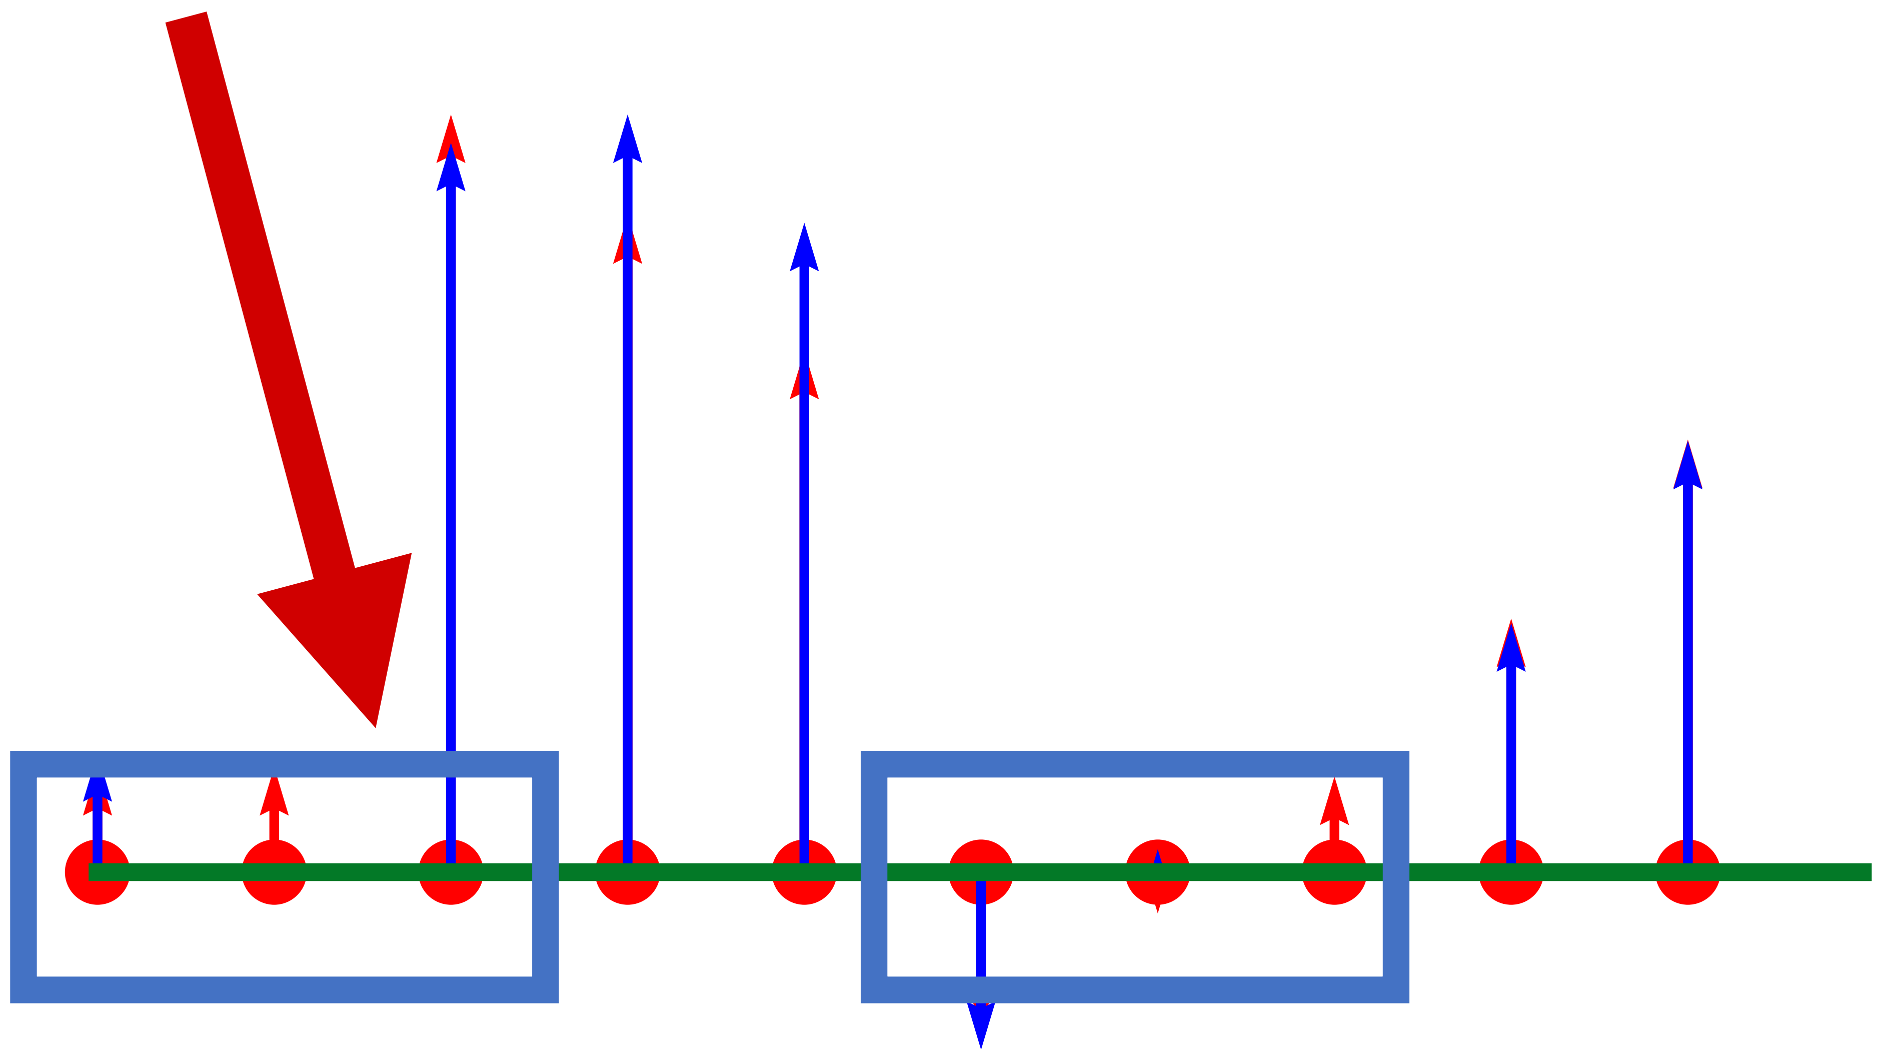

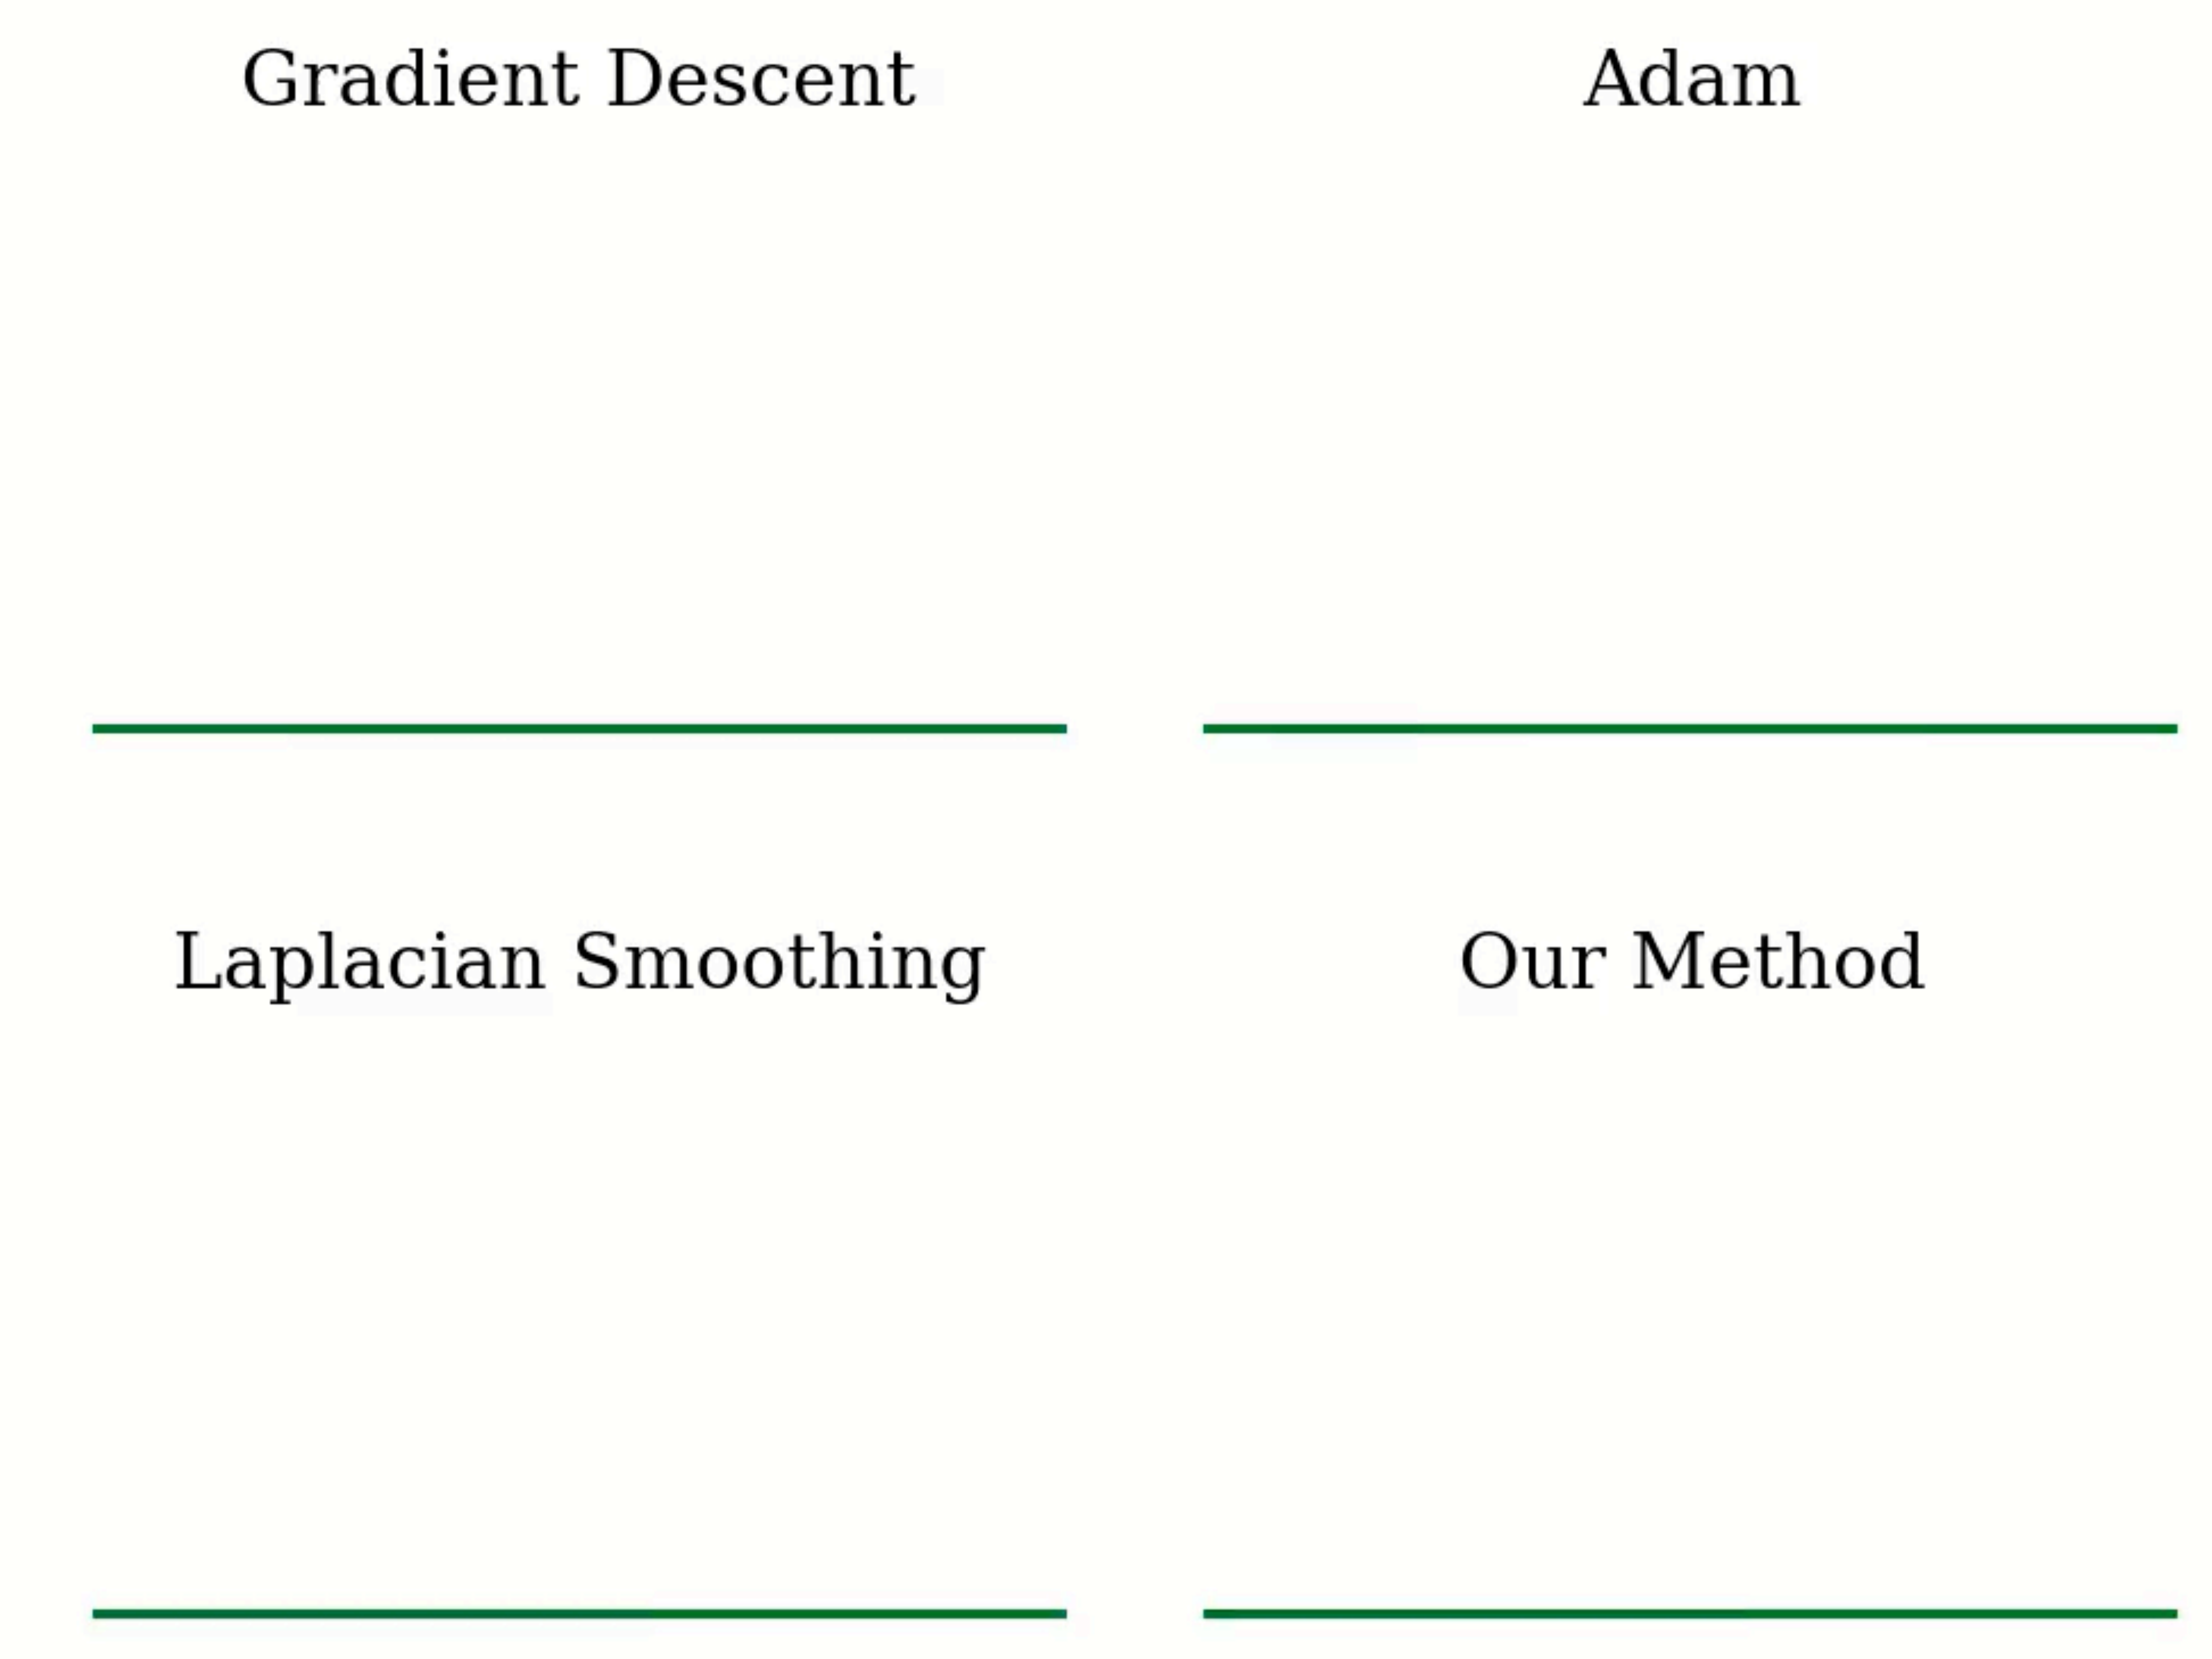
Gradient Descent
Adam
Laplacian Smoothing
Our Method

Implementation

Cross-bilateral filter for textures and volumes

- Less than 4% overhead

For mesh, generalizes Large Steps [Nicolet et al. 2021]

- Replace Laplacian smoothing with Generalized Bilateral Filter [Solomon et al. 2014]
- Less than 10% overhead

Results

Texture

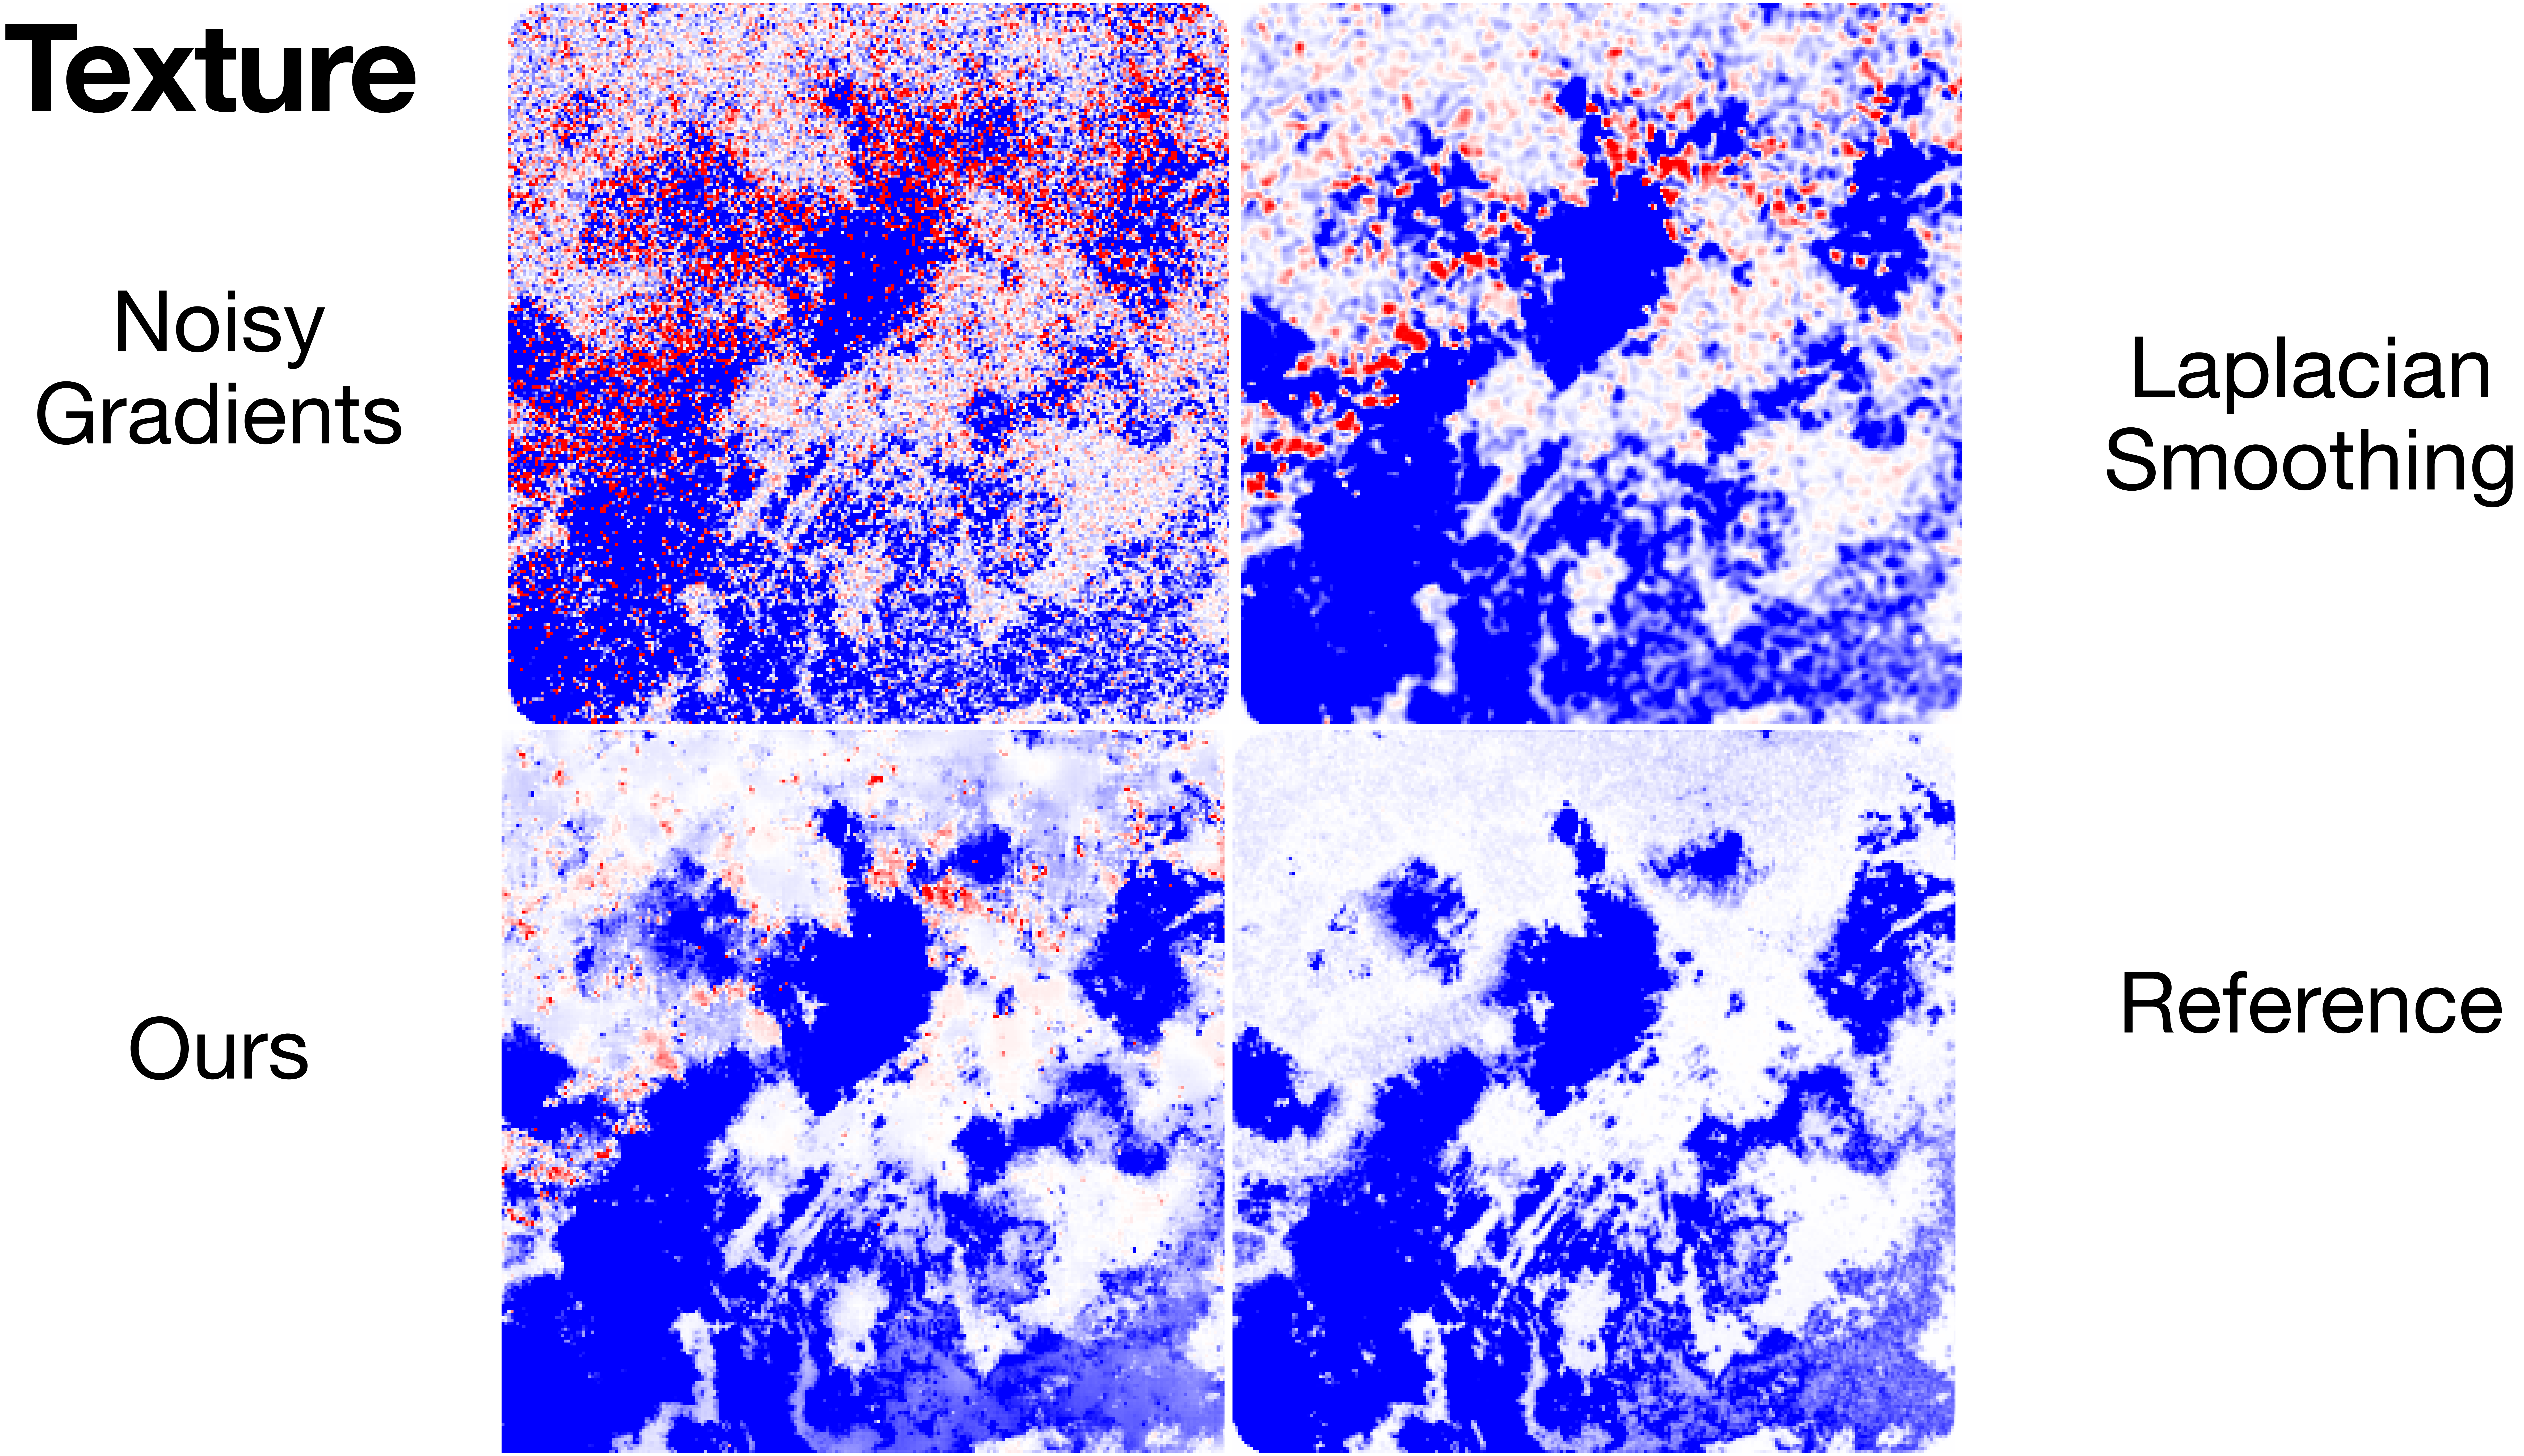

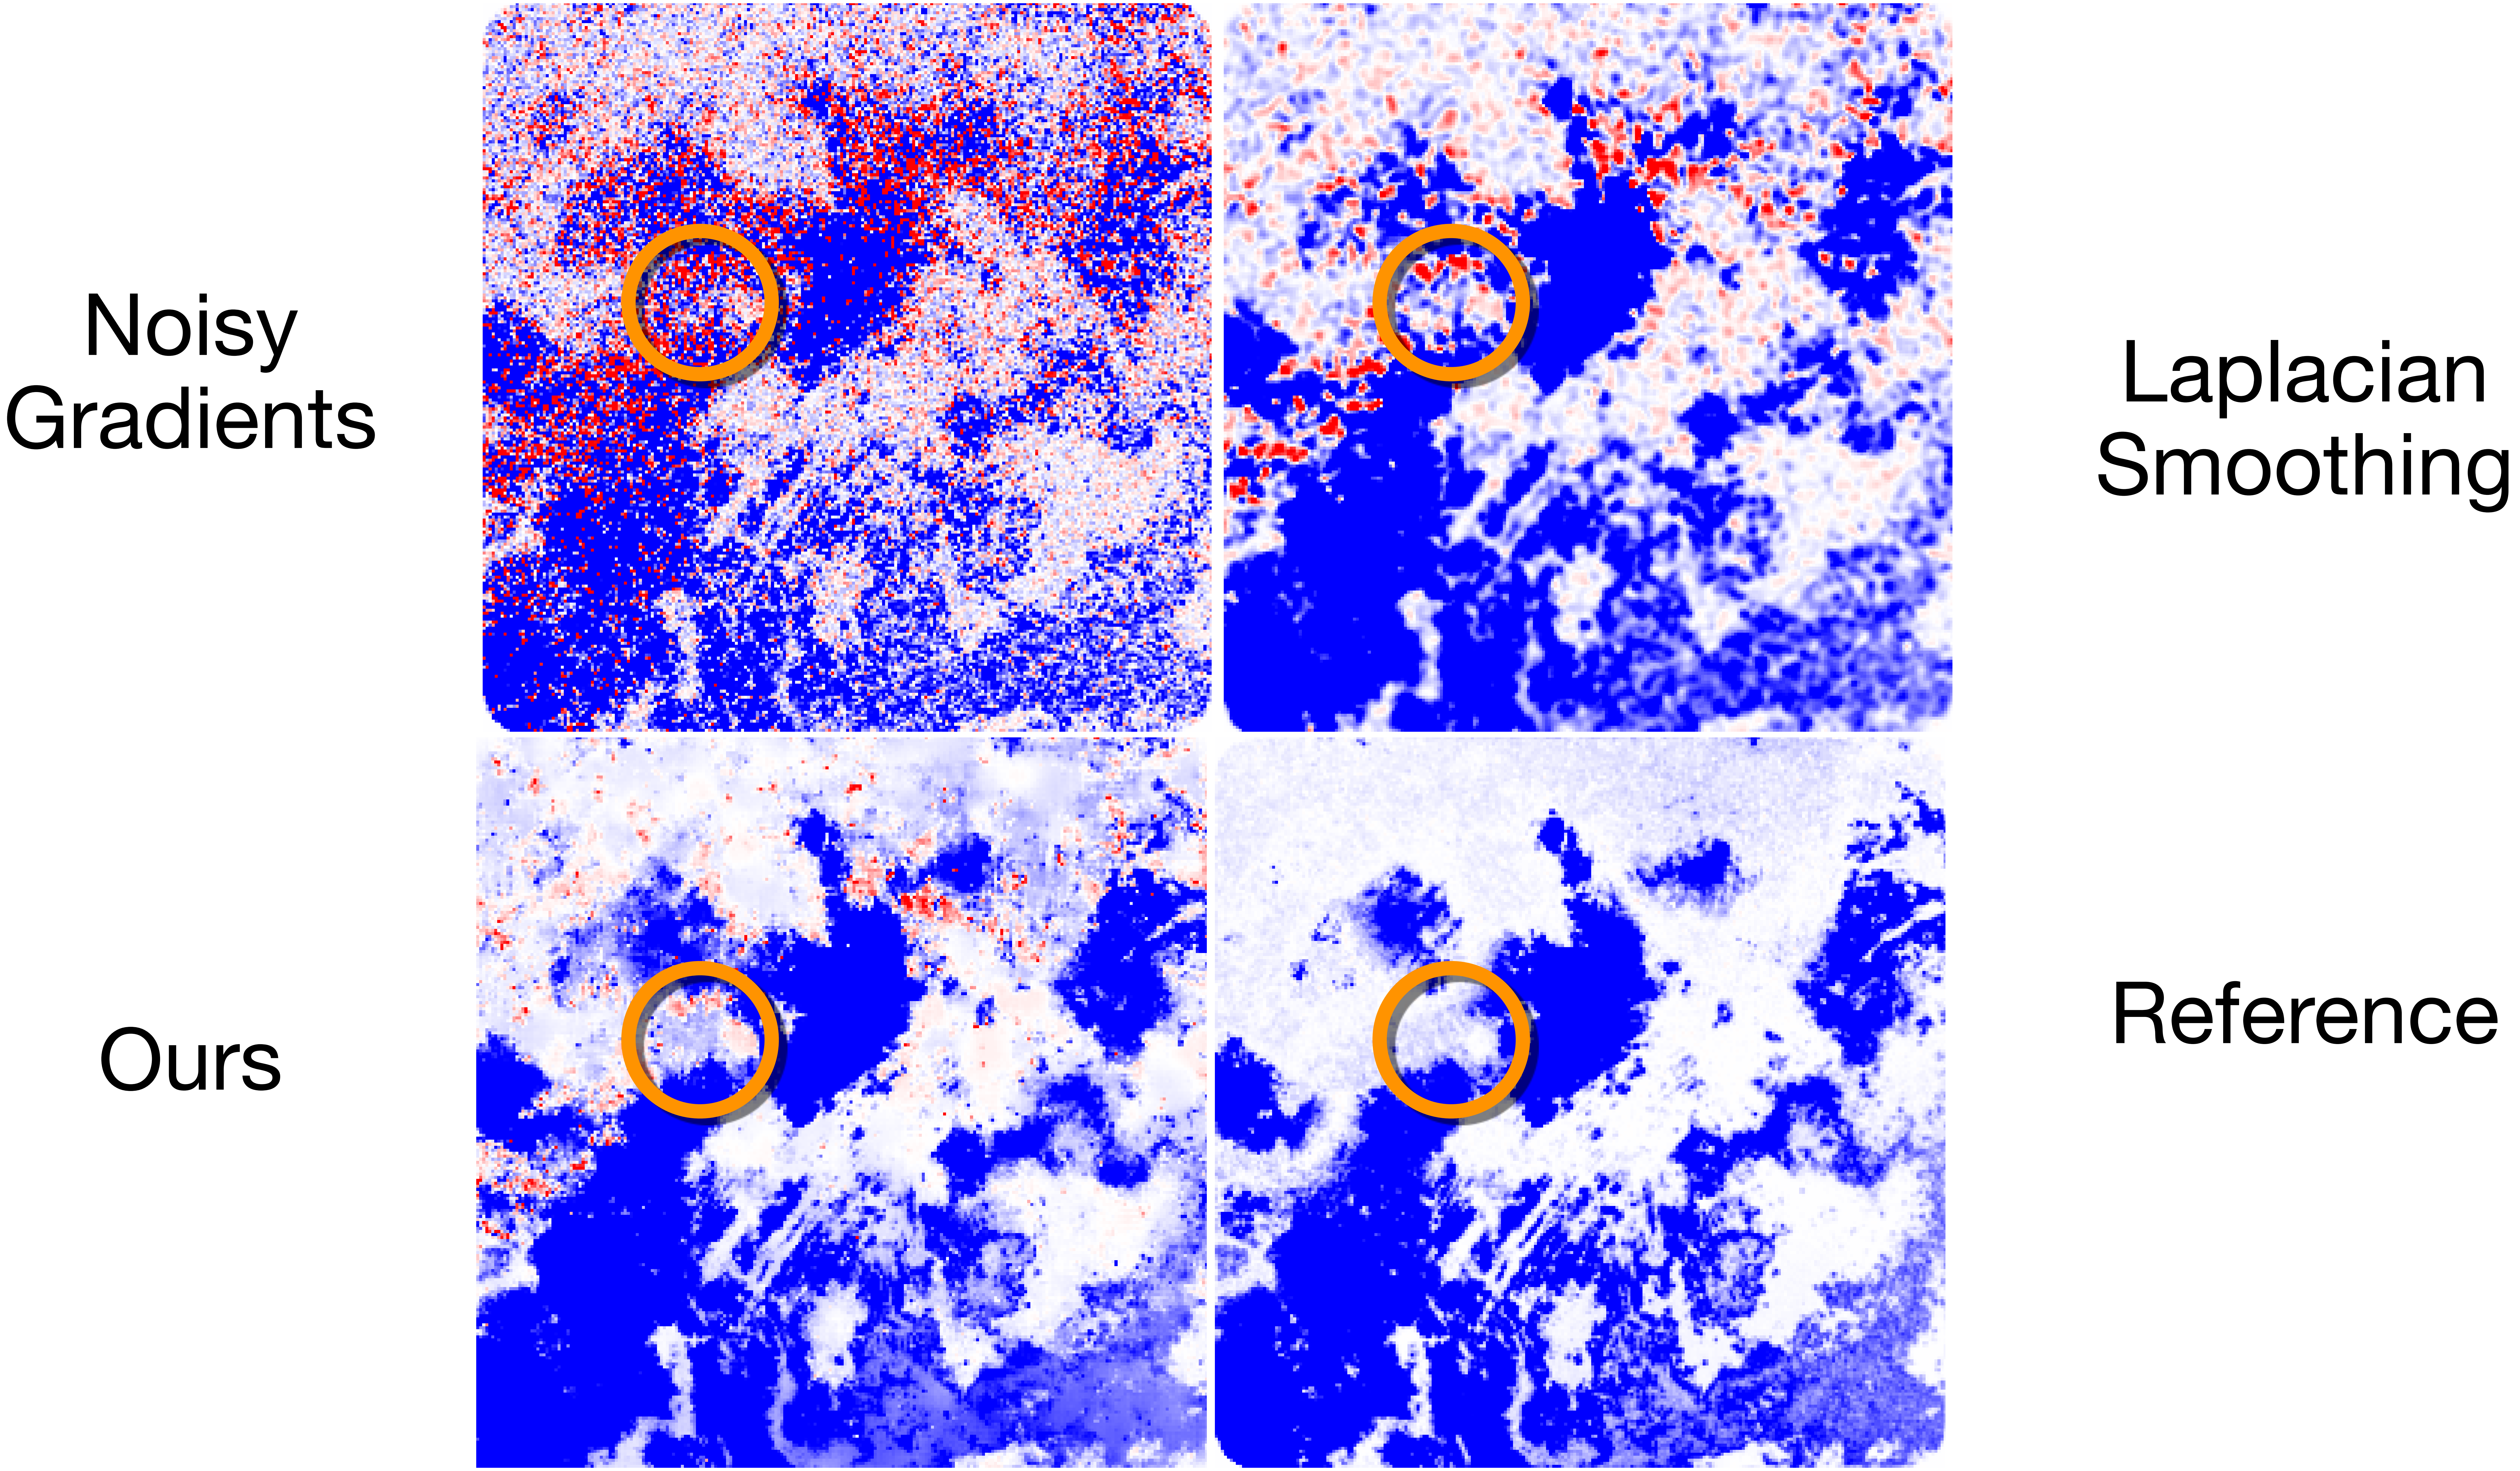
Noisy Gradients
Laplacian Smoothing
Ours
Reference

Volume Recovery

ADAM Ours

Mesh Recovery

Large Steps
[Nicolet et al 2021]

Ours

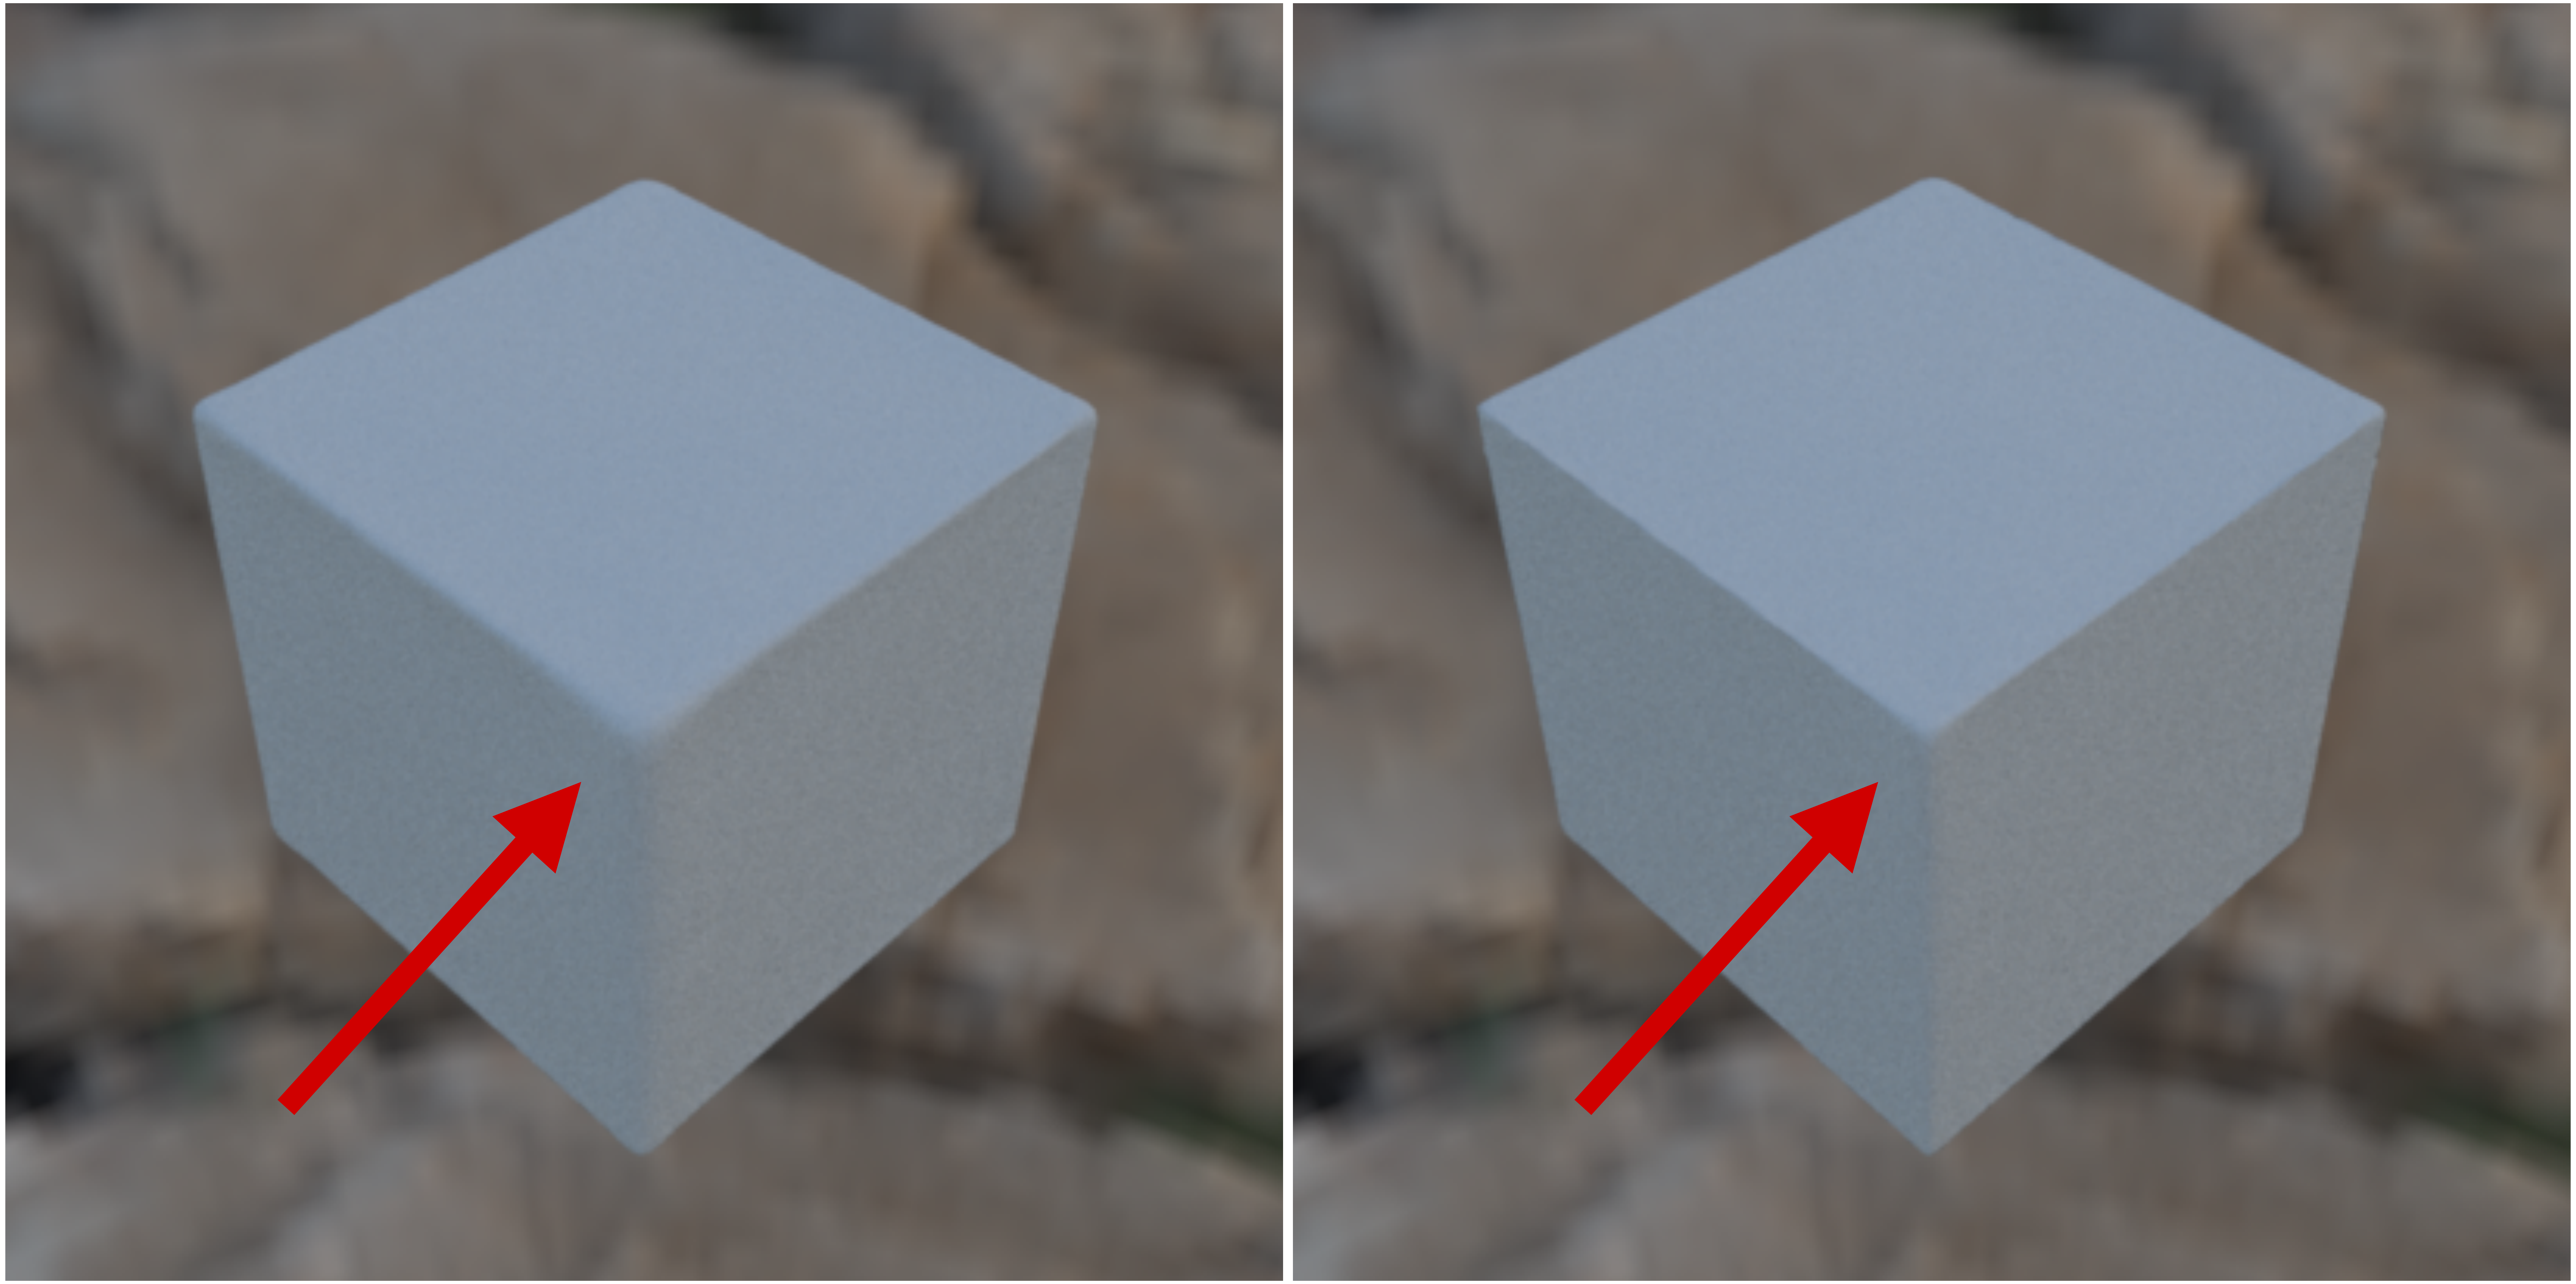

Conclusion

Takeaways

- Filters over **space**
- Filters over **time**
- Preserve **edges**

See the paper for

- Anisotropic preconditioning
- Pre-filtering vs Post-filtering
- Cross-bilateral filter vs Bilateral filter
- Performance Considerations
- And more!

Future Directions

More than inverse rendering:

- Inverse simulation
- Inverse PDE
- And more!

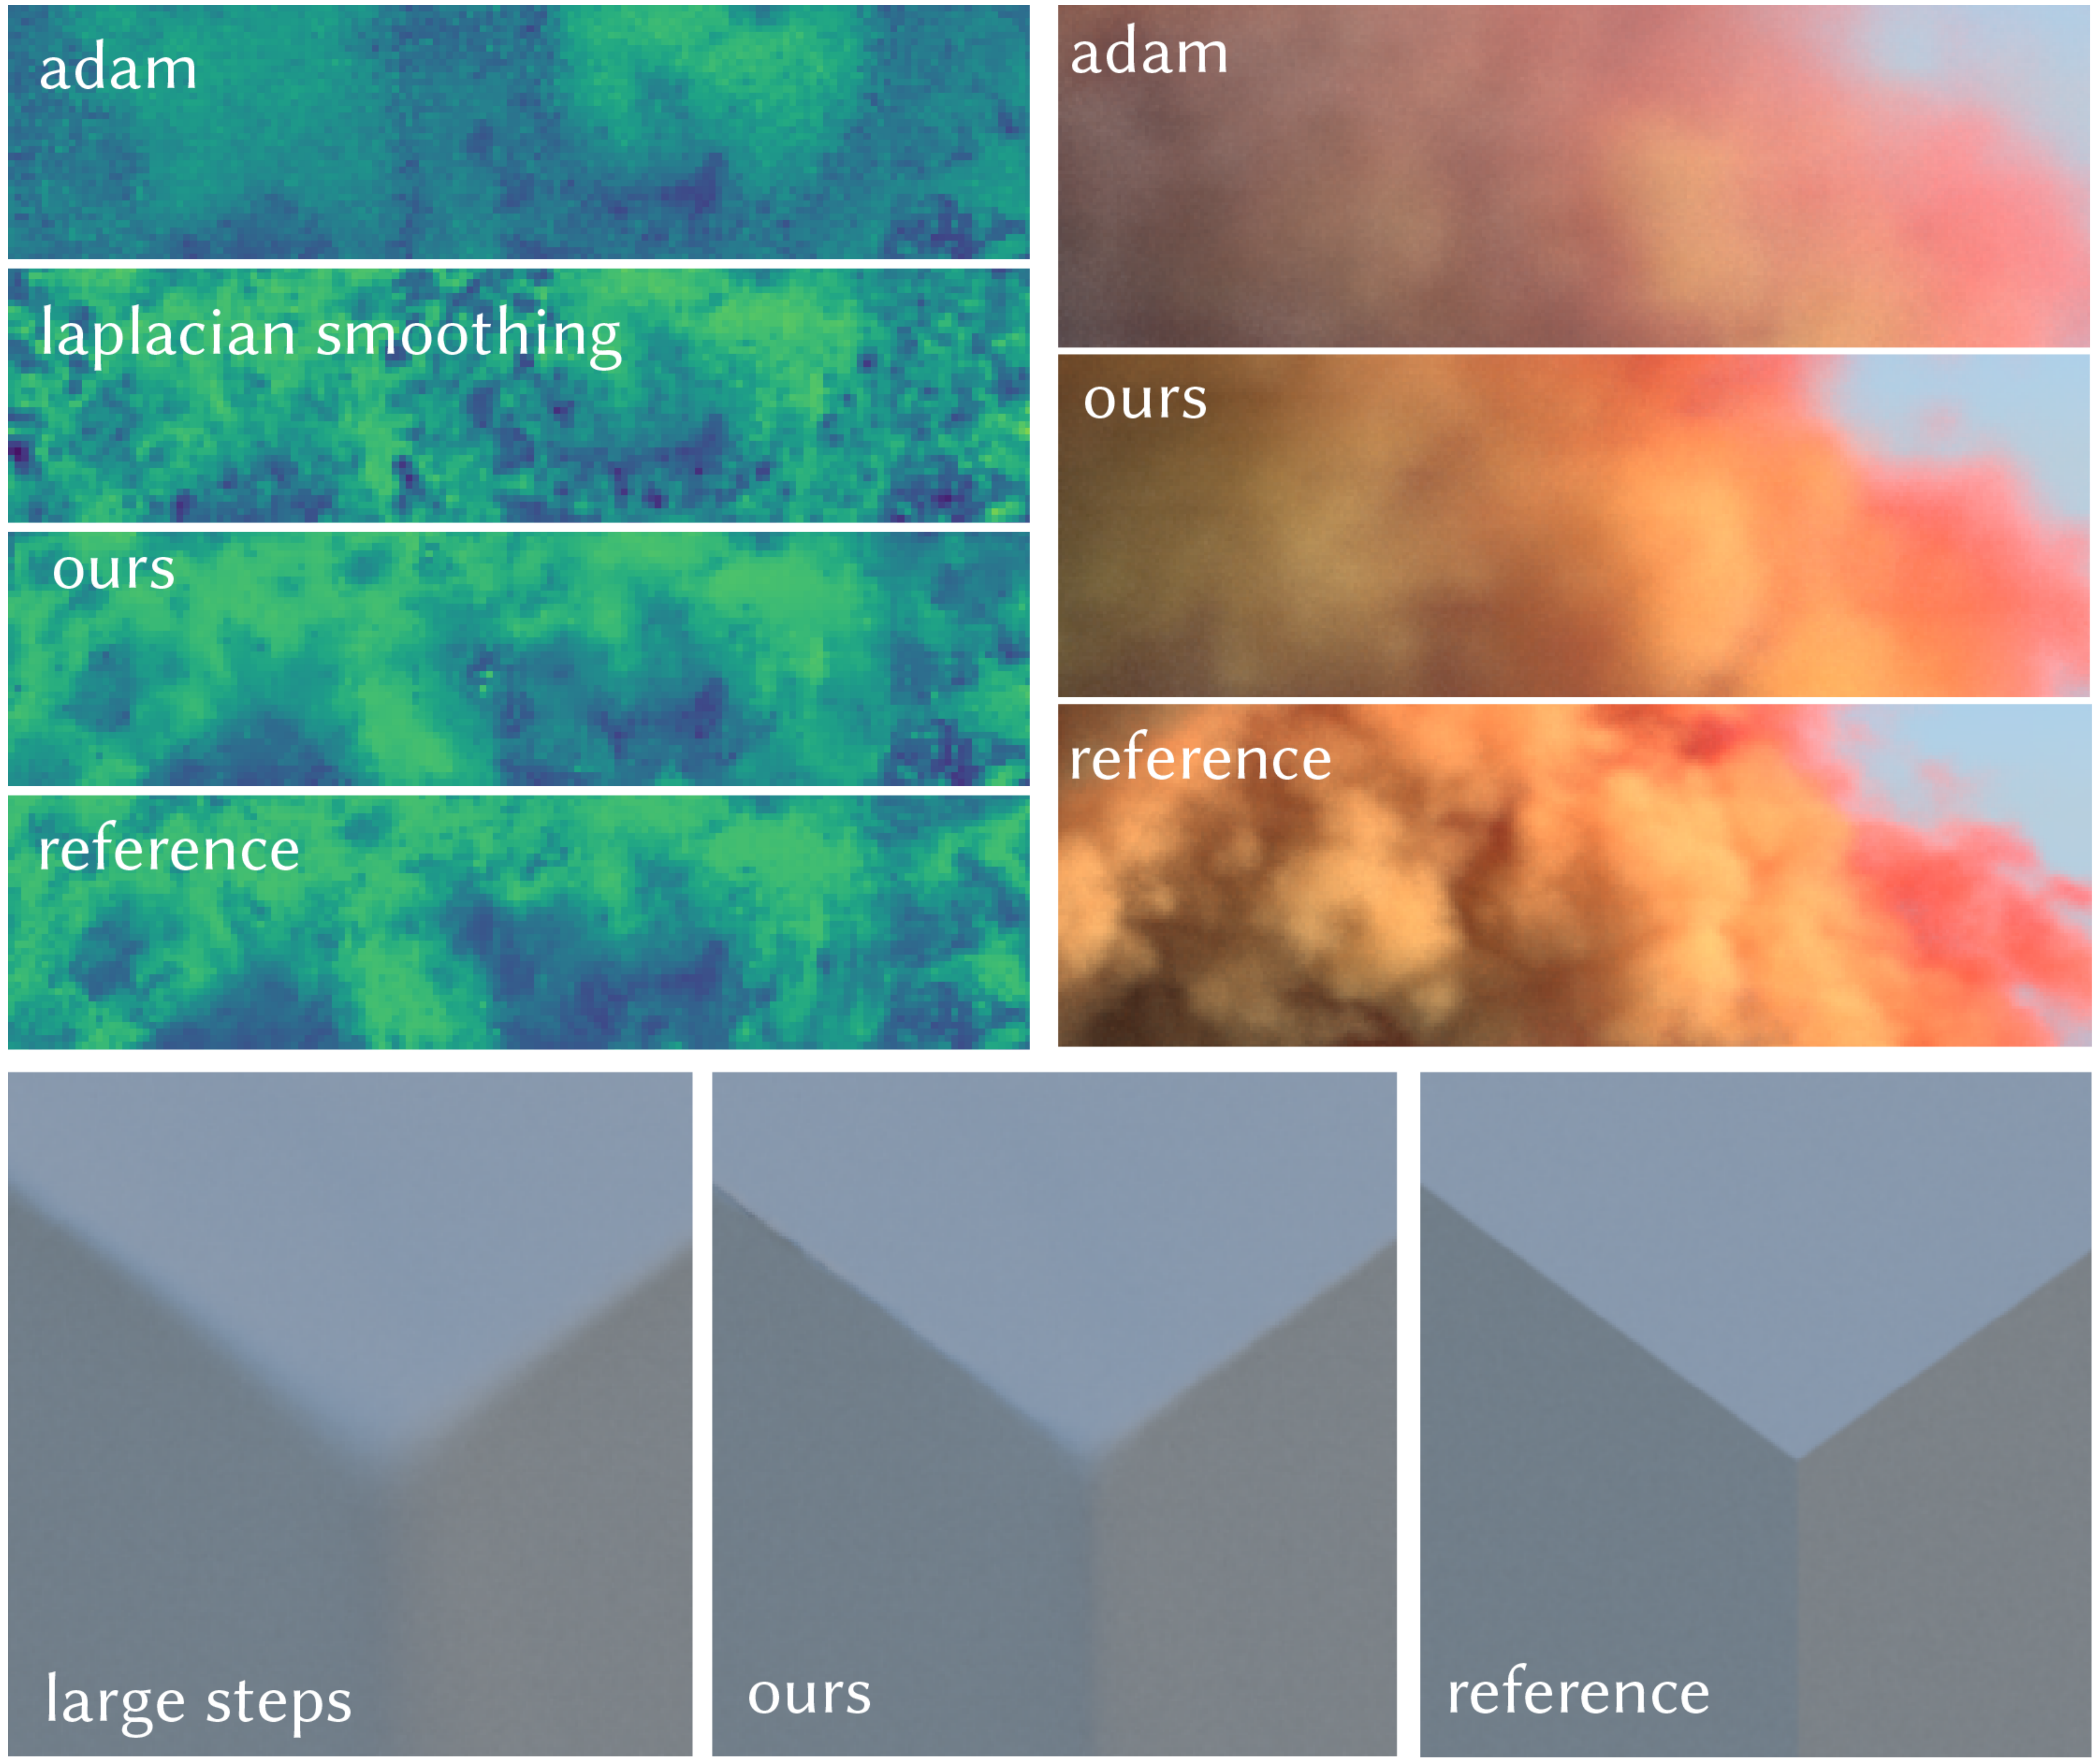

weschang.com/publications/stadam
Code available!

Wesley Chang, Xuanda Yang, Yash Belhe
{wec022, xuy008, ybelhe}@ucsd.edu