Quantifying the Noise Sensitivity of the Wasserstein Metric for Images

Erik Lager[1] Gilles Mordant[2] Amit Moscovich[1]

[1] Tel Aviv University [2] Yale University

Signed Wasserstein discrepancies preserve image geometry under heavy pixel-wise noise, **W_2 error grows like $\sqrt{\sigma}$, not the linear Euclidean rate.**

Summary

We study how Wasserstein metrics for images behave under pixel-wise additive noise, treating images as discrete measures on a fixed grid. Because noise can create negative pixel values, we work with a signed Wasserstein discrepancy. We prove in-expectation bounds under zero-mean Gaussian noise and show empirically that Wasserstein distances preserve image geometry in regimes where the Euclidean distance is dominated by noise.

Problem Setting

An image is a real-valued signal on an $n \times n$ grid G_n with $m = n^2$ pixels. Additive noise is added to two clean images μ, ν:

$$\mu_\varepsilon = \mu + \varepsilon_\mu, \qquad \nu_\varepsilon = \nu + \varepsilon_\nu.$$

The Euclidean (L_2) distance is local, it compares pixels independently and is quickly overwhelmed by noise. Wasserstein distances instead incorporate the spatial geometry of the grid through a ground cost, making them natural candidates for noisy imaging.

Signed Wasserstein Discrepancy

Additive Gaussian noise introduces negative pixel values. Using the Jordan decompositions:

$$\mu = \mu_+ - \mu_-, \qquad \nu = \nu_+ - \nu_-.$$

Define the signed discrepancy (Mainini 2012):

$$W_p^{\pm}(\mu, \nu) := W_p(\mu_+ + \nu_-, \nu_+ + \mu_-).$$

This way, signed, noise-corrupted images can be compared with standard optimal transport between two non-negative measures.

Zero-Sum Gaussian Noise Model

To isolate noise sensitivity from global mass-rescaling effects, the noise vector $N \sim \mathcal{N}(0, \Sigma)$ is zero-sum across the grid:

$$\Sigma_{ij} = \begin{cases} \sigma^2, & i = j \\ -\dfrac{\sigma^2}{m-1}, & i \neq j \end{cases}.$$

This preserves marginal variance σ^2 whose sum vanishes:

$$N_i \sim \mathcal{N}(0, \sigma^2), \qquad \sum_{i=1}^{m} N_i = 0.$$

Single Image results

For zero-sum noise $\varepsilon_1, \varepsilon_2$, $p > 1$ and $C_p = \frac{2^{\frac{1}{p}-1}}{\sqrt{\pi}} \Gamma(\frac{1}{2p} + \frac{1}{2})$:

$$C_p \cdot n^{\frac{2}{p}-1} \cdot \sigma^{1/p} \leq \mathbb{E}\left[\left(W_p^{\pm}(\mu + \varepsilon_1, \mu + \varepsilon_2)\right)\right] \leq \left(\frac{4\sqrt{2}}{\sqrt{\pi}} n\sigma\right)^{1/p}.$$

The Wasserstein-2 error grows sublinearly in σ while the Euclidean distance grows linearly.

Two image results

Similarly, for zero-sum noise $\varepsilon_\mu, \varepsilon_\nu$, $p > 1$ and two measures μ, ν:

$$\mathbb{E}[W_p^{\pm}(\mu + \varepsilon_\mu, \nu + \varepsilon_\nu)] \leq \left(\frac{\sqrt{2}}{2}\right)^{1-\frac{1}{p}} \cdot W_1(\mu, \nu)^{1/p} + \frac{\sqrt{2}}{2}\left(\frac{4}{\sqrt{\pi}} n \log_2(n) + \frac{1}{\sqrt{\pi}} n\right)^{1/p} \sigma^{1/p}.$$

W_1 Unique Bound

For two distinct clean images, the noisy signed W_1 discrepancy stays close to the clean one:

$$\mathbb{E}[W_1^{\pm}(\mu_\varepsilon, \nu_\varepsilon) - W_1^{\pm}(\mu, \nu)] \leq \frac{4n\log_2 n + n}{\sqrt{\pi}}\sigma + \frac{\sqrt{2}}{n}.$$

The inter-image geometry is preserved under controlled additive noise.

Inter-Image Robustness

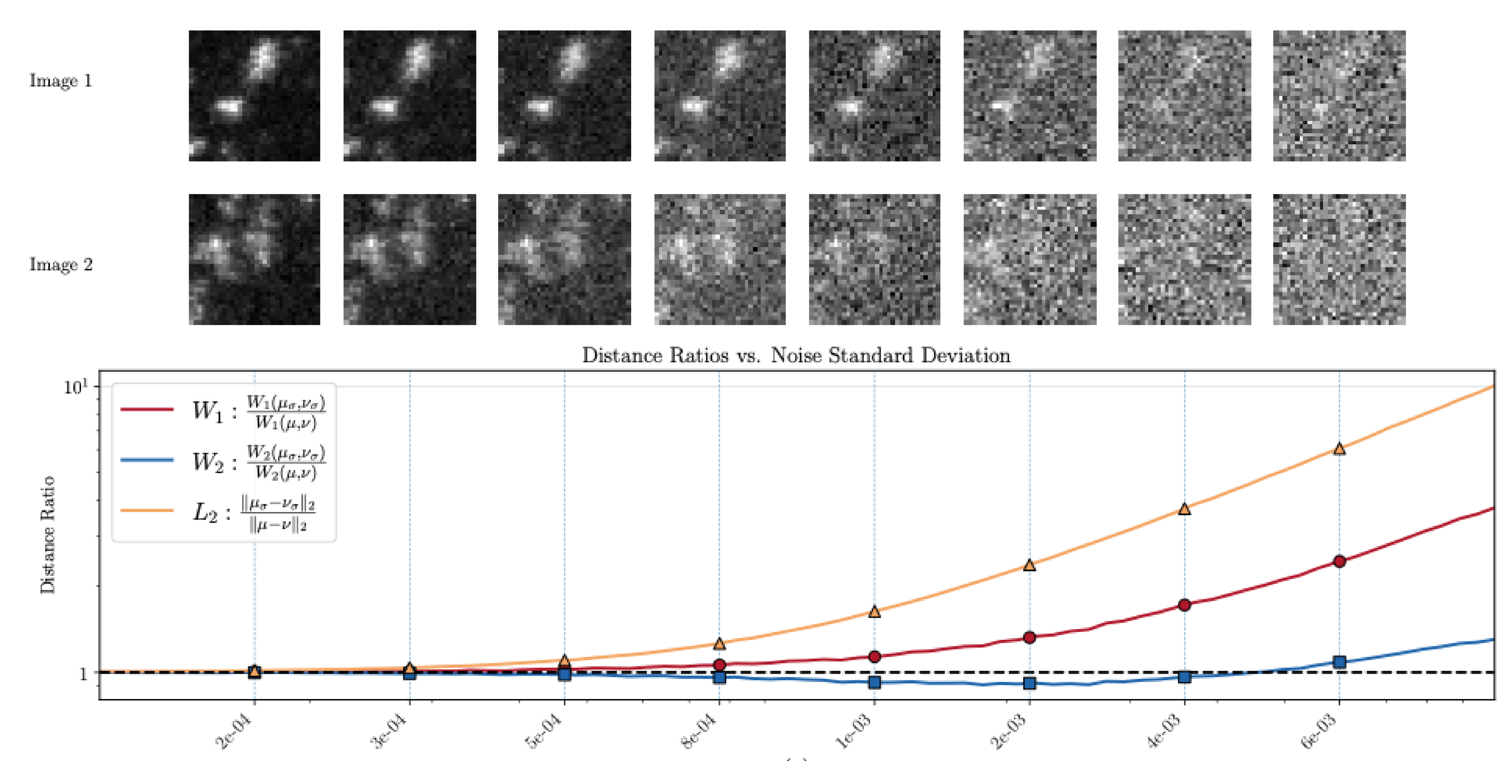

Fig. 1: *Distance ratios for a noisy image pair vs. noise level. A ratio near 1 means the noisy distance still reflects the clean-image distance.*

Cryo-EM Case Study

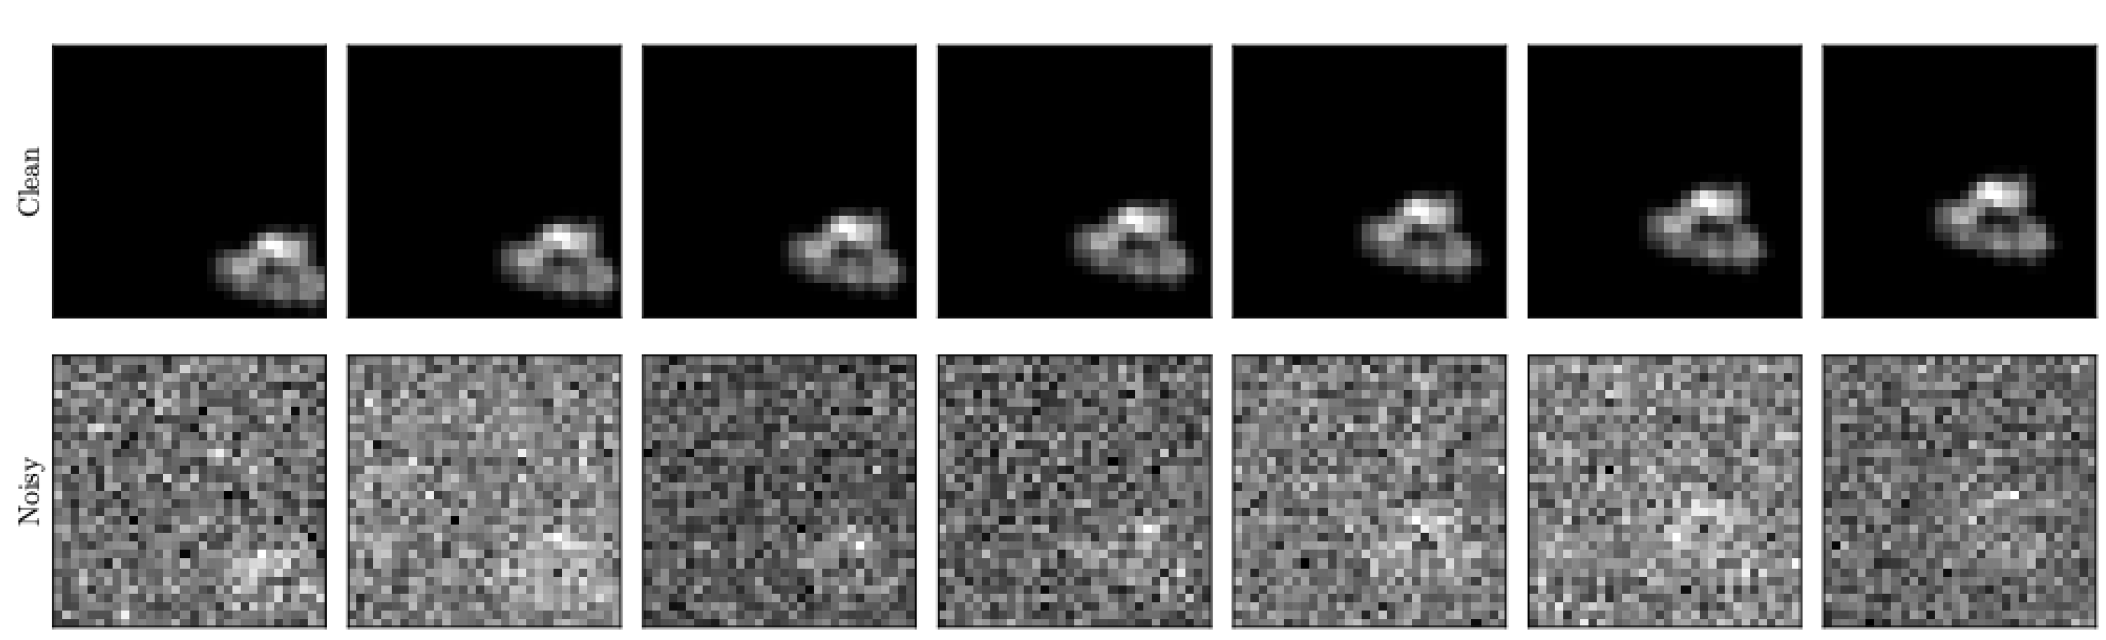

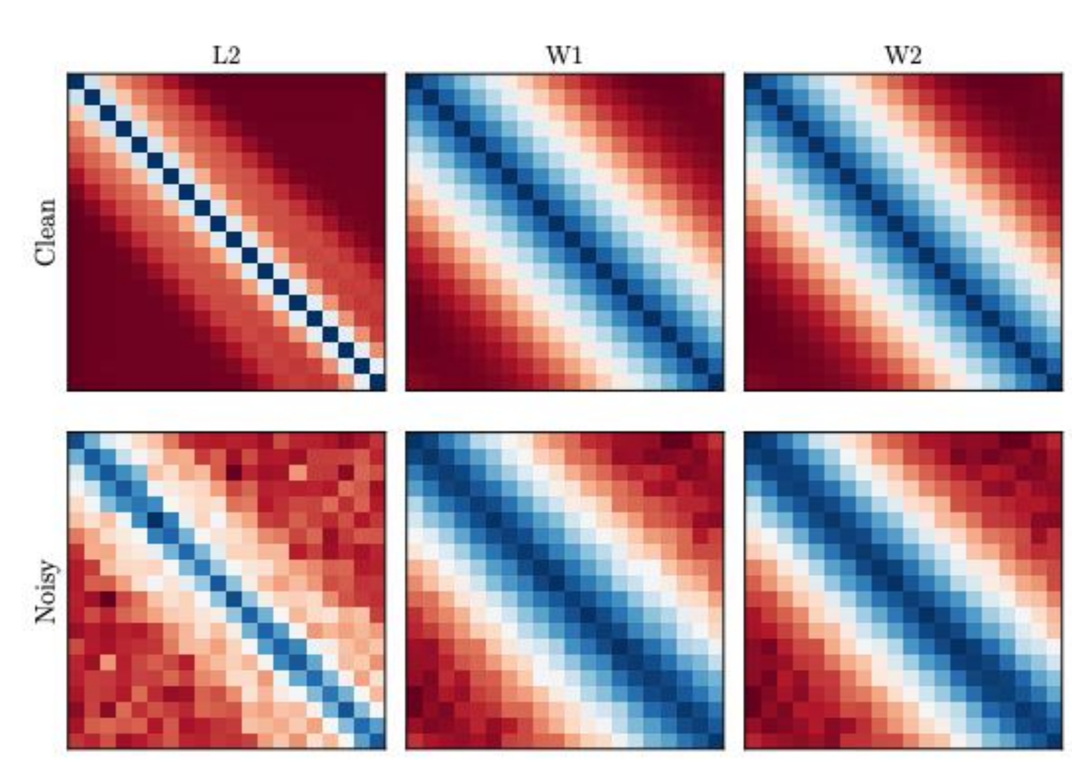

Figs 6–7: *Sample of clean vs. noisy E. coli Hsp90 projections (left) and pairwise distance matrices (right).* Under heavy noise L_2 loses the clean diagonal structure, while W_1 and W_2 preserve the global geometry

Key Takeaways

- Finite-sample expectation bounds for signed Wasserstein discrepancies under pixel-wise Gaussian noise.
- For p = 2 the noise-induced error scales as $\sqrt{\sigma}$, sublinear in σ.
- Experiments confirm the scaling on microscopy images.

Code & figures
github.com/warik21/Quantifying-wasserstein-noise-sensitivity

References

- Erik Lager, Gilles Mordant, Amit Moscovich. “Quantifying the noise sensitivity of the Wasserstein metric for images”. ICML (2026).
- Edoardo Mainini. “A description of transport cost for signed measures”. Journal of Mathematical Sciences (2012).
- Jonathan Weed, Francis Bach. “Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance”. Bernoulli (2019).