


Physical Alignment

A Geometric Diagnostic for Reliable
Scientific Representations

Aleix Seguí and Wes Armour


Scientific Machine Learning Conference 2026

Bath, United Kingdom • 15 September 2026



Scientific analysis depends on derived quantities

A scientific field is compressed, reconstructed, and then used to compute an observable.




Which physical observables can be preserved under a bitrate constraint?


Examples of scientific observables

Fluid dynamics

PDEBench (Takamoto et al., 2022)




Velocity




Vorticity

Cosmology

Nyx (Almgren et al., 2013)




Density




Baryon fraction

Electron microscopy

MICrONS / minnie65
(MICrONS Consortium, 2025)



Image



Contour
features

$$Q(\mathbf{v}) = \partial_y v_x - \partial_x v_y$$


$$Q_{\text{bar}}(\rho_b, \rho_{\text{dm}}) = \frac{\rho_b}{\rho_b + \rho_{\text{dm}} + \varepsilon}$$

$$Q_{\text{cont}}(I) = \sqrt{(w_x \partial_x I)^2 + (w_y \partial_y I)^2}$$

The choice of Q specifies the scientific information we ask the representation to retain.

Velocity and vorticity errors at the same bitrate

Example: two models trained at 0.85 bits per sample, with different physics weights.



The physics loss reallocates error between the velocity field and its vorticity.

A local Gaussian approximation to quantisation


Around each code $z = f_\phi(x)$, we model compression as a small random perturbation.

$$\hat{Z} = Z + \eta, \quad \eta \mid X = x \sim \mathcal{N}(0, \Sigma)$$

$$\mathbb{E}[\eta \mid X = x] = 0, \quad \mathbb{E}[\eta\eta^\top \mid X = x] = \Sigma$$

The covariance Σ describes how compression uncertainty is distributed across latent directions.

This is a local analytical assumption about the channel.



Long and short axes correspond to larger and smaller noise variances.

The local rate proxy depends on the covariance

The entropy model p_ψ supplies a variational rate proxy for the noisy latent code.


$$R_{\text{var}} = \mathbb{E}_X \left[\text{KL} \left(q_\Sigma(\hat{Z} \mid X) \parallel p_\psi(\hat{Z}) \right) \right]$$

A local expansion gives the Σ -dependent terms:

$$R_{\text{sur}}(\Sigma; z) = \underbrace{\frac{1}{2} \text{Tr}(H_R \Sigma)}_{\text{variance placement}} - \underbrace{\frac{1}{2} \log \det \Sigma}_{\text{uncertainty volume}}$$

$H_R \geq 0$ is the local entropy-model curvature.
Uncertainty along a high-curvature direction has a larger trace cost.

The surrogate omits additive constants and higher-order terms in the local expansion.



Both ellipses have the same uncertainty volume. Their orientation changes the rate cost.


Propagation of small variations through the decoder

We linearise the decoder and observable at each sample to study error propagation.

Latent perturbation

Field error

Observable error



$$\delta x \approx J_g(z) \eta$$

$$\delta Q \approx J_Q(x) J_g(z) \eta$$

We study uncertainty-induced error; deterministic reconstruction bias is not considered here.

Three local metrics describe rate and distortion

A rate-distortion objective, represented in latent space, induces three sensitivity factors.

Rate

$$H_R(z)$$

Entropy-model curvature controls the rate cost of a perturbation.

$$d_X(\Sigma) = \text{Tr}(G_{\text{eff}}\Sigma)$$

Field fidelity

$$G_{\text{eff}} = J_g^\top H_D J_g$$

H_D weights field errors; J_g maps them back to latent directions.

Physical observable

$$W_{\text{eff}} = J_g^\top J_Q^\top H_\ell J_Q J_g$$

H_ℓ weights observable errors after the composed map $J_Q J_g$.

$$d_Q(\Sigma) = \text{Tr}(W_{\text{eff}}\Sigma)$$

A tangent-space rate–distortion law

At a fixed sample, choose the noise covariance subject to both distortion budgets.

Theorem 1

Optimal covariance

$$\begin{aligned} \min_{\Sigma > 0} \quad & \text{Tr}(H_R \Sigma) - \log \det \Sigma \\ \text{subject to} \quad & d_Q(\Sigma) \leq \varepsilon_Q, \\ & d_X(\Sigma) \leq \varepsilon_X. \end{aligned}$$

$$(\Sigma^\star)^{-1} = H_R + \alpha W_{\text{eff}} + \gamma G_{\text{eff}}$$

The multipliers $\alpha, \gamma \geq 0$ reflect the observable and field-error budgets.

Directions with larger combined curvature receive less compression noise.

Comparing sensitivities at the same local rate cost


For $H_R > 0$, rate whitening removes the anisotropy introduced by the entropy model.

$$\tilde{W} = H_R^{-1/2} W_{\text{eff}} H_R^{-1/2},$$

$$\tilde{G} = H_R^{-1/2} G_{\text{eff}} H_R^{-1/2}.$$


These metrics measure observable and field sensitivity per unit of local rate cost.

We can now compare the directions in which each objective needs precision.



Shared eigenbasis

Field sensitivity




Rotated eigenbases

Observable sensitivity

The fixed-rate trade-off between observable and field error

Matched-rate experiments



Each curve has its own fixed bitrate.

Theorem 2: a local no-free-lunch result

At a fixed rate budget ρ , let Σ_{Phys} and Σ_{MSE} be the unique optimal allocations for the two losses.

$$H_R > 0, \quad [\tilde{W}, \tilde{G}] \neq 0$$

$$d_Q(\Sigma_{\text{Phys}}) < d_Q(\Sigma_{\text{MSE}})$$

$$\implies d_X(\Sigma_{\text{Phys}}) > d_X(\Sigma_{\text{MSE}}).$$

Better observable fidelity at fixed bitrate comes at a cost in field fidelity.

Measuring physical alignment with principal angles


We propose a diagnostic that compares the principal subspaces of $\tilde{W}$ and $\tilde{G}$.

- 1 Let $U_{W,k}$ and $U_{G,k}$ contain the top- k orthonormal eigenvectors of the two metrics.
- 2 Measure how much these two subspaces overlap:

$$\text{Align}_k(x) = \frac{1}{k} \left\| U_{W,k}^\top U_{G,k} \right\|_F^2 = \frac{1}{k} \sum_{i=1}^k \cos^2 \theta_i$$

θ_i are the principal angles between the two subspaces.

- 3 Average over samples; report k and how much sensitivity the selected modes capture (trace coverage).



$$k = 1 : \quad \text{Align}_1 = \cos^2 \theta$$

1: coincident subspaces.
0: orthogonal subspaces.

Training with a physics-weighted rate–distortion loss

The objective used for training

$$\begin{aligned} \min_{\phi, \theta, \psi} \mathcal{L}_{\text{train}} = & \mathbb{E}_X [R_{\psi}(\hat{Z}) \\ & + \lambda(1 - \beta)D_X(X, \hat{X}) \\ & + \lambda\beta D_Q(Q(X), Q(\hat{X}))]. \end{aligned}$$


$\beta = 0$: field MSE only.

$\beta = 1$: observable loss only.

λ sets the rate–distortion balance.


Comparisons are made at matched target bitrates.

Vorticity: alignment at different ranks



The overall increase is seen for $k = 8, 16, 32$.

Alignment helps explain empirical trade-offs



Coarse structure

Higher alignment is associated with small changes in field fidelity.


Contour features

Lower alignment is associated with a sharper field-observable trade-off.

Each point is an operating point; marker size indicates rate.

Concentrated sensitivity can reduce the required rate

Only directions that strongly affect the observable need tight control (lower perturbation error).



Physics-aware models are more efficient in a rate-distortion sense.

Conclusions

1 **A local allocation law**

Rate, field fidelity and observable sensitivity determine the optimal noise covariance.

2 **A fixed-rate trade-off**

Under the theorem's assumptions, improving physical fidelity requires more pointwise field error.

3 **A computable diagnostic**

Dominant-eigenspace overlap helps explain the trade-offs measured in trained compressors.

Acknowledgements

Fellowship support

“la Caixa” Foundation

Funder ID 100010434 • Fellowship code LCF/BQ/PFA25/110000

Travel support

Linacre College and the Department of Engineering Science

University of Oxford

Thanks to Ryan Cherian for his feedback and encouragement.

Slide preparation

ChatGPT assisted with slide design and LaTeX formatting.