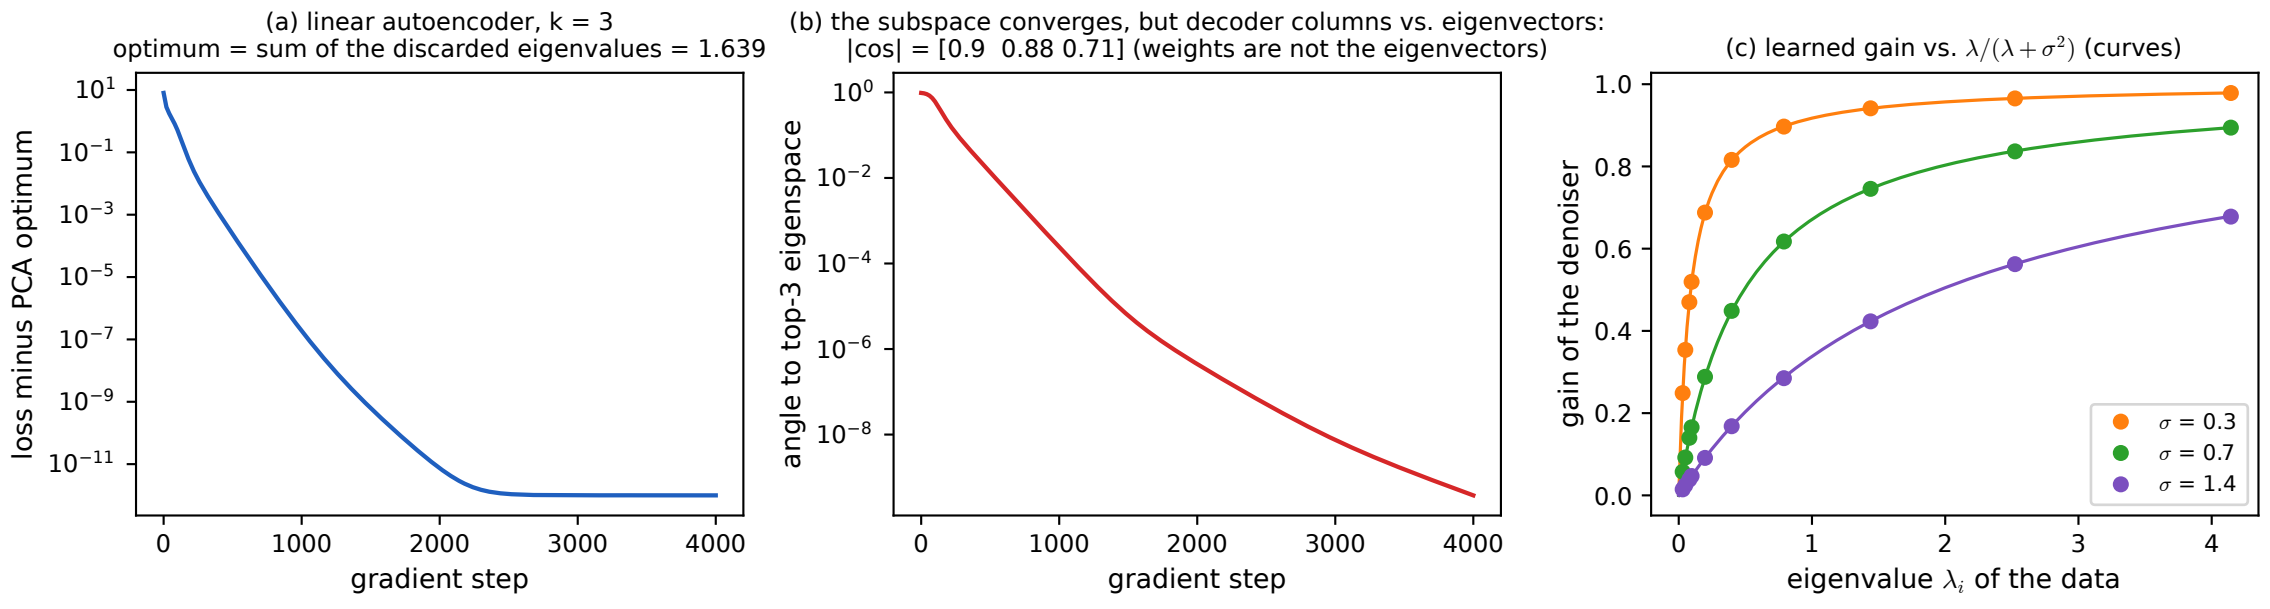

(a) linear autoencoder, k = 3
optimum = sum of the discarded eigenvalues = 1.639
loss minus PCA optimum
gradient step
(b) the subspace converges, but decoder columns vs. eigenvectors:
|cos| = [0.9 0.88 0.71] (weights are not the eigenvectors)
angle to top-3 eigenspace
gradient step
(c) learned gain vs. $\lambda/(\lambda+\sigma^2)$ (curves)
gain of the denoiser
eigenvalue λ_i of the data
σ = 0.3
σ = 0.7
σ = 1.4