



A New Spin on Color Quantization

Samy Lakhal, Alexandre Darmon, Michael Benzaquen

► To cite this version:

Samy Lakhal, Alexandre Darmon, Michael Benzaquen. A New Spin on Color Quantization. Journal of Statistical Mechanics: Theory and Experiment, 2023, pp.033401. hal-03797225

HAL Id: hal-03797225

<https://hal.science/hal-03797225>

Submitted on 4 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A New Spin on Color Quantization

Samy Lakhal,^{1,2,3} Alexandre Darmon,⁴ and Michael Benzaquen^{1,3,5,*}

¹*LadHyX, UMR CNRS 7646, Ecole Polytechnique, 91128 Palaiseau Cedex, France*

²*Institut Jean Le Rond d'Alembert, UMR CNRS 7190, Sorbonne Université, 75005 Paris, France*

³*Chair of EconophysiX & Complex Systems, Ecole Polytechnique, 91128 Palaiseau Cedex, France*

⁴*Art in Research, 33 rue Censier, 75005 Paris, France*

⁵*Capital Fund Management, 23 rue de l'Université, 75007 Paris, France*

(Dated: April 28, 2022)

We address the problem of image color quantization using a Maximum Entropy based approach. We argue that adding thermal noise to the system yields better visual impressions than that obtained from a simple energy minimization. To quantify this observation, we introduce the coarse-grained quantization error, and seek the optimal temperature which minimizes this new observable. By comparing images with different structural properties, we show that the optimal temperature is a good proxy for complexity at different scales. Finally, having shown that the convoluted error is a key observable, we directly minimize it using a Monte Carlo algorithm to generate a new series of quantized images. Adopting an original approach based on the informativity of finite size samples, we are able to determine the optimal convolution parameter leading to the best visuals.

INTRODUCTION

In physics, many problems can be formulated in terms of energy minimization. However, for complex systems with a large number of degrees of freedom, analytical minimizers are often difficult to find, and the ground state is seldom representative of the true physical state of the system at hand. To overcome such issues, a classical method in statistical thermodynamics is to slightly relax the energy minimization constraint and introduce a probabilistic model relying on entropy maximization [1, 2]. Such an approach has allowed for the exploration of sub-optimal solutions with thermal noise and led to the emergence of historical results on phase transitions, e.g. for Ising models [3, 4]. This method has since been popularized in various fields, for example in biology for inference problems [5] or in computer science for classifiers (see e.g. the softmax function [6]) and annealing procedures [7].

Let us now consider the problem of field quantization and its application for images, that is color quantization. It consists in choosing a set of authorized states called the color palette and then projecting each pixel of the original image onto this palette. This method is naturally very relevant for compression and other problems involving digital image processing. The first step consists in finding the most convenient color palette from the original histogram using thresholding levels [8, 9] or clustering methods [10, 11]. The second step is usually done by introducing a suitable dithering procedure [12–14] in order to reduce threshold artifacts and improve the overall visual quality of the quantized image.

Here, we shall focus on the second step by constraining the color palette *a priori*. We adopt a field theory approach, based on the exploration of simple observables at and around optimality using thermal noise. We first pro-

vide a method to explore new solutions to the quantization problem with a Maximum Entropy based approach. We show that, in the specific case of color quantization and with the simplest error measure, this method yields surprisingly good visuals when varying the temperature of the system. To quantify this observation, we introduce a new observable, the convoluted error, to look for optimal thermal noise levels regarding the overall image quality. We then confront the results for images with different structural characteristics, and show that the optimal temperatures are a good proxy for image complexity. We finally implement a Monte Carlo algorithm to directly minimize the convoluted error, and, following a novel approach on the informativity of finite size samples [15–17], we determine the convolution parameter providing the best visuals.

I. STATE QUANTIZATION

We consider a field $h(\mathbf{r}) \subset F$ and want to build its optimal quantized version $\hat{h}(\mathbf{r}) \subset \hat{F}$ where $\hat{F}$ is a subset of F . To do so, one usually minimizes a loss function between $\hat{h}$ and h . A first natural choice for the loss function is a site-wise measure of the quantization error $\mathcal{L}(h, \hat{h}) = \sum_{\mathbf{r}} \mathcal{L}_F(h(\mathbf{r}), \hat{h}(\mathbf{r}))$ where $\mathcal{L}_F$ is the loss for each site, usually Euclidean or logical. With such a definition, the field $\hat{h}^*$ minimizing $\mathcal{L}(h, \hat{h})$ is simply obtained by replacing each original data with the closest state in $\hat{F}$. Note that this is what commonly happens during the sampling of a continuous signal with an instrument, such as a camera projecting colors in the RGB space [18]. There are cases where this simple quantization process leads to unsatisfying $\hat{h}^*$ fields deviating too much from the original data. For example, quantizing a continuous white noise on a grid with a threshold level artificially generates correlated samples of the site percolation problem [19–21].

* michael.benzaquen@polytechnique.edu

To engineer more relevant loss measures, our idea is to explore suboptimal configurations around $\hat{h}^*$ using the Maximum Entropy approach mentioned above [1]. This allows one to define the most agnostic – that is the most entropic – classes of distributions with given constraints, such as normalization. In the context of state quantization we look for the distributions ($\mathcal{P}$) over quantized fields that maximize the following functional:

$$J(\mathcal{P}) = S(\mathcal{P}) - \mu \left[\sum_{\hat{h}} \mathcal{P}(\hat{h}) - 1 \right] - \lambda \sum_{\hat{h}} \mathcal{P}(\hat{h}) \left[\mathcal{L}(h, \hat{h}) - \mathcal{L}(h, \hat{h}^*) \right], \quad (1)$$

where the first term $S(\mathcal{P}) = -\sum_{\hat{h}} \mathcal{P}(\hat{h}) \log \mathcal{P}(\hat{h})$ is the distribution entropy, the second term is the normalization constraint, the last one a constraint on quantization error, and μ and λ their respective Lagrange multipliers. Differentiating Eq. (1) with respect to $\mathcal{P}(\hat{h})$ and μ allows to enforce normalization while leaving λ as a free parameter. This leads to:

$$\mathcal{P}_{h,\lambda}(\hat{h}) = \frac{1}{Z_{h,\lambda}} e^{-\lambda \mathcal{L}(h, \hat{h})}, \quad (2)$$

where $Z_{h,\lambda}$ is the partition function. By setting $T := 1/\lambda$, we recover a Boltzmann-like distribution where the loss function plays the role of the energy. In extreme cases such as $T \rightarrow 0$ and $T \rightarrow +\infty$, we respectively recover the Dirac delta distribution centred on $\hat{h}^*$ and the uniform distribution. As in the study of any Hamiltonian system, increasing the temperature softens the energy minimization constraint and is therefore the opportunity to test other basic observables.

II. COLOR QUANTIZATION

We now apply the above formalism in the classical image processing problem of grayscale quantization. The goal is to reduce the amount of shades taken by the pixels in an image, usually described with 256 levels. In this context, h and $\hat{h}$ respectively correspond to the original and reduced images, while F and $\hat{F}$ respectively correspond to the initial and quantized sets of grayscale levels. To define the loss function we use the naive Squared Error, obtained from squaring the Euclidean distance:

$$\mathcal{L}(h, \hat{h}) = \|h - \hat{h}\|_2^2 = \sum_{\mathbf{r}} [h(\mathbf{r}) - \hat{h}(\mathbf{r})]^2. \quad (3)$$

This loss function being pixel-wise separable, so is the corresponding distribution (Eq. (2)):

$$\mathcal{P}_{h,T}(\hat{h}) = \prod_{\mathbf{r}} p_{h(\mathbf{r}),T}(\hat{h}(\mathbf{r})) = \prod_{\mathbf{r}} \frac{1}{z_{h(\mathbf{r}),T}} e^{-[h(\mathbf{r}) - \hat{h}(\mathbf{r})]^2/T}, \quad (4)$$




FIG. 1. Influence of thermal noise on color quantization. (a) Original benchmark image. (b-d) Quantized versions of the original image generated at low ($T \simeq 0$), intermediate ($T = 0.3$) & high ($T = 1$) temperatures using Eq. (4). (e) Evolution of the rescaled Mean Convolved Squared Error between the original and quantized images (Eq. (5)) with temperature for different values of the convolution parameter α (Eq. (6)).

where $z_{h(\mathbf{r}),T}$ is the partition function of the marginal distribution $p_{h(\mathbf{r}),T}$.

To test our method, we use a classical benchmark image shown in Fig. 1(a), and we sample quantized versions by using the marginal distributions defined in Eq. (4). The authorized colors are chosen as black and white, meaning that $\hat{F} = \{0, 255\}$. Figure 1(b) was generated at $T \rightarrow 0$, equivalent to the naive minimization $\hat{h}^*$ presented in Sec. I, where any texture in the $[0, 127]$ or $[128, 255]$ intervals is simply replaced by black or white pixels respectively. The image, although still recognizable, displays thresholding artifacts such as contouring effects for shaded textures and suppresses a vast amount

of details. In Fig. 1(c), we introduce thermal noise with intermediate temperature $T = 0.3$, leading to a more interesting visual. Parts of the lighter and darker shades are reconstructed and other details like contours now accurately correspond to the real physical features of the objects, and no longer to fluctuations around the threshold value. In Fig. 1(d), one can see that a higher thermal noise level no longer represents a positive contribution, as one excessively randomizes the pixel attribution rule.

Figure 1(c) thus interestingly appears to be a better quantized version than Fig. 1(b)&(d), especially when looking at it from a distance. Precisely, taking a step back has the effect of coarse-graining/convoluting the image and erasing the small-scale fluctuations created by thermal noise. To quantify such an observation, we define the Convolved Squared Error $\mathcal{L}_\theta$ comparing the original and quantized fields after a convolution through a given kernel θ :

$$\mathcal{L}_\theta(h, \hat{h}) = \|(h - \hat{h}) \otimes \theta\|_2^2. \quad (5)$$

Among the many classes of kernels commonly used in image processing, we choose a power-law kernel of the form $\theta_\gamma(\mathbf{r}) \propto \|\mathbf{r}\|^{-\gamma}$, for the sake of physical interpretability and mathematical tractability. Its Fourier transform is also a power-law:

$$\tilde{\theta}_\alpha(\mathbf{k}) \propto \|\mathbf{k}\|^{-\alpha}, \quad (6)$$

with $\alpha = d/2 - \gamma$, where $d = 2$ is the space dimension.

When $\alpha = 0$, the kernel in the direct space is narrow and leaves the image invariant. As α increases, the convolution operation replaces each pixel value with its local average of the field. Figure 1(e) displays the Mean Convolved Squared Error $\mathbb{E}_{\mathcal{P}_{h,T}}[\mathcal{L}_\theta]$ as a function of temperature (see Appendix A for the details of the computation of $\mathbb{E}_{\mathcal{P}_{h,T}}[\mathcal{L}_\theta]$). As expected, the unconvolved Squared Error ($\alpha = 0$, solid line) increases monotonously with T . For higher values of α , however, a local minimum appears at a finite temperature T_α^* . Note that Fig. 1(c) was generated with a noise level T close to the minima displayed in Fig. 1(e), thereby confirming the intuition that the Mean Convolved Squared Error is a relevant observable for the color quantization problem.

III. VISUAL COMPLEXITY

The optimal temperature T_α^* naturally depends on the color histogram of the original image, but also on the spatial arrangement of its pixels. To quantify this last statement, we compare our benchmark image with its transformation through a histogram-invariant operation. We use a procedure which shuffles a fraction of the total pixels at random positions. In Figs. 2(a1) and (a2), we respectively display the original image and its shuffled version, and we plot the temperatures T_α^* as function of α in Fig. 2(a3). For low values of α , the kernel is too narrow and the convolution has almost no effect, unsurprisingly leading to $T_\alpha^* = 0$ for both images.




FIG. 2. Evolution of the optimal temperature T_α^* with the convolution parameter α for different images. (a) Comparison between the original image and a shuffled version generated with a randomizing procedure. (b) Comparison between two benchmark images (Peppers & Mandrill) displaying structural features at different scales.

Then, both temperatures monotonously increase with α until they meet at a plateau where the evolution is independent of the spatial distribution of pixels. Interestingly, we observe that T_α^* is systematically higher for the original image, meaning that the shuffling procedure has strongly affected its structural properties. Indeed, natural images such as the peppers present intelligible patterns and strong spatial regularities, far from the random and uncorrelated rearrangement that the shuffling procedure creates. This behaviour is somehow reminiscent of several classical physical systems such as the Random Field Ising Model (RFIM) for which irregularities lower the critical temperature [22, 23] (see Appendix B for more details on the link between the Convolved Squared Error and the RFIM Hamiltonian).




FIG. 3. Monte Carlo image generation. (a) Original benchmark image. (b-d) Monte Carlo simulations for $\alpha = 0.02$, $\alpha = 0.05$, $\alpha = 0.5$ respectively. Images were initialized with $\hat{h}^*$ and the simulation ran until the loss function reached stability.

In Fig. 2(b), we test another approach by comparing the evolution of T_α^* with α for two different benchmark images, each presenting interesting visual features at different scales. Figure 2(b1) presents less small scale features than Fig. 2(b2), resulting in a higher optimal temperature for low values of α , see Fig. 2(b3) (the inset shows the difference of temperatures ΔT_α^* between the two images). This tendency reverses at higher scales. Using this procedure we compared a number of other types of natural images (forests, fields, buildings, landscapes), as well as classes of simple abstract textures like those presented in [24], with the same conclusions. This supports the intuition that the temperatures T_α^* may be used as measures of multi-scale visual complexities and as such, consistent input features for aesthetic assessment algorithms [25–29].

IV. MONTE CARLO IMAGE GENERATION

As argued above, the Convolved Squared Error $\mathcal{L}_\theta$ in Eq. (5) is a highly relevant observable for the color quantization problem. As such, it seems reasonable to use it as our effective energy function. An idea, see e.g. [14], is to directly generate images minimizing this error function using a Monte Carlo algorithm. This algorithmic approach is necessary as there is no explicit minimum of $\mathcal{L}_\theta(h, \cdot)$.

Here we implement a simplified Monte Carlo algorithm

where we use the power-law kernel defined in Eq. (6) (rather than the Gaussian window with simulated annealing procedures used in [14]). Images were initialized with the solution $\hat{h}^*$, and the algorithm ran until convergence of the loss functions. For algorithmic efficiency, we use the Parsival equality on Eq. (5) to obtain $\mathcal{L}_\theta(h, \hat{h}) = \|\mathcal{F}[h - \hat{h}](\mathbf{k}) \cdot k^{-2\alpha}\|_2^2$, with $\mathcal{F}$ the Fourier transform calculated using the FFT algorithm (thereby assuming periodic boundary conditions). Figure 3(a) displays again the original image for reference, and Figs. 3(b-d) the images corresponding respectively to $\alpha = 0.02$, $\alpha = 0.05$, $\alpha = 0.5$. For Fig. 3(b), the kernel function is narrow: small-scale details like the shadows on the peppers are faithfully reproduced, while leaving large areas of uniform color. Increasing α helps removing the latter artefact and improves the overall visual impression when looked at from a distance, see Fig. 3(c). However, too large convolution windows yield images lacking small-scale accuracy, see Fig. 3(d). A compromise shall thus be found in order to generate the most faithful quantized image, that is the optimal α providing the best trade-off throughout different scales.

V. OPTIMAL COLOR QUANTIZATION

To find the most suitable parameter α , we use a recent information theory based approach [15–17] where the spatial resolution becomes a tunable parameter of the system. First, one defines the image as a sample $(\mathbf{r}, \hat{h}(\mathbf{r}))_{\mathbf{r}}$ of pixels. Each tuple being unique, this description is far from being the most efficient in the context of information theory. Indeed, information could for example be saved by accounting for neighboring pixels with the same color value. One thus modifies the spatial resolution by considering a new sample $(G^\ell(\mathbf{r}), \hat{h}(\mathbf{r}))_{\mathbf{r}}$ where $G^\ell(\mathbf{r})$ is the new position on a coarser grid of step $\ell \in \{1, \dots, L\}$, with L the size of the original image. Several tuples can now be in the same state s and one can look at their occurrence – or degeneracy level – k inside each block. Finally, one constructs two measures of entropy $\hat{H}^\ell[s]$ and $\hat{H}^\ell[k]$, respectively assessing the heterogeneity in the data and the heterogeneity in the data redundancy.

In Fig. 4(a), we vary ℓ and plot $\hat{H}^\ell[k]$ as function of $\hat{H}^\ell[s]$ for images generated with different α . Of most interest to us here is the right part of the graphs, which corresponds to small values of ℓ , and for which the concavity of the curves is very α -dependent. Indeed, one can show – see Appendix C – that the local slope $\mu = d\hat{H}^\ell[k]/d\hat{H}^\ell[s]$ actually corresponds to the trade-off rate between relevance and resolution as data is compressed. The idea is then to choose the convolution parameter such that this trade-off is as stable as possible across all scales of observations, meaning that μ should be close to -1 and as constant as possible when varying ℓ . In other terms, the right part of the graph should be as linear as possible with slope -1. In Fig. 4(b) we plotted the regres-




FIG. 4. Influence of the convolution parameter α on the compressing regime of color quantized images. (a) Plot of $(H^l[s], H^l[k])_\ell$ for low, optimal and high convolution parameters α . (b) Regression coefficient R^2 of the linear fit as function of α . (c) Optimally quantized image, $\alpha = 0.054$.

sion coefficient R^2 obtained from the corresponding linear fit and found that it was maximized for an intermediate value $\alpha = 0.054$. The corresponding image is shown in Fig. 4(c) and is indeed a very good visual compromise.

CONCLUSION

Let us summarize what we have achieved. In the context of color quantization, we confirmed that the naive approach consisting in a simple error minimization does not generally bring satisfactory visuals. To overcome this issue, we introduced thermal noise through a Maximal Entropy based approach and generated quantized images

with more interesting visuals. To quantify this visual impression, we introduced a new observable, called the Convoluted Squared Error, which compares the original and quantized fields after a coarse-graining procedure. Introducing convolution allowed us to find the optimal temperatures minimizing the new observable. Interestingly, we found that such temperatures are a good indicator for complexity at different scales. Moreover, having shown that the Convoluted Squared Error is a highly relevant observable with respect to color quantization, we directly minimized it to generate new images using a Monte Carlo algorithm. In order to find the optimal convolution parameter leading to the best visuals, we used an original approach based on the informativity of finite size samples.

Note that, as mentioned in the introduction, in the present analysis we have chosen the target color palette *ex ante*, focusing on the simplest B&W case. It would be interesting to consider extended color palettes where both the dimension and color values are optimized for a given image. The latter most likely depends on the features of the color histogram of the original image, but also their spatial arrangement. Future research should also be devoted to considering alternatives to the Euclidean distance, such as the Wasserstein distance [30] which should be able to account *a priori* for the local arrangement of the pixels. Finally, another area for improvement is that of the convolution procedure; one might want to test alternative kernel functions which would allow to avoid using periodic boundary conditions.

ACKNOWLEDGMENTS

We warmly thank R. Benichou and J.-P. Bouchaud who contributed to the early stages of this work, as well as M. Marsili whose input on seeking the optimal α was crucial to conclude our analysis. We also thank C. Aubrun, J. Garnier-Brun, P. Mergny and A. Pammi for fruitful discussions. This research was conducted within the Econophysics & Complex Systems Research Chair, under the aegis of the Fondation du Risque, the Fondation de l'Ecole polytechnique, the Ecole polytechnique and Capital Fund Management.

-
- [1] E. T. Jaynes, Physical Review **106**, 620 (1957).
 - [2] E. T. Jaynes, Physical Review **108**, 171 (1957).
 - [3] E. Ising, Zeitschrift für Physik **31**, 253 (1925).
 - [4] L. Onsager, Physical Review **65**, 117 (1944).
 - [5] A. de Martino and D. de Martino, Heliyon **4**, 596 (2018).
 - [6] J. Bridle, in *Advances in Neural Information Processing Systems*, Vol. 2 (Morgan-Kaufmann, 1989).
 - [7] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Science **220**, 671 (1983).
 - [8] P. Heckbert, ACM Siggraph Computer Graphics **16**, 297 (1982).
 - [9] M. Gervautz and W. Purgathofer, in *New Trends in Computer Graphics* (Springer, Berlin, Heidelberg, 1988) pp. 219–231.
 - [10] R. Balasubramanian, C. A. Bouman, and J. P. Allebach, Journal of Electronic Imaging **3**, 45 (1994).

- [11] A. H. Dekker, *Network: Computation in Neural Systems* **5**, 351 (1994).
- [12] R. Ulichney, *Digital Halftoning* (MIT Press, Cambridge, MA, USA, 1987).
- [13] M. Orchard and C. Bouman, *IEEE Transactions on Signal Processing* **39**, 2677 (1991).
- [14] J. Puzicha, M. Held, J. Ketterer, J. Buhmann, and D. Fellner, *IEEE Transactions on Image Processing* **9**, 666 (2000).
- [15] R. J. Cubero, M. Marsili, and Y. Roudi, *Entropy* **20**, 755 (2018).
- [16] R. J. Cubero, M. Marsili, and Y. Roudi, *Journal of Computational Neuroscience* **48**, 85 (2020).
- [17] A. Haimovici and M. Marsili, *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P10013 (2015).
- [18] D. A. Rowlands, *Optical Engineering* **59**, 110801 (2020).
- [19] R. F. Voss, *Journal of Physics A: Mathematical and General* **17**, L373 (1984).
- [20] A. A. Saberi, *Physics Reports Recent advances in percolation theory and its applications*, **578**, 1 (2015).
- [21] V. K. Shante and S. Kirkpatrick, *Advances in Physics* **20**, 325 (1971).
- [22] S. Fishman and A. Aharony, *Journal of Physics C: Solid State Physics* **12**, L729 (1979).
- [23] D. P. Belanger and A. P. Young, *Journal of Magnetism and Magnetic Materials* **100**, 272 (1991).
- [24] S. Lakhal, A. Darmon, J.-P. Bouchaud, and M. Benzaquen, *Physical Review Research* **2**, 022058 (2020).
- [25] G. Valenzise, C. Kang, and F. Dufaux, in *Human Perception of Visual Information: Psychological and Computational Perspectives* (Springer International Publishing, Cham, 2022) pp. 133–181.
- [26] Z. Yuan, Y. Qi, M. Hu, and Q. Li, in *2020 IEEE International Conference on Multimedia Expo Workshops (ICMEW)* (2020) pp. 1–6.
- [27] H. Maître, *Esthétique de la photographie numérique* (ISTE Editions, 2022).
- [28] A. A. Bagrov, I. A. Iakovlev, A. A. Iliasov, M. I. Katsnelson, and V. V. Mazurenko, *Proceedings of the National Academy of Sciences* **117**, 30241 (2020).
- [29] J. McCormack, C. Cruz Gambardella, and A. Lomas, in *Artificial Intelligence in Music, Sound, Art and Design* (Springer International Publishing, Cham, 2021) pp. 203–217.
- [30] L. Rüschendorf, *Probability Theory and Related Fields* **70**, 117 (1985).

Appendix A: Mean (Convolved) Squared Error

Here we derive the expression of the Mean Convolved Squared Error (Fig. 1(e)). Without convolution, the expectation of the loss function reads:

$$\begin{aligned}\mathbb{E}_{\mathcal{P}} [\mathcal{L}(h, \hat{h})] &= \mathbb{E}_{\mathcal{P}} [\|\hat{h} - h\|_2^2] \\ &= \sum_{\mathbf{r}} \mathbb{E}_{\mathcal{P}} \left[\left(\hat{h}(\mathbf{r}) - h(\mathbf{r}) \right)^2 \right],\end{aligned}\quad (\text{A1})$$

where we have used the linearity of $\mathbb{E}_{\mathcal{P}}$. Each term inside the expectation depends on the marginal variable $\hat{h}(\mathbf{r})$ and can be easily computed site-by-site, such that Eq. (A1) can be rewritten as:

$$\mathbb{E}_{\mathcal{P}} [\mathcal{L}(h, \hat{h})] = \sum_{\mathbf{r}} \mathbb{E}_{p_{h(\mathbf{r})}} \left[\left(\hat{h}(\mathbf{r}) - h(\mathbf{r}) \right)^2 \right]. \quad (\text{A2})$$

Applying the same ideas to the Mean Convolved Squared Error yields:

$$\begin{aligned}\mathbb{E}_{\mathcal{P}} [\mathcal{L}_{\theta}(\hat{h}, h)] &= \mathbb{E}_{\mathcal{P}} [\|(\hat{h} - h) \otimes \theta\|_2^2] \\ &= \sum_{\mathbf{r}} \mathbb{E}_{\mathcal{P}} \left[\left((\hat{h} - h) \otimes \theta(\mathbf{r}) \right)^2 \right].\end{aligned}\quad (\text{A3})$$

Developing the square leads to: $\mathbb{E}_{\mathcal{P}}[\mathcal{L}_{\theta}(\hat{h}, h)] =$

$$\begin{aligned}&\sum_{\mathbf{r}, \mathbf{r}_1, \mathbf{r}_2} \theta(\mathbf{r} - \mathbf{r}_1) \theta(\mathbf{r} - \mathbf{r}_2) \times \\ &\mathbb{E}_{\mathcal{P}} \left[\left(\hat{h}(\mathbf{r}_1) - h(\mathbf{r}_1) \right) \left(\hat{h}(\mathbf{r}_2) - h(\mathbf{r}_2) \right) \right].\end{aligned}\quad (\text{A4})$$

The marginal distributions being independent, non-diagonal second order terms (with $\mathbf{r}_1 \neq \mathbf{r}_2$) inside the expectation can be written as products of order 1 expectations. Rewriting the diagonal in terms of squared expectations leads to the following expression:

$$\begin{aligned}\mathbb{E}_{\mathcal{P}} [\mathcal{L}_{\theta}(\hat{h}, h)] &= \|\mathbb{E}_{\mathcal{P}}[\hat{h} - h] \otimes \theta\|_2^2 \\ &+ \left(\mathbb{E}_{\mathcal{P}} [\mathcal{L}(h, \hat{h})] - \|\mathbb{E}_{\mathcal{P}}[\hat{h} - h]\|_2^2 \right) \|\theta\|_2^2,\end{aligned}\quad (\text{A5})$$

where $\mathbb{E}_{\mathcal{P}}[\mathcal{L}(h, \hat{h})]$ is the previously defined Mean Squared Error in Eq. (A2), and $\mathbb{E}_{\mathcal{P}}[\hat{h} - h]$ is the field computed using the marginal distributions on each site.

Appendix B: Analogy with the RFIM

Here we show that the Convolved Squared Error can be rewritten as the Hamiltonian of a Random Field Ising Model (RFIM). Developing the norm in Eq. (5), one obtains:

$$\mathcal{L}_{\theta}(\hat{h}, h) = \|(\hat{h} - h) \otimes \theta\|_2^2 = \sum_{\mathbf{r}} [(\hat{h} - h) \otimes \theta(\mathbf{r})]^2. \quad (\text{B1})$$

Rewriting explicitly the convolution product and defining the interaction matrix J as $J(\mathbf{r}) = \sum_{\mathbf{r}'} \theta(\mathbf{r}') \theta(\mathbf{r}' - \mathbf{r})$, one obtains:

$$\mathcal{L}_{\theta}(\hat{h}, h) = \sum_{\mathbf{r}, \mathbf{r}'} [\hat{h}(\mathbf{r}) - h(\mathbf{r})] J(\mathbf{r}' - \mathbf{r}) [\hat{h}(\mathbf{r}') - h(\mathbf{r}')]. \quad (\text{B2})$$

Since the loss is essentially defined up to a constant, one can discard $\hat{h}$ -independent terms in Eq. (B2), such that $\mathcal{L}_{\theta}(\hat{h}, h) \equiv \mathcal{H}_h[\hat{h}]$ where:

$$\mathcal{H}_h[\hat{h}] := \sum_{\mathbf{r}, \mathbf{r}'} \hat{h}(\mathbf{r}) J(\mathbf{r}' - \mathbf{r}) \hat{h}(\mathbf{r}') - 2 \sum_{\mathbf{r}'} (h \otimes J)(\mathbf{r}') \hat{h}(\mathbf{r}'). \quad (\text{B3})$$

Interestingly, when replacing $\hat{h}$ by a spin field ϕ , one precisely recovers the Antiferromagnetic RFIM Hamiltonian $\mathcal{H}[\phi]$ which, within the bra-ket formalism writes:

$$\mathcal{H}_h[\phi] = \langle \phi | J | \phi \rangle - 2 \langle h | J | \phi \rangle, \quad (\text{B4})$$

where $|h_{\text{eff}}\rangle = 2J|h\rangle$ is the effective external field. Therefore, spins and external field in the Ising model respectively play the same role as the transformed and original images in the color quantization problem. A well-known property of the Antiferromagnetic RFIM, that is interactions encouraging alternations between neighbouring spins, is for example recovered in our observations, see Fig. 2(a).

Appendix C: Resolution/relevance formalism

Here we introduce the formalism used to set the optimal convolution parameter in Sec. V. The idea, detailed in [15–17], consists in studying a data sample for different resolution/compression levels and calculating entropy-based observables.

Let $\mathcal{S} = (s_1, \dots, s_N)$ be a sample of data, and ℓ a compression parameter. One can transform the original sample into a compressed one $\mathcal{S}^\ell = (s_1^\ell, \dots, s_N^\ell)$ such that $\ell = 1$ corresponds to no compression ($\mathcal{S}^1 = \mathcal{S}$) and $\ell = L$ to the totally compressed sample (quasi-degenerate data). At each compression level ℓ , one can easily calculate the number k_s^ℓ of times the state s appears in the sample $\mathcal{S}^\ell$, and the number m_k^ℓ of states appearing k times. One can then define the *resolution* as the entropy of the empirical distribution of states:

$$\hat{H}^\ell[s] = - \sum_s \frac{k_s^\ell}{N} \log \frac{k_s^\ell}{N} = - \sum_k \frac{km_k^\ell}{N} \log \frac{k}{N}. \quad (\text{C1})$$

Another useful observable is the *relevance* defined as the

entropy computed from the states sharing the same occurrence frequency:

$$\hat{H}^\ell[k] = \sum_k \frac{km_k^\ell}{N} \log \frac{km_k^\ell}{N}. \quad (\text{C2})$$

This measure of *relevance* is indeed the most direct way to encode the underlying distribution of the original data. As the latter entropy takes less bits to encode, one has $\hat{H}^\ell[k] < \hat{H}^\ell[s]$, explaining why the data in Fig. 4(a) is under the $y = x$ line. In order to quantify the way information spreads across compression levels, we define the compression rate $\mu_{\ell \rightarrow \ell+1}$ between two successive compression levels as the ratio between loss in relevance and loss in resolution:

$$\mu_{\ell \rightarrow \ell+1} := \frac{\hat{H}^{\ell+1}[k] - \hat{H}^\ell[k]}{\hat{H}^{\ell+1}[s] - \hat{H}^\ell[s]}. \quad (\text{C3})$$

In the right part of Fig. 4(a) and for the first compression steps (low values of ℓ), $\hat{H}[k]$ is a decreasing function of $\hat{H}[s]$, resulting in negative μ . Interestingly, this means that compressing the sample increases the amount of relevant information it contains about the underlying distribution. However, when ℓ increases further, we reach an oversampling regime (left part of Fig. 4(a)) for which further compression reduces both relevance and resolution.

In the case of color quantization, we wish to gain information relevance as we range from low to higher compression in the most scale-invariant way possible. This way, we avoid the scale-specific tradeoff described in Fig. 3. Optimality in that regard is therefore reached by taking μ close to -1 and as constant as possible in the under-sampling regime, by that justifying the linear regression introduced in Sec. V.