


# Introduction to Bayesian Statistics - 9

*Edoardo Milotti*

Università di Trieste and INFN-Sezione di Trieste

# Our next important topic: Bayesian estimates often require complex numerical integrals. How do we confront this problem?



enter the Monte Carlo methods!

1. acceptance-rejection sampling
2. importance sampling
3. statistical bootstrap
4. Bayesian methods in a sampling-resampling perspective
5. Introduction to Markov chains and to Random Walks (RW)
6. Detailed balance and Boltzmann's H-theorem
7. The Gibbs sampler
8. Simulated annealing and the Traveling Salesman Problem (TSP)
9. The Metropolis algorithm
- 10. More on Gibbs sampling**
- 11. Image restoration and Markov Random Fields (MRF)**
- 12. The Metropolis-Hastings algorithm and Markov Chain Monte Carlo (MCMC)**
13. The efficiency of MCMC methods
14. Affine-invariant MCMC algorithms (emcee)

## 10. More on Gibbs sampling

Recall the recipe for Gibbs sampling: we generate a "Gibbs sequence" of random variables

$$Y'_0, X'_0, Y'_1, X'_1, Y'_2, X'_2, \dots, Y'_k, X'_k$$

where one initial value is specified and the others are computed with the rule

$$X'_j \sim f(x \mid Y'_j = y'_j)$$

$$Y'_{j+1} \sim f(y \mid X'_j = x'_j)$$

## Example: bivariate Gaussian distribution

- Bivariate Gaussian likelihood

$$p(y_1, y_2 | \theta_1, \theta_2) = \frac{1}{2\pi \sqrt{|\det V|}} \exp \left[ -\frac{1}{2} (\mathbf{y} - \boldsymbol{\theta})^T V^{-1} (\mathbf{y} - \boldsymbol{\theta}) \right]$$

where the known covariance matrix is

$$V = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

- Posterior pdf from Bayes theorem with improper priors, with a single datapoint  $(y_1, y_2)$


$$p(\theta_1, \theta_2 | y_1, y_2) \propto p(y_1, y_2 | \theta_1, \theta_2)$$

- Expanding, we find

$$p(\theta_1, \theta_2 | y_1, y_2) \sim \exp \left\{ -\frac{1}{2(1-\rho^2)} [(\theta_1 - y_1)^2 - 2\rho(\theta_1 - y_1)(\theta_2 - y_2) + (\theta_2 - y_2)^2] \right\}$$

- Then, from the Gaussian structure of the posterior, we find the marginals



$$p(\theta_1 | \theta_2, y_1, y_2) \sim \exp \left[ -\frac{1}{2(1-\rho^2)} (\theta_1 - (y_1 + \rho(\theta_2 - y_2)))^2 \right]$$
$$p(\theta_2 | \theta_1, y_1, y_2) \sim \exp \left[ -\frac{1}{2(1-\rho^2)} (\theta_2 - (y_2 + \rho(\theta_1 - y_1)))^2 \right]$$




20 initial steps in a Gibbs sampler run:

- orange cross: starting pair
- green dot: position after 20 steps
- red cross: bivariate mean

marginal distributions (second half of the simulation values in a run with 10000 generated pairs)



posterior distribution (second half of the simulation values in a run with 10000 generated pairs)



**Exercise:**

extend this treatment to more than one pair of measured  $y$  values and write a computer program to implement it



# 11. Image restoration and Markov Random Fields (MRF)

The optimization problem.

When dealing with signals, we usually assume that data are corrupted by noise

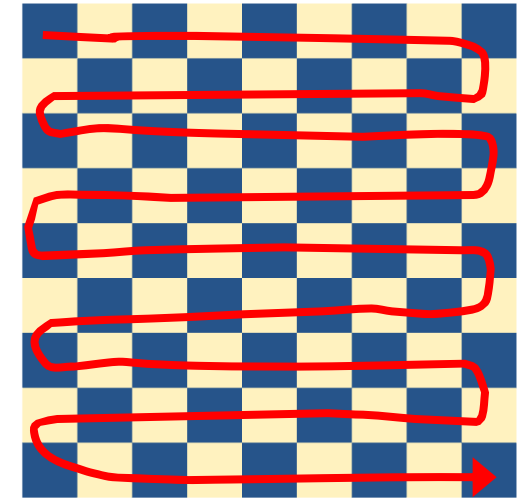
$$d = x + w$$

data vector      signal vector      (vector) noise process

An image can be viewed as a vector, for instance unfolding the sequence of pixels as shown on the right, we obtain the equivalent of a long signal vector.


If there are  $n$  pixels on each side, there are in all  $n^2$  pixels, and if there are  $L$  gray levels, then the number of possible configurations that define an image is

$$N = L^{n^2}$$



$d_{11}$	$d_{12}$	$d_{13}$	$d_{14}$ ...
$d_{21}$	$d_{22}$	$d_{23}$	$d_{24}$ ...
$d_{31}$	$d_{32}$	$d_{33}$	$d_{34}$ ...

pixel map



true  
pixel vector

**$\mathbf{x}$**

*posterior pixel  
distribution*

*likelihood*

*prior pixel  
distribution*

$$P(\mathbf{x}|\mathbf{d}) = \frac{P(\mathbf{d}|\mathbf{x})}{P(\mathbf{d})} P(\mathbf{x}) \propto P(\mathbf{d}|\mathbf{x})P(\mathbf{d})$$

Bayesian estimate of  
true pixel vector from  
observed pixel vector

Given the number of possible configurations

$$N = L^{n^2}$$

we see that even for small image sizes, say  $n = 100$ , with binary levels,  $L = 2$ , we find


$$N = 2^{10^4}$$

therefore, the problem of finding the MAP estimate in a Bayesian context is a hard computational task.

# The MAP estimate depends on the prior distribution

Possible priors:


$P(\mathbf{x})$  flat prior



Maximum Likelihood Estimate  
(MLE)

$$P(\mathbf{x}|\mathbf{d}) \propto P(\mathbf{d}|\mathbf{x})P(\mathbf{d}) \propto P(\mathbf{d}|\mathbf{x})$$

$P(\mathbf{x})$  entropic prior



Maximum Entropy Method  
(MEM)

Notice also that

$$\ln P(\mathbf{x}|\mathbf{d}) \approx \ln P(\mathbf{d}|\mathbf{x}) - [-\ln P(\mathbf{d})]$$

therefore, the MAP estimate is equivalent to maximizing the likelihood with a *penalty function*

$$-\ln P(\mathbf{d})$$

Experiments have been tried with many different penalties, many of them barely justified on probabilistic grounds (or not at all!)

Now, let  $\mathbf{x}$  be the vector of “true values” (uncorrupted intensities of an image, a spectrum, etc. ...), and translate these values into counts

$$n_i = \lfloor \alpha d_i \rfloor$$

( $i = 1, \dots, M$ ). The least informative prior corresponds to a structureless image, and pixelwise it is once again the uniform prior. Then, the probability of one count at the  $i$ -th position is just  $1/M$ .

Likewise, the probability of a given vector of values where the total number of counts is  $N$ , is given by the multinomial probability

$$P(\mathbf{n}) = \frac{N!}{n_1!n_2!\dots n_M!} \left(\frac{1}{M}\right)^N ; \quad \sum_k n_k = N$$

## Using Stirling's approximation

$$n! \approx n^n e^{-n} \quad \ln n! \approx n \ln n - n$$

we find, with the definition  $p_i = \frac{d_i}{\sum_{k=1}^M d_k}$


$$\begin{aligned} \ln P(\mathbf{n}) &\approx (N \ln N - N) - \sum_k (n_k \ln n_k - n_k) \\ &= N \ln N - \sum_k n_k \ln n_k \\ &\approx -\alpha \sum_k d_k \ln d_k + \text{const.} \end{aligned}$$

$$P(\mathbf{n}) \sim \exp \left( -\alpha \sum_k d_k \ln d_k \right) \sim \exp \left( -\alpha' \sum_k p_k \ln p_k \right) = \exp [\alpha' S(\mathbf{d})]$$

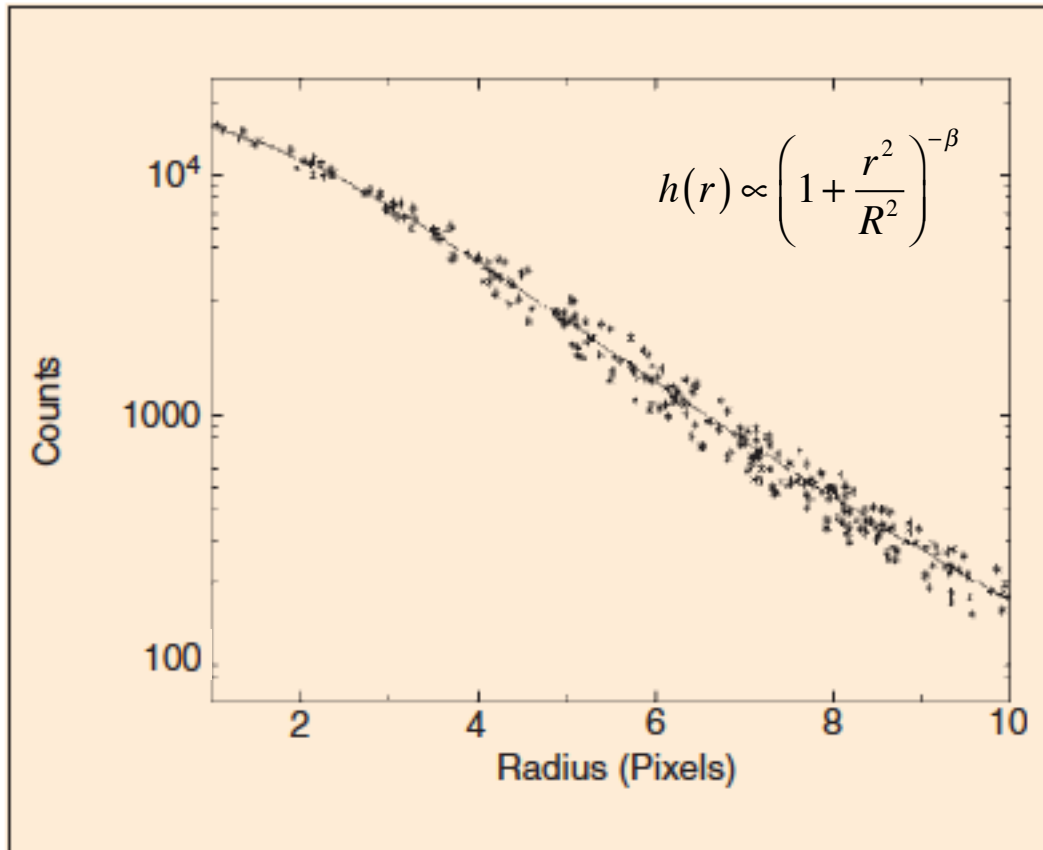
*entropic prior*



# Image likelihood: 1. the observation model

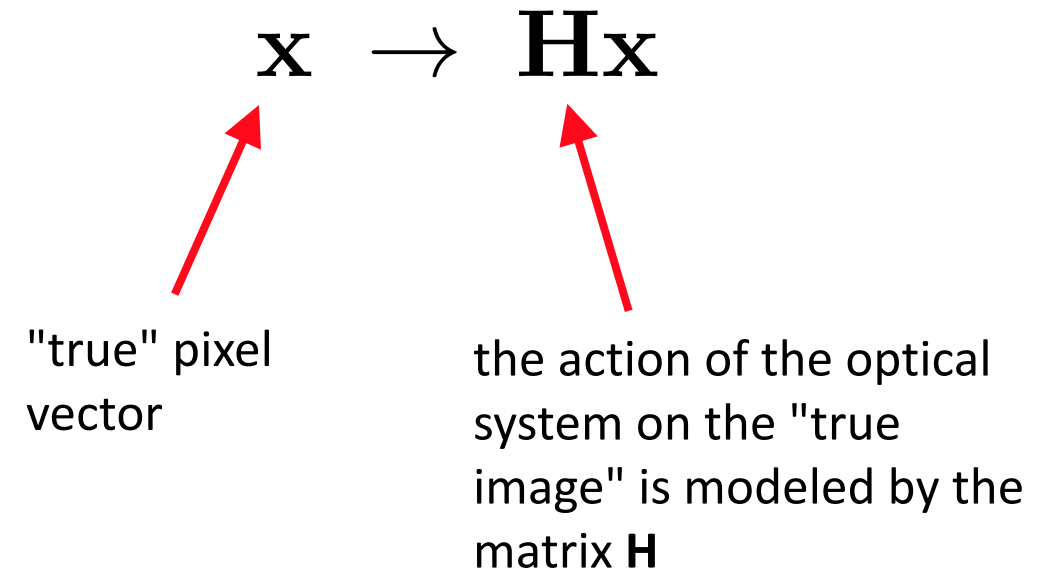








PSF from atmospheric turbulence

In general, the effect of the PSF is modeled by a linear operator



# The Hubble PSF before the first servicing mission



## Image likelihood: 2. the noise model (degradation model)

Gaussian noise model

$$P(\mathbf{x}|\mathbf{d}) \propto \exp \left[ -\frac{(\mathbf{x} - \mathbf{Hd})^2}{2\sigma^2} \right]$$

Poisson noise model

$$P(\mathbf{x}|\mathbf{d}) \propto \prod_n \frac{(\mathbf{Hd})_n^{d_n}}{d_n!} \exp[-(\mathbf{Hd})_n]$$

Poisson noise mostly from detection process, Gaussian noise mostly from electronics or from approximation of Poisson noise. Sometimes we can use the Gaussian approximation of Poisson noise

$$P(\mathbf{d}|\mathbf{x}) \propto \prod_n \frac{(\mathbf{Hx})_n^{d_n}}{d_n!} \exp[-(\mathbf{Hx})_n] \approx \prod_n \exp \left[ -\frac{(d_n - (\mathbf{Hx})_n)^2}{2(\mathbf{Hx})_n} \right] = \exp \left[ -\sum_n \frac{(d_n - (\mathbf{Hx})_n)^2}{2(\mathbf{Hx})_n} \right]$$

# Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images

STUART GEMAN AND DONALD GEMAN

**Abstract**—We make an analogy between images and statistical mechanics systems. Pixel gray levels and the presence and orientation of edges are viewed as states of atoms or molecules in a lattice-like physical system. The assignment of an energy function in the physical system determines its Gibbs distribution. Because of the Gibbs distribution, Markov random field (MRF) equivalence, this assignment also determines an MRF image model. The energy function is a more convenient and natural mechanism for embodying picture attributes than are the local characteristics of the MRF. For a range of degradation mechanisms, including blurring, nonlinear deformations, and multiplicative or additive noise, the posterior distribution is an MRF with a structure

akin to the image model. By the analogy, the posterior distribution defines another (imaginary) physical system. Gradual temperature reduction in the physical system isolates low energy states (“annealing”), or what is the same thing, the most probable states under the Gibbs distribution. The analogous operation under the posterior distribution yields the maximum *a posteriori* (MAP) estimate of the image given the degraded observations. The result is a highly parallel “relaxation” algorithm for MAP estimation. We establish convergence properties of the algorithm and we experiment with some simple pictures, for which good restorations are obtained at low signal-to-noise ratios.

# The Ising model as an example of Markov Random Field



The model describes a system of spins that point only in the +z or -z direction, so that their value can only be  $\pm 1$ .

The Hamiltonian includes only the interaction with the external magnetic field and the interaction between neighboring spins

$$H = -J \sum_{\langle i,j \rangle} \sigma_i \sigma_j - B \sum_i \sigma_i$$

The corresponding lattice magnetization is

$$M = \langle \sigma_i \rangle$$

a quantity that ranges between -1 and +1

# The Bragg-Williams approximation

This is a simple mean-field approximation

$$\langle \sigma_i \sigma_j \rangle \approx \langle \sigma_i \rangle \langle \sigma_j \rangle \quad \text{correlations are ignored}$$

so that the Hamiltonian can be restated in terms of an effective magnetic field

$$H \approx -J \sum_i \sigma_i \sum_{\langle j \rangle_i} \sigma_j - B \sum_i \sigma_i \approx (-Jz \langle \sigma_i \rangle - B) \sum_i \sigma_i = -B_{\text{eff}} \sum_i \sigma_i$$

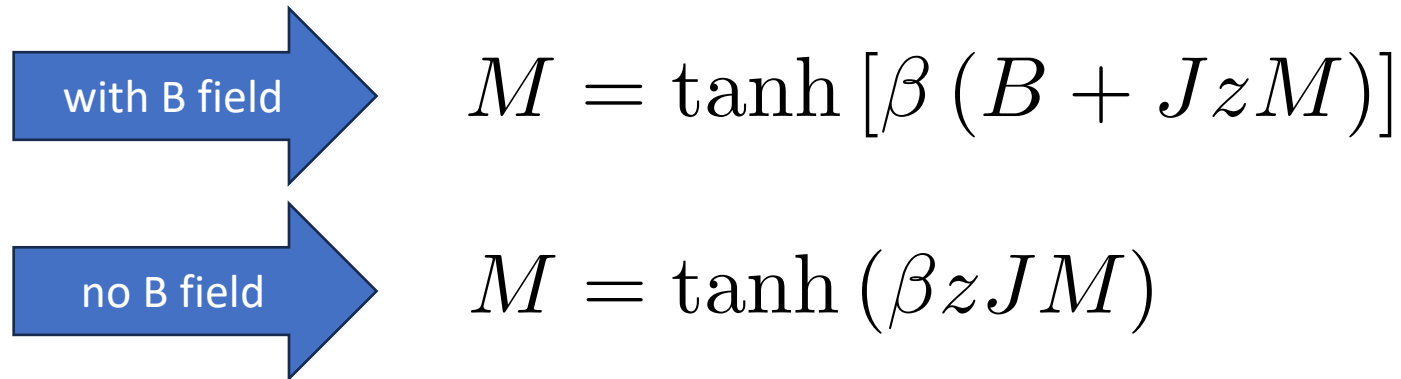
Then, the partition function is

$$\begin{aligned} Z &= \sum_{\text{configurations}} \exp \left( \frac{B_{\text{eff}} \sum_i \sigma_i}{kT} \right) = \prod_i \left( e^{B_{\text{eff}}/kT} + e^{-B_{\text{eff}}/kT} \right) \\ &= [2 \cosh (B_{\text{eff}}/kT)]^N = [2 \cosh (\beta B_{\text{eff}})]^N \end{aligned}$$

Therefore, the magnetization can be obtained as follows


$$\begin{aligned} M &= \frac{1}{NZ} \sum_{\text{configurations}} \sigma_i e^{\beta B_{\text{eff}} \sigma_i} = \frac{1}{NZ} kT \frac{\partial}{\partial B_{\text{eff}}} \sum_{\text{configurations}} e^{\beta B_{\text{eff}} \sigma_i} \\ &= kT \frac{\partial}{\partial B_{\text{eff}}} \ln Z = \tanh(\beta B_{\text{eff}}) \\ &= \tanh[\beta(B + JzM)] \end{aligned}$$

i.e., the magnetization is the solution of the nonlinear equation




with B field  $M = \tanh[\beta(B + JzM)]$

no B field  $M = \tanh(\beta zJM)$



above critical temperature




below critical temperature

Critical temperature

$$T_c = \frac{zJ}{k}$$

$$M = \tanh(\beta z J M)$$





# 90 years of the Ising model

Ernst Ising's analysis of the one-dimensional variant of his eponymous model (*Z. Phys* **31**, 253–258; 1925) is an unusual paper in the history of early twentieth-century physics. Its central result — demonstrating that a linear chain of two-state spins cannot undergo a phase transition at finite temperature — is correct, if somewhat trivial compared with other physics breakthroughs published in the 1920s. But it is Ising's fateful extension of his conclusions to two and three dimensions that proved spectacularly wrong and, paradoxically, earned him an enduring association with the model that now bears his name.

A possible reason for Ising's unexpected celebrity is that his erroneous conclusions betray a superficial understanding of what turned out to be some of the deepest and far-reaching problems to be addressed in twentieth-century physics. The Hamiltonian of the model is simple to write down — it describes a network of spins interacting with each other through a coupling that only applies if the spins are next to each other — but the physics it displays is rich and non-trivial: not only does it provide an intuitive device for illustrating the essential features of phase transitions and critical phenomena, it neatly encapsulates the main traits of



© AIP EMILIO SEGRE VISUAL ARCHIVES

the many-body problem that has come to dominate areas such as condensed-matter physics. The broader class of spin models it belongs to was used to uncover concepts such as universality, renormalization, symmetry-breaking and emergence. Ising can perhaps be forgiven for not predicting all of that.


Famously, the two-dimensional version for the model was solved analytically by Lars Onsager in the early 1940s (*Phys. Rev.* **65**, 117; 1944), a result that is rightly considered a towering achievement among many significant contributions made over the years by the likes of Peierls, Bethe, Yang, Kadanoff (see page 995) Fisher and Wilson, just to name a handful. But the

three-dimensional lattice has never been solved exactly, in spite of a multitude of attempts and false dawns — including a claim by John Maddox (who would later become the editor of *Nature*) made at a conference in Paris in 1952.

Although the 3D model is thought by some to be analytically intractable (and has also been claimed to belong to the NP-complete category of computational decision problems), progress has continued and recent numerical techniques based on conformal field theory have shed further light on the structure of the problem (*J. Stat. Phys.* **157**, 869–914; 2014). Nevertheless, the real value of the Ising model and its many derivatives lies precisely in the complexity they encapsulate. These have found use in fields as disparate as condensed-matter physics, physical chemistry, neuroscience and, more broadly, the study of so-called complex systems.

Ising studied a deceptively simple model that, unknown to him at the time, captures the essential physics of an extremely wide category of problems. He may have been wrong in his 1925 paper, but he tripped over a veritable physics goldmine.

ANDREA TARONI



Quench of an Ising system on a two-dimensional square lattice ( $500 \times 500$ ) with inverse temperature  $\beta = 10$ , starting from a random configuration

(from [https://en.wikipedia.org/wiki/Ising\\_model](https://en.wikipedia.org/wiki/Ising_model))

**2. Bayesian paradigm.** In real scenes, neighboring pixels typically have similar intensities, boundaries are usually smooth and often straight, textures, although sometimes random locally, define spatially homogeneous regions, and objects, such as grass, tree trunks, branches and leaves, have preferred relations and orientations. Our approach to picture processing is to articulate such regularities mathematically, and then to exploit them in a statistical framework to make inferences. The regularities are rarely deterministic; instead, they describe correlations and likelihoods. This leads us to the Bayesian formulation, in which prior expectations are formally represented by a probability distribution. Thus we design a distribution (a “prior”) on relevant scene attributes to capture the tendencies and constraints that characterize the scenes of interest. Picture processing is then guided by this prior distribution, which, if properly conceived, enormously limits the plausible restorations and interpretations.

from Geman & Graffigne, Proc. Int. Congress of Mathematicians, Berkeley 1986


# The Markov property of Markov Random Fields


The Ising model is an example of Markov Random Field (MRF). **What do we mean by "Markov property" in this case?**

This property corresponds to the locality of the spin interactions and the dependence of the probability of finding a certain spin with a given value only on the state of the neighboring spins.

$$P(X_s = x_s | X_r = x_r, r \in S, r \neq s) = P(X_s = x_s | X_r = x_r, r \in G_s)$$

 graph node at  $s$  takes on this value


 index  $r$  belongs to index set  $S$

 index  $r$  belongs to the set of neighbors of  $s$

## The "physical" likelihood determined by Markov Random Fields

With the assumption that the image structure behaves like the magnetization islands in an Ising spin system, we find that the likelihood is given by the Maxwell-Boltzmann distribution with the given configuration energy

$$P(\mathbf{d}|\mathbf{x}) = \frac{1}{Z_n} \exp[-\beta H_n(x, d)]$$

 noise Hamiltonian

Therefore, with a prior specified as

$$P(\mathbf{x}) = \frac{1}{Z_0} \exp[-\beta H_0(x)]$$

We find the posterior

$$P(\mathbf{x}|\mathbf{d}) = \frac{1}{Z} \exp[-\beta H(x, d)] \quad \text{with} \quad H(x, d) = H_n(x, d) + H_0(x)$$

## Choice of the noise energy

The noise energy must contain information on the background but also on the image data.


It would be natural to start with something like





$$H = -J \sum_{\langle i,j \rangle} x_i x_j - \sum_i y_i x_i$$

where the  $y$  represent the image data and act like a sort of local magnetic field. However, this does not take into account the natural fluctuations of the image data. Here we model these fluctuations with a quadratic term (which makes the final model Gaussian), which leads to the following Hamiltonian

$$H = -J \sum_{\langle i,j \rangle} x_i x_j - \frac{1}{2\sigma^2} \sum_i (y_i - x_i)^2 \sim -J \sum_{\langle i,j \rangle} x_i x_j - \frac{1}{\sigma^2} \sum_i y_i x_i$$


An undirected graphical model representing a Markov random field for image de-noising, in which  $x_i$  is a binary variable denoting the state of pixel  $i$  in the unknown noise-free image, and  $y_i$  denotes the corresponding value of pixel  $i$  in the observed noisy image.






**Figure 8.30** Illustration of image de-noising using a Markov random field. The top row shows the original binary image on the left and the corrupted image after randomly changing 10% of the pixels on the right. The bottom row shows the restored images obtained using iterated conditional models (ICM) on the left and using the graph-cut algorithm on the right. ICM produces an image where 96% of the pixels agree with the original image, whereas the corresponding number for graph-cut is 99%.


This method is not restricted to square lattices



(a)



(b)



(c)

FIGURE 1. Three typical graphs supporting MRF-based models for image analysis: (a) rectangular lattice with first-order neighborhood system; (b) non-regular planar graph associated to an image partition; (c) quad-tree. For each graph, the grey nodes are the neighbors of the white one.



## 12. The Metropolis-Hastings algorithm and Markov Chain Monte Carlo

In our analysis of the Metropolis algorithm, we found that

$$T(C \rightarrow C') = \min \left[ 1, \exp \left( -\frac{(E' - E)}{kT} \right) \right]$$

Moreover, we found that the algorithm preserves detailed balance

$$P(C)T(C \rightarrow C') = P(C')T(C' \rightarrow C)$$

where  $P(C)$  is the stationary probability of configuration  $C$ . Indeed, at equilibrium we found that, if  $E' > E$ ,


$$P(C) \exp \left( -\frac{(E' - E)}{kT} \right) = P(C')$$

$$\frac{P(C')}{P(C)} = \exp \left( -\frac{(E' - E)}{kT} \right) \quad \leftarrow \text{ Boltzmann's distribution}$$

In summary

$$P(C) \exp\left(-\frac{(E' - E)}{kT}\right) = P(C')$$

$$\frac{P(C')}{P(C)} = \exp\left(-\frac{(E' - E)}{kT}\right)$$



$$T(C \rightarrow C') = \min\left[1, \frac{P(C')}{P(C)}\right]$$

$$T(C \rightarrow C') = \min\left[1, \exp\left(-\frac{(E' - E)}{kT}\right)\right]$$

*This definition of the transition probability is the starting point for an important further step, the Metropolis-Hastings algorithm.*

we define the transition probability – which includes a proposal function  $q$  –

$$P(\mathbf{x} \rightarrow \mathbf{y}) = q(\mathbf{x}, \mathbf{y})\alpha(\mathbf{x}, \mathbf{y})$$

and the target density

$$\pi(\mathbf{x})$$




and we take the state

$$\mathbf{X} = \mathbf{X}_n$$



next we choose randomly another state  $\mathbf{y}$  and we accept it  $(\mathbf{y} \rightarrow \mathbf{X}_{n+1})$  with probability

$$\alpha(\mathbf{x}, \mathbf{y}) = \min \left\{ 1, \frac{\pi(\mathbf{y})q(\mathbf{y}, \mathbf{x})}{\pi(\mathbf{x})q(\mathbf{x}, \mathbf{y})} \right\}$$



Note that if the proposal function  $q$  is symmetrical, then the acceptance probability takes on the simpler form

$$\alpha(\mathbf{x}, \mathbf{y}) = \min \left\{ 1, \frac{\pi(\mathbf{y}) q(\mathbf{y}, \mathbf{x})}{\pi(\mathbf{x}) q(\mathbf{x}, \mathbf{y})} \right\} \rightarrow \min \left\{ 1, \frac{\pi(\mathbf{y})}{\pi(\mathbf{x})} \right\}$$


and it depends on the target density only.

The M-H algorithm defines a Markov chain, and it is easy to show that **detailed balance holds**. The transition probability is

$$P(\mathbf{x} \rightarrow \mathbf{y}) = q(\mathbf{x}, \mathbf{y}) \alpha(\mathbf{x}, \mathbf{y}) = q(\mathbf{x}, \mathbf{y}) \min \left\{ 1, \frac{\pi(\mathbf{y}) q(\mathbf{y}, \mathbf{x})}{\pi(\mathbf{x}) q(\mathbf{x}, \mathbf{y})} \right\}$$

• case  $\frac{\pi(\mathbf{y}) q(\mathbf{y}, \mathbf{x})}{\pi(\mathbf{x}) q(\mathbf{x}, \mathbf{y})} \geq 1$

$$\begin{aligned} \Rightarrow \alpha(\mathbf{x}, \mathbf{y}) = 1; \quad \alpha(\mathbf{y}, \mathbf{x}) = \frac{\pi(\mathbf{x}) q(\mathbf{x}, \mathbf{y})}{\pi(\mathbf{y}) q(\mathbf{y}, \mathbf{x})} &\Rightarrow \begin{aligned} P(\mathbf{x} \rightarrow \mathbf{y}) &= q(\mathbf{x}, \mathbf{y}) \\ P(\mathbf{y} \rightarrow \mathbf{x}) &= q(\mathbf{y}, \mathbf{x}) \frac{\pi(\mathbf{x}) q(\mathbf{x}, \mathbf{y})}{\pi(\mathbf{y}) q(\mathbf{y}, \mathbf{x})} \end{aligned} \end{aligned}$$



$$\begin{aligned} \pi(\mathbf{x}) P(\mathbf{x} \rightarrow \mathbf{y}) &= \pi(\mathbf{x}) q(\mathbf{x}, \mathbf{y}) \\ \pi(\mathbf{y}) P(\mathbf{y} \rightarrow \mathbf{x}) &= \pi(\mathbf{y}) q(\mathbf{y}, \mathbf{x}) \frac{\pi(\mathbf{x}) q(\mathbf{x}, \mathbf{y})}{\pi(\mathbf{y}) q(\mathbf{y}, \mathbf{x})} = \pi(\mathbf{x}) q(\mathbf{x}, \mathbf{y}) \end{aligned}$$

- case  $\frac{\pi(\mathbf{y})q(\mathbf{y},\mathbf{x})}{\pi(\mathbf{x})q(\mathbf{x},\mathbf{y})} < 1$

$$\Rightarrow \alpha(\mathbf{x},\mathbf{y}) = \frac{\pi(\mathbf{y})q(\mathbf{y},\mathbf{x})}{\pi(\mathbf{x})q(\mathbf{x},\mathbf{y})}; \quad \alpha(\mathbf{y},\mathbf{x}) = 1 \quad \Rightarrow \begin{aligned} P(\mathbf{x} \rightarrow \mathbf{y}) &= q(\mathbf{x},\mathbf{y}) \frac{\pi(\mathbf{y})q(\mathbf{y},\mathbf{x})}{\pi(\mathbf{x})q(\mathbf{x},\mathbf{y})} \\ P(\mathbf{y} \rightarrow \mathbf{x}) &= q(\mathbf{y},\mathbf{x}) \end{aligned}$$


$$\Rightarrow \begin{aligned} \pi(\mathbf{x})P(\mathbf{x} \rightarrow \mathbf{y}) &= \pi(\mathbf{x})q(\mathbf{x},\mathbf{y}) \frac{\pi(\mathbf{y})q(\mathbf{y},\mathbf{x})}{\pi(\mathbf{x})q(\mathbf{x},\mathbf{y})} = \pi(\mathbf{y})q(\mathbf{y},\mathbf{x}) \\ \pi(\mathbf{y})P(\mathbf{y} \rightarrow \mathbf{x}) &= \pi(\mathbf{y})q(\mathbf{y},\mathbf{x}) \end{aligned}$$

Detailed balance holds in both cases and therefore  $\pi(\mathbf{x})$  is stationary



The following figure shows a simulation with the MCMC algorithm and the distribution

$$p(x) = \frac{0.6}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) + \frac{0.3}{\sqrt{2\pi}} \exp\left(-\frac{(x-3)^2}{2}\right) + \frac{0.1}{\sqrt{0.5\pi}} \exp\left(-\frac{(x-1)^2}{0.5}\right)$$

(a three-component mixture model)



```
nrmax = 40 000;  
  
xr = Table[0, {nrmax}];  
xr[[1]] = -4;  
  
nr = 1;  
While[nr < nrmax,  
  xtry = xr[[nr]] + RandomReal[NormalDistribution[0, 1]];  
  If[pdf[xtry] / pdf[xr[[nr]]] > RandomReal[], nr++; xr[[nr]] = xtry];  
]
```








## MCMC simulation of a 2D three-component mixture model

$$p(x, y) = \sum_{i=1}^3 \frac{\alpha_i}{\sqrt{2\pi\sigma_i^2}} \exp \left[ -\frac{(x - \mu_{x,i})^2 + (y - \mu_{y,i})^2}{2\sigma_i^2} \right]$$


$$\begin{aligned} \alpha_1 &= 0.5; & \mu_{x,1} &= 0; & \mu_{y,1} &= 0; & \sigma_1 &= 0.3; \\ \alpha_2 &= 0.3; & \mu_{x,2} &= 1; & \mu_{y,2} &= 1.; & \sigma_2 &= 0.5; \\ \alpha_3 &= 0.2; & \mu_{x,3} &= 2; & \mu_{y,3} &= 0.1; & \sigma_3 &= 0.5; \end{aligned}$$





100000 steps




100000 steps







100 steps




1000 steps



4000 steps



10000 steps








Notice that when the peaks are very narrow, the random walker may have problems visiting all of the peaks

$$p(x, y) = \sum_{i=1}^3 \frac{\alpha_i}{\sqrt{2\pi\sigma_i^2}} \exp \left[ -\frac{(x - \mu_{x,i})^2 + (y - \mu_{y,i})^2}{2\sigma_i^2} \right]$$


$$\begin{aligned} \alpha_1 &= 0.5; & \mu_{x,1} &= 0; & \mu_{y,1} &= 0; & \sigma_1 &= 0.0725; \\ \alpha_2 &= 0.3; & \mu_{x,2} &= 1; & \mu_{y,2} &= 1.; & \sigma_2 &= 0.125; \\ \alpha_3 &= 0.2; & \mu_{x,3} &= 2; & \mu_{y,3} &= 0.1; & \sigma_3 &= 0.125; \end{aligned}$$



With isolated, narrow peaks, increasing the number of steps may not suffice





100000 steps, subdivided into 10 parallel chains with random starting points




The starting points of the chains are uniformly distributed in the plot region, however the "regions of influence" of each peak vary considerably.

This leads to more chains being attracted into the lower peaks, with the result that the distribution is somewhat deformed (wrong alpha's in the mixture model)



***Many techniques have been developed to avoid these pitfalls***

# Example of application of the MCMC technique in radiobiology



**Survival curve for HeLa cells in culture exposed to x-rays.** (From Puck TT, Markus PI: Action of x-rays on mammalian cells. *J Exp Med* 103:653-666, 1956)

Phenomenology: the linear-quadratic law

$$S(D) \approx e^{-\alpha D - \beta D^2}$$




Fig. 1. Clonogenic survival curves illustrating the higher efficiency of the carbon ions compared with X-rays [10] (courtesy of the author, dr. Wilma K. Weyrather).

# Example: Target theory


## Simple Poisson model:

Probability of hitting  $n$  times a given target, when the average number of good hits is  $a$ :

$$P(n) = \frac{a^n}{n!} e^{-a}$$

Probability missing the target:  $P(0) = e^{-a}$

Average number of hits:  $a = D/D_0$



$$S(D) = P(0, D) = e^{-D/D_0}$$

## Multitarget model, asymptotic behavior and threshold effect.

If there are multiple targets, say  $n$  targets, all of which must be hit to kill a cell, then the probability of missing at least one of them – i.e., the survival probability – is

$$S(D) = 1 - (1 - e^{-D/D_0})^n$$

then, for large dose


$$S(D) \approx ne^{-D/D_0}$$

i.e.,

$$\ln S(D) \approx \ln n - D/D_0$$

which is a linear relation with intercept  $\ln n$ , and slope  $-1/D_0$ .





Notice that

$$\left[ \frac{d}{dD} e^{-\alpha D - \beta D^2} \right]_{D=0} = (-\alpha - 2\beta D) e^{-\alpha D - \beta D^2} \Big|_{D=0} = -\alpha$$

and that

$$\frac{d}{dD} \left[ 1 - (1 - e^{-D/D_0})^n \right]_{D=0} = -n \frac{e^{-D/D_0}}{D_0} (1 - e^{-D/D_0})^{n-1} \Big|_{D=0} = 0$$


The derivatives differ in the origin, and the multitarget model fails to reproduce the observed linear-quadratic law.

## The RCR (Repairable-Conditionally Repairable Damage) model

In this case the surviving fraction is

$$S = \exp(-aD) + bD \exp(-cD)$$

This is a 3-parameter expression, which is not easy to fit to data when the data set is small.



1a. Simple Gaussian likelihood for the LQ model


$$L(\alpha, \beta) = \prod_k \exp\left(-\frac{(S_k - S(\alpha, \beta))^2}{2\sigma_k^2}\right)$$


1b. Chose exponential priors for the parameters

1c. Complete posterior pdf

$$p(\alpha, \beta | \{S_k\}, I) = \left[ \prod_k \exp\left(-\frac{(S_k - S(\alpha, \beta))^2}{2\sigma_k^2}\right) \right] \exp(-0.1\alpha) \exp(-0.1\beta)$$

1d. Use MCMC to find the MAP estimate (and any moment of the pdf)





2a. Simple Gaussian likelihood for the RCR model

$$L(a,b,c) = \prod_k \exp\left(-\frac{(S_k - S(a,b,c))^2}{2\sigma_k^2}\right)$$


2b. Chose exponential priors for the parameters

2c. Complete posterior pdf

$$p(a,b,c|\{S_k\},I) = \left[ \prod_k \exp\left(-\frac{(S_k - S(a,b,c))^2}{2\sigma_k^2}\right) \right] e^{-0.2a} e^{-0.2b} e^{-0.2c}$$


2d. Use MCMC to find the MAP estimate (and any moment of the pdf)

# Path in (a,b,c) space





Fit showing individual components: unsatisfactory result



## Revise priors to include constraint on derivative

(priors vanish where derivative in the origin is positive)

