

Lecture 15: Non-Gaussian MRFs, Potential and Influence Functions

Chapter 6.1–6.2

Charles A. Bouman

Purdue University

Monday, September 28, 2026

- ➊ **Rewrite the GMRF using neighbor differences**
- ➋ **Replace the quadratic penalty by a potential ρ** ; the result is still an MRF
- ➌ **Potential and influence**: energy and force
- ➍ **Gaussian versus edge-preserving potentials**: bounded versus unbounded influence
- ➎ **Nonconvex potentials** preserve edges but can make the estimate unstable

Reading: Chapter 6.1–6.2, *Foundations of Computational Imaging*.

Today's Big Ideas

- The energy of a GMRF can be rewritten as a **sum of penalties on the differences between neighboring pixels**.
- Replacing the squared difference by a potential ρ that grows more slowly gives the **pairwise Gibbs distribution**, which is **still an MRF**.
- The potential $\rho(\Delta)$ is energy; its derivative $\rho'(\Delta)$ is the **influence**, the force one pixel exerts on another.
- The Gaussian potential has **unbounded influence**, so it smooths edges. An edge-preserving potential **bounds the influence**.
- Bounding the influence requires a **nonconvex** potential, which preserves edges but can make the MAP estimate **unstable**.

Neighborhood systems and pairwise cliques

Chapters 4 and 5 used the GMRF as an image model, but it **does not model edges** well. We generalize it, first by rewriting it in terms of pixel differences. We reuse the neighborhood system and pairwise cliques of Section 4.6.

Neighborhood system

$\partial s \subset S$ is the set of neighbors of s , and $r \in \partial s$ if and only if $s \in \partial r$.

Pairwise cliques

$\mathcal{P} = \{ \{s, r\} : s \in \partial r \}$ is the set of all unordered neighboring pixel pairs $\{s, r\}$.

The zero-mean GMRF distribution

The distribution of a zero-mean GMRF can be written as

$$p(x) = \frac{1}{z} \exp \left\{ -\frac{1}{2} x^t B x \right\} ,$$

where

- $B_{r,s} = 0$ when $r \notin \partial s \cup \{s\}$, so the neighborhood system is the zero pattern of B ;
- the **partition function** z is the normalizing constant, and $X \sim N(\mathbf{0}, B^{-1})$.

We now rewrite this to make the dependence on differences between neighboring pixels explicit.

From the quadratic form to the pairwise cliques

Property: Pairwise quadratic form identity

For any symmetric $B \in \mathbb{R}^{N \times N}$ and any $x \in \mathbb{R}^N$,

$$x^t B x = \sum_{s \in S} a_s x_s^2 + \frac{1}{2} \sum_{s \in S} \sum_{r \in S} b_{s,r} |x_s - x_r|^2,$$

where $a_s = \sum_{r \in S} B_{s,r}$ and $b_{s,r} = -B_{s,r}$.

For a GMRF, $B_{s,r} = 0$ off the neighborhood, and summing over unordered cliques removes the factor 1/2:

$$x^t B x = \sum_{s \in S} a_s x_s^2 + \sum_{\{s,r\} \in \mathcal{P}} b_{s,r} |x_s - x_r|^2.$$

We take $a_s = 0$, which makes the prior invariant to adding a constant to every pixel.

The pairwise GMRF distribution

Dropping the a_s terms gives the **pairwise GMRF** distribution


$$p(x) = \frac{1}{Z} \exp \left\{ - \sum_{\{s,r\} \in \mathcal{P}} b_{s,r} \frac{|x_s - x_r|^2}{2} \right\} . \quad (1)$$

- It is written explicitly in terms of neighbor differences.
- With $b_{s,r} > 0$, a larger squared difference $|x_s - x_r|^2$ makes the image less probable, which is reasonable since nearby pixels should be similar.

Replace the quadratic penalty by a potential function

The GMRF **over-penalizes** neighbor differences: $|x_s - x_r|^2$ becomes very large across a discontinuous boundary. Replace $|x_s - x_r|^2/2$ by $\rho(x_s - x_r)$, which grows more slowly. This defines the **pairwise Gibbs distribution**

$$p(x) = \frac{1}{z} \exp \left\{ - \sum_{\{s,r\} \in \mathcal{P}} b_{s,r} \rho(x_s - x_r) \right\}. \quad (2)$$



Taking $\rho(\Delta) = |\Delta|^2/2$ recovers the GMRF (1), so the GMRF is a special case. A potential that grows more slowly than the quadratic penalizes edges less.

Energy, potential, and influence functions

The **energy function** of the pairwise Gibbs distribution (2) is

$$u(x) = \sum_{\{s,r\} \in \mathcal{P}} b_{s,r} \rho(x_s - x_r), \text{ with } \Delta = x_s - x_r.$$

- **potential function:** $\rho(\Delta)$;
- **influence function:** its derivative $\rho'(\Delta) = \frac{d\rho(\Delta)}{d\Delta}$.

Three properties we expect of ρ (definitions)

zero at origin: $\rho(0) = 0$;

symmetric: $\rho(-\Delta) = \rho(\Delta)$;

monotone increasing: $\rho(|\Delta| + \epsilon) \geq \rho(|\Delta|)$ for all $\epsilon > 0$.

The choice of ρ determines the accuracy of the model and of the MAP estimate. Selecting it is the second half of this lecture.

The pairwise Gibbs distribution is an MRF

Definition: MRF

A random field X_s , $s \in S$, with neighborhood system ∂s , is an **MRF** (Markov random field) if for all s the conditional distribution of X_s depends only on its neighbors:

$$p_{X_s | X_{r \neq s}}(x_s | x_{r \neq s}) = p_{X_s | X_{\partial s}}(x_s | x_{\partial s}) .$$

Property: pairwise Gibbs distributions are MRFs

If X_s has a pairwise Gibbs distribution with cliques $\mathcal{P}$, then it is an MRF with $\partial s = \{r : \{r, s\} \in \mathcal{P}\}$, and

$$p_{X_s | X_{\partial s}}(x_s | x_{\partial s}) = \frac{1}{Z} \exp \left\{ - \sum_{r \in \partial s} b_{s,r} \rho(x_s - x_r) \right\} .$$

Some MRFs are **not** pairwise Gibbs distributions; the general family is larger (Chapter 14). Pairwise cliques suffice for this chapter.

Proof: the conditional distribution of a pixel

Factor $p(x)$ into the part that depends on x_s and the part that does not. The cliques containing s are exactly the pairs $\{s, r\}$ with $r \in \partial s$:

$$p(x) = \frac{1}{z} \exp \left\{ - \sum_{r \in \partial s} b_{s,r} \rho(x_s - x_r) \right\} f(x_{r \neq s}) .$$

The conditional distribution divides by the integral over x_s ; the constants $1/z$ and $f(x_{r \neq s})$ cancel:

$$p_{x_s | x_{r \neq s}}(x_s | x_{r \neq s}) = \frac{\exp \left\{ - \sum_{r \in \partial s} b_{s,r} \rho(x_s - x_r) \right\}}{\int_{\mathbb{R}} \exp \left\{ - \sum_{r \in \partial s} b_{s,r} \rho(x_s - x_r) \right\} dx_s} .$$

This depends on x_s and its neighbors only, so X_s is an MRF. The proof uses nothing about ρ , so it holds for **every** choice of potential function.


The most probable pixel, and force balance

We select ρ by studying the most probable value of a pixel x_s given its neighbors:

$$\hat{x}_s = \arg \min_{x_s \in \mathbb{R}} \sum_{r \in \partial s} b_{s,r} \rho(x_s - x_r) \implies \sum_{r \in \partial s} b_{s,r} \rho'(\hat{x}_s - x_r) = 0 .$$

- **Energy:** $\rho(x_s - x_r)$ is the potential energy of the pair.
- **Force:** $\rho'(\hat{x}_s - x_r)$ is the force x_r exerts on $\hat{x}_s$; at the solution the neighbor forces cancel.

ρ' is called the **influence function** because it sets how much a neighbor influences the estimate.




Two choices of potential function

We compare two potential functions:

$$\begin{array}{ll} \text{Gaussian potential:} & \rho(\Delta) = |\Delta|^2 ; \\ \text{weak-spring potential:} & \rho(\Delta) = \min \{ |\Delta|^2, T^2 \} . \end{array}$$

- The weak-spring potential is due to Blake and Zisserman, and it needs a threshold T (set to one for the plots).
- For each we also plot the influence function $\rho'(\Delta)$, which is the force from a neighboring pixel.

Gaussian and weak-spring: potential and influence



Left: potential $\rho(\Delta)$. Right: influence $\rho'(\Delta)$, the force. Both plotted for $T = 1$.

Bounded versus unbounded influence

Gaussian potential: unbounded influence

$\rho(\Delta) = |\Delta|^2 \implies \rho'(\Delta) = 2\Delta$. The influence grows without bound, so a pixel on the other side of an edge can have an unbounded effect on $\hat{x}_s$. This is the over-smoothing of edges.

Weak spring: bounded influence

$\rho(\Delta) = \min\{|\Delta|^2, T^2\}$ gives an influence that is **zero** for $|\Delta| > T$, so differences larger than T have no effect on $\hat{x}_s$. Across an edge this preserves the detail.

T is the threshold that separates a difference treated as noise from a difference treated as an edge.

Why nonconvex potentials define edges

A potential that is not convex is a **nonconvex function** (convexity is defined in Lecture 16). Nonconvex potentials define edges well.

The local structure

A nonconvex potential typically has


- a region $|\Delta| < T_c$ where $\rho'(\Delta)$ increases (ρ locally convex), and
- a region $|\Delta| > T_c$ where $\rho'(\Delta)$ decreases (ρ locally concave).

These two regions delineate the “edge” and “nonedge” cases.

Every nonconvex potential has an influence that must decrease somewhere. But nonconvex potentials have a disadvantage, shown next.


The disadvantage: an unstable MAP estimate

Noisy Signals



(a)

Unstable Reconstructions



(b)

(a) two slightly different input signals y ; (b) their MAP estimates using the nonconvex weak-spring model. A slight change in y produces a discontinuous jump in $\hat{x}$.

The discontinuity is in the **estimator**, not the image: a small change in y can move $\hat{x}$ by a large amount. This is useful for edge detection, but undesirable when $\hat{x}$ is viewed as an image.

Summary

- The GMRF energy is a sum over cliques, $\sum_{\{s,r\} \in \mathcal{P}} b_{s,r} |x_s - x_r|^2$; replacing $|x_s - x_r|^2/2$ by a potential ρ gives the pairwise Gibbs distribution, which is **still an MRF**.
- $\rho(\Delta)$ is energy and $\rho'(\Delta)$ is influence (force); the most probable pixel is where the neighbor forces balance, $\sum_{r \in \partial s} b_{s,r} \rho'(\hat{x}_s - x_r) = 0$.
- The Gaussian potential $\rho = |\Delta|^2$ has **unbounded** influence and smooths edges; the weak spring $\min\{|\Delta|^2, T^2\}$ has influence zero for $|\Delta| > T$ and **preserves** edges.
- Bounded influence requires a **nonconvex** potential, which preserves edges but can make the MAP estimate unstable.

Reading: Chapter 6.1–6.2. Next lecture: convex potentials (GGMRF, Huber, QGGMRF).