

# Identification of Solar Features via Markov Random Fields

Michael J. Turmon  
M/S 525-3660  
Jet Propulsion Laboratory  
Pasadena, CA 91109

## Abstract

The solar chromosphere, seen in ultraviolet light, consists of roughly three classes: plage (bright magnetic disturbances), network (hot boundaries of convection cells), and background (cooler interiors of such cells). These classes contribute differently to the total amount of ultraviolet radiation reaching Earth's upper atmosphere. It is of scientific interest (e.g., in studying global warming) to relate plage and network area and intensity to this total UV irradiance. We treat this problem as inference of the underlying class of a pixel based on the observed intensity, perhaps seen through a point-spread function. The spatial information known to solar physicists is encoded in a Markov random field (MRF) smoothness prior and used in a Bayesian inference setup. Also of interest is understanding plage shape and evolution. To this end, a plage structure is represented as a planar graph with associated heights forming a spatial 'membership function'. This membership function biases the underlying MRF so that plagues occur in a physically reasonable way. Sample results are shown.

## 1 Introduction

The solar chromosphere, observable in ultraviolet light, roughly consists of three classes: plage (bright magnetic disturbances), network (hot boundaries of convection cells), and background (cooler interiors of these cells). Figure 1 shows a full image and detail examples of the classes. Plages appear as irregular groups of clumps, seldom near the solar poles. It is believed [14, p. 319] that the same magnetic disturbances that are coincident with sunspots also underlie the plage phenomenon. Similar to sunspots, plages experience a cycle of formation and dissipation, starting out as a relatively compact regions and decaying over many days into a diffuse and broken-up cluster [18, p. 317]. The cell-structured network has little contrast with the background, is spatially

homogeneous, and persists for tens of hours. The cells (difficult to see in this halftoned rendering) have a characteristic size, and it is thought that they arise due to convective processes in the plasma making up the solar atmosphere [18, p. 126].

The three classes contribute differently to the ultraviolet (UV) radiation reaching Earth's upper atmosphere, with the plagues and magnetic network giving the largest contribution. While this radiation cannot be sensed directly from the ground, the features giving rise to it can be. Such measurements are inputs to models of solar irradiance [17] which are crucial to understanding phenomena such as global warming and photochemical processes in the upper atmosphere [16].

Also of interest is the evolution of plagues. As mentioned above, a typical sequence has been described: from plage emergence as a shape of relatively smooth boundary, to expansion, and then eventual dissolution as an irregular, tentacled form. However, the understanding is of a qualitative and anecdotal sort (e.g., for related work, [14, p. 284]), and a more quantitative description of anticipated plage shapes and the evolution of plage regions would be of value.

In both sorts of problem described above one reduces a series of images, comprising perhaps 500MB of data, to a time series of areas and intensities or plage descriptions. For example, per-class areas and mean intensities can be represented in about 48 bytes per 9MB image. The description of the features of interest for one plage might take a few KB per image, depending on the activity level. These time series distill the diffuse information in the large data set into a scientifically usable form, at least for the tasks at hand.

### 1.1 Data

The primary source of data for this study is the set of CaII K full-disk spectroheliograms that has been collected, daily or as observing conditions permit, at Sacramento Peak National Solar Observatory in Sunspot, NM from the mid-sixties onward. (The 'K line' is the spec-

<sup>1</sup>This work was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.




Figure 1: The left panel shows a full-disk chromospheric image from 15 July 1992. The film calibration wedge (see text) is visible at top left. In the center is a detail image of the rather decayed plage pair in the northwest quadrant of the sun. The southeast quadrant has a younger, more concentrated plage. On the right is a detail of the center of the disk, with considerable contrast enhancement, showing just the network and background.

tral absorption line at wavelength 393.4nm. Through an involved physical process, the chromosphere can be observed there in isolation from other layers of the solar atmosphere. At most frequencies the visible light emitted from the photosphere dominates.) The images take in the entire Sun as seen from Earth, and they are recorded on photographic films which are stored at the observatory. An interval of these films, from the mid-eighties forward, has been digitized to  $2K \times 2K$  pixels, at which point atmospheric blurring (and to a lesser degree, telescope diffraction) limits resolution.

Other sources of information about solar active regions are available, but have not been used directly in this study. Of most importance are magnetograms, which are high-resolution images of the line-of-sight component of the magnetic field on the sun. Since plages are a reflection of an underlying magnetic disturbance, practiced observers use magnetograms in conjunction with K-line images to determine the extent of a plage. There are also optical images from other spectral lines — notably  $H\alpha$  at 656.3nm, which encodes processes at a higher level in the chromosphere than the K-line images — and ‘continuum’ images which are not spectrally limited and contain photospheric information.

Observations from these sources are taken regularly but they are not coordinated: they are taken at different times and from different parts of the planet. A system called the PSPT (Precision Solar Photometric Telescope) has been developed [5] to standardize the ob-




Figure 2: Schematic of the observation process showing origin of energy in the chromosphere, attenuation by limb-darkening, blurring by PSF, attenuation by atmosphere and inconsistent exposure time, and nonlinear transformation to film density.

servaion regimen; nevertheless, the historical database embodied in the Sac Peak K-line images will remain invaluable to researchers.

## 1.2 Preprocessing

On its way from the sun to the film, the chromospheric emission is transformed by several processes shown in figure 2. In principle these can be incorporated in the inference mechanism, but for simplicity and computational reasons, it is more practical to invert them separately, in reverse order from film backward.

1. Film responds nonlinearly to incident light, say through a monotone function  $\phi$ , so the digitized film density level must be translated into a light intensity. This amounts to application of the monotone transformation  $\phi^{-1}$  individually to each pixel in the image. To recover the transformation, there is a calibrated step wedge in a corner of each image. Each segment of the wedge has been exposed at a known intensity level. This fixes  $\phi$  at (in this case) eight points and allows it to be inverted.
2. Atmospheric clarity and film exposure times differ from day to day, resulting in a multiplicative attenuation  $\alpha$ . The approximate scale factor on each day relative to some reference day can be obtained by finding the intensity of the ‘quiet sun’, or those regions of the sun without any transient activity. Each image is rescaled to compensate for this temporal variation.
3. Apart from the transparency changes described above, the observed image is subject to atmospheric scattering and blurring which varies daily depending on observation conditions. This distortion is well-modeled by gaussian or lorentzian point-spread functions. Currently, the PSF is not taken in to account, although it can be estimated from the image set [15]. (Since the sun is not a point source, this determination is not trivial.)
4. Solar emissions are affected by ‘limb-darkening’, a spatially varying attenuation  $LD(\cdot)$  which is unity at disc center and decreases gradually as one moves to the edge of the disc. This is due to changes in the apparent width of the chromospheric shell as seen from Earth (see fig. 2). Theory suggests [14, p. 133] a polynomial dependence of LD on  $\mu = \cos \theta$ , where the angle  $\theta = 0$  at disc center and  $\pi/2$  at the edge. Consequently we employ several images to estimate the coefficients of a quadratic attenuation function, constrained to  $LD(\mu = 1) = 1$ , and rescale the image accordingly. For details, see for example [13].

The result of these preprocessing steps is a pixel-by-pixel transformation of the original image into a ‘flattened’ image.

## 2 Image Decomposition

First we discuss the problem of partitioning the image into plage, network, and background components. Generally, scientists either apply a threshold across the flattened image to determine plage areas, or manually surround the plages with polygons. The first method, while

simple and objective, ignores all spatial information that is available. The second method clearly uses a large amount of side information possessed by the scientists, but is also highly subjective, difficult to even describe, and hard to repeat.

Due to the strong prior information available to us about the images, we adopt the Bayesian framework of inference of underlying pixel classes based on the observed intensity. Denoting pixel sites  $s = [s_1 s_2]$  in an image domain  $N$ , and defining matrices of class labels  $\mathbf{x} = \{x_s\}_{s \in N}$  and observed intensities  $\mathbf{y}$ , the posterior probability of labels given data is

$$P(\mathbf{x} | \mathbf{y}) = P(\mathbf{y} | \mathbf{x})P(\mathbf{x})/P(\mathbf{y}) \propto P(\mathbf{y} | \mathbf{x})P(\mathbf{x}) . \quad (1)$$

The maximum a posteriori (MAP) rule maximizes this probability:

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} \log P(\mathbf{y} | \mathbf{x}) + \log P(\mathbf{x}) . \quad (2)$$

As pointed out by Ripley [11], the first term is the familiar likelihood function, telling how the data is gotten from the labels; the second is the prior probability of a given labeling. In practice, the first term forces fidelity to the data while the second penalizes unlikely rough labelings.

Prior models may be specified in many ways; we have used the Markov field models introduced by Besag [2] and Geman and Geman [6]. These models are characterized by

$$P(x_s = k | x_{N \setminus \{s\}}) = P(x_s = k | x_{N(s)}) \quad (3)$$

with respect to a neighborhood structure  $N(s) \subset N$ , where  $\setminus$  is set difference. This simply expresses the Markov property that far-off sites do not influence the distribution of labels when the local neighbors are known. To define such a model, one needs to specify these conditional distributions and a typical choice is

$$P(x_s = k | x_{N(s)}) = Z_s^{-1} \exp[-\beta \sum_{s' \in N(s)} 1(x_{s'} \neq k)] \quad (4)$$

where  $N(s)$  is taken as the 8-pixel neighborhood centered around a site  $s$ , and  $Z_s$  is a constant chosen to make the distribution sum to unity. For  $\beta = 0$ , the distribution is uniform on all  $3^{\text{card}(N)}$  labelings, and as  $\beta$  is increased, smoother and smoother labelings are favored. For example, the right-hand panel of figure 3 shows a typical (random) image from the MRF prior  $P(\mathbf{x})$  (no data is used in generating it). While this image does not precisely match the expected plage/network pattern, the match is much better than the field of independent labels at each site shown on the left. We discuss priors more tailored to this application in section 4.




Figure 3: On the left is a typical labeling drawn from the uniform prior on  $\{1, 2, 3\}^N$ . On the right is one drawn from an iterative approximation to the MRF prior with smoothness  $\beta = 0.7$ .

The remaining ingredient is the likelihood

$$P(\mathbf{y} | \mathbf{x}) = \prod_{s \in N} P(y_s | x_s) \quad (5)$$

where we assume that intensities are independent conditional on the labels being known. The three densities  $P(y | x = k)$  can be estimated from labeled data supplied by scientists. It is not surprising that the plage and network intensities have a heavy tail — making a normal distribution inappropriate. Kolmogorov-Smirnov and Cramér-von Mises distributional tests [9] confirm that the lognormal distribution is a good model for the per-class intensities.

After finding these distributions, we may use the Hammersley-Clifford theorem [8, thm. 2.2.1] to rewrite the cost function of (2) as

$$-\sum_{s \in N} \left( \frac{(\log y_s - \mu_{x_s})^2}{2\sigma_{x_s}^2} + \log \sigma_{x_s} \right) - \beta \sum_{s, s' \text{ nbrs}} 1(x_s \neq x_{s'}) \quad (6)$$

where  $N(s)$  determines the neighborhood relation in the last sum. The tradeoff between consistency of each observed intensity with the mean of its assigned class, and agreement of neighboring class labels, is apparent. We also note that if the class variances are identical, and  $\beta = 0$ , we recover the threshold rule currently used in practice.

However, with  $\beta > 0$ , the optimization becomes coupled across sites, and is entirely intractable for our three-class problem. To tackle this problem we have followed the well-known numerical method known as the Gibbs sampler, detailed in [6] (see also [11, p. 99]). In brief, this works by cycling through each site, computing

$P(x_s | y_s, x_{N(s)})$  for each class, and choosing the next label from this distribution. For finite label spaces, the resulting (random) sequence of labelings converges in distribution to the posterior. To extremize the posterior, one ‘sharpens’ the distribution by decreasing a scale parameter slowly to zero, and the resulting labeling is the MAP estimate.

### 3 Initial Results

Sample results are shown in figure 4. The first panel shows a piece of a chromospheric image from January 1980 with a plage in the lower-right corner. Beside this is the corresponding threshold segmentation. The abundant speckle is consistent with the implicit prior that is uniform over all labelings. In the final panel is the MAP segmentation with MRF prior at  $\beta = 0.7$ . The estimate is found by the standard Gibbs sampler approach with temperature lowered in steps over 800 image sweeps. We note that the MAP/MRF segmentation eliminates many of the tiny gaps in the large plage and makes the network structure more apparent.

The segmentations are a milestone in this part of the effort in solar image processing at JPL. The goal is to determine model parameters valid across several months’ worth of images, to find labelings for each such image, and thus obtain a time series of overall areas and average intensities for each of the three classes. This information can then be used as input to existing solar irradiance models and validated against measurements from space-borne UV monitoring instruments such as UARS/Solstice [12].

### 4 Spatial Descriptions

Now we address the second of the concerns raised in the introduction, that of representing and analyzing plage and network shape. The network shape is the simpler because it is homogeneous across the solar sphere. One way to model the network/background relationship is via a binary MRF texture model [4] of the attraction-repulsion type. In such a model the conditional distributions corresponding to (4) become, for  $k = 1, 2$ ,

$$P(x_s = k | x_{N(s)}) = Z_s^{-1} \exp \left[ \alpha_k + \sum_{s' \in N(s)} \beta(s' - s) 1(x_{s'} = k) \right] \quad (7)$$

where the radially symmetric kernel  $\beta(\cdot)$  is positive near the origin and negative near the edge of the neighborhood, and the constants  $\alpha_k$  are chosen to adjust overall




Figure 4: On the left is a detail of the original image, containing network and plage elements. The center panel shows the threshold segmentation and on the right is the MRF segmentation with  $\beta = 0.7$ .

class probabilities. (The background/network interplay of the model (4) is obtained with a kernel of constant value on  $N(s)$ .) The effect is to encourage clumps of equal labels, but to encourage those labels to change at some distance from the center; the kernel essentially acts as a template of class agreements.

In contrast to the essentially pixel-scale characteristics of the magnetic network, plages are high-level phenomena which are not well-captured by pixel-level rules. Following the lead of Grenander [7], we therefore pursue a hierarchical representation of plages. We will find it convenient to embed the pixel sites  $N$  in a bounded continuum  $\bar{N} \subset R^2$ . To represent a plage, or a cluster of related plages, we propose a tent-like structure defined by a triangulated planar graph

$$\begin{aligned}
 G &= (V, E, h) & (8) \\
 V &\subset \bar{N} & \text{a vertex set} \\
 E &\subset \bar{N}^2 & \text{an edge relation} \\
 h: V &\rightarrow [0, 1] & \text{a height function}
 \end{aligned}$$

The height function is extended to all of  $\bar{N}$  as follows. Each point in  $\bar{N}$  is either on an edge of  $G$ , in the interior of exactly one triangle of  $G$ , or outside all of the triangles. In the first two cases, linear interpolation is used to find the height of the point, otherwise the height is taken as zero. This structure is intended to model the “degree of membership” of a given pixel in the plage class, allows the binding of nearby plage regions into one coherent object. We note that, if the height function is thresholded at a given level, the resulting shape is a cluster of regions bounded by (not necessarily convex) polygons. This is similar to the way scientists currently delimit plage regions manually. An example of a plage graph is in figure 5.

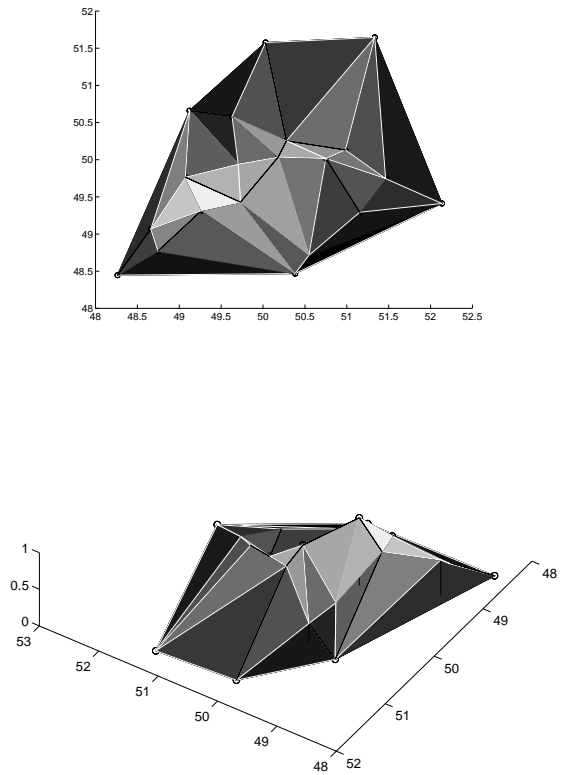


Figure 5: The first panel shows an air view of a ridge structure. The next is a perspective view of the same structure; the vantage point is at the northwest corner looking southeast.




Figure 6: The first panel shows a height function derived from a randomly generated plage graph. In the second is a realization of a random labeling implied by that height function, and finally a simulated solar image.

To define a probability distribution on these structures, we generate each as the Delaunay triangulation [1] of independently chosen points in  $\bar{N}$ . These points comprise  $V$ , and  $E$  is generated mechanically as the Delaunay triangulation of  $V$ . The heights can then be assigned independently to the interior members of  $V$ , and set to zero otherwise. The probability density of such a planar graph is then

$$P(G) = Z^{-1} e^{-\gamma \text{card}(V)} \quad (9)$$

and zero if the graph is not generated as such a triangulation. We have assumed the members of  $V$  are chosen according to the uniform distribution on  $\bar{N}$ , and that the heights are uniform on  $[0, 1]$ . While refinements (such as self-avoiding vertices and correlated heights) are possible, their ultimate effect in the presence of data would be minimal and not worth the added model complexity. Another advantage of this distribution is that additions, deletions, and adjustments of one vertex have a simple effect on the cost, and a local effect on the triangulation and the resulting cost function.

Incorporating the new structure into the existing MRF model is simple. The graph yields an extended height function  $h_G$ , which in turn becomes a spatially-varying bias on the smoothness prior (4):

$$P(x_s = 3 \mid x_{N(s)}, G) = Z_s^{-1} \exp \left[ -\beta \sum_{s' \in N(s)} \mathbf{1}(x_{s'} \neq k) + \phi(h_G(s)) \right] \quad (10)$$

Here  $\phi$  is a smooth, monotone function; below, we have simply used the identity.

The combination of the conditional probability models (7) and (10) allows us to generate random height functions, labelings, and images that are more physically

reasonable than the initial MRF scheme. Figure 6 shows a height function  $h_G(s)$  corresponding to a randomly-chosen ten-vertex graph. A corresponding pixel-level MRF labeling is in the next panel; it was obtained after 100 iterations of the Gibbs sampler applied to (10), again with  $\beta = 0.7$ . The last frame is an observation from the full image model. Intensities have been assigned to each pixel based on the lognormal model fitted to observed histograms and the result blurred with a gaussian PSF of full-width half-max of 7 pixels, reasonable for this image set. These images are intended to give a qualitative understanding of the potential of the models discussed here, as well as crude confirmation that the results are physically reasonable.

Essentially the same procedure outlined in section 2 can be followed for inference. Now the labels together with the height function or plage graph are to be inferred. There are two types of movement in the parameter space: updating a label, which is done by a Gibbs sampler step, and altering the plage graph, which can be done by a Metropolis step [3]. Such a step proceeds by proposing a small parameter change and choosing to make the jump randomly depending on the computed ratio of probabilities. A hierarchical overall updating scheme can be employed, in which the class labels are inferred given the data, and the plage graph gives structure to the class labels. Computationally efficient means of interleaving these steps are now under development.

## 5 Conclusions

Decomposing solar images into their component parts and understanding the evolution of solar active regions are questions of scientific interest, and years-long image

datasets are available from several sources for analysis. Currently, scientists often label images manually, or by thresholding the observed intensities. The use of MRF image priors allows the controlled, objective incorporation of simple kinds of prior knowledge about the spatial coherence of labels. By incorporating these priors in a Bayesian inference setup, images are segmented without the speckle artifacts associated with threshold labeling. To aid scientists in understanding this data set and allow analysis of these images, we have linked the algorithms described in sections 1.2 and 2 to a graphical interface called StarTool (Solar feature Tracking, Analysis, and Representation Tool). The application, built on top of SAOtng package [10], allows interactive loading and examination of images in several formats including the standard astronomical FITS format. Once loaded, images can be preprocessed and labeled by scientists, as well as decomposed via the MRF setup described previously. This provides an extensible interface for scientists to examine, label, and segment solar images.

In an effort to understand the temporal evolution of plage shapes, we have proposed a representation of active regions in terms of a triangulated graph which gives rise to a membership function. The resulting synthesized images appear physically realistic, and methods are being developed to infer the labels and the graph from observed data.

## Acknowledgments

Thanks to Judit Pap of the Department of Physics and Astronomy at UCLA for supplying the data, labeling images, and describing the background of the problem. Thanks are also due to Eric Mjolsness of UCSD for suggestions on image representation and to Saleem Mukhtar of JPL for help with the image processing algorithms.

## References

- [1] F. Aurenhammer. Voronoi diagrams — A survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345–405, 1991.
- [2] J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Jour. Royal Stat. Soc. Ser. B*, 36:192–236, 1974.
- [3] J. Besag, P. Green, D. Higdon, and K. Mengersen. Bayesian computation and stochastic systems. *Statistical Science*, 10:3–66, 1995.
- [4] G. R. Cross and A. K. Jain. Markov random field texture models. *IEEE Trans. Patt. Anal. Mach. Int.*, pages 25–39, 1983.
- [5] P. Foukal and J. Kuhn. NSF initiative: Radiative inputs of the Sun to Earth. In *Solers-22 Conference, Sunspot, NM*, June 1996.
- [6] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Patt. Anal. and Mach. Intell.*, 6:721–741, 1984.
- [7] U. Grenander, Y. Chow, and D. Keenan. *Hands: A Pattern-Theoretic Study of Biological Shapes*. Springer, 1991.
- [8] X. Guyon. *Random Fields on a Network: Modeling, Statistics, and Applications*. Springer, 1995.
- [9] Y. Nikitin. *Asymptotic Efficiency of Nonparametric Tests*. Cambridge Univ., 1995.
- [10] Smithsonian Astrophysical Observatory. The SAOtng software package. Available for ftp at <ftp://sao-ftp.harvard.edu/pub/rd/>.
- [11] B. D. Ripley. *Statistical Inference for Spatial Processes*. Cambridge Univ., 1988.
- [12] G. J. Rottman, T. N. Woods, and T. P. Sparn. Solar stellar irradiance comparison experiment I: Instrument design and operation. *J. Geophys. Res.*, 98:667–677, 1993.
- [13] M. Steinegger, P. N. Brandt, and H. F. Haupt. Sunspot irradiance deficit, facular excess, and the energy balance of solar active regions. *Astron. and Astrophys.*, 310:635–645, 1996.
- [14] M. Stix. *The Sun: An Introduction*. Springer, 1991.
- [15] M. J. Turmon. Empirically determining point-spread functions. Technical report, JPL Machine Learning Systems Group, 1996. Available at <http://www-aig.jpl.nasa.gov/~mjt/>.
- [16] G. L. Withbroe and W. Kalkofen. Solar variability and its terrestrial effects. In *The Sun as a Variable Star: Solar and Stellar Irradiance Variations*, pages 11–19. Cambridge Univ., 1994.
- [17] J. Worden. *A Three Component Proxy Model for the Solar Far Ultraviolet Irradiance*. PhD thesis, Univ. Colorado, Boulder, CO, 1996.
- [18] H. Zirin. *Astrophysics of the Sun*. Cambridge Univ., 1988.