Segmenting Chromospheric Images with Markov Random Fields

Michael J. Turmon (mjt@aig.jpl.nasa.gov)
Jet Propulsion Laboratory, Pasadena, CA 91109
Judit M. Pap
UCLA Dept. of Astronomy, Los Angeles, CA 90095

ABSTRACT The solar chromosphere roughly consists of three types of region: plage, network, and background. Thresholding individual pixel intensities is typically used to identify these regions in solar images. We have incorporated spatial information by using a Bayesian setup with an image prior that prefers spatially coherent labelings; resulting segmentations are more physically reasonable. These priors are a first step in developing an appropriate model for chromospheric images.

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 cells). Plages appear as irregular groups of clumps, seldom near the solar poles. The cell-structured network has little contrast with the background, and is spatially homogeneous. The classes contribute differently to the UV 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 total UV irradiance. To do this, spatially resolved images are needed.

We treat this problem in a Bayesian framework as inference of the underlying pixel classes based on the observed intensity. Denoting pixel sites $s \in N$, 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

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

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

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

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.

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

Prior models may be specified in many ways; we have used the Markov field models introduced by Besag and others for image analysis. These models are defined by the conditional distributions

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

where $N(s)$ is the 8-pixel neighborhood centered around a site s, and $Z(s)$ is a constant chosen to make the distribution sum to unity. The first equality expresses the Markov property that far-off sites do not influence the distribution of labels when the local neighbors are known, while the second favors 'smooth' labelings. Priors more tailored to this application can be built; of some interest is capturing the network structure.

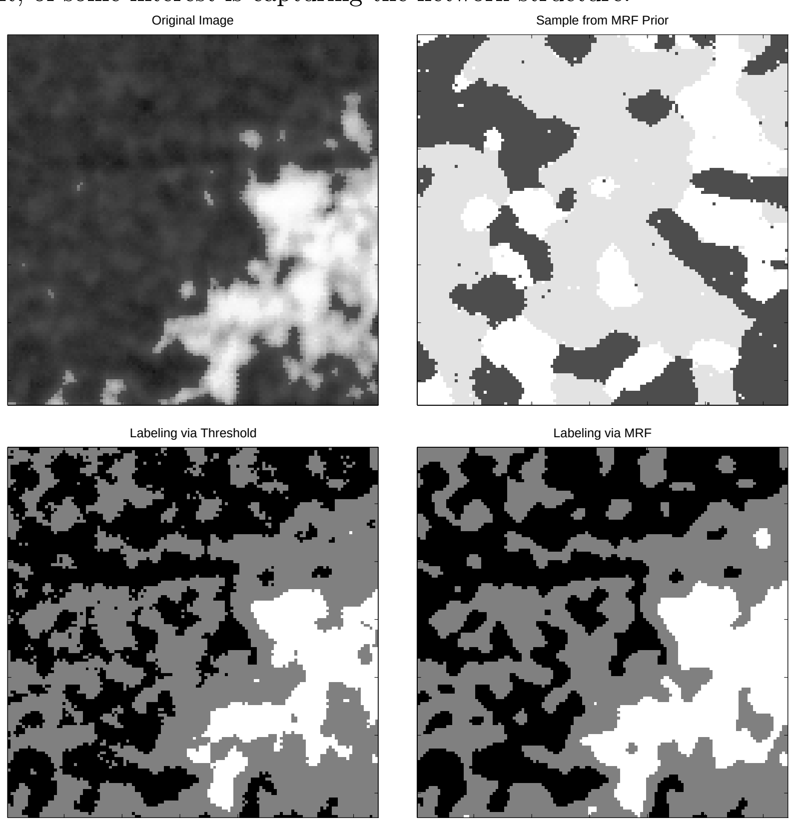

Sample results are shown above. The first panel shows a piece of a chromospheric image from January 1980 with a plage in the lower-right corner. Below this is the corresponding threshold segmentation. The top-right panel shows a typical (random) image from the MRF prior $P(\mathbf{x})$ (no data is used in generating it). While this image does not precisely match any expected plage/network pattern, the match is much better than a field of independent labels at each site. The MAP/MRF segmentation is in the final panel; we note that it has eliminated many of the tiny gaps in the large plage and made the network structure more apparent.