Scott McCloskey COMP 766
Paper Critique: Hummel and Zucker,
"Foundations of Relaxation Labeling Processes"

1 The Big Picture

The biggest contribution of this paper is the presentation of a framework within which labeling problems can be analyzed. They begin by defining the objectives and components of such problems, and rigorously define them. Building on this, they define both support and consistency, allowing for a flexible definition of the former. The objective of constructing a consistent labeling is tackled in later sections, using two parallel methods. By relating the search for a (possibly maximal) consistent labeling to the classical method of gradient descent, they bring a substantial body of knowledge to bear on the problem. They count this is a primary achievement, as it eliminates the need for heuristic or ad hoc methods that - according to the authors - characterize previous labeling methods.

2 The Gory Details

Both the concepts and notation of the paper are introduced in two parallel tracks, one each for the discrete and continuous labellings. For the sake of this summery, I shall present the discrete case first. Essentially, given n objects and m potential labels, a candidate solution takes the form of an n-vector of labellings, each of which is an m-vector of binary values. Moreover, because each object can have only one label assigned to it, each m-vector will have exactly one non-zero value. In order to make life interesting, a set of constraints are given which express the peculiarities of the problem. Between any two objects, these restraints are given as a binary matrix, where, for example, the 3^{rd} element in the 4^{th} row indicates whether the first object can take the third label when the second object can take the fourth label.

In order to build some intuition for this matrix, let's consider the map-coloring problem of graph theory. Recalling the problem's one constraint is that adjacent countries not be the same color (read: have the same label), we would have two distinct sets of constraint matrices: (1) for non-adjacent countries, the matrix would contain all 1s and (2) for bordering countries, the matrix would be all 1s except for the main diagonal, which would be all 0s.

We now have the required information to describe the discrete relaxation process based on the discarding rule. If we start with a labeling matrix containing all ones, we can apply the restrictions - replacing the 1s with 0s - in an iterative fashion. The end result will be a matrix that observes the restrictions, though not necessarily a solution to the problem. It would be possible to under- or over-constrain the problem in such a way that the resulting matrix wouldn't have exactly one label for each object.

Moving along, the authors define the support of a given label at a given object as the number of neighboring objects which *could* be labeled in such a way that the two would be consistent. In terms of the map coloring problem, we can think of the support for coloring a country red as a count of its neighbors who can be colored something other then red.

While the discrete case can be used to explain the map coloring problem, it is insufficient to describe all labeling problems. In order to do so, it is necessary to introduce continuity to both the assignments of labels, as well as to the constraints. In the latter case, we graduate from a simple allow/disallow constraint to one that expresses indifference (zero) or differing magnitudes of approval (positive values), and disapproval (negative). Note that this eliminates the need for the notion of a neighborhood, as non-neighboring objects can have compatibility matrices that contain all zeros.

Once the authors have presented the concepts with respect to the continuous case, they introduce the notion of average local consistency of a labeling $\bar{p}$, called $A(\bar{p})$. After noting that, in general, maximizing this value does not maximize its individual components, the authors present a special case where this does work. Namely, when the constraint matrices are symmetric, it is the case that (local) maxima of $A(\bar{p})$ correspond to consistent labellings as they define them.

The remainder of the paper is spent describing how gradient descent can be used to maximize $A(\bar{p})$, and what the search space looks like. After this, they present a method whereby consistent labellings can be found in the event that the matrices are asymmetric.

3 The Scott Opinion

Stylistically, this paper is very well written. As with any paper whose title is of the form "Foundations of X", there are a large number of definitions and notational issues that need to be handled upfront. These are defined clearly, and in such a way that someone with a particular labeling problem in mind can easily relate to them. The layout of the paper is both clear an explicit, making it easy for the reader to remind himself of certain definitions when they come up later on.

Technically, I'd prefer further details in several areas of the paper. In particular, they don't mention what gradient descent looks like in the discrete case. Given the the discrete solutions are a subset of the continuous solutions, how does gradient descent compare to the discarding rule? Moreover, is there an analogy to the discarding rule in the continuous case? The parallel development of the discrete and continuous problems comes and goes, and in some cases more explicit comparisons would be useful to the reader.

While they might have added more information comparing the discrete and continuous cases, the authors did a good job by leaving out discussion of implementation issues. The paper claims to be concerned with the foundations of a certain type of solutions, and is rightfully limited to a fairly high-level handling of the issues. Given the large body of knowledge about gradient descent, any additional handling of the material would be unnecessary. As a reader moderately interested in the history of science, though, I can't help but to wonder how that body of knowledge has increased since this paper's 1983 publication.

Given that the framework allows labeling problems to be expressed as a search in some high-dimensional space, it seems that a brief mention of other iterative solutions. Simulated Annealing, for instance, is one method that applies to the same types of problems as gradient descent. Without presenting too much detail on the topic (I don't want to be hypocritical with respect to my last point), the characterization of the assignment space might have more explicitly highlighted those characteristics that inform the choice of search method.

While not straying too far from the main thrust of the paper, it would have been nice to include an example of the ad hoc and/or heuristic approaches that were used previously. Since the authors count as their main accomplishment the elimination of such approaches, it seems sensible to mention at least one. Such an example might have been unnecessary for readers with a background in labeling problems, but would be helpful for others.

On the whole, this was an excellent paper. It did not presuppose a familiarity with labeling problems, as is appropriate with a foundations paper, and was appropriately generalized. Given that size in such papers is always limited, the omission of what I've mentioned above is understandable.