What Energy Functions Can Be Minimized via Graph Cuts?

Vladimir Kolmogorov and Ramin Zabih

Computer Science Department, Cornell University, Ithaca, NY 14853
vnk@cs.cornell.edu, rdz@cs.cornell.edu

Abstract. In the last few years, several new algorithms based on graph cuts have been developed to solve energy minimization problems in computer vision. Each of these techniques constructs a graph such that the minimum cut on the graph also minimizes the energy. Yet because these graph constructions are complex and highly specific to a particular energy function, graph cuts have seen limited application to date. In this paper we characterize the energy functions that can be minimized by graph cuts. Our results are restricted to energy functions with binary variables. However, our work generalizes many previous constructions, and is easily applicable to vision problems that involve large numbers of labels, such as stereo, motion, image restoration and scene reconstruction. We present three main results: a necessary condition for any energy function that can be minimized by graph cuts; a sufficient condition for energy functions that can be written as a sum of functions of up to three variables at a time; and a general-purpose construction to minimize such an energy function. Researchers who are considering the use of graph cuts to optimize a particular energy function can use our results to determine if this is possible, and then follow our construction to create the appropriate graph.

1 Introduction and Summary of Results

Many of the problems that arise in early vision can be naturally expressed in terms of energy minimization. The computational task of minimizing the energy is usually quite difficult, as it generally requires minimizing a non-convex function in a space with thousands of dimensions. If the functions have a restricted form they can be solved efficiently using dynamic programming [2]. However, researchers typically have needed to rely on general purpose optimization techniques such as simulated annealing [3,10], which is extremely slow in practice.

In the last few years, however, a new approach has been developed based on graph cuts. The basic technique is to construct a specialized graph for the energy function to be minimized, such that the minimum cut on the graph also minimizes the energy (either globally or locally). The minimum cut in turn can be computed very efficiently by max flow algorithms. These methods have been successfully used for a wide variety of vision problems including image restoration [7,8,12,14], stereo and motion [4,7,8,13,16,20,21], voxel occupancy [23], multi-camera scene reconstruction [18] and medical imaging [5,6,15]. The output of

A. Heyden et al. (Eds.): ECCV 2002, LNCS 2352, pp. 65–81, 2002.

these algorithms is generally a solution with some interesting theoretical quality guarantee. In some cases [7,12,13,14,20] it is the global minimum, in other cases a local minimum in a strong sense [8] that is within a known factor of the global minimum. The experimental results produced by these algorithms are also quite good, as documented in two recent evaluations of stereo algorithms using real imagery with dense ground truth [22,24].

Minimizing an energy function via graph cuts, however, remains a technically difficult problem. Each paper constructs its own graph specifically for its individual energy function, and in some of these cases (especially [8,16,18]) the construction is fairly complex. The goal of this paper is to precisely characterize the class of energy functions that can be minimized via graph cuts, and to give a general-purpose graph construction that minimizes any energy function in this class. Our results play a key role in [18], provide a significant generalization of the energy minimization methods used in [4,5,6,8,12,15,23], and show how to minimize an interesting new class of energy functions.

In this paper we only consider energy functions involving binary-valued variables. At first glance this restriction seems severe, since most work with graph cuts considers energy functions that involve variables with more than two possible values. For example, the algorithms presented in [8] for stereo, motion and image restoration use graph cuts to address the standard pixel labeling problem that arises in early vision. In a pixel labeling problem the variables represent individual pixels, and the possible values for an individual variable represent, e.g., its possible displacements or intensities. However, many of the graph cut methods that handle multiple possible values actually consider a pair of labels at a time. Even though we only address binary-valued variables, our results therefore generalize the algorithms given in [4,5,6,8,12,15,23]. As an example, we will show in section 4.1 how to use our results to solve the pixel-labeling problem, even though the pixels have many possible labels. An additional argument in favor of binary-valued variables is that any cut effectively assigns one of two possible values to each node of the graph. So in a certain sense any energy minimization construction based on graph cuts relies on intermediate binary variables.

1.1 Summary of Our Results

In this paper we consider two classes of energy functions. Let $\{x_1, \ldots, x_n\}$, $x_i \in \{0, 1\}$ be a set of binary-valued variables. We define the class $\mathcal{F}^2$ to be functions that can be written as a sum of functions of up to 2 variables at a time,

$$E(x_1, \ldots, x_n) = \sum_i E^i(x_i) + \sum_{i<j} E^{i,j}(x_i, x_j). \tag{1}$$

We define the class $\mathcal{F}^3$ to be functions that can be written as a sum of functions of up to 3 variables at a time,

$$E(x_1, \ldots, x_n) = \sum_i E^i(x_i) + \sum_{i<j} E^{i,j}(x_i, x_j) + \sum_{i<j<k} E^{i,j,k}(x_i, x_j, x_k). \tag{2}$$

Obviously, the class $\mathcal{F}^2$ is a strict subset of the class $\mathcal{F}^3$.

The main result in this paper is a precise characterization of the functions in $\mathcal{F}^3$ that can be minimized using graph cuts, together with a graph construction for minimizing such functions. Moreover, we give a necessary condition for all other classes which must be met for a function to be minimized via graph cuts.

Our results also identify an interesting class of class of energy functions that have not yet been minimized using graph cuts. All of the previous work with graph cuts involves a neighborhood system that is defined on pairs of pixels. In the language of Markov Random Fields [10,19], these methods consider first-order MRF's. The associated energy functions lie in $\mathcal{F}^2$. Our results allow for the minimization of energy functions in the larger class $\mathcal{F}^3$, and thus for neighborhood systems involve triples of pixels.

1.2 Organization

The rest of the paper is organized as follows. In section 2 we give an overview of graph cuts. In section 3 we formalize the problem that we want to solve. Section 4 contains our main theorem for the class of functions $\mathcal{F}^2$ and shows how it can be used. Section 5 contains our main theorems for other classes. Proofs of our theorems, together with the graph constructions, are deferred to section 6, with some additional details deferred to a technical report [17]. A summary of the actual graph constructions os given in the appendix.

2 Overview of Graph Cuts

Suppose $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed graph with two special vertices (terminals), namely the source s and the sink t. An s-t-cut (or just a cut as we will refer to it later) $C = S, T$ is a partition of vertices in $\mathcal{V}$ into two disjoint sets S and T, such that $s \in S$ and $t \in T$. The cost of the cut is the cut is the sum of costs of all edges that go from S to T:

$$c(S,T) = \sum_{u \in S, v \in T, (u,v) \in \mathcal{E}} c(u,v).$$

The minimum s-t-cut problem is to find a cut C with the smallest cost. Due to the theorem of Ford and Fulkerson [9] this is equivalent to computing the maximum flow from the source to sink. There are many algorithms which solve this problem in polynomial time with small constants [1,11].

It is convenient to denote a cut $C = S, T$ by a labeling f mapping from the set of the nodes $\mathcal{V} - \{s, t\}$ to $\{0, 1\}$ where $f(v) = 0$ means that $v \in S$, and $f(v) = 1$ means that $v \in T$. We will use this notation later.

3 Defining Graph Representability

Let us consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with terminals s and t, thus $\mathcal{V} = \{v_1, \ldots, v_n, s, t\}$. Each cut on $\mathcal{G}$ has some cost; therefore, $\mathcal{G}$ represents the energy

function mapping from all cuts on $\mathcal{G}$ to the set of nonnegative real numbers. Any cut can be described by n binary variables $x_1, \ldots, x_n$ corresponding to nodes in $\mathcal{G}$ (excluding the source and the sink): $x_i = 0$ when $v_i \in S$, and $x_i = 1$ when $v_i \in T$. Therefore, the energy E that $\mathcal{G}$ represents can be viewed as a function of n binary variables: $E(x_1, \ldots, x_n)$ is equal to the cost of the cut defined by the configuration $x_1, \ldots, x_n$ ($x_i \in \{0, 1\}$). Note that the configuration that minimizes E will not change if we add a constant to E.

We can efficiently minimize E by computing the minimum s-t-cut on $\mathcal{G}$. This naturally leads to the question: what is the class of energy functions E for which we can construct a graph that represents E?

We can also generalize our construction. Above we used each node (except the source and the sink) for encoding one binary variable. Instead we can specify a subset $\mathcal{V}_0 = \{v_1, \ldots, v_k\} \subset \mathcal{V} - \{s, t\}$ and introduce variables only for the nodes in this set. Then there may be several cuts corresponding to a configuration $x_1, \ldots, x_k$. If we define the energy $E(x_1, \ldots, x_k)$ as the minimum among the costs of all such cuts, then the minimum s-t-cut on $\mathcal{G}$ will again yield the configuration which minimizes E.

We will summarize graph constructions that we allow in the following definition.

Definition 1. *A function E of n binary variables is called* graph-representable *if there exists a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with terminals s and t and a subset of nodes $V_0 = \{v_1, \ldots, v_n\} \subset \mathcal{V} - \{s, t\}$ such that for any configuration $x_1, \ldots, x_n$ the value of the energy $E(x_1, \ldots, x_n)$ is equal to a constant plus the cost of the minimum s-t-cut among all cuts $C = S, T$ in which $v_i \in S$, if $x_i = 0$, and $v_i \in T$, if $x_i = 1$ ($1 \leq i \leq n$). We say that E is* exactly represented *by $\mathcal{G}$, V_0 if this constant is zero.*

The following lemma is an obvious consequence of this definition.

Lemma 2. *Suppose the energy function E is graph-representable by a graph $\mathcal{G}$ and a subset $\mathcal{V}_0$. Then it is possible to find the exact minimum of E in polynomial time by computing the minimum s-t-cut on $\mathcal{G}$.*

In this paper we will give a complete characterization of the classes $\mathcal{F}^2$ and $\mathcal{F}^3$ in terms of graph representability, and show how to construct graphs for minimizing graph-representable energies within these classes. Moreover, we will give a necessary condition for all other classes which must be met for a function to be graph-representable. Note that it would be suffice to consider only the class $\mathcal{F}^3$ since $\mathcal{F}^2 \subset \mathcal{F}^3$. However, the condition for $\mathcal{F}^2$ is simpler so we will consider it separately.

4 The Class $\mathcal{F}^2$

Our main result for the class $\mathcal{F}^2$ is the following theorem.

Theorem 3. *Let E be a function of n binary variables from the class $\mathcal{F}^2$, i.e. it can be written as the sum*

$$E(x_1, \ldots, x_n) = \sum_i E^i(x_i) + \sum_{i<j} E^{i,j}(x_i, x_j).$$

Then E is graph-representable if and only if each term $E^{i,j}$ satisfies the inequality

$$E^{i,j}(0,0) + E^{i,j}(1,1) \leq E^{i,j}(0,1) + E^{i,j}(1,0).$$

4.1 Example: Pixel-Labeling via Expansion Moves

In this section we show how to apply this theorem to solve the pixel-labeling problem. In this problem, are given the set of pixels $\mathcal{P}$ and the set of labels $\mathcal{L}$. The goal is to find a labeling l (i.e. a mapping from the set of pixels to the set of labels) which minimizes the energy

$$E(l) = \sum_{p \in \mathcal{P}} D_p(l_p) + \sum_{p,q \in \mathcal{N}} V_{p,q}(l_p, l_q)$$

where $\mathcal{N} \subset \mathcal{P} \times \mathcal{P}$ is a neighborhood system on pixels. Without loss of generality we can assume that $\mathcal{N}$ contains only ordered pairs p, q for which $p < q$ (since we can combine two terms $V_{p,q}$ and $V_{q,p}$ into one term). We will show how our method can be used to derive the expansion move algorithm developed in [8].

This problem is shown in [8] to be NP-hard if $|\mathcal{L}| > 2$. [8] gives an approximation algorithm for minimizing this energy. A single step of this algorithm is an operation called an α-expansion. Suppose that we have some current configuration l^0, and we are considering a label $\alpha \in \mathcal{L}$. During the α-expansion operation a pixel p is allowed either to keep its old label l^0_p or to switch to a new label α: $l_p = l^0_p$ or $l_p = \alpha$. The key step in the approximation algorithm presented in [8] is to find the optimal expansion operation, i.e. the one that leads to the largest reduction in the energy E. This step is repeated until there is no choice of α where the optimal expansion operation reduces the energy.

[8] constructs a graph which contains nodes corresponding to pixels in $\mathcal{P}$. The following encoding is used: if $f(p) = 0$ (i.e., the node p is in the source set) then $l_p = l^0_p$; if $f(p) = 1$ (i.e., the node p is in the sink set) then $l_p = \alpha$.

Note that the key technical step in this algorithm can be naturally expressed as minimizing an energy function involving binary variables. The binary variables correspond to pixels, and the energy we wish to minimize can be written formally as

$$E(x_{p_1}, \ldots, x_{p_n}) = \sum_{p \in \mathcal{P}} D_p(l_p(x_p)) + \sum_{p,q \in \mathcal{N}} V_{p,q}(l_p(x_p), l_q(x_q)), \tag{3}$$

where

$$\forall p \in \mathcal{P} \qquad l_p(x_p) = \begin{cases} l^0_p, & x_p = 0 \\ \alpha, & x_p = 1. \end{cases}$$

We can demonstrate the power of our results by deriving an important restriction on this algorithm. In order for the graph cut construction of [8] to work, the function $V_{p,q}$ is required to be a metric. In their paper, it is not clear whether this is an accidental property of the construction (i.e., they leave open the possibility that a more clever graph cut construction may overcome this restriction).

Using our results, we can easily show this is not the case. Specifically, by theorem 1, the energy function given in equation 3 is graph-representable if and only if each term $V_{p,q}$ satisfies the inequality

$$V_{p,q}(l_p(0), l_q(0)) + V_{p,q}(l_p(1), l_q(1)) \leq V_{p,q}(l_p(0), l_q(1)) + V_{p,q}(l_p(1), l_q(0))$$

or

$$V_{p,q}(\beta, \gamma) + V_{p,q}(\alpha, \alpha) \leq V_{p,q}(\beta, \alpha) + V_{p,q}(\alpha, \gamma)$$

where $\beta = l_p^0$, $\gamma = l_q^0$. If $V_{p,q}(\alpha, \alpha) = 0$, then this is the triangle inequality:

$$V_{p,q}(\beta, \gamma) \leq V_{p,q}(\beta, \alpha) + V_{p,q}(\alpha, \gamma)$$

This is exactly the constraint on $V_{p,q}$ that was given in [8].

5 More General Classes of Energy Functions

We begin with several definitions. Suppose we have a function E of n binary variables. If we fix m of these variables then we get a new function E' of $n - m$ binary variables; we will call this function a *projection* of E. The notation for projections is as follows.

Definition 4. *Let $E(x_1, \ldots, x_n)$ be a function of n binary variables, and let I, J be a disjoint partition of the set of indices $\{1, \ldots, n\}$: $I = \{i(1), \ldots, i(m)\}$, $J = \{j(1), \ldots, j(n-m)\}$. Let $\alpha_{i(1)}, \ldots, \alpha_{i(m)}$ be some binary constants. A* projection *$E' = E[x_{i(1)} = \alpha_{i(1)}, \ldots, x_{i(m)} = \alpha_{i(m)}]$ is a function of $n - m$ variables defined by*

$$E'(x_{j(1)}, \ldots, x_{j(n-m)}) = E(x_1, \ldots, x_n),$$

where $x_i = \alpha_i$ for $i \in I$. We say that we fix *variables $x_{i(1)}$, ..., $x_{i(m)}$.*

Now we give a definition of *regular* functions.

Definition 5.

- *All functions of one variable are regular.*
- *A function E of two variables is called regular if $E(0,0) + E(1,1) \leq E(0,1) + E(1,0)$.*
- *A function E of more than two variables is called regular if all projections of E of two variables are regular.*

Now we are ready to formulate our main theorem for $\mathcal{F}^3$.

Theorem 6. *Let E be a function of n binary variables from $\mathcal{F}^3$, i.e. it can be written as the sum*

$$E(x_1, \ldots, x_n) = \sum_i E^i(x_i) + \sum_{i<j} E^{i,j}(x_i, x_j) + \sum_{i<j<k} E^{i,j,k}(x_i, x_j, x_k).$$

Then E is graph-representable if and only if E is regular.

Finally, we give a necessary condition for all other classes.

Theorem 7. *Let E be a function of binary variables. If E is not regular then E is not graph-representable.*

6 Proofs

6.1 Basic Lemmas

Definition 8. *The functional π will be a mapping from the set of all functions (of binary variables) to the set of real numbers which is defined as follows. For a function $E(x_1, \ldots, x_n)$*

$$\pi(E) = \sum_{x_1 \in \{0,1\}, \ldots, x_n \in \{0,1\}} (\Pi_{i=1}^n (-1)^{x_i})\, E(x_1, \ldots, x_n).$$

For example, for a function E of two variables $\pi(E) = E(0,0) - E(0,1) - E(1,0) + E(1,1)$. Note that a function E of two variables is regular if and only if $\pi(E) \leq 0$.

It is trivial to check the following properties of π.

Lemma 9.

- *π is linear, i.e. for a scalar c and two functions E', E'' of n variables $\pi(E' + E'') = \pi(E') + \pi(E'')$ and $\pi(c \cdot E') = c \cdot \pi(E')$.*
- *If E is a function of n variables that does not depend on at least one of the variables then $\pi(E) = 0$.*

The next two lemmas provide "building blocks" for constructing graphs for complex functions.

Lemma 10. *Let $I = \{1, \ldots, n\}$, $I' = \{i'(1), \ldots, i'(n')\} \subset I$, $I'' = \{i''(1), \ldots, i''(n'')\} \subset I$ be sets of indices. If the functions $E'(x_{i'(1)}, \ldots, x_{i'(n')})$ and $E''(x_{i''(1)}, \ldots, x_{i''(n'')})$ are graph-representable, then so is the function*

$$E(x_1, \ldots, x_n) = E'(x_{i'(1)}, \ldots, x_{i'(n')}) + E''(x_{i''(1)}, \ldots, x_{i''(n'')}).$$

Proof. Let us assume for simplicity of notation that E' and E'' are functions of all n variables: $E' = E'(x_1, \dots, x_n)$, $E'' = E''(\alpha_1, \dots, \alpha_n)$. By the definition of graph respresentability, there exist constants K', K'', graphs $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$, $\mathcal{G}'' = (\mathcal{V}'', \mathcal{E}'')$ and the set $\mathcal{V}_0 = \{v_1, \dots, v_n\}$, $\mathcal{V}_0 \subset \mathcal{V}' - \{s, t\}$, $\mathcal{V}_0 \subset \mathcal{V}'' - \{s, t\}$ such that $E' + K'$ is exactly represented by $\mathcal{G}'$, V_0 and $E'' + K''$ is exactly represented by $\mathcal{G}''$, V_0. We can assume that the only common nodes of $\mathcal{G}'$ and $\mathcal{G}''$ are $\mathcal{V}_0 \cup \{s, t\}$. Let us construct the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ as the combined graph of $\mathcal{G}'$ and $\mathcal{G}''$: $\mathcal{V} = \mathcal{V}' \cup \mathcal{V}''$, $\mathcal{E} = \mathcal{E}' \cup \mathcal{E}''$.

It is rather straightforward to show that $\mathcal{G}$, $\mathcal{V}_0$ exactly represent $E+(K'+K'')$ (and, therefore, E is graph-representable). Due to space limitations, the proof is omitted but is contained in a technical report [17].

□

Lemma 11. *Suppose E and E' are two functions of n variables such that*

$$\forall x_1, \dots, x_n \qquad E'(x_1, \dots, x_n) = \begin{cases} E(x_1, \dots, x_n), & x_k = 0 \\ E(x_1, \dots, x_n) + C, & x_k = 1, \end{cases}$$

for some constants k and C ($1 \le k \le n$). Then

- *E' if graph-representable if and only if E is graph-representable;*
- *E' is regular if and only if E is regular.*

Proof. Let us introduce the following function E^C:

$$\forall x_1, \dots, x_n \qquad E^C(x_1, \dots, x_n) = \begin{cases} 0, & x_k = 0 \\ C, & x_k = 1. \end{cases}$$

We need to show that E^C is graph-representable for any C then the first part of the lemma will follow from the lemma 3 since $E' \equiv E+E^C$ and $E \equiv E'+E^{-C}$.

It is easy to construct a graph which represents E^C. The set of nodes in this graph will be $\{v_1, \dots, v_n, s, t\}$ and the set of edges will include the only edge (s, v_k) with the capacity C (if $C \ge 0$) or the edge (v_k, t) with the capacity $-C$ (if $C < 0$). It is trivial to check that this graph exactly represents E^C (in the former case) or $E^C + C$ (in the latter case).

Now let us assume that one of the functions E and E' is regular, for example, E. Consider a projection of E' of two variables:

$$E'[x_{i(1)} = \alpha_{i(1)}, \dots, x_{i(m)} = \alpha_{i(m)}],$$

where $m = n - 2$ and $\{i(1), \dots, i(m)\} \subset \{1, \dots, n\}$. We need to show that this function is regular, i.e. that the functional π of this function is nonpositive. Due to the linearity of π we can write

$$\begin{aligned} &\pi(E'[x_{i(1)} = \alpha_{i(1)}, \dots, x_{i(m)} = \alpha_{i(m)}]) = \\ &= \pi(E[x_{i(1)} = \alpha_{i(1)}, \dots, x_{i(m)} = \alpha_{i(m)}]) \quad + \\ &+ \quad \pi(E^C[x_{i(1)} = \alpha_{i(1)}, \dots, x_{i(m)} = \alpha_{i(m)}]). \end{aligned}$$

The first term is nonpositive by assumption, and the second term is 0 by lemma 2.

□

6.2 Proof of Theorems 1 and 2: The Constructive Part

In this section we will give the constructive part of the proof: given a regular energy function from class $\mathcal{F}^3$ we will show how to construct a graph which represents it. We will do it in three steps. First we will consider regular functions of two variables, then regular functions of three variables and finally regular functions of the form as in the theorem 2.

This will also prove the constructive part of the theorem 1. Indeed, suppose a function is from the class $\mathcal{F}^2$ and each term in the sum satisfies the condition given in the theorem 1 (i.e. regular). Then each term is graph-representable (as we will show in this section) and, hence, the function is graph-representable as well according to the lemma 3.

The other direction of theorems 1 and 2 as well as the theorem 3 will be proven in the section 6.3.

Functions of two variables. Let $E(x_1, x_2)$ be a function of two variables represented by a table

$$E = \begin{array}{|c|c|}\hline E(0,0) & E(0,1) \\ \hline E(1,0) & E(1,1) \\ \hline \end{array}$$

Lemma 4 tells us that we can add a constant to any column or row without affecting theorem 2. Thus, without loss of generality we can consider only functions E of the form

$$E = \begin{array}{|c|c|}\hline 0 & A \\ \hline 0 & 0 \\ \hline \end{array}$$

(we subtracted a constant from the first row to make the upper left element zero, then we subtracted a constant from the second row to make the bottom left element zero, and finally we subtracted a constant from the second column to make the bottom right element zero).

$\pi(E) = -A \leq 0$ since we assumed that E is regular; hence, A is non-negative. Now we can easily construct a graph $\mathcal{G}$ which represents this function. It will have four vertices $\mathcal{V} = \{v_1, v_2, s, t\}$ and one edge $\mathcal{E} = \{(v_1, v_2)\}$ with the cost $c(v_1, v_2) = A$. It is easy to see that $\mathcal{G}$, $\mathcal{V}_0 = \{v_1, v_2\}$ represent E since the only case when the edge (v_1, v_2) is cut (yielding a cost A) is when $v_1 \in S$, $v_2 \in T$, i.e. when $x_1 = 0$, $x_2 = 1$.

Note that we did not introduce any additional nodes for representing binary interactions of binary variables. This is in contrast to the construction in [8] which added auxiliary nodes for representing energies that we just considered. Our construction yields a smaller graph and, thus, the minimum cut can potentially be computed faster.

Functions of three variables. Now let us consider a regular function E of three variables. Let us represent it as a table

$$E = \begin{array}{|c|c|}\hline E(0,0,0) & E(0,0,1) \\ \hline E(0,1,0) & E(0,1,1) \\ \hline E(1,0,0) & E(1,0,1) \\ \hline E(1,1,0) & E(1,1,1) \\ \hline \end{array}$$

Two cases are possible:

Case 1. $\pi(E) \geq 0$. We can apply transformations described in lemma 4 and get the following function:

$$E = \begin{array}{|c|c|}\hline 0 & 0 \\ \hline 0 & A_1 \\ \hline 0 & A_2 \\ \hline A_3 & A_0 \\ \hline \end{array}$$

It is easy to check that these transformations preserve the functional π. Hence, $A = A_0 - (A_1 + A_2 + A_3) = -\pi(E) \leq 0$. By applying the regularity constraint to the projections $E[x_1 = 0]$, $E[x_2 = 0]$, $E[x_3 = 0]$ we also get $A_1 \leq 0$, $A_2 \leq 0$, $A_3 \leq 0$.

We can represent E as the sum of functions

$$E = \begin{array}{|c|c|}\hline 0 & 0 \\ \hline 0 & A_1 \\ \hline 0 & 0 \\ \hline 0 & A_1 \\ \hline \end{array} + \begin{array}{|c|c|}\hline 0 & 0 \\ \hline 0 & 0 \\ \hline 0 & A_2 \\ \hline 0 & A_2 \\ \hline \end{array} + \begin{array}{|c|c|}\hline 0 & 0 \\ \hline 0 & 0 \\ \hline 0 & 0 \\ \hline A_3 & A_3 \\ \hline \end{array} + \begin{array}{|c|c|}\hline 0 & 0 \\ \hline 0 & 0 \\ \hline 0 & 0 \\ \hline 0 & A \\ \hline \end{array}$$

We need to show that all terms here are graph-representable, then lemma 3 will imply that E is graph-representable as well.

The first three terms are regular functions depending only on two variables and thus are graph-representable as was shown in the previous section. Let us consider the last term.

The graph $\mathcal{G}$ that represents this term can be constructed as follows. The set of nodes will contain one auxilary node u: $\mathcal{V} = \{v_1, v_2, v_3, u, s, t\}$. The set of edges will consist of directed edges $\mathcal{E} = \{(v_1, u), (v_2, u), (v_3, u), (u, t)\}$ with capacities $A' = -A$. It is easy to check that $\mathcal{G}$, $\mathcal{V}_0 = \{v_1, v_2, v_3\}$ exactly represent the function $E'(x_1, x_2, x_3) = E(x_1, x_2, x_3) + A'$; the proof is contained in [17].

Case 2. $\pi(E) < 0$. This case is similar to the case 1. We can transform the energy to

$$E = \begin{array}{|c|c|}\hline A_0 & A_3 \\ \hline A_2 & 0 \\ \hline A_1 & 0 \\ \hline 0 & 0 \\ \hline \end{array} = \begin{array}{|c|c|}\hline A_1 & 0 \\ \hline 0 & 0 \\ \hline A_1 & 0 \\ \hline 0 & 0 \\ \hline \end{array} + \begin{array}{|c|c|}\hline A_2 & 0 \\ \hline A_2 & 0 \\ \hline 0 & 0 \\ \hline 0 & 0 \\ \hline \end{array} + \begin{array}{|c|c|}\hline A_3 & A_3 \\ \hline 0 & 0 \\ \hline 0 & 0 \\ \hline 0 & 0 \\ \hline \end{array} + \begin{array}{|c|c|}\hline A & 0 \\ \hline 0 & 0 \\ \hline 0 & 0 \\ \hline 0 & 0 \\ \hline \end{array}$$

where $A = A_0 - (A_1 + A_2 + A_3) = \pi(E) < 0$ and $A_1 \leq 0$, $A_2 \leq 0$, $A_3 \leq 0$ since E is regular. The first three terms are regular functions of two variables and the last term can be represented by the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V} = \{v_1, v_2, v_3, u, s, t\}$ and $\mathcal{E} = \{(u, v_1), (u, v_2), (u, v_3), (s, u)\}$; capacities of all edges are $-A$.

Functions of many variables. Finally let us consider a regular function E which can be written as

$$E(x_1, \ldots, x_n) = \sum_{i<j<k} E^{i,j,k}(x_i, x_j, x_k),$$

where i, j, k are indices from the set $\{1, \ldots, n\}$ (we omitted terms involving functions of one and two variables since they can be viewed as functions of three variables).

Although E is regular, each term in the sum need not necessarily be regular. However we can "regroup" terms in the sum so that each term will be regular (and, thus, graph-representable). This can be done using the following lemma and a trivial induction argument.

Definition 12. *Let $E^{i,j,k}$ be a function of three variables. The functional $N(E^{i,j,k})$ is defined as the number of projections of two variables of $E^{i,j,k}$ with the positive value of the functional π.*

Note that $N(E^{i,j,k}) = 0$ exactly when $E^{i,j,k}$ is regular.

Lemma 13. *Suppose the function E of n variables can be written as*

$$E(x_1, \ldots, x_n) = \sum_{i<j<k} E^{i,j,k}(x_i, x_j, x_k),$$

where some of the terms are not regular. Then it can be written as

$$E(x_1, \ldots, x_n) = \sum_{i<j<k} \tilde{E}^{i,j,k}(x_i, x_j, x_k),$$

where

$$\sum_{i<j<k} N(\tilde{E}^{i,j,k}) < \sum_{i<j<k} N(E^{i,j,k}).$$

Proof. For the simplicity of notation let us assume that the term $E^{1,2,3}$ is not regular and $\pi(E^{1,2,3}[x_3 = 0]) > 0$ or $\pi(E^{1,2,3}[x_3 = 1]) > 0$ (we can ensure this by renaming indices). Let

$$C_k = \max_{\alpha_k \in \{0,1\}} \pi(E^{1,2,k}[x_k = \alpha_k]) \qquad k \in \{4, \ldots, n\}$$

$$C_3 = -\sum_{k=4}^{n} C_k$$

Now we will modify the terms $E^{1,2,3}$, …, $E^{1,2,n}$ as follows:

$$\tilde{E}^{1,2,k} \equiv E^{1,2,k} - R[C_k] \qquad k \in \{3, \ldots, n\}$$

where $R[C]$ is the function of two variables x_1 and x_2 defined by the table

$$R[C] = \begin{array}{|c|c|}\hline 0 & 0 \\ \hline 0 & C \\ \hline \end{array}$$

(other terms are unchanged: $\tilde{E}^{i,j,k} \equiv E^{i,j,k}$, $(i,j) \neq (1,2)$). We have

$$E(x_1, \ldots, x_n) = \sum_{i<j<k} \tilde{E}^{i,j,k}(x_i, x_j, x_k)$$

since $\sum_{k=3}^{n} C_k = 0$ and $\sum_{k=3}^{n} R[C_k] \equiv 0$.

If we consider $R[C]$ as a function of n variables and take a projection of two variables where the two variables that are not fixed are x_i and x_j $(i < j)$, then the functional π will be C, if $(i,j) = (1,2)$, and 0 otherwise since in the latter case a projection actually depends on at most one variable. Hence, the only projections of two variables that could have changed their value of the functional π are $\tilde{E}^{1,2,k}[x_3 = \alpha_3, \ldots, x_n = \alpha_n]$, $k \in \{3, \ldots, n\}$, if we treat $\tilde{E}^{1,2,k}$ as functions of n variables, or $\tilde{E}^{1,2,k}[x_k = \alpha_k]$, if we treat $\tilde{E}^{1,2,k}$ as functions of three variables.

First let us consider terms with $k \in \{4, \ldots, n\}$. We have $\pi(E^{1,2,k}[x_k = \alpha_k]) \leq C_k$, thus

$$\pi(\tilde{E}^{1,2,k}[x_k = \alpha_k]) = \pi(E^{1,2,k}[x_k = \alpha_k]) - \pi(R[C_k][x_k = \alpha_k]) \leq C_k - C_k = 0$$

Therefore we did not introduce any nonregular projections for these terms.

Now let us consider the term $\pi(\tilde{E}^{1,2,3}[x_3 = \alpha_3])$. We can write

$$\pi(\tilde{E}^{1,2,3}[x_3 = \alpha_3]) = \pi(E^{1,2,3}[x_3 = \alpha_3]) - \pi(R[C_3][x_3 = \alpha_3]) =$$
$$= \pi(E^{1,2,3}[x_3 = \alpha_3]) - (-\sum_{k=4}^{n} C_k) = \sum_{k=3}^{n} \pi(E^{1,2,k}[x_k = \alpha_k])$$

where $\alpha_k = \arg\max_{\alpha \in \{0,1\}} \pi(E^{1,2,k}[x_k = \alpha])$, $k \in \{4, \ldots, n\}$. The last expression is just $\pi(E[x_3 = \alpha_3, \ldots, x_n = \alpha_n])$ and is nonpositive since E is regular by assumption. Hence, values $\pi(\tilde{E}^{1,2,3}[x_3 = 0])$ and $\pi(\tilde{E}^{1,2,3}[x_3 = 1])$ are both nonpositive and, therefore, the number of nonregular projections has decreased. □

6.3 Proof of Theorem 3

In this section we will prove a necessary condition for graph representability: if a function of binary variables is graph-representable then it is regular. It will also imply the corresponding directions of the theorems 1 and 2. Note that theorem 1 needs a little bit of reasoning, as follows. Let us consider a graph-representable function E from the class $\mathcal{F}^2$:

$$E(x_1, \ldots, x_n) = \sum_i E^i(x_i) + \sum_{i<j} E^{i,j}(x_i, x_j)$$

E is regular as we will prove in this section. It means that the functional π of any projection of E of two variables is nonpositive. Let us consider a projection where the two variables that are not fixed are x_i and x_j. By lemma 2 the value of the functional π of this projection is equal to $\pi(E^{i,j})$ (all other terms yield zero). Hence, all terms $E^{i,j}$ are regular, i.e. they satisfy the condition in the theorem 1.

Definition 14. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph, $v_1, \ldots, v_k$ be a subset of nodes $\mathcal{V}$ and $\alpha_1, \ldots, \alpha_k$ be binary constants whose values are from $\{0, 1\}$. We will define the graph $\mathcal{G}[x_1 = \alpha_1, \ldots, x_k = \alpha_k]$ as follows. Its nodes will be the same as in $\mathcal{G}$ and its edges will be all edges of $\mathcal{G}$ plus additional edges corresponding to nodes $v_1, \ldots, v_k$: for a node v_i, we add the edge (s, v_i), if $\alpha_i = 0$, or (v_i, t), if $\alpha_i = 1$, with an infinite capacity.*

It should be obvious that these edges enforce constraints $f(v_1) = \alpha_1$, ..., $f(v_k) = \alpha_k$ in the minimum cut on $\mathcal{G}[x_1 = \alpha_1, \ldots, x_k = \alpha_k]$, i.e. if $\alpha_i = 0$ then $v_i \in S$, and if $\alpha_i = 1$ then $v_i \in T$. (If, for example, $\alpha_i = 0$ and $v_i \in T$ then the edge (s, v_i) must be cut yielding an infinite cost, so it would not the minimum cut.)

Now we can give a definition of graph representability which is equivalent to the definition 1. This new definition will be more convenient for the proof.

Definition 15. *We say that the function E of n binary variables is exactly represented by the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and the set $\mathcal{V}_0 \subset \mathcal{V}$ if for any configuration $\alpha_1, \ldots, \alpha_n$ the cost of the minimum cut on $\mathcal{G}[x_1 = \alpha_1, \ldots, x_k = \alpha_k]$ is $E(\alpha_1, \ldots, \alpha_n)$.*

Lemma 16. *Any projection of a graph-representable function is graph-representable.*

Proof. Let E be a graph-representable function of n variables, and the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and the set $\mathcal{V}_0$ represents E. Suppose that we fix variables $x_{i(1)}, \ldots, x_{i(m)}$. It is straightforward to check that the graph $\mathcal{G}[x_{i(1)} = \alpha_{i(1)}, \ldots, x_{i(m)} = \alpha_{i(m)}]$ and the set $\mathcal{V}_0' = \mathcal{V}_0 - \{v_{i(1)}, \ldots, v_{i(m)}\}$ represent the function $E' = E[x_{i(1)} = \alpha_{i(1)}, \ldots, x_{i(m)} = \alpha_{i(m)}]$. □

This lemma implies that it suffices to prove theorem 3 only for energies of two variables.

Let $E(x_1, x_2)$ be a graph-representable function of two variables. Let us represent this function as a table:

$$E = \begin{array}{|c|c|}\hline E(0,0) & E(0,1) \\ \hline E(1,0) & E(1,1) \\ \hline \end{array}$$

Lemma 4 tells us that we can add a constant to any column or row without affecting theorem 3. Thus, without loss of generality we can consider only functions E of the form

$$E = \begin{array}{|c|c|}\hline 0 & 0 \\ \hline 0 & A \\ \hline \end{array}$$

(we subtracted a constant from the first row to make the upper left element zero, then we subtracted a constant from the second row to make the bottom left element zero, and finally we subtracted a constant from the second column to make the upper right element zero).

We need to show that E is regular, i.e. that $\pi(E) = A \leq 0$. Suppose this is not true: $A > 0$.

Suppose the graph $\mathcal{G}$ and the set $\mathcal{V}_0 = \{v_1, v_2\}$ represent E. It means that there is a constant K such that $\mathcal{G}$, $\mathcal{V}_0$ exactly represent $E'(x_1, x_2) = E(x_1, x_2) + K$:

$$E' = \begin{array}{|c|c|}\hline K & K \\ \hline K & K + A \\ \hline \end{array}$$

The cost of the minimum s-t-cut on $\mathcal{G}$ is K (since this cost is just the minimum entry in the table for E'); hence, $K \geq 0$. Thus the value of the maximum flow from s to t in $\mathcal{G}$ is K. Let $\mathcal{G}^0$ be the residual graph obtained from $\mathcal{G}$ after pushing the flow K. Let $E^0(x_1, x_2)$ be the function exactly represented by $\mathcal{G}^0, \mathcal{V}_0$.

By the definition of graph representability, $E'(\alpha_1, \alpha_2)$ is equal to the value of the minimum cut (or maximum flow) on the graph $\mathcal{G}[x_1 = \alpha_1, x_2 = \alpha_2]$. The following sequence of operations shows one possible way to push the maximum flow through this graph.

- First we take the original graph $\mathcal{G}$ and push the flow K; then we get the residual graph $\mathcal{G}^0$. (It is equivalent to pushing flow through $\mathcal{G}[x_1 = \alpha_1, x_2 = \alpha_2]$ where we do not use edges corresponding to constraints $x_1 = \alpha_1$ and $x_2 = \alpha_2$).
- Then we add edges corresponding to these constraints; then we get the graph $\mathcal{G}^0[x_1 = \alpha_1, x_2 = \alpha_2]$.
- Finally we push the maximum flow possible through the graph $\mathcal{G}^0[x_1 = \alpha_1, x_2 = \alpha_2]$; the amount of this flow is $E^0(\alpha_1, \alpha_2)$ according to the definition of graph representability.

The total amount of flow pushed during all steps is $K + E^0(\alpha_1, \alpha_2)$; thus,

$$E'(\alpha_1, \alpha_2) = K + E^0(\alpha_1, \alpha_2)$$

or

$$E(\alpha_1, \alpha_2) = E^0(\alpha_1, \alpha_2)$$

We proved that E is exactly represented by $\mathcal{G}^0$, $\mathcal{V}_0$.

The value of the minimum cut/maximum flow on $\mathcal{G}^0$ is 0 (it is the minimum entry in the table for E); thus, there is no augmenting path from s to t in $\mathcal{G}^0$. However, if we add edges (v_1, t) and (v_2, t) then there will be an augmenting path from s to t in $\mathcal{G}^0[x_1 = \alpha_1, x_2 = \alpha_2]$ since $E(1, 1) = A > 0$. Hence, this augmenting path will contain at least one of these edges and, therefore, either v_1 or v_2 will be in the path. Let P be the part of this path going from the source until v_1 or v_2 is first encountered. Without loss of generality we can assume that it will be v_1. Thus, P is an augmenting path from s to v_1 which does not contain edges that we added, namely (v_1, t) and (v_2, t).

Finally let us consider the graph $\mathcal{G}^0[x_1 = 1, x_2 = 0]$ which is obtained from $\mathcal{G}^0$ by adding edges (v_1, t) and (s, v_2) with infinite capacities. There is an augmenting path $\{P, (v_1, t)\}$ from the source to the sink in this graph; hence, the minimum cut/maximum flow on it greater than zero, or $E(1, 0) > 0$. We get a contradiction.

Acknowledgements. We thank Olga Veksler and Yuri Boykov for their careful reading of this paper, and for valuable comments which greatly improved its readibility. We also thank Ian Jermyn for helping us clarify the paper's motivation. This research was supported by NSF grants IIS-9900115 and CCR-0113371, and by a grant from Microsoft Research.

Appendix: Summary of Graph Constructions

We now summarize the graph constructions used for regular functions. The notation $D(v, c)$ means that we add an edge (s, v) with the weight c if $c > 0$, or an edge (v, t) with the weight $-c$ if $c < 0$.

Regular Functions of One Binary Variable

Recall that all functions of one variable are regular. For a function $E(x_1)$, we construct a graph $\mathcal{G}$ with three vertices $\mathcal{V} = \{v_1, s, t\}$. There is a single edge $D(v_1, E(1) - E(0))$.

Regular Functions of Two Binary Variables

We now show how to construct a graph $\mathcal{G}$ for a regular function $E(x_1, x_2)$ of two variables. It will contain four vertices: $\mathcal{V} = \{v_1, v_2, s, t\}$. The edges $\mathcal{E}$ are given below.

- $D(v_1, E(1,0) - E(0,0))$;
- $D(v_2, E(1,1) - E(1,0))$;
- (v_1, v_2) with the weight $-\pi(E)$.

Regular Functions of Three Binary Variables

We next show how to construct a graph $\mathcal{G}$ for a regular function $E(x_1, x_2, x_3)$ of three variables. It will contain five vertices: $\mathcal{V} = \{v_1, v_2, v_3, u, s, t\}$. If $\pi(E) \geq 0$ then the edges are

- $D(v_1, E(1,0,1) - E(0,0,1))$;
- $D(v_2, E(1,1,0) - E(1,0,0))$;
- $D(v_3, E(0,1,1) - E(0,1,0))$;
- (v_2, v_3) with the weight $-\pi(E[x_1 = 0])$;
- (v_3, v_1) with the weight $-\pi(E[x_2 = 0])$;
- (v_1, v_2) with the weight $-\pi(E[x_3 = 0])$;
- (v_1, u), (v_2, u), (v_3, u), (u, t) with the weight $\pi(E)$.

If $\pi(E) < 0$ then the edges are

- $D(v_1, E(1,1,0) - E(0,1,0))$;
- $D(v_2, E(0,1,1) - E(0,0,1))$;
- $D(v_3, E(1,0,1) - E(1,0,0))$;
- (v_3, v_2) with the weight $-\pi(E[x_1 = 1])$;
- (v_1, v_3) with the weight $-\pi(E[x_2 = 1])$;
- (v_2, v_1) with the weight $-\pi(E[x_3 = 1])$;
- (u, v_1), (u, v_2), (u, v_3), (s, u) with the weight $-\pi(E)$.

References

1. Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
2. Amir Amini, Terry Weymouth, and Ramesh Jain. Using dynamic programming for solving variational problems in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(9):855–867, September 1990.
3. Stephen Barnard. Stochastic stereo matching over scale. *International Journal of Computer Vision*, 3(1):17–32, 1989.
4. S. Birchfield and C. Tomasi. Multiway cut for stereo and motion with slanted surfaces. In *International Conference on Computer Vision*, pages 489–495, 1999.
5. Yuri Boykov and Marie-Pierre Jolly. Interactive organ segmentation using graph cuts. In *Medical Image Computing and Computer-Assisted Intervention*, pages 276–286, 2000.
6. Yuri Boykov and Marie-Pierre Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images. In *International Conference on Computer Vision*, pages I: 105–112, 2001.
7. Yuri Boykov, Olga Veksler, and Ramin Zabih. Markov Random Fields with efficient approximations. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 648–655, 1998.
8. Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, November 2001.
9. L. Ford and D. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
10. S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
11. A. Goldberg and R. Tarjan. A new approach to the maximum flow problem. *Journal of the Association for Computing Machinery*, 35(4):921–940, October 1988.
12. D. Greig, B. Porteous, and A. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society, Series B*, 51(2):271–279, 1989.
13. H. Ishikawa and D. Geiger. Occlusions, discontinuities, and epipolar lines in stereo. In *European Conference on Computer Vision*, pages 232–248, 1998.
14. H. Ishikawa and D. Geiger. Segmentation by grouping junctions. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 125–131, 1998.
15. Junmo Kim, John Fish, Andy Tsai, Cindy Wible, Ala Willsky, and William Wells. Incorporating spatial priors into an information theoretic approach for fMRI data analysis. In *Medical Image Computing and Computer-Assisted Intervention*, pages 62–71, 2000.
16. Vladimir Kolmogorov and Ramin Zabih. Visual correspondence with occlusions using graph cuts. In *International Conference on Computer Vision*, pages 508–515, 2001.
17. Vladimir Kolmogorov and Ramin Zabih. What energy functions can be minimized via graph cuts? Technical report CUCS-TR2001-1857, Cornell Computer Science Department, November 2001.
18. Vladimir Kolmogorov and Ramin Zabih. Multi-camera scene reconstruction via graph cuts. In *European Conference on Computer Vision*, 2002.
19. S. Li. *Markov Random Field Modeling in Computer Vision*. Springer-Verlag, 1995.

20. S. Roy. Stereo without epipolar lines: A maximum flow formulation. *International Journal of Computer Vision*, 1(2):1–15, 1999.
21. S. Roy and I. Cox. A maximum-flow formulation of the n-camera stereo correspondence problem. In *International Conference on Computer Vision*, 1998.
22. Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. Technical Report 81, Microsoft Research, 2001. To appear in *IJCV*. An earlier version appears in CVPR 2001 Workshop on Stereo Vision.
23. Dan Snow, Paul Viola, and Ramin Zabih. Exact voxel occupancy with graph cuts. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 345–352, 2000.
24. Richard Szeliski and Ramin Zabih. An experimental comparison of stereo algorithms. In B. Triggs, A. Zisserman, and R. Szeliski, editors, *Vision Algorithms: Theory and Practice*, number 1883 in LNCS, pages 1–19, Corfu, Greece, September 1999. Springer-Verlag.