To appear in Discrete Applied Mathematics.

The Forbidden Projections of Unate Functions

Aaron Feigelson Lisa Hellerstein *

Abstract

We characterize the forbidden projections of unate Boolean functions. Forbidden projections are analogous to forbidden graph minors. Unate functions have been studied in switching theory and in computational learning theory.

1 Introduction

A major topic in graph and matroid theory has been the characterization of classes of graphs and matroids by forbidden minors (see, e.g., [9]). Analogously, it is possible to characterize some classes of Boolean functions by forbidden *projections*. A minor of a graph is formed by contraction and deletion of a subset of the edges in the graph. A projection of a Boolean function is formed by taking a subset of the input variables of the function and fixing each of them to either 0 or 1.

Surprisingly, despite the interest in graph and matroid minors, there has been almost no work on forbidden projections. One exception is the work of Seymour [8], who has characterized certain classes of clutters by forbidden clutter minors. A clutter is a collection of sets L such that for $A_1, A_2 \in L$, $A_1 \not\subset A_2$. Therefore, the sets of a clutter correspond directly to the minterms of a monotone Boolean function. Seymour's work can be viewed as giving a forbidden projection characterization of certain subclasses of the monotone Boolean functions.

In this paper, we characterize the class of *unate* Boolean functions by forbidden projections. Unateness is a generalization of monotonicity. Let $x \in V$. The function f is *monotone in* x if for every assignment

*Dept. of Electrical Engineering and Computer Science Northwestern University 2145 Sheridan Rd. Evanston, IL 60208-3118. Lisa Hellerstein was partially supported by NSF Grant CCR-9210957.

$\mathbf{a} \in \{0,1\}^V$ that sets x to 0, $f(\mathbf{a}) = 1 \Rightarrow f(\mathbf{a}_{x\leftarrow 1}) = 1$. (Here $\mathbf{a}_{x\leftarrow 1}$ denotes the assignment that is identical to $\mathbf{a}$ except that x is set to 1.) That is, changing x from 0 to 1 in an assignment cannot decrease the value of f. The function f is *anti-monotone in* x if for every assignment $\mathbf{a} \in \{0,1\}^V$ that sets x to 1, $f(\mathbf{a}) = 1 \Rightarrow f(\mathbf{a}_{x\leftarrow 0}) = 1$. The function f is *unate in* x if it is either monotone or anti-monotone in x.

A Boolean function is *monotone* if it is monotone in all its input variables. It is *unate* if it is unate in all its input variables. Unate functions have been studied extensively in switching theory [6, 7]. More recently, they have been exploited in the development of algorithms in computational learning theory [1, 3].

The class of monotone Boolean functions has a simple characterization in terms of forbidden projections. The class consists of exactly those functions that do not have any projections equivalent to $g(x) = \overline{x}$. In contrast, the characterization of unate functions by forbidden projections is significantly more complex.

1.1 Minimally Non-Unate Functions

To characterize the unate Boolean functions by forbidden projections, it suffices to characterize the class of *minimally non-unate* functions. We call this class MNU. A function f is in MNU iff f is non-unate, but every non-trivial projection of f is unate. Every non-unate function has a projection that is in MNU. Therefore, a function f is unate iff no projection of f is in MNU.

One of the simplest functions in MNU is the Boolean *consensus* function. A Boolean function f is a consensus function if $f(\mathbf{a}) = 1$ when $\mathbf{a}$ is the all-0's or all-1's assignment, and $f(\mathbf{a}) = 0$ otherwise. To verify that this function is in MNU, note that any non-trivial projection of f will be inconsistent with at least one of the two satisfying assignments. The resulting projection must be unate since it has at most one satisfying assignment.

A function f is a *generalized-consensus* function if there is some assignment $\mathbf{s}$ such that $f(\mathbf{s}) = f(\overline{\mathbf{s}}) = 1$ (where $\overline{\mathbf{s}}$ is the bit-wise complement of $\mathbf{s}$) and $f(\mathbf{a}) = 0$ on all other assignments $\mathbf{a}$. Generalized consensus functions are also in MNU.

The function $\phi(x,y,z) = \overline{x}\,\overline{z} \vee xy$ is in MNU but is not a generalized-consensus function (because exactly four of the eight possible assignments are satisfying assignments). To verify that ϕ is in MNU, note that $\phi(1,0,0) = 0$ and $\phi(0,0,0) = 1$ so ϕ is not monotone in x, and $\phi(0,1,1) = 0$ and $\phi(1,1,1) = 1$ so ϕ is not anti-monotone in x. However, fixing the value of any variable of ϕ results in a unate projection.

The characterization of MNU is fairly involved. It hinges on the fact that if f is a function in MNU, then there are precisely two complementary assignments that demonstrate the "non-unateness" of all the non-unate variables of f. That is, changing the setting of any non-unate variable x in one of these assignments shows that f is not monotone in x. Changing the setting of x in the other assignment shows that f is not anti-monotone in x. If f is a consensus function, then f has only non-unate variables, and the two complementary assignments are the all-0's and the all-1's assignments. Other functions in MNU contain both unate and non-unate variables. In the function ϕ given above, y and z are unate variables, and x is a non-unate variable. The generalized consensus functions (which include the consensus function) and their negations are an extreme case of the functions in MNU, those with no unate variables.

1.2 A description of the characterization

Our characterization is based on partial orders. Let $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \{0,1\}^V$ be assignments to the variable set V. Define the partial order $(V, \prec_{\mathbf{c}})$ such that $\mathbf{a} \prec_{\mathbf{c}} \mathbf{b}$ iff $\mathbf{a} \neq \mathbf{b}$ and for all $x \in V$, $\mathbf{c}(x) = 0 \Rightarrow \mathbf{a}(x) \leq \mathbf{b}(x)$ and $\mathbf{c}(x) = 1 \Rightarrow \mathbf{b}(x) \leq \mathbf{a}(x)$. If $\mathbf{c}$ is the all 0's assignment, then $\prec_{\mathbf{c}}$ is the standard partial order on the Boolean lattice, and $f : \{0,1\}^V \rightarrow \{0,1\}$ is monotone iff for all $\mathbf{a}, \mathbf{b} \in \{0,1\}^V$, $\mathbf{a} \prec_{\mathbf{c}} \mathbf{b} \Rightarrow f(\mathbf{a}) \leq f(\mathbf{b})$.

A Boolean function $f : \{0,1\}^V \rightarrow \{0,1\}$ is unate iff for some $\mathbf{c} \in \{0,1\}^V$, $\mathbf{a} \prec_{\mathbf{c}} \mathbf{b} \Rightarrow f(\mathbf{a}) \leq f(\mathbf{b})$. In this case we say that $(V, \prec_{\mathbf{c}})$ is a *monotone orientation* of f. Moreover, if for $x \in V$, $\mathbf{c}(x) = 0$, then f is monotone in x, and if $\mathbf{c}(x) = 1$, then f is anti-monotone in x.

For any assignment $\mathbf{b} \in \{0,1\}^V$, and any subset $V' \subseteq V$, let $\mathbf{b}|_{V'}$ denote the assignment $\mathbf{b}$ restricted to the domain V'. That is, $\mathbf{b}|_{V'} : V' \rightarrow \{0,1\}$ such that for all $x \in V'$, $\mathbf{b}|_{V'}(x) = \mathbf{b}(x)$.

We now describe how to construct an arbitrary function $f : \{0,1\}^V \rightarrow \{0,1\}$ in MNU. Start with a set of variables V such that $|V| \geq 2$. Pick two complementary assignments, $\mathbf{s}$ and $\mathbf{t} = \overline{\mathbf{s}}$, to V. Choose a subset $V_U \subsetneq V$ such that $|V_U| \neq 1$, and an assignment $\mathbf{c} \in \{0,1\}^{V_U}$ such that if $|V_U| > 0$ then $\mathbf{c} \neq \mathbf{s}|_{V_U}$ and $\mathbf{c} \neq \mathbf{t}|_{V_U}$. Then choose a unate function g defined on V_U such that g obeys the following properties:

1. $(V_U, \prec_{\mathbf{c}})$ is a monotone orientation of g
2. $g(\mathbf{s}|_{V_U}) = g(\mathbf{t}|_{V_U})$
3. for all $\mathbf{a} \in \{0,1\}^{V_U}$, if $\mathbf{a} \prec_{\mathbf{c}} \mathbf{s}|_{V_U}$ or $\mathbf{a} \prec_{\mathbf{c}} \mathbf{t}|_{V_U}$, then $g(\mathbf{a}) = 0$
   for all $\mathbf{a} \in \{0,1\}^{V_U}$, if $\mathbf{s}|_{V_U} \prec_{\mathbf{c}} \mathbf{a}$ or $\mathbf{t}|_{V_U} \prec_{\mathbf{c}} \mathbf{a}$, then $g(\mathbf{a}) = 1$

Finally, construct $f : \{0,1\}^V \rightarrow \{0,1\}$ as follows. For all $\mathbf{b} \in \{0,1\}^V$,

$$f(\mathbf{b}) = \begin{cases} \overline{g(\mathbf{b}|_{V_U})} & \text{if } \mathbf{b} \in \{\mathbf{s}, \mathbf{t}\} \\ g(\mathbf{b}|_{V_U}) & \text{otherwise} \end{cases}$$

Note that any projection assigning values to all the non-unate variables of f (and leaving the unate variables unassigned) is either equivalent to g, or differs from g at precisely one point, $\mathbf{s}|_{V_U}$ or $\mathbf{t}|_{V_U}$. Any projection that assigns values to all the unate variables of f is either a constant function, AND, OR, NAND, or NOR.

In Section 2.2 we show that a function is in MNU iff it can be constructed by the above method. We also show in Section 2.2 that if f is produced by the construction, then V_U is the set of all the unate variables of f. Thus if f is in MNU, then the number of unate variables of f is not equal to 1. The following functions in MNU illustrate the construction.

Example 1 Consider the function ϕ defined above.

$$\phi(x, y, z) = \overline{x}\,\overline{z} \vee xy$$

To generate this function using the construction, let $V_U = \{y, z\}$, let $\mathbf{s} = (1, 1, 1)$ and $\mathbf{t} = (0, 0, 0)$ be assignments to (x, y, z), and let $\mathbf{c} = (0, 1)$ be an assignment to (y, z). Then define $g(y, z) = y\overline{z}$. It is easy to verify that g is unate and satisfies all the required properties. □

Example 2 Let $V = \{x_1, x_2, \ldots, x_n\}$. Let $V_U = \emptyset$. Let $\mathbf{s} = (0, 0, \ldots, 0)$ and $\mathbf{t} = (1, 1, \ldots, 1)$ be assignments to $(x_1, x_2, \ldots, x_n)$. Let $g \equiv 0$ (since g is defined on $V_U = \emptyset$, it must be a constant function). Let $\mathbf{c}$ be the empty assignment. The resulting function in MNU is as follows.

$$f(0, 0, \ldots, 0) = f(1, 1, \ldots, 1) = 1$$

$$f(x_1, x_2, \ldots, x_n) = 0 \text{ on all other assignments.}$$

This is the consensus function. Note that if $\mathbf{s}$ and $\mathbf{t}$ were chosen differently, the resulting function would be a generalized-consensus function. □

Example 3 Let $V = \{w, x, y, z\}$, $V_U = \{y, z\}$, and $\mathbf{c} = (0, 0)$. Let $\mathbf{s} = (1, 0, 0, 1)$ and $\mathbf{t} = (0, 1, 1, 0)$ be assignments to (w, x, y, z). Let $g(y, z) = y \vee z$. Again, g satisfies the necessary properties. Thus $f(w, x, y, z)$ is defined as follows:

$$f(0,0,0,0) = f(1,0,0,0) = f(0,1,0,0) = f(1,1,0,0) = f(1,0,0,1) = f(0,1,1,0) = 0$$

$$f(w, x, y, z) = 1 \text{ on all other assignments.}$$

□

1.3 The Size of MNU

Recall that there is only one minimally non-monotone function, $f(x) = \overline{x}$. In contrast, the following argument shows that for all n, there are a doubly exponential number of minimally non-unate functions on n variables. We construct a subclass of functions MNU on n variables. Say $|V| = n$ and n is odd (ignore one variable if n is even). Pick a variable $x \in V$, and let $V_U = V \setminus \{x\}$. Let $\mathbf{c}$ be the all-zeroes assignment to V_U. Consider the set H of assignments to V_U that have exactly half of their variables set to one and half set to zero. Let $\mathbf{s}$ and $\mathbf{t}$ be two fixed complementary assignments to V such that $\mathbf{s}|_{V_U}, \mathbf{t}|_{V_U} \in H$. Now consider the set of all functions $g : \{0,1\}^{V_U} \to \{0,1\}$ such that $g(\mathbf{s}|_{V_U}) = g(\mathbf{t}|_{V_U})$, $g(\mathbf{a}) = 0$ if $\mathbf{a}$ has fewer than half its variables set to 1, and $g(\mathbf{a}) = 1$ if $\mathbf{a}$ has more than half its variables set to 1. The output of g can be either 0 or 1 on all assignments in $H \setminus \{\mathbf{s}|_{V_U}, \mathbf{t}|_{V_U}\}$. All such functions g satisfy the properties required of g in the construction. Since $|H| = \binom{n-1}{(n-1)/2}$, the number of such functions g is $2 * 2^{\binom{n-1}{(n-1)/2}-2}$. Since each such g defines a different function f in MNU, it follows by Stirling's approximation that the number of functions on n variables in MNU is at least $2 * 2^{\binom{n-1}{(n-1)/2}-2} = 2^{\Omega\left(\frac{2^n}{\sqrt{n}}\right)}$. The number of these functions that are non-isomorphic under permutations of the variables is at least $2 * 2^{\binom{n-1}{(n-1)/2}-2}/n!$ which is also $2^{\Omega\left(\frac{2^n}{\sqrt{n}}\right)}$.

1.4 Recognition Problems

In graph theory, the study of forbidden minors has yielded important results on graph recognition problems (see e.g., [9]). Below, we discuss the complexity of the recognition problem for unate functions, and for MNU. However, we do not claim that our characterization of unate functions by forbidden projections has implications for the recognition problem for unate functions.

We define the recognition problem for a class of Boolean functions $\mathcal{F}$ as follows: given the truth table of a Boolean function f, determine whether f is in $\mathcal{F}$.

There is a simple algorithm for solving the recognition problem for unate functions that is based on the properties of unate functions given in Section 2.1. The algorithm checks, for every assignment $\mathbf{a}$, if $f(\mathbf{a}_{x\leftarrow 0}) < f(\mathbf{a}_{x\leftarrow 1})$ or $f(\mathbf{a}_{x\leftarrow 0}) > f(\mathbf{a}_{x\leftarrow 1})$. If there exist assignments $\mathbf{b}$ and $\mathbf{c}$ such that $f(\mathbf{b}_{x\leftarrow 0}) < f(\mathbf{b}_{x\leftarrow 1})$ and $f(\mathbf{c}_{x\leftarrow 0}) > f(\mathbf{c}_{x\leftarrow 1})$, then f is non-unate. Otherwise f is unate. This algorithm clearly runs in time polynomial in its size of its input, which is the truth table of f. The exact complexity of the algorithm depends on the form of the input and the model of computation. Suppose assignments are represented by n-bit numbers and the truth table is an array of length 2^n indexed by those numbers. Then if elementary operations on n-bit numbers (such as taking the bit-wise exclusive OR of two of them) take constant time, the algorithm runs in time $O(n2^n)$.

Note that the existence of the assignments $\mathbf{b}$ and $\mathbf{c}$ such that $f(\mathbf{b}_{x\leftarrow 0}) < f(\mathbf{b}_{x\leftarrow 1})$ and $f(\mathbf{c}_{x\leftarrow 0}) > f(\mathbf{c}_{x\leftarrow 1})$ guarantees that f does contain a (possibly trivial) projection in MNU. However, it is not necessary to explicitly find a projection of f in MNU to determine that f is non-unate.

There is also a simple algorithm for solving the recognition problem for the class of MNU functions. It uses the above method to check whether f itself is non-unate. Also, for each variable x, it checks whether the two projections setting x to 0 and 1 respectively are both unate. Clearly f is in MNU iff f is non-unate and all projections checked are unate.

A variation of the recognition problem for Boolean functions takes as input a Boolean formula for a function, rather than a truth table. When $\mathcal{F}$ consists solely of the function that is 0 on all assignments, the variant recognition problem is actually co-SAT, the complement of SAT, and is therefore co-NP complete. The variant recognition problem for unate functions is also co-NP complete. It is in co-NP because a non-deterministic algorithm can guess and verify the existence of assignments $\mathbf{b}_{x\leftarrow 0}$, $\mathbf{b}_{x\leftarrow 1}$, $\mathbf{c}_{x\leftarrow 0}$, and $\mathbf{c}_{x\leftarrow 1}$ described above, thus proving that f is non-unate. It is co-NP hard by the folowing simple reduction from co-SAT. Given a Boolean formula ϕ, construct $g = \phi \wedge \psi$, where $\psi = (x \vee \overline{y})(\overline{x} \vee y)$ is the non-unate parity function, and x,y are not variables in ϕ. If ϕ is not satisfiable, then $g \equiv 0$ and is therefore unate. If ϕ is satisfiable, then g is non-unate because it contains ψ as a projection.

The variant recognition problem for MNU is co-NP hard, by the following reduction from co-SAT. Let ϕ be a Boolean formula on variable set V. Given ϕ, construct $g = (\phi \wedge \alpha) \vee h$, where $\alpha = \overline{x}\,\overline{y}z \vee x\overline{y}\,\overline{z} \vee \overline{x}y\overline{z} \vee xyz$ expresses the (non-unate) parity function on x, y, z, and h is a formula expressing the consensus function on $V \cup \{x, y, z\}$. If ϕ is not satisfiable, then $g \equiv h$ which is in MNU. If ϕ is satisfiable, then it can be shown that projecting g according to the satisfying assignment of ϕ yields a non-unate function, and thus g is not in MNU. The recognition problem for MNU is also in the complexity class D^{P} since it is the intersection of a co-NP problem (unateness of each projection of f) and an NP problem (non-unateness of f). (See [5] for a discussion of D^{P}.) It is unknown if the recognition problem for MNU is D^{P}-complete.

2 Formal Characterization of MNU

2.1 Definitions

We will consider an *assignment* $\mathbf{a} \in \{0,1\}^V$ to be a a function from V to $\{0,1\}$. All functions appearing in this paper have the form $f : \{0,1\}^V \rightarrow \{0,1\}$ for some Boolean variable set V. The word "function" will be used in place of "Boolean function" for the remainder of the paper.

If $x^{(0)} \in \{0,1\}$, let $\mathbf{a}_{x \leftarrow x^{(0)}}$ denote the assignment $\mathbf{b}$ such that $\mathbf{b}(y) = \mathbf{a}(y)$ for all $y \in V - \{x\}$, and $\mathbf{b}(x) = x^{(0)}$. Let $\mathbf{a}_{\overline{x}}$ denote the assignment $\mathbf{a}_{x \leftarrow \overline{\mathbf{a}(x)}}$. For a set of variables $X \subseteq V$, let $\mathbf{a}_{\overline{X}}$ denote the assignment $\mathbf{a}_{\overline{x_1}, \overline{x_2}, \ldots, \overline{x_n}}$ where $X = \{x_1, x_2, \ldots, x_n\}$. Two assignments are *complementary* if they disagree on every variable.

If $f : \{0,1\}^V \rightarrow \{0,1\}$, $\mathbf{a}$ is an assignment to V, $x \in V$, and $f(\mathbf{a}) \neq f(\mathbf{a}_{\overline{x}})$, then $\mathbf{a}$ is *justifying* for x in f. We say that f *depends on* a variable $x \in V$ if there is a justifying assignment for x in f.

A partial assignment is a function $\mathbf{p} : V \rightarrow \{0, 1, \star\}$. If $\mathbf{p}(x) = \star$ we say that x is not assigned by $\mathbf{p}$. If $\mathbf{p}$ is a partial assignment and $\mathbf{a}$ is an assignment, let $\mathbf{p}/\mathbf{a}$ denote the *extension* of $\mathbf{p}$ obtained by setting the unassigned variables according to $\mathbf{a}$, i.e., $(\mathbf{p}/\mathbf{a})(x) = \mathbf{p}(x)$ if $\mathbf{p}(x) \neq \star$, and $(\mathbf{p}/\mathbf{a})(x) = \mathbf{a}(x)$ otherwise.

If $f : \{0,1\}^V \rightarrow \{0,1\}^V$ and $\mathbf{p}$ is a partial assignment to V, we define $f_{\mathbf{p}}$ to be the function derived from f by "hardwiring" the variables assigned values by $\mathbf{p}$. That is, $f_{\mathbf{p}} : \{0,1\}^V \rightarrow \{0,1\}$ such that for all $\mathbf{a} \in \{0,1\}^V$, $f_{\mathbf{p}}(\mathbf{a}) = f(\mathbf{p}/\mathbf{a})$. For any such $f_{\mathbf{p}}$ there is an associated function $\hat{f}_{\mathbf{p}}$ that is equivalent to $f_{\mathbf{p}}$ on a restricted domain. That is, $\hat{f}_{\mathbf{p}} : \{0,1\}^{V_S} \rightarrow \{0,1\}$ where $V_S = \{x \in V \,|\, \mathbf{p}(x) = \star\}$. Therefore, $f_{\mathbf{p}}(\mathbf{a}) = \hat{f}_{\mathbf{p}}(\mathbf{a}|_{V_S})$ for all $\mathbf{a} \in \{0,1\}^V$. The function $\hat{f}_{\mathbf{p}}$ is a *projection* of f. Clearly, a projection $\hat{f}_{\mathbf{p}}$ is unate iff $f_{\mathbf{p}}$ is unate. Where it will cause no confusion, we will also refer to $f_{\mathbf{p}}$ as a projection.

We say a class of functions C is *projection closed* if any projection of a function from C is also in C. The class of all unate functions is projection closed. A function $g : \{0,1\}^W \rightarrow \{0,1\}$ is a *forbidden* projection of C if there is no projection $\hat{f}_{\mathbf{p}} : \{0,1\}^{V'} \rightarrow \{0,1\}$ of a function f in C such that, under an isomorphism from W to V', $\hat{f}_{\mathbf{p}}$ and g are equivalent.

Let x be a non-unate variable of f. Then there are assignments $\mathbf{a}$ and $\mathbf{b}$ such that

$$\begin{aligned} f(\mathbf{a}_{x\leftarrow 0}) &= 0 = f(\mathbf{b}_{x\leftarrow 1}) \\ f(\mathbf{a}_{x\leftarrow 1}) &= 1 = f(\mathbf{b}_{x\leftarrow 0}). \end{aligned}$$

If additionally, $\mathbf{a}(x) = \overline{\mathbf{b}(x)} = 0$, we say that f is non-unate in x *with respect to the ordered pair of assignments* $(\mathbf{a}, \mathbf{b})$. Such an $\mathbf{a}$ and $\mathbf{b}$ will always exist for any f and non-unate x. The definition via an ordered pair merely forces $f(\mathbf{a}) = f(\mathbf{b}) = 0$ and $\mathbf{a}(x) = \overline{\mathbf{b}(x)} = 0$ which will simplify the statement of several lemmas.

Each partial order $(\{0,1\}^V, \prec_{\mathbf{c}})$ is a *lattice*. To represent a Boolean function with respect to $(\{0,1\}^V, \prec_{\mathbf{c}})$, take the Hasse diagram of $(\{0,1\}^V, \prec_{\mathbf{c}})$ and replace each element $\mathbf{a}$ by $f(\mathbf{a})$. We use these representations to illustrate some of our lemmas.

For example, on the standard Boolean lattice, $\begin{smallmatrix} & 1 & \\ 0 & & 0 \\ & 0 & \end{smallmatrix}$ is the binary AND function. Binary OR is $\begin{smallmatrix} & 1 & \\ 1 & & 1 \\ & 0 & \end{smallmatrix}$. The consensus function on three variables is represented as in Figure 1.

2.2 Characterization

Our goal is to provide a set of conditions characterizing functions in MNU. The conditions center around the existence of the function g from which f is defined. The following lemmas are used to derive the conditions and prove their necessity.

LEMMA 1 *Let $f : \{0,1\}^V \to \{0,1\}$ be a function in MNU. Then $|V| > 1$ and f depends on v for all $v \in V$.*

PROOF Every function of one variable is unate. Since f is non-unate, $|V| > 1$. If f did not depend on some variable x, the two projections which fix x to 0 and fix x to 1 would be identical, and both non-unate.

■

LEMMA 2 *Let $f : \{0,1\}^V \to \{0,1\}$ be a function in MNU that is non-unate in x with respect to the ordered assignment pair $(\mathbf{a}, \mathbf{b})$. Then $\mathbf{a} = \overline{\mathbf{b}}$.*

PROOF By definition of $(\mathbf{a}, \mathbf{b})$ we have $\mathbf{a}(x) = \overline{\mathbf{b}(x)}$ so we need only prove that $\mathbf{a}(y) = \overline{\mathbf{b}(y)}$ for all $y \neq x$. Let y be a variable in V such that $y \neq x$. Assume (for contradiction) that $\mathbf{a}(y) = \mathbf{b}(y)$. Without loss of generality, say $\mathbf{a}(y) = \mathbf{b}(y) = 0$. Let $\mathbf{p}$ be the partial assignment setting y to 0 and all other variables to $\star$. Then $f_{\mathbf{p}}(\mathbf{a}) = f(\mathbf{a})$, $f_{\mathbf{p}}(\mathbf{a}_{\overline{x}}) = f(\mathbf{a}_{\overline{x}})$, $f_{\mathbf{p}}(\mathbf{b}) = f(\mathbf{b})$ and $f_{\mathbf{p}}(\mathbf{b}_{\overline{x}}) = f(\mathbf{b}_{\overline{x}})$ since $\mathbf{p}/\mathbf{a} = \mathbf{a}$ and $\mathbf{p}/\mathbf{b} = \mathbf{b}$. Since $y \neq x$, $f_{\mathbf{p}}$ is also non-unate in x:

$$\begin{aligned} f_{\mathbf{p}}(\mathbf{a}_{x\leftarrow 0}) = f(\mathbf{a}_{x\leftarrow 0}) = \quad 0 \quad &= f(\mathbf{b}_{x\leftarrow 1}) = f_{\mathbf{p}}(\mathbf{b}_{x\leftarrow 1}) \\ f_{\mathbf{p}}(\mathbf{a}_{x\leftarrow 1}) = f(\mathbf{a}_{x\leftarrow 1}) = \quad 1 \quad &= f(\mathbf{b}_{x\leftarrow 0}) = f_{\mathbf{p}}(\mathbf{b}_{x\leftarrow 0}). \end{aligned}$$

But then $f_{\mathbf{p}}$ (and hence $\hat{f}_{\mathbf{p}}$) is non-unate, contradicting f's membership in MNU. Therefore $\mathbf{a}(y) = \overline{\mathbf{b}(y)}$.

■

LEMMA 3 *Let* $f : \{0,1\}^V \to \{0,1\}$ *be a function in* MNU*. If* f *is non-unate in* x *with respect to the ordered pair of assignments* $(\mathbf{a},\mathbf{b})$*, then* $\mathbf{a},\mathbf{b},\mathbf{a}_{\overline{x}},\mathbf{b}_{\overline{x}}$ *are the only justifying assignments for* x *and* $(\mathbf{a},\mathbf{b})$ *is unique.*

PROOF Let $\mathbf{c}$ be an assignment which is justifying for x, that is, $f(\mathbf{c}_{x\leftarrow 0}) \neq f(\mathbf{c}_{x\leftarrow 1})$. We must show that $\mathbf{c} \in \{\mathbf{a},\mathbf{b},\mathbf{a}_{\overline{x}},\mathbf{b}_{\overline{x}}\}$.

Suppose $f(\mathbf{c}_{x\leftarrow 0}) = 0$. If $\mathbf{c}(x) = 0$ then f is non-unate with respect to assignment pair $(\mathbf{c},\mathbf{b})$. If $\mathbf{c}(x) = 1$ then f is non-unate with respect to assignment pair $(\mathbf{c}_{\overline{x}},\mathbf{b})$. By Lemma 2, either $\mathbf{c} = \overline{\mathbf{b}}$ or $\mathbf{c}_{\overline{x}} = \overline{\mathbf{b}}$ (so $\mathbf{c} = \overline{\mathbf{b}}_{\overline{x}}$). Since $\mathbf{a} = \overline{\mathbf{b}}$ we have either $\mathbf{c} = \mathbf{a}$ or $\mathbf{c} = \mathbf{a}_{\overline{x}}$.

If $f(\mathbf{c}_{x\leftarrow 0}) = 1$ then f is non-unate with respect to either assignment pair $(\mathbf{a},\mathbf{c})$ or $(\mathbf{a},\mathbf{c}_{\overline{x}})$, which gives either $\mathbf{c} = \mathbf{b}$ or $\mathbf{c} = \mathbf{b}_{\overline{x}}$. It follows immediately that the pair $(\mathbf{a},\mathbf{b})$ is unique. ■

The next two lemmas give properties of the justifying assignments of the non-unate variables. Lemma 4 is necessary to prove Lemma 5.

LEMMA 4 *Let* f *be a function in* MNU *defined on variable set* V *with* $|V| > 2$*. If* f *is non-unate in* x *with respect to assignment pair* $(\mathbf{a},\mathbf{b})$ *and* f *is also non-unate in* $y \neq x$ *with respect to assignment pair* $(\mathbf{c},\mathbf{d})$ *then either* $\{\mathbf{c},\mathbf{d}\} = \{\mathbf{a},\mathbf{b}\}$ *or* $\{\mathbf{c}_{\overline{y}},\mathbf{d}_{\overline{y}}\} = \{\mathbf{a}_{\overline{x}},\mathbf{b}_{\overline{x}}\}$*. Therefore, there are precisely two assignments that are justifying for both* x *and* y*.*

Before proving Lemma 4, note that it is stated only for $|V| > 2$. If $|V| = 2$ then there are only four possible assignments to the variables of V and f must be either XOR or $\overline{XOR}$. In either case, all four assignments are justifying for both x and y. Thus the conclusion of the lemma does not hold for $|V| = 2$.

PROOF From Lemma 3 we have $\{\mathbf{a},\mathbf{b},\mathbf{a}_{\overline{x}},\mathbf{b}_{\overline{x}}\}$ are the only justifying assignments for x, and $\{\mathbf{c},\mathbf{d},\mathbf{c}_{\overline{y}},\mathbf{d}_{\overline{y}}\}$ are the only justifying assignments for y.

Since $|V| > 2$ and by Lemma 2 $\mathbf{a} = \overline{\mathbf{b}}$ and $\mathbf{c} = \overline{\mathbf{d}}$, assignments $\mathbf{a},\mathbf{b},\mathbf{a}_{\overline{x}},\mathbf{b}_{\overline{x}},\mathbf{a}_{\overline{y}},\mathbf{b}_{\overline{y}}$ are all distinct. Therefore, if $\{\mathbf{a},\mathbf{b}\} = \{\mathbf{c},\mathbf{d}\}$ then $\mathbf{a}_{\overline{x}}$ and $\mathbf{b}_{\overline{x}}$ cannot be justifying assignments for y (because $\{\mathbf{a}_{\overline{x}},\mathbf{b}_{\overline{x}}\} \cap \{\mathbf{c},\mathbf{d},\mathbf{c}_{\overline{y}},\mathbf{d}_{\overline{y}}\} = \emptyset$). Similarly, if $\{\mathbf{a}_{\overline{x}},\mathbf{b}_{\overline{x}}\} = \{\mathbf{c}_{\overline{y}},\mathbf{d}_{\overline{y}}\}$ then $\mathbf{a}$ and $\mathbf{b}$ cannot be justifying assignments for y. This shows that there are at most two assignments which are justifying for both x and y.

We now show that there exist at least two assignments which are justifying for both x and y. It suffices to show that either $\mathbf{c}$ or $\mathbf{c}_{\overline{y}}$ is a justifying assignment for x (If $\mathbf{c}$ is justifying for x then $\{\mathbf{a},\mathbf{b}\} = \{\mathbf{c},\mathbf{d}\}$. If $\mathbf{c}_{\overline{y}}$ is justifying for x then $\{\mathbf{a}_{\overline{x}},\mathbf{b}_{\overline{x}}\} = \{\mathbf{c}_{\overline{y}},\mathbf{d}_{\overline{y}}\}$.)

Assume (for contradiction) that neither $\mathbf{c}$ nor $\mathbf{c}_{\overline{y}}$ is justifying for x. We have $\{\mathbf{c}, \mathbf{c}_{\overline{y}}\} \cap \{\mathbf{a}, \mathbf{b}, \mathbf{a}_{\overline{x}}, \mathbf{b}_{\overline{x}}\} = \emptyset$. Also, $f(\mathbf{c}_{\overline{x},y\leftarrow 0}) = f(\mathbf{c}_{y\leftarrow 0})$, lest $\mathbf{c}$ or $\mathbf{c}_{\overline{y}}$ be justifying for x. Likewise, $f(\mathbf{c}_{\overline{x},y\leftarrow 1}) = f(\mathbf{c}_{y\leftarrow 1})$.

Let $\mathbf{p}$ be the partial assignment on V where $\mathbf{p}(x) = \overline{\mathbf{c}(x)}$, $(= \mathbf{d}(x))$, and $\mathbf{p}(z) = \star$ for all $z \neq x$. Then

$$\begin{aligned} f_{\mathbf{p}}(\mathbf{c}_{\overline{x},y\leftarrow 0}) = f(\mathbf{c}_{\overline{x},y\leftarrow 0}) = f(\mathbf{c}_{y\leftarrow 0}) = \quad 0 \quad = f(\mathbf{d}_{y\leftarrow 1}) = f_{\mathbf{p}}(\mathbf{d}_{y\leftarrow 1}) \\ f_{\mathbf{p}}(\mathbf{c}_{\overline{x},y\leftarrow 1}) = f(\mathbf{c}_{\overline{x},y\leftarrow 1}) = f(\mathbf{c}_{y\leftarrow 1}) = \quad 1 \quad = f(\mathbf{d}_{y\leftarrow 0}) = f_{\mathbf{p}}(\mathbf{d}_{y\leftarrow 0}). \end{aligned}$$

But then y is a non-unate variable of $f_{\mathbf{p}}$, and thus $\hat{f}_{\mathbf{p}}$ is a non-unate non-trivial projection of f, contradicting $f \in MNU$. Therefore, either $\mathbf{c}$ or $\mathbf{c}_{\overline{y}}$ must be a justifying assignment for x. ■

The previous lemma concerned a pair x, y of non-unate variables of f. We generalize to all the non-unate variables of f as follows.

LEMMA 5 (Star Lemma) *Let f be a function in MNU defined on variable set V, $|V| > 2$. If V_N is the set of non-unate variables of f, then there are exactly two assignments $\mathbf{s}$ and $\mathbf{t}$ that are justifying for all variables in V_N, and $\mathbf{t} = \overline{\mathbf{s}}$.*

PROOF It suffices to show that any three non-unate variables share two justifying assignments. Suppose $\mathbf{s}$ and $\mathbf{t}$ are both justifying assignments for x, y, and z. If x, y, and ν share two justifying assignments, those assignments must also be $\mathbf{s}$ and $\mathbf{t}$, since these are the only two assignments which are justifying for both x and y by Lemma 4.

Say that f is non-unate in x with respect to $(\mathbf{a}, \mathbf{b})$, non-unate in y with respect to $(\mathbf{c}, \mathbf{d})$, and non-unate in z with respect to $(\mathbf{q}, \mathbf{r})$. From three applications of Lemma 4 we have the following three facts:

1. Either $\{\mathbf{a}, \mathbf{b}\} = \{\mathbf{c}, \mathbf{d}\}$ or $\{\mathbf{a}_{\overline{x}}, \mathbf{b}_{\overline{x}}\} = \{\mathbf{c}_{\overline{y}}, \mathbf{d}_{\overline{y}}\}$
2. Either $\{\mathbf{c}, \mathbf{d}\} = \{\mathbf{q}, \mathbf{r}\}$ or $\{\mathbf{c}_{\overline{y}}, \mathbf{d}_{\overline{y}}\} = \{\mathbf{q}_{\overline{z}}, \mathbf{r}_{\overline{z}}\}$
3. Either $\{\mathbf{q}, \mathbf{r}\} = \{\mathbf{a}, \mathbf{b}\}$ or $\{\mathbf{q}_{\overline{z}}, \mathbf{r}_{\overline{z}}\} = \{\mathbf{a}_{\overline{x}}, \mathbf{b}_{\overline{x}}\}$.

It follows that either the left column is true or the right column is true (otherwise an inconsistency results). In either case, we have that x, y, and z share two justifying assignments. We label these assignments $\mathbf{s}$ and $\mathbf{t}$. Since $\mathbf{a} = \overline{\mathbf{b}}$ and $\mathbf{a}_{\overline{x}} = \overline{\mathbf{b}_{\overline{x}}}$, we have that $\mathbf{s} = \overline{\mathbf{t}}$. ■

We call these common justifying assignments the *Star assignments* since in the Hasse diagram associated with f, each such assignment is labelled differently from its neighbors, and the subgraph consisting of the assignment and its neighbors form a star graph.

For the remaining lemmas, we will suppose that f is a function in MNU defined on V, and $|V| > 2$. Let $V_U \subsetneq V$ be the set of unate variables of f. Let $V_N = V - V_U$ be the set of non-unate variables of f. Let $\mathcal{P} = \{\mathbf{p} \in \{0,1,\star\}^V \mid \forall x \in V_U, \mathbf{p}(x) = \star; \forall y \in V_N, \mathbf{p}(y) \in \{0,1\}\}$. That is, $\mathcal{P}$ is the set of partial assignments which fix the non-unate variables of f and leave the unate variables assigned. Clearly, $|\mathcal{P}| = 2^{|V_N|}$.

LEMMA 6 *There exists some assignment* $\mathbf{c} \in \{0,1\}^{V_U}$ *such that for all* $\mathbf{p} \in \mathcal{P}$, $(V_U, \prec_{\mathbf{c}})$ *is a monotone orientation of* $\hat{f}_{\mathbf{p}}$.

PROOF We construct $\mathbf{c}$ as follows. Since the variables of V_U are each unate in f, by the definition of unateness that for each $v \in V_U$ at least one of the following two cases holds:

$$\begin{array}{ll} I. & f(\mathbf{a}_{v\leftarrow 0}) = 1 \Rightarrow f(\mathbf{a}_{v\leftarrow 1}) = 1 \text{ for all } \mathbf{a} \in \{0,1\}^V \text{ or} \\ II. & f(\mathbf{a}_{v\leftarrow 0}) = 0 \Rightarrow f(\mathbf{a}_{v\leftarrow 1}) = 0 \text{ for all } \mathbf{a} \in \{0,1\}^V. \end{array}$$

In fact, by Lemma 1, *exactly* one of these two cases holds, for if both held, f would not depend on v.

For each $v \in V_U$, if case I holds, let $\mathbf{c}(v) = 0$. If case II holds, let $\mathbf{c}(v) = 1$. Let $\mathbf{p} \in \mathcal{P}$. For all $\mathbf{a} \in \{0,1\}^V$, $f_{\mathbf{p}}(\mathbf{a}) = f(\mathbf{p}/\mathbf{a})$. Therefore $f_{\mathbf{p}}$ is unate and $(V_U, \prec_{\mathbf{c}})$ is a monotone orientation of $\hat{f}_{\mathbf{p}}$. ■

For a function f with unate variables V_U, we call the assignment $\mathbf{c}$ as constructed above the *monotone setting* of f. The function f is completely determined by the $2^{|V_N|}$ projections $\hat{f}_{\mathbf{p}}$ for $\mathbf{p} \in \mathcal{P}$. By Lemma 6, $(V, \prec_{\mathbf{c}})$ is a monotone orientation for each $\hat{f}_{\mathbf{p}}$. For each $\mathbf{p} \in \mathcal{P}$, let $L_{\mathbf{p}}$ be (the Hasse diagram of) the lattice $(V, \prec_{\mathbf{c}})$, labelled according to $\hat{f}_{\mathbf{p}}$. The next lemma shows that all such labelled lattices are equal, with the exception of the two lattices $\hat{f}_{\mathbf{p_s}}$ and $\hat{f}_{\mathbf{p_t}}$ where $\mathbf{p_s}$ assigns the variables in V_N according to $\mathbf{s}$, and $\mathbf{p_t}$ assigns the variables in V_N according to $\mathbf{t}$. The labelled lattice $L_{\mathbf{p_s}}$ differs from the other lattices exactly on its label for the assignment $\mathbf{s}|_{V_U}$, and the labelled lattice $L_{\mathbf{p_t}}$ differs from the other lattices exactly on its label for the assignment $\mathbf{t}|_{V_U}$ (see Figure 2).

LEMMA 7 *Let* $\mathbf{s}$ *and* $\mathbf{t}$ *be the Star assignments of* f. *Then for all* $\mathbf{p}, \mathbf{q} \in \mathcal{P}$, *and for any assignment* $\mathbf{a} \in \{0,1\}^{V_U}$ *such that* $\mathbf{a} \neq \mathbf{s}|_{V_U}$ *and* $\mathbf{a} \neq \mathbf{t}|_{V_U}$,

$$\hat{f}_{\mathbf{p}}(\mathbf{a}) = \hat{f}_{\mathbf{q}}(\mathbf{a}).$$

PROOF Let $\mathbf{b} \in \{0,1\}^V$ be such that $\mathbf{a} = \mathbf{b}|_{V_U}$. Say $\mathbf{p}$ and $\mathbf{q}$ differ on m variables of V_N, $v_1, v_2, \ldots, v_m$. Then we can "walk" from $\mathbf{p}$ to $\mathbf{q}$ in m steps by flipping bits (changing the assignment of v_i as i goes from 1 to m).

$$\mathbf{p} = \mathbf{p}_0 \rightarrow \mathbf{p}_1 \rightarrow \mathbf{p}_2 \rightarrow \ldots \rightarrow \mathbf{p}_m = \mathbf{q}$$

where $\mathbf{p}_i = \mathbf{p}_{\overline{v_1}, \overline{v_2}, \ldots, \overline{v_i}}$ for $0 \leq i \leq m$.

Because $\mathbf{b}$ differs from both $\mathbf{s}$ and $\mathbf{t}$ on variables of V_U and we are only flipping variables of V_N, we know for each $i \geq 0$ that $(\mathbf{p}_i/\mathbf{b}) \neq \mathbf{s}$ and $(\mathbf{p}_i/\mathbf{b}) \neq \mathbf{t}$. Thus $\mathbf{p}_i/\mathbf{b}$ cannot be a justifying assignment for v_i by Lemma 3 and the Star Lemma, forcing $f_{\mathbf{p}_i}(\mathbf{b}) = f_{\mathbf{p}_{i-1}}(\mathbf{b})$. So

$$f_{\mathbf{p}}(\mathbf{b}) = f_{\mathbf{p}_0}(\mathbf{b}) = \ldots = f_{\mathbf{p}_m}(\mathbf{b}) = f_{\mathbf{q}}(\mathbf{b}).$$

Since $\hat{f}_{\mathbf{p}}(\mathbf{a}) = f_{\mathbf{p}}(\mathbf{b})$ and $f_{\mathbf{q}}(\mathbf{b}) = \hat{f}_{\mathbf{q}}(\mathbf{a})$, it follows that $\hat{f}_{\mathbf{p}}(\mathbf{a}) = \hat{f}_{\mathbf{q}}(\mathbf{a})$. ■

LEMMA 8 *Let* $\mathbf{s}$ *and* $\mathbf{t}$ *be the Star assignments of* f. *Then for any assignment* $\mathbf{b} \in \{0,1\}^V \setminus \{\mathbf{s},\mathbf{t}\}$ *such that* $\mathbf{b}|_{V_U} = \mathbf{s}|_{V_U}$ *or* $\mathbf{b}|_{V_U} = \mathbf{t}|_{V_U}$,

$$f(\mathbf{s}) = f(\mathbf{t}) \neq f(\mathbf{b}).$$

PROOF Let $\mathbf{b} \in \{0,1\}^V \setminus \{\mathbf{s},\mathbf{t}\}$ such that $\mathbf{b}|_{V_U} = \mathbf{s}|_{V_U}$. Say $\mathbf{s}$ and $\mathbf{b}$ differ in non-unate variables $\nu_1, \nu_2, \ldots, \nu_m$. Then $f(\mathbf{s}_{\overline{\nu_1}}) \neq f(\mathbf{s})$ by the Star Lemma. Since neither $\mathbf{s}_{\overline{\nu_1}}$ nor $\mathbf{s}_{\overline{\nu_1},\overline{\nu_2}}$ equal $\mathbf{s}$ or $\mathbf{t}$ (by the assumption that $|V| > 2$), Lemma 3 implies that $f(\mathbf{s}_{\overline{\nu_1}}) = f(\mathbf{s}_{\overline{\nu_1},\overline{\nu_2}})$. We walk to $\mathbf{b}$ using the same reasoning:

$$\overline{f(\mathbf{s})} = f(\mathbf{s}_{\overline{\nu_1}}) = f(\mathbf{s}_{\overline{\nu_1},\overline{\nu_2}}) = \ldots = f(\mathbf{s}_{\overline{\nu_1},\overline{\nu_2},\ldots,\overline{\nu_m}}) = f(\mathbf{b}).$$

By the same argument, $\overline{f(\mathbf{t})} = f(\mathbf{b})$. ■

LEMMA 9 *Let* $\mathbf{c} \in \{0,1\}^{V_U}$ *be the monotone setting of* f. *Let* $\mathbf{s}$ *be a Star assignment of* f. *Then for any partial assignment* $\mathbf{p} \in \mathcal{P}$

$$\hat{f}_{\mathbf{p}}(\mathbf{a}) = 1, \forall \mathbf{a} \in \{0,1\}^{V_U} \textit{ such that } \mathbf{s}|_{V_U} \prec_{\mathbf{c}} \mathbf{a} \textit{ or } \mathbf{t}|_{V_U} \prec_{\mathbf{c}} \mathbf{a}$$
$$\hat{f}_{\mathbf{p}}(\mathbf{a}) = 0, \forall \mathbf{a} \in \{0,1\}^{V_U} \textit{ such that } \mathbf{a} \prec_{\mathbf{c}} \mathbf{s}|_{V_U} \textit{ or } \mathbf{a} \prec_{\mathbf{c}} \mathbf{t}|_{V_U}$$

PROOF (See Figure 3.) Let $\mathbf{p} \in \mathcal{P}$. Since $\hat{f}_{\mathbf{p}}$ is unate with monotone orientation $(V_U, \prec_{\mathbf{c}})$, if $\hat{f}_{\mathbf{p}}(\mathbf{a}_0) = 1$ for some assignment $\mathbf{a}_0$ then $\hat{f}_{\mathbf{p}}(\mathbf{a}) = 1$ for all $\mathbf{a}_0 \prec_{\mathbf{c}} \mathbf{a}$. Similarly, if $\hat{f}_{\mathbf{p}}(\mathbf{a}_0) = 0$ then $\hat{f}_{\mathbf{p}}(\mathbf{a}) = 0$ for all $\mathbf{a}$ such that $\mathbf{a} \prec_{\mathbf{c}} \mathbf{a}_0$. Lemma 7 implies that $\hat{f}_{\mathbf{p}}(\mathbf{a}) = \hat{f}_{\mathbf{q}}(\mathbf{a})$ for all $\mathbf{a}$ such that $\mathbf{s}|_{V_U} \prec_{\mathbf{c}} \mathbf{a}$, $\mathbf{t}|_{V_U} \prec_{\mathbf{c}} \mathbf{a}$, $\mathbf{a} \prec_{\mathbf{c}} \mathbf{s}|_{V_U}$ or $\mathbf{a} \prec_{\mathbf{c}} \mathbf{t}|_{V_U}$. Therefore, it suffices to show that there are partial assignments $\mathbf{p}_1$, $\mathbf{p}_2$, $\mathbf{p}_3$ and $\mathbf{p}_4$ in $\mathcal{P}$ such that

$$f_{\mathbf{p}_1}(\mathbf{s}) = 0, \quad f_{\mathbf{p}_2}(\mathbf{t}) = 1$$
$$f_{\mathbf{p}_3}(\mathbf{s}) = 1, \quad f_{\mathbf{p}_4}(\mathbf{t}) = 0.$$

By Lemma 8, $f(\mathbf{s}) = f(\mathbf{t})$. Without loss of generality, assume that $f(\mathbf{s}) = f(\mathbf{t}) = 0$. Then let $\mathbf{p}_1$ be the partial assignment fixing the variables of V_N as they appear in $\mathbf{s}$, and let $\mathbf{p}_4 = \overline{\mathbf{p}_1}$. Now $\mathbf{p}_1/\mathbf{s} = \mathbf{s}$ and $\mathbf{p}_4/\mathbf{t} = \mathbf{t}$ so $f_{\mathbf{p}_1}(\mathbf{s}) = f(\mathbf{s}) = 0$ and $f_{\mathbf{p}_4}(\mathbf{t}) = f(\mathbf{t}) = 0$. Let $\mathbf{p}_3$ differ from $\mathbf{p}_1$ in precisely one variable

$v \in V_N$. Since $\mathbf{p}_1/\mathbf{s} = \mathbf{s}$ is a justifying assignment for v by the Star Lemma, $f_{\mathbf{p}_3}(\mathbf{s}) = 1$. Simliarly, let $\mathbf{p}_2$ differ in one variable from $\mathbf{p}_4$ so that $f_{\mathbf{p}_2}(\mathbf{s}) = 1$. ■

We present the characterization of MNU in the following theorem.

THEOREM 1 *A Boolean function $f : \{0,1\}^V \to \{0,1\}$ is minimally non-unate if and only if there exists a subset $V_U \subset V$, an assignment $\mathbf{c} \in \{0,1\}^{V_U}$, two complementary assignments $\mathbf{s}, \mathbf{t} \in \{0,1\}^V$, and a Boolean function g defined on V_U such that:*

a) $(V, \prec_{\mathbf{c}})$ is a monotone orientation of g

b) $g(\mathbf{s}|_{V_U}) = g(\mathbf{t}|_{V_U})$

c) for all $\mathbf{a} \in \{0,1\}^{V_U}$, if $\mathbf{a} \prec_{\mathbf{c}} \mathbf{s}|_{V_U}$ or $\mathbf{a} \prec_{\mathbf{c}} \mathbf{t}|_{V_U}$, then $g(\mathbf{a}) = 0$
for all $\mathbf{a} \in \{0,1\}^{V_U}$, if $\mathbf{s}|_{V_U} \prec_{\mathbf{c}} \mathbf{a}$ or $\mathbf{t}|_{V_U} \prec_{\mathbf{c}} \mathbf{a}$, then $g(\mathbf{a}) = 1$

d) for any assignment $\mathbf{b} \in \{0,1\}^V$

$$f(\mathbf{b}) = \begin{cases} \overline{g(\mathbf{b}|_{V_U})} & \textit{if } \mathbf{b} \in \{\mathbf{s}, \mathbf{t}\} \\ g(\mathbf{b}|_{V_U}) & \textit{otherwise} \end{cases}$$

e) V_U is the set of unate variables of f

Before presenting the proof of the theorem, we discuss its relation to the construction given in the introduction. Condition e), stating that V_U is the set of unate variables of f, is included in the theorem for clarity; it follows from the other conditions. Let $v \in V_U$. If $\mathbf{b}, \mathbf{b}_{\overline{v}} \in \{0,1\}^V$, then by properties a) and c), and the definition of f, $\mathbf{b}|_{V_U} \prec_{\mathbf{c}} \mathbf{b}_{\overline{v}}|_{V_U}$ implies that $f(\mathbf{b}) \leq f(\mathbf{b}_{\overline{v}})$. Therefore, f is unate in v; it is monotone in v if $\mathbf{c}(v) = 0$, and anti-monotone in v if $\mathbf{c}(v) = 1$. If $v \in V - V_U$, then by property b) and the definition of f, $f(\mathbf{s}_{\overline{v}}) \neq f(\mathbf{s}) = f(\mathbf{t}_{\overline{v}}) \neq f(\mathbf{t})$. Since $\mathbf{s}$ and $\mathbf{t}$ are complementary, it follows that f is not unate in v. Thus the theorem still holds if Condition e) is deleted. The conditions cannot all hold if $|V_U| \neq 0$ and $\mathbf{c} \in \{\mathbf{s}|_{V_U}, \mathbf{t}|_{V_U}\}$ because condition b) requires that $g(\mathbf{s}|_{V_U}) = g(\mathbf{t}|_{V_U})$, and in this case condition c) implies that $g(\mathbf{s}|_{V_U}) \neq g(\mathbf{t}|_{V_U})$. A consequence is that $|V_U| \neq 1$ for any function in MNU. The construction given in the introduction follows immediately from the theorem and these observations.

PROOF Let $f : \{0,1\}^V$. The theorem is vacuously true if $|V| \leq 1$. Suppose $|V| = 2$, and f obeys the conditions of the theorem. As argued above, $|V_U| \neq 1$, and therefore since V_U is a proper subset of V, $|V_U| = 0$. In this case f is either XOR or $\overline{XOR}$. XOR is a generalized consensus function, and both it and $\overline{XOR}$ are in MNU. Conversely, if $|V| = 2$ and f is in MNU, f must be equal to either XOR or $\overline{XOR}$, and in either case f satisfies the conditions of the theorem.

Assume now that $|V| > 2$.

$f \in MNU \Rightarrow f$ satisfies Conditions

Let V_U be the unate variables of f, and let $V_N = V - V_U$ be the non-unate variables of f. As in the proof of the lemmas, let $\mathcal{P}$ be the set of partial assignments to V that assign values to variables in V_N and leave the variables in V_U unassigned. Let $\mathbf{s}$ and $\mathbf{t}$ be the Star assignments of f as described in Lemma 5.

Suppose $|V_U| = \emptyset$. By Lemma 8, $f(\mathbf{s}) \neq f(\mathbf{t})$ and for any $\mathbf{a} \in \{0,1\}^V \setminus \{\mathbf{s}, \mathbf{t}\}$, $f(\mathbf{a}) \neq f(\mathbf{s}) = f(\mathbf{t})$. Therefore f is either a generalized consensus function or the negation of such a function. In either case, f is in MNU.

Now assume $|V_U| > 0$. Let $\mathbf{c}$ be the monotone setting of f (defined in the proof of Lemma 6). We define g as follows. For all $\mathbf{a} \in \{0,1\}^V$,

$$g(\mathbf{a}|_{V_U}) = \begin{cases} f(\mathbf{a}) & \text{if } \mathbf{a} \notin \{\mathbf{s}, \mathbf{t}\} \\ \overline{f(\mathbf{a})} & \text{if } \mathbf{a} \in \{\mathbf{s}, \mathbf{t}\}. \end{cases}$$

We now prove that g is well-defined. Suppose $\mathbf{a}|_{V_U} = \mathbf{b}|_{V_U}$ but $\mathbf{a} \neq \mathbf{b}$. We need to show that the definition of $g(\mathbf{a}|_{V_U})$ is consistent with the definition of $g(\mathbf{b}|_{V_U})$. There are two cases.

First, suppose $\mathbf{a}|_{V_U} \notin \{\mathbf{s}|_{V_U}, \mathbf{t}|_{V_U}\}$. Then $\mathbf{a}, \mathbf{b} \notin \{\mathbf{s}, \mathbf{t}\}$, and applying the definition of g to $\mathbf{a}$ and $\mathbf{b}$ yields $g(\mathbf{a}|_{V_U}) = f(\mathbf{a})$ and $g(\mathbf{b}|_{V_U}) = f(\mathbf{b})$. We need to show $f(\mathbf{a}) = f(\mathbf{b})$. Let $\mathbf{p_a}, \mathbf{p_b} \in \mathcal{P}$ be the partial assignments that set the variables in V_N according to $\mathbf{a}$ and $\mathbf{b}$ respectively. Then, $\mathbf{a}/\mathbf{p_a} = \mathbf{a}$ and $\mathbf{b}/\mathbf{p_b} = \mathbf{b}$. Since $\mathbf{a}|_{V_U} = \mathbf{b}|_{V_U}$, $\mathbf{b}/\mathbf{p_a} = \mathbf{a}$ and $\mathbf{a}/\mathbf{p_b} = \mathbf{b}$. Lemma 7 implies $f_{\mathbf{p_b}}(\mathbf{a}) = f_{\mathbf{p_a}}(\mathbf{a}) = f(\mathbf{a})$. Since $f_{\mathbf{p_b}}(\mathbf{a}) = f(\mathbf{p_b}/\mathbf{a}) = f(\mathbf{b})$, we have $f(\mathbf{a}) = f(\mathbf{b})$ as desired.

For the second case, assume $\mathbf{a}|_{V_U} \in \{\mathbf{s}|_{V_U}, \mathbf{t}|_{V_U}\}$ (and recall that $\mathbf{a}|_{V_U} = \mathbf{b}|_{V_U}$). Without loss of generality, assume $\mathbf{a}|_{V_U} = \mathbf{s}|_{V_U}$. If neither $\mathbf{a}$ nor $\mathbf{b}$ is equal to $\mathbf{s}$, then applying the definition of g to $\mathbf{a}$ and $\mathbf{b}$ yields $g(\mathbf{a}|_{V_U}) = f(\mathbf{a})$ and $g(\mathbf{b}|_{V_U}) = f(\mathbf{b})$. By Lemma 8 we have $f(\mathbf{a}) \neq f(\mathbf{s})$ and $f(\mathbf{b}) \neq f(\mathbf{s})$, so $f(\mathbf{a}) = f(\mathbf{b})$ as desired. Suppose however that $\mathbf{a} = \mathbf{s}$ (the case where $\mathbf{b} = \mathbf{s}$ is symmetric). Because $\mathbf{a}|_{V_U} = \mathbf{b}|_{V_U}$ and $\mathbf{t} = \overline{\mathbf{s}}$, $\mathbf{b} \neq \mathbf{t}$. Also, $\mathbf{b} \neq \mathbf{s}$ because $\mathbf{a} \neq \mathbf{b}$. By Lemma 8 $f(\mathbf{a}) = f(\mathbf{s}) \neq f(\mathbf{b})$. Since $\mathbf{a} = \mathbf{s}$, $g(\mathbf{a}|_{V_U})$ is defined to be equal to $\overline{f(\mathbf{a})}$, and since $\mathbf{b} \notin \{\mathbf{s}, \mathbf{t}\}$, $g(\mathbf{b}|_{V_U})$ is defined to be equal to $f(\mathbf{b})$. So the definition of g is also consistent here. This completes the proof of the consistency of g.

Let $\mathbf{p} \in \mathcal{P}$ such that $\mathbf{p}(x) = \mathbf{s}(x)$ for some $x \in V_N$. Note that since $\mathbf{p}(x) = \mathbf{s}(x)$ we have $f_{\mathbf{p}}(\mathbf{a}) = g(\mathbf{a}|_{V_U})$ for all $\mathbf{a} \in \{0,1\}^V \setminus \{\mathbf{s}\}$. We now show that all the conditions hold.

a). By Lemma 6, $(V_U, \prec_{\mathbf{c}})$ is a monotone orientation of $\hat{f}_{\mathbf{p}}$. The function g can differ from $\hat{f}_{\mathbf{p}}$ only on $\mathbf{s}|_{V_U}$. By Lemma 9

$$\hat{f}_{\mathbf{p}}(\mathbf{a}) = 1, \forall \mathbf{a} \in \{0,1\}^{V_U} \text{ such that } \mathbf{s}|_{V_U} \prec_{\mathbf{c}} \mathbf{a} \text{ or } \mathbf{t}|_{V_U} \prec_{\mathbf{c}} \mathbf{a}$$
$$\hat{f}_{\mathbf{p}}(\mathbf{a}) = 0, \forall \mathbf{a} \in \{0,1\}^{V_U} \text{ such that } \mathbf{a} \prec_{\mathbf{c}} \mathbf{s}|_{V_U} \text{ or } \mathbf{a} \prec_{\mathbf{c}} \mathbf{t}|_{V_U}.$$

Thus, $(V_U, \prec_{\mathbf{c}})$ is also a monotone orientation for g.

b). $g(\mathbf{s}|_{V_U}) = \overline{f(\mathbf{s})}$ and $\overline{f(\mathbf{t})} = g(\mathbf{t}|_{V_U})$ by the definition of g. By the Star Lemma, $f(\mathbf{s}) = f(\mathbf{t})$, so $g(\mathbf{s}|_{V_U}) = g(\mathbf{t}|_{V_U})$.

c). Follows from Lemma 9 because $(\mathbf{p}/\mathbf{s})|_{V_U} = \mathbf{s}|_{V_U}$, $(\mathbf{p}/\mathbf{t})|_{V_U} = \mathbf{t}|_{V_U}$, and for all $\mathbf{b} \in \{0,1\}^V \setminus \{\mathbf{s}, \mathbf{t}\}$, $g(\mathbf{b}|_{V_U}) = f_{\mathbf{p}}(\mathbf{b})$.

d). From definition of g.

e). From definition of V_U.

f satisfies Conditions $\Rightarrow f \in MNU$

Suppose there exist $V_U, \mathbf{c}, \mathbf{s}, \mathbf{t}$ and g as described in the theorem. We show that $f \in MNU$. Applying conditions d) and b), we see that

$$\begin{aligned} f(\mathbf{s}) &= \overline{g(\mathbf{s}|_{V_U})} && \text{from d)} \\ &= \overline{g(\mathbf{t}|_{V_U})} && \text{from b)} \\ &= f(\mathbf{t}) && \text{from d).} \end{aligned}$$

Say $\nu \in V_N$. Then

$$\begin{aligned} f(\mathbf{s}_{\overline{\nu}}) &= g(\mathbf{s}_{\overline{\nu}}|_{V_U}) && \text{from d)} \\ &= g(\mathbf{s}|_{V_U}) && \text{since } \nu \notin V_U \\ &= \overline{f(\mathbf{s})} && \text{from d).} \end{aligned}$$

Likewise, $f(\mathbf{t}_{\overline{\nu}}) = \overline{f(\mathbf{t})}$.

Without loss of generality, assume $\mathbf{s}(\nu) = 0$ and $f(\mathbf{s}) = 0$. Then, since $\mathbf{s}$ and $\mathbf{t}$ are complementary, $\mathbf{t}(\nu) = 1$ and we have

$$\begin{aligned} f(\mathbf{s}) = f(\mathbf{s}_{\nu \leftarrow 0}) = \quad 0 \quad &= f(\mathbf{t}_{\nu \leftarrow 1}) = f(\mathbf{t}) \\ f(\mathbf{s}_{\overline{\nu}}) = f(\mathbf{s}_{\nu \leftarrow 1}) = \quad 1 \quad &= f(\mathbf{t}_{\nu \leftarrow 0}) = f(\mathbf{t}_{\overline{\nu}}) \end{aligned}$$

so f is a non-unate function.

To show that every non-trivial projection of f is unate, it suffices to show the unateness of the projections produced by fixing the value of exactly one variable (because then all projections of such unate functions are unate). Consider any partial assignment $\mathbf{p}$ which fixes a single variable $x \in V$ and leaves all other variables undetermined. We show that $\hat{f}_{\mathbf{p}}$ is unate in each of its variables and hence is a unate function.

For any variable $\mu \in V_U \setminus \{x\}$, condition a) implies that μ is a unate variable in f ($\mathbf{c}$ will serve as the monotone setting of each $\hat{f}_{\mathbf{p}}$, so it must be the monotone setting for f). A unate variable remains unate under any projection of the function, so $\hat{f}_{\mathbf{p}}$ is unate in μ. Now let $\nu \in V_N \setminus \{x\}$. For any assignment $\mathbf{a}$ to V such that $\mathbf{a} \notin \{\mathbf{s}, \mathbf{t}, \mathbf{s}_{\overline{\nu}}, \mathbf{t}_{\overline{\nu}}\}$ we know from condition e) that $f(\mathbf{a}) = g(\mathbf{a}|_{V_U}) = g(\mathbf{a}_{\overline{\nu}}|_{V_U}) = f(\mathbf{a}_{\overline{\nu}})$. Thus the only justifying assignments for ν are $\{\mathbf{s}, \mathbf{t}, \mathbf{s}_{\overline{\nu}}, \mathbf{t}_{\overline{\nu}}\}$. Since $\mathbf{s}$ and $\mathbf{t}$ are complementary, we have either $\mathbf{p}(x) = \mathbf{s}(x)$ or $\mathbf{p}(x) = \mathbf{t}(x)$. Without loss of generality, say $\mathbf{p}(x) = \mathbf{s}(x) = \overline{\mathbf{t}(x)}$. Then

$$\begin{aligned} f_{\mathbf{p}}(\mathbf{t}) &= f(\mathbf{t}_{\overline{x}}) && \text{since } \mathbf{p}/\mathbf{t} = \mathbf{t}_{\overline{x}} \\ &= f(\mathbf{t}_{\overline{x},\overline{\nu}}) && \text{since } \mathbf{t}_{\overline{x}} \text{ is not justifying for } \nu \\ &= f_{\mathbf{p}}(\mathbf{t}_{\overline{\nu}}). \end{aligned}$$

So $\mathbf{t}$ and $\mathbf{t}_{\overline{\nu}}$ are not justifying for ν in $f_{\mathbf{p}}$, leaving $\mathbf{s}$ and $\mathbf{s}_{\overline{\nu}}$ as the only possible justifying assignments for ν in $f_{\mathbf{p}}$. This implies that $f_{\mathbf{p}}$ is either monotone or anti-monotone, i.e., unate, in ν.

Therefore, since f is non-unate but any non-trivial projection of f is unate, $f \in MNU$. ■

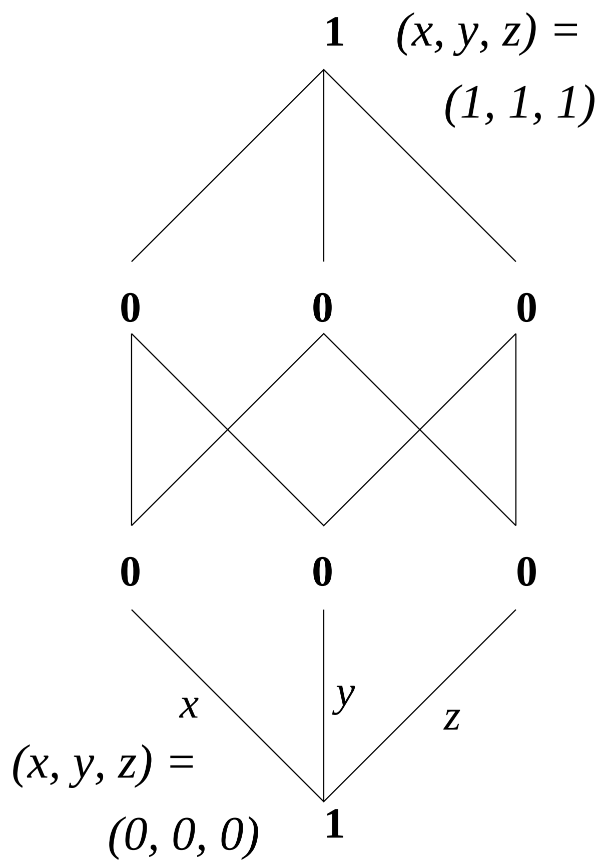

Fig. 1: Diagram of the consensus function

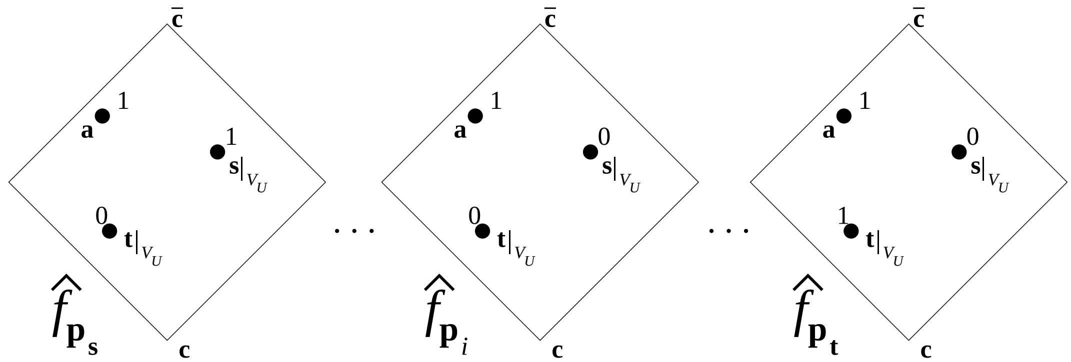

Fig. 2: Lemmas 7 and 8. All intermediate lattices $\hat{f}_{\mathbf{p}_i}$ are identical. Lattices $\hat{f}_{\mathbf{p}_\mathbf{s}}$ and $\hat{f}_{\mathbf{p}_\mathbf{t}}$ differ from $\hat{f}_{\mathbf{p}_i}$ only on assignments $\mathbf{s}|_{V_U}$ and $\mathbf{t}|_{V_U}$.

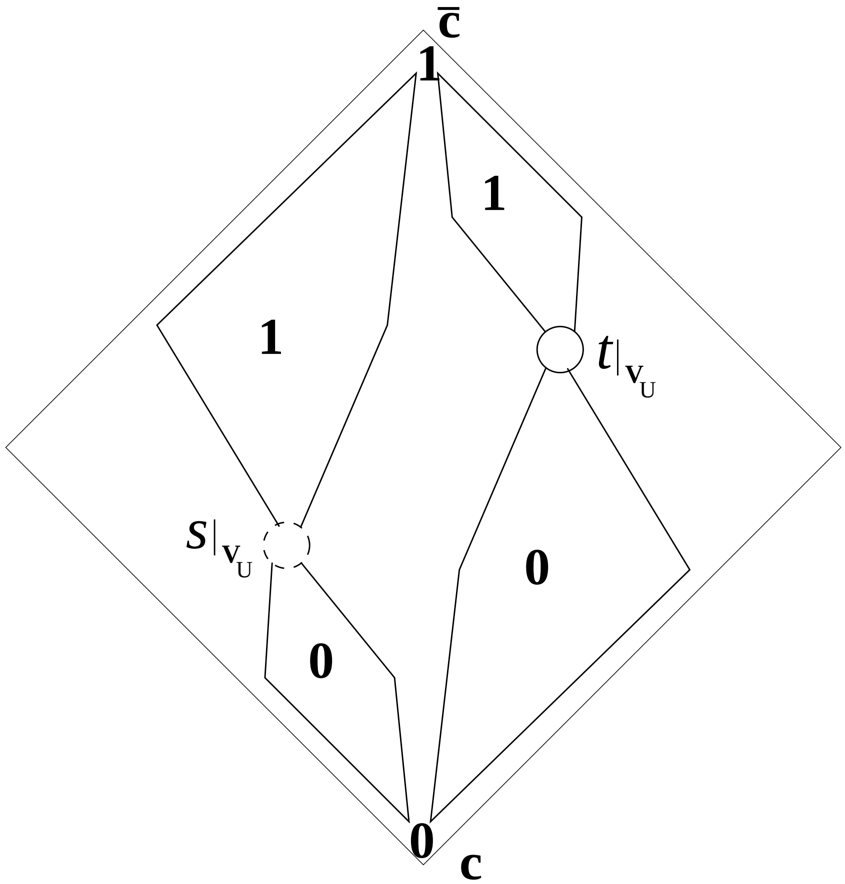

Fig. 3: Lemma 9. The lattice of $\hat{f}_{\mathbf{P}}$.

References

[1] D. Angluin, L. Hellerstein, and M. Karpinski, Learning read-once formulas with queries, Journal of the ACM (1993) 185–210.

[2] T. Brylawski, Constructions, in: N. White, ed., Theory of Matroids (Cambridge University Press, 1986) 127–223.

[3] N. Bshouty, Exact Learning via the Monotone Theory, Proceedings of the 34th Annual IEEE Symposium on the Foundations of Computer Science (1993) 302–311.

[4] L. Hellerstein, K. Pillaipakkamnatt, V. Raghavan, and D. Wilkins, How many queries are needed to learn? Proceedings of the Twenty Seventh Annual ACM Symposium on Theory of Computing (1995) 190–199.

[5] D. S. Johnson, A Catalog of Complexity Classes, J. van Leeuwen, ed., Handbook of Theoretical Computer Science Vol. A (MIT Press, 1990) 67–161.

[6] Z. Kohavi, Switching and Finite Automata Theory (McGraw-Hill, New York, 1978).

[7] R. McNaughton, Unate truth functions, IRE Transactions on Electronic Computers (March 1961) 1–6.

[8] P. D. Seymour, The forbidden minors of binary clutters, J. London Math. Soc. (2), 12 (1976) 356–360.

[9] J. van Leeuwen, Graph Algorithms, J. van Leeuwen, ed., Handbook of Theoretical Computer Science Vol. A (MIT Press, 1990) 525–631.