On the Semantics of Local Characterizations for Linear-Invariant Properties*

Arnab Bhattacharyya†
MIT CSAIL
abhatt@mit.edu

Elena Grigorescu‡
Georgia Tech
elena@cc.gatech.edu

Jakob Nordström§
KTH Royal Institute of Technology
jakobn@kth.se

Ning Xie¶
MIT CSAIL
ningxie@csail.mit.edu

April 13, 2011

Abstract

A property of functions on a vector space is said to be *linear-invariant* if it is closed under linear transformations of the domain. Linear-invariant properties are some of the most well-studied properties in the field of property testing. Testable linear-invariant properties can always be characterized by so-called local constraints, and of late there has been a rapidly developing body of research investigating the testability of linear-invariant properties in terms of their descriptions using such local constraints. One problematic aspect that has been largely ignored in this line of research, however, is that syntactically distinct local characterizations need not at all correspond to semantically distinct properties. In fact, there are known fairly dramatic examples where seemingly infinite families of properties collapse into a small finite set that was already well-understood.

In this work, we therefore initiate a systematic study of the *semantics* of local characterizations of linear-invariant properties. For such properties the local characterizations have an especially nice structure in terms of forbidden patterns on linearly dependent sets of vectors, which can be encoded formally as *matroid constraints*. We develop techniques for determining, given two such matroid constraints, whether these constraints encode identical or distinct properties, and show for a fairly broad class of properties that these techniques provide necessary and sufficient conditions for deciding between the two cases. We use these tools to show that recent (syntactic) testability results indeed provide an infinite number of infinite strict hierarchies of (semantically) distinct testable locally characterized linear-invariant properties.

1 Introduction

Call a function $f : \{0,1\}^n \to \{0,1\}$ *linear* if it is a homomorphism between the groups $\{0,1\}^n$ and $\{0,1\}$ with the group operations being bitwise XOR addition. Then f is linear if and only if it satisfies the "local

*A weaker set of preliminary results from this paper was previously reported in [BGNX10].

†Supported in part by a DOE Computational Science Graduate Fellowship and NSF Awards 0514771, 0728645, and 0732334.

‡Supported in part by NSF award CCR-0829672 and NSF award 1019343 to the Computing Research Association for the Computing Innovation Fellowship Program.

§Part of this work performed while at MIT supported by grants from the Royal Swedish Academy of Sciences, the Ericsson Research Foundation, the Sweden-America Foundation, the Foundation Olle Engkvist Byggmästare, the Sven and Dagmar Salén Foundation, and the Foundation Blanceflor Boncompagni-Ludovisi, née Bildt.

¶Supported by NSF Awards 0514771, 0728645 and 0732334.

constraints" $f(\mathbf{x}_1 + \mathbf{x}_2) = f(\mathbf{x}_1) + f(\mathbf{x}_2)$ for all $\mathbf{x}_1, \mathbf{x}_2 \in \{0,1\}^n$. But linearity can be characterized by many other families of local constraints. For instance, $f(\mathbf{x}_1+\mathbf{x}_2+\mathbf{x}_3+\mathbf{x}_4) = f(\mathbf{x}_1)+f(\mathbf{x}_2)+f(\mathbf{x}_3)+f(\mathbf{x}_4)$ for all $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4 \in \{0,1\}^n$ also holds exactly when f is linear. A natural question is how we can find all the families of local constraints that characterize a particular property of a function. Also, given two families of constraints, how do we know if they characterize the same or distinct properties?

These questions lie at the heart of a fundamental problem of *property testing*, namely to determine what makes a property testable. The field of property testing, as initiated by [BLR93, BFL91] and defined formally by [RS96, GGR98], is the study of algorithms that query their input a very small number of times and with high probability decide correctly whether it satisfies a given property or is "far" from satisfying that property. A property is called *testable*, or sometimes *strongly testable* or *locally testable*, if the number of queries can be made independent of the size of the object without affecting the correctness probability. Perhaps quite surprisingly, it has been found that a large number of natural properties meet this strong requirement; see e.g. the surveys [Rub06, Ron09, Sud10] for more information. The aforementioned fundamental problem is to find a combinatorial description that captures these testable properties.

To explain the connection to local constraints, consider a property which has a so-called one-sided tester that always accepts if the input satisfies the property and rejects with some nonzero probability otherwise. This is the case for most natural testable properties, and especially for testable algebraic properties, which are the focus of this work. It is well-known [KS08, GR09] that any such property must *necessarily* have a characterization in terms of local constraints, in the sense we described above for linearity. It has proven particularly productive to use such characterizations in terms of local constraints to understand testability of properties. Indeed, this approach was suggested even in the early seminal work [RS96], and the project of finding combinatorial descriptions of the testable graph properties was successfully completed in [AFNS06] employing local constraints. In the setting of algebraic properties, and more specifically linear-invariant properties (which we will define shortly), there has also been a sequence of recent results [Gre05, BCSX09, Sha09, KSV10, BGS10, CSX11], showing testability in terms of local characterizations.

However, one issue that has been largely ignored in this context is that one and the same property can have many different local characterizations. The example of linearity above illustrates this situation. A quite dramatic generalization of this example was given in [BCSX09], where it was shown that any property of functions $f : \{0,1\}^n \to \{0,1\}$ that is generated by applying any non-monotone constraint on $f(\mathbf{x}_1), \ldots, f(\mathbf{x}_k), f(\mathbf{x}_1 + \cdots + \mathbf{x}_k)$ for all $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \{0,1\}^n$ in fact collapses into one of a small finite set of already well-studied properties, regardless of k and of the exact form of the constraint placed on the evaluations of the function! This result clearly demonstrates the need for a better understanding of the correspondence between local constraints and the properties they characterize. If we are to analyze the testability of properties through their local characterizations, we must also be able to understand the *semantics* of local constraints. It is not sufficient to have just the testability results for two properties characterized by two different families of local constraints, but we also want to understand whether these two properties are truly different, meaning that the testability of one is not automatically implied by the testability of the other, or whether they are essentially the same.

The purpose of the current work is to initiate a systematic study of these issues in the context of algebraic properties. Let us start by making some of the above discussion formal.

1.1 Local Characterizations, Linear Invariance, and Matroid Freeness

Consider properties of functions mapping some domain D to some range R. We let $\{\mathsf{D} \to \mathsf{R}\}$ denote the set of all such functions, and we describe a property $\mathcal{P}$ by the set of functions $\mathcal{P} \subseteq \{\mathsf{D} \to \mathsf{R}\}$ satisfying the property. For a positive integer k, a *k-local constraint* $C = (\mathbf{a}_1, \ldots, \mathbf{a}_k; S)$ is given by k elements $\mathbf{a}_1, \ldots, \mathbf{a}_k \in \mathsf{D}$ and a set $S \subseteq \mathsf{R}^k$. A function $f : \mathsf{D} \to \mathsf{R}$ is said to satisfy the constraint C if $(f(\mathbf{a}_1), \ldots, f(\mathbf{a}_k)) \notin S$, and a property $\mathcal{P} \subseteq \{\mathsf{D} \to \mathsf{R}\}$ satisfies C if every function $f \in \mathcal{P}$ does so. A collection of k-local constraints $C_1, \ldots, C_m$ *k-locally characterizes* the property $\mathcal{P}$ if $f \in \mathcal{P}$ holds exactly when f satisfies C_i for every $i \in [m]$. If k is not a function of the size of D, the property is said to be *locally characterized*.

In the following, we will consider the domain $\mathsf{D} = \mathbb{F}^n$ for some finite field $\mathbb{F}$. A property $\mathcal{P} \subseteq \{\mathbb{F}^n \to \mathsf{R}\}$ is said to be *linear-invariant* if for every function $f \in \mathcal{P}$ and for every $\mathbb{F}$-linear map $L : \mathbb{F}^n \to \mathbb{F}^n$ the function $f \circ L$ is also in $\mathcal{P}$, where $\circ$ denotes function composition $(f \circ L)(x) = f(L(x))$. Linear invariance has been found to be a natural way to abstract the key features of some of the most well-studied algebraic properties such as linearity [BLR93], being a polynomial of low degree [RS96, AKK+05], being a homogeneous polynomial [KS08], having a small number of nonzero Fourier coefficients [GOS+09], et cetera. For linear-invariant properties, local characterizations have an especially nice structure. To see this, observe that specifying a property to be linear-invariant also enforces a symmetry among the local constraints satisfied by the property. If a linear-invariant property $\mathcal{P}$ satisfies $C = (\mathbf{a}_1, \ldots, \mathbf{a}_k; S)$, then it must also conform to the constraint $C \circ L = (L(\mathbf{a}_1), \ldots, L(\mathbf{a}_k); S)$ for any linear map $L : \mathbb{F}^n \to \mathbb{F}^n$. Thus, $\mathcal{P}$ must satisfy all constraints in the *orbit* of C, i.e., the family of constraints $\{C \circ L \mid \text{linear } L : \mathbb{F}^n \to \mathbb{F}^n\}$. It is straightforward to verify that one can encode the orbit of a constraint $C = (\mathbf{a}_1, \ldots, \mathbf{a}_k; S)$ by a tuple $(\mathbf{v}_1, \ldots, \mathbf{v}_k; S)$ for vectors $\mathbf{v}_i$ in the smaller space $\mathbb{F}^k$ such that the orbit of C equals $\{(L(\mathbf{v}_1), \ldots, L(\mathbf{v}_k); S) \mid \text{linear } L : \mathbb{F}^k \to \mathbb{F}^n\}$. For this, the exact identity of the elements $\mathbf{v}_1, \ldots, \mathbf{v}_k$ is not important – the only thing that matters is the linear dependencies between them. Hence, it is convenient to think of them as the representation of a *linear matroid*[1] $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$. Furthermore, it is not hard to show that if a linear-invariant property is locally characterized, then it can be characterized by the orbit of a *single* constraint (see Proposition 1 of [BCSX10]). We thus have the following fact.

Fact 1.1. *If a linear-invariant property $\mathcal{P} \subseteq \{\mathbb{F}^n \to \mathsf{R}\}$ is k-locally characterized, there exists a linear matroid $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_K\} \subseteq \mathbb{F}^K$ and a set $S \subseteq \mathsf{R}^K$ for some $K = O(\exp(k))$ such that a function $f : \mathbb{F}^n \to \mathsf{R}$ is in $\mathcal{P}$ if and only if $(f(L(\mathbf{v}_1)), \ldots, f(L(\mathbf{v}_K))) \notin S$ for all $\mathbb{F}$-linear maps $L : \mathbb{F}^K \to \mathbb{F}^n$.*

In this case, we say that $\mathcal{P}$ is characterized by (M, S)*-freeness*. In this work, we study the following problem: *Given matroids M_1, M_2 and sets S_1, S_2, what is the relationship between the properties characterized by (M_1, S_1)-freeness and (M_2, S_2)-freeness?* Is there a set of combinatorial conditions that, given M_1, M_2, S_1, and S_2, allows us to decide whether the two corresponding properties are "very close" to or "very far" from each other? In general, this is a formidable question to answer, but we are able to resolve it in some natural (and interesting) restricted cases.

The main restriction that we impose on the properties is that we require they be characterized as (M, S)-freeness with $|S| = 1$. Note that if $S = S_1 \cup S_2$, we can express (M, S)-freeness as the intersection of (M, S_1)-freeness and (M, S_2)-freeness. Thus, in this sense the properties studied here are the building blocks of all locally characterized linear-invariant properties. To keep the notation concise, let us introduce some new nomenclature for these so-called *matroid freeness* properties.

Definition 1.2 (Matroid freeness). Given integers $k \geq r \geq 1$, a set $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ of k vectors in $\mathbb{F}^r$, and a string $\sigma \in \mathsf{R}^k$, we say that a function $f : \mathbb{F}^n \to \mathsf{R}$ is (M, σ)*-free* if there does not exist any linear map $L : \mathbb{F}^r \to \mathbb{F}^n$ such that $f(L(\mathbf{v}_i)) = \sigma_i$ for all $i \in [k]$. Otherwise, if such an L exists we say f *contains* (M, σ) *at* L.

We often refer to σ as the *label* and the pair (M, σ) as a *labeled matroid*. We write $\mathcal{M}_{\mathbb{F},\mathsf{R}}[\neg\sigma]$ to denote the subset of (M, σ)-free functions in $\bigcup_{n \in \mathbb{N}^+} \{\mathbb{F}^n \to \mathsf{R}\}$. Usually, however, the domain $\mathbb{F}^n$ and range R of the functions will be clear from context, and therefore we will drop the indices and just write $\mathcal{M}[\neg\sigma]$.

The testability of matroid freeness properties has been an area of intense research recently, and serves as an important source of motivation for our work on the semantics of these properties. We discuss this next.

[1]The formal definition of a matroid is not too important in this context. So although we provide a definition in Appendix A for completeness, in the rest of this paper the reader can just think of a matroid as a set of elements in a vector space over $\mathbb{F}$.

1.2 Testability of Matroid Freeness Properties

Given a property $\mathcal{P}$ of functions in $\{\mathbb{F}^n \to \mathsf{R}\}$, we say that $f : \mathbb{F}^n \to \mathsf{R}$ is δ-*far* from $\mathcal{P}$ if $\min_{g\in\mathcal{P}} \Pr_{\mathbf{x}\in\mathbb{F}^n}[f(\mathbf{x}) \neq g(\mathbf{x})] > \delta$ and δ-*close* otherwise. $\mathcal{P}$ is said to be *testable* (with one-sided error) if there is a function $q : (0,1) \to \mathbb{Z}^+$ and an algorithm T that, given as input a parameter $\delta \in (0,1)$ and oracle access to a function $f : \mathbb{F}^n \to \mathsf{R}$, makes at most $q(\delta)$ queries to the oracle for f, always accepts if $f \in \mathcal{P}$ and rejects with probability at least $2/3$ if f is δ-far from $\mathcal{P}$.

The study of matroid freeness properties is a fundamental part of the general program of characterizing the testable linear-invariant properties. It is believed that the fact that the properties are endowed with linear-invariance will allow for clean rules that determine testability. In fact, if we restrict ourselves to proximity-oblivious testable properties[2] as defined in [GR09], then it is conjectured that linear-invariant properties are testable if and only if they are locally characterized. Furthermore, it is known that a tester for a testable linear-invariant property can be assumed, without loss of generality, to be checking for (M, S)-freeness for some matroid M and some set S. Thus, our study of the relationship between local constraints and the properties they characterize may be viewed as a study of what property is actually being tested by a given test. Note that in this sense our work runs counter to the direction of most property testing research, where the primary concern is instead understanding what tests can be used for testing a given property. We refer to [GK10, Sud10] for more detailed discussions of invariance, testability and local characterizations.

In what follows, let us focus on $\mathbb{F} = \mathbb{F}_2$ and $\mathsf{R} = \{0,1\}$ since this is the setting in which testability of matroid freeness properties have been systematically investigated. This line of work was initiated by Green [Gre05], who showed that $\mathcal{M}[\neg\sigma]$ is testable for the matroid $M = \{\mathbf{e}_1, \ldots, \mathbf{e}_k, \sum_{i\in[k]} \mathbf{e}_i\}$ (where $\mathbf{e}_i$ denotes the i^{th} unit vector) and for σ being the all-ones string, henceforth denoted $\sigma = 1^*$. This was generalized in [BCSX09] to show that if M is any *graphic matroid* (see Section 2), then $\mathcal{M}[\neg 1^*]$ is testable. It is easy to check that such a property is monotone, i.e., if f is $(M, 1^*)$-free and f' is obtained from f by flipping $f(\mathbf{x})$ from 1 to 0 at any point $\mathbf{x} \in \mathbb{F}_2^n$, then f' is also $(M, 1^*)$-free. The papers [KSV10] and [Sha09] independently showed that the restriction to graphic matroids could be dropped, again assuming $\sigma = 1^*$. Very recently, [BGS10] proved that for arbitrary σ, but now reintroducing the requirement that M be graphic,[3] the non-monotone property $\mathcal{M}[\neg\sigma]$ (and any intersection of such properties) is testable.

This sequence of testability results for increasing sets of matroids M and labels σ seems to hold out the hope for a complete characterization of all testable linear-invariant properties in a not too distant future. However, there is a potentially very problematic gap in results such as those in [BGS10] that can make such hopes very misleading! The issue is that although these results establish testability of (M, σ)-freeness for an infinite syntactic collection of labeled matroids (M, σ), it is not at all clear that this provides an infinite collection of new, semantically distinct properties not previously known to be testable. Indeed, as was observed in the very first lines of this paper, it can certainly happen that the properties $\mathcal{M}^1[\neg\sigma^1]$ and $\mathcal{M}^2[\neg\sigma^2]$ are identical for two different labeled matroids (M^1, σ^1) and (M^2, σ^2). A little more subtly, it could also be the case that for distinct (M^1, σ^1) and (M^2, σ^2), any function which is (M^1, σ^1)-free is very close to being (M^2, σ^2)-free, so that even though the properties are not identical, any one-sided tester for (M^1, σ^1)-freeness can be easily modified to be a tester for (M^2, σ^2)-freeness. A third pitfall is that it could be the case for three distinct labeled matroids (M^1, σ^1), (M^2, σ^2) and (M^3, σ^3) that the property $\mathcal{M}^3[\neg\sigma^3]$ is equal to (or very close to) the union of the properties $\mathcal{M}^1[\neg\sigma^1]$ and $\mathcal{M}^2[\neg\sigma^2]$. In this case, testability of (M^3, σ^3)-freeness is trivially guaranteed by the testability of (M^1, σ^1)-freeness and (M^2, σ^2)-freeness. Thus, for the above cited result of [BGS10] to be nontrivial, one should ensure that the properties covered

[2] Proximity oblivious testable properties have testers which have q equal to a constant, independent of δ, but which reject with some nonzero probability for every $\delta > 0$.

[3] It should be noted that strictly speaking, the properties in [BGS10, KSV10, Sha09] are described in terms of forbidding solutions to systems of linear equations rather then satisfying local matroid constraints. These two formulations are essentially equivalent, however, as explained in Appendix A.

in that result are not close to the union of properties already previously known to be testable.

It should be stressed again that these concerns are far from hypothetical. As mentioned earlier, it was shown in [BCSX09] that if $M = \{\mathbf{e}_1, \ldots, \mathbf{e}_k, \sum_{i=1}^{k} \mathbf{e}_i\}$ for any k, then while there is an infinite hierarchy of distinct properties when $\sigma = 0^*$ or $\sigma = 1^*$, it turns out that $\mathcal{M}[\neg\sigma]$ when $\sigma \notin \{0^*, 1^*\}$ always degenerates to one of a finite set of properties that have all been known to be testable ever since [BLR93]. It is a natural question whether it could be the case more generally that non-monotone matroid freeness properties degenerate to one of a small set of already well-studied properties. This is posed as an open problem in [BCSX09], and resolving this question was an important motivation for this work.

1.3 Summary of Our Results

Very briefly, given two labeled matroids (M, σ) and (N, τ), we establish necessary and sufficient conditions for when the two properties $\mathcal{M}[\neg\sigma]$ and $\mathcal{N}[\neg\tau]$ are identical or distinct, provided that the constraints meet certain structural conditions. We then go on to show the existence of labeled matroids that satisfy these conditions. Finally, we use these results to rule out the aforementioned objections about testability results for matroid freeness properties by exhibiting infinite hierarchies of distinct non-monotone testable matroid freeness properties. We now describe these results in some more detail.

The main tool we use to show separations between matroid freeness properties is the notion of a *labeled matroid homomorphism*. Just as the notions of graph homomorphisms and its variants are helpful in counting occurrences of (induced) subgraphs inside graphs (see [AS06] for a survey), labeled matroid homomorphisms allow us to count the number of times a given labeled matroid is "contained" in a function. More precisely, we define a labeled matroid homomorphism ϕ from (M, σ) to (N, τ) to be a map ϕ that *(i)* is linear, *(ii)* maps vectors of M to vectors of N, and *(iii)* preserves labels in the sense that the σ-label of any vector $\mathbf{v}$ in M equals the τ-label of $\mathbf{w} = \phi(\mathbf{v})$ in N. Observe that since ϕ is linear, if some vectors of M are linearly dependent, then their images in N are also linearly dependent.

Dichotomy theorems. It is not hard to show that if there is a labeled matroid homomorphism from (M, σ) to (N, τ), then any (M, σ)-free function is also (N, τ)-free. It is reasonable to ask whether the fact that (M, σ) does *not* map homomorphically into (N, τ) can also provide some information about the relationship between the two properties of (M, σ)-freeness and (N, τ)-freeness. If we are very optimistically inclined, we might even inquire whether the existence or non-existence of homomorphisms *exactly* determines the relationship between these two properties in the sense that (M, σ)-freeness is far in Hamming distance from being contained in (N, τ)-freeness in this latter case. Somewhat surprisingly, it turns out that this is in fact true for monotone matroid freeness properties.

Theorem 1.3 (First main theorem (informal)). *For any two labeled matroids $(M, 1^*)$ and $(N, 1^*)$ it holds that either $\mathcal{M}[\neg 1^*]$ is contained in $\mathcal{N}[\neg 1^*]$ or $\mathcal{M}[\neg 1^*]$ is "well separated" from $\mathcal{N}[\neg 1^*]$ in the sense that there is a function that is $(M, 1^*)$-free but far from being $(N, 1^*)$-free. The first case applies if there exists a labeled matroid homomorphism from $(M, 1^*)$ to $(N, 1^*)$, and the second case applies otherwise.*

Notice that this implies a strong dichotomy: one of the two cases in Theorem 1.3 must hold, and it can never be the case that the two properties are distinct but close in a property testing sense. (See Theorem 3.6 for the formal statement of Theorem 1.3.) It should be noted that some limited results of the same flavor were proven for specific graphic matroids and monotone patterns in [BCSX09], and our techniques are inspired by that paper. However, our results are much stronger in that they apply to arbitrary (non-graphic) linear matroids and provide an exact criterion for when two labeled matroids with monotone labels are distinct.

An obvious next question is whether this dichotomy, and the characterization in terms of labeled matroid homomorphisms, hold not only for monotone properties but also for matroid freeness properties in general. We provide two answers to this question.

Our first answer is that we identify a fairly broad class of examples where the dichotomy in terms of homomorphisms applies also in the non-monotone case. We restrict ourselves in the theorem statement below to *graphic* matroid freeness properties, which, as mentioned previously, are of special interest since they are essentially the properties shown to be testable in [BGS10],[4] although our results hold in slightly larger generality and in particular extend even to properties not currently known to be testable.

Theorem 1.4 (Second main theorem (informal)). *If M and N are graphic matroids and σ and τ are non-monotone patterns that satisfy certain structural conditions, then it holds that either (M, σ)-freeness is contained in (N, τ)-freeness or the former property is "well separated" from the latter, and this is exactly determined by the existence or non-existence of a labeled matroid homomorphism from (M, σ) to (N, τ).*

We refer to Theorem 3.6 and Theorem 3.7 for formal instantiations of Theorem 1.4. Our focus in Theorem 1.4 is on matroids over complete graphs K_d. Our technical contribution lies in using the structure of the complete graph to argue that if (M, σ) does not embed homomorphically into (N, τ), where M, N are graphic submatroids of the complete graph and σ, τ have the needed properties as will be described below, then it is possible to pack into a function many copies of (N, τ), i.e., many violations of (N, τ)-freeness, while still keeping the function (M, σ)-free.

Our second answer is that the dichotomy as stated above does *not* hold in general for arbitrary non-monotone matroid freeness properties. Namely, we exhibit two labeled (graphic) matroids (M, σ) and (N, τ) such that (M, σ) does not embed into (N, τ) and yet $\mathcal{M}[\neg\sigma]$ is contained in $\mathcal{N}[\neg\tau]$. However, it can easily be shown that (M, σ) "almost" embeds into (N, τ) in the sense that there is mapping if we are also allowed to send vectors in M to the all-zero vector $\mathbf{0}$ (which is not a member of N), so our construction naturally leads to a more refined version of the dichotomy question that remains open.

The next main theorem tackles an issue brought up in the discussion about testability above. We would like to demonstrate that a given matroid freeness property is not close to the union of a collection of other matroid freeness properties. For instance, we might want to separate a property newly discovered to be testable from a union of properties already known to be so. We provide necessary and sufficient criteria for such a separation but again in a restricted setting. One interesting case in which the theorem below can be applied is when M is the complete graph matroid. (See Theorem 3.9 for the formal statement.)

Theorem 1.5 (Third main theorem (informal)). *Suppose that $\{(N^1, \tau^1), \ldots, (N^t, \tau^t)\}$ is a collection of labeled binary matroids and that $M \subseteq \mathbb{F}_2^k$ is a binary matroid containing $F_k^{\leq 2}$. Then either $\mathcal{M}[\neg 1^*]$ is contained in $\bigcup_{i=1}^t \mathcal{N}^i[\neg\tau^i]$ or $\mathcal{M}[\neg 1^*]$ is "well-separated" from this union. This is exactly determined by whether or not there exists a labeled matroid homomorphism from $(M, 1^*)$ to some (N^i, τ^i).*

Hierarchy Theorems. We apply these dichotomy results to rule out the concerns discussed earlier about degeneracy of matroid freeness properties.

Theorem 1.6 (Fourth main theorem (informal)). *There are infinite hierarchies of testable, non-monotone and monotone matroid freeness properties such that for each hierarchy, consecutive properties in the hierarchy are well separated. Furthermore, it is not the case that any one of the properties equals (or is close to) the union of some subset of the other properties, and the properties are not close to the well-studied class of low-degree polynomials.*

Thus, in particular, Theorem 1.6 allows us to conclude that the properties shown to be testable in [BGS10] do indeed constitute a new, infinitely large class of testable linear-invariant properties. We prove two slightly incomparable formalizations of the above theorem, and they yield two hierarchies with many

[4]To be precise, [BGS10] shows testability for a strictly larger class of so-called complexity-1 matroids, but we do not want to get into too much technicalities here.

"well separations" involving a large number of graphic matroid freeness properties. (The formal statements are described in Theorems 5.4 and 5.5.) The proof of these theorems turn out to be fairly involved, and there are a surprising number of intricacies involved in the relationship between the different properties. Due to the space constraints all proofs of the results discussed above, and even some of the formal definitions, have to be deferred to the full-length version of this paper (attached as an appendix).

1.4 Organization of This Paper

In Section 2, we provide some necessary background and elaborate on the motivation behind this work. Section 3 is devoted to establishing the dichotomy theorems. In Section 4, we prove the non-existence of labeled matroid homomorphisms, which can then be combined with the results from Section 3 to establish the infinite hierarchies of testable non-monotone properties in Section 5. We defer the proof of Theorem 5.5 to Appendix B. In Section 6 we describe a counterexample to the question of whether the non-existence of labeled homomorphisms determine containment between the respective properties, which leads to a new intriguing open problem as discussed above. We conclude in Section 7 by discussion some further questions left open by our work. Some background material, which might be useful although not necessary to understand the rest of the paper, is presented in Appendix A for completeness.

2 Preliminaries and Motivation

Let $\mathbb{N} = \{0, 1, \ldots\}$ denote the set of natural numbers and let $\mathbb{N}^+ = \mathbb{N} \setminus \{0\}$. Let $n \geq 1$ be a natural number. We write $[n]$ to denote the set $\{1, 2, \ldots, n\}$. We write $\mathbb{F}$ to denote a (finite) field.

2.1 Properties of Functions

Recall that a *property* $\mathcal{P}$ is a subset of functions from domain(s) D_n to range(s) R_n, that is $\mathcal{P} = \bigcup_{n \in \mathbb{N}^+} \mathcal{P}_n$ where $\mathcal{P}_n \subseteq \{\mathsf{D}_n \to \mathsf{R}_n\}$. It is customary to suppress n in this notation. Throughout this paper, we will have $\mathsf{D}_n = \mathbb{F}^n$ and $\mathsf{R}_n = \mathsf{R}$ for some fixed R, often $\mathsf{R} = \{0, 1\}$. Also, unless otherwise stated calligraphic letters $\mathcal{P}, \mathcal{Q}$, et cetera will denote properties. Recall also that the *(relative) distance* between functions f and g is the probability $\Pr_{x \in \mathsf{D}}[f(x) \neq g(x)]$ that they differ on some x drawn uniformly at random from D, and that f is δ-far from a propety $\mathcal{P}$ if it has distance at least δ to every function $g \in \mathcal{P}$. The following two definitions capture the notion of two properties being "well separated" from each other.

Definition 2.1 (δ-separated). For two properties $\mathcal{P}, \mathcal{Q} \subseteq \bigcup_{n \in \mathbb{N}^+} \{\mathbb{F}^n \to \mathsf{R}\}$, we say that $\mathcal{Q}$ is *δ-separated* from $\mathcal{P}$ if for infinitely many n there are functions $f_n : \mathbb{F}^n \to \mathsf{R}$ that are in $\mathcal{Q}$ but are δ-far from being in $\mathcal{P}$ (where $\delta > 0$ is fixed and in particular independent of n).

When the exact value of the parameter δ is not of importance (which is the case most of the time), we will simply say that $\mathcal{Q}$ is *well separated* from $\mathcal{P}$, and we say that $\mathcal{P}$ and $\mathcal{Q}$ are *mutually well separated* if there are separations in both directions. For brevity, we will sometimes use the notation $\mathcal{Q} \nsubseteq_\delta \mathcal{P}$ to signify that $\mathcal{Q}$ is δ-separated from $\mathcal{P}$ and $\mathcal{P} \neq_\delta \mathcal{Q}$ to signify that the properties are mutually well separated.

Definition 2.2 (δ-strictly contained). For two properties $\mathcal{P}, \mathcal{Q} \subseteq \bigcup_{n \in \mathbb{N}^+} \{\mathbb{F}^n \to \mathsf{R}\}$, we say that $\mathcal{P}$ is *δ-strictly contained*, or just *strictly contained*, in $\mathcal{Q}$ if $\mathcal{P} \subseteq \mathcal{Q}$ but $\mathcal{Q}$ is δ-separated from $\mathcal{P}$. We will use the notation $\mathcal{P} \subset_\delta \mathcal{Q}$ to signify that $\mathcal{P}$ is strictly contained in $\mathcal{Q}$.

2.2 Matroids and Labeled Homomorphisms

Turning next to matroids, for the purposes of this paper the reader can think of a *linear matroid* M as a set of vectors $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ in $\mathbb{F}^r$ for $r \leq k$. We will often write N to denote some other matroid $\{\mathbf{w}_1, \ldots, \mathbf{w}_\ell\}$

over $\mathbb{F}^s$ for $s \leq \ell$, and unless otherwise stated $\mathbf{v}_i$ is assumed to denote a vector in M and $\mathbf{w}_j$ to denote a vector in N.

We let $\mathbf{e}_1, \mathbf{e}_2, \ldots$ denote the unit vectors in the ambient space, i.e., $\mathbf{e}_i$ is 1 at coordinate i and 0 everywhere else. Sometimes when we need to distinguish basis vectors of different matroids we will also write $\mathbf{f}_1, \mathbf{f}_2, \ldots$ to denote unit vectors with respect to some other matroid. The *support set* of a vector $\mathbf{v}$, denoted by $\mathrm{supp}(\mathbf{v})$, is defined to the set of coordinates at which $\mathbf{v}$ is non-zero: $\mathrm{supp}(\mathbf{v}) = \{i \in [n] : v_i \neq 0\}$. The *weight* $|\mathbf{v}|$ of a vector $\mathbf{v}$ is the number of non-zero coordinates in $\mathbf{v}$, i.e., $|\mathbf{v}| = |\mathrm{supp}(\mathbf{v})|$.

We write $\mathbf{1}_r$ to denote the vector in $\mathbb{F}^r$ with all coordinates equal to one and $\mathbf{0}_r$ to denote the vector with zeros in all coordinates. In what follows below, the ambient vector space $\mathbb{F}^r$ will usually be clear from context, in which case we will drop the dimension subindex and write just $\mathbf{1}$ and $\mathbf{0}$. For any vector $\mathbf{v} \in \mathbb{F}_2^r$ we write $\overline{\mathbf{v}}$ to denote the vector $\mathbf{1}_r + \mathbf{v}$, i.e., the vector that is zero in some position i if and only if $v_i = 1$, and we will refer to $\overline{\mathbf{v}}$ as the *complement* of $\mathbf{v}$. Note that for every $\mathbf{v}$ we have $|\mathbf{v}| + |\overline{\mathbf{v}}| = r$. We will sometimes write $\mathbf{u} = \mathbf{v}|\mathbf{w}$ to denote that we decompose the vector $\mathbf{u}$ into the two vectors $\mathbf{v}$ and $\mathbf{w}$ such that $\mathbf{u}$ is the concatenation of $\mathbf{v}$ and $\mathbf{w}$.

We write $\sigma = \langle \sigma_1, \ldots, \sigma_k \rangle \in \mathsf{R}^k$ and $\tau = \langle \tau_1, \ldots, \tau_\ell \rangle \in \mathsf{R}^\ell$ to denote labels or patterns corresponding to the matroids M and N respectively. As mentioned in the introduction, we will refer to the pair (M, σ) as a *labeled matroid*, that is, a matroid $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ with each vector $\mathbf{v}_i$ labeled by $\sigma_i \in \mathsf{R}$ for $i = 1, \ldots, k$. We use fraktur letters $\mathfrak{M}$ and $\mathfrak{N}$ to denote sets of labeled matroids. We let f and g denote functions $\mathbb{F}^n \to \mathsf{R}$, where we will often, but not always, have $\mathbb{F} = \mathbb{F}_2 = \mathrm{GF}(2)$ and $\mathsf{R} = \{0, 1\}$. Note that k, ℓ, r, s are all fixed while we think of n as going to infinity.

In what follows, it will be convenient to have a compact notation for matroid freeness properties. Thus, we will write $\mathcal{M}_{\mathbb{F},\mathsf{R}}[\neg\sigma]$ to denote the subset of (M, σ)-free functions in $\bigcup_{n \in \mathbb{N}^+} \{\mathbb{F}^n \to \mathsf{R}\}$. Usually, however, the domain $\mathbb{F}^n$ and range R of the functions will be clear from context in this paper, and therefore we will drop the subindices and write just $\mathcal{M}[\neg\sigma]$. Notice that (M, σ) denotes a matroid M labeled by the pattern σ, whereas $\mathcal{M}[\neg\sigma]$ written with a calligraphic letter denotes the *property* of all functions f that are free from any occurrences of the labeled matroid (M, σ). The way to read this notation is that whenever we pick a set of vectors in $\mathbb{F}^n$ with linear dependencies as specified by M and evaluate f on these points, we do *not* see the pattern σ.

For $L : \mathbb{F}^r \to \mathbb{F}^n$ a linear transformation and $f : \mathbb{F}^n \to \mathsf{R}$ a function, we will sometimes use the notation $f(L(M))$ to denote the pattern obtained by evaluating f at $L(\mathbf{v}_1), L(\mathbf{v}_2), \ldots, L(\mathbf{v}_k)$. So, for instance $f(L(M)) = \sigma$ means that $f(L(\mathbf{v}_i)) = \sigma_i$ for all $i = 1, 2, \ldots, k$, or in other words that we do see the pattern σ when f is evaluated on $L(M)$.

A central notion of this work will be that of a *matroid homomorphism*.

Definition 2.3 (Matroid homomorphism [BCSX09]). Let $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ and $N = \{\mathbf{w}_1, \ldots, \mathbf{w}_\ell\}$ be two matroids with $M \subseteq \mathbb{F}^r$ and $N \subseteq \mathbb{F}^s$. A *matroid homomorphism* $\phi : M \to N$ is a $\mathbb{F}$-linear map[5] from $\mathbb{F}^r$ to $\mathbb{F}^s$ such that $\phi(\mathbf{v}_i) \in \{\mathbf{w}_1, \ldots, \mathbf{w}_\ell\}$ for every $1 \leq i \leq k$. We will also say that ϕ is an *embedding* of the matroid M into the matroid N, or that M *embeds* into N.

In contrast to [BCSX09], we want to be able to study not only monotone but also non-monotone matroid freeness properties, i.e., properties characterized by matroid constraints (M, σ) where we can have $\sigma \notin \{0^r, 1^r\}$. In order to do so, we need the following generalization of Definition 2.3 to arbitrary matroid constraints.

Definition 2.4 (Labeled matroid homomorphism). A *labeled matroid homomorphism* $\phi : (M, \sigma) \to (N, \tau)$ is a matroid homomorphism from M to N which in addition preserves labels in the sense that if $\phi(\mathbf{v}_i) = \mathbf{w}_j$, then $\sigma_i = \tau_j$. If there exists a labeled homomorphism from (M, σ) to (N, τ), we say that (M, σ) *embeds* into (N, τ) and write $(M, \sigma) \hookrightarrow (N, \tau)$; otherwise, we write $(M, \sigma) \not\hookrightarrow (N, \tau)$.

[5] $\phi : \mathbb{F}^r \to \mathbb{F}^s$ is $\mathbb{F}$-linear if for any $a, b \in \mathbb{F}$, $x, y \in \mathbb{F}^r$ it holds that $\phi(ax + by) = a\phi(x) + b\phi(y)$

We say that M is a *submatroid* of N if there is an injective matroid homomorphism mapping M into N. If in addition this homomorphism is surjective we say that M and N are *isomorphic.* In what follows, we will usually implicitly consider all isomorphic matroids to be the same. The concepts of submatroids and isomorphic matroids extend to labeled matroids in the natural way by considering matroid homomorphisms that are also label-preserving.

A matroid $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ is said to be *graphic* if there exists a graph G with k edges for which these edges can be associated with the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ in M in such away that any subset of vectors $S \subseteq \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ is linearly dependent if and only if the associated set of edges contains a cycle. In this case, we denote M by $M(G)$. (Also notice that we write $\mathbf{v}$ and $\mathbf{e}$ for vectors to distinguish them from vertices v and edges e in graphs.) We require the graph G to be simple; that is, G has no self-loops or parallel edges. It is a well-known fact that graphic matroids always can be represented as *binary matroids*, i.e., linear matroids over $\mathbb{F}_2$.

Let us now fix $\mathsf{R} = \{0, 1\}$ for the rest of this section. We can visualize a graphic matroid constraint $(M(G), \sigma)$ as the graph G with $0/1$-labels σ_i on its edges. In what follows, we will sometimes identify $(M(G), \sigma)$ with this *labeled graph* (G, σ). We say that (G, σ) is a *labeled subgraph* of (H, τ) if G is a subgraph of H such that the edge labels of the common edges coincide. For graphic matroid constraints $(M(G), \sigma)$ and $(M(H), \tau)$, a labeled matroid homomorphism is simply a mapping of edges in G to edges with the same labels in H such that cycles in G map to cycles in H. Clearly, if (G, σ) is a labeled subgraph (H, τ), then the embedding of the vertices of G in H induces a matroid homomorphism in the natural way. These are not the only labeled graphic homomorphisms, however. In particular, a matroid homomorphism need not map edges incident to the same vertex in G to incident edges in H.

2.3 Matroid representations

As long as we are only studying monotone patterns, we need not worry too much about exactly how our linear matroids are represented. When we want to discuss (M, σ)-freeness for a non-monotone pattern σ, however, we have to specify exactly how M is represented in order to make sure that the matroid constraint is well-defined.

Definition 2.5 (Standard binary matroid representation). We say that a binary matroid $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ is in *standard representation* if there is a $d < k$ such that $\mathbf{v}_i = \mathbf{e}_i$ for $i = 1, \ldots, d$ and the rest of the vectors are linear combinations of these vectors, i.e., for all $j > d$ we have $\mathbf{v}_j = \sum_{i \in I_j} \mathbf{e}_i$ for index sets $I_j \subseteq [d]$, where $\mathbf{v}_j$, $j > d$, are enumerated in lexicographical order with respect to I_j.

It is easy to see that any matroid has a standard linear representation, namely by associating unit vectors to a basis of the matroid and then representing the other elements by appropriate linear combinations of the basis elements. When we talk about matroid constraints from now on, it will always be for linear matroids in standard representation.

We remark that the standard representation does not uniquely specify matroid constraints (M, σ) — for a fixed M in standard representation there can be several distinct patterns $\sigma^1, \sigma^2, \ldots$ such that (M, σ^i) all represent the same labeled matroid. For instance, for the complete graph K_d on d vertices, all labeled matroids $\left(M(K_d), 1^{d_1}01^{d_2}\right)$ are easily verified to be the same for all $d_1 + d_2 = \binom{d}{2} - 1$. However, it is not the case that for all $d_1 + d_2 + d_3 = \binom{d}{2} - 2$ the labeled matroids $\left(M(K_d), 1^{d_1}01^{d_2}01^{d_3}\right)$ are all the same. On the contrary, one can prove (although we do not do so here) that one gets two different cases depending on whether the two 0-labeled edges are incident to a common vertex or not. The point is that the standard representation of M gives us at least *one* well-defined description of the labeled matroid, whereas the example just discussed shows that without such a representation a non-monotone labeled matroid (M, σ) could be ambiguous.

In this work, we will focus on a particular class of matroid freeness properties, the understanding of which is arguably fundamental to the broader study of the space of testable linear-invariant properties. Before explaining what we mean by this we introduce some useful definitions.

Definition 2.6 (Full and partial matroids). The *full (linear binary) matroid* F_d of dimension d is $F_d = \{\sum_{i \in I} \mathbf{e}_i \big| \emptyset \neq I \subseteq [d]\}$. The *partial matroid of weight* w is the matroid $F_d^{\leq w} = \{\sum_{i \in I} \mathbf{e}_i \big| \emptyset \neq I \subseteq [d],\ |I| \leq w\}$.

It was shown in [BGS10] that any linear-invariant property in $\mathbb{F}_2^n \to \{0,1\}$ that is testable with a one-sided tester can be written as an $\mathfrak{M}$-freeness property for a possible infinite collection $\mathfrak{M}$ of matroid constraints.[6] In other words, any such property can be characterized as the intersection of a possibly infinite number of (M,σ)-freeness properties. Moreover, each (M,σ)-freeness property can be written as the intersection of a finite collection of (F_d,σ')-freeness properties, where F_d is the *full linear matroid* defined above.

Namely, suppose that the matroid M lives in $\mathbb{F}_2^d$ and let us for simplicity assume that it consists of the first $k \leq 2^d - 1$ vectors in the standard representation. Then it is straightforward to verify that (M,σ)-freeness is exactly the intersection of all $(F_d, \sigma\tau)$-freeness properties, where τ ranges over all patterns in $\{0,1\}^{2^d-(k+1)}$ and $\sigma\tau$ denotes concatenation. (This corresponds to that a violation of (M,σ)-freeness occurs as soon as the first k vectors in F_d are mapped to points in $\mathbb{F}^n$ evaluating to the pattern $\sigma \in \{0,1\}^k$, regardless of what the evaluation pattern looks like for the rest of the points.) As another example of the expressive power of matroid constraints, note that low-degree polynomials (with constant term zero), can be specified as the intersection of all (F_d,τ)-freeness properties, where d is fixed and σ ranges over all patterns in $\{0,1\}^{2^d-1}$ of odd parity.

It follows from the preceding paragraph that all linear-invariant properties in $\mathbb{F}_2^n \to \{0,1\}$ that are one-sided-testable are collections of full linear matroid freeness properties, so matroid constraints (F_d,σ) can be seen to be the building blocks of all one-sided-testable linear-invariant properties. In the other direction, [BGS10] established that any collection $\mathfrak{M}$ of *graphic* matroid freeness properties is testable by a one-sided tester. It is again not hard to see that for any graphic matroid contraint $(M(G),\sigma)$, where G is a graph on d vertices, one can represent $(M(G),\sigma)$-freeness as the intersection of all $(M(K_d),\sigma\tau)$-freeness properties where σ labels the edges of G and τ ranges over all possible labels on the edges in K_d not present in G.

Thus, to gain an understanding of the semantics of matroid freeness properties in general, and of the testability results in [BGS10] in particular, a necessary first step is to comprehend graphic matroid constraints $(M(K_d),\sigma)$. Furthermore, even this first step appears to be a challenging problem in its own right. Therefore, in what follows we will mostly focus on matroids over complete graphs. For the rest of this paper, we fix the representation of such matroids as follows.

Definition 2.7 (Standard representation of complete graph matroids). We choose the $d-1$ independent basis vectors of $M(K_d)$ to be the $d-1$ edges incident to some (arbitrarily chosen but) fixed vertex. The $\binom{d}{2}$ vectors in $M(K_d)$ will then consist of all the $d-1$ weight-1 vectors and all the $\binom{d-1}{2}$ weight-2 vectors in $\mathbb{F}_2^{d-1}$. Moreover, these $\binom{d}{2}$ vectors are always ordered lexicographically as

$$M(K_d) = \{\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_{d-1}, \mathbf{e}_1+\mathbf{e}_2, \mathbf{e}_1+\mathbf{e}_3, \ldots, \mathbf{e}_1+\mathbf{e}_{d-1}, \mathbf{e}_2+\mathbf{e}_3, \mathbf{e}_2+\mathbf{e}_4, \ldots, \mathbf{e}_{d-2}+\mathbf{e}_{d-1}\} \ .$$

Notice that $M(K_d)$ is an example of a partial matroid as defined above, namely $M(K_d) = F_{d-1}^{\leq 2}$.

See Figure 1 for an illustration of the standard matroid representation of $M(K_5)$. We make the observation, which will be used later, that for any fixed vector $\mathbf{v} \in M(K_d)$ we can find a standard matroid basis that that contains $\mathbf{v}$ and is on the form in Definition 2.7, i.e., where every non-basis matroid vectors is a sum

[6] Modulo a standard technical assumption that is not relevant to this discussion — we refer to [BGS10] for the precise statement.

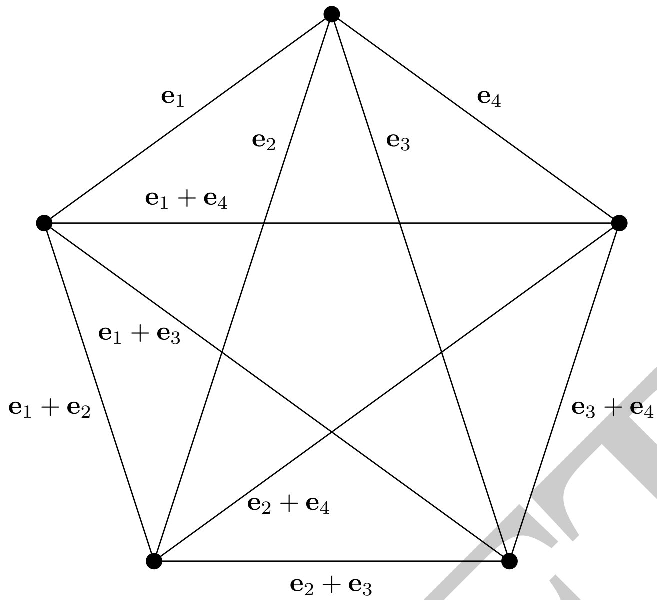

Figure 1: The standard matroid representation for $M(K_5)$.

of two basis vectors. For the weight-1 vectors this is by definition, but it also holds for weight-2 vectors by symmetry. Namely, if we fix a vector $\mathbf{e}_i + \mathbf{e}_j$ for $i < j$, it is immediate to verify that the set of vectors $\{\mathbf{e}_i, \mathbf{e}_i + \mathbf{e}_1, \dots \mathbf{e}_i + \mathbf{e}_{i-1}, \mathbf{e}_i + \mathbf{e}_{i+1} \dots \mathbf{e}_i + \mathbf{e}_j, \dots \mathbf{e}_i + \mathbf{e}_{d-1}\}$ and their pairwise sums generate $M(K_d)$. Another easy way of seeing this is perhaps to take a look at Figure 1 and make a "proof by picture."

As a final notational convention, we make explicit our use of wildcards in patterns, allowing the use of * when the meaning is clear from context. For instance, we will write 1^* to denote the all-ones pattern, and 0^d1^* denotes the pattern with ones everywhere except in the first d positions (relative to some fixed representation of the matroid in question).

3 A Method for Proving Distinctness of Matroid Freeness Properties

In this section, we develop a method to determine the relations between matroid freeness properties by way of labeled matroid homomorphisms. In brief, we establish a connection between the existence of an embedding from (M, σ) to (N, τ) on the one hand and the distance between the property of being (M, σ)-free and the property of being (N, τ) -free on the other. Namely, we show that if there exists a labeled homomorphism between (M, σ) and (N, τ) then (N, τ)-freeness contains (M, σ)-freeness, and that otherwise, at least in some specific cases which we characterize, these properties are well separated in the sense of Definition 2.1. We find this quite surprising, since it is not at all clear a priori how to amplify non-existence of embeddings into statistical distance between the associated properties. We note that it remains an intriguing open problem in what generality this connection holds for any pair of binary matroids (M, σ) and (N, τ).

Let us start with the lemma establishing the easy direction of the connection between labeled matroid homomorphisms and containment of matroid freeness properties. Recall that $\mathcal{M}[\neg\sigma]$ and $\mathcal{N}[\neg\tau]$ denote the properties consisting of all functions $\bigcup_{n\in\mathbb{N}^+}\{\mathbb{F}^n \to \mathsf{R}\}$ that are free from occurrences of (M, σ) and (N, τ), respectively. Furthermore, let us write $\mathcal{K}_d[\neg\tau]$ to denote the property of being free from $(K_d, \tau) = (M(K_d), \tau)$ and $\mathcal{F}_d^{\leq w}[\neg\tau]$ the property of being free from $(F_d^{\leq w}, \tau)$. Recall also that $\mathcal{M}[\neg\sigma] \nsubseteq_\delta \mathcal{N}[\neg\tau]$ denotes that $\mathcal{M}[\neg\sigma]$ is well separated from $\mathcal{N}[\neg\tau]$ and $\mathcal{M}[\neg\sigma] \neq_\delta \mathcal{N}[\neg\tau]$ denotes that the properties are mutually well separated.

Lemma 3.1. *If M and N are any linear matroids such that there is a labeled matroid homomorphism from (M, σ) to (N, τ), then $\mathcal{M}[\neg\sigma] \subseteq \mathcal{N}[\neg\tau]$*

Before proving the lemma, we state a simple corollary that will be useful later on.

Corollary 3.2. *If (G, σ) is a labeled subgraph of (H, τ) and $f : \mathbb{F}_2^n \to \mathsf{R}$ is a $(M(G), \sigma)$-free function, then f is also $(M(H), \tau)$-free.*

Proof of Corollary 3.2. If (G, σ) is a labeled subgraph of (H, τ), then in particular there is a labeled matroid homomorphism from $(M(G), \sigma)$ to $(M(H), \tau)$, namely the one induced by the embedding of the vertices of G in H. The claim now follows by Lemma 3.1. □

Proof of Lemma 3.1. Let $\phi : M \to N$ be a labeled matroid homomorphism. Suppose that $f : \mathbb{F}^n \to \mathsf{R}$ is not (N, τ)-free and in particular that f contains (N, τ) at a linear map $L : N \to \mathbb{F}^n$. We claim that this implies that f contains (M, σ) at the linear map $L \circ \phi : M \to \mathbb{F}^n$. By assumption, for all $j \in [m]$ we have $f(L(\mathbf{w}_j)) = \tau_j$. Suppose $\phi(\mathbf{v}_i) = \mathbf{w}_{j_i}$. Then by definition $\sigma_i = \tau_{j_i}$ since ϕ preserves labels, and for all $i \in [k]$ it clearly holds that $f((L \circ \phi)(\mathbf{v}_i)) = f(L(\phi(\mathbf{v}_i))) = f(L(\mathbf{w}_{i_j}))) = \tau_{j_i} = \sigma_i$, establishing the claim. □

Lemma 3.1 provides a method of arguing that some syntactically different properties are in fact identical. Consider for example $\mathcal{K}_3[\neg 1^*]$ and $\mathcal{K}_4[\neg 1^*]$. Since $(K_3, 1^*)$ is a labeled subgraph of $(K_4, 1^*)$, clearly $\mathcal{K}_3[\neg 1^*] \subseteq \mathcal{K}_4[\neg 1^*]$. Perhaps somewhat counter-intuitively, we can also show that there is a labeled matroid homomorphism from $(K_d, 1^*)$ to $(K_3, 1^*)$ and hence the containment holds in the other direction as well. To see this, write K_3 in standard representation over $\mathbf{e}_1, \mathbf{e}_2$ and K_4 in standard representation over $\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3$. Define $\phi : K_4 \to K_3$ by $\phi(\mathbf{f}_1) = \mathbf{e}_1$, $\phi(\mathbf{f}_2) = \mathbf{e}_2$, and $\phi(\mathbf{f}_3) = \mathbf{e}_1 + \mathbf{e}_2$ and extend it to all of K_4 by linearity. We leave it to the reader to verify that $\phi(\mathbf{f}_i + \mathbf{f}_j) \in K_3$ for all $1 \leq i < j \leq 3$. Since ϕ is trivially label preserving when all vectors are 1-labeled, it follows that ϕ is a labeled matroid homomorphism. We write this down as a proposition for reference.

Proposition 3.3. *The labeled matroid $(M(K_4), 1^*)$ embeds into $(M(K_3), 1^*)$, so $\mathcal{K}_3[\neg 1^*] = \mathcal{K}_4[\neg 1^*]$.*

To provide some more intuition, we give another hopefully instructive example, this time of non-identical properties. Observe that any function which is $(M(K_3), 011)$-free is also $(M(K_4), 011111)$-free by Corollary 3.2, since $(K_3, 011)$ is clearly a labeled subgraph of $(K_4, 011111)$. Also, it is not too hard to show that $\mathcal{K}_3[\neg 011]$ and $\mathcal{K}_4[\neg 011111]$ are not *exactly* the same. To see this, fix any $\mathbf{y} \in \mathbb{F}_2^n \setminus \{\mathbf{0}\}$ and consider the function $f_{\mathbf{y}} : \mathbb{F}_2^n \to \{0, 1\}$ defined by $f_{\mathbf{y}}(\mathbf{x}) = 1$ if $\mathbf{x} = \mathbf{y}$ and $f_{\mathbf{y}}(\mathbf{x}) = 0$ otherwise. We want to argue that $f_{\mathbf{y}}$ is $(M(K_4), 011111)$-free but not $(M(K_3), 011)$-free. Let again $M(K_3)$ be represented over unit vectors $\mathbf{e}_1, \mathbf{e}_2$ and $M(K_4)$ over unit vectors $\mathbf{f}_1, \mathbf{f}_2, \mathbf{f}_3$. The linear map $L_1 : M(K_3) \to \mathbb{F}_2^n$ sending $\mathbf{e}_1$ to $\mathbf{0}$ and $\mathbf{e}_2$ to $\mathbf{y}$ gives a pattern $\langle 011 \rangle$ for $M(K_3)$. Now suppose that $f_{\mathbf{y}}$ would contain $(M(K_4), 011111)$ at some linear map $L_2 : M(K_4) \to \mathbb{F}_2^n$. Then in particular $f_{\mathbf{y}}(L_2(\mathbf{f}_2)) = f_{\mathbf{y}}(L_2(\mathbf{f}_3)) = 1$, so we must have $L_2(\mathbf{f}_2) = L_2(\mathbf{f}_3) = \mathbf{y}$. But then $f_{\mathbf{y}}(L_2(\mathbf{f}_2 + \mathbf{f}_3)) = f_{\mathbf{y}}(L_2(\mathbf{f}_2) + L_2(\mathbf{f}_3)) = f_{\mathbf{y}}(\mathbf{0}) = 0 \neq 1$. Contradiction. Hence $f_{\mathbf{y}}$ is $(M(K_4), 011111)$-free.

Note that the above argument shows that $\mathcal{K}_3[\neg 011]$ and $\mathcal{K}_4[\neg 011111]$ are distinct, but does not rule out that the two properties are "essentially" the same in that they are very close in Hamming distance. However, it follows as a corollary of results that we will prove later in this paper (see Lemma 5.2) that not only does $\mathcal{K}_4[\neg 011111]$ contain $\mathcal{K}_3[\neg 011]$, but this containment is strict in the sense of Definition 2.2.

We next consider the other, harder, direction in the correspondence between labeled matroid homomorphisms and matroid freeness properties.

3.1 The High-Level Idea

We want to reduce the problem of separating (M, σ)-freeness from (N, τ)-freeness to a question about matroid homomorphisms. Our goal is to prove a statement of the following kind:

> *Suppose that (M, σ) and (N, τ) are matroid constraints such that there is no labeled matroid homomorphism from (M, σ) to (N, τ). Then the property of (M, σ)-freeness is δ-separated from that of (N, τ)-freeness.*

To prove such a statement, we need to exhibit an infinite family of functions $f_n : \mathbb{F}^n \to \mathsf{R}$ for $n \to \infty$ that is (M, σ)-free but far from (N, τ)-free. The general outline of the argument is as follows:

1. First, we define a "canonical function" $f_{(N,\tau)} : \mathbb{F}^n \to \mathsf{R}$ that encodes the structure of the labeled matroid (N, τ). More precisely, suppose that $N = \{\mathbf{w}_1, \ldots, \mathbf{w}_\ell\} \subseteq \mathbb{F}^s$. Then $f_{(N,\tau)}$ is constructed by splitting $\mathbf{x} \in \mathbb{F}^n$ into $\mathbf{y}|\mathbf{z}$ for $\mathbf{y} \in \mathbb{F}^s$ and $\mathbf{z} \in \mathbb{F}^{n-s}$ and letting $f_{(N,\tau)}(\mathbf{x})$ be (in some sense) the indicator function for whether $\mathbf{y} \in \mathbb{F}^s$ is a vector in N and if so what label it has.

2. Then, we prove that $f_{(N,\tau)}$ is dense in instances of the matroid pattern (N, τ) and has to be changed in many positions to become (N, τ)-free.

3. Finally, we assume that $f_{(N,\tau)}$ contains an instance of the matroid (M, σ) as witnessed by the linear transformation $L : M \to \mathbb{F}^n$. Then we want to argue that composing L with the projection $\pi : \mathbb{F}^n \to \mathbb{F}^s$ that maps $\mathbf{x} = \mathbf{y}|\mathbf{z}$ to $\mathbf{y}$, we obtain a labeled matroid homomorphism $\pi \circ L$ from (M, σ) to (N, τ). But this contradicts the assumption that there is no such homomorphism.

To construct a function that is dense in a pattern is relatively straightforward, and we achieve this (essentially) by a padding argument. The hard part is the third and final step in the argument. Notice that $\pi \circ L$ is linear by construction, but we are only guaranteed that it maps M to $\mathbb{F}^s$, not into N. In general, most vectors in $\mathbb{F}^s$ are *not* in N, so we somehow have to make sure that we land only in this subset of vectors. Furthermore, if it holds that $(\pi \circ L)(\mathbf{v}_i) = \mathbf{w}_j$, then we must make sure that the labels σ_i and τ_j agree. We remark that in fact, it is not at all clear what it actually means that $f_{(N,\tau)}$ should be an "indicator function" for (N, τ). The function $f_{(N,\tau)}$ has to map all of $\mathbb{F}^n$ to R, and in general for each value $\tau_j \in \mathsf{R}$ there will be some vector $\mathbf{w}_j \in N$ such that τ_j is the correct value. However, all the (majority of) vectors $\mathbf{x} \in \mathbb{F}^n$ that do *not* correspond to vectors in N also have to map somewhere in R, and we need to detect that when such a vector maps to τ_i, this does not indicate that $\mathbf{x}$ is a vector in N labeled by τ_i. This is the tricky part, and indeed we do not know how to accomplish this for completely general labeled linear matroids (M, σ) and (N, τ). However, by imposing some structural restrictions on our matroids, we can still derive theorems of the same type that yield strong results when applied in the right way.

3.2 Canonical Functions for Labeled Matroids

We now formalize the intuition from the previous section and generalize the definition of a canonical function for one matroid to a canonical function that encodes a collection of matroids.

Definition 3.4 (Matroid set canonical function). Suppose that $\mathfrak{N} = \{(N^1, \tau^1), (N^2, \tau^2), \ldots, (N^t, \tau^t)\}$ is a set of labeled matroids for $N^i = \{\mathbf{w}_1^i, \ldots, \mathbf{w}_{\ell_i}^i\} \subseteq \mathbb{F}^{s_i}$ and $\tau^i = \langle \tau_1^i \cdots \tau_{\ell_i}^i \rangle$ for $\tau_j^i \in \mathsf{R}$. Suppose that $n > s = \sum_{i=1}^t s_i$ and write the vector $\mathbf{x} \in \mathbb{F}^n$ as $\mathbf{x} = \mathbf{y}_1|\mathbf{y}_2|\ldots|\mathbf{y}_t|\mathbf{z}$ for $\mathbf{y}_i \in \mathbb{F}^{s_i}$ and $\mathbf{z} \in \mathbb{F}^{n-s}$. Let $b \in \mathsf{R}$ and let $S_\mathbf{0}$ and $S_{i,\mathbf{w}}$ be some sets in $\mathbb{F}^{n-s}$ for all $i \in [t]$ and $\mathbf{w} \in N^i$. Then the *matroid set canonical* *function* $f^b_{\mathfrak{N}} : \mathbb{F}^n \to \mathsf{R}$ for $\mathfrak{N}$ (with respect to b and the sets $S_{\mathbf{0}}, S_{i,\mathbf{w}}$) is defined by

$$f^b_{\mathfrak{N}}(\mathbf{x}) = f^b_{\mathfrak{N}}(\mathbf{y}_1|\mathbf{y}_2|\dots|\mathbf{y}_t|\mathbf{z}) = \begin{cases} 0 & \text{if } \mathbf{y}_i = \mathbf{0} \text{ for all } i \in [t] \text{ and } \mathbf{z} \in S_{\mathbf{0}}; \\ \tau^i_j & \text{if } \mathbf{y}_i = \mathbf{w}^i_j \in N^i, \, \mathbf{z} \in S_{i,\mathbf{w}_j}, \text{ and } \mathbf{y}_{i'} = \mathbf{0} \text{ for all } i' \neq i; \\ b & \text{otherwise.} \end{cases}$$

If the set $\mathfrak{N} = \{(N, \tau)\}$ consists of just a single matroid (N, τ), we will write $f^b_{(N,\tau)}$ to denote a corresponding canonical function.

Note that $b \in \mathsf{R}$ denotes a "padding value." Loosely speaking, our canonical function evaluates to b on points that do not correspond to vectors in matroids N^i.

The function $f^b_{\mathfrak{N}}$ encodes $\mathfrak{N}$ in the sense that it is dense in the matroid pattern (N^i, τ^i) for all $i \in [t]$, as well as in any (M, σ) that maps homomorphically into (N^i, τ^i).

Lemma 3.5. *Let $f^b_{\mathfrak{N}} : \mathbb{F}^n \to \mathsf{R}$ be a canonical function for $\mathfrak{N} = \{(N^1, \tau^1), (N^2, \tau^2), \dots, (N^t, \tau^t)\}$ as in Definition 3.4 such that all sets $S_{i,\mathbf{w}}$ are identical and form a linear subspace of $\mathbb{F}^{n-s}$ of size $|\mathbb{F}|^{n-s-\mathrm{O}(1)}$. If $(M, \sigma) \hookrightarrow (N^l, \tau^l)$ for some $l \in [t]$ then $f^b_{\mathfrak{N}}$ is δ-far from being (M, σ)-free, where δ is some universal constant independent of n. Furthermore, if $S_0 = \mathbb{F}^{n-s}$ and $S_{i,\mathbf{w}}$ are instead randomly and independently chosen as half of the elements in $\mathbb{F}^{n-s}$, then $f^b_{\mathfrak{N}}$ is δ-far from being (M, σ)-free with high probility.*

Proof. We proceed proving the first part of the lemma. Assume w.l.o.g. $(M, \sigma) \hookrightarrow (N^1, \tau^1)$. Let $\phi : (M, \sigma) \to (N^1, \tau^1)$ be a homomorphism. For simplicity of notation let $\ell = \ell_1$, $k = s_1$ and $S = S_{1,\mathbf{w}}$. Suppose that $N^1 = \{\mathbf{w}_1, \dots, \mathbf{w}_\ell\}$ for $\mathbf{w}_j \in \mathbb{F}^k$, and let $L' : \mathbb{F}^k \to \mathbb{F}^{n-s+k}$ be any linear transformation sending $\mathbf{w}_j$ to $\mathbf{w_j}|\mathbf{z_j}$ for arbitrary $\mathbf{z}_j \in S \subset \mathbb{F}^{n-s}$. Let $L : \mathbb{F}^k \to \mathbb{F}^n$ be the linear transformation extending L' trivially by mapping $\mathbf{w}_j$ to $\mathbf{w}_j|\mathbf{0}|\mathbf{z}_j$ with $\mathbf{0} \in \mathbb{F}^{s-k}$.

Then for all $\mathbf{v}_i \in M$, if $\phi(\mathbf{v}_i) = \mathbf{w}_j$ it is easy to verify that $f^b_{\mathfrak{N}}\big((L \circ \phi)(\mathbf{v}_i)\big) = f^b_{\mathfrak{N}}\big((L(\mathbf{w}_j)\big) = \tau_j = \sigma_i$, where the last equality holds since ϕ preserves labels. Hence, f^b_N contains (M, σ) at $L \circ \phi$.

The proof of δ-farness closely follows a similar argument in [BCSX09]. Set $q = |\mathbb{F}|$. Suppose $|S| \leq q^{n-s-c}$ (for some $c = \mathrm{O}(1)$). Let g be a function such that $\mathrm{dist}(f^b_{\mathfrak{N}}, g) = \delta < q^{-s-k-c}$. Let f_1 and g_1 be the restrictions of $f^b_{\mathfrak{N}}$ and g respectively, to the subspace $\mathbb{F}^k \times \mathbb{F}^{n-s}$ (corresponding to the projection of $\mathbb{F}^n$ to the matroid N^1 and the padding subspace). We will show that g_1 contains (M, σ) at a linear transformation $L_1 : M \to \mathbb{F}^k \times \mathbb{F}^{n-s}$, and then extend L_1 as before to a linear transformation $L : \mathbb{F}^k \to \mathbb{F}^n$ by padding it with zeros on the subspaces corresponding to matroids N^i, $i \geq 2$. This will imply that g contains (M, σ) at L and conclude that if g is δ-close then there is a matroid homomorphism between (M, σ) and (N^1, τ^1).

Define $\delta_1 = \mathrm{dist}(f_1, g_1)$ and notice that $\delta_1 \leq q^{s-k}\delta$. Now for each $\mathbf{w} \in N^1$, since $|S| \leq q^{n-s-c}$ it follows that $\Pr_{\mathbf{y} \in \mathbb{F}^k, \mathbf{z} \in S}[f_1(\mathbf{y}|\mathbf{z}) \neq g_1(\mathbf{y}|\mathbf{z})] \leq q^c \delta_1$. For $\mathbf{w}_i \in N^1$ let $\delta_i = \Pr_{\mathbf{z} \in S}[f_1(\mathbf{w_i}|\mathbf{z}) \neq g_1(\mathbf{w_i}|\mathbf{z})]$. Notice that

$$\delta_i \leq q^k \Pr_{\mathbf{y} \in \mathbb{F}^k, \mathbf{z} \in S}[f_1(\mathbf{y}|\mathbf{z}) \neq g_1(\mathbf{y}|\mathbf{z})] \leq q^{c+k}\delta_1.$$

Hence $\sum \delta_i \leq q^{2k+c}\delta_1 < 1$.

Now consider a random linear map $\tilde{L}_1 : \mathbb{F}^k \to S$, and its extension $\tilde{L} : \mathbb{F}^k \to \mathbb{F}^{n-s+k}$ given by $\tilde{L}(\mathbf{y}) = \mathbf{y}|\mathbf{\tilde{L}_1}(\mathbf{y})$. For every non-zero $\mathbf{y}$ and in particular for $\mathbf{y} \in N^1$, we have that $\tilde{L}_1(\mathbf{y})$ is distributed uniformly over the subspace S. Thus, for any fixed $i \in [\ell]$, we have

$$\Pr_{\tilde{L}_1}[g_1(\tilde{L}(\mathbf{w}_i)) \neq \tau_i] = \Pr_{\tilde{L}_1}[g_1(\tilde{L}(\mathbf{w}_i)) \neq f_1(\tilde{L}(\mathbf{w}_i))] \leq \delta_i \quad . \tag{3.1}$$

By the union bound, we get that

$$\Pr_{\tilde{L}_1}[\exists\, i \text{ such that } g(\tilde{L}(\mathbf{w}_i)) \neq \tau_i] \leq \sum_i \delta_i < 1 \quad . \tag{3.2}$$

In other words, there exists a linear map $\tilde{L}_1$ (and thus $\tilde{L}$) such that for every i, $g(\tilde{L}(\mathbf{w}_i)) = \tau_i$ and so g contains (N, τ) at $\tilde{L}$ and hence (M, σ) at the linear map $L' = \tilde{L} \circ \phi$.

To prove the second part of the lemma, we note that one can carry out the analysis of the last paragraph in the above proof with probability over $\tilde{L}_1$ and all $S_{\mathbf{w}}$ (instead of just $\tilde{L}_1$) and then apply a version of the Markov inequality. It is also easy to see that if τ_j is distinct from the padding value b, then different choices of $S_{\mathbf{w}_j}$ yield distinct functions. In this way, we can obtain a very large family of canonical functions that are (N, τ)-free but far from being (M, σ)-free. □

3.3 Three Dichotomy Theorems

In Section 3.2, we carried out the first two steps in the proof outline in Section 3.1. We now present two classes of pairs of labeled matroids (M, σ) and (N, τ) for which we can also successfully complete the crucial third step, and thus establish a dichotomy in the sense that if containment does not hold between the two properties, then they must be well separated.

Theorem 3.6 (First dichotomy theorem). *Let M, N be any linear matroids and let τ be any pattern for N. Then $\mathcal{M}[\neg 1^*] \subseteq \mathcal{N}[\neg\tau]$ if and only if there is a labeled matroid homomorphism from $(M, 1^*)$ to (N, τ); otherwise well-separation $\mathcal{M}[\neg 1^*] \nsubseteq_\delta \mathcal{N}[\neg\tau]$ holds.*

Before proving this theorem, we note that as a corollary it immediately yields our first main result in Theorem 1.3 on page 5 providing a full characterization of monotone matroid freeness properties. This follows simply by setting $\tau = 1^*$.

Proof of Theorem 3.6. The "if" part of the claim is Lemma 3.1. For the "only if" direction, let us assume that $(M, 1^*)$-freeness is contained in (N, τ)-freeness. We need to prove that this implies that there is a labeled matroid homomorphism from $(M, 1^*)$ to (N, τ).

To this end, consider a canonical function $f^0_{(N,\tau)}$ for (N, τ) padded with zeros and with the $S_{\mathbf{v}}$-sets independently and randomly chosen to be half of the available points. We know from Lemma 3.5 that such a function $f^0_{(N,\tau)}$ is not (N, τ)-free. Hence if $(M, 1^*)$-freeness is contained in (N, τ)-freeness it cannot be $(M, 1^*)$-free either, so suppose it contains $(M, 1^*)$ at the linear transformation $L : M \to \mathbb{F}^n$. We claim that if we let π be the projection that maps $\mathbf{x} = \mathbf{y}|\mathbf{z}$ to $\mathbf{y}$, then $\pi \circ L$ must be a labeled matroid homomorphism from $(M, 1^*)$ to (N, τ).

To verify the claim, notice first that the map $\pi \circ L$ is clearly linear. We need to check that it sends every vector $\mathbf{v}_i \in M$ to some vector $\mathbf{w}_j$ in N and in addition that the labels of the vectors are preserved. But since M has the monotone pattern $\langle 1^* \rangle$ the label is always 1; hence, by assumption we have $f^0_{(N,\tau)}(L(\mathbf{v}_i)) = 1$ for all i. It follows from the way the canonical function was constructed in Definition 3.4 that we must have $L(\mathbf{v}_i) = \mathbf{y}_i|\mathbf{z}_i$ where $\mathbf{y}_i = \mathbf{w}_j$ for some $\mathbf{w}_j \in N$ labeled by $\tau_j = 1$, since these are the only vectors for which $f^0_{(N,\tau)}$ evaluates to 1. Thus, $\pi \circ L$ is a labeled matroid homomorphism from $(M, 1^*)$ to (N, τ), which establishes the claim. As a final note, we point out that the structure of the $S_{\mathbf{v}}$-sets did not matter at all in this proof. □

We can use Theorem 3.6 not only to separate monotone properties from non-monotone ones, but also to separate two non-monotone properties from each other. Namely, let (N^1, τ^1) and (N^2, τ^2) be non-monotone labeled matroids such that (N^1, τ^1) is a submatroid of (N^2, τ^2) (for instance, by being a labeled subgraph). Suppose furthermore that we can find a monotone matroid $(M, 1^*)$ having the property that $(M, 1^*) \hookrightarrow (N^2, \tau^2)$ but $(M, 1^*) \not\hookrightarrow (N^1, \tau^1)$. Then it follows from Lemma 3.1 and Theorem 3.6 that (N^1, τ^1)-freeness must be strictly contained in (N^2, τ^2)-freeness. We will see examples of such results in Section 5.

However, while this is already considerably stronger than the monotone separation results in [BCSX09], it is still not quite satisfactory. The problem is that results obtained in this manner do not show that non-monotonicity adds anything essential to matroid freeness properties. For all that we know, it might be the case that (N^1, τ^1)-freeness is identical to $(M, 1^*)$-freeness so that the only essential constraint is the monotone one and the non-monotone constraints are just syntactic sugar. Our second dichotomy theorem, while being more restricted in the structural conditions it places on the matroids, is also much more powerful in that it directly separates non-monotone properties without going via monotone ones.

Theorem 3.7 (Second dichotomy theorem). *Let N be a matroid in $\mathbb{F}_2^d$ containing partial matroid $F_{d'}^{\leq w}$ of weight $w \geq 2$ as a submatroid on the first $d' \leq d$ basis vectors, and let $\langle 0^{d'}1^* \rangle$ be the pattern for N that gives 0-labels to the vectors $\mathbf{e}_1, \ldots, \mathbf{e}_{d'}$ and 1-labels to all other vectors.*

Then $\mathcal{F}_c^{\leq w}[\neg 0^c 1^*] \subseteq \mathcal{N}[\neg 0^{d'} 1^*]$ if and only if there is a labeled matroid homomorphism from $(F_c^{\leq w}, 0^c 1^*)$ to $(N, 0^{d'} 1^*)$; otherwise well-separation $\mathcal{F}_c^{\leq w}[\neg 0^c 1^*] \not\subseteq_\delta \mathcal{N}[\neg 0^{d'} 1^*]$ holds.

The two theorems 3.6 and 3.7 together constitute the formal version of our second main result in Theorem 1.4. We remark that in contrast to Theorem 3.6, in Theorem 3.7 we are only considering the case when the range of the functions in our properties is $\mathsf{R} = \{0, 1\}$.

Proof of Theorem 3.7. The "if" direction is again Lemma 3.1. For the "only if" direction, suppose that $F_c^{\leq w}$-freeness is contained in $(N, 0^{d'}1^*)$-freeness. Consider a canonical function $f = f^1_{(N, 0^{d'}1^*)}$ where all $S_{\mathbf{w}}$-sets are chosen to be the same (large) linear subspace $S \subseteq \mathbb{F}_2^{n-d}$ for some arbitrary choice of such a subspace. Observe that we are now padding the canonical function with ones, as opposed to the zero-padding in the proof of Theorem 3.6.

We know from Lemma 3.5 that f is far from being $(N, 0^{d'}1^*)$-free. Hence, by our assumption it cannot be $F_c^{\leq w}$-free either. Suppose that f contains $F_c^{\leq w}$ at $L : F_c^{\leq w} \to \mathbb{F}_2^n$. Let π be the projection that maps $\mathbf{x} = \mathbf{y}|\mathbf{z} \in \mathbb{F}_2^n$ to $\mathbf{y} \in \mathbb{F}_2^d$. We want to argue that $\pi \circ L$ must be a labeled matroid homomorphism from $(F_c^{\leq w}, 0^c 1^*)$ to $(N, 0^{d'}1^*)$.

Let us first focus on the basis vectors in $F_c^{\leq w}$, which we will denote $\mathbf{f}_1, \ldots, \mathbf{f}_c$ and which are all 0-labeled. Since f applied on the image of $F_c^{\leq w}$ under L evaluates to the pattern $\langle 0^c 1^* \rangle$, we have $f(L(\mathbf{f}_i)) = 0$ for all $i \in [c]$. Looking at the definition of f, this means that $L(\mathbf{f}_i) = \mathbf{y}_i|\mathbf{z}_i$ where either $\mathbf{y}_i = \mathbf{e}_l$ for some $\mathbf{e}_l \in N$, $l \leq d'$, or else $\mathbf{y}_i = 0$. We also note for the record, since we will need it later in the proof, that we must have $\mathbf{z}_i \in S$ in both of these cases.

Clearly, if $L(\mathbf{f}_i) = \mathbf{y}_i|\mathbf{z}_i$ for $\mathbf{y}_i = \mathbf{0}$, the linear map $\pi \circ L$ is no matroid homomorphism (since $\mathbf{0} \notin N$) and the construction breaks down. We claim, however, that this can never happen. Given this claim, which willl be established at the end of the proof, all basis vectors $\mathbf{f}_i \in F_c^{\leq w}$ must then be mapped by L to $\mathbf{e}_{l_i}|\mathbf{z}_i$ for some $\mathbf{e}_{l_i} \in N$, $l_i \leq d'$ and some $\mathbf{z}_i \in S$.

Consider now the weight-2 vectors in $F_c^{\leq w}$, i.e., sums $\mathbf{f}_i + \mathbf{f}_j$. By linearity we have $(\pi \circ L)(\mathbf{f}_i + \mathbf{f}_j) = \mathbf{e}_{l_i} + \mathbf{e}_{l_j} \in \mathbb{F}^d$ for $l_i, l_j \leq d'$. Again we have two cases. If $\mathbf{e}_{l_i} \neq \mathbf{e}_{l_j}$, then by the assumptions in the statement of the theorem we have that $\mathbf{e}_{l_i} + \mathbf{e}_{l_j}$ is a vector in N labeled by 1 as desired. If, however, $\mathbf{e}_{l_i} = \mathbf{e}_{l_j}$, then $\mathbf{f}_i + \mathbf{f}_j$ gets mapped to $\mathbf{0}$ and the construction breaks down. This cannot happen, however, since it would imply that $f(L(\mathbf{f}_i + \mathbf{f}_j)) = f(\mathbf{0}|\mathbf{z}_i + \mathbf{z}_j) = 0$. (Again for the record, this holds because $\mathbf{z}_i, \mathbf{z}_j \in S$ implies that also $\mathbf{z}_i + \mathbf{z}_j \in S$, since S is a linear subspace). But $f(L(\mathbf{f}_i + \mathbf{f}_j)) = 0 \neq 1$ contradicts the assumption that f evaluates to the pattern $\langle 0^c 1^* \rangle$ on the image of $F_c^{\leq w}$ under L. But from this it follows that since all basis vectors in $F_c^{\leq w}$ get mapped to distinct 0-labeled basis vectors in N, and since all w'-wise sums of these basis vectors for $w' \leq w$ are also in N and are labeled by 1 by assumption, all vectors in $F_c^{\leq w}$ must in fact map into N in a label-preserving way. Hence, $\pi \circ L$ is a labeled matroid homomorphism from $F_c^{\leq w}$ to N.

It remains to prove the claim that $L(\mathbf{f}_i) \neq \mathbf{0}|\mathbf{z}_i$ for all basis vectors $\mathbf{f}_i \in F_c^{\leq w}$. Suppose on the contrary that there is a vector $\mathbf{f}_i$ such that $L(\mathbf{f}_i) = \mathbf{0}|\mathbf{z}_i$. Fix some other basis vector $\mathbf{f}_j$, $j \neq i$, in $F_c^{\leq w}$ that is mapped by L to $\mathbf{y}_j|\mathbf{z}_j$ for $\mathbf{y}_j \in \{\mathbf{e}_1, \ldots, \mathbf{e}_{d'}\} \cup \{\mathbf{0}\}$, and consider $L(\mathbf{f}_i + \mathbf{f}_j) = L(\mathbf{f}_i) + L(\mathbf{f}_j) = \mathbf{y}_j|\mathbf{z}_i + \mathbf{z}_j$. By assumption $f(L(\mathbf{f}_i)) = f(L(\mathbf{f}_j)) = 0$, which by the definition of f implies that $\mathbf{z}_i, \mathbf{z}_j \in S$. This in turn means that $f(L(\mathbf{f}_i + \mathbf{f}_j)) = f(L(\mathbf{f}_j)) = 0$. But this is again a contradiction to the assumption that f evaluates to the pattern $\langle 0^c 1^* \rangle$ on $L(F_c^{\leq w})$, which requires that $f(L(\mathbf{f}_i + \mathbf{f}_j)) = 1$. The claim follows, and the proof of the theorem is complete. □

We remark that, unlike the proof of Theorem 3.6, here we crucially use the fact that all $S_{\mathbf{w}}$-sets are identically the same linear subspace and so we cannot replace these sets by, for instance, random subsets.

There is a simple but very useful way of strengthening Theorem 3.7 that we present next.

Corollary 3.8. *Let M be a binary linear matroids containing $F_c^{\leq 2}$ as a submatroid, and suppose σ is any pattern for M such that $\mathbf{f}_1, \ldots, \mathbf{f}_c \in F_c^{\leq 2}$ are all 0-labeled and $\mathbf{f}_i + \mathbf{f}_j$, $1 \leq i < j \leq c$ are all 1-labeled, but where the other labels are arbitrary. Let N be another binary linear matroids containing $F_d^{\leq 2}$ and let τ be a pattern for N that labels the vectors $\mathbf{e}_1, \ldots, \mathbf{e}_d \in F_d^{\leq 2}$ by 0 and where all other vectors of N are 1-labeled. Then if $c > d$ it holds that $\mathcal{M}[\neg\sigma] \nsubseteq_\delta \mathcal{N}[\neg\tau]$.*

Proof. Consider a canonical function $f = f^1_{(N,\tau)}$ chosen exactly as in Theorem 3.7. We know that f is dense in violations of (N, τ)-freeness. Assume that the function contains (M, σ) at L. Then clearly f also contains the submatroid $(F_c^{\leq 2}, 0^c 1^*)$ at L restricted to $F_c^{\leq 2}$. If so, going through the proof of Theorem 3.7 we can extract a labeled homomorphism ϕ of $F_c^{\leq 2}$ to N. In this embedding, all 0-labeled basis vectors $\mathbf{f}_i$ and $\mathbf{f}_j$, $i \neq j$, in $F_c^{\leq 2}$ must be sent to pairwise distinct vectors by ϕ, since otherwise there is no way we could have $\phi(\mathbf{f}_i) + \phi(\mathbf{f}_i) = \phi(\mathbf{f}_i + \mathbf{f}_j) \in N$. But there are not enough distinct 0-labeled vectors in (N, τ) to achieve this, since $c > d$. □

We note for future reference that in Corollary 3.8 we need to make no assumptions about c or d being greater than some positive constant.

Using Definition 3.4 we can also separate a matroid freeness property from a union of other matroid freeness properties.

Theorem 3.9 (Third dichotomy theorem). *Let $\mathfrak{N} = \{(N^1, \tau^1), (N^2, \tau^2), \ldots, (N^t, \tau^t)\}$ be any set of labeled binary matroids and suppose that $M \subseteq \mathbb{F}_2^k$ is some other matroid that contains $F_k^{\leq 2}$ as a submatroid and is such that $(M, 1^*) \not\to (N^i, \tau^i)$ for all $i \in [t]$. Then the property of $(M, 1^*)$-freeness is δ-separated from the union of all (N^i, τ^i)-freeness properties for $i \in [t]$, i.e., $\mathcal{M}[\neg 1^*] \nsubseteq_\delta \bigcup_{i=1}^t \mathcal{N}^i[\neg\tau^i]$.*

Proof. Consider the canonical function $f^0_{\mathfrak{N}}$. Recall that by Lemma 3.5 $f^0_{\mathfrak{N}}$ is δ-far from $\bigcup_{i=1}^t \mathcal{N}^i[\neg\tau^i]$. Suppose now that $f^0_{\mathfrak{N}}$ contains $(M, 1^*)$ at the linear transformation $L : M \to \mathbb{F}^n$. We claim that there must exist some $\ell \in [t]$ such that L maps M to vectors $\mathbf{0}|\mathbf{y}_j|\mathbf{0}|\mathbf{z}$, where $\mathbf{z} \in S_{\ell,\mathbf{w}}$ and $\mathbf{y}_j \in N^\ell$ and $\mathbf{0}$ here indicates that all the other components not corresponding to the subspace of N^ℓ are null. Given this claim, one can now restrict L to this subspace and apply (the proof of) Theorem 3.6 on $(M, 1^*)$ and (N^ℓ, τ^ℓ) to extract a labeled matroid homomorphism. (Note that this proof works regardless of the structure of the $S_{\ell,\mathbf{w}}$-sets.)

We now proceed to prove the claim. Let the basis for M be $\mathbf{f}_1, \ldots, \mathbf{f}_d$ and note that by assumption M contains all pairwise sums $\mathbf{f}_i + \mathbf{f}_j$. Since we are seeing a pattern $\langle 1^* \rangle$, the transformation L must map all elements in M to vectors of the form $\mathbf{0}|\mathbf{y}_i|\mathbf{0}|\mathbf{z}_i$ for some $i \in [t]$, where $\mathbf{y}_i$ is a nonzero vector in the subspace corresponding to the ith matroid N^i. Suppose that two basis vectors $\mathbf{f}_i$ and $\mathbf{f}_j$ are mapped to vectors $L(\mathbf{f}_i) = \mathbf{0}|\mathbf{y}_i|\mathbf{0}|\mathbf{z}_i$ and $L(\mathbf{f}_j) = \mathbf{0}|\mathbf{y}_j|\mathbf{0}|\mathbf{z}_j$, where $\mathbf{y}_i$ and $\mathbf{y}_j$ belong to subspaces corresponding to distinct matroids N^i, and N^j respectively. But this cannot be, because if so $L(\mathbf{f}_i + \mathbf{f}_j) = L(\mathbf{f}_i) + L(\mathbf{f}_j)$ would be a vector that is nonzero in two subspaces, and hence by construction $f^0_{\mathfrak{N}}(L(\mathbf{f}_i + \mathbf{f}_j)) = 0$, contrary to our assumption. Thus all basis vectors $\mathbf{f}_i$ map to vectors $\mathbf{0}|\mathbf{y}_i|\mathbf{0}|\mathbf{z}$ where $\mathbf{y}_i \in N^\ell$, for some unique $\ell \in [t]$. Since every vector $\mathbf{v} \in M$ is just a linear combination of the basis vectors, the linear transformation L must map M to vectors $\mathbf{0}|\mathbf{y}|\mathbf{0}|\mathbf{z}$ where again $\mathbf{y} \in N^{\ell}$. (Note that this second step, too, works regardless of the structure of the $S_{i,\mathbf{w}}$-sets, so we can pick them randomly here and get so many canonical functions that some are far from low-degree polynomials.)

□

4 Some Labeled Graphic Matroid Non-Homomorphisms

In order for the method developed in the previous section to be useful, we need to find (families of) labeled matroids that do not embed homomorphically into each other. In this section, we establish such matroid non-homomorphism results for graphic matroids. Recall that for labeled graphic matroids $(M(G), \sigma)$ and $(M(H), \tau)$, which we will from now identify with their underlying labeled graphs (G, σ) and (H, τ) for ease of notation, the matroid vectors correspond to edges in the graphs, and a labeled matroid homomorphism is a mapping of edges to edges that preserves labels and cycles.

We first show that there is no way to embed K_d into K_{d-1} for $d \geq 5$. We remark that as we saw in Proposition 3.3, there is in fact a homomorphism from $(K_4, 1^*)$ to $(K_3, 1^*)$. Thus, the condition $d \geq 5$ is necessary.

Lemma 4.1. *For all $d \geq 5$, there is no labeled matroid homomorphism from (K_d, σ) to (K_{d-1}, τ) for any patterns σ and τ.*

To prove Lemma 4.1, we ignore the patterns σ and τ and instead argue directly that regardless of what these patterns look like, there can exist no edge homomorphism from K_d to K_{d-1} that preserves cycle structure. This argument rests on two simple but very useful claims.

To state these claims, we recall from Definition 2.7 that the standard representation of K_d is to fix some vertex v and pick as a basis the vectors corresponding to all edges e incident to this vertex. Even once we have fixed such a vertex and associated its edges with unit vectors $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_{d-1}$, we can get another essentially equivalent basis by fixing any other vertex and looking at the vectors corresponding to edges incident to that vertex instead, as explained after Definition 2.7. We will refer to any such basis, which in this section we identify with the corresponding set of edges, as a *standard form basis*.

Claim 4.2. For any $c, d \geq 3$, if $(K_d, \sigma) \hookrightarrow (K_c, \tau)$, then it holds that any two incident edges in K_d must map to distinct edges in K_c.

Claim 4.3. For any $d \geq 5$ and $c \geq 4$, if $(K_d, \sigma) \hookrightarrow (K_c, \tau)$, then all edges in any standard form basis of K_d must map to edges in K_c that are all incident to one common vertex.

Given these two claims, Lemma 4.1 follows immediately by a pigeonhole argument: since a standard form basis in K_d has $d-1$ edges while all vertices in K_{d-1} only has $d-2$ incident edges, the claims 4.2 and 4.3 cannot possibly both hold simultaneously. This proves the lemma by contradiction.

It remains to establish the claims, and we do so next. Note again that we write v to denote a vertex and e to denote an edge, whereas vectors are denoted $\mathbf{e}, \mathbf{f}, \mathbf{v}, \mathbf{w}$, et cetera. In what follows below, we will go freely back and forth between edge and vector representation of the graphic matroids.

The first claim is more or less immediate.

Proof of Claim 4.2. Let $e_1 = (v_i, v_j)$ and $e_2 = (v_i, v_k)$ be two incident edges in K_d and suppose that they map to the same edge in K_c. Now e_1 and e_2 form a cycle in K_d together with $e_3 = (v_j, v_k)$, but since K_c does not have self-loops there is no way to map e_3 to an edge in K_c so that this cycle is preserved. (Or, reasoning in terms of matroid vectors in a linear space over $\mathbb{F}_2$, since the vectors for e_1 and e_3 map to the same vector in K_c and hence cancel, there is no way to map the third vector to a non-zero vector in K_3 so that the sum of the images of all three vectors cancel.) □

The second claim is not much harder, but requires a little more work.

Proof of Claim 4.3. Suppose K_d embeds into K_c by a linear map ϕ and let $\{\mathbf{f}_1, \ldots, \mathbf{f}_{d-1}\}$ be the basis vectors of K_d and $\{\mathbf{e}_1, \ldots, \mathbf{e}_{c-1}\}$ be the basis vectors of K_c. We show by induction on k, $1 \leq k \leq d-1$, that $\{\mathbf{f}_1, \ldots, \mathbf{f}_k\}$ must map to (distinct) edges incident to some common vertex v in K_c.

Without loss of generality, assume $\phi(\mathbf{f}_1) = \mathbf{e}_1$ (if $\mathbf{f}_1$ would map to any other vector we just make a basis change in K_c as explained after Definition 2.7). By Claim 4.2 we have $\phi(\mathbf{f}_2) \neq \mathbf{e}_1$. Also note that $\phi(\mathbf{f}_2) \neq \mathbf{e}_i + \mathbf{e}_j$ with $1 < i < j$, because this would imply $\phi(\mathbf{f}_1 + \mathbf{f}_2) = \mathbf{e}_1 + \mathbf{e}_i + \mathbf{e}_j$, which is weight-3 vector that is not a member of K_c. (In terms of edges, this would correspond to incident edges $e_1 = (v_i, v_j)$ and $e_2 = (v_i, v_k)$ in K_d mapping to non-incident edges in K_c, but if so there would be no way the map ϕ could preserve the cycle of e_1 and e_2 with $e_3 = (v_j, v_k)$.)

Therefore we are left with two cases: $\phi(\mathbf{f}_2) = \mathbf{e}_i$ or $\phi(\mathbf{f}_2) = \mathbf{e}_1 + \mathbf{e}_i$ for some $i > 1$, and because of symmetry we may choose $i = 2$ without loss of generality. Let us analyze these two cases.[7]

Case 1 ($\phi(\mathbf{f}_2) = \mathbf{e}_2$): Consider where ϕ can send $\mathbf{f}_3$. Clearly $\phi(\mathbf{f}_3) \notin \{\mathbf{e}_1, \mathbf{e}_2\} = \{\phi(\mathbf{f}_1), \phi(\mathbf{f}_2)\}$ by the distinctness in Claim 4.2. Also, we claim that $\phi(\mathbf{f}_3) \neq \mathbf{e}_1 + \mathbf{e}_2$. To see this, observe that if $\{\phi(\mathbf{f}_1), \phi(\mathbf{f}_2), \phi(\mathbf{f}_3)\} = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1 + \mathbf{e}_2\}$, then ϕ would have to map $\mathbf{f}_4$ to some vector outside of this set by distinctness, but if so it in turns follows that $\phi(\mathbf{f}_3 + \mathbf{f}_4) = \phi(\mathbf{f}_3) + \phi(\mathbf{f}_4)$ would be a vector of weight at least 3. (Notice that here we crucially use $d \geq 5$). To conclude, we argue that in fact $\phi(\mathbf{f}_3) \neq \mathbf{e}_i + \mathbf{e}_j$ for any $i < j$ with $j > 2$. For if $i = 1$, say, it follows just as above that $\phi(\mathbf{f}_2 + \mathbf{f}_3) = \phi(\mathbf{f}_2) + \phi(\mathbf{f}_3)$ would be a weight-3 vector. Hence we must have $\phi(\mathbf{f}_3) = \mathbf{e}_i$ for some i, which we may without loss of generality set to $i = 3$.

Case 2 ($\phi(\mathbf{f}_2) = \mathbf{e}_1 + \mathbf{e}_2$): In this case we have $\phi(\mathbf{f}_3) \neq \mathbf{e}_1$ by distinctness, and it must also hold that $\phi(\mathbf{f}_3) \neq \mathbf{e}_2$, since otherwise we would have $\{\phi(\mathbf{f}_1), \phi(\mathbf{f}_2), \phi(\mathbf{f}_3)\} = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1 + \mathbf{e}_2\}$ which can be ruled out by the same argument as in case 1. Furthermore, $\phi(\mathbf{f}_3) \neq \mathbf{e}_i$ for $i > 2$, since otherwise $\phi(\mathbf{f}_2 + \mathbf{f}_3)$ would be a weight-3 vector. Hence, $\phi(\mathbf{f}_3) = \mathbf{e}_i + \mathbf{e}_j$ for some $i < j$. But then we must have $i = 1$, since otherwise $\phi(\mathbf{f}_1 + \mathbf{f}_3)$ would have weight 3. We have proven that in this case as well, the vectors $\phi(\mathbf{f}_1) = \mathbf{e}_1$, $\phi(\mathbf{f}_2) = \mathbf{e}_1 + \mathbf{e}_2$, and $\phi(\mathbf{f}_3) = \mathbf{e}_1 + \mathbf{e}_j$ must correspond to edges incident to a common vertex.

Now we proceed to the inductive step. Suppose ϕ maps the vectors $\{\mathbf{f}_1, \ldots, \mathbf{f}_k\}$ to k edges incident to a single vertex in K_c for some $k \geq 4$. Without loss of generality, we assume that the vectors are mapped to $\mathbf{e}_1, \ldots, \mathbf{e}_k$ in K_c (because again, since we know that the edges are incident we are free to make a basis change in K_c so that we get the standard basis in terms of unit vectors). Consider the image of edge $\mathbf{f}_{k+1}$ in K_c. Clearly, $\phi(\mathbf{f}_{k+1}) \neq \mathbf{e}_i$ for any $i \in [k]$ by distinctness. It is also easy to see that $\phi(\mathbf{f}_{k+1})$ cannot be any weight-2 vector. For if we would have $\phi(\mathbf{f}_{k+1}) = \mathbf{e}_i + \mathbf{e}_j$, then since $k \geq 3$ there would exist some $l \in [k]$ with $l \notin \{i, j\}$ such that $\phi(\mathbf{f}_l + \mathbf{f}_{k+1}) = \mathbf{e}_i + \mathbf{e}_j + \mathbf{e}_l \notin K_c$, which is a contradiction. Therefore, $\phi(\mathbf{f}_{k+1}) = \mathbf{e}_i$ for some $i > k$. This completes the induction step and thus finishes the proof of the claim. □

For our second lemma, we observe that we can also use Claim 4.3 to study labeled homomorphisms from K_d into itself and rule out that any two distinct patterns $\langle 0^{c_1} 1^* \rangle$ and $\langle 0^{c_2} 1^* \rangle$ for $c_1 \neq c_2$ can be mapped homomorphically into one another. Again we need the condition $d \geq 5$ here, since otherwise we could use the construction in Proposition 3.3 to embed, for instance, $(K_4, 1^*)$ into $(K_4, 0^c 1^*)$ for $0 < c \leq 3$.

Lemma 4.4. *For all $d \geq 5$ and $c_1 \neq c_2$, there is no labeled matroid homomorphism from $(K_d, 0^{c_1} 1^*)$ to $(K_d, 0^{c_2} 1^*)$.*

[7] In fact, the attentive reader might have noted here that without loss of generality we can restrict ourselves to only one case and fix $\phi(\mathbf{f}_2) = \mathbf{e}_2$. This is so since in the other case we can again make a basis change in K_c as described after Definition 2.7 to get a new standard basis containing $\mathbf{e}_1$ and $\mathbf{e}_1 + \mathbf{e}_2$. However, we believe that a formal case analysis as given in this proof, although strictly speaking unnecessary, is easier to follow.

Proof. Fix any basis of $(K_d, 0^{c_1}1^*)$. It follows from Claim 4.3 that if there were a homomorphism from $(K_d, 0^{c_1}1^*)$ to $(K_d, 0^{c_2}1^*)$, then the basis vectors of $(K_d, 0^{c_1}1^*)$ must map into a basis of $(K_d, 0^{c_2}1^*)$. This in turn clearly implies that the homomorphism is one-to-one and onto, so it is an isomorphism between $(K_d, 0^{c_1}1^*)$ and $(K_d, 0^{c_2}1^*)$. But this is impossible, since the number of 0-labeled vectors in the two matroids are different. □

5 Infinite Hierarchies of Well Separated Matroid Freeness Properties

We have finally reached the point where we can put all the material in Sections 3 and 4 together and prove the existence of infinite hierarchies of (M, σ)-freeness properties as claimed in the introduction. We remark that the techniques we have developed could be used to yield many different such hierarchies. For brevity and concreteness, we will focus below on two specific results that illustrate this general point.

First, we will prove a sequence of lemmas leading up to Theorem 5.4, which is one instantiation of our fourth main result claimed informally in Theorem 1.6. Theorem 5.4 exhibits two nested hierarchies of properties where each pair of properties on the same level is separated from each other and from properties on previous levels in a very strong sense.

Second, we prove a more general result in Theorem 5.5 that establishes even an infinite number of nested property hierarchies. The proof of this second theorem turns out to be much more intricate, however, and therefore cannot be presented in full in this section. Also, although we believe that the separation between the property hierarchies should be as strong as in Theorem 5.4, we are currently unable to prove this as numerous technical obstacles arise during the course of the proof.

As the first step towards establishing Theorem 5.4, let us prove that monotone $M(K_d)$-freeness properties form a strict hierarchy. Note that we will continue the mild abuse of notation introduced in Section 4 by identifying a graphic matroid $(M(G), \sigma)$ and its underlying labeled graph (G, σ).

Lemma 5.1. *For $d \geq 4$, the $\mathcal{K}_d[\neg 1^*]$ properties form an infinite hierarchy of strictly contained properties.*

Proof. Since $(K_d, 1^*)$ is a labeled subgraph of $(K_{d+1}, 1^*)$ we have that $\mathcal{K}_d[\neg 1^*]$ is contained in $\mathcal{K}_{d+1}[\neg 1^*]$ by Corollary 3.2. Since there is no homomorphism from $(K_{d+1}, 1^*)$ to $(K_d, 1^*)$ for $d \geq 4$ according to Lemma 4.1, we conclude from Theorem 3.6 that the containment must be strict in a property testing sense. □

We remark that this lemma improves on a similar theorem of [BCSX09], which could only show separation between the graphic matroids of K_d and $K_{\binom{d}{2}+2}$. But we can strengthen Lemma 5.1 even further as follows.

Lemma 5.2. *For $d \geq 4$ and any sequence $\{c_d\}_{d=4}^{\infty}$ with $1 \leq c_d < d$, the $\mathcal{K}_d[\neg 0^{c_d}1^*]$ properties form an infinite sequence of properties such that $\mathcal{K}_{d+1}[\neg 0^{c_{d+1}}1^*]$ is δ-separated from $\mathcal{K}_d[\neg 0^{c_d}1^*]$. If in addition it holds for the c_d-sequence that $c_d \leq c_{d+1} \leq c_d + 1$, then we get an infinite hierarchy of strictly contained properties.*

Proof. $(K_d, 1^*)$ does not embed in $(K_d, 0^{c_d}1^*)$ by Lemma 4.4. However, $(K_d, 1^*)$ is a labeled subgraph of $(K_{d+1}, 0^{c_{d+1}}1^*)$, since if we throw away the unique vertex in $(K_{d+1}, 0^{c_{d+1}}1^*)$ incident to all 0-labeled edges what we have left is exactly $(K_d, 1^*)$. It follows that the function $f^0_{(K_d, 0^{c_d}1^*)}$, which is far from being $(K_d, 0^{c_d}1^*)$-free, is $(K_d, 1^*)$-free by Theorem 3.6 and hence $(K_{d+1}, 0^{c_{d+1}}1^*)$-free by Corollary 3.2

If it furthermore holds that $c_d \leq c_{d+1} \leq c_d + 1$, then $(K_d, 0^{c_d}1^*)$ is a labeled subgraph of $(K_{d+1}, 0^{c_{d+1}}1^*)$, which gives containment of the corresponding matroid freeness properties (again by Corollary 3.2). □

Observe that Lemma 5.2 is indeed a strengthening of Lemma 5.1. This is so since the functions $f^0_{(K_d,0^{c_d}1^*)}$, which are $(K_d,1^*)$-free but far from $(K_{d-1},1^*)$-free, witness that $\mathcal{K}_d[\neg 1^*]$ is δ-separated from $\mathcal{K}_{d-1}[\neg 1^*]$ (and containment in the other direction is obvious since $(K_{d-1},1^*)$ is a subgraph of $(K_d,1^*)$).

Notice, however, that as discussed after the proof of Theorem 3.6, the way we establish Lemma 5.2 does not ensure that non-monotone matroid freeness properties are nontrivial. We might worry that perhaps all the non-monotone properties coincide with the intermediate monotone properties of $\mathcal{K}_d[\neg 1^*]$ used to obtain the separation. The next lemma provides some assurance to us by conclusively ruling out this possibility.

Lemma 5.3. *For $d \geq 4$, it holds that $\mathcal{K}_{d+1}[\neg 0^d 1^*]$ is δ-separated from the union of $\mathcal{K}_d[\neg 0^{d-1}1^*]$ and $\mathcal{K}_d[\neg 1^*]$.*

Proof. Note first that this lemma is conceptually different from the preceding ones, since here we need to find a function that is $(K_{d+1},0^d1^*)$-free but is simultaneously far from being $(K_d,0^{d-1}1^*)$-free and $(K_d,1^*)$-free. On the face of it, such cases are not covered by the techniques in Sections 3 and 4, which only relates pair of labeled matroids. However, there is a way to get around this obstacle by finding an "intermediate" labeled matroid such that $(K_d,0^{d-1}1^*)$ and $(K_d,1^*)$ both embed into this matroid but $(K_{d+1},0^d1^*)$ does not. We pick this intermediate matroid to be $(K_{d+1},0^{d-1}1^*)$, that is, the graphic matroid over the complete graph on $d+1$ vertices that have all edges *but one* in the standard basis labeled by 0 and has 1-labels everywhere else.

It is easy to check that $(K_d,0^{d-1}1^*)$ and $(K_d,1^*)$ are both labeled subgraphs of $(K_{d+1},0^{d-1}1^*)$. Consequently, we can appeal to Lemma 3.5 to conclude that the canonical function $f^1_{(K_{d+1},0^{d-1}1^*)}$ is dense in violations of both $(K_d,0^{d-1}1^*)$-freeness and $(K_d,1^*)$-freeness. However, Lemma 4.4 shows that $(K_{d+1},0^d1^*)$ does not embed homomorphically into $(K_{d+1},0^{d-1}1^*)$, and therefore $f^1_{(K_{d+1},0^{d-1}1^*)}$ must be $(K_{d+1},0^d1^*)$-free according to Theorem 3.7. The lemma follows. □

Combining all of these lemmas, we can prove that non-monotone graphic matroid freeness properties provide infinite hierarchies of strictly contained properties. The reader might be helped in parsing the next theorem and its proof by looking at the Venn diagram-style illustration in Figure 2.

Theorem 5.4 (First hierarchy theorem). *Let $\mathcal{K}_d[\neg 1^*]$ denote the set of all $(K_d,1^*)$-free functions and $\mathcal{K}_d[\neg 0^{d-1}1^*]$ the set of all $(K_d,0^{d-1}1^*)$-free functions in $\bigcup_{n\in\mathbb{N}^+}\{\mathbb{F}_2^n \to \{0,1\}\}$. Then the following holds:*

1. *The properties $\mathcal{K}_d[\neg 1^*]$ for $d \geq 4$ form an infinite hierarchy of strictly contained properties.*
2. *The properties $\mathcal{K}_d[\neg 0^{d-1}1^*]$ for $d \geq 4$ form an infinite hierarchy of strictly contained properties.*
3. *All of the above properties intersect in the sense that $\bigcap_{d=4}^{\infty}(\mathcal{K}_d[\neg 1^*] \cap \mathcal{K}_d[\neg 0^{d-1}1^*]) \neq \emptyset$.*
4. *$\left(\bigcap_{d=4}^{\infty}\mathcal{K}_d[\neg 0^{d-1}1^*]\right) \setminus \left(\bigcup_{d=4}^{\infty}\mathcal{K}_d[\neg 1^*]\right) \neq \emptyset$, and in fact the first intersection of properties is δ-separated from the second union.*
5. *For $d \geq 4$, $\mathcal{K}_d[\neg 1^*] \cup \mathcal{K}_d[\neg 0^{d-1}1^*]$ is strictly contained in $\mathcal{K}_{d+1}[\neg 0^d1^*]$.*
6. *For $d \geq 4$, $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^{d-1}1^*]$ are mutually well separated from one another.*
7. *For $d \geq 5$, all the properties $\mathcal{K}_d[\neg 1^*]$, $\mathcal{K}_d[\neg 0^{d-1}1^*]$, $\mathcal{K}_d[\neg 1^*] \setminus \mathcal{K}_{d-1}[\neg 1^*]$, and $\mathcal{K}_d[\neg 0^{d-1}1^*] \setminus \mathcal{K}_{d-1}[\neg 0^{d-2}1^*]$ are far from being low-degree polynomials.*

Proof. The proofs of Claims 1 and 2 were given in Lemma 5.1 and Lemma 5.2, respectively. As was discussed after the proof of Lemma 5.2, both of these hierarchies are witnessed by the functions $g_d = f^0_{(K_d,0^{c_d}1^*)}$ as plotted schematically in Figure 2. To show claim 3, we observe that the constant function

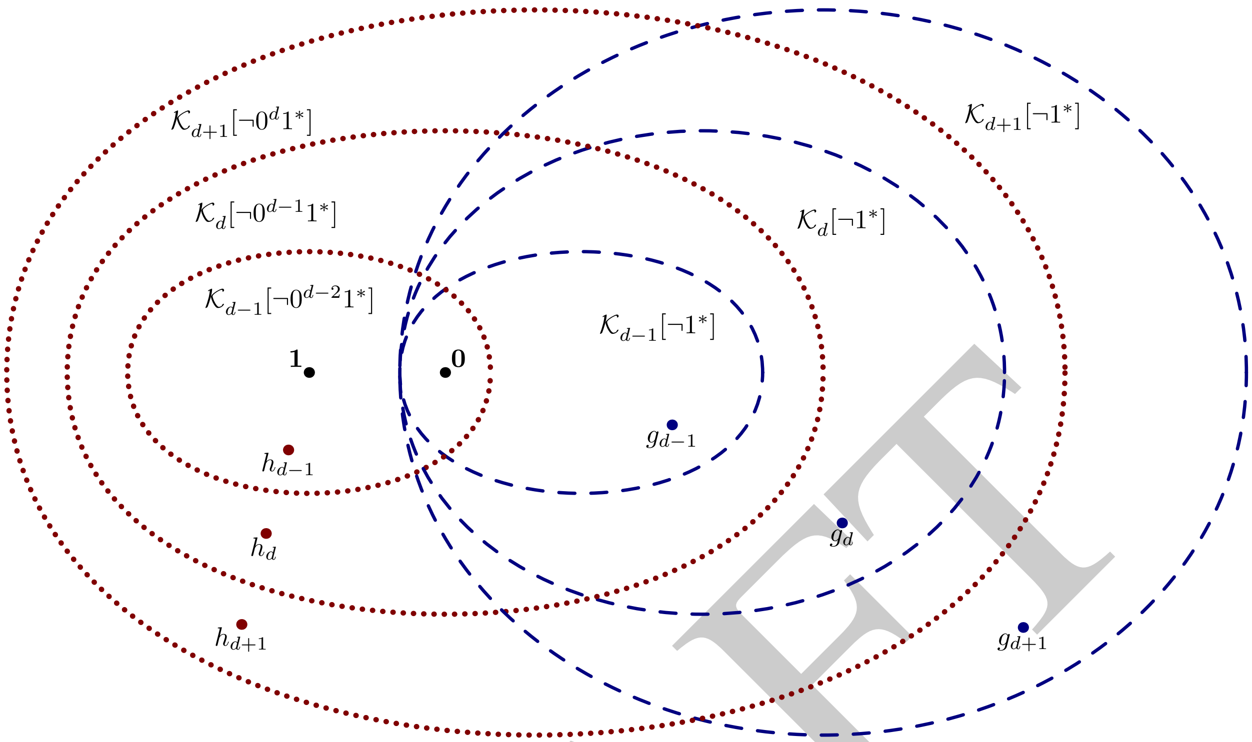

Figure 2: Illustration of hierarchies of properties $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^{d-1} 1^*]$ and separating functions.

$\mathbf{0} : \mathbb{F}_2^n \to \{0, 1\}$ sending all points to 0 belongs to all of the properties $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^{d-1} 1^*]$. Claim 4 similarly follows since the constant function $\mathbf{1} : \mathbb{F}_2^n \to \{0, 1\}$ sending all points to 1 must be $(K_d, 0^{d-1} 1^*)$-free for all d simply by virtue of not having any zeros, while it is far from $(K_d, 1^*)$-free for all d for exactly the same reason. Claim 5 was established in Lemma 5.3, using the functions $h_d = f^1_{(K_d, 0^{d-2} 1^*)}$, and taken together, the functions g_d and h_d can be seen to witness the mutual separations in claim 6.

Consider finally claim 7. Recall that what we want to say is that the properties $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^{d-1} 1^*]$ are "new" properties not known to be testable before. Note that by necessity, such a statement must be somewhat informal — unless we can provide a full enumeration of all testable properties and separate our new properties from all of them via some kind of diagonalization argument, which arguably seems neither feasible nor particularly reasonable. But what seems natural to do is to prove formally that matroid freeness properties are not identical to the "usual suspects", which in this case would seem to be low-degree polynomials.

We remark that one can first make the easy observation that it cannot possibly be the case that *all* of the properties $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^{d-1} 1^*]$ are low-degree polynomials. If they were, there would be no way they could nest and intersect in the way shown in Figure 2, since low-degree polynomials just form one strict hierarchy of concentric circles with respect to degree. We want to prove something stronger, however, namely that *none* of the properties $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^{d-1} 1^*]$ can be just low-degree polynomials.

By picking large, random sets $S_{\mathbf{w}}$ in Definition 3.4 we get a huge family of canonical functions instead of just one, and Lemma 3.5 and Theorem 3.6 still hold.

What this means is that we can think of every witnessing function g_d in Figure 2 as being a large, dense cloud of such functions, and these functions are simply far too many to all be low-degree polynomials, or even to be close to low-degree polynomials. This establishes the claim, and the theorem follows. □

The attentive reader might have noticed that there is one natural piece missing in Theorem 5.4, namely the claim that $\mathcal{K}_{d+1}[\neg 0^d 1^*] \setminus (\mathcal{K}_d[\neg 1^*] \cup \mathcal{K}_d[\neg 0^{d-1} 1^*])$ is also far from being just low-degree polynomials. This seems very likely to be the case but we are currently unable to prove this. It is easy to prove that there are polynomials of very high degree in $\mathcal{K}_{d+1}[\neg 0^d 1^*] \setminus (\mathcal{K}_d[\neg 1^*] \cup \mathcal{K}_d[\neg 0^{d-1} 1^*])$. The way to see this is to define $h_d = f^1_{(K_d, 0^{d-2}1^*)}$ as in Definition 3.4 except that we choose $S_{\mathbf{w}} = S$ to be a very small subspace, say of constant dimension. Then h_d will evaluate to 1 everywhere except at a constant number of points, and therefore it cannot possibly be a low-degree polynomial. However, all such h_d are also close to the constant function evaluating to 1 everywhere, which has very low degree indeed. One natural idea is instead to pick the set S randomly to get a large number of functions h_d, and then argue that they are so many that here must be examples of such functions that are far from being low-degree polynomials. Unfortunately, this does not work. The proof of Theorem 3.7 turns out to be surprisingly delicate, and provably requires S to be a subspace. We still strongly believe that the properties $\mathcal{K}_{d+1}[\neg 0^d 1^*] \setminus (\mathcal{K}_d[\neg 1^*] \cup \mathcal{K}_d[\neg 0^{d-1} 1^*])$ are far from low-degree polynomials for all $d \geq 4$, but it seems new techniques would be required to establish such a claim.

We have seen in Theorem 5.4 that the two properties of $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^{d-1} 1^*]$ are distinct. As described before, a very natural question is whether it holds in general that two distinct labeled matroids (M, σ) and (N, τ) must encode distinct matroid freeness properties.

Notice that we have to make this question more precise in order for it to make sense. Clearly, there are many ways of encoding the same property. For instance, if we just permute the vertices of K_d by some permutation π and let τ be the corresponding permutation of the pattern σ on the edges reordered according to π, then $\mathcal{K}_d[\neg\sigma]$ is identical to $\mathcal{K}_d[\neg\tau]$ although σ and τ might look syntactically different. This is a somewhat artificial example, however, since in this case (K_d, σ) and (K_d, τ) are isomorphic as labeled matroids, so obviously they must encode the same property.

A more reasonable way of phrasing the question is to ask whether it is the case that labeled matroids that are "not the same" in some more fundamental, non-syntactical sense must encode different properties. In particular, a very natural question is whether the properties of (M, σ)-freeness and (N, τ)-freness are distinct, or even mutually well separated, whenever (M, σ) and (N, τ) do not embed into one another homomorphically. Recall that this is a necessary condition in the sense that if $(M, \sigma) \hookrightarrow (N, \tau)$, then (M, σ)-freeness is contained in (N, τ)-freness. For the particular case of $\mathcal{K}_d[\neg 1^*]$ versus $\mathcal{K}_d[\neg 0^{d-1} 1^*]$, it also seemed to be sufficient in the sense that the proof of Theorem 5.4 relied heavily on the fact that the two matroids do not embed homomorphically into one another.

In general, this question is still way out of reach for two arbitrary labeled binary matroids (M, σ) and (N, τ), and even for two graphic matroids (K_d, σ) and (K_d, τ) with arbitrary patterns σ and τ. However, we establish in Theorem 5.5 below that for a fairly broad class of (K_d, σ)-freness properties the question of whether they are distinct or not is determined precisely by the existence or non-existence of matroid homomorphisms between them.

Namely, we can show that for any two labeled matroids (K_d, σ) and (K_d, τ) which have 1-labels everywhere except on a subset of some standard basis of K_d, and where the number of 0-labeled vectors is distinct in σ and τ, it holds that $\mathcal{K}_d[\neg\sigma]$ and $\mathcal{K}_d[\neg\tau]$ are distinct in a property testing sense. We remark that without loss of generality we can write such matroids as $(K_d, 0^c 1^*)$ for $0 \leq c \leq d-1$, since we can fix the standard basis to be $\mathbf{e}_1, \ldots, \mathbf{e}_{d-1}$ and permute the basis vectors so that all 0-labeled vectors come first in the pattern. Note that this result significantly broadens the result in Theorem 5.4, which only relates the two special cases of $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^{d-1} 1^*]$

As an example, we have that $\mathcal{K}_d[\neg 1^*]$ and $\mathcal{K}_d[\neg 0^c 1^*]$ for $c > 0$ are pairwise mutually well separated. In fact, in this particular case it is easy to show that this holds for all patterns of zeros and ones, and in particular for all $c > 0$ and not just for $0 < c \leq d-1$. A natural question is what happens when we start to combine these properties. Could it for instance be that the union $\bigcup_{c>1} \mathcal{K}_d[\neg 0^c 1^*]$ of all $(K_d, 0^c 1^*)$-freeness properties completely covers $(K_d, 1^*)$-freeness? This possibility cannot be ruled out by the results above, but using Theorem 3.9 we can also show that the property of $(K_d, 1^*)$-freeness in fact is not contained even in the union of all $(K_d, 0^c1^*)$-freeness properties.

The formal statement of the results described above follows in Theorem 5.5. Since this theorem relates a large number of properties, it is not possible to provide a Venn diagram of the relations. However, in Figure 3 we attempt to sketch how some of the properties are related. Please observe that the figure does not provide a complete picture of relations since such a picture would become hopelessly cluttered, but we hope that the example relations illustrated can nevertheless help the reader get a sense of what Theorem 5.5 is saying.

Theorem 5.5 (Second hierarchy theorem). *Let $\mathcal{K}_d[\neg 0^c1^*]$ denote the set of all $(K_d, 0^c1^*)$-free functions in $\bigcup_{n\in\mathbb{N}^+}\{\mathbb{F}_2^n \to \{0,1\}\}$ for $d \geq 5$. Then we have the following relations:*

1. *For $0 \leq c_1 < c_2 \leq d-1$ it holds that $\mathcal{K}_d[\neg 0^{c_2}1^*] \not\subseteq_\delta \mathcal{K}_d[\neg 0^{c_1}1^*]$.*
2. *For $0 \leq c_1 < c_2 \leq d-1$, $c_1 \neq 1$, it holds that $\mathcal{K}_d[\neg 0^{c_1}1^*] \not\subseteq_\delta \mathcal{K}_d[\neg 0^{c_2}1^*]$.*
3. *$\mathcal{K}_d[\neg 01^*] \not\subseteq_\delta \mathcal{K}_d[\neg 0^{d-1}1^*]$, but for $1 < c < d-1$ it holds that $\mathcal{K}_d[\neg 01^*] \subset_\delta \mathcal{K}_d[\neg 0^c1^*]$.*
4. *For all $0 \leq c_1 \leq d-1$ and $0 \leq c_2 \leq d$, it holds that $\mathcal{K}_{d+1}[\neg 0^{c_2}1^*] \not\subseteq_\delta \mathcal{K}_d[\neg 0^{c_1}1^*]$.*
5. *For $0 \leq c \leq d-1$, it holds that $\mathcal{K}_d[\neg 0^c1^*] \subset_\delta \mathcal{K}_{d+1}[\neg 0^c1^*]$ and $\mathcal{K}_d[\neg 0^c1^*] \subset_\delta \mathcal{K}_{d+1}[\neg 0^{c+1}1^*]$.*
6. *It holds that $\mathcal{K}_d[\neg 1^*] \not\subseteq_\delta \bigcup_{c=1}^{d(d-1)/2} \mathcal{K}_d[\neg 0^c1^*]$, and furthermore $\mathcal{K}_d[\neg 1^*] \setminus \bigcup_{c=1}^{d(d-1)/2} \mathcal{K}_d[\neg 0^c1^*]$ is provably far from being the property of low-degree polynomials.*

Observe that one way of obtaining the infinite number of nested property hierarchies claimed in the beginning of this section is to consider chains $\mathcal{K}_d[\neg 0^{c_d}1^*] \subset_\delta \mathcal{K}_{d+1}[\neg 0^{c_{d+1}}1^*] \subset_\delta \mathcal{K}_{d+2}[\neg 0^{c_{d+2}}1^*] \subset_\delta \ldots$ where at each level we pick either $c_{d+1} = c_d$ or $c_{d+1} = c_d + 1$. As can be seen in Figure 3, all distinct choices of starting value c_d and increments 0/1 will give rise to different strict hierarchies.

As we mentioned in the beginning of this section, the proof of Theorem 5.5 requires quite a lot of work, and is also quite lengthy. We therefore defer the full proof to Appendix B. One of the problems in the proof is that although the canonical functions in Definition 3.4 are a very useful tool, it turns out that they are not powerful enough to give us all the results that we need. We conclude this section by providing an example of this, showing first an instance where the canonical functions fail, and then proving a simple special case in Theorem 5.5 that gives some of the flavor or how these difficulties can be overcome. Again, the general argument turns out to be much more complicated and therefore is not presented in this section.

Hence, by way of example, consider the properties $\mathcal{K}_d[\neg 0^{d-4}1^*]$ and $\mathcal{K}_d[\neg 0^{d-1}1^*]$, which are claimed in Theorem 5.5 to be mutually well separated. It is straightforward to show that $\mathcal{K}_d[\neg 0^{d-1}1^*]$ is δ-separated from $\mathcal{K}_d[\neg 0^{d-4}1^*]$ This follows from the fact that $(K_d, 0^{d-1}1^*) \not\hookrightarrow (K_d, 0^{d-4}1^*)$ as shown in Lemma 4.4, and the separation is established by studying the canonical function $f^0_{(K_d, 0^{d-4}1^*)}$. In the other direction, we also know from Lemma 4.4 that $(K_d, 0^{d-4}1^*) \not\hookrightarrow (K_d, 0^{d-1}1^*)$. However, we cannot use this non-embedding fact and canonical functions $f^b_{(K_d, 0^{d-1}1^*)}$ to prove $\mathcal{K}_d[\neg 0^{d-4}1^*] \not\subseteq_\delta \mathcal{K}_c[\neg 0^{d-1}1^*]$.

To see this, suppose first that we pad the canonical function $f^b_{(K_d, 0^{d-1}1^*)}$ with $b = 1$. In this case we can just map the 0-labeled basis vectors in $(K_d, 0^{d-4}1^*)$ to 0-labeled basis vectors in $(K_d, 0^{d-1}1^*)$, and then send the three 1-labeled basis vectors to three different vectors in $\mathbb{F}_2^n$ of weight roughly $n/3$ and with disjoint support. Then by linearity, all pairwise sums of basis vectors will map either to weight-2 vectors, which are 1-labeled since the encode the matroid $(K_d, 0^{d-1}1^*)$, or to vectors of weight at least $n/3$, which are 1-labeled by padding. Hence, $f^1_{(K_d, 0^{d-1}1^*)}$ contains the matroid pattern $(K_d, 0^{d-4}1^*)$. Also, it can also be shown that if we pad with $b = 0$, then $f^0_{(K_d, 0^{d-1}1^*)}$ is not $(K_d, 0^{d-4}1^*)$-free either. Namely, the reader

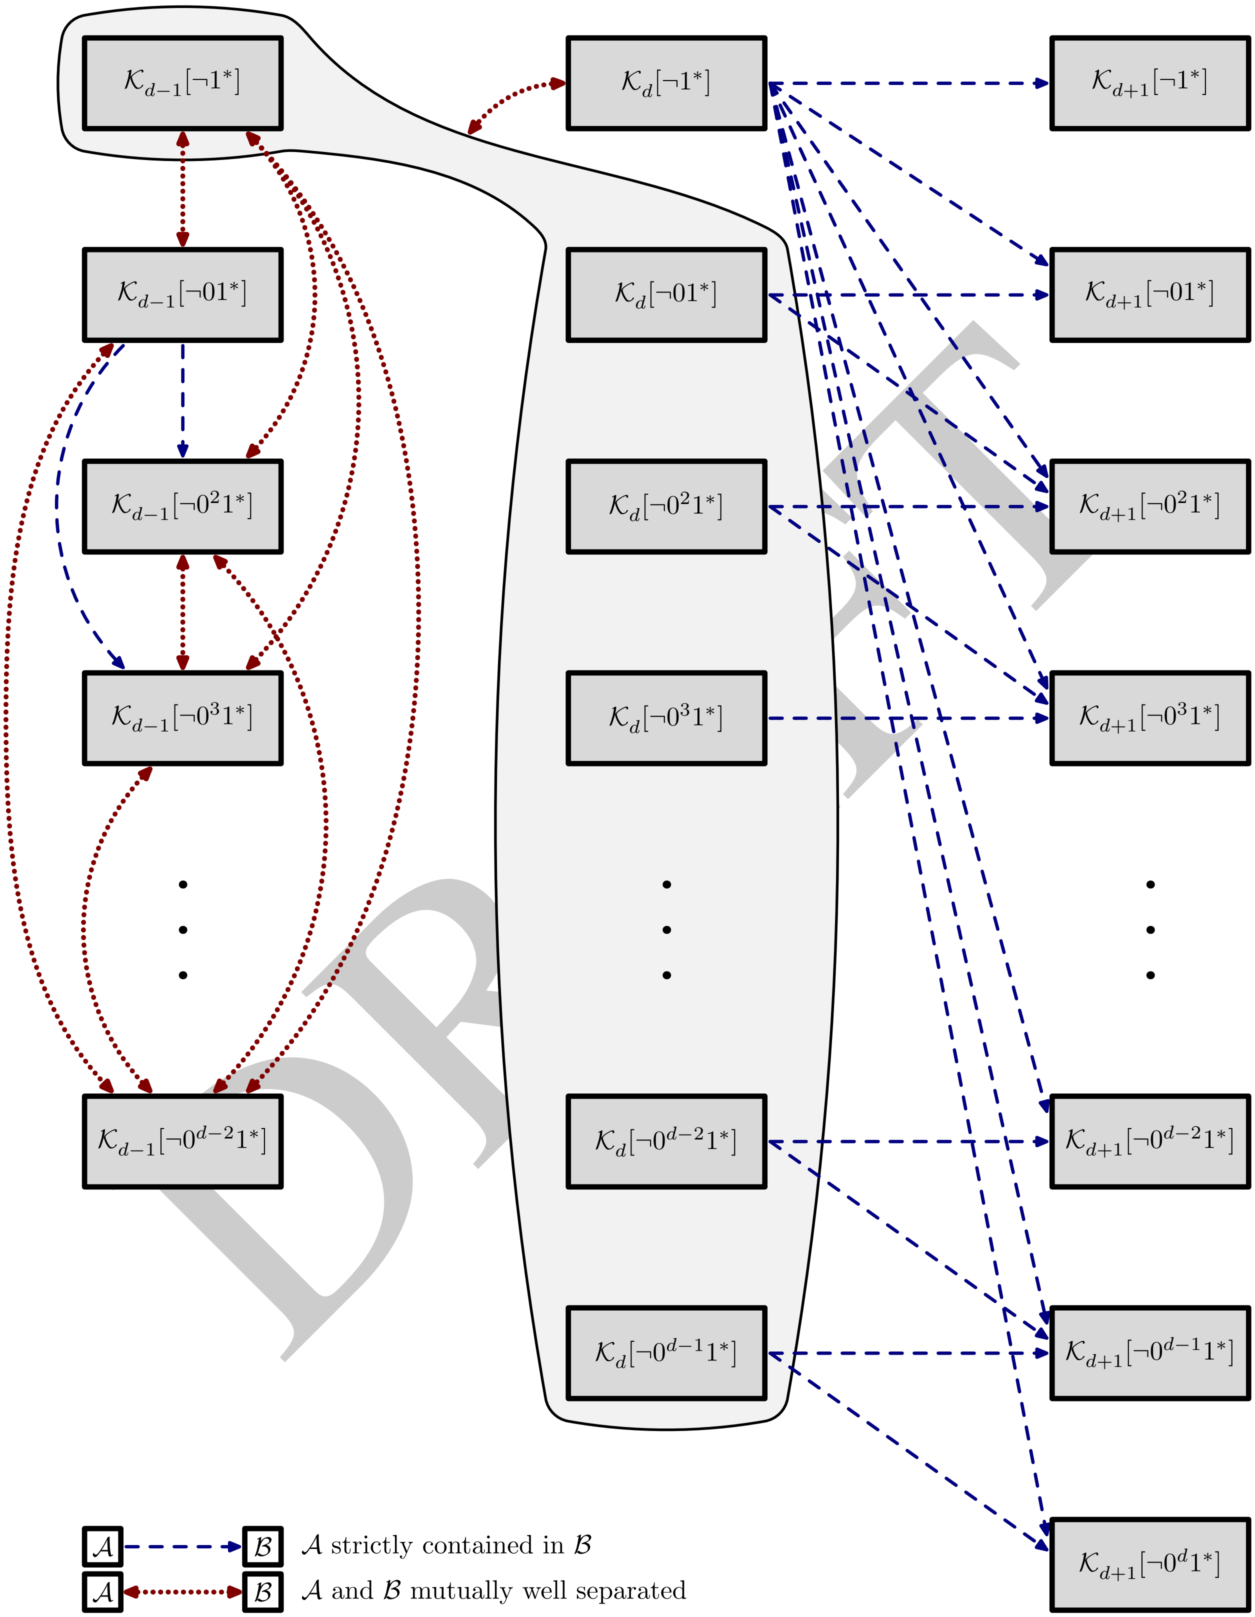

Figure 3: Illustration of (some) relations between properties $\mathcal{K}_d[\neg 0^c 1^*]$ for all $d \geq 5$ and $c = 0, 1, \ldots, d-1$.

can verify that if we take $L : \mathbb{F}_2^{d-1} \to \mathbb{F}_2^n$ to be the linear mapping that sends the 1-labeled basis vectors $\mathbf{f}_{d-3}, \mathbf{f}_{d-2}, \mathbf{f}_{d-1}$ in $(K_d, 0^{d-4}1^*)$ to $L(\mathbf{f}_{d-3}) = \mathbf{e}_1 + \mathbf{e}_2$, $L(\mathbf{f}_{d-2}) = \mathbf{e}_1 + \mathbf{e}_3$, and $L(\mathbf{f}_{d-1}) = \mathbf{e}_2 + \mathbf{e}_3$ in $\mathbb{F}_2^n$, and that maps the 0-labeled vectors $\mathbf{f}_i$, $1 \le i \le d-4$ to $L(\mathbf{f}_i) = \mathbf{e}_1 + \mathbf{e}_2 + \mathbf{e}_3 + \mathbf{e}_{i+3}$, then $f^0_{(K_d, 0^{d-1}1^*)}$ contains the matroid pattern $(K_d, 0^{d-4}1^*)$ at L. So both flavors of the canonical function fail to separate these two properties.

However, there is another simple function that establishes the separation $\mathcal{K}_d[\neg 0^{d-4}1^*] \nsubseteq_\delta \mathcal{K}_c[\neg 0^{d-1}1^*]$, as follows from the next lemma.

Lemma 5.6. $\mathcal{K}_{d_1}[\neg 0^{c_1}1^*] \nsubseteq_\delta \mathcal{K}_d[\neg 0^{d-1}1^*]$ *for* $d_1, d \ge 3$ *and* $1 \le c_1 < d_1 - 1$.

Proof. Consider the function $g(\mathbf{x}) = g(y|\mathbf{z})$ for $y \in \mathbb{F}_2$ and $\mathbf{z} \in \mathbb{F}_2^{n-1}$ defined by $g(y|\mathbf{z}) = 1 - y$. This function is far from being $(K_d, 0^{d-1}1^*)$-free — just consider any linear transformation L that sends $\mathbf{e}_i$ to $1|\mathbf{z}_i$. (The formal proof is the same as in Lemma 3.5, only much simpler.)

Now consider $(K_{d_1}, 0^{c_1}1^*)$. By assumption, $\mathbf{e}_1 \in K_{d_1}$ is 0-labeled and $\mathbf{e}_{d_1-1} \in K_{d_1}$ is 1-labeled, so in order to get $f(L(K_{d_1})) = \langle 0^{c_1}1^* \rangle$ we must have $L(\mathbf{e}_1) = 1|\mathbf{z}_1$ and $L(\mathbf{e}_{d_1-1}) = 0|\mathbf{z}_2$. But if so, $L(\mathbf{e}_1 + \mathbf{e}_{d_1-1}) = 1|(\mathbf{z}_1 + \mathbf{z}_2)$ which is a 0-labeled point. Hence, no linear tranformation can send K_{d_1} to a set of vectors exhibiting the pattern $\langle 0^{c_1}1^* \rangle$, so g is $(K_{d_1}, 0^{c_1}1^*)$-free. □

Note that as a special case, Lemma 5.6 establishes our claim that $\mathcal{K}_d[\neg 0^{d-4}1^*]$ is well separated from $\mathcal{K}_d[\neg 0^{d-1}1^*]$

At an intuitive level, it still seems that the reason that $\mathcal{K}_d[\neg 0^{d-4}1^*]$ and $\mathcal{K}_d[\neg 0^{d-1}1^*]$ are mutually well separated is that the matroids $\mathcal{K}_d[\neg 0^{d-4}1^*]$ and $\mathcal{K}_d[\neg 0^{d-1}1^*]$ do not map homomorphically into each other. But as we have seen we cannot prove this by constructing functions from which we can extract explicit homomorphisms from the assumption that they are not separated, so the canonical functions machinery we have developed is not powerful enough to show that the separation follows from the non-embedding.

6 Homomorphisms Do Not Determine Matroid Properties

One of the motivating questions of this work has been whether containment for non-monotone matroid freeness properties always is determined by the existence or non-existence of homomorphisms.

The examples in [BCSX09] as well as the examples exhibited in this paper provide a fairly broad class of matroid freeness properties for which the answer is "yes." We note that in one direction, it is clear that if (M, σ) embeds into (N, τ), then (M, σ)-freeness is contained in (N, τ)-freeness. Thus, the way to *disprove* this connection this would be to find two labelled matroids (M, σ) and (N, τ) such that on the one hand $(M, \sigma) \not\hookrightarrow (N, \tau)$ but on the other hand it still holds that (M, σ)-freeness is contained in (N, τ)-freeness. In this section, we exhibit such pairs of labelled matroids. More precisely, our result is as follows.

Theorem 6.1. *It is* not *the case that* (M, σ)*-freeness is contained in* (N, τ)*-freeness if and only if* (M, σ) *embeds into* (N, τ)*. In particular, for* $d \ge 5$ *and* $1 < c < d - 1$ *it holds that* $(K_d, 01^*) \not\hookrightarrow (K_d, 0^c1^*)$ *but nevertheless we have* $\mathcal{K}_d[\neg 01^*] \subseteq \mathcal{K}_d[\neg 0^c1^*]$.

Proof. By Lemma Lemma 4.4 $(K_d, 01^*) \not\hookrightarrow (K_d, 0^c1^*)$. We will now show that if $f : \mathbb{F}_2^n \to \{0, 1\}$ is a function that is not $(K_d, 0^c1^*)$-free, then f is not $(K_d, 01^*)$-free either. This proves that $\mathcal{K}_d[\neg 01^*] \subseteq \mathcal{K}_d[\neg 0^c1^*]$.

Thus, suppose that f contains $(K_d, 0^c1^*)$ at the linear transformation $L : K_d \to \mathbb{F}_2^n$ That is, we have $f(L(\mathbf{e}_i)) = 0$ for $1 \le i \le c$, $f(L(\mathbf{e}_j)) = 1$ for $c < j < d$, and $f(L(\mathbf{e}_i + \mathbf{e}_j)) = f(L(\mathbf{e}_i) + L(\mathbf{e}_j)) = 1$ for $1 \le i < j < d$. We divide the proof into two cases depending on the value of f at the all-zero vector.

1. $f(\mathbf{0}_n) = 0$: Then if we let the linear transformation $L_0^* : K_d \to \mathbb{F}_2^n$ be defined by

$$
\begin{aligned}
L_0^*(\mathbf{e}_1) &= \mathbf{0} \\
L_0^*(\mathbf{e}_2) &= L(\mathbf{e}_1 + \mathbf{e}_2) \\
L_0^*(\mathbf{e}_3) &= L(\mathbf{e}_1 + \mathbf{e}_3) \\
&\vdots \\
L_0^*(\mathbf{e}_{d-1}) &= L(\mathbf{e}_1 + \mathbf{e}_{d-1})
\end{aligned}
\tag{6.1}
$$

   we have that $f(L_0^*(\mathbf{e}_1)) = f(\mathbf{0}_n) = 0$, for $2 \leq j < d$ that $f(L_0^*(\mathbf{e}_j)) = f(L(\mathbf{e}_1 + \mathbf{e}_j)) = 1$ and $f(L_0^*(\mathbf{e}_1 + \mathbf{e}_j)) = f(L_0^*(\mathbf{e}_1) + L_0^*(\mathbf{e}_j)) = f(\mathbf{0}_n + L(\mathbf{e}_1 + \mathbf{e}_j)) = 1$, and for $2 \leq i < j < d$ that $f(L_0^*(\mathbf{e}_i + \mathbf{e}_j)) = f(L(\mathbf{e}_1 + \mathbf{e}_i) + L(\mathbf{e}_1 + \mathbf{e}_j)) = f(L(\mathbf{e}_i + \mathbf{e}_j)) = 1$. In other words, f contains $(K_d, 01^*)$ at L_0^*.

2. $f(\mathbf{0}_n) = 1$: In this case let $L_1^* : K_d \to \mathbb{F}_2^n$ be defined by

$$
\begin{aligned}
L_1^*(\mathbf{e}_1) &= L(\mathbf{e}_1) && \text{and} \\
L_1^*(\mathbf{e}_j) &= L(\mathbf{e}_{d-1}) && \text{for all } 2 \leq j < d.
\end{aligned}
\tag{6.2}
$$

   By assumption, we have $f(L(K_d)) = \langle 0^c 1^* \rangle$. Since the vector $\mathbf{e}_1$ is 0-labelled by $(K_d, 0^c 1^*)$ and the vector $\mathbf{e}_{d-1}$ is 1-labelled, this means that $f(L(\mathbf{e}_1)) = 0$ and $f(L(\mathbf{e}_{d-1})) = 1$. From this we conclude that it holds that $f(L_1^*(\mathbf{e}_1)) = 0$ and for $2 \leq j < d$ that $f(L_1^*(\mathbf{e}_j)) = 1$ and $f(L_1^*(\mathbf{e}_1 + \mathbf{e}_j)) = f(L(\mathbf{e}_1 + \mathbf{e}_{d-1})) = 1$. Finally, for $2 \leq i < j < d$ we have $f(L_1^*(\mathbf{e}_i + \mathbf{e}_j)) = f(L(\mathbf{e}_{d-1}) + L(\mathbf{e}_{d-1})) = f(\mathbf{0}_n) = 1$. We have shown that $f(L_1^*(K_d)) = \langle 01^* \rangle$, so f contains $(K_d, 01^*)$ at L_1^*.

Thus, in both cases we can find a linear transformation at which f contains $(K_d, 01^*)$, and the theorem follows. □

The proof of Theorem 6.1 motivates the following definition.

Definition 6.2 (Zero-augmented matroid). Let $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\} \subseteq \mathbb{F}^r$ be any linear matroid. Then we define the *degenerate* or *zero-augmented version* $M^{\mathbf{0}}$ of M to be the set of vectors $M^{\mathbf{0}} = \{\mathbf{0}_r, \mathbf{v}_1, \ldots, \mathbf{v}_k\}$.

Note that $M^{\mathbf{0}}$ is not a matroid in the ordinary sense since it contains the all-zero vector.

It is easy to see that for any labelled matroid (M, σ) the property of (M, σ)-freeness can be written as $\mathcal{M}[\neg\sigma] = \bigcap_{b \in \mathsf{R}} \mathcal{M}^{\mathbf{0}}[\neg b\sigma]$. This is simply saying that any function f that is (M, σ)-free must be so regardless of the value of f at $\mathbf{0}$. The concept of matroid homomorphisms extends to zero-augmented matroids without any modifications. Using this generalized definition it is also straightforward to show that if for two labelled matroids (M, σ) and (N, τ) it is the case for all $b \in \mathsf{R}$ that $(M, \sigma) \hookrightarrow (N^{\mathbf{0}}, b\tau)$, then it holds that $\mathcal{M}[\neg\sigma] \subseteq \mathcal{N}[\neg\tau]$.

In fact, this is exactly what is going on in the proof of Theorem 6.1. The two cases in the case analysis establish that $(K_d, 01^*) \hookrightarrow (K_d^{\mathbf{0}}, 00^c 1^*)$ and $(K_d, 01^*) \hookrightarrow (K_d^{\mathbf{0}}, 10^c 1^*)$, respectively, from which it follows that $\mathcal{K}_d[\neg 01^*] \subseteq \bigcap_{b \in \{0,1\}} \mathcal{K}_d^{\mathbf{0}}[\neg b 0^c 1^*] = \mathcal{K}_d[\neg 0^c 1^*]$. Note that this also shows that while it is clearly true that if $(M, \sigma) \hookrightarrow (N, \tau)$, then we must also (trivially) have $(M, \sigma) \hookrightarrow (N^{\mathbf{0}}, b\tau)$, for all $b \in \mathsf{R}$, the opposite direction does not hold. That is, it can very well be the case that $(M, \sigma) \hookrightarrow (N^{\mathbf{0}}, b\tau)$, for all $b \in \mathsf{R}$, while it still holds that $(M, \sigma) \not\hookrightarrow (N, \tau)$.

We conclude from this discussion that while Theorem 6.1 clearly answers the question whether it is true that (M, σ)-freeness is contained in (N, τ)-freeness if and only if (M, σ) embeds into (N, τ) in the negative, a more important contribution of this theorem is perhaps that it reveals that it was arguably not quite the right question to ask in the first place. Just focusing exclusively on the vectors in the matroids (M, σ) embeds into (N, τ) ignores the important role that the $\mathbf{0}$-vector can play (as is evident in the proof of Theorem 6.1). Therefore, a more refined version of the question is as follows.

Open Problem 1. *Is it true or not that (M, σ)-freeness is contained in (N, τ)-freeness if and only if it holds for all $b \in \mathsf{R}$ that (M, σ) embeds into $(N^{\mathbf{0}}, b\tau)$?*

We remark that, as noted above, the "if" direction of this statement clearly holds, so the interesting question is what can be said about the "only if" direction. Again, the way to prove that the answer is "no" would be to find two labelled matroids (M, σ) and (N, τ) such that $(M, \sigma) \not\rightarrow (N^{\mathbf{0}}, b\tau)$ for some $b \in \mathsf{R}$ but nevertheless $\mathcal{M}[\neg\sigma] \subseteq \mathcal{N}[\neg\tau]$. The main focus of this paper is on $\mathbb{F} = \mathbb{F}_2$ and $\mathsf{R} = \{0, 1\}$ and this is the case for which we would mainly want to see Open Problem 1 resolved, but any results for other $\mathbb{F}$ or R would also be of interest.

7 Concluding Remarks

Motivated by questions raised in [BCSX09] and the recent testability results in [BGS10], in this paper we have studied the semantics of matroid freeness properties, and in particular the problem of determining when two syntactically different matroid constraints in fact also encode *semantically* different properties. We have developed a new method for comparing matroid freeness constraints based on the concept of *labeled matroid homomorphisms*, and have shown that for a surprisingly broad class of matroid freeness properties this method exactly characterizes the relation between two matroid freeness properties. Even more, when the method works, it in fact establishes a strong dichotomy in the sense that either one property must be contained in the other or the properties are strictly distinct in a property testing sense. As a consequence, we established that results in [BGS10] do indeed provide infinite hierarchies of new properties not known to have been testable before.

Our work raises many interesting questions which we believe merit further study.

- Perhaps the most obvious open problem is in what generality our method of characterizing matroid freeness properties in terms of labeled homomorphisms can be made to work, and in particular whether it can be extended to arbitrary labeled graphic matroids, or even arbitrary linear matroids. That is, is it always true that (M, σ)-freeness is contained in (N, τ)-freeness if and only if there is a labeled matroid homomorphism from (M, σ) to (N, τ), and that the two properties must be well separated otherwise? In fact, in view of the results in Section 6, we have refined this question as explained in Open Problem 1. We prove that for fairly broad classes of matroid freeness properties the answer to the refined question is "yes."

  As was explained in Section 2, complete graphic matroids $(M(K_d), \sigma)$ can be seen to be building blocks for all labeled graph matroid freeness properties. Thus, a first step towards the resolution of this question might be to understand $(M(K_d), \sigma)$-freeness for any pattern σ, and then study how intersections of such properties behave.

- Leaving aside the issue of labeled matroid homomorphisms, another fundamental open problem is whether there must always hold a dichotomy between containment and δ-separation for matroid freeness properties. If this would turn out to be the case, the next question is whether such a dichotomy would extend even further to arbitrary linear-invariant properties or even affine invariant properties.

- Turning to some of the specifics of our results, our Theorem 5.5 proves separations between $(K_d, 0^c 1^*)$-freeness properties when $c < d$. An immediate question is whether this can be extended to the case where $c \geq d$ as well. Perhaps c up to $2d - 3$ should be doable by (modifications of) our current techniques, since then one still retains fairly good control over where the 0-labelled basis vectors can map.

- We understand single matroid freeness properties a little bit now. Also, we understand the other extreme of an intersection of very many matroid freeness properties well in certain cases - recall from Section 2 that polynomials of degree d (with constant term zero), can be specified as the intersection of all $(F_d, \sigma\tau)$-freeness properties, where d is fixed and σ ranges over all patterns in $\{0,1\}^{2^d-1}$ of odd parity. But can we say anything about properties "in between" these two extremes? What happens if we take, say, $d = 2$ and forbid patterns of weight $1, 3, 7$ but allow weight 5? Clearly we get a potentially larger property, but which property is it? Is it another Reed-Muller code, or is it something else?

- Of course, deciding the testability of (M, σ)-freeness for matroids M of any complexity is a huge open question (if admittedly somewhat orthogonal to the theme of the current paper).

- A special case: Can we develop techniques to understand the semantics and/or testability of properties of set operations (say $\mathcal{M}[\neg\sigma] \cap \mathcal{N}[\neg\tau]$) for two matroids (M, σ) and (N, τ)?

We conclude this section with a discussion on the first open problem mentioned above and some further connections with testing graph properties. The core of this problem is determining when (M, σ)-freeness and (N, τ)-freeness are identical properties for two syntactically different labeled matroids (M, σ) and (N, τ). If we look at graphic matroids, one observation is that blowing up the underlying graph does not change the property. Formally, for a graph G and for a positive integer t, define the *order-t blowup* of G to be the graph $G^{(t)}$ obtained by replacing each vertex of G by an independent set of size t and each edge in G by the complete bipartite graph $K_{t,t}$. Furthermore, if an edge of G is labeled by an element of $\{0, 1\}$, then use that label for all edges in the associated complete bipartite graph in the blowup graph; if σ was the original labeling for the edges, we call the new labeling for the blowup graph $\sigma^{(t)}$.

The fact that graph blow-ups preserve matroid freeness properties, which we write down as Proposition 7.1 below, is similar in flavor to the Erdős-Stone theorem from extremal graph theory. The Erdős-Stone theorem essentially says that for any graph G and integer $t \geq 1$, G-freeeness (i.e., not containing G as an induced subgraph) and $G^{(t)}$-freeness are not δ-separated for any constant $\delta > 0$; see, for example, [Die05]. However, the proof of the analogous statement in the matroid freeness case turns out to be much simpler, because a matroid homomorphism is not required to be injective whereas the subgraph relationship in graphs is an injection.

Proposition 7.1. *Given a graph G on m edges and a string $\sigma \in \{0,1\}^m$, suppose H is a subgraph of $G^{(t)}$ for some $t \geq 1$ that contains at least one copy of G. Also, suppose τ is the restriction of $\sigma^{(t)}$ to the edges of H. Then, $(M(G), \sigma)$-freeness is identical to $(M(H), \tau)$-freeness.*

The proof is straightforward. The fact that $(M(G), \sigma)$-freeness is contained in $(M(H), \tau)$-freeness follows from Corollary 3.2. The other direction holds because the map that takes each edge of H to the edge of G from where it originated is a labeled matroid homomorphism from $(M(H), \tau)$ to $(M(G), \sigma)$, and so we can apply Lemma 3.1.

It should be noted that Proposition 7.1 is *not* a characterization of equality even for monotone graphic matroid freeness properties. For instance, while K_4 is easily seen not to be a subgraph of any blowup of K_3, it nevertheless holds that $(M(K_3), 1^*)$-freeness and $(M(K_4), 1^*)$-freeness are identical properties as shown in Proposition 3.3. What our dichotomy theorems in Section 3.3 establish is that for all monotone properties and a nontrivial subclass of non-monotone properties, equality of properties corresponds exactly to existence of matroid homomorphisms in both directions. As noted above, the question whether such a correspondence holds in general for any non-monotone matroid freeness properties remains wide open.

Acknowledgments

We are grateful to Madhu Sudan for suggesting that we investigate the problems studied in this work, for giving many helpful insights and comments along the way, and for challenging us to redouble our efforts in order to strengthen our first, preliminary results. We also want to thank Asaf Shapira for his collaboration in the early stages of this research, as well as for his continuing useful advice.

References

[AFNS06] Noga Alon, Eldar Fischer, Ilan Newman, and Asaf Shapira. A combinatorial characterization of the testable graph properties: it's all about regularity. In *STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 251–260, 2006.

[AKK+05] Noga Alon, Tali Kaufman, Michael Krivelevich, Simon Litsyn, and Dana Ron. Testing Reed-Muller codes. *IEEE Transactions on Information Theory*, 51(11):4032–4039, 2005.

[AS06] Noga Alon and Asaf Shapira. Homomorphisms in Graph Property Testing. *Topics in Discrete Mathematics*, pages 281–313, 2006.

[BCSX09] Arnab Bhattacharyya, Victor Chen, Madhu Sudan, and Ning Xie. Testing linear-invariant non-linear properties. In *Symposium on Theoretical Aspects of Computer Science*, pages 135–146, 2009.

[BCSX10] Arnab Bhattacharyya, Victor Chen, Madhu Sudan, and Ning Xie. Testing linear-invariant non-linear properties: A short report. Technical Report 10-116, Electronic Colloquium in Computational Complexity, July 2010.

[BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1(1):3–40, 1991.

[BGNX10] Arnab Bhattacharyya, Elena Grigorescu, Jakob Nordström, and Ning Xie. Separations of matroid freeness properties. Technical Report 10-136, Electronic Colloquium in Computational Complexity, August 2010.

[BGS10] Arnab Bhattacharyya, Elena Grigorescu, and Asaf Shapira. A unified framework for testing linear-invariant properties. In *Proc. 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 478–487, 2010.

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comp. Sys. Sci.*, 47:549–595, 1993. Earlier version in STOC'90.

[CSX11] Victor Chen, Madhu Sudan, and Ning Xie. Property testing via set-theoretic operations. In *Proc. 2nd Innovations in Computer Science*, pages 211–222, 2011. Preprint available at http://www.eccc.uni-trier.de/report/2010/156/.

[Die05] Reinhard Diestel. *Graph Theory (Graduate Texts in Mathematics)*. Springer, August 2005.

[GGR98] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45:653–750, 1998.

[GK10] Oded Goldreich and Tali Kaufman. Proximity oblivious testing and the role of invariances. *Electronic Colloquium in Computational Complexity*, 10-058, April 2010.

[GOS+09] Parikshit Gopalan, Ryan O'Donnell, Rocco A. Servedio, Amir Shpilka, and Karl Wimmer. Testing Fourier dimensionality and sparsity. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP '09)*, pages 500–512, 2009.

[GR09] Oded Goldreich and Dana Ron. On proximity oblivious testing. In *Proc. 41st Annual ACM Symposium on the Theory of Computing*, pages 141–150, 2009.

[Gre05] Ben Green. A Szemerédi-type regularity lemma in abelian groups, with applications. *Geom. Funct. Anal.*, 15(2):340–376, 2005.

[KS08] Tali Kaufman and Madhu Sudan. Algebraic property testing: the role of invariance. In *Proc. 40th Annual ACM Symposium on the Theory of Computing*, pages 403–412, 2008.

[KSV10] Daniel Král', Oriol Serra, and Lluís Vena. A removal lemma for systems of linear equations over finite fields. *Israel Journal of Mathematics (to appear)*, 2010. Preprint available at http://arxiv.org/abs/0809.1846.

[Oxl03] James Oxley. What is a matroid? Technical Report 2002-9, Louisiana State University Mathematics Electronic Preprint Series, 2003. Revised version available at http://www.math.lsu.edu/~oxley/.

[Ron09] Dana Ron. Algorithmic and analysis techniques in property testing. *Foundations and Trends in Theoretical Computer Science*, 5(2):73–205, 2009.

[RS96] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM J. on Comput.*, 25:252–271, 1996.

[Rub06] Ronitt Rubinfeld. Sublinear time algorithms. In *Proceedings of International Congress of Mathematicians 2006*, volume 3, pages 1095–1110, 2006.

[Sha09] Asaf Shapira. Green's conjecture and testing linear-invariant properties. In *Proc. 41st Annual ACM Symposium on the Theory of Computing*, pages 159–166, 2009.

[Sud10] Madhu Sudan. Invariance in property testing. Technical Report 10-051, Electronic Colloquium in Computational Complexity, March 2010.

[Wil73] Robin J. Wilson. An introduction to matroid theory. *American Mathematical Monthly*, 80:500–525, 1973.

A Matroids, Matroid Freeness and Systems of Linear Equations

Let us start this appendix by presenting a formal definition of what a matroid is. There are many equivalent ways of defining matroids, and for some of the different formulations it is in fact nontrivial to show that they are equivalent. We will use the definition presented next, and refer the reader to, for instance, [Oxl03, Wil73] for more background on matroid theory.

Definition A.1 (Matroid). A *matroid* M is a finite set S, along with a set $\mathcal{I}$ of subsets of S, such that:

1. The empty set is in $\mathcal{I}$.

2. If X is in $\mathcal{I}$, then every subset of X is also in $\mathcal{I}$.

3. If X and Y are both in $\mathcal{I}$ and $|X| = |Y| + 1$, then there exists an element $x \in X \setminus Y$ such that $Y \cup \{x\}$ is in $\mathcal{I}$.

The set S is called the *ground set* of M, and the set $\mathcal{I}$ is the collection of *independent sets* of M. Those subsets of S which are not in $\mathcal{I}$ are called *dependent*. A maximal independent set — that is, an independent set X which becomes dependent on adding any element of S — is called a *basis* for the matroid. It is a basic result of matroid theory that any two bases of a matroid M must have the same number of elements. This number is called the *rank* of M.

Two important classes of matroids, which we briefly discuss next, are *linear matroids* and *graphic matroids*.

We say that a matroid M on a ground set $S = \{x_1, \ldots, x_k\}$ is a *linear matroid*, or *vector matroid*, if there is a field $\mathbb{F}$ and vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ in $\mathbb{F}^k$ such that any subset $\{x_i \mid i \in T\}$ indexed by $T \subseteq [k]$ is independent if and only if the corresponding vectors $\{\mathbf{v}_i \mid i \in T\}$ form a linearly independent set. A matroid is *binary* if it is linear with $\mathbb{F} = \mathbb{F}_2$. Note that when we are interested in property testing of linear-invariant properties, the only matroid freeness properties that really make sense to consider are those of linear matroids.

Given an undirected graph G, we can let S be the set of edges $E(G)$ of G and $\mathcal{I}$ consist of the subsets of $S = E(G)$ that do not contain any cycles in the graph. Then $M = (S, \mathcal{I})$ can be shown to be a matroid, which we refer to as the *graphic matroid* $M(G)$ over G. Any graphic matroid $M(G)$ can be represented as a binary matroid. One way of seeing this is to consider the incidence matrix of G and let $\mathbf{v}_i$ be the rows corresponding to the edges. Then any cycle in G will correspond to a (subset of) vectors summing to zero. Another possibility is to fix any spanning tree of G and let the edges $e_1, e_2, \ldots$ in this spanning tree T correspond to unit vectors $\mathbf{e}_1, \mathbf{e}_2, \ldots$. Then any edge e not in the spanning tree T will correspond to the sum of the vectors for the unique minimal set of edges in T that together with e yields a cycle.

As the reader can see, in this paper we used the latter approach with spanning trees emanating from a single, unique vertex to get our standard representation for $M(K_d)$ (Definition 2.7). Another possibility would have been to use the incidence matrix representation. Note that this would have given a very nice and symmetric representation with all vectors in $M(K_d)$ having Hamming weight 2, and with a basis corresponding to fixing some coordinate j and requiring that all the basis vectors have a 1 in this coordinate. Of course, our standard representation is just taking such a basis and "puncturing" it by deleting this j^{th} coordinate from all vectors. For our purposes, it turns out to be convenient to have a representation where all basis vectors and weight 1 and all other vectors had weight 2. However, we want to point out that this is not the only way of thinking about $M(K_d)$, and that it might be interesting when trying to generalize our results to investigate whether the representation with all vectors of uniform weight 2 might be a more fruitful way of looking at $M(K_d)$.

Let us now explain how one can see that the matroid-freeness representation of properties employed in [BCSX09] and the current paper on the one hand, and the system of linear equations representation of properties used in [KSV10, Sha09, BGS10] on the other, are essentially equivalent. This equivalence is in some sense folklore knowledge, but since it is a priori not entirely obvious, and since we have not seen it actually written down anywhere, we give an explicit exposition of the correspondence here for completeness. Note, however, that we will only discuss monotone properties below. Non-monotone properties have also been formulated using systems of linear equations, in this context most notably in [BGS10], but since the notation is a bit heavier and the ideas are essentially the same, we ignore the issue of non-monotonicity in this appendix for the sake of simplicity.

The *system of linear equations representation* is the following. Let $\mathbb{K}$ be a field. Let k and ℓ be fixed integers with $k < \ell$. Let $A\mathbf{x} = \mathbf{b}$ be a system of k linear equations in ℓ variables, where $A \in \mathbb{K}^{k \times \ell}$ and $\mathbf{b} \in \mathbb{K}^k$. We say that a set $S \subseteq \mathbb{K}$ is $(A, \mathbf{b})$-*free* if it contains no solution to $A\mathbf{x} = \mathbf{b}$; that is, S is $(A, \mathbf{b})$-free if there is no vector $\mathbf{x} \in S^\ell$ that satisfies all of the k equations in $A\mathbf{x} = \mathbf{b}$.

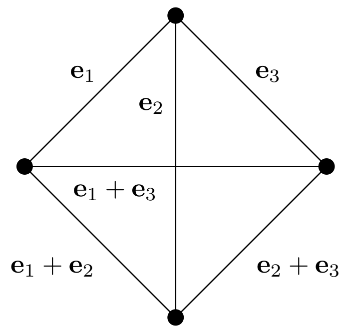

(a) $M(K_4) = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_1{+}\mathbf{e}_2, \mathbf{e}_1{+}\mathbf{e}_3, \mathbf{e}_2{+}\mathbf{e}_3\}$.

$$\begin{bmatrix} 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 \end{bmatrix}$$

(b) Matrix encoding $\mathbf{v}_4 = \mathbf{e}_1 + \mathbf{e}_2$, $\mathbf{v}_5 = \mathbf{e}_1 + \mathbf{e}_3$, and $\mathbf{v}_6 = \mathbf{e}_2 + \mathbf{e}_3$.

Figure 4: The binary graphic matroid $M(K_4)$ and the corresponding linear equation system matrix.

When considered as a property testing problem, we usually first pick a finite field $\mathbb{F}$ and then take $\mathbb{K} = \mathbb{F}^n$. For properties that arise naturally in mathematics and computer science, it is usually the case that the property can be specified uniformly for all n using a finite description. Thus, it is of particular interest to consider the case when A and $\mathbf{b}$ have finite descriptions. Thus, A is usually taken as having entries over $\mathbb{F}$, not $\mathbb{K}$. Furthermore, in order for the properties to be linear-invariant we take $\mathbf{b} = \mathbf{0}$. (We note, however, that the results in [KSV10, Sha09] also hold when we can pick A (non-uniformly) as any matrix in $\mathbb{K}^{k\times\ell}$ and for any field $\mathbb{K}$, not just $\mathbb{K} = \mathbb{F}^n$, and that in these papers we can also have $\mathbf{b} \neq \mathbf{0}$.)

As a first simple example, the system of linear equations representation corresponding to $(M(C_\ell), 1^*)$-freeness, where C_ℓ is the cycle of length ℓ, consists of just a single linear equation $\sum_{i=1}^{\ell} x_i = 0$. Therefore we have $A = [1\ 1\ 1\ \cdots\ 1]$ and $\mathbf{b} = \mathbf{0}$, encoding that the sum of ℓ vectors is zero. Another simple, but less trivial, example is that of $(K_4, 1^*)$-freeness. Figure 4 shows the graphic matroid $M(K_4)$ and the matrix A its corresponding representation as a system of linear equations.

Let us now consider $(M, 1^*)$-freeness for a general linear matroid $M = \{\mathbf{v}_1, \ldots, \mathbf{v}_{\ell-k}, \mathbf{v}_{\ell-k+1}, \ldots, \mathbf{v}_\ell\}$, where $\{\mathbf{v}_1, \ldots, \mathbf{v}_{\ell-k}\}$ form a basis for the matroid and each of the vectors in $\{\mathbf{v}_{\ell-k+1}, \ldots, \mathbf{v}_\ell\}$ can be written as a linear combination of the first $\ell - k$ vectors $\mathbf{v}_{\ell-k+i} = -\sum_{j=1}^{\ell-k} B_{ij}\mathbf{v}_j$ with coefficients $B_{i,j} \in \mathbb{F}$. Without loss of generality, we can think of the first $\ell - k$ vectors as being $\mathbf{v}_1 = \mathbf{e}_1, \ldots \mathbf{v}_{\ell-k} = \mathbf{e}_{\ell-k}$. To transform $(M, 1^*)$-freeness into a system of linear equations representation, we construct a matrix $A \in \mathbb{F}^{k\times\ell}$ in which the i^{th} row consists of the coefficients of the linear equations describing $\mathbf{v}_{\ell-k+i}$. Specifically, for every $1 \le i \le k$ and $1 \le j \le \ell$, $A_{ij} = B_{ij}$ when $1 \le j \le \ell - k$; $A_{ij} = 1$ if $j = \ell - k + i$ and $A_{ij} = 0$ otherwise.

To go in the other direction and transform a system of linear equations into matroid freeness representation, we may assume without loss of generality that the matrix A has rank k (otherwise we may delete the redundant rows). Then by permuting the columns of A and appropriately changing the basis of $\mathbb{K}$, we can transform A into the form $[B|I_k]$, where I_k is a k-by-k identity matrix and B is a k-by-$(\ell - k)$ matrix. Now $B = \{B_{ij}\}$ is exactly the matrix we defined above which contains the coefficients of k linear combinations of the non-basis vectors in the matroid in terms of the $\ell - k$ basis vectors, so it is immediate to recover the matroid freeness representation from this matrix.

B Proof of Second Hierarchy Theorem

In this section we present the proof of Theorem 5.5 on page 24.

B.1 Proof of item 1

Lemma B.1. *Let $d \geq 5$. Then for all $0 \leq c_1 < c_2 < d$ it holds that $\mathcal{K}_d[\neg 0^{c_2}1^*] \nsubseteq_\delta \mathcal{K}_d[\neg 0^{c_1}1^*]$.*

Proof. We need to show that there is a Boolean function g which is $(K_d, 0^{c_2}1^*)$-free and is δ-far from being $(K_d, 0^{c_1}1^*)$-free. If $c_2 = 1$, then $c_1 = 0$. In this case, the all-one function on $d-1$ bits is clearly δ-far from being $(K_d, 1^*)$-free but is $(K_d, 01^*)$-free.

Now let $c_2 \geq 2$. Consider the canonical function $f^1_{(K_d, 0^{c_1}1^*)}$, which is clearly δ-far from being $(K_d, 0^{c_1}1^*)$-free for some constant δ. Suppose $f^1_{(K_d, 0^{c_1}1^*)}$ contains $(K_d, 0^{c_2}1^*)$, then it follows immediately that $f^1_{(K_d, 0^{c_1}1^*)}$ also contains $(K_{c_2+1}, 0^{c_2}1^*)$. Note that all the 0-labeled basis vectors in $(K_{c_2+1}, 0^{c_2}1^*)$ must map to pairwise distinct vectors, since otherwise their sum would map to $\mathbf{0}$, which is labeled 0. However, there simply aren't enough 0-labeled vectors in $(K_d, 0^{c_1}1^*)$ for such a mapping, we thus reach a contradiction. That is, the canonical function $f^1_{(K_d, 0^{c_1}1^*)}$ is $(K_d, 0^{c_2}1^*)$-free and hence the theorem is proved. □

B.2 Proof of item 2: general case

We will need to do quite a lot of work before the proofs of the theorems above are completed. Some arguments will recur in slightly different forms throughout the proofs, so we want to make a general definition here and prove a proposition that will take care of all of these different but very similar cases.

Definition B.2 (Zero vector indicator function). Let us say that a function $f : \mathbb{F}_2^n \to \mathsf{R}$ is a *zero vector indicator function* if we can write it as $f(\mathbf{x}) = f(\mathbf{y}|\mathbf{z})$ for $\mathbf{y} \in \mathbb{F}_2^m$ and $\mathbf{z} \in \mathbb{F}_2^{n-m}$ so that it holds that $f(\mathbf{0}|\mathbf{z}) = 0$ for any $\mathbf{z}$.

For instance, the canonical function in Definition 3.4 is a zero vector indicator function if $S_\mathbf{0} = \mathbb{F}_2^{n-K}$ or if the padding value is $b = 0$.

Proposition B.3. *Let (M, σ) be a labeled matroid containing $F_d^{\leq 2} = \{\mathbf{e}_i, \mathbf{e}_i + \mathbf{e}_j \mid 1 \leq i, j \leq d, i \neq j\}$ as a submatroid with all $\mathbf{e}_i$ labeled by 0 and all $\mathbf{e}_i + \mathbf{e}_j$ labeled by 1. Suppose that $f : \mathbb{F}_2^n \to \mathsf{R}$ is a zero vector indicator function containing (M, σ) at the linear map L, and for every $i \in [d]$ we write $L(\mathbf{e}_i) = \mathbf{y}_i|\mathbf{z}_i$. Then we must have $\mathbf{y}_i \neq \mathbf{y}_j$ for all $i \neq j$.*

Proof. This is true because otherwise the linear transformation L would map some sum $\mathbf{e}_i + \mathbf{e}_j$ to the zero vector, which is 0-labeled in f, a contradiction. □

Let us now state the main lemma to be proven in this section.

Lemma B.4. *Let $d \geq 5$. Then $\mathcal{K}_d[\neg 0^{c_1}1^*] \nsubseteq_\delta \mathcal{K}_d[\neg 0^{c_2}1^*]$ for $0 \leq c_1 < c_2 < d$ provided that $c_1 \notin \{1, d-4\}$.*

We will spend the rest of this section proving this lemma. As will become clear form the proof, the perhaps seemingly somewhat arbitrary restriction $c_1 \notin \{1, d-4\}$ will be crucial for our argument.

In what follows, we will consistently write $\mathbf{e}_1, \ldots, \mathbf{e}_{d-1}$ to denote the standard basis vectors for the matroid $(K_d, 0^{c_2}1^*)$ and $\mathbf{f}_1, \ldots, \mathbf{f}_{d-1}$ to denote the standard basis vectors for $(K_d, 0^{c_1}1^*)$. To separate the properties in Lemma B.4, we will consider canonical functions as in Definition 3.4. More specifically, we will use f_{sep} to denote the canonical function $f^0_{(K_d, 0^{c_2}1^*)}$, where all sets $S_\mathbf{v}$ are chosen to be the whole subspace $\mathbb{F}_2^{n-d+1}$.

We first note that f_{sep} chosen in this way will always be dense in the pattern $(K_d, 0^{c_2}1^*)$.

Proposition B.5. *The function f_{sep} specified above is δ-far from being $(K_d, 0^{c_2}1^*)$-free for some absolute constant δ.*

Proof. This follows immediately from Lemma 3.5. □

Our goal is now to show that f_{sep} is $(K_d, 0^{c_1}1^*)$-free. To prove this, we assume that f_{sep} is not $(K_d, 0^{c_1}1^*)$-free and derive a contradiction. We present the proof as a sequence of simple claims that follow from the assumption that there exists a linear transformation $L : \mathbb{F}_2^{d-1} \to \mathbb{F}_2^n$ such that f_{sep} contains $(K_d, 0^{c_1}1^*)$ at L.

Suppose L maps the basis vector $\mathbf{f}_i$ to $L(\mathbf{f}_i) = \mathbf{y}_i|\mathbf{z}_i$, where $\mathbf{y}_i \in \mathbb{F}_2^{d-1}$ and $\mathbf{z}_i \in \mathbb{F}_2^{n-d+1}$. Let us write $L'(\mathbf{f}_i) = \mathbf{y}_i$ to denote the projection of this map to the part of coordinates which we care about most.

Claim B.6. L' maps all the basis vectors $\mathbf{f}_i$, $1 \le i \le d-1$, to pairwise distinct vectors.

Proof. This is Proposition B.3. Note that we can apply this proposition since we are padding with $b = 0$ in the canonical function f_{sep}. □

Claim B.7. L' maps every 1-labeled basis vector $\mathbf{f}_i$, $c_1 < i \le d-1$, to a vector $\mathbf{y}_i$ with $|\mathbf{y}_i| = 2$.

Proof. Let us narrow down the possible options for the weights of the vectors $\mathbf{y}_i = L'(\mathbf{f}_i)$ step by step.

Firstly, by construction we have that the L'-part of the linear transformation must result in weights $|L'(\mathbf{f}_i)| \in \{1, 2\}$ for all i, $c_1 < i \le d-1$ since that is where the vectors are where f_{sep} can evaluate to 1.

A second easy observation is that at least one of the vectors $\mathbf{f}_i$, $c_1 < i \le d-1$, must be mapped by L' to some weight-2 vector. This follows since by construction there are not enough 1-labeled unit vectors $\mathbf{e}_j$ on which $f_{\text{sep}}(\mathbf{e}_j|\mathbf{z})$ can evaluate to 1 to satisfy the distinctness requirements in Claim B.6.

But this in turn implies that there are at most two vectors in $\{\mathbf{f}_i \mid c_1 < i \le d-1\}$ that can map to weight-1 vectors. To see this, let $\mathbf{f}_{c_1+1}$ be a basis vector such that that $L'(\mathbf{f}_{c_1+1}) = \mathbf{e}_i + \mathbf{e}_j$ (we just argued that there is at least one such basis vector) and suppose in order to derive a contradiction that there are three or more basis vectors being mapped by L' to (distinct) weight-1 vectors. If so, one of these vectors, say $\mathbf{f}_{c_1+2}$, must get mapped to $\mathbf{e}_\ell$ for some ℓ such that $i \ne \ell \ne j$. But then by linearity, $\mathbf{f}_{c_1+1} + \mathbf{f}_{c_1+2}$ would get mapped by L' to $\mathbf{e}_i + \mathbf{e}_j + \mathbf{e}_\ell$, which is a weight-3 vector and hence must corresponds to a point at which f_{sep} evaluates to 0. This is contrary to the assumption that f_{sep} contains $(K_d, 0^{c_1}1^*)$ at the linear transformation L.

So far, we have shown that there can be at most two basis vectors $\mathbf{f}_i$, $c_1 < i \le d-1$, that are mapped by L' to weight-1 vectors. Now we rule out the possibility that there is any such vector by a case analysis.

Case 1: Suppose there is exactly one 1-labelled basis vector $\mathbf{f}_i$, $c_1 < i \le d-1$, that maps to some weight-1 vector. Without loss of generality, say that $L'(\mathbf{f}_{c_1+1}) = \mathbf{e}_{d-1}$. Then for $c_1+2 \le i \le d-1$, all $L'(\mathbf{f}_i)$ must be weight-2 vectors. Moreover, we must have $d-1 \in \operatorname{supp}(L'(\mathbf{f}_i))$ for all the other 1-labelled basis vectors, as otherwise $L'(\mathbf{f}_{c_1+1} + \mathbf{f}_i)$ would be a weight-3 vector and f_{sep} would evaluate to 0 on $L(\mathbf{f}_{c_1+1} + \mathbf{f}_i)$ contrary to assumption. It follows that $L'(\mathbf{f}_i) = \mathbf{e}_{j_i} + \mathbf{e}_{d-1}$ for every i, $c_1+2 \le i \le d-1$, where the indices j_i are all distinct by Claim B.6. Moreover, it must hold that $j_i \in \{c_2+1, \ldots, d-2\}$ for all j_i since $L'(\mathbf{f}_{c_1+1} + \mathbf{f}_i)$ is a weight-1 vector but f_{sep} nevertheless evaluates to 1 on $L(\mathbf{f}_{c_1+1} + \mathbf{f}_i)$. But by the construction of f_{sep} there are strictly less than $d-2-c_1$ distinct choices for j_i while we have $d-2-c_1$ basis vectors to take care of (namely, $\mathbf{f}_{c_1+2}, \ldots, \mathbf{f}_{d-1}$), each requiring a distinct j_i. This is a contradiction.

Case 2: There are exactly two 1-labelled basis vectors that are mapped to weight-1 vectors by L'. Notice that this case cannot apply for $c_1 \ge d-3$ since we proved that there is at least one 1-labelled basis vector mapping to a weight-2 vector, so we have $c_1 < d-3$.

Let us postpone the case $c_1 = d-4$, and first assume $c_1 < d-4$. In this case there are at least four 1-labeled basis vector in $(K_d, 0^{c_1}1^*)$. Consequently, there are at least two basis vectors that map to some weight-2 vector. Now suppose without loss of generality that $L'(\mathbf{f}_{d-1}) = \mathbf{e}_{d-1}$ and $L'(\mathbf{f}_{d-2}) = \mathbf{e}_{d-2}$.

Let $\mathbf{y}_{d-3} = L'(\mathbf{f}_{d-3})$ and $\mathbf{y}_{d-4} = L'(\mathbf{f}_{d-4})$ be two weight-2 vectors to which L' maps $\mathbf{f}_{d-3}$ and $\mathbf{f}_{d-4}$. Note that the sum of a weight-1 and weight-2 vector has weight either 1 or 3. In order to make both $\mathbf{y}_{d-3} + \mathbf{e}_{d-1}$ and $\mathbf{y}_{d-3} + \mathbf{e}_{d-2}$ into weight-1 vectors, it must be the case that $\mathbf{y}_{d-3} = \mathbf{e}_{d-1} + \mathbf{e}_{d-2}$. However then, at least one of $\mathbf{y}_{d-4} + \mathbf{e}_{d-1}$ and $\mathbf{y}_{d-4} + \mathbf{e}_{d-2}$ would be a weight-3 vector, which is a contradiction.

Finally, we look at the case that $c_1 = d - 4$, i.e., when there are three 1-labelled basis vectors $\mathbf{f}_i$ and we have two basis vectors mapping to weight-1 vectors and one basis vector mapping to a weight-2 vector. Suppose without loss of generality that $|L'(\mathbf{f}_{d-1})| = 2$ and $|L'(\mathbf{f}_{d-2})| = |L'(\mathbf{f}_{d-3})| = 1$. The support of $L'(\mathbf{f}_{d-1})$ must intersect with the supports of both weight-1 vectors to which the other basis vectors are mapped, since otherwise the image of their sum would be a weight-3 vector. Therefore, we may without loss of generality write $L'(\mathbf{f}_{d-1}) = \mathbf{e}_{d-1} + \mathbf{e}_{d-2}$, $L'(\mathbf{f}_{d-2}) = \mathbf{e}_{d-1}$ and $L'(\mathbf{f}_{d-3}) = \mathbf{e}_{d-2}$. Note since $c_1 < c_2$, we infer that the only 1-labeled weight-1 vectors are $\mathbf{e}_{d-1}$ and $\mathbf{e}_{d-2}$.

Now for each $1 \leq i \leq d - 4$, write

$$L'(\mathbf{f}_i) = \sum_{j \in S_i} \mathbf{e}_j + \sum_{k \in T_i} \mathbf{e}_k, \tag{B.1}$$

where $S_i \subseteq \{d-2, d-1\}$ and $T_i \subseteq \{1, \ldots, d-3\}$. First, note that there is no i such that $S_i = \emptyset$. This is so since if $S_i = \emptyset$ for some i, then $T_i = \emptyset$ as required by the fact that $|L'(\mathbf{f}_i) + L'(\mathbf{f}_{d-1})| \leq 2$. But by Claim B.9, this is impossible.

Second, there is no $1 \leq i \leq d - 4$ such that $S_i = \{d - 1\}$ (or $S = \{d - 2\}$, by symmetry). This is because $|L'(\mathbf{f}_i) + L'(\mathbf{f}_{d-3})| \leq 2$ requires that $T_i = \emptyset$, but this in turn implies that $L'(\mathbf{f}_i)$ is is a vector where f_{sep} evaluates to 1, a contradiction. What about $S_i = \{d - 2, d - 1\}$? Since $\mathbf{0}$ is 0-labeled by f_{sep}, it must hold that $L'(\mathbf{f}_i) + L'(\mathbf{f}_{d-1}) \neq \mathbf{0}$ which in turn requires that $T_i \neq \emptyset$. Now the conditions $|L'(\mathbf{f}_i) + L'(\mathbf{f}_{d-2})| \leq 2$ and $|L'(\mathbf{f}_i) + L'(\mathbf{f}_{d-3})| \leq 2$ force $|T_i| = 1$. But then, $L'(\mathbf{f}_i) + L'(\mathbf{f}_{d-1}) = \mathbf{e}_\ell$ for some $1 \leq \ell \leq d - 3$, which is clearly a vector where f_{sep} has to evaluate to 0, as we showed before.

To conclude, there is at least $d - 4 \geq 1$ 0-labeled vectors but there is no place to map these vectors, thus we reach a contradiction. □

We point out that Claim B.7 means that in what follows, when dealing with the 1-labelled basis vectors $\mathbf{f}_i$ in $(K_d, 0^{c_1}1^*)$ and where they are mapped by L', we will not need to consider the unit vectors $\mathbf{e}_j$ on which $f_{\text{sep}}(\mathbf{e}_j|\mathbf{z}) = 1$. This is so since vectors from $(K_d, 0^{c_1}1^*)$ will never get mapped by L' to such vectors $\mathbf{e}_j$ by Claim B.7. Furthermore, the claim also yields the following conclusion.

Claim B.8. For all $c_1 < i < j \leq d - 1$, we have $|\text{supp}(L'(\mathbf{f}_i)) \cap \text{supp}(L'(\mathbf{f}_j))| = 1$. That is, the vector supports of every pair of 1-labeled basis vectors in $(K_d, 0^{c_1}1^*)$ under the linear map L' intersect in exactly one coordinate.

Proof. Suppose this is not the case, then the weights of the pairwise sums would be 4, but all weight-4 vectors correspond to points where f_{sep} evaluates to 0. □

Let us now turn to the 0-labeled basis vectors $\mathbf{f}_i$, $1 \leq i \leq c_1$, in $(K_d, 0^{c_1}1^*)$. We remark in advance that in order to prove the next claim, we will need to use our assumption $c_1 \neq 1$, i.e., that there are either zero or at least two 0-labeled basis vectors in $(K_d, 0^{c_1}1^*)$.

Claim B.9. For every $1 \leq i \leq c_1$, it holds that $L'(\mathbf{f}_i) \neq \mathbf{0}$.

Proof. If $c_1 = 0$ then the claim is vacuously true, and $c_1 = 1$ has been ruled out by our assumptions. Thus, let $c_1 \geq 2$ and suppose that there is some basis vector, say $\mathbf{f}_1$, such that $L'(\mathbf{f}_1) = \mathbf{0}$. By assumption we also have that $\mathbf{f}_2$ gets mapped by L' to some point $\mathbf{y}_2$ such that $f_{\text{sep}}(L'(\mathbf{f}_2)) = 0$. But by the definition of f_{sep}, this would imply that $f_{\text{sep}}(L'(\mathbf{f}_i + \mathbf{f}_j)) = f_{\text{sep}}(\mathbf{y}_1) = 0$, a contradiction. ☐

Our final claim needed to establish Lemma B.4 says that every pair of one 0-labelled and one 1-labelled vector in $(K_d, 0^{c_1}1^*)$ are mapped by L' to vectors with intersecting supports.

Claim B.10. For all $\mathbf{f}_i$, $1 \leq i \leq c_1$, and all $\mathbf{f}_j$, $c_1 < j \leq d-1$, it holds that the supports of $\mathbf{y}_i = L'(\mathbf{f}_i)$ and $\mathbf{y}_j = L'(\mathbf{f}_j)$ intersect with each other.

Proof. Recall that by Claim B.7, for each j, $c_1 < j \leq d-1$, we have $|\mathbf{y}_j| = 2$, and that by Claim B.9, for each i, $1 \leq i \leq c_1$, we have $|\mathbf{y}_i| \geq 1$. Therefore, if the supports of these two vectors would not intersect, this would imply $|\mathbf{y}_i + \mathbf{y}_j| \geq 3$. But all such vectors correspond to points where f_{sep} evaluates to 0, a contradiction. ☐

Summing up, we now know about the 1-labelled basis vectors $\mathbf{f}_i$, $c_1 < i \leq d-1$, in $(K_d, 0^{c_1}1^*)$ that for all $\mathbf{y}_i = L'(\mathbf{f}_i)$ we have $|\mathbf{y}_i| = 2$ and all pairwise sums $\mathbf{y}_i + \mathbf{y}_j = L'(\mathbf{f}_i + \mathbf{f}_j)$ for $c_1 < i < j < d$ satisfy $|\mathbf{y}_i + \mathbf{y}_j| = 2$. Let us use this to prove Lemma B.4 by a case analysis for the value of c_1.

Case 1 ($c_1 = d-2$)**:** Note that in this case we must have $c_2 = d-1$, i.e., all basis vectors in $(K_d, 0^{c_2}1^*)$ are labelled by 0 and $f_{\text{sep}}(\mathbf{f}_i|\mathbf{z}_i) = 0$ for $1 \leq i \leq d-1$ and any $\mathbf{z}_i$. Thus, $L'(\mathbf{f}_{d-1})$ must be a weight-2 vector if f_{sep} is to evaluate to 1 on $L(\mathbf{f}_{d-1})$, and without loss of generality (because of the symmetry of $(K_d, 0^{c_2}1^*)$) we can write $L'(\mathbf{f}_{d-1}) = 11|\mathbf{0}$.

We claim that no vector in $\{L'(\mathbf{f}_1), \ldots, L'(\mathbf{f}_{d-2})\}$ can have weight 1. To see this, observe that this would imply that the sum of such a vector with $L'(\mathbf{f}_{d-1})$ would be a vector of weight either 1 or 3, which in both cases corresponds to points at which f_{sep} evaluates to 0. Also, it is clear that no vector in $\{L'(\mathbf{f}_1), \ldots, L'(\mathbf{f}_{d-2})\}$ has weight 2, since $f_{\text{sep}}(\mathbf{y}|\mathbf{z}) = 1$ for all weight-2 vectors $\mathbf{y}$ and any $\mathbf{z}$. Finally, no $L'(\mathbf{f}_i)$ has weight 0. For if $|L'(\mathbf{f}_i)| = 0$ and $|L'(\mathbf{f}_j)| \geq 3$ for some $1 \leq i, j \leq d-2$ then the sum $\mathbf{f}_i + \mathbf{f}_j$ must get mapped by L' to a vector of weight at least 3, i.e., corresponding to a point where f_{sep} evaluates to 0 contrary to assumption. Therefore, all the vectors $\{L'(\mathbf{f}_1), \ldots, L'(\mathbf{f}_{d-2})\}$ have weights at least 3.

Notice that if for some $\mathbf{f}_i$, $1 \leq i \leq d-2$, we had $L'(\mathbf{f}_i) = \mathbf{e}_1 + \sum_{j \in S} \mathbf{e}_j$, with $S \subseteq \{3, \ldots, d-1\}$ and $|S| \geq 2$, then $L'(\mathbf{f}_i + \mathbf{f}_{d-1})$ would be a weight-3 vector. For the same reason, no vector is of the form $L'(\mathbf{f}_i) = \mathbf{e}_2 + \sum_{j \in S} \mathbf{e}_j$. So every vector in the set $\{L'(\mathbf{f}_1), \ldots, L'(\mathbf{f}_{d-2})\}$ can be written as $L'(\mathbf{f}_i) = \mathbf{e}_1 + \mathbf{e}_2 + \sum_{j \in S_i} \mathbf{e}_j$. But if L should send $\mathbf{f}_i + \mathbf{f}_{d-1}$ to a point where $f_{\text{sep}}(L(\mathbf{f}_i + \mathbf{f}_{d-1})) = 1$, then this requires that each S_i is of size exactly 2. And the condition $f_{\text{sep}}(L(\mathbf{f}_i + \mathbf{f}_j)) = 1$, requires $|S_i \cap S_j| = 2$ for every $1 \leq i < j \leq d-2$. Combining the two conditions together implies that all sets S_i are identical, but since we have $d - 2 > 2$ distinct S_i, this gives a contradiction.

Case 2 ($c_1 = d-3$)**:** By Claims B.7 and B.8, we can write the images of the two 1-labelled basis vectors in $(K_d, 0^{c_1}1^*)$, possibly after reordering the coordinates, as $L'(\mathbf{f}_{d-1}) = 110|\mathbf{0}$ and $L'(\mathbf{f}_{d-2}) = 101|\mathbf{0}$. Observe that since $c_1 = d-3$, we have $c_2 = d-2$ or $d-1$, i.e., there is at most one 1-labeled weight-1 vector in $(K_d, 0^{c_2}1^*)$.

For every $1 \leq i \leq d-3$, let us write

$$L'(\mathbf{f}_i) = \sum_{j \in S_i} \mathbf{e}_j + \sum_{k \in T_i} \mathbf{e}_k, \tag{B.2}$$

where $S_i \subseteq \{1,2,3\}$ and $T_i \subseteq \{4,\dots,d-1\}$ for every $1 \le i \le d-3$, and let us check all possible configurations for S_i. That is, we partition all the $d-3$ vectors $L'(\mathbf{f}_i)$, $1 \le i \le d-3$, into $2^3 = 8$ disjoint classes according to their support subsets S_i. By Claim B.10, we have $S_i \ne \emptyset$. That is, the class with $S_i = \emptyset$ is empty.

It must also hold that $S_i \ne \{2\}$ since otherwise $L'(\mathbf{f}_i + \mathbf{f}_{d-2})$ would be a vector of weight at least 3. For a symmetric reason we have $S_i \ne \{3\}$. If $S_i = \{2,3\}$, then we must have $T_i = \emptyset$ to ensure that $L'(\mathbf{f}_{d-1} + \mathbf{f}_i)$ is a weight-2 vector. However, this makes $f_{\text{sep}}(L(\mathbf{f}_i)) = 1$, a contradiction.

What if $S_i = \{1,2\}$ (or $S_i = \{1,3\}$)? Observe that then we have $L'(\mathbf{f}_{d-1} + \mathbf{f}_i) = \sum_{k \in T_i} \mathbf{e}_k$ and $L'(\mathbf{f}_{d-2} + \mathbf{f}_i) = \mathbf{e}_2 + \mathbf{e}_3 + \sum_{k \in T_i} \mathbf{e}_k$ by the linearity of L'. The canonical function should evaluate to 1 on both corresponding points. But the first point requires $|T_i| \ge 1$ while the second point needs $T_i = \emptyset$. Thus these configurations can be ruled out.

Suppose that $S_i = \{1\}$. Since we have at most one 1-labeled weight-1 vector in $(K_d, 0^{c_2}1^*)$, it follows that we cannot have $T_i = \emptyset$ for every $1 \le i \le d-3$ (since otherwise the sums of $L'(\mathbf{f}_i)$ with $L'(\mathbf{f}_{d-1})$ and $L'(\mathbf{f}_{d-2})$ would be two distinct weight-1 vectors corresponding to points where f_{sep} would evaluate to 1, which is contrary to the construction of this function). On the other hand, if $|T_i| \ge 2$ for some i, then $L'(\mathbf{f}_{d-1}) + L'(\mathbf{f}_i)$ would be a vector of weight at least 3, which is also impossible for the same reason. But if $|T_i| = 1$, then $L'(\mathbf{f}_i)$ would be a weight-2 vector which is labeled 1 in f_{sep}. Therefore, the class with $S_i = \{1\}$ must also be empty.

The last possibility for S_i is $S_i = \{1,2,3\}$. By the same reasoning above, since there is at most one weight-1 vector in $(K_d, 0^{c_2}1^*)$, we have $T_i \ne \emptyset$ for every $1 \le i \le d-3$. Once again, follow the same argument as for the case that $S_i = \{1\}$, we conclude that $|T_i| \le 1$ for some $1 \le i \le d-3$. Consequently, all T_i must be singletons, i.e., $T_i = \{j_i\}$ for some $4 \le j_i \le d-1$. Note that there are only $d-4$ distinct indices in the set $\{4,\dots,d-1\}$. Therefore, the class with $S_i = \{1,2,3\}$ contains at most $d-4$ vectors.

To conclude, we have $d-3$ distinct vectors to map but there are only $d-4$ places for them, a contradiction.

Case 3 ($c_1 \le d-5$): Since $c_1 \le d-5$, there are at least four 1-labeled basis vectors $\mathbf{f}_{d-1}, \mathbf{f}_{d-2}, \mathbf{f}_{d-3}, \mathbf{f}_{d-4}$ in $(K_d, 0^{c_1}1^*)$. By Claims B.7 and B.8, we may write $L'(\mathbf{f}_{d-1}) = 1100|\mathbf{0}$ and $L'(\mathbf{f}_{d-2}) = 1010|\mathbf{0}$ without loss of generality (possibly after reordering coordinates). By symmetry, there are two possibilities for $L'(\mathbf{f}_{d-3})$: $0110|\mathbf{0}$ and $1001|\mathbf{0}$. But if we take $L'(\mathbf{f}_{d-3}) = 0110|\mathbf{0}$, then there is nowhere L' can map $\mathbf{f}_{d-4}$ so that the conditions in Claims B.7 and B.8 are both satisfied. Therefore, we must have $L'(\mathbf{f}_{d-3}) = 1001|\mathbf{0}$. This argument readily extends to more than four vectors, so it holds in general (possibly after reordering coordinates) that we may write $L'(\mathbf{f}_{d-i}) = \mathbf{e}_1 + \mathbf{e}_{i+1}$ for every $1 \le i \le d-1-c_1$. Let $D = d - c_1$. For every $1 \le i \le c_1$, let

$$L'(\mathbf{f}_i) = \sum_{j \in S_i} \mathbf{e}_j + \sum_{k \in T_i} \mathbf{e}_k,$$

where $S_i \subseteq [D]$ and $T_i \subseteq \{D+1,\dots,d-1\}$.

First, by Claim B.10, $S_i \ne \emptyset$. If $1 \notin S_i$, then by Claim B.10, it must be the case that $S_i = \{2,\dots,D\}$. But since $c_1 \le d-5$, so $D \ge 5$, therefore $|L'(\mathbf{f}_{d-1}) + L'(\mathbf{f}_i)| \ge D-2 \ge 3$, a contradiction. For the same reason, we can not have $S_i = [D]$ either. If $\{1\} \subsetneq S_i \subsetneq [D]$, then either $|S_i| = 2$ and $|T_i| \ge 1$ (by the distinctness requirement of Claim B.6) or $|S_i| \ge 3$. There must exist some index $j+1 \in \{2,\dots,D\}$ such that $j+1 \notin S_i$. Consequently, $L'(\mathbf{f}_i) + L'(\mathbf{f}_j)$ would be a vector of weight at least 3. Therefore, the only possibility left is $S_i = \{1\}$ for every $1 \le i \le c_1$. It follows that each T_i must be a singleton from the set $\{D+1,\dots,d-1\}$. Once again, a simple counting argument shows that there are only $c_1 - 1$ such singletons but we have c_1 0-labeled vectors to accommodate.

This completes the case analysis, and Lemma B.4 follows.

B.3 Proof of item 2: the special case of $c_1 = d - 4$

The case when $c_1 = d - 4$ requires some special treatment. The full proof involves some tedious case analysis. To make the proof simpler but at the same time contains all the main ideas, we prove here for the slightly weaker form which holds for all $d \geq 18$. Since the functions used for separation results are slightly different, we treat the cases of $c_2 = d - 2$ and $c_2 = d - 3$ separately in Section B.3.1 and Section B.3.2, respectively.

B.3.1 Proof for the Subcase $c_2 = d - 2$

Lemma B.11. *Let $d \geq 17$ be a fixed integer. Then $\mathcal{K}_d[\neg 0^{d-4}1^*] \nsubseteq_\delta \mathcal{K}_d[\neg 0^{d-2}1^*]$.*

Let $k = \lceil \frac{d}{2} \rceil$. we consider the following canonical function as in Definition 3.4: we write $g : \mathbb{F}_2^k \to \{0,1\}$ to denote the following function:

$$g(\mathbf{y}) = \begin{cases} 1 & \text{if } \mathbf{y} = \mathbf{e}_k \text{ or } |\mathbf{y}| \in \{2, k-2, k\}; \\ 0 & \text{otherwise,} \end{cases} \tag{B.3}$$

and we pick all of the sets $S_\mathbf{v}$ to be the entire subspace $\mathbb{F}_2^{n-k}$.

So apart from $\mathbf{e}_k$, the function g is one precisely if the input is some pairwise sum $\mathbf{e}_i + \mathbf{e}_j$, the complement $\overline{\mathbf{e}_i + \mathbf{e}_j}$ of such a sum (where i and j are two distinct integers in $[k]$), or the all-one vector $\mathbf{1}$.

In the rest of this section, we will stick to the following notations. Let $\{\mathbf{f}_1, \ldots, \mathbf{f}_{d-1}\}$ denote the basis vectors of the matroid $M(K_d)$ in the standard representation. If a linear transformation L that maps the basis vector $\mathbf{f}_i$ to $L(\mathbf{f}_i) = \mathbf{y}_i | \mathbf{z}_i$, where $\mathbf{y}_i \in \mathbb{F}_2^k$ and $\mathbf{z}_i \in \mathbb{F}_2^{n-k}$, then we write $L'(\mathbf{f}_i) = \mathbf{y}_i$ for the projection of this map to the part of coordinates which we care about most.

Proposition B.12. *The function g defined in (B.3) contains some $(K_d, 0^{d-2}1^*)$-patterns.*

Proof. This can be easily seen by considering the linear map L whose projection L' maps the basis vectors $\{\mathbf{f}_1, \ldots, \mathbf{f}_{d-1}\}$ in $M(K_d)$ as follows: for $1 \leq i \leq d-2$, if $i = 2\ell - 1$, then $L'(\mathbf{f}_i) = \mathbf{e}_\ell$; if $i = 2\ell$, then $L'(\mathbf{f}_i) = \overline{\mathbf{e}_\ell}$. Finally we set $L'(\mathbf{f}_{d-1}) = \mathbf{e}_k$. □

Lemma B.13. *The function g defined in (B.3) is $(K_d, 0^{d-4}1^*)$-free.*

Proof. We prove the lemma by contradiction. Suppose that there is a linear map L such that g contains $(0^{d-4}1^*)$-pattern at L's projection L'. First we observe that, since $g(\mathbf{0}) = 0$, for every pair of distinct i and j, $1 \leq i \neq j \leq d-1$, we have $L'(\mathbf{f}_i) \neq L'(\mathbf{f}_j)$. In other words, L must map the $d-1$ basis vectors to $d-1$ distinct vectors in $\mathbb{F}_2^k$. Note that this is simply a somewhat weaker version of Proposition B.3.

Claim B.14. L must map the $d-1$ basis vectors of $M(K_d)$ to $d-1$ distinct vectors in $\mathbb{F}_2^k$.

Another simple fact is that the zero vector is not the in images of the basis vectors in $M(K_d)$ under the linear map L.

Claim B.15. There is no basis vector of $M(K_d)$ that maps to $\mathbf{0}$.

Proof. Clearly, the 1-labeled basis vectors can not map to $\mathbf{0}$ as they are labeled 0 in g. Since $d - 4 \geq 2$, there are at least two 0-labeled basis vectors in $M(K_d)$, and by Claim B.14, there is at most one of them can map to $\mathbf{0}$. Suppose without loss of generality that $L'(\mathbf{f}_1) = \mathbf{0}$ and $L'(\mathbf{f}_2) \neq \mathbf{0}$. But then, on the one hand, $g(L'(\mathbf{f}_2)) = 0$ since $\mathbf{f}_2$ is a 0-labeled basis vector; on the other hand, $g(L'(\mathbf{f}_1 + \mathbf{f}_2)) = g(L'(\mathbf{f}_2)) = 1$, since their sum is labeled 1 in g. We thus reach a contradiction. □

Now we consider the set $\{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$. By assumption, g evaluates to 1 at these three vectors. Note that it also holds by assumption that L maps all 0-labeled basis vectors $\mathbf{f}_i$ for $1 \le i \le d-4$ to points such that $g(L'(\mathbf{f}_i)) = 0$. What we will prove is that there is no way to achieve this in such a way that it also holds that all pairwise sums of basis vectors $\mathbf{f}_i + \mathbf{f}_j$ for $1 \le i < j < d$ map to 1-labeled points, i.e., for any $i < j$, $g(L'(\mathbf{f}_i + \mathbf{f}_j)) = g(L'(\mathbf{f}_i) + L'(\mathbf{f}_j)) = 1$. But this contradicts the assumption that we are seeing $(K_d, 0^{d-4}1^*)$ at L.

The proof is by a case analysis. Let us start by ruling out some special cases. Our first observation is that the all-one vector $\mathbf{1}$ cannot be in the set.

Claim B.16. We have that $\mathbf{1} \notin \{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$.

Proof. For the purpose of contradiction, suppose $L'(\mathbf{f}_{d-1}) = \mathbf{1}$ and consider where L could map the 0-labeled basis vectors $\mathbf{f}_i$ for $1 \le i \le d-4$. Note that if $L'(\mathbf{f}_i)$ is a 0-labeled vector of weight ℓ, where $0 \le \ell \le k$, then the sum of $L'(\mathbf{f}_i)$ and $L'(\mathbf{f}_{d-1})$ must be a 1-labeled weight-$(k-\ell)$ vector. Recall that all the 1-labeled vectors in g are of weights either 1, 2, $k-2$ or k, therefore the weights of $L'(\mathbf{f}_i)$ can only be 0, 2, $k-2$ or $k-1$. First, by Claim B.15, we rule out the only weight-0 vector $\mathbf{0}$. Second, we cannot have $|L'(\mathbf{f}_i)| \in \{2, k-2\}$ because all the weight-2 and weight-$(k-2)$ vectors are labeled 1 in g. Thus, the only possibility left is that all $L'(f_i)$ are weight-$(k-1)$ vectors. But by Claim B.14, these 0-labeled basis vectors map to $d-4 \ge 2$ distinct weight-$(k-1)$ vectors. Consequently, their sums with $L'(\mathbf{f}_{d-1}) = \mathbf{1}$ would map to at least 2 distinct weight-1 vectors which would need to be labeled with 1. However, by our construction there is only one such vector (namely $\mathbf{e}_k$). □

Our next observation is that $\mathbf{e}_k$ cannot be in the set either.

Claim B.17. We have that $\mathbf{e}_k \notin \{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$.

Proof. For the sake of contradiction, assume without loss of generality that $L'(\mathbf{f}_{d-1}) = \mathbf{e}_k$. Then which vector could $L'(\mathbf{f}_{d-2})$ be? It must be distinct from $\mathbf{e}_k$ by Claim B.14 and cannot be $\mathbf{1}$ by Claim B.16. So we must have $|L'(\mathbf{f}_{d-2})| \in \{2, k-2\}$. If $L'(\mathbf{f}_{d-2})$ were a weight-2 vector, then $L'(\mathbf{f}_{d-1} + \mathbf{f}_{d-2}) = L'(\mathbf{f}_{d-1}) + L'(\mathbf{f}_{d-2})$ would be either a weight-3 or a weight-1 vector distinct from $\mathbf{e}_k$, and this vector would have a 0-label which is a contradiction. If on the other hand $L'(\mathbf{f}_{d-2})$ were a weight-$(k-2)$ vector, then $L'(\mathbf{f}_{d-1} + \mathbf{f}_{d-2})$ would be either a weight-$(k-3)$ vector or a weight-$(k-1)$ vector. But there is no vector of these two weights in g that is labeled 1, and we reach a contradiction again. □

Therefore, we are left with the cases that the set $\{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$ consists of three vectors whose weights are either 2 or $k-2$. To simplify the proof, let us introduce a new notation for the purpose of the proofs in this section only. We define the *reduced support* of a vector $\mathbf{v} \in \mathbb{F}_2^k$, denoted by $\mathrm{supp}_R(\mathbf{v})$ as: $\mathrm{supp}_R(\mathbf{v}) = \{j \in [k] : \mathbf{v}_j = 1\}$ if $|\mathbf{v}| \le \lfloor \frac{k}{2} \rfloor$, and $\mathrm{supp}_R(\mathbf{v}) = \mathrm{supp}_R(\overline{\mathbf{v}})$ otherwise. In other words, we are going to view a vector $\mathbf{v}$ and its complement $\overline{\mathbf{v}}$ as having the same reduced support, which is defined to be the "ordinary" support of the vector among $\mathbf{v}$ and $\overline{\mathbf{v}}$ whose weight is smaller.

Also for simplicity, we write $\{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\} = \{\mathbf{u}, \mathbf{v}, \mathbf{w}\}$. An important observation is, the supports of these three vectors $\mathbf{u}, \mathbf{v}$ and $\mathbf{w}$ must be pairwise intersecting. This is because otherwise a non-intersecting pair of vectors would have a weight-4 or weight-$(k-4)$ vector[8] as their sum vector. But all these vectors are labeled 0 in g, a contradiction.

If two of the vectors in $\{\mathbf{u}, \mathbf{v}, \mathbf{w}\}$ have the same reduced support set, then, by Claim B.14 (and for the same reason, we cannot have all three vectors having the same reduced support set), these two vectors are complements with each other. This is analyzed in Case 1 below. Now if the reduced support sets of these three vectors are three distinct subsets of $[k]$, each of size 2 and pairwise intersecting, then there are only two possibilities, which are treated in Case 2 and Case 3 below, respectively.

[8] By our assumption, $d \ge 17$, so $k \ge 9$ and hence $k-4 > 4$, therefore the two weights do not overlap.

Case 1: $\mathbf{w} = \overline{\mathbf{u}}$. Without loss of generality, we assume $\mathbf{u} = \mathbf{e}_i + \mathbf{e}_j$ and $\mathrm{supp}_R(\mathbf{v}) = \{i, \ell\}$ (that is, either $\mathbf{v} = \mathbf{e}_i + \mathbf{e}_\ell$ or $\mathbf{v} = \mathbf{e}_i + \mathbf{e}_\ell + 1^k$), where $1 \le i < j < \ell \le k$. Now we consider the set $T = \{L'(\mathbf{f}_1), L'(\mathbf{f}_2), \ldots, L'(\mathbf{f}_{d-4})\}$ and write $\mathbf{x}$ for an arbitrary element in T. By assumption, $g(\mathbf{x}) = 0$ and the sums of $\mathbf{x}$ with $\mathbf{u}, \mathbf{v}$ and $\mathbf{w}$ must all evaluate to 1 in g. First, $|\mathrm{supp}_R(\mathbf{x})| \neq 1$, since otherwise sums between $\mathbf{x}$ and $\mathbf{u}, \mathbf{v}$ and $\mathbf{w}$ would be three distinct vectors with reduced support set sizes either 1 or 3, but there is only one such vector (namely $\mathbf{e}_k$) evaluates to 1 in g. For the same reason, $|\mathrm{supp}_R(\mathbf{x})|$ cannot be an odd number. Second, note that $\mathbf{x} \neq \mathbf{0}_k$ by Claim B.15. Third, $|\mathrm{supp}_R(\mathbf{x})| \neq 2$ as all vectors of weight 2 and weight $(k-2)$ evaluate to 1 in g. Finally, observe that it is impossible that $|\mathrm{supp}_R(\mathbf{x})| \ge 6$. To see this, first note that this is impossible for $k < 12$ by our definition of reduced support sets. If, on the other hand, $k \ge 12$, then the support size of $\mathbf{x} + \mathbf{u}$ (or $\overline{\mathbf{x} + \mathbf{u}}$) is at least 4 and is at most $k/2 + 2 < k - 2$, but there is no vector labeled 1 in g of weight in this range. So we conclude that every vector in T must have reduced support size exactly 4. Furthermore, we must also have that $\mathrm{supp}_R(\mathbf{u}) \subseteq \mathrm{supp}_R(\mathbf{x})$ and $\mathrm{supp}_R(\mathbf{v}) \subseteq \mathrm{supp}_R(\mathbf{x})$, as otherwise their sum would be a vector of reduced support size at least 4 and at most 6, but there is no vector labeled 1 of these weights in g for $k \ge 9$. It follows that the reduced support of each vector in T is of the form $\{i, j, \ell, m\}$, where i, k, ℓ, m are four distinct integers. There are $d-4$ vectors in T, so we need at least $\lceil \frac{d-4}{2} \rceil = k - 2$ distinct integers m (since for every vector $\mathbf{x}$ of weight-4, $\mathbf{x}$ and $\overline{\mathbf{x}}$ are two distinct vectors but having the same reduced support set), requiring at least $(k-2) + 3 = k + 1$ distinct integers in the set $[k]$, a contradiction.

Case 2: $\mathrm{supp}_R(\mathbf{u}) = \{i, j\}$, $\mathrm{supp}_R(\mathbf{v}) = \{i, \ell\}$ and $\mathrm{supp}_R(\mathbf{w}) = \{j, \ell\}$, where $1 \le i < j < \ell \le k$. The arguments are identical to Case 1: that is, the only possibility is the reduced support of each vector in T is of the form $\{i, j, \ell, m\}$, where i, j, ℓ, m are four distinct integers, but there are not enough integers in $[k]$ to accommodate $k + 1$ distinct integers.

Case 3: $\mathrm{supp}_R(\mathbf{u}) = \{i, j\}$, $\mathrm{supp}_R(\mathbf{v}) = \{i, \ell\}$ and $\mathrm{supp}_R(\mathbf{w}) = \{i, m\}$, where the indices are four distinct integers $1 \le i < j < \ell < m \le k$. Once again, we can use the same arguments as in Case 1 to show that the reduced support of each vector x in T has weight 4 and satisfies $\mathrm{supp}_R(\mathbf{u}) \subseteq \mathrm{supp}_R(\mathbf{x})$, $\mathrm{supp}_R(\mathbf{v}) \subseteq \mathrm{supp}_R(\mathbf{x})$ and $\mathrm{supp}_R(\mathbf{w}) \subseteq \mathrm{supp}_R(\mathbf{x})$. Clearly there is only one reduced support set (and hence at most two vectors) satisfying all these conditions, namely $\mathrm{supp}_R(\mathbf{x}) = \{i, j, \ell, m\}$, but we have $d - 4 > 2$ vectors in T, a contradiction.

This completes the case analysis and thus proves Lemma B.13. □

Lemma B.11 follows directly from Proposition B.12 and Lemma B.13.

B.3.2 Proof for the Subcase $c_2 = d - 3$

Lemma B.18. *Let $d \ge 18$ be a fixed integer, then $(K_d, 0^{d-4}1^*)$-free $\nsubseteq (K_d, 0^{d-3}1^*)$-free.*

The proof of Lemma B.18 is very much the same as the proof Lemma B.11, so we only describe the difference.

Let $k = \lceil \frac{d-1}{2} \rceil$. Now we let $g_1 : \mathbb{F}_2^k \to \{0, 1\}$ defined by:

$$g_1(y) = \begin{cases} 1 & \text{if } y = \mathbf{e}_k, \overline{\mathbf{e}_k}, \text{ or } |\mathbf{y}| \in \{2, k-2, k\}; \\ 0 & \text{otherwise.} \end{cases} \tag{B.4}$$

Proposition B.19. *The function g_1 defined in (B.4) contains some $(K_d, 0^{d-3}1^*)$-patterns.*

Proof. Function g_1 contains some $(K_d, 0^{d-3}1^*)$-patterns can be checked by considering the linear map L'' that maps the basis vectors $\{\mathbf{f}_1, \ldots, \mathbf{f}_{d-1}\}$ in $M(K_d)$ as follows: For $1 \le i \le d-3$, if $i = 2\ell - 1$, then $L''(\mathbf{f}_i) = \mathbf{e}_\ell$; if $i = 2\ell$, then $L''(\mathbf{f}_i) = \overline{\mathbf{e}_\ell}$. Finally we set $L''(\mathbf{f}_{d-2}) = \mathbf{e}_k$ and $L''(\mathbf{f}_{d-1}) = \overline{\mathbf{e}_k}$. □

Lemma B.20. *The function g_1 defined in (B.4) is $(K_d, 0^{d-4}1^*)$-free.*

Proof. We will follow the proof of Lemma B.13. First, Claim B.14 and Claim B.15 follow directly from the same proofs without any modification.

Now we consider the set $\{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$. We still have Claim B.16 but will need to modify the proof slightly. Namely, we need to also consider the case that $d-4$ vectors, $\{L'(\mathbf{f}_1), \dots, L'(\mathbf{f}_{d-4})\}$, are weight-1 vectors. However, by symmetry, the same argument for the case that these vectors are weight-$(k-1)$ vectors applies, so Claim B.16 still holds.

Claim B.21. We have that $\mathbf{e}_k \notin \{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$ and $\overline{\mathbf{e}_k} \notin \{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$.

Proof. For the sake of contradiction, assume without loss of generality that $L'(\mathbf{f}_{d-1}) = \mathbf{e}_k$. Then which vector could $L'(\mathbf{f}_{d-2})$ be? It must be distinct from $\mathbf{e}_k$ by Claim B.14 and cannot be $\mathbf{1}$ by Claim B.16. If $L'(\mathbf{f}_{d-2})$ were a weight-2 vector, then $L'(\mathbf{f}_{d-1} + \mathbf{f}_{d-2}) = L'(\mathbf{f}_{d-1}) + L'(\mathbf{f}_{d-2})$ would be either a weight-3 or a weight-1 vector distinct from $\mathbf{e}_k$, and this vector would have a 0-label which is a contradiction. If on the other hand $L'(\mathbf{f}_{d-2})$ were a weight-$(k-2)$ vector, then $L'(\mathbf{f}_{d-1} + \mathbf{f}_{d-2})$ would be either a weight-$(k-3)$ vector or a weight-$(k-1)$ vector not equal to $\overline{\mathbf{e}_k}$. But there are no such vectors in g that are labeled 1, and we reach a contradiction again. Therefore, the only possibility left is $L'(\mathbf{f}_{d-2}) = \overline{\mathbf{e}_k}$. But now we have no place for $L'(\mathbf{f}_{d-3})$, as we can apply the same reasoning above for $L'(\mathbf{f}_{d-2})$ to $L'(\mathbf{f}_{d-3})$. We thus prove that $\mathbf{e}_k \notin \{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$.

Now by symmetry, the same argument shows that $\overline{\mathbf{e}_k} \notin \{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$. ☐

Therefore, we are left with the cases that the set $\{L'(\mathbf{f}_{d-3}), L'(\mathbf{f}_{d-2}), L'(\mathbf{f}_{d-1})\}$ consists of three vectors whose weights are either 2 or $k-2$. Finally, we repeat the same case analysis as in the proof of Lemma B.11. The only difference is we need the fact that $\lceil \frac{d-4}{2} \rceil \geq k-2$, which is clearly true as we defined k to be $\lceil \frac{d-1}{2} \rceil$. ☐

Now Lemma B.18 follows directly from Proposition B.19 and Lemma B.20.

B.4 Proof of item 3

Lemma B.22. *We have $\mathcal{K}_d[\neg 01^*] \not\subseteq_\delta \mathcal{K}_d[\neg 0^{d-1}1^*]$ and for any $1 < c \leq d-1$ it holds that $\mathcal{K}_d[\neg 01^*] \subset_\delta \mathcal{K}_d[\neg 0^c 1^*]$.*

The first part of item 3 is proven in Section 6 and the second part is proven in Lemma 5.6.

B.5 Proof of item 4

Lemma B.23. *For any $c_1 \leq d-1$ and $c_2 \leq d$, $\mathcal{K}_{d+1}[\neg 0^{c_2}1^*] \not\subseteq_\delta \mathcal{K}_d[\neg 0^{c_1}1^*]$.*

Proof. Consider the function $f^0_{(K_d, 0^{c_1}1^*)}$, which is far from being $(K_d, 0^{c_1}1^*)$-free. Suppose that this function contains $(K_{d+1}, 0^{c_2}1^*)$ at some linear transformation $L : K_{d+1} \to \mathbb{F}_2^n$. Then it is straightforward to verify that the function must also contain $(K_d, 1^*)$ pattern. To see this, note that since $c_2 \leq d$, there exists at least one 1-labeled weigth-1 vector in the $(K_{d+1}, 0^{c_2}1^*)$ pattern. Suppose, without loss of generality, that $\mathbf{e}_d$ is one such vector. Then let $\phi : K_d \to K_{d+1}$ which maps $\mathbf{e}_i$ to $\mathbf{f}_i = \mathbf{e}_i + \mathbf{e}_d$ for all $1 \leq i \leq d-1$. Then it is easy to check that $f^0_{(K_d, 0^{c_1}1^*)}$ contains $(K_d, 1^*)$ at $L \circ \phi$. But now by the proof of the First Dichotomy Theorem, we infer that there exists an embedding from $(K_d, 1^*)$ into $(K_d, 0^c 1^*)$, which is a contradiction since no such embedding exists. ☐

B.6 Proof of item 5

Lemma B.24. *For any* $0 \leq c \leq d-1$*, it holds that* $\mathcal{K}_d[\neg 0^c 1^*] \subset_\delta \mathcal{K}_{d+1}[\neg 0^c 1^*]$ *and* $\mathcal{K}_d[\neg 0^c 1^*] \subset_\delta \mathcal{K}_{d+1}[\neg 0^{c+1} 1^*]$.

Proof. The Lemma follows immediately from item 4 and subgraph containment. □

B.7 Proof of item 6

Lemma B.25. *Suppose* $d \geq 5$ *and let* $\mathfrak{N} = \{(K_{d_1}, 0^{c_1}1^*), (K_{d_2}, 0^{c_2}1^*), \ldots, (K_{d_t}, 0^{c_t}1^*)\}$ *be any finite set of labeled matroids over complete graphs such that* $d_1 < d$ *or* $c_i > 0$ *for all* $i \in [t]$*. Then the property of* $(K_d, 1^*)$*-freeness is well separated from the union of* $(K_{d_i}, 0^{c_i}1^*)$*-freeness properties for all* $i \in [t]$*, and the difference is provably far from being low-degree polynomials.*

Proof. Apply the Third Dichotomy Theorem, Theorem 3.9 with $M = (K_d, 1^*)$-freeness and $\mathfrak{N} = \{(K_{d_1}, 0^{c_1}1^*), (K_{d_2}, 0^{c_2}1^*), \ldots, (K_{d_t}, 0^{c_t}1^*)\}$. □