THE GEOMETRY OF PARTIAL FITNESS ORDERS AND AN EFFICIENT METHOD FOR DETECTING GENETIC INTERACTIONS

CAITLIN LIENKAEMPER*[1], LISA LAMBERTI*[2], JAMES DRAIN[3], NIKO BEERENWINKEL[2], AND ALEX GAVRYUSHKIN[✉2]

Abstract. We present an efficient computational approach for detecting genetic interactions from fitness comparison data together with a geometric interpretation using polyhedral cones associated to partial orderings. Genetic interactions are defined by linear forms with integer coefficients in the fitness variables assigned to genotypes. These forms generalize several popular approaches to study interactions, including Fourier-Walsh coefficients, interaction coordinates, and circuits. We assume that fitness measurements come with high uncertainty or are even unavailable, as is the case for many empirical studies, and derive interactions only from comparisons of genotypes with respect to their fitness, i.e. from partial fitness orders. We present a characterization of the class of partial fitness orders that imply interactions, using a graph-theoretic approach. Our characterization then yields an efficient algorithm for testing the condition when certain genetic interactions, such as sign epistasis, are implied. This provides an exponential improvement of the best previously known method. We also present a geometric interpretation of our characterization, which provides the basis for statistical analysis of partial fitness orders and genetic interactions.

Genetic interactions – or dependence of the fitness effect produced by a set of mutations on the genetic background – play an important role in determining evolutionary trajectories of populations. For instance, a set of individually beneficial mutations may exhibit diminishing returns: while the combined effect of all mutations is beneficial, it is not as beneficial as one would expect given the individual mutations, for example the sum of all single mutation effects. Diminishing returns have been shown to slow down the pace of adaptation (Chou et al. 2011). Similarly, the combined effect of a set of mutations can be stronger than the effect expected from the individual mutations. This so-called synergistic epistasis between deleterious mutations has been observed for instance in human populations, where it affects the distribution of deleterious alleles in the genome (Sohail et al. 2017). Genetic interactions can affect not just the magnitude of the effect of a mutation, but also the sign of the effect. That is, a particular mutation may have a beneficial effect in one genetic background and a deleterious effect in another. This type of

[1]Department of Mathematics, Pennsylvania State University

[2]Department of Biosystems Science and Engineering, ETH Zürich, Basel, Switzerland; SIB Swiss Institute of Bioinformatics, Basel, Switzerland

[3]Department of Mathematics, University of California, Los Angeles

E-mail addresses: clienkaemper@psu.edu, lisa.lamberti@bsse.ethz.ch, jamesdrain@ucla.edu, niko.beerenwinkel@bsse.ethz.ch, ✉alex@gavruskin.com.

Date: August 25, 2017.

*Equal contribution.

interaction, often termed sign epistasis, can act to constrain evolutionary trajectories by requiring that some mutations occur before others (Gong, Suchard, and Bloom 2013; Kvitek and Sherlock 2011; Weinreich, Watson, and Chao 2005).

In this work, we are concerned with detecting the presence of genetic interactions from limited genotype-phenotype data. Empirical data sets produced in interaction studies typically provide replicate measurements of a trait, such as antibiotic or other forms of resistance, for various genotypes. The trait is often chosen to measure (approximately) fitness, so the measurements can be used to infer genetic interactions in the fitness landscape. However, in many cases the association between the trait and fitness is unclear. Here, we make a conservative assumption that the only reliable information about fitness carried by the trait is the (partial) ranking of genotypes. Specifically, we assume that the only signal available in the data is of the form "genotype g_1 has higher fitness than genotype g_2", for various pairs of genotypes. We call this type of data a *partial fitness order*, which is a partial order of genotypes with respect to their fitness. Data of this type arises for example by considering certain directed acyclic graphs where the vertices are given by genotypes and where the edges between genotypes are directed towards the genotype with higher fitness.

Even though we assume that we cannot access accurate fitness measurements for all genotypes in a given genotype space, we are still able to detect interactions implied by the partial fitness order. Specifically, this means we are able to detect sign epistasis. Even when complete data about fitness values is available, a method to detect interactions from a partial order alone provides additional insight. For instance, even when an observed set of fitness values implies interaction, the rank order of the values may not. In this case, we may be able to conclude diminishing returns or synergistic epistasis and refute sign epistasis. Alternately, if we are interested only in a certain set of fitness comparisons (for instance, comparisons between mutational neighbors), we might want to determine whether this set of comparisons is sufficient to imply interaction. These observations will have implications for the evolutionary trajectories of the system and provide insight as to the mechanism of the interaction (Weinreich, Watson, and Chao 2005). A detailed analysis of the class of linear orders that imply interactions including relevant literature and data, can be found in (Crona, Gavryushkin, et al. 2017).

In this paper, we build on the results obtained by Crona, Gavryushkin, et al. (2017) and settle the most general case of ranking data, when only a partial fitness order is available. Specifically, we present a graph-theoretic characterization of the class of partial fitness orders that imply genetic interactions (Section 2). Second, we use our characterization to derive an efficient algorithm for testing the condition when a given partial order implies such an interaction (Section 2). The complexity of our algorithm is $\mathcal{O}(n^3)$, where n is the number of genotypes under consideration. This provides an exponential improvement of the best previously known method that involves iterating through the list of all possible linear extensions of the given partial order (Crona, Gavryushkin, et al. 2017). Third, we count all partial orders for up to 8 labeled elements which imply an interaction (Section 3), and fourth, we provide a geometric interpretation of our characterization, useful for statistical analysis of partial fitness orders that imply genetic interactions (Section 4). Finally, we use public data sets to demonstrate our methods (Section 5) and conclude by describing possible future directions for this line of research (Section 6).

1. Technical introduction

In this section we introduce all necessary terminology and notations and describe a basic statistical approach to probabilistic inference of (higher-order) interactions from partial fitness orders, which further motivates our results on partial fitness orders.

Consider a system $\mathcal{G}$ consisting of k n-locus biallelic genotypes, where k and n are positive integer numbers. The system $\mathcal{G}$ can be seen as a set of k binary sequences of length n. We denote the fitness of a genotype $g \in \mathcal{G}$ by w_g, which is a non-negative real number. All fitness measurements $(w_g)_{g\in\mathcal{G}}$ together then determine what is called the *fitness landscape* associated to $\mathcal{G}$. We refer to the tuple of all the fitness measurements defining a fitness landscape by $W = (w_g)_{g\in\mathcal{G}}$.

We define genetic interactions using linear forms with integer coefficients in the following way. Let f be a linear form in variables W with integer coefficients which sum to 0, that is,

$$f(W) = \sum_{1\leq i\leq t} c_i w_{p_i} - \sum_{1\leq j\leq s} d_j w_{n_j},$$

such that

$$\sum_{1\leq i\leq t} c_i - \sum_{1\leq j\leq s} d_j = 0,$$

where all c_i, d_j, t, and s are positive integer numbers such that $t + s = k$ and $\mathcal{G} = \{p_1, \ldots, p_t, n_1, \ldots, n_s\}$. We then say that the fitness landscape *has positive f-interaction* if $f(W) > 0$. When $f(W) < 0$, we say that the fitness landscape *has negative f-interaction*. From now on, we will focus on determining positive f-interactions unless we explicitly state otherwise. Everything in this paper extends straightforwardly to the negative case. This definition of f-interaction generalizes several well-known and widely used notions of gene interaction, including Fourier-Walsh coefficients, interaction coordinates, and circuit interactions (Weinreich, Lan, et al. 2013; Beerenwinkel, Pachter, and Sturmfels 2007).

Estimating the probability of f-interaction in an empirical setting amounts to quantifying the uncertainty of $f(W) > 0$ given some comparative fitness data. Although in this setting the fitness values w_g are unavailable, the comparative fitness data typically allows to deduce a partial fitness order. For example, this type of data could arise from replicate measurements of any quantitative trait monotonic with respect to fitness. Another example of such data is fitness comparison data produced in competition or survival experiments.

Sometimes a partial fitness order is enough to deduce f-interactions and hence estimate the probability of interaction. We use the following definition to address such situations. We say that a partial fitness order $\mathcal{P} = (\mathcal{G}, \prec)$ *implies positive f-interaction* if $f(W) > 0$ whenever W satisfy the partial order $\mathcal{P}$, that is, $w_g < w_h$ for all $g \prec h$.

Then the probability of f-interaction given comparative fitness data can be estimated by considering the probability support of all partial orders that imply f-interaction. However, such an estimation would require the condition of whether or not a partial order $\mathcal{P}$ implies f-interaction to be checked routinely for different partial orders $\mathcal{P}$. Hence the complexity of this condition might become a computational bottleneck in practice. In this paper we resolve this problem by designing an efficient polynomial-time algorithm for checking whether or not a partial order implies f-interaction.

As observed earlier, our work is motivated by and relies on previous results contained in (Crona, Gavryushkin, et al. 2017). We now review their results about how to detect whether a rank order implies f-interaction. In the terminology used in our paper, this situation corresponds to the case when the partial fitness order is a linear order (also called a total order or a rank order). We then study arbitrary partial orders, and present an efficient algorithm for determining when a partial order implies f-interaction.

For a linear form f as above, define a function ϕ^f from the set of (strict) linear orders on the set of genotypes to words over the alphabet $\{P, N\}$ as follows. The linear order on k genotypes

$$g_1 \succ g_2 \succ \cdots \succ g_k$$

is mapped to a word ω consisting of c_i copies of P if w_{g_i} has a positive coefficient c_i in f; and d_i copies of N if w_{g_i} has a negative coefficient d_i in f. The subwords of ω corresponding to different g_i's appear according to the linear order (from highest to lowest). See (Crona, Gavryushkin, et al. 2017, Section A.2) for various examples of this construction.

The following result from (Crona, Gavryushkin, et al. 2017) depends upon the characterization of linear orders which imply interaction in terms of Dyck words.

Definition 1. Let Σ be an alphabet consisting of two letters P and N. A *Dyck word* is a word ω consisting of an equal number of P's and N's such that every prefix of ω contains at least as many P's as N's.

For example, $PPNNPN$ and $PNPNPN$ are Dyck words, but $PPNNNP$ is not, since the prefix $PPNNN$ contains more N's than P's. Clearly the definition of Dyck word does not depend on the choice of alphabet one considers. Dyck words arise in a number of contexts in discrete mathematics. For instance, they are in bijection with full binary trees with $n+1$ leaves. If the symbols P and N are replaced with "(" and ")", a string is a Dyck word only if all parenthesis are correctly matched. For more on Dyck words, see (Stanley 2001).

Theorem 1 (Crona, Gavryushkin, et al. (2017))**.** *A linear order $\mathcal{L} = (\mathcal{G}, \prec)$ on the set of genotypes implies positive f-interaction if and only if $\phi^f(\mathcal{L})$ is a Dyck word that starts with P.*

The statement of this theorem can equivalently be formulated by saying that a linear order $\mathcal{L}$ implies positive f-interaction if and only if there exists a partition of the set of all genotypes $\mathcal{G}$ into pairs (p_i, n_j) such that $p_i \succ n_j$ for all i, j, where each p_i appears in c_i pairs and each n_j in d_j pairs (see Crona, Gavryushkin, et al. 2017). As we will see in the next section, this formulation can be generalized to arbitrary partial orders.

2. Efficient method to infer interactions from partial orders

In this section we generalize Theorem 1 and characterize partial orders which imply f-interactions. We then apply our characterization to design an efficient algorithm for testing the condition of whether or not a partial order implies f-interaction.

Theorem 2. *A partial order $\mathcal{P} = (\mathcal{G}, \prec)$ implies positive f-interaction if and only if there exists a partition of the set of all genotypes $\mathcal{G}$ into pairs (p_i, n_j) such that $p_i \succ n_j$ for all i, j, where each p_i appears in c_i pairs and each n_j in d_j pairs.*

Before we proceed with the proof, let us illustrate the result with an example. Let

$$W = (w_{001}, w_{010}, w_{100}, w_{111}, w_{000})$$

and consider the linear form:

$$f(W) = w_{001} + w_{010} + w_{100} - w_{111} - 2w_{000},$$

which measures the deviation of the triple mutant's fitness from its expectation based on single mutants only, and corresponds to one of the circuit interactions characterized in (Beerenwinkel, Pachter, and Sturmfels 2007). We see that there is positive f-interaction implied by the order of the fitness values if and only if for each fitness value with a positive coefficient in f, there is a strictly smaller fitness value with negative coefficient. This smaller fitness value might not necessarily be unique, as it is the case here for w_{000}.

Proof. The "only if" implication is clear and follows from (Crona, Gavryushkin, et al. 2017). We provide some details for convenience of the reader.

Recall that

$$f(W) = \sum_{1\leq i\leq t} c_i w_{p_i} - \sum_{1\leq j\leq s} d_j w_{n_j} = \sum_{\substack{1\leq i\leq t\\ 1\leq j\leq s}} \Big(c_i w_{p_i} - d_j w_{n_j}\Big).$$

If there is a partition for the partial order $\mathcal{P}$ as claimed and the indices i, j are arranged according to the partition, each difference in the last sum is strictly greater than 0, thus $f(W) > 0$.

In order to prove the converse statement, we prove its contrapositive. That is, we prove that if there is no such partition of the set of all genotypes, f-interaction is not implied.

Construct the bipartite graph $G = (\mathbf{P} \sqcup \mathbf{N}, E)$, where $\mathbf{P}$ contains c_i copies of p_i for each i and $\mathbf{N}$ contains d_j copies of n_j for each j, such that $(n_j, p_i) \in E$ whenever $n_j \prec p_i$ in the partial order $\mathcal{P}$. Note that the number of elements in $\mathbf{P}$ and $\mathbf{N}$ coincide, since the coefficients of f sum to zero. Observe that a partition of the set of all genotypes $\mathcal{G}$ into pairs (p_i, n_j) such that $n_j \prec p_i$ for all i, j as in the theorem is exactly a perfect matching in G. By Hall's theorem, there exists a perfect matching in G if and only if for each $S \subset \mathbf{N}$, we have

$$|S| \leq |N_G(S)|,$$

where $N_G(S)$ is the set of vertices in G adjacent to a vertex in S (Diestel 2016). Suppose that there is no partition of the genotypes, and thus no perfect matching in G. Then there is some subset $S \subset \mathbf{N}$ such that

$$|S| > |N_G(S)|.$$

Using this subset S, we construct a linear extension of the partial order $\mathcal{P}$ such that the linear extension does not map to a Dyck word under the map ϕ^f, introduced above.

Note that the only elements of $\mathbf{P}$ which must be greater than an element of S in a linear order compatible with $\mathcal{P}$ are those in $N_G(S)$. Therefore, we can construct a linear order $\mathcal{L}$ compatible with $\mathcal{P}$ such that the only elements of $\mathbf{P}$ which are greater than an element of S are those in $N_G(S)$. To do so, reindex $\mathbf{N}$ so that $S = \{n_1, \ldots, n_j\}$ where the order $n_1 \prec \ldots \prec n_j$ is compatible with $\mathcal{P}$. Let X_1 be the set $\{x \mid n_1 \prec x\}$, that is, the set of elements in $\mathcal{P}$ ordered above n_1. Note that

$X_1 \subset \mathbf{N} \cup N_G(S)$, since if $n_1 \prec x$ and $x \in \mathbf{P}$, then by construction, $x \in N_G(S)$. Consider a linear extension of $\mathcal{P}$ on X_1 and list the elements of this extension from the greatest to the smallest. Append n_1 to the end of this list. Now, we repeat this process with the remaining elements of S: we let X_2 be the set of elements which are required by $\mathcal{P}$ to be greater than n_2, but are not contained in X_1. That is, $X_2 = \{x \mid n_2 \prec x\} \setminus X_1$. Again, $X_2 \subseteq \mathbf{N} \cup N_G(S)$. Extend the previous list by appending a linear extension of $\mathcal{P}$ on X_2 followed by n_2 thus maintaining the property that the list is sorted in a linear order compatible with $\mathcal{P}$. We continue this process for the rest of elements in S. To complete the construction, we sort all remaining elements in an order compatible with $\mathcal{P}$ and append them to the end of the list. This is possible, since any element which is not included in $S, X_1, \ldots, X_j$ is smaller then all elements of S and all elements of $X_1 \cup \ldots \cup X_j$. Therefore, our list is a valid linear order $\mathcal{L}$ compatible with $\mathcal{P}$. However, this linear order does not map to a Dyck word under ϕ^f, since the smallest prefix of $\phi^f(\mathcal{L})$ which contains the images of all elements of S will only contain images of elements of $\mathbf{N} \cup N_G(S)$. Since $|N_G(S)| < |S|$, this prefix will contain more N's than P's. Therefore, by Theorem 3 of (Crona, Gavryushkin, et al. 2017), this partial order does not imply positive f-interaction. □

A variant of Theorem 2 was independently obtained in (Crona and Luo 2017).

We now proceed by designing an efficient algorithm for detecting whether or not a partial fitness order $\mathcal{P}$ implies f-interaction. Recall that $f(W) = \sum_{1 \le i \le t} c_i w_{p_i} - \sum_{1 \le j \le s} d_j w_{n_j}$ and $\sum_{1 \le i \le t} c_i = \sum_{1 \le j \le s} d_j$. In the next result, denote $m = \sum_{1 \le i \le t} c_i$.

Theorem 3. *There exists an $\mathcal{O}(m^3)$ algorithm to determine whether a partial order $\mathcal{P} = (\mathcal{G}, \prec)$ implies positive f-interaction.*

Proof. We use the notation introduced in the proof of Theorem 2. That theorem implies that a partial order implies that $f(W) > 0$ if and only if there exists a partition of the set $P' \cup N'$ into pairs (p_i, n_j) such that $p_i \succ n_j$. As noted in the proof of Theorem 2, this condition is equivalent to the existence of a perfect matching in the bipartite graph $G = (\mathbf{P} \sqcup \mathbf{N}, E)$ such that $(n_j, p_i) \in E$ whenever $p_i \succ n_j$. Thus, we can detect whether the poset implies $f(W) > 0$ by checking whether there exists a perfect matching in G.

Using the Hopcroft-Karp algorithm (Hopcroft and Karp 1971), we can find a perfect matching in a bipartite graph with m vertices in time $\mathcal{O}(m^{5/2})$. Assume we are presented the partially ordered set $\mathcal{P}$ as a directed acyclic graph. Then we can construct the bipartite graph G in time $\mathcal{O}(m^3)$ using the following algorithm: for each $p \in \mathbf{P}$, use breadth first search to find all elements of $\mathbf{N}$ which are reachable along a directed path from p, and add these to p's list of neighbors. The worst-case complexity of this step is $\mathcal{O}(m^2)$. Repeat this process for each of the m vertices to obtain a representation of G as an adjacency list – that is, as a list containing a list of neighboring vertices for each vertex. Thus, since we can construct the graph in $\mathcal{O}(m^3)$ time and determine whether a matching exists in the graph in $\mathcal{O}(m^{5/2})$ time, and $m^3 > m^{5/2}$, we can determine whether the partial order determines the sign of a linear form in $\mathcal{O}(m^3)$ time. □

Note that the worst case complexity is a function not of the number of genotypes, but of $\sum_{1 \le i \le t} c_i$. Since genotypes are included in $\mathbf{P}$ or $\mathbf{N}$ multiple times if their

coefficient in the linear form f is an integer other than ± 1, $\mathbf{P}$ and $\mathbf{N}$ can in general be larger than $\mathcal{G}$. Thus, the complexity of our algorithm for computing whether a partial order on a set of genotypes implies f-interaction depends not just on the number of genotypes, but also on the coefficients of f.

The two upper level steps of our algorithm are

- constructing the bipartite graph and
- detecting whether a perfect matching exists.

We think that the upper complexity bound $\mathcal{O}(m^3)$ achieved in this paper can be improved. In particular, an appropriate choice of data structures can reduce the complexity of constructing the bipartite graph, which is the most demanding step in our algorithm. However, this improvement goes beyond the scope of this work, as we found the algorithm's performance sufficient on typical data sets used in higher-order interaction studies (see Crona, Gavryushkin, et al. 2017, for a survey of such data sets). See `https://github.com/gavruskin/fitlands/tree/posets/posets.ipynb`, where our implementation of the algorithm designed in Theorem 3 can be accessed.

3. Counting partial orders that imply interaction

In this section, we determine the proportion of partial orders on $k = 2$, 4, 6, and 8 genotypes that imply positive f-interaction for any linear form f with k coefficients 1 and -1 which sum to 0. Our results are summarized in Table 1.

k	positive interactions	no interactions	proportion implying interaction
2	1	1	0.67
4	31	157	0.28
6	10,876	108,271	0.17
8	22,217,743	387,287,893	0.10

Table 1. For $k \in \{2, 4, 6, 8\}$, the number of partial orders implying positive f-interaction, the number of partial orders not implying f-interaction, and the proportion of partial orders implying either positive or negative f-interaction, truncated to the second digit. Note that the number of partial orders implying negative f-interaction is, by symmetry, equal to the number of partial orders implying positive f-interaction. Hence, the proportion is obtained by dividing twice the number of partial orders implying positive interaction by the total number of partial orders.

We carried out our computations in the open-source mathematics software system Sage (`http://www.sagemath.org`) using the source code available at `https://github.com/gavruskin/fitlands/tree/posets/`. To make the computations possible, we employed the automorphism groups (Stein 2008) of our partial orders in the following way.

Definition 2. An *automorphism* of a directed graph $G = (V, E)$ is a bijection $\sigma : V \to V$ such that if $(u, v) \in E$, then $(\sigma(u), \sigma(v)) \in E$. The set of automorphisms of a finite direct graph forms a group, denoted by $\mathrm{Aut}(G)$, under the operation of composition.

In our enumeration, we first identify the poset with its Hasse diagram. We then observe that the number of different ways to label a Hasse diagram with k elements is $\frac{k!}{|\mathrm{Aut}(\mathcal{P})|}$. To see this, take a labeling of the elements on the poset. Each automorphism of the Hasse diagram gives a permutation of the labels which does not change the partial order. Likewise, any permutation of the labels which does not change the partial order is an automorphism of the Hasse diagram. Thus, each relabeling of the poset elements falls into an equivalence class of $|\mathrm{Aut}(\mathcal{P})|$ relabelings that do not change the partial order. Thus, since there are $k!$ ways to relabel the elements, there are $\frac{k!}{|\mathrm{Aut}(\mathcal{P})|}$ non equivalent ways to label the poset.

As an example, consider the partial order

$$\mathcal{P} = (\{00, 01, 10, 11\}, \{11 \succ 01, 11 \succ 10, 10 \succ 00, 01 \succ 00\})$$

with Hasse diagram depicted on the left of the following figure:

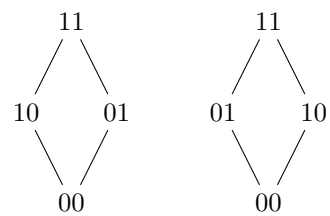

The relabeling to the Hasse diagram on the right yields the same partial order. Thus, this permutation switching 10 and 01 is an automorphism of the graph. However, any other permutation of the labels changes the partial order. Thus, this poset has an automorphism group of cardinality 2. Then since there are $4! = 24$ ways to relabel the elements of $\{00, 01, 10, 11\}$, there are $4!/2 = 12$ different labelings of this poset.

The implementation of our algorithm proceeds as follows: Using a built in function in Sage, we produced a list of partial orders on k elements up to isomorphism. Since whether a partial order implies f-interaction depends on the element labels, not just the structure of the partial order, we have to distinguish between different labeled partial orders within each partial order automorphism class. To achieve this, for each partial order (up to automorphism), we check for f-interaction with each partition of the labels used by Sage into two classes of equal size. We let one class correspond to elements with coefficient 1 and one class correspond to elements with coefficient -1. Then each choice of partition corresponds to $((k/2)!)^2$ possible reassignments of the labels, since we can permute the elements within the class with coefficient 1 and the class with coefficient -1 without changing the partition. We compute the cardinality of the automorphism group for each partial order, and add $\frac{1}{|\mathrm{Aut}(\mathcal{P})|}$ to our count of partial orders that imply positive, negative, or no f-interaction respectively.

We note that as we increase k, the proportion of partial orders implying f-interaction decreases. The sums of the values we obtained for the number of partial orders implying positive f-interaction, the number of partial orders implying negative f-interaction, and the number of partial orders not implying any f-interaction appears in OEIS (*The On-Line Encyclopedia of Integer Sequences* 2017) as the number of partial orders on k labeled elements, as expected. However, none of the columns of our table appear in OEIS. Further, each number is either prime or has a prime factorization including fairly large primes. Thus, the number of partial

orders, the number of partial orders implying f-interaction, and the number of partial orders not implying f-interaction do not appear to be given by any simple formula.

The final computation, for $k = 8$, took approximately two hours to complete on the SageMath Cloud server. The main factor affecting the performance of these computations is that the number of partial orders on k elements is super-exponential in the number of elements (*The On-Line Encyclopedia of Integer Sequences* 2017). Therefore, we did not count the number of partial orders on k elements which imply f-interaction for $k \geq 10$. Since we can still quickly check whether individual partial orders on $k \geq 10$ elements imply f-interaction, sampling methods can be used to estimate how the proportion of partial orders which imply interaction changes as we increase k.

Our computational observation that the fraction of partial orders which imply f-interaction approaches 0 as k increases is true in general:

Theorem 4. *As k approaches infinity, the fraction of partial orders on $k \leq 2^n$ elements which imply f-interaction approaches zero.*

Proof. Consider a partial order on k elements chosen uniformly at random and take a linear extension of this order also uniformly at random. Since each linear order is consistent with the same number of partial orders, this method produces a uniformly random linear order. Since the probability that a random linear order implies interaction goes to zero as k goes to infinity (Crona, Gavryushkin, et al. 2017, Corollary 2), the probability that the partial order we started with implies interaction goes to zero as k goes to infinity. Thus, the fraction of partial orders on k elements which imply f-interaction approaches zero as k approaches infinity. □

4. Geometric interpretation

In this section, we present a geometric interpretation of Theorem 2, and derive some geometric, combinatorial, and statistical results. We denote the space $\mathbb{R}^{|\mathcal{G}|}$ where $|\mathcal{G}|$ is the number of elements in $\mathcal{G}$ simply by $\mathbb{R}^{\mathcal{G}}$ and index the coordinate axes by the unknown fitness values w_g's, where $g \in \mathcal{G}$. Since we assume fitness values to be nonnegative, we work in the positive orthants of $\mathbb{R}^{\mathcal{G}}$, which we denote by $\mathbb{R}^{\mathcal{G}}_{\geq 0}$. Thus, every point of $\mathbb{R}^{\mathcal{G}}_{\geq 0}$ corresponds to a measurement of a fitness value for each genotype. The geometric interpretation of our results about partial orders uses the language of polyhedral cones.

Definition 3. A *convex polyhedral cone* in $\mathbb{R}^k$ is a subset of $\mathbb{R}^k$ cut out by a finite number of linear inequalities.

Like all cones, convex polyhedral cones are closed under the operations of taking nonnegative linear combinations of their elements. For the remainder of this paper, we will refer to convex polyhedral cones simply as cones. Let $\mathcal{P} = (\mathcal{G}, \prec)$ be a partial order (which may be a total order). Then the set of points $x \in \mathbb{R}^{\mathcal{G}}_{\geq 0}$ whose coordinate order satisfies $\mathcal{P}$ defines a cone. To see this, note that this set of points is cut out by a finite number of linear inequalities defined by the partial order. Hence this set is a convex polyhedral cone by definition. We call a cone associated to a partial order in this way an *order cone* and denote it by $C_{\mathcal{P}}$. Further, a linear form f in the variables w_g's defines a hyperplane through the origin, where f takes the value 0. The linear form is positive in the half space on one side of this hyperplane

and negative on the other. Let $H_{f,+}$ be the half space on which f is positive, H_f be the hyperplane on which f is zero, and $H_{f,-}$ be the half space on which f is negative. Then Theorem 2 has the following geometric interpretation. Recall that $f(W) = \sum_{1\leq i\leq t} c_i w_{p_i} - \sum_{1\leq j\leq s} d_j w_{n_j}$ is a linear form with integer coefficients that sum to zero, that is $\sum_{1\leq i\leq t} c_i = \sum_{1\leq j\leq s} d_j$.

Theorem 5. *The cone $C_{\mathcal{P}} \subset H_{f,+}$ if and only if there exists a partition of the set of all genotypes $\mathcal{G}$ into pairs (p_i, n_j) such that $p_i \succ n_j$ for all i, j, where each p_i appears in c_i pairs and each n_j in d_j pairs.*

Proof. Note that $\mathcal{P}$ implies positive f-interaction if and only if $C_{\mathcal{P}} \subset H_{f,+}$. Then the claim follows from Theorem 2. □

Now, we prove the following theorem which is important for statistical analysis of interactions from fitness comparison data, for example, in the analysis of a sample from the probability distribution over the space of partial fitness orders with the aim to quantify the uncertainty of interactions. Specifically, we show that given two uncertain fitness measurements which give distinct linear orders $\mathcal{A}$ and $\mathcal{B}$, either the linear orders differ by an adjacent transposition, that is, by a permutation exchanging two entries in a given linear order, or there exist fitness measurements corresponding to a sequence of linear orders, which are intermediate between $\mathcal{A}$ and $\mathcal{B}$.

Theorem 6. *Let $U \subset \mathbb{R}^k$ be a path connected, open set which has a nonempty intersection with the cones $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$, where $\mathcal{A}$ and $\mathcal{B}$ are linear orders on k elements. Then there exist linear orders $\mathcal{L}_1 = \mathcal{A}, \ldots, \mathcal{L}_n = \mathcal{B}$ such that $\mathcal{L}_i$ and $\mathcal{L}_{i+1}$ differ by one adjacent transposition and $U \cap C_{\mathcal{L}_i} \neq \varnothing$ for each $1 \leq i < n$.*

Proof. We denote cones $C_{\mathcal{L}_i}$ by simply C_i, for all i, throughout the proof. Note that the cones C_i and C_j have a $(k-1)$-dimensional face as their intersection if and only if $\mathcal{L}_i$ and $\mathcal{L}_j$ differ by a single adjacent transposition. Further, note that any path which passes from C_i to C_j without passing through the interior of any other cone must pass through the intersection of the boundaries of C_i and C_j. Either this boundary is an $(k-1)$-dimensional face, and $\mathcal{L}_i$ and $\mathcal{L}_j$ differ by an adjacent transposition, or we pass through a lower dimensional face. In the latter case, this means a neighborhood of the point where we pass through the boundary contains all cones that intersect at this point. Thus, a neighborhood of this point contains a sequence of order cones of linear orders which differ by one adjacent transposition each.

Now, consider a path from a point in $U \cap C_{\mathcal{A}}$ to a point in $U \cap C_{\mathcal{B}}$. Since U is open and path connected, it contains some path of this form, as well as a neighborhood around the path. Consider the sequence of cones this path passes through – by the observation above, we know that if the path passes from C_i to C_j, either $\mathcal{L}_i$ and $\mathcal{L}_j$ differ by an adjacent transposition, or a neighborhood of the point where we pass from C_i to C_j intersects the cones of a sequence of linear orders which differ by one adjacent transposition each. Thus, in either case, a path from $C_{\mathcal{A}}$ to $C_{\mathcal{B}}$ has a neighborhood that intersects the order cones of a sequence of linear orders $\mathcal{L}_1, \ldots, \mathcal{L}_n$ such that $\mathcal{L}_1 = \mathcal{A}$, $\mathcal{L}_n = \mathcal{B}$, and $\mathcal{L}_i$ and $\mathcal{L}_{i+1}$ differ by one adjacent transposition for each $1 \leq i < n$. These linear orders satisfy the claim of the theorem. □

We demonstrate Theorem 6 in Figure 1, where we depicted six regions corresponding to six slices of order cones in $\mathbb{R}^3_{\geq 0}$. Each cone is associated to a liner order of three elements, $x = w_{g_i}, y = w_{g_j}, z = w_{g_k}$. Points inside the cones might correspond to different fitness values satisfying the linear order defining the given cone. Thus, for example a point inside $C_{x>y>z}$ specifies three fitness values associated to three genotypes $g_i, g_j, g_k \in \mathcal{G}$ and such that $w_{g_i} > w_{g_j} > w_{g_k}$.

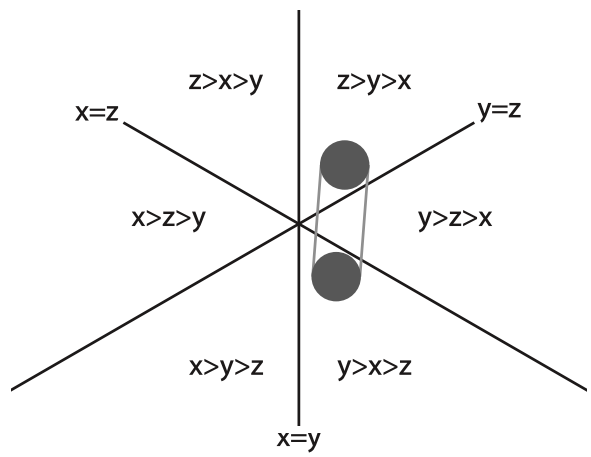

FIGURE 1. We show a slice of $\mathbb{R}^3_{\geq 0}$, demonstrating that $\mathbb{R}^3_{\geq 0}$ is divided into six cones corresponding to the six possible orders of the elements x, y, z. The convex hull of a ball in the order cone $C_{z>y>x}$ and a ball in the order cone $C_{y>x>z}$ passes through the order cone $C_{y>z>x}$. Note that the order $y > z > x$ differs by one adjacent transposition from both $z > y > x$ and $y > x > z$.

As a corollary to Theorem 6, we obtain the following result.

Corollary 1. *Suppose $\mathcal{A}$ and $\mathcal{B}$ are linear orders which imply positive f-interaction. Then there exists a sequence of linear orders $\mathcal{L}_1 = \mathcal{A}, \ldots, \mathcal{L}_k = \mathcal{B}$ such that $\mathcal{L}_i$ and $\mathcal{L}_{i+1}$ for $1 \leq i < n$ all differ by one adjacent transposition and each $\mathcal{L}_i$ either implies positive interaction or does not imply f-interaction.*

Proof. The half-space $H_{f,+}$ is a connected open set which has a nonempty intersection with $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$. Thus, this result follows from Theorem 5. □

Spaces with complicated geometries and combinatorics are known to cause significant difficulties for statistical analysis (Billera, Holmes, and Vogtmann 2001). For example, the space of trees, which is a particular instance of the space of partial orders, required deep mathematical advances to understand basic statistics, such as confidence regions and convex hulls, over the space (Billera, Holmes, and Vogtmann 2001; Gavryushkin and Drummond 2016). Advances in the geometry and combinatorics of such spaces are the stepping stone for efficient statistical methods such as Markov Chain Monte Carlo (Gavryushkin, Whidden, and Matsen 2017; Dinh et al. 2017). In a similar vein, we expect that the approach of this section allows to efficiently study probability distributions over the space of partial fitness orders given empirical fitness measurements. Hence our results provide a theoretical foundation for statistical analysis of partial fitness orders given data.

5. Applications

In this section we illustrate how our results can be used in fitness-based genetic interactions studies. We will use the following notion.

Definition 4. Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be partial orders. Their *intersection* is the partial order $\mathcal{P}_1 \cap \mathcal{P}_2$ such that $x \prec y$ in $\mathcal{P}_1 \cap \mathcal{P}_2$ if and only if $x \prec y$ is in both $\mathcal{P}_1$ and $\mathcal{P}_2$.

To illustrate the above definition consider the partial orders $\mathcal{P}_1 = w_{00} > w_{10} > w_{01}$ and $\mathcal{P}_2 = w_{10} > w_{01} > w_{00}$. Then $\mathcal{P}_1 \cap \mathcal{P}_2$ is the partial order given by $w_{10} > w_{01}$.

Malaria. As a first application, we consider the following three linear fitness orders inferred and analyzed by Ogbunugafor and Hartl (2016) (see also Crona, Gavryushkin, et al. 2017). These fitness orders have been obtained by measuring the growth rate of the parasite *Plasmodium vivax* exposed to various concentrations of the antimalarial drug pyrimethamine.

$$\mathcal{L}_1 : 111 \succ 011 \succ 001 \succ 101 \succ 010 \succ 100 \succ 110 \succ 000$$
$$\mathcal{L}_2 : 111 \succ 011 \succ 001 \succ 010 \succ 100 \succ 101 \succ 110 \succ 000$$
$$\mathcal{L}_3 : 111 \succ 011 \succ 010 \succ 001 \succ 100 \succ 110 \succ 101 \succ 000$$

As noted in (Crona, Gavryushkin, et al. 2017), each of these linear orders implies negative total 3-way interaction, which in our terms is f-interaction where

$$f(W) = w_{000} + w_{011} + w_{101} + w_{110} - w_{001} - w_{010} - w_{100} - w_{111}.$$

Using our approach developed in this paper, we can further analyze and strengthen this conclusion. To do so, we discard the contradictions among the three rank orders, which correspond to three different drug concentrations, and consider only the shared rankings of the genotypes across different concentrations, that is, the partial fitness order $\mathcal{P}$ obtained by taking the intersection $\mathcal{L}_1 \cap \mathcal{L}_2 \cap \mathcal{L}_3$. This intersection has Hasse diagram:

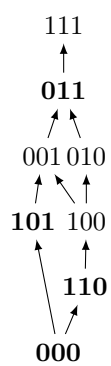

We then apply Theorem 2 and conclude that the obtained partial order $\mathcal{P}$ implies negative f-interaction. Indeed, the three bottom genotypes with even number of ones (depicted in bold) can be matched to the three genotypes with even number of ones above them, and matching 011 to 111 completes the construction of a perfect matching in this graph. This means that even if the only comparisons we are confident about are those that agree across the three drug concentrations, we can still conclude that this system exhibits negative total 3-way interaction. Hence any tuple of fitness values W which satisfies every fitness comparison found in each

linear order will imply interaction. Further, in this case our results from Section 4 imply that any point in the convex hull of the points we have measured will imply negative 3-way interaction. This region allows to incorporate the uncertainties of the measurements and provides a region within which those measurements can vary without breaking the conclusion of interaction. Thus, our results allow to translate the uncertainties at the level of growth rate measurements to the level of interactions.

We conclude by noting that the poset intersection of multiple linear orders which imply f-interaction needs not to imply f-interaction. As an example, consider the linear orders

$$\mathcal{L}_4 : 111 \succ 011 \succ 001 \succ 101 \succ 010 \succ 100 \succ 110 \succ 000$$
$$\mathcal{L}_5 : 111 \succ 011 \succ 010 \succ 101 \succ 001 \succ 100 \succ 110 \succ 000$$

Note that both linear orders imply negative total 3-way interaction – they both map to the Dyck word $NPNPNNPP$. However, the partial order $\mathcal{L}_4 \cap \mathcal{L}_5 = 111 \succ 011 \succ 001, 010, 101 \succ 100 \succ 110 \succ 000$ does not imply total 3-way interaction, since it can be extended to the following linear order:

$$\mathcal{L}_6 : 111 \succ 011 \succ 101 \succ 010 \succ 001 \succ 100 \succ 110 \succ 000,$$

which maps to a non-Dyck word $NPPNNNPP$.

TEM β-lactamase. As a second application, we consider the data produced in the antibiotic resistance study of the TEM-family of β-lactamase (Mira et al. 2015). In the table below we display the average growth rates of 16 genotypes grown in the antibiotic AMP in 12 replicates. The 16 genotypes include the wild type and all combinations of amino acid substitutions in TEM-50.

Genotype	**0000**	**1000**	**0100**	**0010**	**0001**	**1100**	**1010**	**1001**
Average growth rate	1.851	1.570	2.024	1.948	2.082	2.186	0.051	2.165
Genotype	**0110**	**0101**	**0011**	**1110**	**1101**	**1011**	**0111**	**1111**
Average growth rate	2.033	2.198	2.434	0.088	2.322	0.083	0.034	2.821

TABLE 2. Average growth rates of the 16 genotypes grown in the antibiotic AMP.

By using these growth rate as a measure of fitness, we deduce that this system has positive total 4-way interaction by directly computing:

$$f(W) = w_{0000} + w_{0011} + w_{0101} + w_{0110} + w_{1001} + w_{1010} + w_{1100} + w_{1111} - (w_{0001} + w_{0010} + w_{0100} + w_{1000} + w_{0111} + w_{1011} + w_{1101} + w_{1110})$$

and noting that $f(W) > 0$ for the values in Table 2.

In the following, we demonstrate how our approach can be used to analyze the uncertainty of the conclusion that the system has positive total 4-way interaction. We start by ranking the genotypes according to their growth rates. This yields the linear order described in the following table (see Table 5 in (Mira et al. 2015)): In the following we assume that the average growth rates of these 16 genotypes

Genotype	**0000**	**1000**	**0100**	**0010**	**0001**	**1100**	**1010**	**1001**
Rank	11	12	9	10	7	5	15	6
Genotype	**0110**	**0101**	**0011**	**1110**	**1101**	**1011**	**0111**	**1111**
Rank	8	4	2	13	3	14	16	1

TABLE 3. Ranking of the 16 genotypes grown in the antibiotic AMP according to their average growth rates listed in Table 2.

vary but that they always preserve the ranking listed in Table 3. This assumption implies that the average growth rates map to the word

$$\omega = PPNPPPNPNNPNNNPN$$

according to the mapping described in Section 1. Reading ω from left to right, we notice that ω is a Dyck word. Hence, Theorem 1 can be used to confirm that the system has positive interaction. However, comparing the ranking and the average growth rates yields the following observation. The difference between the average growth rates assigned to the genotypes 0101, 1100, and 1001 as well as to the pairs of genotypes 0110, 0100 and 1110, 1011 is smaller comparing to the differences between all other average growth rates assigned to the remaining genotypes. See Figure 2 for a visual summary of this observation. These small differences in the

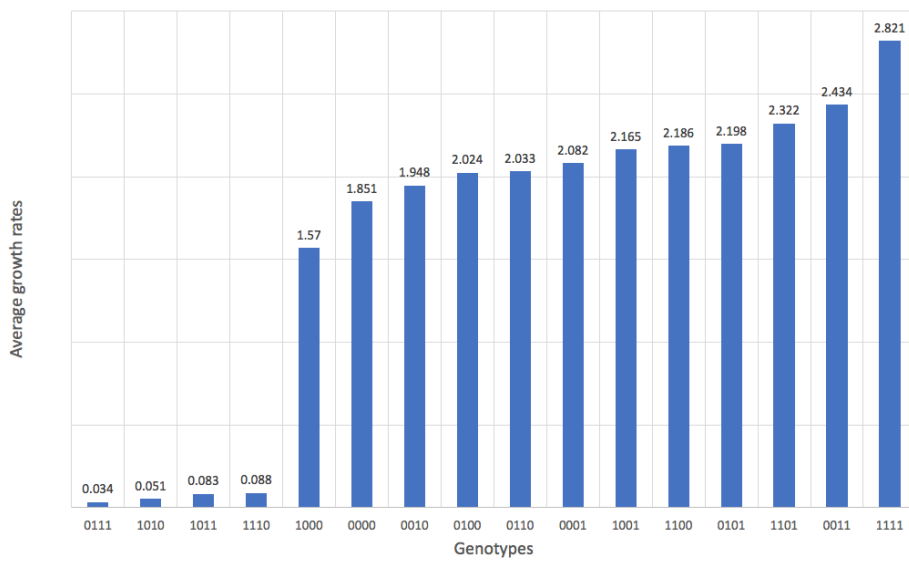

FIGURE 2. Average growth rates of 16 genotypes grown in the antibiotic AMP.

growth rates highlight the difficulty of finding accurate and robust linear orders among the genotypes.

$i = j$	p_i	n_j	$w_{p_i} - w_{n_j}$
1	1111	1101	0.499
2	0011	0001	0.352
3	0101	0100	0.174
4	1100	0010	0.238
5	1001	1000	0.595
6	0110	1110	1.945
7	0000	1011	1.768
8	1010	0111	0.017

TABLE 4. Perfect matching of genotypes obtained from the partial fitness order according to the average growth rate.

In contrast, the partial order approach, developed in this work, confirms again that the system has positive total 4-way interaction using fewer and more pronounced comparisons. To illustrate this point, we follow the notations from Theorem 2 and consider the following 8 genotypes p_i's and 8 genotypes n_j's as in Table 4. The linear form $f(W)$ is as above. The rows in Table 4 indicate a perfect matching among the two sets of genotypes. Thus, for example the first row indicates that the average growth rates associated to 1101 and 1111 are such that $w_{1101} > w_{1111}$. Similarly, for the other rows. From these 8 comparisons of the type $w_{p_i} > w_{n_j}$ alone, one deduces that the system has positive total 4-way interaction. Moreover, comparing with the actual average growth rates from Table 2 one can observe that the 8 differences $w_{p_i} - w_{n_j}$ (see third column in Table 4) are bigger then the differences between the averages growth rates of the critical genotypes mentioned above. Since these differences are more significant, they provide a more reliable conclusion in the empirical setting. Finally, notice that any other perfect matching between the 16 genotypes satisfying $w_{p_i} > w_{n_j}$ would equally well yield to the same conclusion. In summary, an advantage of the partial order approach is that it relies on fewer pairwise inequalities, which makes the approach more practical.

6. Discussion and future directions

Understand genetic interactions from fitness measurements associated to genotypes represents a major challenge in evolutionary biology. In this work, we have focused on the case where only fitness comparisons between certain genotypes are available. This is a common assumption in practice, as there might be more uncertainty or cost involved in deducing some fitness measurements than others. Furthermore, this approach allows to exclude the measurements that have high uncertainty.

In this setting, we present a new algorithm to detect genetic interactions from partial fitness orders in an efficient way. Moreover, we derive a geometric characterization of the class of partial orders that imply interactions. This description, involving the geometry of convex polyhedral cones, provides a solid framework to develop statistical analysis of genetic interactions from partial fitness data.

Our work inspires a number of questions which remain open. First, while we are able to characterize the set of partial orders which imply interaction, our characterization is in terms of a matching in a separate graph which we can construct from our partial order. Thus, it is not immediately clear how to relate properties of a partial order to the probability that this partial order implies interaction. For instance, how does the number and type of incomparable pairs in linear orders affect whether interaction is implied?

This work opens up a number of possible approaches to quantifying interactions in genomic data sets, along with the uncertainties that come from the stochastic nature of fitness data. In particular, now that we have an efficient test for sign epistasis, we are well placed to estimate the prevalence of sign epistasis in the biological world and to investigate the mechanisms which cause sign epistasis.

Acknowledgments

We are grateful to Bernd Sturmfels for hosting us at the Max Planck Institute for Mathematics in the Sciences in Leipzig, where this work started.

References

Beerenwinkel N, Pachter L, and Sturmfels B (2007). "Epistasis and Shapes of Fitness Landscapes". *Statistica Sinica* 17, pp. 1317–1342.

Billera LJ, Holmes SP, and Vogtmann K (2001). "Geometry of the Space of Phylogenetic Trees". *Adv. Appl. Math.* 27.4, pp. 733–767. ISSN: 0196-8858. DOI: 10.1006/aama.2001.0759.

Chou HH, Chiu HC, Delaney NF, Segrè D, and Marx CJ (2011). "Diminishing returns epistasis among beneficial mutations decelerates adaptation". *Science* 332.6034, pp. 1190–1192.

Crona K, Gavryushkin A, Greene D, and Beerenwinkel N (2017). "Inferring Genetic Interactions From Comparative Fitness Data". *bioRxiv; eLife* under revision. DOI: 10.1101/137372.

Crona K and Luo M (2017). "Higher order epistasis and fitness peaks". arXiv (accessed on August 8, 2017): 1708.02063 (q-bio.QM).

Diestel R (2016). *Graph Theory*. Springer Verlag.

Dinh V, Bilge A, Zhang C, and Matsen IV FA (2017). "Probabilistic Path Hamiltonian Monte Carlo". arXiv: 1702.07814 [q-bio.PE].

Gavryushkin A and Drummond AJ (2016). "The space of ultrametric phylogenetic trees". *J. Theor. Biol.* 403, pp. 197–208. ISSN: 0022-5193. DOI: 10.1016/j.jtbi.2016.05.001.

Gavryushkin A, Whidden C, and Matsen 4th FA (2017). "The combinatorics of discrete time-trees: theory and open problems". en. *J. Math. Biol.* ISSN: 0303-6812, 1432-1416. DOI: 10.1007/s00285-017-1167-9.

Gong LI, Suchard MA, and Bloom JD (2013). "Stability-mediated epistasis constrains the evolution of an influenza protein". *eLife* 2, e00631.

Hopcroft JE and Karp RM (1971). "An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs". *SIAM J. Comput.* 2.4, pp. 225–231. DOI: 10.1137/0202019.

Kvitek DJ and Sherlock G (2011). "Reciprocal sign epistasis between frequently experimentally evolved adaptive mutations causes a rugged fitness landscape". *PLoS genetics* 7.4, e1002056.

Mira PM et al. (2015). "Rational design of antibiotic treatment plans: a treatment strategy for managing evolution and reversing resistance". en. *PLoS One* 10.5, e0122283. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0122283.

Ogbunugafor CB and Hartl D (2016). "A pivot mutation impedes reverse evolution across an adaptive landscape for drug resistance in Plasmodium vivax". en. *Malar. J.* 15, p. 40. ISSN: 1475-2875. DOI: 10.1186/s12936-016-1090-3.

Sohail M et al. (2017). "Negative selection in humans and fruit flies involves synergistic epistasis". *Science* 356.6337, pp. 539–542.

Stanley RP (2001). *Enumerative Combinatorics, Volume 2*. Cambridge University Press.

Stein W et al. (2008). "Sage: Open source mathematical software". *7 December 2009*.

The On-Line Encyclopedia of Integer Sequences (2017). https://oeis.org/. Accessed: 2017-8-4. URL: https://oeis.org/.

Weinreich DM, Lan Y, Wylie CS, and Heckendorn RB (2013). "Should evolutionary geneticists worry about higher-order epistasis?" en. *Curr. Opin. Genet. Dev.* 23.6, pp. 700–707. ISSN: 0959-437X, 1879-0380. DOI: 10.1016/j.gde.2013.10.007.

Weinreich DM, Watson RA, and Chao L (2005). "Perspective: Sign epistasis and genetic constraint on evolutionary trajectories". en. *Evolution* 59.6, pp. 1165–1174.