Published in final edited form as:

IEEE Trans Evol Comput. 2016 April ; 20(2): 263–274. doi:10.1109/TEVC.2015.2454857.

Tunably Rugged Landscapes With Known Maximum and Minimum

Narine Manukyan,

Department of Computer Science, University of Vermont, Burlington, VT 05401 USA

Margaret J. Eppstein, and

Department of Computer Science, University of Vermont, Burlington, VT 05401 USA

Jeffrey S. Buzas

Department and Mathematics and Statistics, University of Vermont, Burlington, VT 05401 USA

Abstract

We propose NM landscapes as a new class of tunably rugged benchmark problems. NM landscapes are well defined on alphabets of any arity, including both discrete and real-valued alphabets, include epistasis in a natural and transparent manner, are proven to have known value and location of the global maximum and, with some additional constraints, are proven to also have a known global minimum. Empirical studies are used to illustrate that, when coefficients are selected from a recommended distribution, the ruggedness of NM landscapes is smoothly tunable and correlates with several measures of search difficulty. We discuss why these properties make NM landscapes preferable to both NK landscapes and Walsh polynomials as benchmark landscape models with tunable epistasis.

Index Terms

Benchmark landscapes; fitness landscapes; NK landscapes; Walsh polynomials

I. Introduction

Simulated landscapes are widely used for evaluating search strategies, where the goal is to find the landscape location with maximum fitness value [1], [2]. Without loss of generality and for notational simplicity, we assume function maximization, rather than minimization, throughout this paper.

NK landscapes [1] have been classic benchmarks for generating landscapes with epistatic interactions. They are described by two parameters: 1) N specifies the number of binary features and 2) K specifies that the maximum degree of epistatic interactions among the features is $K + 1$ [3]. NK landscapes have been used in many applications (see [4]–[8]) and widely studied in theory (see [9]–[13]), as they can generate landscapes with tunable

ruggedness by varying K . However, NK landscapes have several limitations. Buzas and Dinitz [11] recently showed that the expected number of local peaks in NK landscapes rises in large discrete jumps as K is increased, but actually decreases as a function of the number of interaction terms for a given K (Fig. 1, red lines). Additionally, the problem of finding the location and value of the global optimum of unrestricted NK landscapes with $K > 1$ is NP-complete [10] (although for restricted classes one can use dynamic programming [10], [14] or approximation algorithms [10]). NK landscapes have only been defined for binary alphabets.

Walsh polynomials are a superset of NK landscapes that overcome some of the limitations of NK landscapes. For example, they allow more explicit control over which interaction terms are present. The problem of finding the global maximum value of a Walsh polynomial is also NP-complete, although a restricted subset of Walsh polynomials has a known global maximum [15]. However, even in this case finding the global minimum is still NP-complete, preventing proper normalization by the range of fitnesses in the landscapes. As with NK landscapes, Walsh polynomials are only defined for binary alphabets.

In this paper, we introduce a different, more flexible subset of general interaction models that we dub NM landscapes. Like NK landscapes and Walsh polynomials, NM landscapes incorporate epistatic feature interactions. However, NM landscapes also: 1) include epistasis in a natural and transparent manner; 2) have known value and location of the maximum fitness; 3) work with alphabets of any arity, including discrete and real-valued alphabets; 4) with additional constraints have known value and location of the minimum fitness; and 5) when coefficients are chosen properly, have relatively smoothly tunable ruggedness. In Section II, we introduce the general class of parametric interaction models and Walsh polynomials, then in Section III we define NM landscapes and prove the properties 1)–4) above. In Sections IV and V, we describe experiments and results that demonstrate property 5) above. In Section VI, we discuss the importance of these properties and point out several advantages of NM landscapes as benchmark problems for studying search in tunably rugged landscapes.

II. Interaction Models and Walsh Polynomials

Walsh polynomials provide a mathematical framework for defining any real-valued function on bit strings [16], [17]. A Walsh polynomial has the following form:

$$f(\mathbf{y}) = \sum_{j=0}^{2^q-1} \omega_j \psi_j(\mathbf{y}) \quad (1)$$

where q is the length of the bit string $\mathbf{y}$, each bit $y_i \in \{0, 1\}$, and each $\omega_j \in \mathbb{R}$. The Walsh function $\psi_j(\mathbf{y})$ corresponding to the j th partition is defined as

$$\psi_j(\mathbf{y}) = \begin{cases} 1, & \text{if } \mathbf{y} \wedge j_2 \text{ has even parity} \\ -1, & \text{otherwise} \end{cases} \quad (2)$$

where j_2 denotes the binary representation of j .

NK landscapes are a subset of Walsh polynomials. Walsh polynomials have a one-to-one correspondence with the more general class of general parametric interaction models, when such models are restricted to binary alphabets [17].

A fitness landscape F can be defined for N features using a general parametric interaction model of the form

$$F(\mathbf{x}) = \sum_{k=1}^m \beta_{U_k} \prod_{i \in U_k} x_i \quad (3)$$

where m is the number of terms, and each of m coefficients $\beta_{U_k} \in \mathbb{R}$. For $k = 1, \dots, m$, $U_k \subseteq \{1, 2, \dots, N\}$, where U_k is a set of indices of the features in the k th term, and the length $|U_k|$ is the order of the interaction. We adopt the convention that when $U_k = \emptyset$, $\prod_{j \in U_k} x_j \equiv 1$. If the parametric interaction model is defined on a binary alphabet, we adopt the convention that binary values are represented as $x_i \in \{-1, 1\}$ (rather than $\{0, 1\}$, as in Walsh polynomials). However, general parametric interaction models are also well defined for discrete valued features with higher arities as well as for real-valued alphabets and provide a more intuitive way of representing epistatic interactions among features.

A more readable notation for (3) is as

$$F(\mathbf{x}) = \beta_0 + \sum_{i=1}^N \beta_i x_i + \sum_{i=1}^{N-1} \sum_{j=i+1}^N \beta_{i,j} x_i x_j + \sum_{i=1}^{N-2} \sum_{j=i+1}^{N-1} \sum_{k=j+1}^N \beta_{i,j,k} x_i x_j x_k + H \quad (4)$$

where we only explicitly show up to third-order interactions and H represents higher order interactions up to some maximum order $M \leq N$. Note that some β_{U_k} parameters may be zero, so not all terms need be present.

For example, consider a model with $N = 2$ loci and $U_1 = \emptyset$, $U_2 = \{1\}$, $U_3 = \{2\}$, and $U_4 = \{1, 2\}$. The interaction model for this example is

$$F(\mathbf{x}) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_{1,2} x_1 x_2 \quad (5)$$

where β_0 is the average value of all fitnesses in the landscape, β_1 and β_2 are the coefficients of the main effects of the binary features x_1 and x_2 , and $\beta_{1,2}$ is the coefficient of the second order epistatic interaction $x_1 x_2$. The Walsh polynomial corresponding to (5) is

$$f(\mathbf{y}) = \omega_0 \psi_0(\mathbf{y}) + \omega_1 \psi_1(\mathbf{y}) + \omega_2 \psi_2(\mathbf{y}) + \omega_3 \psi_3(\mathbf{y}) = \beta_0 \psi_0(\mathbf{y}) - \beta_1 \psi_1(\mathbf{y}) - \beta_2 \psi_2(\mathbf{y}) + \beta_{1,2} \psi_3(\mathbf{y}) \quad (6)$$

where

$$y_i = \begin{cases} 1, & \text{when } x_i = 1 \\ 0, & \text{when } x_i = -1. \end{cases} \quad (7)$$

Notice that there is a one-to-one correspondence of each term in (5) with each term in (6) but the signs of the coefficients are different. Specifically, for the example above

$$\beta_0 = \omega_0, \quad \beta_1 = -\omega_1, \quad \beta_2 = -\omega_2, \quad \beta_{1,2} = \omega_3. \quad (8)$$

A random point selected in the search space of a Walsh polynomial can be forced to be the global maximum by properly adjusting the sign of each of the nonzero Walsh coefficients, with the maximum fitness value equal to the sum of the absolute values of all Walsh coefficients [15]. However, the location and the value of the global minimum remains unknown.

General parametric interaction models are the standard models used in statistics to study effects of multiple features on an outcome (see [18]). They are easy to define and the interactions are transparent and easy to interpret (unlike in NK landscapes and Walsh polynomials). For example, the interaction terms present in (5) are clearly evident, whereas the Walsh functions ψ_i in (6) obscure this. To date, general parametric interaction models have received very little attention in the evolutionary computation literature, with notable exceptions [19]–[21].

Buzas and Dinitz [11] showed that for every NK landscape with a given K one can create an equivalent parametric interaction model, where the maximum order of interactions is $K + 1$. They show that the NK algorithm dictates that the interaction model contain all main effects and subinteractions contained in higher order interactions. For example, if a nonzero interaction coefficient $\beta_{i,j,k}$ is present in an NK landscape, then there will generally be nonzero coefficients $\beta_i, \beta_j, \beta_k, \beta_{i,j}, \beta_{i,k}, \beta_{j,k}$ (there is an infinitesimally small probability that one or more of these coefficients may be zero). For the classic NK model where K is constant and $K \ll N$, main effect coefficients have the largest expected magnitude, second-order interactions have larger expected magnitude than third-order interactions, and so on [11]. Thus, NK landscapes are a very restricted subset of Walsh polynomials and the more general class of parametric interaction models.

III. NM Landscapes

The class of Walsh polynomials is a subset of the larger class of general interaction models. Here, we introduce a different subset of general interaction models called NM landscapes, where N is the number of features and all interactions in the model are of order M .

Definition 1

NM models comprise the set of all general interaction models specified by (3), with the added constraints that: 1) all coefficients β_{U_k} are non-negative; 2) each feature value x_i ranges from negative to positive values; and 3) the absolute value of the lower bound of the range is less than or equal to the upper bound of the range of x_i .

In this paper, each β_{U_k} is randomly created as

$$\beta_{U_k} = e^{-\text{abs}(\mathcal{N}(0, \sigma))} \quad (9)$$

where $\mathcal{N}(0, \sigma)$ is a random number drawn from a Gaussian distribution with 0 mean and standard deviation of σ , which results in fitnesses that are symmetric around 0 (Fig. 2). As the value of σ increases, the means and standard deviations of the coefficients decrease, which results in a smaller range of fitness values and increasing clumping of fitness values (Fig. 2). In contrast, when coefficients are drawn from a uniform distribution in the range $[0, 1]$, the fitnesses are skewed left (Fig. 3). *NM* landscapes offer several desirable properties, as described in the following.

Proposition 1

NM landscapes with a binary alphabet have a known global maximum.

Proof—By Definition 1, $\beta_{U_k} > 0$ for all nonzero terms. Thus, the maximum possible value for each term ($\beta_{U_k} \prod_{j \in U_k} x_j$) in an *NM* landscape with a binary alphabet $x_i \in \{-1, 1\}$ is achieved when

$$x_i = 1, \quad \forall i = 1, \dots, n \quad (10)$$

and the value of the global maximum is

$$F_{\max} = \beta_0 + \sum_{i=1}^N \beta_i + \sum_{i=1}^{N-1} \sum_{j=i+1}^N \beta_{i,j} + \sum_{i=1}^{N-2} \sum_{j=i+1}^{N-1} \sum_{k=j+1}^N \beta_{i,j,k} + \sum_{\forall \beta_{U_h}} \beta_{U_h} \quad (11)$$

where β_{U_h} are the coefficients of all the remaining higher order interactions. Note that the calculation of (11) has time complexity of $O(m)$, where m is the number of terms in the model.

Proposition 2

NM landscapes can be defined on discrete alphabets of any arity or on real-valued alphabets, and the value and location of a global maximum is independent of the discretization of the alphabet.

Proof—By Definition 1, all coefficients are non-negative, therefore the maximum $F_{\max}$ of an NM landscape with any discrete or real-valued alphabet $x_{a,b}$ defined in the range $[-a, b]$, where $0 < a < b$, occurs when $x_i = b, \forall i = 1, \dots, N$ and its value is

$$F_{\max} = \beta_0 + \sum_{i=1}^N \beta_i b + \sum_{i=1}^{(N-1)} \sum_{j=i+1}^N \beta_{i,j} b^2 + \sum_{i=1}^{(N-2)} \sum_{j=i+1}^{(N-1)} \sum_{k=j+1}^N \beta_{i,j,k} b^3 + \sum_{\forall \beta_{U_h}} \beta_{U_h} b^{|U_h|} \quad (12)$$

where β_{U_h} are the coefficients of all the remaining higher order interactions and the lengths $|U_h|$ are the orders of these interactions. Thus, the magnitude of $F_{\max}$ is a function of b , but is independent of the arity of the alphabet.

For alphabets where $a = b$, β_0 represents the mean value of the landscape. We note that it is trivial to extend this proposition and proof to NM landscapes with heterogeneous alphabets (i.e., different ranges and/or arities for each feature variable), as long as the lower bound for each feature is negative, the upper bound is positive, and the absolute value of the lower bound is less than the upper bound. However, for notational simplicity we only demonstrate the proof for homogeneous alphabets. We refer to the above described general NM landscapes as Type I NM landscapes.

We conjecture that changing the arity of features in NM landscapes does not change the number, locations, or values of the local peaks or global minima, because higher arity alphabets simply sample the same landscape at a higher resolution that interpolates between the local peaks (empirical data, not shown, supports this conjecture).

Note that interaction models with all non-negative interaction coefficients β_{U_k} , but no negative feature values, generate unimodal landscapes. However, since alphabets in NM landscapes are defined to include both negative and positive features, NM landscapes have multiple local optima whenever any interaction terms are included.

Proposition 3

NM landscapes that include all main effects have exactly one global maximum.

Proof—By Proposition 1, a maximum fitness $F_{\max}$ of an NM landscape is achieved at point $\mathbf{x} = [b, b, \dots, b]$. Let $\mathbf{y} = [y_1, y_2, \dots, y_n]$, where $\mathbf{y} \neq \mathbf{x}$ (i.e., there exists at least one i such that $y_i < b$). Since each $x_i \in [-a, b]$ and $a < b$ (by Definition 1) then $\beta_i * y_i < \beta_i * b$ and the value of the interaction model at point $\mathbf{y}$ will be strictly less than the value at point $\mathbf{x}$. Thus, $\mathbf{x}$ is the only global maximum.

Proposition 4

NM landscapes that include only even order terms with alphabets in the range $[-b, b]$ are symmetric and have exactly two global maxima at maximal distance apart in feature space.

Proof—Since $x_i^{2t} = (-x_i)^{2t}, \forall t \in \mathbb{N}$, then for all NM landscapes with only even order terms, $F(\mathbf{x}) = F(-\mathbf{x})$ for each pair of points $\mathbf{x} = [x_1, x_2, \dots, x_N]$ and $-\mathbf{x} = [-x_1, -x_2, \dots, -x_N]$. Thus,

NM landscapes with only even order interactions and alphabets in $[-b, b]$ range are symmetric and the two global maxima are at locations $[b, b, \dots, b]$ and $[-b, -b, \dots, -b]$, which are the maximum distance away from each other in the feature space.

When the value of the global maximum of the landscape is known one can partially normalize fitnesses to the range $[-1, 1]$ using the following formula:

$$F = \frac{F - F_{\min}}{F_{\max} - F_{\min}}. \quad (13)$$

However, proper normalization of fitnesses to the interval $[0, 1]$ also requires prior knowledge of the global minimum of the landscape, as

$$F = \frac{F - F_{\min}}{F_{\max} - F_{\min}}. \quad (14)$$

To this end, we define subsets of NM landscapes that have a known global minimum. While there are many ways to do this, below we present two such subsets.

Proposition 5

NM landscapes that include only main effects with odd indices (e.g., x_1, x_3, x_5 , etc.) and any terms with an odd number of odd indices (e.g., $x_1x_2, x_1x_3x_5, x_1x_3x_6x_7$, etc.) and alphabets in the range $[-1, 1]$ have a global minimum located at point $[-1, 1, -1, 1 \dots]$. For example, models of this form including up to $M=3$ order terms are given by

$$F(x) = \beta_0 + \sum_{i \text{ odd}}^N \beta_i x_i + \sum_{\substack{i \text{ odd} \\ j \text{ even}}}^N \beta_{i,j} x_i x_j + \sum_{\substack{i \text{ odd} \\ j \text{ even} \\ k \text{ even}}}^N \beta_{i,j,k} x_i x_j x_k + \sum_{\substack{i \text{ odd} \\ j \text{ odd} \\ k \text{ odd}}}^N \beta_{i,j,k} x_i x_j x_k. \quad (15)$$

Proof—At the point $[-1, 1, -1, 1 \dots]$ all terms with an odd number of odd indices will have a negative sign, as the product of an odd number of negative numbers is negative. Thus, this point is the global minimum of the landscape with value

$$F_{\min} = - \left(\beta_0 + \sum_{i \text{ odd}}^N \beta_i + \sum_{\substack{j \text{ odd} \\ j \text{ even}}}^N \beta_{i,j} + \dots \right) \quad (16)$$

(where only terms up through second order are explicitly shown above).

We refer to the *NM* landscapes defined in Proposition 5 as Type II *NM* landscapes. Note that Type II *NM* landscapes can easily be extended to alphabets in the range $[-a, b]$, where $a \leq b$, although for notational simplicity we have limited the range to $[-1, 1]$ in the above proof.

Proposition 6

NM landscapes with only odd order terms and alphabets in the range $[-a, b]$, where $a \leq b$, have a global minimum located at $[-a, -a, \dots, -a]$.

Proof—By Definition 1, $\beta_{U_k} > 0$ for all nonzero terms, $x_i \in [-a, b] \forall i$, and $a \leq b$. Therefore, the value of each term $T_k = \beta_{U_k} * x_{U_k}$ has to be $-|a|^{|U_k|}$. When all the features $x_i = -a$, $T_k = -|a|^{|U_k|}$. Therefore, the point $[-a, -a, \dots, -a]$ is a global minimum with value

$$F_{\min} = - \left(\beta_0 + \sum_{i=1}^N \beta_i a + \sum_{i=1}^{(N-2)} \sum_{j=i+1}^{(N-1)} \sum_{k=j+1}^N \beta_{i,j,k} a^3 + \sum_{\forall \beta_{U_h}} \beta_{U_h} a^{|U_h|} \right). \quad (17)$$

We refer to the *NM* landscapes defined in Proposition 6 as Type III *NM* landscapes. When $a = b$ the global maximum and minimum of Type III *NM* landscapes have the same absolute value, but opposite signs. Because *NM* landscapes allow only non-negative coefficients but require both positive and negative feature values, we are thus able to construct *NM* landscapes with known maximum and known minimum, enabling normalization of fitnesses to the range $[0, 1]$ by (14). In contrast, Walsh polynomials allow both positive and negative coefficients, but have only non-negative feature values. Thus, while it is possible to manipulate the signs of the Walsh coefficients to specify the location of the global maximum [15], the global minimum of a Walsh polynomial is still unknown, even if one restricts the order of the interactions as in Types II or III *NM* landscapes.

IV. Experiments

A. Ruggedness

We illustrate how ruggedness changes on binary *NM* landscapes with coefficients drawn from the distribution in (9). Since we assess the ruggedness of these models using exhaustive search, we limit our experiments to $N \leq 15$.

In one set of experiments, we generated random Type I “master” *NM* models, including terms for all N main effects and the $\sum_{i=1}^N \binom{N}{i}$ possible interaction terms (e.g., for $N=10$ there are 1023 overall terms; 1013 interaction terms plus ten main effects). We then systematically created subsets of each of these master models that include an increasing number m of terms from the master model, as follows. We started with a base model that includes all main effects. Random second-order terms were then added in groups of ten (or less if there are not ten left). After we had included all of the second-order terms, we began adding randomly selected groups of ten third-order terms, and so on, until the single N -order interaction term was included. We performed these incremental explorations of 100 master *NM* models for each of $N=10$ with $\sigma=10$, and $N=15$ with $\sigma \in \{15, 30, 100\}$.

In another set of experiments we similarly created 100 master Type II *NM* landscapes according to (15), with $N = 10$ and $\sigma = 10$. We created increasing subsets from the master models, as described above for the Type I landscapes.

We computed two standard measures of landscape ruggedness [22], [23]: 1) we counted the number of local peaks (where a local peak is defined as any point whose fitness value is greater than that of all of its neighbors) and 2) we computed the lag 1 autocorrelation of random walks through the landscapes, where lag 1 autocorrelation is defined as the correlation coefficient between adjacent fitness values in a random walk.

B. Distribution of Fitnesses and Local Peaks

We generated representative *NM* landscapes with $N = 10$ and $\sigma = 10$ for each of $M \in \{1, 2, 3, 4, 6, 10\}$ for both Types I and II *NM* landscapes. We visualize these landscapes by plotting the fitnesses of all points in the landscape as a function of their distances (in feature space) to the global optimum, indicating which are local optima.

C. Basin of Attraction of Global Optimum

We assessed the size of the basin of attraction of the global maximum of Types I and III *NM* and *NK* landscapes for different values of $K = M - 1 \in \{1, 9\}$ and $N = 10$. The fitness matrix of *NK* landscapes is generated from random uniform numbers in the $[0, 1]$ range. We calculated the size of the global basin of attraction as a weighted sum of the points in the landscape that can reach the global maximum using only hill climbing. Each point was weighted based on the percentage of its immediate neighbors with higher fitnesses that were also in the basin of attraction of the global maximum.

D. Searchability of the Landscapes

We assessed how searchable *NM* landscapes are using simple genetic algorithms (GAs). In all the experiments we used a GA with $N = 32$, $\sigma = 32$, population size 256, uniform crossover rate 0.7, uniform mutation and the number of random seeds of 32 (these parameter values were selected to be the same as in [15]).

We studied search on Type I *NM* landscapes with $M = 2$ and $P \in \{0, 0.1, \dots, 1\}$ proportions of all possible second-order interactions. We studied the search on Type III *NM* landscapes with $M \in \{1, 3, 5\}$ including all possible main effects and odd order interactions of order M .

V. Results

The number of local peaks L in *NM* landscapes increases relatively smoothly as we increase the number of terms (m) in both Types I and II *NM* landscapes (i.e., the regions between the vertical lines on Fig. 4) and as we increase the maximum order of interactions M (i.e., as we cross a vertical line on Fig. 4).

Note that the average number of local peaks for a given m in both Types I and III landscapes is on the same order of magnitude as the expected number of local peaks in *NK* models with the same N and $K + 1 = M$ [compare Fig. 4(a) and (c) to Fig. 1].

Similarly, the lag 1 autocorrelation of random walks through both Types I and III NM landscapes decreases relatively smoothly as the number of terms m is increased in models with a given M , as well as when the maximum order of interactions M is increased [Fig. 4(b) and (d)], where lower autocorrelation corresponds to greater ruggedness. Notice that, especially for small M , the increase in ruggedness (as measured by both the number of local peaks and the lag 1 autocorrelation) asymptotically slows as the number of terms m increases (Fig. 4).

The larger the σ values in (9), the smaller the range of fitness values in the landscape (Fig. 2), resulting in larger standard deviations of both the number of local peaks in the landscape (shown in Fig. 5 for $N=15$) and the autocorrelation (not shown).

We show the fitnesses of all points in representative NM landscapes with $N=10$ and $\sigma=10$ as a function of their distances in feature space to the global maximum for both Type I (Fig. 6) and Type II (Fig. 7) NM landscapes. The global maximum is indicated by the leftmost red $\times$ in each panel and the remaining red $\times$'s are suboptimal local peaks. As we increase the maximum order of interactions M , the fitness difference between the global maximum and other points in the landscape increases; this effect is amplified for Type II NM landscapes (Fig. 7) relative to Type I NM landscapes (Fig. 6).

In both Types I and II NM landscapes the distance in feature space between the global maximum and the nearest local peak generally decreases with increasing M and the sizes of the basin of attraction for the global maximum decreases (Fig. 8). Our results show that NK and NM landscapes have similar sizes of the basin of attraction for the global maximum for small and large $K=(M-1)$. However, the size of the basin of attraction for the global maximum of both Types I and II NM landscapes decreases with increasing M rapidly for $M=5$ then levels out, while for NK landscapes the decrease is more gradual (see Fig. 8).

The difficulty of a GA search on NM landscapes also increases with increasing m and M , by several measures of difficulty. When the maximum order of interactions $M=2$ and the proportion P of all the possible second-order interactions increases from 0 to 1 in 0.1 increments, our results show that the mean of the best fitnesses found by the GA decreases, although above $P=0.7$ there is little if any further change in difficulty (Fig. 9). We speculate that this might correspond to the periodically asymptotic pattern in the ruggedness noted previously as the maximum number of terms m approaches the maximum possible for a given M [Fig. 4(a) and (b)]. Results are shown for only the first 30 generations, after which no further improvement was observed. Histograms of the Hamming distances between the best solutions found by the GA and the global maximum are shown for 32 runs of the GA, for different proportions of the possible second-order interactions (Fig. 10). For unimodal landscapes ($M=1$), the GA found the global optimum in all 32 runs (Fig. 10). For more rugged landscapes the global optimum was also found in some runs, and surprisingly the proportion of times it was found increased from $P=0.2$ to $P=1$. However, as the ruggedness increased, those runs in which the best individuals were suboptimal generally became stuck farther and farther from the global optimum (note how the distributions become increasingly spread out to the right, as you view the panels in Fig. 10 from top to bottom).

When fitnesses are not normalized, a higher maximum order of interactions M results in higher raw fitnesses [Fig. 11(a)]. This is due to the fact that summing more interaction terms result in higher ranges of fitness (Figs. 6 and 7). However, when fitnesses are properly normalized by (14) to the range $[0, 1]$, increasing the maximum order of interactions in the model decreases the values of the best individuals' fitnesses found [Fig. 11(b)], reflecting the fact that GA search becomes more difficult at higher M .

While the proportion of times that a GA was able to find the global maximum out of 32 runs decreased as the maximum order of interactions M increased [Fig. 12(a)], the means and standard deviations of the distances between the best solutions found by the GA and the global maximum increased [Fig. 12(b)]. Normalizing by (13), rather than (14) exaggerates both the apparent relative generational increase in fitnesses in the GA and the variance in fitnesses across different random landscapes with the same m and M (Fig. 11). This illustrates how knowing the global minimum can help to accurately assess the mean and standard deviation of the relative increase in fitnesses and fairly compare search results on different NM landscapes.

VI. Discussion

In this paper, we introduce NM landscapes, which are parametric interaction models that: 1) have non-negative coefficients and 2) are well defined for feature alphabets of any arity (from binary to real-valued), as long as 3) the minimum value in the alphabet is negative with absolute value less than or equal to the positive maximum. This combination of constraints ensures that a global maximum is located at the point where all decision variables have their maximum value, with the optimal value equal to the sum of the model coefficients. By further restricting which combinations of interactions are present, various subsets of NM models also have known location and value of the global minimum (we illustrate two such subsets, which we refer to as Types II and III NM landscapes). By using an appropriate non-negative distribution for the coefficients, the resulting NM landscape models have relatively smoothly tunable ruggedness. Epistatic terms are transparently represented in NM landscapes, making it trivial to control or analyze exactly which terms and interactions are present. In the following, we discuss why these various aspects of NM landscapes are valuable, and how they offer advantages over NK landscapes and Walsh polynomials as epistatic benchmark problems.

A. Value of Finely Tunable Epistasis

Although NK landscapes have been widely used as benchmark problems with varying degrees of epistasis, there are many potential applications that require more fine control over which terms are present or absent.

For example, this paper was originally motivated by some of our previous research in comparing search strategies for healthcare improvement [2], [24]. In the context of clinical fitness landscapes, it is not reasonable to assume that all features have only main effects (corresponding to $K = 0$ in NK landscapes) as there are many known interactions between various practices and/or treatments in the real world (see [25], [26]). However, it is also not reasonable to assume that every feature interacts with at least one other feature

(corresponding to $K = 1$). Rather, we sought to explore the performance of different clinical quality improvement strategies (including randomized controlled trials and team quality improvement collaboratives) in more realistic clinical fitness landscapes where all features had main effects but varying numbers of second-order interactions were also present.

Alternatively, in some application domains one may wish to model purely epistatic landscapes in which no main effects are present. For example, in complex diseases there may be little if any association between single genes and incidence of disease [27]. Similarly, the electrical grid is explicitly ensured to be stable with respect to the loss of any one component, but interactions between two or more component outages can lead to large cascading failures [28]. For these types of applications, we and others have been seeking algorithms that are capable of detecting purely epistatic interactions (see [28]–[30]). To test these algorithms, one must be able to create benchmark landscapes where there are interaction terms but no main effects.

Classic NK landscapes cannot model landscapes between $K = 0$ and $K = 1$, nor can they model landscapes with no main effects or where the strengths of the main effects are smaller than the strengths of interaction terms [11]. In contrast, general interaction models (including NM landscapes) easily allow fine control over exactly which terms are present or absent and one can easily specify different magnitudes of coefficients for different terms. This is also possible using Walsh polynomials, although the notation is not as simple or transparent.

In the experiments shown here, we provide evidence that increasing the number m and/or maximum order M of interactions increases the ruggedness of NM landscapes with coefficients generated using (9) with $\sigma = N$, as measured by number of local peaks and the lag 1 autocorrelation of random walks through the landscapes (Fig. 4). We also assessed the difficulty of these landscapes by several different measures of search difficulty, including size of the basin of attraction of the global optimum (Fig. 8), final normalized best fitnesses found with a GA (Figs. 9 and 11), distances from the global optimum of suboptimal final best fitnesses found by a GA [Figs. 10 and 12(a)], and proportion of times a GA was able to find the global optimum [Fig. 12(b)].

B. Value of Fitness Normalization

Since the range of possible fitness values varies so much between rugged landscapes (as illustrated in Figs. 6 and 7), it is important to normalize fitnesses to a consistent range if one desires to compare fitness values on different landscapes (Fig. 11), or to assess the variability of a search strategy on landscapes with a given m and M (Fig. 9). In [2] and [24], we used logistic transforms of general parametric interaction models with unknown maxima to model search on clinical fitness landscapes with varying numbers of second-order interactions. While the logistic function successfully bounds the transformed fitnesses to the open interval $(0, 1)$, it also has the side effect of compressing high fitness values to the degree that there is very little difference between the fitnesses of the optimal peak and many suboptimal peaks. This may be a realistic assumption in health care (where there may be many possible combinations of clinical practices that yield good results), but for applications where such compression is not ideal it may be more appropriate to normalize fitnesses to

values $\in [0, 1]$ using (13), which requires knowing the global maximum, or even better to the closed interval $[0, 1]$ using (14), which also requires knowing the global minimum. *NM* landscapes enable these types of normalization, as discussed in the following sections.

C. Value of Knowing the Global Maximum

Knowing the best possible fitness offers obvious benefits, including: 1) one can terminate a search as soon as the known optimal value is found, potentially saving significant computation time; 2) one can compare methods by assessing the frequency with which the search strategies are able to find the global maximum; 3) one can tell if a stalled search has found the global optimum or is stuck on a local optimum; and 4) one can normalize fitnesses to be $\in [0, 1]$ using (13). Knowing the location of the global maximum in feature space offers obvious additional benefits [31] including: 1) one can track the evolving distances of solutions to the global optimum as the search progresses, which could potentially inform ways to improve the search process; 2) one can compare the distances (in feature space) from the best final solution to the global optimum; 3) one can assess the difficulty of the fitness landscape by assessing the correlation of fitness values encountered on a random walk with the distances to the global optimum; and 4) one can empirically explore a landscape near the global optimum in order to assess the size and shape of its basin of attraction.

For arbitrary epistatic landscapes (including *NK* landscapes, general parametric interaction models, and Walsh polynomials) finding the global maximum is NP complete. However, there are restricted subsets of these for which the global maximum is known. For example, in Walsh polynomials one can select an arbitrary point and then adjust the signs of the coefficients to force this to be the global maximum [15]. In *NM* landscapes both the location and value of the global maximum is trivially known.

D. Value of Knowing the Global Minimum

While fitnesses can be partially normalized to values $\in [0, 1]$ with (13) (as in Fig. 9), this can still be misleading, since the range of fitness values has not been properly accounted for. It is thus preferable to normalize to values in the closed interval $[0, 1]$ with (14), as in Fig. 11. For example, in Fig. 13 we illustrate how both increase in relative fitnesses over time and the variability of fitnesses on different landscapes with the same maximum order *M* are over-estimated when normalizing by (13), which only requires that the maximum possible fitness value be known, relative to when the data is normalized by (14), which requires that both the maximum and minimum possible fitness values be known.

Finding the global minimum is NP complete in *NK* landscapes and Walsh polynomials. However, in certain subsets of *NM* landscapes (e.g., Types II and III *NM* landscapes) the value and location of the global minimum is trivially known, enabling proper normalization of fitnesses.

E. Value of Arbitrary Arity of the Alphabet

Both *NK* landscapes and Walsh polynomials are defined for combinatorial problems with binary alphabets [1], [16], [17]. There are also a variety of benchmark problems with tunable

difficulty for real-valued alphabets (see [32], [33]). However, in some applications it would be desirable to have one type of model with tunable ruggedness that could be applied to binary alphabets, integer alphabets, real-valued alphabets, or heterogeneous alphabets. For example, in real clinical fitness landscapes, decision variables can have a variety of arities ranging from binary (e.g., whether or not a certain practice is performed) to real-valued (e.g., the amount or duration of application of a particular treatment) [24].

NM landscapes are well defined for alphabets of all arities (including mixed arities); changing the arity does not change the location or value of the global maximum or minimum.

F. Value of Transparency of Interactions

Various researchers are working on developing algorithms to try to detect which interactions are present in fitness landscapes and use these inferred interactions to guide the search (e.g., the linkage tree GA [34]). Being able to easily control exactly which feature interactions are present and also know the relative strengths of these interactions would facilitate the testing and validation of such approaches, as one could easily see whether the algorithm was properly estimating interaction terms.

NK landscapes offer little control over which interactions are to be included, and once generated it is nonobvious which interaction terms are present or what their coefficients values are (without significant effort [11]). Walsh polynomials present a framework where specific interaction terms can be included or excluded from the model, but the notation can be confusing and obfuscates which terms are present [see the example in (6)]. In *NM* landscapes, the interaction terms and their coefficients are obvious, since this is how interaction models are defined [see the example in (5)].

VII. Conclusion

We propose a new class of fitness landscapes with tunable degrees of epistasis, referred to as *NM* landscapes. All *NM* landscapes have a known global optimum, various subsets of *NM* landscapes have a known global minimum (thus permitting proper normalization of fitness values), the ruggedness and search difficulty of *NM* landscapes can be made to be relatively smoothly tunable, *NM* landscapes are well defined on alphabets of any arity, and which epistatic interactions are included in a particular instantiation of an *NM* landscape is easily controlled or analyzed. In summary, *NM* landscapes are a simple but powerful class of models that offer several benefits over *NK* landscapes and Walsh polynomials as benchmark models with tunable epistasis. MATLAB code for generating *NM* landscapes is available online (<http://www.cs.uvm.edu/~meppstei/NMlandscapes>).

Acknowledgments

This work was supported by the National Institutes of Health Eunice Kennedy Shriver National Institute of Child Health and Human Development Award under Grant 1R21HD068296.

References

1. Kauffman SA, Weinberger ED. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *J. Theor. Biol.* 1989; 141(2):211–245. [PubMed: 2632988]
2. Eppstein MJ, Horbar J, Buzas J, Kauffman S. Searching the clinical fitness landscape. *PLoS ONE*. 2012; 7(11) Art. ID e49901.
3. Kauffman, S. *The Origins of Order: Self Organization and Selection in Evolution*. New York, NY, USA: Oxford Univ. Press; 1993.
4. Giannoccaro, I. *Behavioral Issues in Operations Management*. London, U.K.: Springer; 2013. Complex systems methodologies for behavioural research in operations management: NK fitness landscape; p. 23–47.
5. Moraglio, A., Togelius, J. *Proc. 11th Annu. Conf. Genet. Evol. Comput.* Montreal, QC, Canada: 2009. Geometric differential evolution; p. 1705–1712.
6. Rowe W, et al. Analysis of a complete DNA—Protein affinity landscape. *J. Royal Soc. Interf.* 2010; 7(44):397–408.
7. Tomko, N., Harvey, I., Philippides, A. *SmartData*. New York, NY, USA: Springer; 2013. Unconstrain the population: The benefits of horizontal gene transfer in genetic algorithms; p. 117–127.
8. Arnaldo I, Contreras I, Millán-Ruiz D, Hidalgo JI, Krasnogor N. Matching island topologies to problem structure in parallel evolutionary algorithms. *Soft Comput.* Jul.2013 17:1209–1225.
9. Pelikan, M. *Proc. 10th Annu. Conf. Genet. Evol. Comput.* Atlanta, GA, USA: 2008. Analysis of estimation of distribution algorithms and genetic algorithms on NK landscapes; p. 1033–1040.
10. Wright AH, Thompson RK, Zhang J. The computational complexity of N-K fitness functions. *IEEE Trans. Evol. Comput.* Nov.2000 4(4):373–379.
11. Buzas J, Dinitz J. An analysis of NK landscapes: Interaction structure, statistical properties and expected number of local optima. *IEEE Trans. Evol. Comput.* Dec.2014 18(6):807–818.
12. Hordijk, W. *Recent Advances in the Theory and Application of Fitness Landscapes*. Berlin, Germany: Springer; 2014. Correlation analysis of coupled fitness landscapes; p. 369–393.
13. Liaw, RT., Ting, CK. *Proc. IEEE Symp. Found. Comput. Intell. (FOCI)*. Singapore: 2013. Effect of model complexity for estimation of distribution algorithm in NK landscapes; p. 76–83.
14. Gao Y, Culberson JC. An analysis of phase transition in NK landscapes. *J. Artif. Intell. Res.* 2002; 17(1):309–332.
15. Tanese, R. Ph.D. dissertation. *Comput. Sci. Eng., Univ. Michigan*; Ann Arbor, MI, USA: 1989. Distributed genetic algorithms for function optimization.
16. Forrest, S., Mitchell, M. *Proc. 4th Int. Conf. Genet. Algorithms*. San Mateo, CA, USA: 1991. The performance of genetic algorithms on Walsh polynomials: Some anomalous results and their explanation; p. 182–189.
17. Kallel, L., Naudts, B., Reeves, CR. *Theoretical Aspects of Evolutionary Computing*. Berlin, Germany: Springer; 2001. Properties of fitness functions and search landscapes; p. 175–206.
18. Montgomery, DC. *Design and Analysis of Experiments*. Vol. 7. New York, NY, USA: Wiley; 1984.
19. Reeves, CR. *Parallel Problem Solving From Nature PPSN VI*. Berlin, Germany: Springer; 2000. Experiments with tuneable fitness landscapes; p. 139–148.
20. Reeves, CR., Wright, CC. *Proc. 6th Int. Conf. Genet. Algorithms*. San Francisco, CA, USA: 1995. Epistasis in genetic algorithms: An experimental design perspective; p. 217–224.
21. Reeves, CR., Wright, CC. *Foundations of Genetic Algorithms*. Vol. 3. Burlington, MA, USA: Morgan Kaufmann; 1995. An experimental design perspective on genetic algorithms.
22. Weinberger E. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biol. Cybern.* 1990; 63(5):325–336.
23. Vassilev VK, Fogarty TC, Miller JF. Information characteristics and the structure of landscapes. *Evol. Comput.* 2000; 8(1):31–60. [PubMed: 10753230]
24. Manukyan N, Eppstein MJ, Horbar JD. Team learning for healthcare quality improvement. *IEEE Access*. Aug.2013 1:545–557. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6587516>. [PubMed: 25360395]

25. Johnell K, Klarin I. The relationship between number of drugs and potential drug-drug interactions in the elderly: A study of over 600 000 elderly patients from the Swedish prescribed drug register. *Drug Safety*. 2007; 30(10):911–918. [PubMed: 17867728]
26. Dechartres A, Boutron I, Trinquart L, Charles P, Ravaud P. Single-center trials show larger treatment effects than multicenter trials: Evidence from a meta-epidemiologic study. *Ann. Intern. Med*. 2011; 155(1):39–51. [PubMed: 21727292]
27. Moore JH. The ubiquitous nature of epistasis in determining susceptibility to common human diseases. *Human Hered*. 2003; 56(1–3):73–82.
28. Eppstein MJ, Hines P. A ‘random chemistry’ algorithm for identifying collections of multiple contingencies that initiate cascading failure. *IEEE Trans. Power Syst*. Aug.2012 27(3):1698–1705.
29. Eppstein, MJ., Haake, P. Proc. IEEE Symp. Comput. Intell. Bioinform. Comput. Biol. (CIBCB). Sun Valley, ID, USA: 2008. Very large scale ReliefF for genome-wide association analysis; p. 112–119.
30. Urbanowicz, RJ., Moore, JH. Proc. 12th Annu. Conf. Genet. Evol. Comput. Portland, OR, USA: 2010. The application of michigan-style learning classifiersystems to address genetic heterogeneity and epistasisin association studies; p. 195–202.
31. Jones, T., Forrest, S. Proc. ICGA. Vol. 95. Pittsburgh, PA, USA: 1995. Fitness distance correlation as a measure of problem difficulty for genetic algorithms; p. 184–192.
32. Caamaño, P., Prieto, A., Becerra, JA., Bellas, F., Duro, RJ. *Neural Information Processing. Theory and Algorithms*. Berlin, Germany: Springer; 2010. Real-valued multimodal fitness landscape characterization for evolution; p. 567–574.
33. Rönkkönen J, Li X, Kyrki V, Lampinen J. A framework for generating tunable test functions for multimodal optimization. *Soft Comput*. 2011; 15(9):1689–1706.
34. Thierens, D., Bosman, PA. Proc. 15th Annu. Conf. Genet. Evol. Comput. Conf. Amsterdam, The Netherlands: 2013. Hierarchical problem solving with the linkage tree genetic algorithm; p. 877–884.

Biographies

Narine Manukyan received the B.S. degree in applied mathematics and informatics from Yerevan State University, Yerevan, Armenia, in 2008, and the M.S. and Ph.D. degrees in computer science from University of Vermont, Burlington, VT, USA, in 2011 and 2014, respectively.

She is Assistant Professor at Champlain College and Data Scientist with Faraday, Inc., Burlington. Her research interests include artificial intelligence, evolutionary computation, complex networks, data mining, machine learning, and modeling complex systems.

Margaret J. Eppstein received the B.S. degree in zoology from Michigan State University, East Lansing, MI, USA, in 1978, and the M.S. degree in computer science and the Ph.D. degree in environmental engineering from University of Vermont, Burlington, VT, USA, in 1983 and 1997, respectively.

She is Full Professor and Chair of the Department of Computer Science, University of Vermont. She has been with the Computer Science Faculty, University of Vermont, since 1983 and was Founding Director of the Vermont Complex Systems Center from 2006 to 2010. Her research interests include complex systems analysis and modeling in a wide variety of application domains, including biological, environmental, technological, and social systems.

Jeffrey S. Buzas received the Ph.D. degree in statistics from North Carolina State University, Raleigh, NC, USA, in 1993.

He is a Professor and the Chair of the Department of Mathematics and Statistics, and Director of the Statistics Program, at the University of Vermont, Burlington, VT, USA. His research interests in evolutionary computation was the result of collaborations with computer scientists and physicians on how to model the search for better medical practice and treatments as a search on fitness landscapes.

Fig. 1.

Number of local peaks in NK landscapes with $N=10$, as a function of the number of terms in the equivalent parametric interaction model (m , bottom x -axis) for $K = \{1, 2, \dots, 9\}$ (top x -axis). The black dots show empirical results of ten random landscapes generated for each value of K ; red lines show the expected number of peaks (L) of these same landscapes computed according to the formula given in [11]. The inset shows a magnification of the $K=3$ results.

Fig. 2.

Histograms of all 1024 fitnesses in NM landscapes for $M=2$, $N=10$, and coefficients drawn from (9) with σ as indicated above each panel.

Fig. 3. Histograms of all 1024 fitnesses in binary NM landscapes for $M = 2$, $N = 10$, and coefficients drawn from a uniform distribution in the range $[0, 1]$.

Fig. 4. Number of local peaks as a function of the number of terms m (x -axis) and order of interactions M (labels near top), for (a)–(b) Type I and (c)–(d) Type II binary NM landscapes with $N=10$ and $\sigma=10$. The gray area shows one standard deviation and black lines show the means for 100 random NM landscapes. (a) and (c) show the number of local peaks, (b) and (d) show the lag 1 autocorrelation for the Types I and II NM landscapes, respectively.

Fig. 5. Number of local peaks as a function of the number of terms m (x -axis) and order of interactions M (labels near top), for Type II NM landscapes with $N=15$ and $\sigma \in \{15, 30, 100\}$. The shaded areas show one standard deviation and the lines show the means for 100 random Type II NM landscapes.

Fig. 6.

Visualization of fitnesses of all the points in representative Type I binary NM landscapes with $N=10$, $\sigma=10$ versus their distances from the global optimum in feature space for (a) $M=1$, (b) $M=2$, (c) $M=3$, (d) $M=4$, (e) $M=6$, and (f) $M=10$. In these models, all possible interactions for orders M were included.

Fig. 7.

Visualization of fitnesses of all the points in representative Type II binary NM landscapes with $N=10$, $\sigma=10$ versus their distances from the global optimum in feature space for (a) $M=1$, (b) $M=2$, (c) $M=3$, (d) $M=4$, (e) $M=6$, and (f) $M=10$. In these models, all possible interactions for orders M were included.

Fig. 8. Size of the basin of attraction for the global optimum for NK , Types I and II NM landscapes as a function of the maximum order of the interactions ($K = M - 1$).

Fig. 9.

Average of the best fitnesses in populations of 256 agents shown for 30 generations of a GA search performed on 32 random NM landscapes with $M=2$, $N=32$, and different proportions (P) of second-order interactions.

Fig. 10.

Histograms of the distances between the best solution at generation 30 and the global maximum for 32 random Type I NM binary landscapes with $M=2$ and $N=32$, and with proportions (P) of second-order interactions from top to bottom, as indicated.

(a)

(b)

Fig. 11.

Mean of the best fitnesses found by the GA over 30 generations (x -axis) for 32 random Type III NM landscapes with $M \in [1, 3, 5]$ and $N = 32$ when (a) fitnesses are not normalized, (b) fitnesses are normalized by (14).

Fig. 12.

Results of search with a GA on 32 random NM landscapes with $M = [1, 3, 5]$ and $N = 32$.

(a) Proportion of the times the search found the global maximum out of 32 runs. (b) Mean and the standard deviation of the distances between the best solution found by a GA and the global maximum, when the global maximum was not found. The dashed line indicates that at $M = 1$ all runs found the global maximum.

Fig. 13.

Mean and the standard deviation of the best fitnesses found by a GA on NM landscapes with $M=3$ and $N=32$, when fitnesses are either normalized by (13) (red dashed lines) or by (14) (black lines).