

# Neutrality in Evolutionary Algorithms...

## What do we know?

Received: date / Accepted: date

**Abstract** Over the last years, the effects of neutrality have attracted the attention of many researchers in the Evolutionary Algorithms (EAs) community. A mutation from one gene to another is considered as neutral if this modification does not affect the phenotype. This article provides a general overview on the work carried out on neutrality in EAs. Using as a framework the origin of neutrality and its study in different paradigms of EAs (e.g., Genetic Algorithms, Genetic Programming), we discuss the most significant works and findings on this topic. This work points towards open issues, which the community needs to address.

**Keywords** Neutrality · Phenotypic Mutation Rates · Problem Hardness · Genotype-Phenotype Mappings · Evolutionary Algorithms.

### 1 Introduction

Evolutionary Computation (EC) systems are inspired by the theory of natural evolution [14]. The theory argues that through the process of selection, organisms become adapted to their environments and this is the result of accumulative beneficial mutations. However, in the late 1960s, Kimura [48] put forward the theory that the majority of evolutionary changes at molecular level are the result of random fixation of selectively neutral mutations. In other words, the mutations that take place in the evolutionary process are neither advantageous nor disadvantageous to the survival of individuals. Kimura's theory, called *neutral theory of molecular evolution* or more frequently called "*neutral theory*", considers a mutation from one gene to another as neutral if this modification does not affect the phenotype.

Kimura's theory was highly criticised by the biology research community when he proposed it because the neutral theory was considered to be opposed

---

Address(es) of author(s) should be given

to the Darwinian Theory. This perception, however, was corrected by Kimura as he stated “*The theory (neutral theory) does not deny the role of natural selection in determining the course of adaptive evolution, but it assumes that only a minute fraction of DNA changes in evolution are adaptive in nature, while the great majority of phenotypically silent molecular substitutions exert no significant influence on survival and reproduction and drift randomly through the species*” [49].

The Neutral theory has inspired researchers from the EC community to incorporate neutrality in their systems in the hope that it can aid evolution.

The effects of neutrality, an area that over the last years has been studied widely in the EC community as we will see in the next paragraphs, on evolutionary search have been considered in a number of studies, the results of which, however have been contradictory. Despite of this, there are some works that we consider are relevant for the understanding of neutrality in EC. Thus, the goal of this paper is to provide a summary of these works and that this can be used as a main reference in the area of neutrality in EC.

For this purpose, we start this work by presenting the scenario that helped Kimura to shape his well-known theory of molecular evolution. We focus our attention on studies carried out using Genetic Algorithms and Genetic Programming.

This paper is organised as follows. In the next section, the origins of the neutral theory are presented. In Section 3, previous work on neutrality from a biological point of view is surveyed. Theoretical works which study neutrality are presented in Section 4. In Section 5, a review of the literature on neutrality in GAs is provided and in Section 6 previous work on neutrality in GP is surveyed. Finally, in Section 7 we discuss some open issues of neutrality in EAs and we also summarise the key points of this article.

## 2 Origins of the Neutral Theory

In [49], Kimura referred to various works that helped him shape his theory. In the following paragraphs, a brief description of these works will be presented to set a background and to illustrate how neutrality was proposed within the framework of evolution.

Kimura started by explaining the work of Lamarck [13,71]. Lamarck was perhaps the first to propose an explanation, in the field of biology, of how evolution takes place. Lamarck suggested that the use or disuse of the parts of an organism is transmitted to offspring. To explain this, he used the well-known example of giraffes (i.e., their long legs and necks were suggested to be the result of generations of stretching them to reach the leaves of tall trees). Later, Weismann [71] argued that Lamarck’s idea was flawed and to show this, he conducted an experiment where he cut off the tails of mice for twenty two successive generations. He did not find any change in the structure of the tails (i.e., shorter tails) in the final generation.

Years later, Darwin proposed his influential Theory of Natural Selection [14]. In his work, he provided arguments to explain evolution. He pointed out that better or fitter individuals will be naturally selected to survive. His findings can be summarised in his own words “*As many more individuals of each species are born than can possibly survive, and as consequently there is a frequently recurring struggle for existence, it follows that any being, if it vary in any manner profitable to itself, under the complex and sometimes varying conditions of life, will have a better chance of survival and thus be naturally selected. From the strong principle of inheritance, any selected variety will tend to propagate its new and modified form.*” [14].

Mendel [71] proposed three laws to explain the inheritance from parents to offspring through generations. These complemented Darwin’s theory. Later on, the mathematical theory of population genetics was developed by some researchers like Wright [98] studied the population consequences of Mendel’s laws. Fisher [23] also presented mathematical work synthesising Darwinian natural selection and Mendelian’s heredity laws. Fisher’s main contribution was to use stochastic methods in population genetics. More specifically, he used the notion of random fluctuation of gene frequency from generation to generation.

## 2.1 Genes and Mutations at the Molecular Level

Before molecular data became available, studies of evolution were conducted at phenotypic level and it was impossible to corroborate ideas by connecting them to changes at molecular level. It was not until the mid-1960s that molecular data became available and the first studies of amino acid sequences were carried out. This allowed Kimura to determine the evolutionary rates of nucleotide substitutions, so Kimura was able to carry out his investigations at much finer scale than previously done.

In his work, Kimura pointed out that there are two types of gene mutations: “... (1) *replacements of one nucleotide base for another, and (2) structural changes consisting of deletions and insertions of one or more nucleotides bases as well as transpositions and inversions of larger DNA segments.*” [49]. Moreover, Kimura stated that the former type of mutation is the most frequent type of mutation found at molecular level.

Analysing Haldane’s studies [37], Kimura stated that Haldane’s estimations regarding amino acids substitutions were too conservative by far and that mutations were occurring more frequently than previously thought. However, when Kimura estimated the new mutation rate, he found out that if this new mutation rate will take place in nature, then no organism would be able to survive. This was a key finding enabling Kimura to state his well-known neutral theory of molecular evolution: “... *a majority of nucleotide substitutions in the course of evolution must be the result of random fixation of selectively neutral or nearly neutral mutants rather than positive Darwinian selection, and many of the enzyme polymorphisms are selectively neutral and*

**Table 1** Brief summary on the origins of the neutral theory of molecular evolution proposed by Kimura [49].

<i>Contribution</i>	<i>Short Description or Definition</i>
Mutations at molecular level	Kimura [49] found out that Haldane’s estimations regarding amino acids substitutions were too conservative [37]. This, in consequence, opened the door to Kimura’s study in neutral mutations.
Neutral theory of molecular evolution	Kimura [49] considers a mutation from one gene to another as neutral if this modification does not affect the phenotype.
Neutral mutations	An independent work carried out by King and Jukes [50] supported the arguments first raised by Kimura. In their work, King and Jukes argued that most protein evolution is due to neutral mutations and genetic drift.

*maintained by the balance between mutational input and random extinction.*” [49]. This finding was supported by a similar work independently carried out by King and Jukes [50].

Kimura’s neutral theory has inspired many researchers to study neutrality and there is a large body of research papers in this area. In the next sections, previous work in different fields is presented to give an idea of the results regarding the presence of neutrality in evolutionary search.

## 2.2 Summary

In this section, we have presented some relevant works that helped Kimura to establish his well-known theory called “Neutral theory of molecular evolution” (more frequently called “neutral theory”). We started with a general overview of the work presented first by Lamarck [13,71], followed by the influential theory of natural selection as proposed by Darwin [14]. When Kimura proposed his theory, he was highly criticised because many researchers considered his theory to be opposed to the Darwinian Theory. This, however, was corrected by Kimura by stating that his theory should be seen as a complement to the Darwinian Theory. We also briefly summarised the work conducted by Haldane [37]. This work was a key finding that helped Kimura to realise that neutral mutations take place more frequently than previously thought. Kimura’s finding were independently corroborated by King and Jukes [50]. Table 1 shows a brief summary of some works that helped Kimura to shaping his theory.

## 3 Previous Work on Neutrality in Biology

In [97], Wilke et al. emphasised the role of mutation in selecting for flatter landscapes. In particular, the authors made an effort to show how selection favours genotypes, interconnected by mutation, whose average replication rate

is highest. As shown by the Wilke et al., this was achieved when using low mutation rates. Interestingly, the authors also were able to show that when the mutation rate was increased substantially, these genotypes occupied lower fitness peaks. However, they also noticed that these genotypes were located in flatter regions in the fitness landscape, indicating that neutrality provides robustness.

Schuster et al. [73] focused their attention on the distribution of RNA secondary structures (i.e., mapping from the RNA sequence to RNA secondary structure). This analysis of the distribution was carried out by studying the frequencies of occurrence for individual shapes and for this purpose the authors represented samples of RNA secondary structures as a tree-graph. The results showed that common shapes (these are structures that can be accessed from an arbitrary sequence by a number of mutations smaller than the chain length) are less frequently found than rare structures. Furthermore, Schuster et al. found that sequences leading to the same structure are randomly distributed.

Schuster [74] extended his previous investigation [73] suggesting that evolution cannot be understood without studying the genotype-phenotype mapping. He carried out molecular evolution experiments with RNA molecules that led him to conclude that the existence of selectively neutral phenotypes plays an important role in the success of evolution. Schuster, however, also pointed out that neutrality is not an indispensable element for evolution. Instead neutrality should be considered as a medium that, under certain circumstances, might help evolution.

Huynen pointed out that the existence of large amounts of redundancy is a key search strategy in natural evolution [41]. The author used the mapping from RNA sequence to RNA secondary structure to carry out his research. The mapping in this sequence presents extensive amounts of redundancy. Although, on average, RNA landscapes are very rugged, the existence of neutral paths allows smooth exploration. A similar point has been made in [40]. To illustrate the existence of neutral paths, Huynen et al. performed neutral walks (a brief description of how they work is presented in the following paragraph) on the network of RNA secondary structures to measure the total number of new structures encountered by neutral mutations. The authors called this the *rate of innovation* and this helped them to show that the number of new structures found by neutral mutations increased linearly over time.

In [5, 41, 69, 76] the authors used neutral random walks. The algorithm to perform neutral walks is given in [70] (see Algorithm 1 for a full description) and works as follows: (a) start with a random solution, (b) generate all its neighbours and (c) choose a neutral neighbour that results in an increase in the distance from the starting point. This process is repeated until no further distance can be increased.

Fontana and Schuster [25] pointed out that protein folding induces very complex topologies. They used a sample RNA sequence and mapped it into an RNA secondary structure. They carried out an experiment which consisted in defining a specific target shape and observing how the population evolved towards that shape. Interestingly, they found that during the process, there

---

**Algorithm 1:** Neutral Random Walk

---

```

input: landscape
 $x_0 \leftarrow$  random configuration
walk  $\leftarrow x_0$ 
 $d \leftarrow 0$ 
 $\xi \leftarrow$  neutral neighbours ( $x_0$ )
while  $\xi \neq \emptyset$  do
  randomise the order of the list  $\xi$ 
  search for a  $y \in \xi$  such that  $d(x_0, y) > d$ 
  if found  $y$  then
    append  $y$  to walk
     $\xi \leftarrow$  neutral neighbours ( $y$ )
     $d \leftarrow d(x_0, y)$ 
  else
     $\xi \leftarrow \emptyset$ 
return walk

```

---

were flat periods (clearly referring to neutrality) where no apparent adaptive progress was taking place. However, there were also sudden approaches towards the defined target. Moreover, the authors pointed out that the dominant shapes were changing during periods of no apparent progress.

In [65], Nimwegen et al. mentioned that throughout the evolutionary process, neutrality appears automatically. They focused their attention on how the population tends to move (i.e., evolve) through neutral networks (this concept will be discussed in Section 5). Nimwegen and co-workers suggested that the population does not move purely randomly through these networks. Instead, the majority of individuals tend to migrate and stay at highly connected (i.e., with a high number of neutral neighbours) parts of the network, resulting in phenotypes that are relatively robust against mutations (i.e., thanks to neutrality, the phenotype remains unchanged when mutations are affecting the genotype). It should be noticed that it is normally accepted that a solution  $s'$  is considered to be a neighbour of a solution  $s$  if  $s'$  is one Hamming distance away from  $s$ , the set of  $s$  neighbours is denoted by  $V(s)$ . Thus, a *neutral neighbour* of  $s$  is a neighbour of the same fitness.

In the same vein, Wagner [95] argued that the presence of neutrality in a system makes it more robust against mutations. Moreover, Wagner stated that neutrality should be viewed as an element that offers evolvability in the sense that it can help to discover new phenotypes. He pointed out that neutrality in itself cannot offer any benefit because, by definition, a neutral mutation at genotype level does not change the phenotypic expression. Wagner, however, stated that through evolution, neutrality provides new adaptations, so it could be of help in that it allows evolutionary search to visit areas previously unexplored.

**Table 2** Brief summary of some works developed in the area of biology.

<i>Keywords</i>	<i>Short Description or Definition</i>
RNA secondary structures	It refers to the mapping from the RNA sequence to RNA secondary structure. These structures have been widely used in various works (e.g., [73,74,40,41]). Studies varies from analysing distribution of RNA secondary structures [73], existence of selectively neutral phenotypes [74], measure of redundancy [40,41], among others .
Neutral random walks	The use of neutral random walks have been widely used [5,41,69,76]. It has been used, normally, to gain insight of the neutral landscape's features. A full description of the algorithm can be found in Algorithm 1.
Neutral networks	It is normally accepted that a solution $s'$ is considered to a neighbour of a solution $s$ if $s'$ is one Hamming distance away from $s$ . Thus a neutral neighbour of $s$ is a neighbour of the same fitness [65].

### 3.1 Summary

As the reader might be aware, there are many interesting papers in the area of biology dedicated to study and analyse the effects of neutrality. In this section, we have made an effort to summaries few works on this area, selecting those that we think have been relevant or inspired in the area of evolutionary computation. This section started describing the use of RNA secondary structures on the analysis of neutrality. Then, we briefly described some works using two of the most-well known tools used to analyse neutrality called neutral random walks and neutral networks. Table 2 presents a brief summary of some key works on neutrality in biology.

## 4 Theoretical Work on Neutrality

### 4.1 Biological Based

In [69], Reidys et al. focused their attention on neutral networks and used a mathematical model of genotype-phenotype mapping to analyse them. In contrast with other works that state that elements form a neutral network if they are one Hamming distance away from each other, Reidys and co-workers suggested that identical phenotypic structures form a neutral network if these structures exceed a certain threshold value. In their studies, the authors used the RNA secondary structure because of its high degree of redundancy (i.e., there are many more sequences than structures). The relationship between RNA sequence and secondary structure is seen as a mapping from sequence space into shape space [73].

---

## 4.2 Artificial Based

In [82], Toussaint and Igel pointed out that standard approaches to self-adaptation in evolutionary algorithms [21] are a basic and explicit example of the benefit of neutrality. In these approaches the genome is augmented with strategy parameters which typically describe the mutation distribution (e.g., the mutation rate). These are neutral parts of the genome which are co-adapted during evolution so as to induce better search distributions. The point of view developed in [82] suggests that the core aspect of neutrality is that different genomes in a neutral set provide a variety of different mutation distributions from which evolution may select in a self-adaptive way. Interestingly, theoretical work on the evolution of strategy parameters [8] can thus be re-interpreted as theoretical results on the evolution of neutral traits.

This line of thought was further formalised by Toussaint in [81]. Given a fixed genotype-phenotype mapping one can investigate the variety of mutation distributions induced by different genomes in a neutral set. In their work, the authors introduced and formalised *trivial neutrality* that is a form of neutrality and the phenotypic projections (i.e., phenotypic mutation distributions) are constant over each neutral set. Toussaint showed that trivial neutrality is a necessary and sufficient condition for compatibility with phenotypic projection of a mutation-selection GA, i.e., whether one or another representative of a neutral set is present in a population does not influence the evolution of phenotypes. Intuitively this means that, in the case of trivial neutrality, neutral traits have no effect on phenotypic evolution. In the case of non-trivial neutrality, different genotypes in a neutral set induce different phenotypic distributions, which imply a selection between equivalent genotypes similar to the selection of strategy parameters in self-adaptive EAs. Toussaint interpreted this as the underlying mechanism of the evolution of genetic representations.

In [54,55], Lehre and Haddow proposed a simple mapping called *2PDOL*. Basically the idea of this mapping consists of rewriting a given expression by expanding each element into two symbols (i.e., *2* in *2PDOL* refers to this property). Lehre and Haddow pointed out that the proposed mapping was inspired by the simplicity of the RNA secondary structure folding which has been shown to be successful in finding a specific shape target [25]. Using this mapping, the authors defined two types of neutrality: *step k-neutrality* and *remaining neutrality*. The former refers to neutrality that takes place during the mapping process whereas the latter refers to neutrality that can be seen once the mapping process has taken place. Using the *step k-neutrality*, Lehre and Haddow found that the amount of neutrality varies at the genotype level and showed, at least for their studies, how for complex phenotypes (they used Kolmogorov complexity as a measure) the amount of neutrality is low. Conversely, they found that for phenotypes with low complexity the amount of neutrality is high. Moreover, Lehre and Haddow showed how the size of the neutral networks using *2PDOL* varies from small to very large.

In [24], Fonseca and Correia developed two redundant representations using different approaches based on mathematical tools. They focused their attention

on the properties highlighted in [72] and mentioned that some of Rothlauf and Goldberg's findings disagree with the results found in [24]. That is, Rothlauf and Goldberg (see Section 5) reported that when using synonymously redundant representation, the connectivity between phenotypes is not increased. Fonseca and Correia, however, stated that this is not necessarily true. They reported that in their proposed representations the connectivity between phenotypes tends to increase with the number of redundant bits. Moreover, they argued that such connectivity is observed even with very little redundancy, so the belief that large amounts of neutrality must be present to aid evolution [76] (this work will be explained in the following section) should be analysed in more detail.

In [16], Doerr et al. analysed in detail the run-time, which was defined by the authors as the number of constructed solutions until an optimal solution has been obtained for the first time, of the (1+1) EA using the model proposed in our previous works [31–33]. Basically the idea presented in [31–33] was to allow having a neutral layer, of constant fitness, identically distributed in the whole search space. In their investigation, Doerr et al. used three problems: OneMax problem (where the idea is to maximise a bitstring) and two deceptive problems: (a) one with a single global optimum and a single local optimum and, (b) one with a single global optimum and two local optima. The authors showed that when neutrality is added in the search space, the run-time for the first two problems increases exponentially, whereas for the last problem (a deceptive problem with a single global optimum and two local optima) neutrality significantly reduces the run-time from exponential to polynomial. This shows how neutrality could be more beneficial depending on the degree of neutrality (i.e., the fitness defined in the neutral layer) and the type of problem used. Their findings, interestingly, correspond to the results firstly shown in [31–33].

In [94], Verel et al. presented a similar analogy of “connecting” two points in the search space as expressed in our previous work [31,32], as described previously. In their work, the authors presented an algorithm, called scuba search, that allows to move from one point to another in the search space when there is no gradient information. As the authors pointed out, this algorithm is expensive in terms of exploring part of the search space to get information that could guide evolution. They also proposed a more generic algorithm that tries to overcome this situation.

In the following paragraphs, we will present some works where neutrality has been studied using GAs and GP. Before doing so, we can say that researchers have added neutrality in very different ways. Generally speaking, however, we can say that neutrality is added (i.e., increasing the search space without correspondingly the solution space, assigning constant fitness to most individuals in the population, etc.) in a way that after an individual is mutated, the resulting individual remains the same at the genotype level

**Table 3** Brief summary of some theoretical works on neutrality.

<i>Keywords</i>	<i>Short summary</i>
Self-adaptation	According to [82], standard approaches to self-adaptation are an example of the benefits of neutrality. This implies that theoretical work on the evolution of strategy parameters [81] can be re-interpreted as theoretical results on the evolution of neutral traits.
Trivial neutrality	This type of neutrality is present if the phenotypic mutation projections (the phenotypic mutation distributions) are constant over each neutral set []. This means, that this type of neutrality, neutral traits have no effect on phenotypic evolution
Synonymously redundant representation	The work first presented by Rothlauf [72] (described in Section 5) was further analysed by Fonseca and Correia [24]. In their work, the authors mentioned a disagreement between their results and Rothlauf's findings. That is, Fonseca and Correia mentioned that the the degree of connectivity between phenotypes increases with the number of redundant bits.
Tunneling properties	In [31], we proposed for the first time one of the simplest form of neutrality, called 'constant neutrality' which consists of of a neutral network of constant fitness identically distributed in the whole search space. To better imagine the effects of neutrality in evolutionary search, constant neutrality was treated as a tunnel between two points of attraction. This work has been further been explored in [16,94,32].
Run-time	Neutrality has also been studying by analysing the run-time of the algorithm in the presence of various degrees of neutrality and different types of problems. The results reported in [16] perfectly agree with other results presented in [31–33].

### 4.3 Summary

In this section, we made an effort to connecting those studies of neutrality in biology (presented in Section 3) to theoretical studies developed in EC. We started this section by summarising works that tried to calculate real mutation rates in the presence of neutrality. Then, we presented works that clearly stated the benefits of neutrality in EC, in particular, in self-adaptation [82,81]. Some results that contradict other works were also covered in this section (e.g., the work developed by [24] shows some degree of contraction with the work developed in [72]). We also covered simple approaches (e.g., constant neutrality [31]) that helped to better understand how neutrality can be seen as a tunnel between two points of attraction. We finished this section by presenting how the effects of neutrality can also be studied by analysing the run-time of an algorithm, as shown in [16]. A brief summary of these works is presented in Table 3.

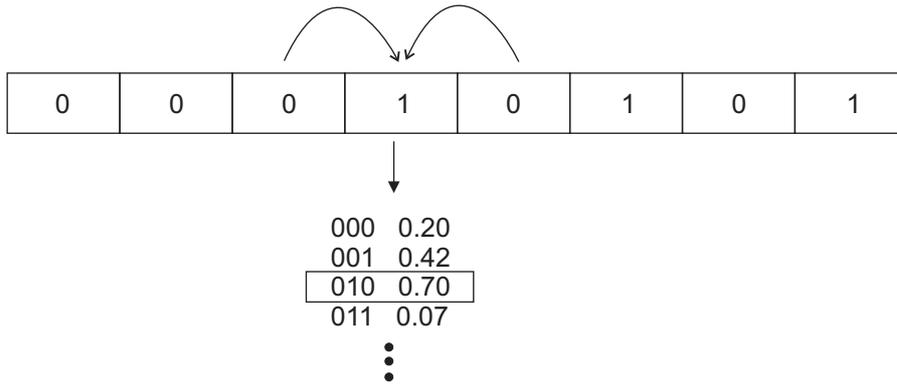
## 5 Previous Work on Neutrality in Genetic Algorithms

According to Harvey and Thompson [38], improvements of fitness can occur during evolution in GAs even in the presence of a converged population. Furthermore, they claimed: “*If a non-neutral genetic encoding (one which generates a non-neutral fitness landscape), with binary genotypes of length  $n$ , is modified by the addition of  $g$  extra redundant loci, then each phenotype will now be represented by  $2^g$  points in genotype space instead of just one. These points will form a connected neutral network. However nothing will have been gained by this exercise – we shall term this type of redundancy useless junk*” [38]. To understand this it is necessary to define a fitness landscape. This was first introduced in biology by [98]. This concept has dominated the way geneticists think about biological evolution and has been adopted within the EC community. In simple terms, a fitness landscape can be seen as a plot where each point on the horizontal axis represents all the genes in an individual corresponding to that point. The fitness of that individual is plotted as the height against the vertical axis. Thus, a fitness landscape is a representation of a search space which may contain peaks, valleys, hills and plateaus.

Also, Harvey and Thompson [38] introduced, probably for the first time the concept of *neutral networks*. The original definition of a neutral network was defined a set of points in the search space which fitness is the same. Sometimes, neutral networks are also defined as points in the search space that are connected through neutral point-mutations where the fitness is the same for all the points in such network. This concept has been regarded as a key element in neutrality as shown in [45, 83]. In [45], the authors provided a more formal definition of a neutral network. Also, in their work, in the context of Terraced NK Landscapes [63] (this will be explained later in this section), Katada and Ohkura pointed out that landscapes with a higher degree of neutrality have the larger sizes of neutral networks. They also indicated that all networks have some portals to the networks of higher fitness. According to the authors, this means that all neutral networks are, somehow, connected.

The line of thought presented by [38] was further explored years later in [6], where Barnett highly criticised what he called the “traditional GA view” regarding convergence. He stated that the lack of improvement in fitness is not due to the population being trapped in local optima. In many real-world problems it is possible (or even common) to have a large amount of neutrality, so the dynamics of evolution must be seen in terms of navigating among neutral networks that eventually will lead to higher-fit neutral networks. So, in problems with these features the problem of premature convergence would not exist.

Some useful techniques have also been proposed to analyse the neutrality present in fitness landscapes using GAs as shown by Katada et al. [46]. In their work, the authors noticed how the mutation rate plays a key element in evolutionary search and, in specific, they focused their attention on its effects when neutrality is present. More specifically, Katada and collaborators noticed how during evolution and the presence of neutrality, the population



**Fig. 1** An  $NK$  system with  $N = 8$  and  $K = 2$ . Epistatic interactions are shown for the fourth locus whose fitness contribution is 0.70. In this example, epistatic interactions are adjacent loci.

can be one of two periods: transient periods that refers when fitness tends to improve and equilibrium periods where the population exhibits a more less constant fitness behaviour. The equilibrium period last longer compared to the transient period. The authors made an effort to vary the length of these periods by varying the mutation rate during the evolutionary process. They were able to show, that this is possible, but one should pay attention to a maximum threshold for the mutation rate because a high mutation rate could harm evolution by losing fit individuals.

As we will see in the next sections, the work on neutrality developed in the area of GAs is quite vast. We have decided to categorise these works in two main parts: fitness landscapes proposed to the analysis of neutrality, followed by the addition of neutrality by the use of genotype-phenotype mappings (including some criticism by using this approach).

### 5.1 Fitness Landscapes Proposed to Analyse Neutrality

The  $NK$  landscape [47] was developed to investigate how the ruggedness (e.g., a rugged landscape has an irregular topography consisting of numerous peaks (i.e., local optima)) of a landscape changes according to the degree of epistasis. The latter takes place when the action of one gene is modified by one or more genes. Thus, the fitness of individuals depends upon the interaction of a number of their genes. An  $NK$  landscape is defined by two parameters:  $N$  which represents the number of genes that an organism has and  $K$  which represents the number of epistatic interactions between genes. Each gene makes a contribution to the total fitness of the system based on  $K + 1$  values: its own and those of the  $K$  components to which it is linked. Epistatic interactions can be either with random loci or restricted to adjacent loci. Each of the  $2^{K+1}$  possible combinations of component values is mapped to a fitness contribution drawn uniformly at random from the range  $[0, 1]$ . Figure 1 depicts an example.

The total fitness of a system is the average of the  $N$  fitness contributions. Thus, in this type of landscape it is very unlikely to find a neutral mutation.

Inspired by this type of landscape, Barnett [5] introduced the  $NKp$  landscape where the parameter  $p$  controls the degree of neutrality present in the landscape. In this landscape the values to the fitness tables are assigned to 0 with probability  $p$ . In case an entry is not set to 0, then a value drawn uniformly at random from the range  $[0, 1]$  is assigned. Thus, when  $p = 1$ , the  $NKp$  landscape is a flat landscape. Barnett performed some theoretical and empirical studies and pointed out that the addition of neutrality in this type of landscape has minimal or null effects on the ruggedness of the landscape. Moreover, Barnett claimed that the amount of neutrality decreases as the fitness increases, so the degree of neutrality present in the evolutionary search is not constant.

Designing a fitness landscape based on  $NK$  landscapes to vary the degree of neutrality was also explored in [62]. Newman and Engelhardt proposed the  $NKq$  landscapes which are very similar to the  $NKp$  landscape [5] the main difference being that the fitness contribution of the genes is in the range of  $[0, q]$  but the final fitness is normalised in the range of  $[0, 1]$ .

In [34], Geard et al. compared the  $NK$ ,  $NKp$  and  $NKq$  landscapes and found that there are several big differences in the distribution of neutral mutations when neutrality is artificially added. For instance, they reported that the  $NKq$  landscapes are qualitatively similar to the  $NK$  landscapes, while the opposite is true for the  $NKp$  landscapes where the presence of neutrality seems to make the sampling of detrimental mutations more likely.

In the context of economics, Lobo and co-workers [56] proposed a type of fitness landscape called *technological landscape*. This type of landscape, inspired by the  $NK$  landscape, allowed them to study some effects of neutrality. A technological landscape is tuned by natural number  $M$ . More specifically the fitness of a solution is rounded so it can take  $M$  different values. In their work, Lobo et al. argued that neutral networks can be seen as elements that allow individuals to improve their fitness by moving through the search space. According to the authors, this depends on the accessibility [25] of the neutral networks near the global optimum. Furthermore, the authors stated that neutrality offers robustness and innovation, so it is a desirable feature in a system. Lobo et al., mentioned that in the presence of a rugged landscape, a high degree of neutrality will help the search to find better solutions, rather than in the absence of neutrality. The opposite will occur for a smooth landscape. That is, the presence of neutrality will be detrimental in finding a global solution for this type of landscape.

In [7], Beaudoin et al. proposed another type of fitness landscape denominated  $ND$  landscape, where  $N$  is the length of the genome and  $D$  is what they called the neutral degree distribution. They divided the search space into  $D$  neutral networks by giving each of them a different fitness value (details of the algorithm can be found in [7, page 508]). The authors claimed that the other types of landscapes (i.e.,  $NKp$ ,  $NKq$  and *Technological*) used in the literature to analyse neutrality do not consider the distribution of neutrality which, is

in fact, a key feature in evolution. They used neutral degree distributions and fitness distance correlation [42] to analyse how neutrality affects evolution in this type of landscape. For this purpose, they built a *ND* landscape with deceptive features [36] and found that when the problem is difficult, the addition of neutrality makes the problem easier while the opposite happens when the problem is easy. The results reported in [7] are particularly interesting because they match perfectly our own independent findings [31–33, 66, 26].

## 5.2 Further Comments on Fitness Landscapes with Neutrality

In the last paragraphs, we have presented some fitness landscapes that haven't been proposed or used to better understand the effects of neutrality in evolutionary search. There are, however, other elements that researchers have used to analyse neutrality.

Nimwegen et al. [83] studied the dynamical features observed of a population using the royal road problem [60]. In their work, the authors studied several elements to understand the behaviour of a population during generations. In particular, they were interested in understanding metastability (e.g., period of stasis followed by brief periods of rapid change in evolutionary processes behaviour). For this purpose, Nimwegen et al. defined some elements that helped them in their studies. These include *epochal evolution* which is defined as the period of time where the system seems stabilise (clearly indicating the presence of neutrality) on some feature distribution followed by a brief burst of change, *speed of innovation* that refers how often the population moves towards better spaces (fitter areas), and *fitness fluctuation amplitude* that measures the “jumps” performed by a population through generations.

Smith et al. [78] also proposed the use of other elements to understand the effects of neutrality. They were inspired by the fact that most researchers use a single global metric to understand the properties of a fitness landscape. The authors stated that even when these global measures can be of use, it will be more beneficial to have measures that focus on specific areas of the search space. Thus, they proposed the use of fitness *evolvability portraits* (this is calculated by averaging evolvability over a population of equal fitness) that can be used to compare both the ruggedness and neutrality in a set of tunable rugged and tunable neutral landscapes. Smith et al. showed how the method used in their work is beneficial in detailing features of the search space. Moreover, as they stated in their work, the method used by them should be seen as a complement to other methods, rather than considering it as the only method that can be used to “visualise” properties of the landscape (e.g., ruggedness).

Other measures have been proposed by Vassilev et al. [92]. The authors made a similar argument arose by Nimwegen et al. [83], the analysis of the fitness landscape should be considered as an ensemble of various objects, which are characterised by several elements such as size, form and distribution. In other words, they also considered the study of the landscape by decomposing it. To do so, the authors proposed three new information analysis of fitness

landscapes: *information content*, *partial information content* and *information stability*. As explained by the authors, the first two are based on measures of the amount of information contained in the ensemble of objects, where each of them consists of a point in the landscape and its nearest neighbours. More specifically, information content can be seen as the amount of information needed to reconstruct the system exactly. This measure is an estimate of the variety of shapes in the ensemble. So, this can be used to measure the ruggedness of the landscape. Partial information content relates to the modality encountered on the landscape path. Finally, information stability refers as the highest possible difference in the fitness values of two neighbouring points.

Other element to study neutrality is the Nei's standard genetic distance as shown by Katada and Ohkura [44]. In their work, the authors highlighted some benefits of the Nei's distance when the the mutation rate per locus is small. These include that the distance increases approx. linearly over generations in the presence of neutrality, the distance increases when neutrality increases, the distance decreases with the increase of ruggedness in landscape with neutrality. To test their idea, Katada and Ohkura used *NK* landscapes (as described in the previous paragraphs) and found out that the Nei's distance was able to measure the amount of neutrality present in the fitness landscape. They further confirmed their results by using robotics simulations and compared their analysis with extensive experimental analysis. Details of how to calculate the Nei's distance can be found in [44].

In Section 6, we further present other works that propose the use of different methods to analyse the degree of neutrality in a landscape in the context of Genetic Programming [52,67].

As we have seen, researchers have proposed and used different types of fitness landscapes for the study of neutrality. There are other elements, however, that researchers have also used, as we will see in the following section. In particular by the use of genotype-phenotype mappings.

### 5.3 Genotype-Phenotype Mappings - Definitions and Properties

In [79], Smith et al. used a complex genotype-phenotype mapping in the context of a visual shape recognition task to control the navigation of a robot and studied the population's behaviour during neutral phases (i.e., periods where fitness remains constant). They focused their attention on the evolvability of the population defined as the ability of individuals both to produce fitter individuals and to not produce less fit individuals. Empirically, Smith et al. concluded that during neutral phases, the population is not doing anything useful. The authors did not deny that neutrality could provide a buffer to move the population towards better places but, at the same time, they argued that the presence of neutrality alone does not allow the evolving of the population faster than in its absence. So, they concluded that the presence of neutrality is not beneficial for evolutionary search. This line of thought was further investigated in [80] where Smith et al. focused their attention on the population's

dynamics to point out that when the fitness does not change, this is due to the presence of neutrality rather than to the population getting stuck in local optima.

In [76], Shipman argued that neutrality is beneficial if neutral networks are spread over the search space and if there is a high degree of connectivity between them. Shipman mentioned that these neutral networks are sets of genotypes at the same level of fitness. Furthermore, he stated that these networks have a high degree of connectivity if they are connected by single point mutations. To study the effects of neutrality, he used a hybrid approach of neural networks and GAs to conduct his research. Shipman tested his approach on a robotic task where the goal was to control a robot that navigates through a corridor. He found that the presence of a high degree of neutrality helped to discover many more phenotypes and eventually to sample higher fitness. That is, the presence of neutrality has a beneficial impact in evolutionary search because it helps to discover fitter solutions. However, Shipman also concluded that he could not see any benefit in the number of generations required to find optimal solutions (i.e., the presence of neutrality does not help to find solutions faster).

In [75,77], Schackleton et al. illustrated that neutrality can be artificially added to the evolutionary search with the use of genotype-phenotype mappings. To illustrate this, the authors proposed five different types of mappings. Let us briefly describe them. The first mapping, called *static random mapping*, consisted in defining a genotype of length 30 which is mapped to a phenotype of 16 bits. The mapping used was randomly initialised and remained static afterwards. The second mapping, called *trivial voting mapping*, consisted in taking 3 bits at genotype level to represent one bit a phenotypic level. The bit is set to 1 if the majority of the 3 bits voted in favour, 0 otherwise. The third mapping, called *standard voting mapping*, is a variation of the previous mapping. The main difference is that the set of bits at genotype level can overlap, so when a single point mutation takes place multiple phenotypic bits could simultaneously change. In the fourth mapping, called *cellular automaton mapping*, each of the phenotypic bits was associated with a truth table. Three adjacent bits were used as inputs in the truth table and the corresponding output determined the new state (i.e., phenotypic bit). Finally, the fifth mapping, called *random Boolean network* (RBN) is a variation of the previous mapping. The main difference is that the 3 bits can be at any positions, so it is necessary to encode those positions at the genotype level. The authors noted that the amount of redundancy plays a key role in evolution. Moreover, they observed that some mappings (i.e., standard voting, cellular automaton and RBN) were more beneficial than others (i.e., trivial voting).

In [19,20], Ebner et al. extended the previous investigation of [75,77]. For this purpose, the authors analysed the effects of the RBN mapping and the cellular automaton (both described previously) in the context of what they called *phenotype-species mapping*. In their work, Ebner et al. emphasised some benefits gained in the presence of neutral mappings. As detailed in their work, the authors explained how neutral networks (defined and discussed in Sec-

tion 5) induced by mappings sustain high mutation rates. Moreover, neutral networks help the population to spread “randomly” throughout the search space which could represent an advantage if the environment changes. Also, the authors stated that one of the main benefits of a redundant mapping is that this type of mapping offer diversity. Thus allowing the population to get stuck in local optima.

In [96], Weicker and Weicker focused their attention on the consequences of redundancy. They stated that there are four areas where redundancy can be found: coding based, representation based, conceptual redundancy and technical redundancy. Coding based redundancy takes place when the size of the search space does not match the size of the genotype space. Representation based redundancy is caused by structural reasons (i.e., it is caused by either the problem considered or the optimisation technique used). Conceptual redundancy takes place when redundancy is added using gene interactions (i.e., [77,75]). Technical redundancy takes place when, for instance, a decoder is used to improve or repair a given solution. Weicker and Weicker focused their studies on the latter two cases using a decoder method and a diploid encoding on a binary representation. The former method works by defining “1” as an element for inclusion and “0” as an element to be omitted. In the diploid encoding, each individual is formed by two complete candidate solutions and an extra bit which defines the active solution. The diploid encoding was inspired by the work reported in [15] where the authors proposed the *structured Genetic Algorithm* which is a special case of diploid encoding, where an extra bit switches between two complete candidate solutions. As can be seen, Weicker and Weicker used very different ways of adding neutrality and analysed a particular point: how the presence of redundancy converts local optima into plateau points. They reported that, in both cases, this happens and that a decoder method can find good solutions quicker than a diploid encoding. The authors, however, were unable to explain why this happened.

In [9,10], Chow proposed the use of a population of individuals that are composed by two chromosomes: a *data chromosome* and a *mapping chromosome*. The former chromosome is one that contains the genotypic expression (i.e., 0s and 1s for binary strings) whereas the latter chromosome stores bit locations as integers and these determine the position of the values of the data chromosome. Both chromosomes undergo separate genetic operations. For the data chromosome the traditional genetic operators (i.e., crossover and mutation) are used, while for the mapping chromosome, the permutation operator is used to alter the bit ordering. Neutrality is present in this type of mapping because when the permutation operator takes place, there is no guarantee that all the bits in the data chromosome participate in the genotype-phenotype mapping. That is, there is a possibility that more than one integer is repeated in the mapping chromosome. Using a hybrid algorithm (i.e., GA receiving feedback from a local hill climbing), Chow tested his approach on trap functions [35,36,43] and reported that this type of neutrality had a beneficial impact in evolutionary search.

In [72], Rothlauf and Goldberg stated that redundancy is a common element found in any EC system and emphasised that the effects of redundancy in evolutionary search depend basically on the nature of redundancy. They identified some properties of redundant representations: (a) a redundant representation is *uniform* if all phenotypes can be obtained by the same number of genotypes, (b) a redundant representation is *synonymously* redundant (an example of this type of redundancy is the trivial voting mapping proposed in [75,77]) if the genotypes that map to the same phenotype are part of a neutral network (i.e., they are close to each other), (c) a redundant representation presents high *locality* if neighbouring genotypes map to neighbouring phenotypes and finally, (d) a redundant representation presents high *connectivity* if the number of phenotypes which are accessible from a phenotype by one bit-flip mutation is high. They also mentioned that non-synonymously redundant representations (examples of this type of redundancy are the cellular automaton and the RBN described in [75,77] and criticised in [51], see below) are those where two genotypes representing the same phenotype are very different from each other. Thus, in synonymously redundant representations, genetic operators work well and the search is smoother than in the non-synonymously redundant representations where the search operators show a poor performance. So, in this type of redundancy the search behaves like random search.

### 5.3.1 Criticism to Some Genotype-Phenotype Mappings

In [51], Knowles and Watson criticised the usefulness of neutrality when added via a mapping function. In particular, they focused their attention on the RBN mapping proposed and studied in [19,75,77] and measured its performance using the rate of fitness increase. Knowles and Watson used GAs and Hill-Climbing on three different problems free of neutrality to compare the performance obtained when RBN was and was not used. They showed that the performance of the search algorithms used in their experiments was better in the absence of neutrality. Moreover, they suggested that the RBN mapping leads to a random exploration in the search space, so it is difficult to imagine how evolutionary search can gain anything from using this type of mapping.

## 5.4 Summary

In this section we have presented some works using genetic algorithms to study and analyse the effects of neutrality in evolutionary search. We started by reviewing works that defined key concepts in neutrality (e.g., neutral networks [38]). Then, we focused our attention on some fitness landscapes that have been proposed and used to study neutrality (mostly all of them inspired by the *NK* landscape proposed in [47]). We followed our summaries on some “tools” that researchers have been using in the study of neutrality. Table 4 shows a brief summary of some works that have used GAs for the study and analysis of the effects of neutrality.

**Table 4** Brief summary of some works using on neutrality using Genetic Algorithms.

<i>Keywords</i>	<i>Short summary</i>
Neutral networks	It is defined as a set of points in the search space which fitness is the same [38]. Sometimes neutral networks are also defined as points in the search space that are connected through neutral point-mutations where the fitness is the same for all the points in such network.
Fitness Landscapes	As acknowledge by several researchers, the effects of neutrality are hard to understand for different reason, such as the use of different operators, representations and more, as pointed out in [31]. Thus, researchers have used different fitness landscapes such as: <i>NK</i> landscapes [47] (Figure 1 depicts an example), <i>KKp</i> [5], <i>NKq</i> [62], <i>technological</i> [56], and <i>ND</i> [7] landscapes. All these type of landscapes have in common that is relatively easy to control the amount of some properties, such as neutrality, present in the landscape. Intuitively this means that, these type of landscape offer a control environment to studying some effects.
“Tools” to study neutrality	Genetic Algorithms have widely been used to study the effects of neutrality because of their “simplicity”. Researchers have proposed and studied different aspects of neutrality by using different “tools” or “metrics”, such as epochal evolution, speed of innovation, fitness fluctuations [83], information content, partial information content and information stability [92], Nei’s standard genetic distance [44], phenotypic mutation rates [66].
Genotype-Phenotype mappings	Several methods have also been proposed to add neutrality in the evolutionary search. In particular by the use of genotype-phenotype mappings. Works include the use of constant neutrality [31], static random mapping, trivial voting mapping, standard voting mapping, cellular automaton mapping and Boolean network [75,77], bit-wise neutrality [66].

## 6 Previous Work on Neutrality in Genetic Programming

In [2], Banzhaf mentioned that constrained optimisation problems (i.e., problems where a potential solution is judged by its fitness and that also must obey certain restrictions) can be handled using a genotype-phenotype mapping. To do so, he proposed a mapping called *Binary Genetic Programming*. He argued that this mapping allows one to use unrestricted search operators in the genotype space (i.e., search space) while at the same time the feasibility of solutions in the phenotype space (i.e., solution space) is guaranteed. The latter is accomplished thanks to a correction step that takes place in case a solution is not feasible. Briefly, Banzhaf’s mapping was based on the use of a transition table that was composed of 32 codes (i.e., 5-bit binary coding) to each of which a corresponding symbol was assigned (i.e., a function or a terminal). So, a bitstring of length  $n$  (i.e., where  $n$  is a multiple of 5) can be

translated into an expression of  $n/5$  functions and terminals. Given that there were fewer functions and terminals than different combinations of codings, this introduces neutrality. Banzhaf showed that this type of mapping was useful in a constraint optimisation problem. However, his conclusions are the result of several steps which were not analysed in detail.

In [18], Ebner suggested that part of the success of GP in finding solutions is due to fact of neutrality induced by “junk” code or introns (which are defined as “parts of the genotype that emerge as a result of the evolution of individuals with a variable length representation and have no influence on the survival of the individual” [4]). In his work, Ebner mentioned that finding a specific individual in the GP search space is almost impossible, but finding a specific results it is not. To explain this, he mentioned that the presence of introns helps to have multiple individuals leading to the same result. Thus, the presence of introns in GP could be beneficial.

In [93], Vasseliv and Miller explored the effects of adding artificial neutrality to Miller’s approach, called *Cartesian Genetic Programming* (CGP) [57], and studied its effects using an evolvable hardware problem (i.e., three-bit multiplier). CGP uses a genotype-phenotype mapping (an integer string coding) that allows the presence of inactive code (i.e., this is, according to the authors, how neutrality is artificially added). This is a representation of programs in which functions are joined by connections. The authors claimed that neutrality helps avoid getting stuck in local optima. Moreover, Vassilev and Miller performed neutral random walks [70] to show that the amount of neutrality decreases when the search is close to the global optimum.

This line of research was further extended in [99], where Yu and Miller argued that in the traditional GP representation there are two types of *implicit* neutrality (i.e., neutrality that is already present in the evolutionary search): *functional redundancy* and *introns*. The former refers to the fact that multiple genotypes can represent the same phenotype whereas the latter refers to the fact that there are parts of a program that are semantically redundant. This means that if a genetic transformation removes that redundancy from the genotype, the program’s behaviour will remain the same (i.e., suppose that we are in the presence of the genotype  $G_1 = (NOT(NOT(AND(OR X_1 X_2)X_1)))$  by removing  $(NOT, NOT)$ , the resulting new genotype will be  $G'_1 = (AND(OR X_1 X_2)X_1)$  so it is clear that  $G_1$  and  $G'_1$  will compute the same result). The authors argued that this type of neutrality (i.e., implicit neutrality) is difficult to identify and control during evolution, so they used CGP to add *explicit* neutrality (where the authors were referring to inactive code). They tested their approach on the even-3-parity problem. CGP allowed the authors to use Hamming distance to measure the amount of neutrality present in the evolutionary search. They found that the more neutrality is present during evolution, the higher the percentage of success in finding the global optimum irrespective of the mutation rate used. Thus, they concluded, neutrality is fundamental to improve evolvability. the phenotype level (e.g., a node that is multiplied by zero). They carried out extensive empirical experiments that helped them to conclude that the best

performance (using an evolutionary strategy on the even-3-parity and 2-bit multiplier) was achieved with high degrees of neutrality. However, they were unable to give a clear explanation for what happened as they mentioned “*Detailed further study is required to ascertain how the redundancy is utilized and interacts during evolution to assess problem solving.*” [58].

Yu and Miller [100] developed a mathematical framework to study the implicit neutrality found in the OneMax problem when using a simple GA. They focused their attention on the amount of neutrality in each of the fitness classes (in OneMax there are  $\ell+1$  fitness classes for a chromosome of length  $\ell$ ). The number of chromosomes with a given fitness in the search space is given by a binomial coefficient. So, the largest number of chromosomes corresponds to the fitness  $\ell/2$  (assuming, for simplicity, that  $\ell$  is even). The situation changed dramatically when they used CGP [57]. As the authors pointed out, the amount of neutrality is highly dependent on the representation. Based on their empirical findings, Yu and Miller found that neutrality had a beneficial impact on this unimodal landscape and claimed that it provides a buffer to absorb destructive mutations.

### 6.1 Work Inspired by Genotype-Phenotype Mappings

In [17], Downing used Binary Decision Diagrams (BDDs), which are a special case of Binary Decision Trees (BDT) (where each non-terminal element is associated with a Boolean variable which has *if* and *else* children). BDDs are presented as directed and acyclic graphs. This allowed Downing to define four different neutral mutations and one adaptive mutation. When a neutral mutation takes place, the fitness of the new individual is copied from its parent, so there is no need to calculate the fitness value for the generated child (as pointed out by Downing, this is an advantage over Yu and Miller’s approach [99] where each time a neutral mutation takes place, the fitness of a new individual must be calculated). Given the nature of BDD, Downing tested his proposed approach on Boolean problems (i.e., even- $n$ -parity Boolean functions, where  $7 \leq n \leq 17$ ) obtaining excellent performance (i.e., 100% success rate for all the even- $n$ -parity functions). As stated by Downing, one should be careful to interpret the results bearing in mind that not only is neutrality present but also the representation allows modularisation, so the results are a mixture of several ingredients.

### 6.2 Criticism to Explicit Neutrality in GP

In [12], Collins claimed that the conclusions reported in [101], i.e., that neutrality is beneficial, are flawed. Collins started his analysis by highlighting that the use of a Boolean parity problem is a strange choice given that the problem in itself is neutral. In particular, Yu and Miller used two different function sets formed by either *EQ* and *XOR* or simply *EQ*, so this problem has a

needle-in-haystack property: there are only two possible fitness values. So, the effects of neutrality are harder to analyse using this type of problem. Moreover, Collins focused his attention on the results found for the even-12-parity Boolean problem and pointed out that the CGP representation used in [101] favours shorter sequences than those yielding solutions for this Boolean problem. Although [101] reported good results (i.e., 55% of success in finding the solutions), Collins proved that random search has even better performance. Collins used various sampling methods to corroborate his findings and concluded that the effects of neutrality are more difficult to comprehend than previously thought.

### 6.3 Analysis of Landscapes

In [3], Banzhaf and Leier exhaustively analysed the search space of a Boolean function problem using only the *AND* function to illustrate how there are many more common phenotypes than uncommon phenotypes (i.e., phenotypes that represent the global solution) in the search space and the latter can be made accessible by the presence of neutrality. According to Banzhaf and Leier, neutral networks can be of use only if they are highly intertwined. This property allows the search to move quickly from one neutral network to another and, eventually, to sample the global optimum.

Recently, Vanneschi and co-workers [88] defined and used several measures to analyse the neutrality present in some Boolean parity fitness landscapes. Let us briefly describe them. Firstly, they defined the *average neutrality ratio*  $\bar{r}$  of neutral network as the mean of the neutrality ratios of the individuals contained in a network. Secondly, the authors defined the *average  $\Delta$ -fitness of the neutral network* which is the average fitness obtained after applying a mutation to an individual contained in the neutral network. Thirdly, they defined the *Non Improvable (NI) Solutions ratio* ( $r_{ni}$ ) which is the number of solutions that are generated by mutation operators and that are not fitter than the original solution. Fourthly, Vanneschi et al. defined the *Non Worsenable (NW) Solutions ratio* ( $r_{nw}$ ) which is the number of solutions that are generated by mutation operators and that are not worse than the original solution.

In their initial analysis, the authors focused their attention on small landscapes (i.e., the even-2-parity function and set a maximum depth of 3), and used two different function sets,  $F_1 = \{XOR\}$  and  $F_2 = \{NAND, NOT\}$ . Finally, they defined three mutations which are a variation of the inflate and deflate mutation operators originally defined in [84] to study the effects of neutrality on evolutionary search. Vanneschi et al. pointed out that the effects of neutrality strongly depend on a landscape's features. To prove this, they used the proposed measures to show that in certain landscapes such as the one induced by the function set  $\{NAND\}$  it is very unlikely that neutrality will help to improve fitness, where the opposite is true in the landscapes defined by  $\{XOR, NOT\}$ .

In [86], Vanneschi summarised some “tools” that people have used to study different properties of the fitness landscape, including neutrality. These include measuring neutrality in Cartesian gp using the Hamming distance, a pioneer study performed by Yu and Miller [99] (as described previously in Section 6), fitness distance correlation used in GP (tree-like structures) [84]. Other “tools” mentioned by Vanneschi include the negative slope coefficient [87] that can be calculated even without knowing the optimas genotypes and it does not (explicitly) use any distance. In [87], the authors reported good results on some GP benchmark problems. Successively in [85] these results have been extended to some real-like applications. Then in [68] the authors gave a formal model of the fitness proportional negative slope coefficient and a more rigorous justification of it. Finally, in [91], the authors pointed out the limitations of this approach.

## 6.4 Summary

In this section, we covered several works in the area of GP that have made an effort to better understand the effects of neutrality in evolutionary search. Contrary to the studies carried out using GAs where there is a substantial amount of work (as summarised in the previous section), there are fewer works using GP. Probably one reason is the difficulty that the tree-like structure representation used in typical GP presents. In section, we started reviewing works that have used some form of genotype-phenotype mapping. We then presented some criticism to this kind of approaches (i.e., genotype-phenotype mapping). We finalised this section by summarising some measures that researchers have proposed in the study of neutrality in GP. Table 5 presents a brief summary of some works done on the study of neutrality in GP.

## 7 Open Issues and Conclusions

There are many techniques available in the specialised literature that allow computer systems to learn (i.e., decision trees, neural networks, Bayesian learning, evolutionary algorithms and more [61]). This work is primarily based on the use of *Evolutionary Computation* (EC) [1, 22] systems (also known as *evolutionary algorithms*). Genetic Algorithms (GAs) [39] and Genetic Programming (GP) [52] are the methods widely analysed and summarised in this work.

The Genetic Algorithm is a widely used form of the evolutionary algorithm. Originally, GAs were conceived by Holland as a means of studying adaptive behaviour. They, however, are largely used as function optimisation techniques. GP has perhaps the richest representation among all the paradigms in EC systems. GP can be used to evolve computer programs to solve problems automatically without having to tell them explicitly how [52, 53, 67].

Despite the proved effectiveness of EC systems, there are limitations in such systems and researchers have been interested in making them more powerful by

**Table 5** Brief summary of some works using on neutrality using Genetic Programming.

<i>Keywords</i>	<i>Short summary</i>
Genotype-Phenotype mappings	Just as in GAs, the use of genotype-phenotype mappings has also been used in GP. Different approaches have been used. Some studies include the use of Binary GP [2], Cartesian GP [57], Constant neutrality [27], Uniform GP [29].
“Tools” to study neutrality	In GP, the study of neutrality, and in fact, of any effect seen in GP evolution is hard to measure given the typical representation used in GP (i.e., tree-like structure). There are, however, some “tools” that have been defined to study some neutrality effects. These include the use of the average neutrality ratio, average fitness of the neutral network, non improvable solutions ratio, and ‘worsenable’ solutions [88]. Another tool that has been used to analyse the hardness of a problem with and without the presence of neutrality (artificially induced) is the fitness distance correlation (fdc) [42]. The use of fdc in GP has been reported in [11,89,84]. For fixed-length representations in GP (e.g., Cartesian GP), the Hamming distance has also been used [99] in the study of neutrality.

using different elements. One of these elements is *neutrality* (*the neutral theory of molecular evolution* [48,49]) which the EC community has incorporated in their systems in the hope that it can aid evolution. Briefly, neutrality considers a mutation from one gene to another as neutral if this modification does not affect the fitness of an individual. A more detailed description of neutrality is presented in Section 2.

Neutrality has attracted the interest of researchers in the EC community. However, the results reported in relation to the benefits of neutrality in evolutionary search are contradictory, as we have mentioned throughout this work. For instance, in *“Finding Needles in Haystacks is not Hard with Neutrality”* [101], Yu and Miller performed runs using the well-known Cartesian GP (CGP) representation [57,59] and also used the even- $n$ -parity Boolean functions with different degrees of difficulty ( $n = \{5, 8, 10, 12\}$ ). They compared performance when neutrality was present and in its absence and reported that the performance of their system was better when neutrality was present.

A few years later, Collins claimed the opposite and presented the paper entitled *“Finding Needles in Haystacks is Harder with Neutrality”* [12]. He further explored the idea presented by Yu and Miller and explained that the choice of this type of problems is unusual and in fact not suitable for analysing neutrality using CGP. This is because both the landscape and the form of the representation used have a high degree of neutrality and these make the drawing of general conclusions on the effects of neutrality difficult.

These works (both nominated as best papers in their conference tracks!) are just two examples of many publications available in the specialised literature which show controversial results on neutrality.

It is clear that a large amount of work has been done towards understanding the effects of neutrality in evolutionary search. However, there are some areas that in our view need to be addressed. In the following paragraphs, some potential areas to be explored are discussed.

### 7.1 Form of Neutrality

As shown in Section 5, there is a large number of works on the use of different genotype-phenotype mappings in GAs to add neutrality in evolutionary search. Possibly, the main reason for doing this, it is the relative “simplicity” of GAs. This has allowed researchers to perform a more detail analysis of the effects of neutrality. In fact, many works have defined and used several tools to performing this exhaustive analysis as mentioned in Sections 5 and 6, and briefly summarised in Tables 4 and 5.

Genotype-phenotype mappings have also been used in GP to add neutrality into the search space, as described in Section 6. Of course, neutrality’s effects are harder to analyse using this kind of representation (tree-like structures). However, it is possible to analyse it in detail under some circumstances that allow to have more less control of the system. So, it would be interested exploring new forms of encoding functions (in bit-wise neutrality), or using bitwise-type (introduced and explored in Section 5) of neutrality in GP. Likewise, it would be interesting using something like degree neutrality (introduced and explained using GP in Section 6) but in GAs (e.g., using variable length strings to encode fixed length ones).

### 7.2 Mathematical Frameworks

Often, it is common to see that researchers report experimental results without a support of statistical information nor mathematical frameworks. In our view, this is one of the key elements missing in the understanding on the effects of neutrality. Put it in other way, the lack of mathematical frameworks developed within the frame of understanding the effects of neutrality in evolutionary search has had a big impact on the confusion of it.

It will be highly desired if research can be further explored into theoretical work to corroborate empirical evidence. For instance, we need to do theory of *fitness distance correlation* and phenotypic mutation rates for GP. We need to understand phenotypic crossover rates both for GP and GAs.

### 7.3 Prediction on the Effect of Neutrality

Throughout the paper, we have mentioned different benchmark problems that have been used in the study and analysis of neutrality in evolutionary search. There are some cases where the effects, whether beneficial or detrimental, of neutrality are clearer. Intuitively this means that the fitness landscape and

the type of neutrality play an important role in determining the effects of neutrality in the problem.

That is why, we believe that one needs to find good ways of predicting [30, 28, 29, 42, 90] when the addition of neutrality (and what form of neutrality) can be beneficial in practical situations, e.g., when facing an unknown problem.

#### 7.4 Final Comments

We have presented a survey of neutrality. To do so, we have summarised works where neutrality has been analysed, starting from its origins in biology and finishing with the most recent studies including the use of Genetic Algorithms and Genetic Programming, where both empirical experiments and mathematical frameworks have been used in an attempt to explain the effects of neutrality.

Throughout this work, we have highlighted the fact that there are no conclusive results on neutrality in EAs. The confusion regarding neutrality is due to several reasons. These include the following:

- many studies have based their conclusions on performance statistics (i.e., on whether or not a system with neutrality could solve a particular problem faster than a system without neutrality) rather than a more in-depth analysis (i.e., measure of hardness, population dynamics, etc.),
- studies often consider problems, representations and search algorithms that are relatively complex and, so, results represent the compositions of multiple effects (e.g., bloat or spurious attractors in GP),
- there is not a single definition of neutrality, and different studies have added neutrality to problems in radically different ways,
- very often studies focused their attention on particular properties of neutrality without properly defining them,
- the features of a problems landscape change when neutrality is artificially added, but rarely has an effort been made to understand in exactly what ways.

#### Acknowledgments

We would like to thank the editor, associate editor and reviewers for their fair and useful comments and ideas. The paper has been considerably strengthened thanks to their feedback.

This research is based upon works supported by Science Foundation Ireland under Grant No. 08/IN.1/I1868.

#### References

1. Thomas Bäck, David B. Fogel, and Zbigniew Michalewicz, editors. *Evolutionary Computation 1: Basic Algorithms and Operators*. IOP Publishing Ltd., Bristol, UK, 1999.

2. W. Banzhaf. Genotype-Phenotype-Mapping and Neutral Variation: A Case Study in Genetic Programming. In Yuval Davidor, Hans-Paul Schwefel, and Reinhard Manner, editors, *PPSN III: Proceedings of the 3rd International Conference on Parallel Problem Solving from Nature*, pages 322–332, Jerusalem, Israel, 1994. Springer-Verlag.
3. W. Banzhaf and A. Leier. Evolution on Neutral Networks in Genetic Programming. In T. Yu, R. Riolo, and B. Worzel, editors, *Genetic Programming - Theory and Applications III*. Kluwer Academic, 207–221, 2006.
4. Wolfgang Banzhaf, Frank D. Francone, Robert E. Keller, and Peter Nordin. *Genetic programming: an introduction: on the automatic evolution of computer programs and its applications*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1998.
5. Lionel Barnett. Ruggedness and Neutrality – the NKp Family of Fitness Landscapes. In Christoph Adami, Richard K. Belew, Hiroaki Kitano, and Charles Taylor, editors, *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*, pages 18–27, Cambridge, MA, 1998. MIT Press.
6. Lionel Barnett. Netcrawling-Optimal Evolutionary Search with Neutral Networks. In *Proceedings of the 2001 Congress on Evolutionary Computation*, pages 30–37. IEEE Press, 2001.
7. William Beaudoin, Sebasiten Verel, Philippe Collard, and Cathy Escazut. Deceptiveness and Neutrality The ND Family of Fitness Landscapes. In Maarten Keijzer, Mike Cattolico, Dirk Arnold, Vladan Babovic, Christian Blum, Peter Bosman, Martin V. Butz, Carlos Coello Coello, Dipankar Dasgupta, Sevan G. Ficici, James Foster, Arturo Hernandez-Aguirre, Greg Hornby, Hod Lipson, Phil McMinn, Jason Moore, Guenther Raidl, Franz Rothlauf, Conor Ryan, and Dirk Thierens, editors, *GECCO 2006: Proceedings of the 2006 Conference on Genetic and Evolutionary Computation*, volume 1, pages 507–514, Seattle, WA, USA, 8-12 July 2006. ACM Press.
8. H. Beyer. *The Theory of Evolution Strategies*. Springer Verlag, 2001.
9. R. Chow. Effects of Phenotypic Feedback and the Coupling of Genotypic and Phenotypic Spaces in Genetic Searches. In *Proceedings of the 2004 IEEE Congress on Evolutionary Computation (CEC-2004)*, volume 1, pages 242–249, Portland, Oregon, USA, 2004. IEEE.
10. R. Chow. Evolving Genotype to Phenotype Mappings with a Multiple-Chromosome Genetic Algorithm. In Kalyanmoy Deb, Riccardo Poli, Wolfgang Banzhaf, Hans-Georg Beyer, Edmund K. Burke, Paul J. Darwen, Dipankar Dasgupta, Dario Floreano, James A. Foster, Mark Harman, Owen Holland, Pier Luca Lanzi, Lee Spector, Andrea Tettamanzi, Dirk Thierens, and Andrew M. Tyrrell, editors, *GECCO 2004: Proceedings of the 2004 Conference on Genetic and Evolutionary Computation*, volume 1 of *Lecture Notes in Computer Science*, pages 1006–1017, Seattle, WA, USA, 26-30 June 2004. Springer.
11. Manuel Clergue, Philippe Collard, Marco Tomassini, and Leonardo Vanneschi. Fitness Distance Correlation and Problem Difficulty for Genetic Programming. In William B. Langdon, Erick Cantú-Paz, Keith E. Mathias, Rajkumar Roy, David Davis, Riccardo Poli, Karthik Balakrishnan, Vasant Honavar, Günter Rudolph, Joachim Wegener, Larry Bull, Mitchell A. Potter, Alan C. Schultz, Julian F. Miller, Edmund K. Burke, and Natasa Jonoska, editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2002*, pages 724–732, New York, 9-13 July 2002. Morgan Kaufmann Publishers.
12. M. Collins. Finding Needles in Haystacks is Harder with Neutrality. In Hans-Georg Beyer, Una-May O’Reilly, Dirk V. Arnold, Wolfgang Banzhaf, Christian Blum, Eric W. Bonabeau, Erick Cantu-Paz, Dipankar Dasgupta, Kalyanmoy Deb, James A. Foster, Edwin D. de Jong, Hod Lipson, Xavier Llorca, Spiros Mancoridis, Martin Pelikan, Guenther R. Raidl, Terence Soule, Andy M. Tyrrell, Jean-Paul Watson, and Eckart Zitzler, editors, *GECCO 2005: Proceedings of the 2005 Conference on Genetic and Evolutionary Computation*, volume 2, pages 1613–1618, Washington DC, USA, 25-29 June 2005. ACM Press.
13. Pietro Corsi. *The Age of Lamarck Evolutionary Theories in France*. University of California Press, USA, 1988.
14. Charles Darwin. *On The Origin of Species by Means of Natural Selection*. John Murray, 1859.

15. Dipankar Dasgupta and Douglas R. McGregor. Nonstationary Function Optimization using the Structured Genetic Algorithm. In R. Manner and B. Manderick, editors, *PPSN II: Proceedings of the 2nd International Conference on Parallel Problem Solving from Nature*, pages 145–154, Brussels, Belgium, 1992. Elsevier.
16. Benjamin Doerr, Michael Gnewuch, Nils Hebbinghaus, and Frank Neumann. A rigorous view on neutrality. In *IEEE Congress on Evolutionary Computation*, pages 2591–2597. IEEE, 2007.
17. R. M. Downing. Evolving Binary Decision Diagrams Using Implicit Neutrality. In *Proceedings of Congress on Evolutionary Computation (CEC 2005)*, volume 3, pages 2107–2113, Edinburgh, Scotland, 2005. IEEE Press.
18. M. Ebner. On the search space of genetic programming and its relation to nature’s search space. In *Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on*, 1999.
19. Marc Ebner, Patrick Langguth, Jurgen Albert, Mark Shackleton, and Rob Shipman. On Neutral Networks and Evolvability. In *Proceedings of the 2001 IEEE Congress on Evolutionary Computation*, pages 1–8. IEEE Press, 27-30May 2001.
20. Marc Ebner, Mark Shackleton, and Rob Shipman. How Neutral Networks Influence Evolvability. *Complexity*, 7(2):19–33, 2001.
21. A. E. Eiben, Robert Hinterding, and Zbigniew Michalewicz. Parameter control in evolutionary algorithms. *IEEE Trans. Evolutionary Computation*, 3(2):124–141, 1999.
22. Agoston E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer Verlag, 2003.
23. R. A. Fisher. On the Dominance Ratio. In *Proceedings of the Royal Society of Edinburgh*, volume 42, pages 321–341, 1922.
24. Carlos Fonseca and Marisol Correia. Developing Redundant Binary Representations for Genetic Search. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation (CEC 2005)*, pages 372–379, Edinburgh, 2-4 September 2005. IEEE.
25. W. Fontana and P Schuster. Continuity in Evolution: On the Nature of Transitions. *Science*, 280:1431–1452, 1998.
26. Edgar Galván-López. Effects of Neutrality on Evolutionary Search. In *EvoPhD - Second European Graduate Student Workshop on Evolutionary Computation*, Valencia, Spain, 11–13 April 2007.
27. Edgar Galván-López, Stephen Dignum, and Riccardo Poli. The Effects of Constant Neutrality on Performance and Problem Hardness in GP. In M. O’Neill, L. Vanneschi, S. Gustafson, A. I. Esparcia Alcazar, I. De Falco, A. Della Cioppa, and E. Tarantino, editors, *EuroGP 2008 - 11th European Conference on Genetic Programming*, volume 4971 of *LNCS*, pages 312–324, Napoli, Italy, 26–28 March 2008. Springer.
28. Edgar Galván-López, James McDermott, Michael O’Neill, and Anthony Brabazon. Defining locality in genetic programming to predict performance. In *2010 IEEE World Congress on Computational Intelligence*, pages 1828–1835, Barcelona, Spain, 18-23 July 2010. IEEE Computational Intelligence Society, IEEE Press.
29. Edgar Galván-López, James McDermott, Michael O’Neill, and Anthony Brabazon. Defining locality in problem hardness in genetic programming. *Genetic Programming and Evolvable Machines*, 2010.
30. Edgar Galván-López, James McDermott, Michael O’Neill, and Anthony Brabazon. Towards an understanding of locality in genetic programming. In *GECCO ’10: Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 901–908, New York, NY, USA, 2010. ACM.
31. Edgar Galván-López and Riccardo Poli. An Empirical Investigation of How and Why Neutrality Affects Evolutionary Search. In Maarten Keijzer, Mike Cattolico, Dirk Arnold, Vladan Babovic, Christian Blum, Peter Bosman, Martin V. Butz, Carlos Coello Coello, Dipankar Dasgupta, Sevan G. Ficici, James Foster, Arturo Hernandez-Aguirre, Greg Hornby, Hod Lipson, Phil McMinn, Jason Moore, Guenther Raidl, Franz Rothlauf, Conor Ryan, and Dirk Thierens, editors, *GECCO 2006: Proceedings of the 2006 Conference on Genetic and Evolutionary Computation*, pages 1149–1156, Seattle, WA, USA, 8-12 July 2006. ACM Press.
32. Edgar Galván-López and Riccardo Poli. Some Steps Towards Understanding How Neutrality Affects Evolutionary Search. In Thomas Philip Runarsson, Hans-Georg

- Beyer, Edmund Burke, Juan J. Merelo-Guervós, L. Darrell Whitley, and Xin Yao, editors, *Parallel Problem Solving from Nature (PPSN IX)*. 9th International Conference, volume 4193 of *LNCS*, pages 778–787, Reykjavik, Iceland, 9–13 September 2006. Springer-Verlag.
33. Edgar Galván-López and Riccardo Poli. The effects of constant and bit-wise neutrality on problem hardness, fitness distance correlation and phenotypic mutation rates. *IEEE Transactions on Evolutionary Computation*, 2010.
  34. N. Geard, J. Wiles, J. Hallinan, B. Tonkes, and B. Skellett. A Comparison of Neutral Landscapes – NK, NKp and NKq. In D.B. Fogel, M.A. El-Sharkawi, X. Yao, G. Greenwood, H. Iba, P. Marrow, and M. Shackleton, editors, *Proceedings of Congress on Evolutionary Computation (CEC 2002)*, volume 1, pages 205–210. IEEE Press, 2002.
  35. David E. Goldberg. Construction of High-Order Deceptive Functions Using Low-Order Walsh Coefficients. *Ann. Math. Artif. Intell.*, 5(1):35–47, 1992.
  36. David E. Goldberg, Kalyanmoy Deb, and Jeffrey Horn. Massive Multimodality, Deception, and Genetic Algorithms. In R. Männer and B. Manderick, editors, *PPSN II: Proceedings of the 2nd International Conference on Parallel Problem Solving from Nature*, pages 37–48, Amsterdam, 1992. Elsevier Science Publishers, B. V.
  37. J. B. S. Haldane. The Cost of Natural Selection. *J. Genet.*, 55:511–524, 1957.
  38. I. Harvey and A. Thompson. Through the Labyrinth Evolution Finds A Way: A Silicon ridge. In Tetsuya Higuchi, Masaya Iwata, and Weixin Liu, editors, *Proceedings of the First International Conference on Evolvable Systems: From Biology to Hardware*, volume 1259, pages 406–422. Springer-Verlag, 1996.
  39. J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.
  40. M. Huynen, P. Stadler, and W. Fontana. Smoothness Within Ruggedness: The Role of Neutrality in Adaptation. *Proc. Natl. Acad. Sci. (USA)*, 93:397–401, 1996.
  41. M. A. Huynen. Exploring Phenotype Space Through Neutral Evolution. *Molecular Evolution*, 43:165–169, 1996.
  42. Terry Jones. *Evolutionary Algorithms, Fitness Landscapes and Search*. PhD thesis, University of New Mexico, Albuquerque, 1995.
  43. H. Kargupta, K. Deb, and D. Goldberg. Ordering Genetic Algorithms and Deception. In R. Männer and B. Manderick, editors, *PPSN II: Proceedings of the 2nd International Conference on Parallel Problem Solving from Nature*, pages 49–58, Amsterdam, 1992. Elsevier Science Publishers, B. V.
  44. Y. Katada and K. Ohkura. Estimating the degree of neutrality in fitness landscapes by the nei’s standard genetic distance – an application to evolutionary robotics – . In Gary G. Yen, Simon M. Lucas, Gary Fogel, Graham Kendall, Ralf Salomon, Byoung-Tak Zhang, Carlos A. Coello Coello, and Thomas Philip Runarsson, editors, *Proceedings of the 2006 IEEE Congress on Evolutionary Computation*, pages 483–490, Vancouver, BC, Canada, 16–21 July 2006. IEEE Press.
  45. Yoshiaki Katada and Kazuhiro Ohkura. Analysis on topologies of fitness landscapes with both neutrality and ruggedness based on neutral networks. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation, GECCO ’09*, pages 1855–1856, New York, NY, USA, 2009. ACM.
  46. Yoshiaki Katada, Kazuhiro Ohkura, and Kanji Ueda. An approach to evolutionary robotics using a genetic algorithm with a variable mutation rate strategy. In Xin Yao, Edmund K. Burke, José Antonio Lozano, Jim Smith, Juan J. Merelo Guervós, John A. Bullinaria, Jonathan E. Rowe, Peter Tiño, Ata Kabán, and Hans-Paul Schwefel, editors, *PPSN*, volume 3242 of *Lecture Notes in Computer Science*, pages 952–961. Springer, 2004.
  47. S. A. Kauffman. *The Origins of Order – Organization and Selection in Evolution*. Oxford University Press, New York, 1993.
  48. Motoo Kimura. Evolutionary rate at the molecular level. In *Nature*, volume 217, pages 624–626, 1968.
  49. Motoo Kimura. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, Cambridge, UK, 1983.
  50. J. L. King and T. H. Jukes. Non-Darwinian Evolution. *Science*, 164:788–798, 1969.

51. Joshua D. Knowles and Richard A. Watson. On the Utility of Redundant Encodings in Mutation-Based Evolutionary Search. In Juan J. Merelo Guervós, Panagiotis Adamidis, Hans-Georg Beyer, José Luis Fernández-Villacañas Martín, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature - PPSN VII: 7th International Conference*, pages 88–98, Granada, Spain, 2002. Springer Verlag.
52. John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. The MIT Press, Cambridge, Massachusetts, 1992.
53. William B. Langdon and Riccardo Poli. *Foundations of Genetic Programming*. Springer, Berlin, 2002.
54. P. K. Lehre and P. C. Haddow. Accessibility Between Neutral Networks in Indirect Genotype-Phenotype Mappings. In *The 2005 IEEE Congress on Evolutionary Computation (CEC 2005)*, volume 1, pages 419–426, Edinburgh, UK, 2005. IEEE Press.
55. Per Kristian Lehre and Pauline Catriona Haddow. Phenotypic Complexity and Local Variations in Neutral Degree. *BioSystems*, 87(2-3):233–42, 2006.
56. J. Lobo, J. H. Miller, and W. Fontana. Neutrality in Technological Landscapes. In *Santa Fe Working Paper*, 2004.
57. Julian F. Miller. An Empirical Study of the Efficiency of Learning Boolean Functions Using a Cartesian Genetic Approach. In Wolfgang Banzhaf, Jason M. Daida, A. E. Eiben, Max H. Garzon, Vasant Honavar, Mark J. Jakiela, and Robert E. Smith, editors, *Proceedings of the Genetic and Evolutionary Computation Conference GECCO'99*, volume 2, pages 1135–1142, Orlando, Florida, 13–17 July 1999. Morgan Kaufmann.
58. Julian F. Miller and Stephen L. Smith. Redundancy and computational efficiency in cartesian genetic programming. *IEEE Transactions on Evolutionary Computation*, 10(2):167–174, April 2006.
59. Julian F. Miller and Peter Thomson. Cartesian genetic programming. In Riccardo Poli, Wolfgang Banzhaf, William Langdon, Julian Miller, Peter Nordin, and Terence Fogarty, editors, *Third European Conference on Genetic Programming EuroGP 2000*, volume 1802 of *LNCS*, pages 121–132, Edinburgh, 15–16 April 2000. Springer-Verlag.
60. Melanie Mitchell, Stephanie Forrest, and John H. Holland. The royal road for genetic algorithms: Fitness landscapes and GA performance. In Francisco J. Varela and Paul Bourguine, editors, *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 245–254, Cambridge, MA, 1992. MIT Press.
61. Tom M. Mitchell. *Machine Learning*. Mc Graw Hill, New York, 1996.
62. M. Newman and R. Engelhardt. Effects of Selective Neutrality on the Evolution of Molecular Species. *Proceedings of the Royal Society of London*, 265(1403):1333–1338, 1998.
63. M. E. J. Newman and Robin Engelhardt. Effects of neutral selection on the evolution of molecular species. Working Papers 98-01-001, Santa Fe Institute, January 1998.
64. M. E. J. Newman and Robin Engelhardt. Effects of Selective Neutrality on the Evolution of Molecular Species. In *Proceedings of the Royal Society of London*, volume 265, pages 1333–1338, 1998.
65. E. Van Nimwegen, J. P. Crutchfield, and M. Huynen. Neutral Evolution of Mutational Robustness. *Proc. Natl. Acad. Sci. USA*, 96(17):9716–9720, 1999.
66. Riccardo Poli and Edgar Galván-López. On The Effects of Bit-Wise Neutrality on Fitness Distance Correlation, Phenotypic Mutation Rates and Problem Hardness. In C. R. Stephens, M. Toussaint, D. Whitley, and P.F. Stadler, editors, *Foundations of Genetic Algorithms IX*, Lecture Notes in Computer Science, pages 138–164, Mexico city, Mexico, 8–11 January 2007. Springer-Verlag.
67. Riccardo Poli, William B. Langdon, and Nicholas Freitag McPhee. *A Field Guide to Genetic Programming*. Published via <http://lulu.com> and freely available at <http://www.gp-field-guide.org.uk>, 2008. (With contributions by J. R. Koza).
68. Riccardo Poli and Leonardo Vanneschi. Fitness-proportional negative slope coefficient as a hardness measure for genetic algorithms. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, GECCO '07, pages 1335–1342, New York, NY, USA, 2007. ACM.
69. C. Reidys, P. F. Stadler, and P. Schuster. Generic Properties of Combinatory Maps – Neutral Networks of RNA Secondary Structures. *Bull. Math. Biol.*, 59:339–397, 1997.

70. Christian M. Reidys and Peter F. Stadler. Neutrality in Fitness Landscapes. *Applied Mathematics and Computation*, 117(2–3):321–350, 2001.
71. Mark Ridley. *Evolution*. Blackwell Publishing, USA, 2003.
72. Franz Rothlauf and David Goldberg. Redundant Representations in Evolutionary Algorithms. *Evolutionary Computation*, 11(4):381–415, 2003.
73. P. Schuster, W. Fontana, P. F. Stadler, and I. L. Hofacker. From Sequences to Shapes and Back: A Case Study in RNA Secondary Structures. *Royal Society of London Proceedings Series B*, 255:279–284, March 1994.
74. Peter Schuster. Genotypes with Phenotypes: Adventures in an RNA Toy World. *Biophysical Chemistry*, 66(2):75–110, 1997.
75. M. A. Shackleton, R. Shipman, and M. Ebner. An Investigation of Redundant Genotype-Phenotype Mappings and Their Role in Evolutionary Search. In A. Zalzalá, C. Fonseca, J. H. Kim, and A. Smith, editors, *Proceedings of the International Congress on Evolutionary Computation (CEC 2000)*, pages 493–500. IEEE Press, 2000.
76. R. Shipman. Genetic Redundancy: Desirable or Problematic for Evolutionary Adaptation. In A. Dobnikar, N. C. Steele, D. W. Pearson, and R. F. Albrecht, editors, *4th International Conference on Artificial Neural Networks and Genetic Algorithms (ICANNGA'99)*, pages 337–344. Springer-Verlag, 1999.
77. Rob Shipman, Mark Shackleton, Marc Ebner, and Richard Watson. Neutral Search Spaces for Artificial Evolution: A Lesson From Life. In M. Bedau, S. Rasmussen, J. McCaskill, and N. Packard, editors, *Artificial Life: Proceedings of the Seventh International Conference on Artificial Life*, pages 162–169. MIT Press, 2000.
78. Tom Smith, Phil Husbands, Paul Layzell, and Michael O'Shea. Fitness landscapes and evolvability. *Evol. Comput.*, 10:1–34, March 2002.
79. Tom Smith, Phil Husbands, and Michael O'Shea. Neutral Networks and Evolvability with Complex Genotype-Phenotype Mapping. In *ECAL '01: Proceedings of the 6th European Conference on Advances in Artificial Life*, pages 272–281, London, UK, 2001. Springer-Verlag.
80. Tom Smith, Phil Husbands, and Michael O'Shea. Neutral Networks in an Evolutionary Robotics Search Space. In *Proceedings on Evolutionary Computation 2001*, pages 136–145. IEEE Press, 2001.
81. M. Toussaint. On the Evolution of Phenotypic Exploration Distributions. In Kenneth A. De Jong, Riccardo Poli, and Jonathan Rowe, editors, *Foundations of Genetic Algorithms 7 (FOGA 2003)*, pages 169–182. Morgan Kaufmann, 2003.
82. Marc Toussaint and Christian Igel. Neutrality: A necessity for self-adaptation. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2002)*, pages 1354–1359, 2002.
83. Erik van Nimwegen, James P. Crutchfield, and Melanie Mitchell. Statistical dynamics of the royal road genetic algorithm. *Theor. Comput. Sci.*, 229:41–102, November 1999.
84. Leonardo Vanneschi. *Theory and Practice for Efficient Genetic Programming*. PhD thesis, Faculty of Science, University of Lausanne, Switzerland, 2004.
85. Leonardo Vanneschi. Investigating problem hardness of real life applications. In Riolo et al., editor, *Genetic Programming Theory and Practice V*, chapter 7, pages 107–124. Springer US, 2007.
86. Leonardo Vanneschi. Fitness landscapes and problem hardness in genetic programming. In *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers, GECCO '09*, pages 3657–3684, New York, NY, USA, 2009. ACM.
87. Leonardo Vanneschi, Manual Clergue, Phillippe Collard, Marco Tomassini, and S. Verel. Fitness clouds and problem hardness in genetic programming. In *EuroGP, LNCS*, pages 690–701. Springer, 2004.
88. Leonardo Vanneschi, Yuri Pirola, Philippe Collard, Marco Tomassini, Sebastian Verel, and Giacarulo Mauri. A Quantitative Study of Neutrality in GP Boolean Landscapes. In Maarten Keijzer, Mike Cattolico, Dirk Arnold, Vladan Babovic, Christian Blum, Peter Bosman, Martin V. Butz, Carlos Coello Coello, Dipankar Dasgupta, Sevan G. Ficici, James Foster, Arturo Hernandez-Aguirre, Greg Hornby, Hod Lipson, Phil McMinn, Jason Moore, Guenther Raidl, Franz Rothlauf, Conor Ryan, and Dirk Thierens, editors, *GECCO 2006: Proceedings of the 2006 Conference on Genetic and Evolutionary Computation*, volume 1, pages 895–902, Seattle, WA, USA, 8–12 July 2006. ACM Press.

89. Leonardo Vanneschi, Marco Tomassini, Philippe Collard, and Manuel Clergue. Fitness Distance Correlation in Structural Mutation Genetic Programming. In Conor Ryan, Terence Soule, Maarten Keijzer, Edward P. K. Tsang, Riccardo Poli, and Ernesto Costa, editors, *Proceedings of the sixth European Conference on Genetic Programming, EuroGP 2003*, volume 2610 of *LNCS*, pages 455–464, Essex, 14–16 April 2003. Springer-Verlag.
90. Leonardo Vanneschi, Marco Tomassini, Philippe Collard, Sébastien Vérel, Yuri Pirola, and Giancarlo Mauri. A comprehensive view of fitness landscapes with neutrality and fitness clouds. In *EuroGP'07: Proceedings of the 10th European conference on Genetic programming*, pages 241–250, Berlin, Heidelberg, 2007. Springer-Verlag.
91. Leonardo Vanneschi, Andrea Valsecchi, and Riccardo Poli. Limitations of the fitness-proportional negative slope coefficient as a difficulty measure. In *GECCO '09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 1877–1878, New York, NY, USA, 2009. ACM.
92. Vesselin K. Vassilev, Terence C. Fogarty, and Julian F. Miller. Information characteristics and the structure of landscapes. *Evol. Comput.*, 8:31–60, March 2000.
93. Vesselin K. Vassilev and Julian F. Miller. The Advantages of Landscape Neutrality in Digital Circuit Evolution. In *ICES '00: Proceedings of the Third International Conference on Evolvable Systems*, pages 252–263, London, UK, 2000. Springer-Verlag.
94. Sébastien Vérel, Philippe Collard, and Manuel Clergue. Scuba search : when selection meets innovation. *CoRR*, abs/0707.0643, 2007.
95. Andreas Wagner. Robustness, Evolvability and Neutrality. *FEBS Letters*, 579(8):1772–1778, 2005.
96. K. Weicker and N. Weicker. Burden and Benefits of Redundancy. In W. Martin and W. Spears, editors, *Foundations of Genetic Algorithms 6*, pages 313–333, San Francisco, 2000. Morgan Kaufmann.
97. C. O. Wilke, J.L. Wang, C. Ofria, R.E. Lenski, and C. Adami. Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature*, 412:331–333, 2001.
98. Sewall Wright. The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution. In D. F. Jones, editor, *Proceedings of the Sixth International Congress on Genetics*, volume 1, pages 356–366, 1932.
99. Tina Yu and Julian Miller. Neutrality and the Evolvability of Boolean Function Landscape. In Julian F. Miller, Marco Tomassini, Pier Luca Lanzi, Conor Ryan, Andrea G. B. Tettamanzi, and William B. Langdon, editors, *Genetic Programming, Proceedings of EuroGP'2001*, volume 2038, pages 204–217, Lake Como, Italy, 18–20 2001. Springer-Verlag.
100. Tina Yu and Julian Miller. The Role of Neutral and Adaptive Mutation in an Evolutionary Search on the OneMax Problem. In William B. Langdon, Erick Cantú-Paz, Keith E. Mathias, Rajkumar Roy, David Davis, Riccardo Poli, Karthik Balakrishnan, Vasant Honavar, Günter Rudolph, Joachim Wegener, Larry Bull, Mitchell A. Potter, Alan C. Schultz, Julian F. Miller, Edmund K. Burke, and Natasa Jonoska, editors, *Late Breaking Papers at the Genetic and Evolutionary Computation Conference (GECCO-2002)*, New York, 9–13 July 2002. Morgan Kaufmann Publishers.
101. Tina Yu and Julian F. Miller. Finding Needles in Haystacks is not Hard with Neutrality. In James A. Foster, Evelyne Lutton, Julian F. Miller, Conor Ryan, and Andrea Tettamanzi, editors, *Genetic Programming Proceedings of the 5th European Conference, EuroGP 2002*, volume 2278 of *LNCS*, pages 13–25, Kinsale, Ireland, 3–5 April 2002. Springer-Verlag.