

Evolving Balanced Nonlinear Boolean Functions with Cellular Genetic Algorithms

Luigi Rovito¹, Andrea De Lorenzo¹, Mauro Castelli^{2,*}, Luca Manzoni³

¹ Department of Engineering and Architecture, University of Trieste, via Valerio 12, 34128, Trieste, Italy

² NOVA Information Management School (NOVA IMS), Universidade Nova de Lisboa, 1070-312 Lisbon, Portugal

³ Department of Mathematics, Informatics and Geosciences, University of Trieste, via Weiss 2, 34128, Trieste, Italy

ARTICLE INFO

Article history:

Received 1 May 2026

Received in revised form 26 July 2026

Accepted 10 September 2026

Available online 18 September 2026

Keywords:

Evolutionary Computation; Cybersecurity; Genetic Algorithms; Boolean Functions; Cellular Evolutionary Algorithms; Cryptography

ABSTRACT

Boolean functions are essential building blocks in the design of secure stream ciphers. Finding cryptographically strong Boolean functions is, however, a notoriously difficult optimization problem, owing to the size of the search space and the combinatorial nature of the task. Evolutionary Algorithms (EAs) have proven effective at constructing Boolean functions that are simultaneously balanced and highly nonlinear, two properties that are crucial for resisting correlation and statistical attacks. Nevertheless, these and other heuristic approaches often converge prematurely to local optima because of the complexity of the fitness landscape. Cellular Automata (CA) have been used extensively to represent and analyze Boolean functions, and spatially structured populations arranged on cellular toroidal grids have been shown, in other evolutionary paradigms, to slow the propagation of dominant individuals and to encourage broader exploration. The two ideas have not been brought together: cellular population structures have not previously been used to guide the search for balanced, nonlinear Boolean functions. We therefore investigate the integration of cellular population structures into Genetic Algorithms (GAs) for evolving balanced truth tables with high nonlinearity. We evolve Boolean functions of 8 to 16 variables and compare different neighborhood radii of the cellular grid. Our results show that, for larger numbers of variables, small neighborhoods improve the nonlinearity of the evolved functions relative to a non-spatial GA, and that this improvement is statistically significant. We emphasize that the contribution of this work is an analysis of the search and diversity dynamics induced by the cellular structure: the nonlinearity values obtained remain below the best-known values for balanced Boolean functions, and the study is not intended as a step forward in the construction of cryptographically deployable functions.

1. Introduction

Boolean functions are fundamental components of symmetric cryptography, where they are used to build secure stream ciphers. The security of such systems depends directly on the cryptographic strength of the Boolean functions employed: weak functions can be exploited through correlation attacks, linear approximation, or other statistical attacks [1, 2, 3]. Constructing Boolean functions that simultaneously satisfy several cryptographic criteria (such as balancedness, nonlinearity, and correla-

*Corresponding author.

E-mail address: mcastelli@novaims.unl.pt

<https://doi.org/10.31181/dma412026185>

© The Author(s) 2026 | [Creative Commons Attribution 4.0 International License](#)

tion immunity) is, however, a computationally hard problem, and no efficient constructive algorithm is known that covers all of these criteria at once [2, 3]. This is precisely what motivates the use of heuristic and evolutionary search techniques for the task.

Every n -variable Boolean function, where n is the number of inputs, can be represented by its truth table, that is, a binary string of length 2^n . The total number of n -variable Boolean functions is therefore 2^{2^n} , which grows doubly exponentially. A large search space alone does not imply computational hardness; the difficulty of constructing secure Boolean functions stems rather from the lack of efficient constructive algorithms for building highly secure functions [2]. The landscape is large, discrete, and multimodal, which makes exhaustive search infeasible even for $n > 5$, whereas stream ciphers typically require $n \geq 13$ [2, 3]. Moreover, even where theoretical upper bounds on properties such as nonlinearity are known, those bounds have not yet been attained for large n [3].

Against this background, Evolutionary Algorithms (EAs) have emerged as promising tools for discovering secure Boolean functions with strong cryptographic properties. Early work showed that Genetic Algorithms (GAs) [4, 5] and Genetic Programming (GP) [6, 7] can successfully evolve Boolean functions with good properties. Evolutionary search in this domain nonetheless remains challenging: the landscape is rugged and highly sensitive to small changes in the truth table, which leads to premature convergence, stagnation, and loss of diversity.

Despite the demonstrated effectiveness of GAs and GP on this task, and despite ample evidence that spatially structured, cellular populations mitigate premature convergence and preserve diversity in other evolutionary paradigms such as Geometric Semantic Genetic Programming (GSGP) and the Semantic Learning algorithm based on Inflate and deflate Mutation (SLIM) [8, 9], the integration of cellular population structures into GA-based search for balanced, nonlinear Boolean functions has not, to the best of our knowledge, been investigated before. This is a specific and, in our view, timely gap, given how susceptible this problem is to premature convergence and to loss of diversity.

Several studies have used Cellular Automatas (CAs) to model, represent, and analyze Boolean functions [10]. A CA is a classical nature-inspired computational model consisting of an m -dimensional grid of cells, each of which is in one of a finite number of states [11]. At every time step, all cells update their internal state synchronously according to a local rule that can observe only a neighborhood of the cell. Whereas these approaches use cellular models to encode Boolean functions, the idea of embedding cellular toroidal grids within EAs in order to influence the evolutionary process itself has not been explored for Boolean function optimization. Such spatially structured populations can increase diversity by slowing the propagation of dominant individuals and by promoting broader exploration of the search space [12, 13, 14].

Motivated by recent applications of cellular structures to EAs and GSGP [8, 9], we investigate whether integrating cellular structures into GAs can improve the search for balanced, nonlinear Boolean functions. Balancedness is essential because unbalanced functions introduce statistical biases that an attacker can exploit. Nonlinearity provides resilience against fast-correlation attacks and linear approximations [3].

We take nonlinearity as the optimization objective for two reasons. First, it is one of the most fundamental and most rigorously characterized cryptographic properties, and it admits a precise optimization target bounded by known theoretical limits (Section 2.2). Second, balancedness can be enforced as a hard, structurally preserved constraint within a truth-table-based GA representation (Section 4.1), rather than treated as a secondary, soft objective within the fitness function, as is usually necessary with tree-based (GP) representations. We adopt a cellular GA specifically because spatially structured populations have been shown to control selection pressure and to delay the propagation of dominant individuals in other evolutionary settings [12, 15], which addresses precisely the premature convergence and diversity loss that are known to affect evolutionary search for cryptographically strong Boolean functions.

Our aim is to understand how the spatial organization of the population influences the search, and whether interactions confined to small neighborhoods make that search more effective. We evolve Boolean functions of dimensionality $n = 8$ to $n = 16$ on a toroidal grid with varying neighborhood radii, and analyze the effect of the radius on the resulting nonlinearity and on the search dynamics.

We stress that the contribution of this work lies in the systematic integration and empirical characterization of cellular population structures within a GA operating on the truth-table representation of Boolean functions, and not in the introduction of new evolutionary operators or representations. Specifically, our contribution is threefold: (i) to the best of our knowledge, this is the first application of toroidal cellular population structures (previously shown to be effective in other evolutionary paradigms such as GSGP and SLIM [8, 9]) to the problem of evolving balanced, nonlinear Boolean functions through truth-table permutation; (ii) we provide a systematic study of the interaction between neighborhood radius, selection probability, and problem dimensionality ($n = 8$ to 16), which has not previously been characterized for this problem; and (iii) we complement the performance comparison with a diversity-based analysis using Moran's I [16], which explains the observed results mechanistically rather than reporting aggregate performance alone. This sets our work apart from earlier GA [4, 17, 18] and GP [7, 19] studies of Boolean functions, which do not use spatial population structuring [20], and from earlier work on cellular evolutionary algorithms, which has been applied to GSGP and SLIM but not to the evolution of Boolean functions by means of a GA.

To avoid any ambiguity, we state explicitly what is and what is not methodologically new. The cellular, spatially structured toroidal population model is not itself novel: it has been studied in GAs [15], in GP [21], and in Geometric Semantic GP [8]. What is new here is threefold. First, the mechanism is transferred to a search space that no previous cellular study has considered, namely the space of balanced truth tables explored under balancedness-preserving variation operators, in which neighborhood-local selection interacts with a hard structural invariant rather than with an unconstrained genotype. Second, we identify an interaction between neighborhood radius, selection probability, and problem dimensionality that is specific to this landscape and does not follow from existing results on cellular EAs, which were obtained on continuous symbolic-regression benchmarks with semantic operators. Third, we establish a quantitative rather than merely visual link between the persistence of spatial diversity and the quality of the final solution (Section 6.2). The paper should therefore be read as an empirical characterization of cellular search dynamics on a cryptographically motivated combinatorial landscape, and not as a proposal for a new evolutionary operator or representation. It is precisely because the representation and the operators are held fixed that the effect of the population topology can be isolated.

The Python 3.12 code used in the experiments is publicly available online*.

2. Background

This section introduces the notation used in the remainder of the paper and recalls the notions about Boolean functions and their cryptographic properties that we shall need.

2.1 Notation

We denote with n the number of variables of a Boolean function and with f a generic Boolean function. Let $\mathbb{B} = \{0, 1\}$, let $\oplus$ denote the sum operator (XOR) defined over $\mathbb{B}$, and let $\wedge$ denote the product operator (AND). Let $\mathbf{w} \cdot \mathbf{x} = \bigoplus_{i=1}^n \mathbf{w}_i \wedge \mathbf{x}_i$ be the inner product defined over $\mathbb{B}^n$. Let $w_H : \mathbb{B}^n \rightarrow \mathbb{N}$ be the Hamming weight of a binary string, that is, the number of non-zero values it contains. The Hamming weight of f is obtained by applying the Hamming weight to Ω_f , the truth table of f . Let $d_H : \mathbb{B}^n \times \mathbb{B}^n \rightarrow \mathbb{N}$ be the Hamming distance between two binary strings, that is, the number of positions in which the two strings differ (equivalently, the Hamming weight of their bitwise

*<https://github.com/lurovi/EvoNonLinearBoolPermutations>

XOR).

2.2 Boolean functions and their properties

A Boolean function $f : \mathbb{B}^n \rightarrow \mathbb{B}$ maps n binary inputs to a binary output and admits several representations, including its truth table, its algebraic normal form, and its Walsh spectrum $\mathbf{S}_f \in \mathbb{R}^{2^n}$. Given $f : \mathbb{B}^n \rightarrow \mathbb{B}$, its polar form is the function $\hat{f} : \mathbb{B}^n \rightarrow \{-1, 1\}$ defined by $\hat{f}(\mathbf{x}) = (-1)^{f(\mathbf{x})}$. The Walsh spectrum, obtained through the Walsh-Hadamard Transform (WHT) $\hat{F} : \mathbb{B}^n \rightarrow \mathbb{R}$ of the *polar (bipolar) representation* of f ,

$$\hat{F}(\mathbf{w}) = \sum_{\mathbf{x} \in \mathbb{B}^n} \hat{f}(\mathbf{x}) (-1)^{\mathbf{w} \cdot \mathbf{x}} = \sum_{\mathbf{x} \in \mathbb{B}^n} (-1)^{f(\mathbf{x}) \oplus (\mathbf{w} \cdot \mathbf{x})}, \quad (1)$$

makes it possible to measure key cryptographic indicators such as nonlinearity and resiliency.

The inverse transform can be used to reconstruct the polar representation $\hat{f}$ of a Boolean function from its Walsh spectrum $\mathbf{S}_f$; the function f is then recovered as $f(\mathbf{x}) = \frac{1 - \hat{f}(\mathbf{x})}{2}$. When the inverse transform is applied to a generic Walsh spectrum, such as that of a plateaued function [22], what is obtained is a pseudo-Boolean function in polar form $g : \mathbb{B}^n \rightarrow \mathbb{R}$, from which the nearest Boolean function must then be computed.

Among the most relevant properties, *balancedness* requires the truth table to contain equally many zeros and ones. Balancedness prevents the statistical bias that an attacker could otherwise exploit. The *nonlinearity* of f is defined as the minimum Hamming distance to the set of all affine functions,

$$\bar{\ell}(f) = 2^{n-1} - \frac{1}{2} \max_{\mathbf{w} \in \mathbb{B}^n} |\hat{F}(\mathbf{w})| \quad (2)$$

and it quantifies resistance to fast-correlation attacks. For even n , nonlinearity is bounded by the covering radius bound $U_n^{\text{crb}} = 2^{n-1} - 2^{\frac{n}{2}-1}$ [23], whereas for odd n the maximum nonlinearity of an n -variable Boolean function lies between $L_n^{\bar{\ell}} = 2^{n-1} - 2^{\frac{n-1}{2}}$ and $U_n^{\bar{\ell}} = 2 \lfloor 2^{n-2} - 2^{\frac{n-2}{2}} \rfloor$ [24]. Boolean functions that attain the covering radius bound for even n are called bent functions [25, 26]. Bent functions have optimal nonlinearity, but by construction they cannot be balanced.

Other important characteristics include *resiliency* (t), which allows a function to resist correlation attacks, and *algebraic degree* (d), the degree of the multivariate polynomial that represents the algebraic normal form of f ; a high algebraic degree makes attacks based on the Berlekamp-Massey algorithm less effective [27]. Tarannikov's bound states that $\bar{\ell} \leq 2^{n-1} - 2^{t+1}$ [28], while Siegenthaler's bound states that $d \leq n - t - 1$ [29]. The fact that several security properties are constrained by bounds of this kind illustrates how difficult it is to optimize all of them at once.

3. Related Work

This section reviews the state of the art in the search for Boolean functions and in cellular methods for optimization.

3.1 Optimization of Boolean functions

Building Boolean functions with strong cryptographic properties is a difficult problem that has been addressed by means of algebraic constructions [3, 30], heuristic techniques [31, 32], and EAs [33]. In particular, Carlet et al. [34] proposed a method for evolving constructions of nonlinear, balanced functions in which GP [35] rediscovered the indirect sum construction. The same authors also combined EAs with established algebraic constructions [36], while Picek et al. [37] used GP to evolve algebraic constructions of bent functions. Picek et al. [38] likewise explored the integration of algebraic and evolutionary approaches in the search for secure Boolean functions.

GAs [39] were first used to evolve balanced Boolean functions directly, by permuting their truth

tables [4, 5]. The truth-table encoding has since been used to generate bent functions [20], to identify balanced functions satisfying various criteria such as nonlinearity, algebraic degree, and correlation immunity [18], and to discover nonlinear functions by combining evolutionary search with hill climbing [17] and with hybrid local search [40]. In this context, Manzoni et al. [41] also proposed a family of balancedness-preserving crossover operators for GAs.

A second encoding of Boolean functions that can be evolved with a GA is the Walsh spectrum. Clark et al. [32] pioneered this spectral inversion technique, using simulated annealing to permute three-valued Walsh spectra in order to discover plateaued Boolean functions. Mariot et al. [42] built on that idea and used GAs to evolve plateaued Boolean functions directly. Mariot et al. [43] further implemented a particle swarm optimizer to generate Boolean functions with desirable cryptographic properties.

A substantial body of Evolutionary Computation (EC) literature represents Boolean functions as trees and evolves them with GP [44, 45, 46]. An early attempt to use GP to evolve Boolean functions in cryptographic settings was made by Picek et al. [47] and was subsequently extended in a number of follow-up studies [6, 7]. This work showed that evolving syntax trees with GP generally yields higher nonlinearity than evolving truth tables directly. Additional constraints such as balancedness must, however, be incorporated explicitly into the fitness function, since they cannot be encoded in the representation itself. More recently, Mariot et al. [48] used GAs and GP to construct perfectly balanced Boolean functions. Spectral inversion, introduced by Clark et al. [32] and later combined with evolutionary search by Mariot et al. [42], remains the principal alternative to the direct evolution of truth tables.

Further studies have used EAs to evolve homogeneous bent functions [49], self-dual bent functions [50], and rotation-symmetric functions [51, 52]. Notably, Carlet et al. have recently compared the nonlinearity achieved by different Boolean encodings for functions of odd size [53] and have used several encodings to evolve Boolean functions with five-valued spectra [54].

3.2 Cellular-inspired optimization

Introducing spatial and cellular structures into EAs has already proven promising [12], particularly in GAs [15] and in GP [21]. A further widely used approach to preserving diversity is speciation, which divides the population into niches of individuals that share structural traits [55, 56, 57].

Earlier work has shed light on why spatial structures matter: it has shown that they control the takeover time, that is, the number of generations a dominant individual needs to spread through the whole population, and it has analyzed the effect of selection pressure in EAs with and without an imposed cellular structure [13, 58].

Recent work on cellular structures has also considered GSGP [59], a variant of GP that uses semantic operators and has been shown to outperform GP on a range of tasks [59, 60]. GSGP has two main drawbacks: exponential tree growth and premature convergence [60]. To address the first of these, Vanneschi et al. introduced the SLIM [61], a variant of GSGP that can produce offspring smaller than their parents. In this setting, recent work has successfully applied cellular structures to the populations of both GSGP [8] and SLIM [9], showing that cellular variants can outperform vanilla GSGP and certain variants of SLIM by improving diversity in the early part of the evolutionary run.

4. Proposed Method

We propose a cellular GA for evolving balanced Boolean functions with high nonlinearity. All algorithmic components are described below; the specific hyper-parameter values are reported in Sec. 5.

4.1 Problem setting and representation

We consider Boolean functions of n variables represented explicitly by their truth tables, that is, by binary strings of length 2^n . All candidate solutions are constrained to be balanced, so that their

truth tables contain exactly 2^{n-1} zeros and 2^{n-1} ones.

The population is initialized by generating several random permutations of a fixed balanced truth table, which guarantees that every initial individual satisfies the balancedness constraint. Throughout the evolutionary run, all variation operators (crossover and mutation) are designed to preserve this property.

Note that any balanced truth table can serve as a starting point for building another balanced truth table by random permutation.

4.2 Fitness function

The optimization objective is to maximize nonlinearity. For each f , the WHT is applied to its truth table to obtain the Walsh spectrum, from which the nonlinearity is derived as described in Sec. 2.2.

To improve the guidance of the search and to facilitate transitions between nonlinearity levels, we adopt the smoothed fitness function

$$\tilde{\ell}(f) = \bar{\ell}(f) + \frac{2^n - A_{\mathbf{s}_f}}{2^n}, \quad (3)$$

where $A_{\mathbf{s}_f}$ denotes the number of occurrences of the maximal absolute coefficient in the spectrum of f .

Since a higher nonlinearity corresponds to a smaller maximal spectral magnitude, minimizing $A_{\mathbf{s}_f}$ drives the search towards higher nonlinearity levels. The second term lies in $[0, 1)$, since $1 \leq A_{\mathbf{s}_f} \leq 2^n$, so that $\lfloor \tilde{\ell}(f) \rfloor = \bar{\ell}(f)$ and the smoothing acts only as a tie-breaker within a nonlinearity level. This formulation of the fitness function has been used successfully in earlier work [33, 53, 54].

4.3 Cellular selection mechanism

We use a cellular framework inspired by [8, 9] to control selection pressure and the diffusion of information.

Individuals are arranged on an m -dimensional toroidal grid in which each cell holds one individual, assigned at random during initialization, and interacts only with the cells in its neighborhood. The neighborhood of radius r of a cell i is the Moore neighborhood $\mathcal{N}_i$, that is, a hypercube of side length $2r + 1$ centered on i . The toroidal topology ensures that all neighborhoods have the same size. We denote this structure by $\mathcal{T}_r^m$.

Selection and variation are performed locally. For each cell i , a subset $\mathcal{S} \subseteq \mathcal{N}_i$ is formed by including each neighbor independently with probability $p \in (0, 1]$. Rank-based selection is then applied to $\mathcal{S}$. If crossover is applied, the offspring is generated by crossing the two best individuals in $\mathcal{S}$; otherwise, the best individual in $\mathcal{S}$ is mutated. The resulting offspring replaces the individual in cell i .

4.4 Crossover

We use a position-based crossover with count repair. Given two balanced parents T_1 and T_2 , half of the positions are selected at random and copied from T_1 into the offspring, and the remaining positions are filled with the corresponding bits of T_2 . Because this operation may violate balancedness, a repair step follows: if the child has too many ones (respectively, zeros), a matching number of ones (respectively, zeros) is flipped to zeros (respectively, ones) at random.

4.5 Mutation

Mutation is implemented as a k -swap operator. At each mutation event, k disjoint pairs of positions are selected at random and their values are swapped. The parameter k is drawn uniformly from a predefined integer interval $[k_1, k_2]$, and balancedness is preserved by construction.

4.6 Overall algorithm

Algorithm 1 summarizes the proposed method. Crossover and mutation are mutually exclusive: if crossover is performed, mutation is not, and vice versa.

Algorithm 1 Cellular GA for Boolean function optimization

```

1: Initialize a population of balanced truth tables on a  $\mathcal{T}_r^m$  grid
2: Evaluate the fitness of all individuals (generation 0)
3: for  $j$  in  $[1..\text{generations}]$  do
4:   for each cell  $i$  in the grid do
5:     Construct the neighborhood  $\mathcal{N}_i$ 
6:     Sample a subset  $\mathcal{S} \subseteq \mathcal{N}_i$  with probability  $p$ 
7:     Select the best individual(s) from  $\mathcal{S}$ 
8:     if crossover is applied then
9:       Generate the offspring by position-based crossover with repair
10:    else
11:      Generate the offspring by  $k$ -swap mutation
12:    end if
13:    Replace the individual in cell  $i$  with the offspring
14:  end for
15:  Evaluate the new population (generation  $j$ )
16:  Replace the worst-fitness individual with the best-so-far individual (elitism)
17: end for
18: return the best Boolean function found
    
```

When sampling $\mathcal{S} \subseteq \mathcal{N}_i$ with probability p , we always ensure that $\mathcal{S}$ contains at least two individuals. If this constraint is not met after sampling, one or two individuals drawn at random from $\mathcal{N}_i$ are added as required. This guarantees that crossover or mutation can always be performed.

5. Experimental Setup

Our goal is to assess whether a cellular configuration can improve the nonlinearity of the evolved Boolean functions relative to a standard GA. The parameters used in this study are not intended to be optimal settings for Boolean function optimization; they have been set mainly to values that are standard in evolutionary optimization. This choice is consistent with the primary goal of the work, which is to analyze the effect of cellular selection and of the diversity preservation it provides.

We study Boolean functions with $n \in \{8, \dots, 16\}$ and perform 50 independent runs of each configuration. The population size is 100 and the population is evolved for 1000 generations, as in [8, 9]. As earlier work has shown [41, 42], a small population evolved for many generations yields better solutions in this setting than a large population evolved for fewer generations.

The population is initialized by generating several random permutations of a fixed balanced truth table made of alternating 1 and 0 bits.

The crossover and mutation probabilities are both fixed at 0.5, and the two operators are mutually exclusive. The mutation parameters k_1 and k_2 are set to 1 and 5, respectively. This interval was chosen so as to avoid excessively disruptive mutations, which would change the truth table too heavily and destroy the structure accumulated during the run.

We emphasize that the population size, the number of generations, the mutation interval, and the dimensionality of the grid were held fixed throughout this study, in line with our own earlier work on cellular structures in other evolutionary paradigms [9] and with earlier GA-based studies of Boolean functions [4, 7, 41]. The only factors that vary across configurations are therefore the cellular parameters under investigation, namely the neighborhood radius r and the selection probability p . In this

way, the effect of the cellular population structure is not confounded with the effect of jointly tuning the remaining hyper-parameters for optimal performance, which is beyond the scope of this study.

The specific values were chosen as follows. The population size (100) and the generation budget (1000) reproduce the setting used in the studies on cellular EAs on which this work builds [8], and are consistent with the finding, established in the Boolean function literature [4, 7] and confirmed again in Section 6.3, that for a fixed number of fitness evaluations a small population evolved for many generations dominates a large population evolved for few generations. The mutation interval $[1, 5]$ limits each mutation to at most five transpositions out of 2^n truth-table positions, that is, to a relative perturbation of at most $5 \cdot 2^{1-n}$, which keeps mutation non-disruptive across the whole range of n considered here; this is precisely the regime in which cellular structures have been reported to preserve, rather than destroy, spatial correlation [9]. The 10×10 square toroidal grid is the smallest two-dimensional lattice compatible with a population of 100 that keeps the neighborhood uniform along both dimensions and admits radii up to $r = 3$ without wrap-around overlap. Section 6.3 varies the population size and the generation budget at a constant evaluation budget and shows that the qualitative conclusions of Section 6.1 are preserved once the neighborhood size is rescaled accordingly.

We first compare two non-cellular baselines. The first is a standard GA with tournament selection (tournament size 4). The second extends the first with explicit elimination of duplicates: whenever the variation operators produce an offspring identical (in terms of its truth table) to one already produced in the current generation, that offspring is discarded and variation is retried up to 10 times. If no unique offspring is obtained after these retries, the duplicate is accepted. This mechanism promotes population diversity by limiting the number of identical individuals. Note that both baselines are identical to Algorithm 1 except for the absence of the grid and, consequently, for the selection scheme (tournament selection instead of the cellular strategy). The elimination of duplicates used in the second baseline is not used in our cellular GA, since diversity there should be preserved naturally.

We compare the best nonlinearity values obtained across the repetitions using the Wilcoxon–Mann–Whitney test [62] with $\alpha = 0.05$. For $n = 8$ and $n = 11$, the diversity-preserving GA (with 10 retries) significantly outperforms the vanilla GA; for all other values of n , the two are statistically equivalent. Since the GA with elimination of duplicates performs at least as well as the standard one, we adopt it as our baseline and denote it by $\mathcal{T}^0$, as it uses no toroidal structure.

For the cellular configurations, we use a 10×10 two-dimensional toroidal grid. We test configurations with $p \in \{0.25, 0.5, 0.75, 1.0\}$ and $r \in \{1, 2, 3\}$. The toroidal configurations with the three radii are denoted by $\mathcal{T}_1^2$, $\mathcal{T}_2^2$, and $\mathcal{T}_3^2$, respectively. The values of p were chosen to span a range from weak to strong selection. In particular, $p = 1.0$ corresponds to deterministic selection over the full neighborhood. Selection is performed within local neighborhoods as described in Sec. 4, with no explicit elimination of duplicates. Unless stated otherwise, all parameters are kept fixed across experiments, so as to isolate the effect of the cellular selection mechanism rather than to maximize performance through tuning.

We analyze five key measures: (i) the nonlinearity $\bar{\ell}$ of the best solution found, which is the cryptographic property to be maximized (balancedness being guaranteed by design); (ii) the median, over the population, of the fitness, i.e. the smoothed nonlinearity $\tilde{\ell}$; (iii) the average pairwise Hamming distance (H) between the truth tables of the individuals in the population, used as a complementary, non-spatial measure of genotypic diversity; (iv) the average pairwise Euclidean distance (E) between the Walsh spectra of the individuals in the population, used as a complementary, non-spatial measure of semantic diversity; and (v) the global Moran’s I (I) [16], a measure of spatial autocorrelation that quantifies diversity in neighborhood-based populations.

The statistic I measures how similar neighboring individuals are. Representing each individual by

its semantic vector, namely its Walsh spectrum, we compute

$$I = \frac{M}{W} \frac{\sum_{i=1}^M \sum_{j=1}^M \mathbf{w}_{ij} (\mathbf{x}^i - \bar{\mathbf{x}})(\mathbf{x}^j - \bar{\mathbf{x}})}{\sum_{i=1}^M (\mathbf{x}^i - \bar{\mathbf{x}})^2}, \quad (4)$$

where M is the population size, $\mathbf{w}_{ij} \in \mathbb{R}$ is the (i, j) -th entry of the neighborhood matrix $\mathbf{w} \in \mathbb{R}^{M \times M}$, $\mathbf{w}_{ii} = 0$ for $1 \leq i \leq M$, $W = \sum_{i=1}^M \sum_{j=1}^M \mathbf{w}_{ij}$, $\mathbf{x}^i$ is the semantic vector of the i -th individual in the population, and $\bar{\mathbf{x}}$ is the mean of the semantic vectors of all individuals in the population. Here $\mathbf{w}_{ij} = 1$ if $i \neq j$ and the j -th individual belongs to the neighborhood of i , and $\mathbf{w}_{ij} = 0$ otherwise. In the non-cellular methods, the whole population forms a single neighborhood.

The values of I lie in $[-1, 1]$: values close to 1 indicate strong positive spatial autocorrelation, that is, semantically similar individuals clustered together, whereas values close to -1 indicate negative autocorrelation. Values around 0 indicate a random spatial pattern. A positive I means that neighborhoods contain semantically similar individuals that differ from those in other regions of the grid, which indicates that the toroidal structure is effectively maintaining spatial diversity across the population.

All experiments were run on a HP ProLiant DL560 Gen10 with quad Intel Xeon Gold 6140 CPU (18 cores) and 1 TB of RAM.

6. Results and Discussion

This section reports the experimental results in terms of the best nonlinearity achieved, and then examines the effect of the cellular structure on population diversity and on the distribution of fitness.

Section 6.1 reports the distribution of the best nonlinearity values obtained by each configuration (Figure 1), reports confidence intervals (Table 1), compares the best nonlinearity values with the globally best-known values for balanced Boolean functions (Table 2) and against the values reported in the literature for alternative evolutionary paradigms (Table 3), summarizes the corresponding statistical analysis, and shows the convergence trends of the best-so-far nonlinearity (Figure 2). Section 6.2 analyzes the median fitness of the whole population (Figure 3), the spatial autocorrelation measured by Moran's I (Figure 4), the average pairwise Hamming and Euclidean distances (Figure 5 and Figure 6), the relationship between the persistence of diversity and the quality of the final solution (Table 4), and the spatial distribution of fitness over the grid (Figure 7). Section 6.3 reports the scaling analysis obtained with a larger population and a reduced generation budget (Figure 8).

6.1 Best nonlinearity

We now compare the best nonlinearity values obtained by the cellular configurations with those of the baseline. Statistical comparisons use the Wilcoxon–Mann–Whitney test [62] with $\alpha = 0.05$, and p -values adjusted by the Holm–Bonferroni correction [63], following best practices in the EC literature [64, 65]. For each configuration, a pairwise test against the baseline determines whether the cellular setup attains significantly higher nonlinearity (12 pairwise tests per n). The results are reported in Fig. 1, where Cliff's Delta [66] values identify the configurations that outperform the baseline. Table 1 reports 95% confidence intervals for the best nonlinearity values of the baseline and of the best cellular configuration according to Fig. 1, which we discuss next.

The box plots show that, for a small number of variables, cellular structures bring no improvement over the baseline. As the search space grows, that is, as n increases, large neighborhoods ($r = 2$ or $r = 3$) combined with high selection pressure ($p = 0.75$ or $p = 1.0$) tend to degrade performance, which suggests that wide interaction ranges and strong selection pressure can hinder the discovery of better solutions. Conversely, for $n \geq 13$ the configuration with radius $r = 1$ and selection probability $p = 0.5$ consistently produces statistically better solutions than the baseline, with Cliff's Delta values above 0.45. Even on complex fitness landscapes, therefore, local interaction can improve the search, provided that the neighborhood is small and the selection pressure moderate.




Fig. 1. Distribution, across repetitions, of the best final nonlinearity values. Cliff's Delta values annotate the cellular methods that are better than the non-cellular baseline.

Table 1

95% confidence intervals (Student's t) of the best nonlinearity values found across the 50 repetitions of $\mathcal{T}^0$ and $\mathcal{T}_1^2$ with $p = 0.5$. For $n = 9$, all runs attain the same nonlinearity value (236).

n	95% CI of best $\bar{l}$	
	$\mathcal{T}^0$	$\mathcal{T}_1^2$
8	[115.52, 115.92]	[114.81, 115.43]
9	[236.00, 236.00]	[236.00, 236.00]
10	[479.81, 480.03]	[479.88, 480.04]
11	[974.46, 975.30]	[973.73, 974.59]
12	[1969.80, 1970.84]	[1970.61, 1971.47]
13	[3976.01, 3977.03]	[3977.89, 3978.83]
14	[8007.94, 8009.66]	[8010.34, 8011.98]
15	[16106.09, 16108.23]	[16110.39, 16112.49]
16	[32348.74, 32351.10]	[32355.08, 32357.32]

We observe that a particular ratio between neighborhood size and population size ($9/100 = 0.09$ in our case) appears to be what allows the cellular toroidal grid to improve on the baseline. This should be taken into account when the population size is scaled up, since the neighborhood size may have to be scaled up with it. We return to this point in Section 6.3.

Tab. 2 compares the best-known nonlinearity values for balanced functions with those obtained

by $\mathcal{T}^0$ and by $\mathcal{T}_1^2$ with $p = 0.5$, the best cellular setting.

Table 2

Median and maximum of the best nonlinearity values found across the 50 repetitions of $\mathcal{T}^0$ and of $\mathcal{T}_1^2$ with $p = 0.5$, compared with the globally best-known nonlinearity values for balanced Boolean functions.

n	Best-known $\bar{\ell}$	$\mathcal{T}^0$		$\mathcal{T}_1^2$	
		Max	Median	Max	Median
8	116	116	116	116	116
9	240	236	236	236	236
10	492	480	480	480	480
11	992	978	976	976	974
12	2010	1974	1970	1974	1972
13	4036	3982	3976	3982	3978
14	8120	8016	8008	8016	8012
15	16 272	16 116	16 107	16 120	16 112
16	32 638	32 358	32 350	32 364	32 356

The cellular configuration outperforms the baseline for the larger values of n . For $n \in \{15, 16\}$ in particular, the cellular setup reaches nonlinearity levels that the baseline never attains. This suggests that the performance gap between cellular and non-cellular evolution widens as n increases, which is of interest because large values of n constitute the practically relevant regime [2, 3].

Despite this improvement, the nonlinearity values achieved by the cellular GA remain below the best-known values reported in the literature for balanced Boolean functions (Table 2). We attribute this gap to two main factors.

First, as noted in Sec. 5, the hyper-parameters used in this study (population size, number of generations, and mutation interval) were deliberately fixed to standard values rather than tuned for performance, in order to isolate the effect of the cellular population structure; the best-known values reported in Tab. 2, by contrast, are typically the result of extensive tuning accumulated over several independent studies.

Second, many of the best-known constructions exploit algebraic structure, for example the indirect sum construction or constructions based on bent functions [34, 36], which is not directly accessible to the structure-agnostic search over the space of balanced truth tables that we adopted here. Closing this gap through hyper-parameter optimization, through hybridization with algebraic constructions, or both, is left for future work (Section 7).

It is worth situating these results with respect to methods based on other evolutionary paradigms and representations. Approaches based on GP [6, 7], on algebraic constructions [34], or on hybrid evolutionary and algebraic techniques [36, 43] often report higher nonlinearity values than those obtained here (Table 2). This is to be expected, since such approaches either exploit algebraic structure that a permutation-based search over truth tables cannot access, or treat balancedness as a soft, penalty-based fitness criterion rather than as a structural invariant of the representation, as it is in the present work (Section 4.1). A direct experimental comparison with these paradigms would therefore confound the effect of the toroidal population structure investigated here with effects due to representation and constraint handling. We consequently regard such a comparison as an important but methodologically distinct direction for future work (Section 7).

To make this positioning quantitative rather than narrative, Table 3 collects the results reported in the literature for the principal alternative methodological families that have been applied to this problem, together with the representation used, the constraint-handling strategy, the presence or absence of local-search refinement, and the range of function sizes actually studied in each case. Each

row reports the range of function sizes actually studied in the cited work and the best result reported there for balanced functions, together with the size at which that result was obtained. The sizes are not aligned across rows because the cited works do not study the same sizes; forcing a common set of columns would leave most cells empty and would misrepresent the literature. The setting adopted in the present work, in which balancedness is enforced as a hard structural invariant and nonlinearity is the sole optimization objective, is uncommon in this literature: most of the cited works either impose additional cryptographic criteria alongside nonlinearity or relax balancedness altogether. The rows are therefore not directly commensurable, and we present them as context rather than as a like-for-like comparison.

Table 3

Table collecting the results reported in the literature for the principal alternative methodological families that have been applied to Boolean functions optimization.

Work	Method family	Representation	Balancedness	Local search	n studied	Best reported $\bar{\ell}$ (at n)	Additional criteria
[4]	GA, truth table	Truth table	Structural	Yes	8–12	112(8), 232(9), 476(10), 976(11), 1972(12)	CI(1)
[32]	Spectral inversion (SA / GA)	Walsh spectrum	Indirect	No	6–9	No balanced solutions obtained for $n \geq 8$	Plateaued, resiliency
[43]	Particle swarm optimization	Truth table (swap-based)	Structural	Yes	8–12	116(8), 236(9), 480(10), 976(11), 1972(12)	CI(1), PC(1) or absolute indicator
[6, 7]	Genetic programming	Syntax tree	Fitness penalty	No	Mainly 8	116(8)	Multiple design criteria
[53]	GP and GA, odd sizes	Tree / bit string	Not enforced	Optional	7, 9, 11, 13	GP: 240(9), 992(11), 4032(13); GA: 236(9), 978(11), 3980(13)	Nonlinearity only
[34, 36]	EA-assisted algebraic constructions	Construction	Structural	No	9–18	240(9), 492(10), 992(11), 2008(12), 4032(13), 8120(14), 16256(15), 32608(16), 65280(17), 130753(18)	Nonlinearity only
This work	Cellular GA	Truth table	Structural	No	8–16	116(8), 236(9), 480(10), 978(11), 1974(12), 3982(13), 8016(14), 16120(15), 32364(16)	Nonlinearity only

Two observations follow. First, the residual gap with respect to the best-known values is systematically associated with a method’s ability to exploit algebraic structure, whether through explicit constructions or through hybrid algebraic and evolutionary schemes, and with the use of local-search refinement, rather than with the presence or absence of a population structure.

Second, and directly relevant to the scope of the present study, most of the evolutionary approaches in this literature report results for function sizes up to $n = 12$ or below, and comparable figures for $n = 13, 14, 15$, and 16 are largely absent. The regime in which the cellular structure is beneficial in our experiments is therefore also the regime that prior evolutionary work covers least, which is one reason for studying it. Because these methods differ from ours in representation, in constraint handling, and in computational budget, Table 3 is offered as context and not as a controlled comparison; the controlled comparison is the one in Figure 1, in which the representation and the operators are identical across all methods and only the population topology varies.

Finally, the plots in Fig. 2 show that, for $n \geq 14$, the configurations with radius 1 begin to outperform the baseline before the midpoint of the run, which reinforces the trends observed in the statistical analysis.

6.2 Population fitness and diversity

We also analyze the behavior of the population as a whole. Figure 3 shows the trend of the median fitness of the population over the generations, which reflects the behavior of all individuals rather than that of the best one alone.

For large n , this trend closely follows the one in Fig. 2, while persistent variability remains within the population. To gain further insight, we analyze diversity by means of Moran’s I , as in [8, 9].




Fig. 2. Trend of the best-so-far nonlinearity (median across the repetitions) for $p = 0.5$. The shaded area denotes the inter-quartile range.

Figure 4 tracks the trend of I over the generations and shows that, for radius 1 and $n \geq 14$ in particular, the cellular grid maintains diversity during the first half of the run. This suggests that local neighborhoods slow the spread of dominant individuals and thereby preserve the spatial organization of the population. Eventually, dominant solutions propagate through all neighborhoods and I converges to zero. For $n \in \{15, 16\}$, $\mathcal{T}_1^2$ maintains higher values of I even in the later generations.

This behavior is consistent with established theoretical results on selection pressure in spatially structured EAs. The takeover time has been shown to grow as the neighborhood size decreases: populations that are not arranged on a lattice exhibit fast takeover, whereas spatially structured populations with small neighborhoods exhibit substantially slower takeover [13, 58]. Since $r = 1$ defines the smallest Moore neighborhood on our toroidal grid, it minimizes the rate at which high-fitness genotypes diffuse through the population, which explains why positive values of Moran's I persist longer for $r = 1$ than for $r = 2$ and $r = 3$ (Figure 4). A moderate selection probability ($p = 0.5$), as opposed to deterministic full-neighborhood selection ($p = 1.0$), further reduces the effective selection intensity. Together, these two effects explain why this configuration strikes the best balance between exploration and exploitation on the rugged, multimodal nonlinearity landscape underlying the problem.

To corroborate the diversity trends captured by Moran's I , we additionally computed the average pairwise Hamming distance between the truth tables of the individuals in the population over the




Fig. 3. Trend of the median fitness of the whole population (median across the repetitions) for $p = 0.5$. The shaded area denotes the inter-quartile range.

generations (Fig. 5), as well as the average pairwise Euclidean distance between their Walsh spectra (Fig. 6). The trends visible in both plots confirm that cellular configurations with small neighborhood radii maintain higher genotypic and semantic diversity over a larger fraction of the run than the non-cellular baseline, in line with the analysis based on spatial autocorrelation above. This is especially clear for $r = 1$ and $n \geq 13$, consistently with the statistical findings in Fig. 1. The additional analysis in Fig. 5 and Fig. 6 thus shows that the spatial arrangement provided by the cellular method with a small neighborhood radius does improve population diversity.

To quantify the relationship between diversity and best discovered nonlinearity directly, we summarize each diversity trajectory (computed separately for I , H , and E) by its *median* value over the first half of the run (generations 1–500), and we correlate this summary with the final best-so-far nonlinearity of the same repetition, separately for each method and each n . The resulting Pearson correlations (ρ) are reported in Tab. 4. The Pearson correlation is computed across the 50 repetitions of $\mathcal{T}^0$, $\mathcal{T}_1^2$, $\mathcal{T}_2^2$, and $\mathcal{T}_3^2$ with $p = 0.5$, for each n . Correlations marked with * remain significant ($\alpha = 0.05$) after Holm–Bonferroni correction for multiple comparisons within each row. A dash (“–”) indicates that the diversity metric, the final nonlinearity, or both are (roughly) constant across all 50 repetitions for that configuration, leaving the correlation coefficient undefined. This occurs when every repetition reaches the same best-known nonlinearity value (e.g., $n = 9$), and when diversity has collapsed to the same value in every repetition within the first 500 generations, which is common for




Fig. 4. Moran's I of the population (median across the repetitions) over the generations, for $p = 0.5$. The shaded area denotes the inter-quartile range.

$\mathcal{T}_2^2$ and $\mathcal{T}_3^2$ under the Hamming- and Euclidean-distance metrics.

The relationship holds most consistently for $\mathcal{T}_1^2$, the configuration identified in Sec. 6.1 as the strongest cellular setting: the correlation between early Moran's I and final nonlinearity is positive for $n \geq 11$. The non-spatial diversity measures (H and E) show a similar pattern for both the baseline and $\mathcal{T}_1^2$ when $n \geq 11$. We find no consistent within-method relationship for $\mathcal{T}_2^2$ or $\mathcal{T}_3^2$, nor for $n \leq 10$, where the search space is small enough that most repetitions reach the best-known or a near-best-known nonlinearity value regardless of their diversity trajectory (cf. the degenerate confidence interval reported for $n = 9$ in Tab. 1). We therefore read Tab. 4 as partial, method- and n -specific quantitative support for the mechanism illustrated in the line plots, concentrated in $\mathcal{T}_1^2$ and in the larger values of n where the cellular structure helps (Sec. 6.1), rather than as evidence of a relationship holding uniformly across all cellular configurations.

Overall, the cellular method that most reliably outperforms the baseline, $\mathcal{T}_1^2$, is also the one for which Tab. 4 shows the most consistent within-method link between early diversity persistence and final nonlinearity. For $n < 13$, diversity vanishes earlier, which explains why the cellular approach offers no advantage in smaller search spaces. Earlier studies [9] found that cellular diversity is best preserved with non-disruptive variation operators, and that the grid structure collapses otherwise. Here, the balancedness-preserving operators produce offspring that are semantically similar to their parents, which allows the grid to sustain spatial correlations and thus to enhance diversity.




Fig. 5. Average pairwise Hamming distance (H) between the truth tables of the individuals in the population, across repetitions and generations, for $p = 0.5$. The shaded area is the inter-quartile range.

Taken together, the findings in this analysis indicate that a cellular structure slows the spread of dominant individuals in the early generations. This increases the chances of accumulating several distinct high-quality solutions, which in turn lead, by the end of the run, to functions with higher nonlinearity than those obtained by the baseline. Conversely, the absence of a cellular structure can lead to premature stagnation, because the early spread of dominant individuals means that potentially interesting regions of the search space are never adequately explored. The baseline may therefore reach higher nonlinearity levels than the cellular method during the first part of the run, but, because the early exploration was insufficient, it fails to improve on those levels later on.

To illustrate the spatial dynamics further, we visualize population fitness by means of heat maps. Figure 7 shows one example for $n = 16$ and $p = 0.5$. In the baseline, fitness is distributed at random over the grid, whereas for large radii the population becomes homogeneous. With radius 1, distinct zones of high and low fitness appear early on, and the dominant individuals spread only gradually. This configuration (T_1^2) achieves the highest nonlinearity at the end of the run for $n = 16$.

6.3 Effect of population size and neighborhood scaling

The analysis in Section 6.1 and Section 6.2 was carried out with a population of 100 individuals arranged on a 10×10 toroidal grid. Since the ratio of neighborhood size to population size has emerged above as the quantity that appears to govern the benefit of the cellular structure, we now vary the




Fig. 6. Average pairwise Euclidean distance (E) between the Walsh spectra of the individuals in the population, across repetitions and generations, for $p = 0.5$. The shaded area is the inter-quartile range.

population size while holding the evaluation budget constant. We set the population size to 900 and the number of generations to 111, so that the total number of fitness evaluations (99 900) matches, up to rounding, the 100 000 evaluations of the main campaign, and we arrange the population on a 30×30 toroidal grid, which is square as in the main study. All remaining parameters are unchanged. Figure 8 reports the distribution of the best nonlinearity values obtained in this setting, an analysis analogous to the one in Fig. 1.

Two findings emerge from the figure. First, the nonlinearity values obtained are substantially and consistently worse than those in Fig. 1, for the baseline and for the cellular variants alike. For $n = 16$, for instance, the median of the baseline in Fig. 1 is 32 350, and $\mathcal{T}_1^2$ is consistently better than that under moderate selection pressure, whereas in Fig. 8 the same value is the best nonlinearity attained by any method across all repetitions. This confirms the well-established finding [7, 41, 42] that, for this particular problem and for a fixed total number of fitness evaluations, a small population evolved for many generations is preferable to a large population evolved for few.

Second, most of the cellular variants outperform the baseline for every value of n , which confirms that the cellular algorithm is effective regardless of the population size. One further observation is nonetheless in order. In contrast to what we saw in Fig. 1, the consistently successful cellular variants in Fig. 8 are those with $r = 3$, that is, $\mathcal{T}_3^2$. At first sight this may appear to contradict the main analysis, but it is in fact consistent with our earlier observation about neighborhood size.

Table 4

Pearson correlation (ρ) between the median value of each diversity metric over the first half of the run (generations 1–500) and the final best-so-far nonlinearity, computed across the 50 repetitions with $p = 0.5$.

	n	$\mathcal{T}^0$	$\mathcal{T}_1^2$	$\mathcal{T}_2^2$	$\mathcal{T}_3^2$
I	8	–	0.20	0.32	0.06
	9	–	–	0.15	0.10
	10	–0.20	0.16	0.24	0.22
	11	0.28	0.47*	0.31	0.08
	12	–	0.51*	0.31	0.03
	13	–	0.34	0.26	0.04
	14	–	0.35	0.22	0.12
	15	0.33	0.29	0.40*	0.10
	16	0.20	0.44*	0.32*	0.35*
H	8	–0.30	0.13	–	–
	9	–	–	–	–
	10	–0.24	0.15	–	–
	11	0.37*	0.50*	–	–
	12	0.44*	0.40*	–	–
	13	0.35*	0.36*	–	–
	14	0.37*	0.40*	0.20	–
	15	0.34	0.22	0.33	–
	16	0.29*	0.48*	0.39*	0.34*
E	8	–0.30	0.13	–	–
	9	–	–	–	–
	10	–0.24	0.16	–	–
	11	0.37*	0.51*	–	–
	12	0.45*	0.43*	–	–
	13	0.35	0.35	–	–
	14	0.37*	0.39*	0.20	–
	15	0.36*	0.22	0.33	–
	16	0.30*	0.48*	0.39*	0.34*

Because the population is much larger, the neighborhood size has to be rescaled if the cellular variant is to remain effective at high n . Here the neighborhood size of $\mathcal{T}_3^2$ is 49, and the ratio of this value to the population size (900) is 0.05, which is of the same order as the ratio of 0.09 observed in the main analysis for $\mathcal{T}_1^2$. For smaller n , the baseline attains consistently low nonlinearity values, and even cellular variants with a smaller radius outperform it. Taken together, the two settings indicate that the neighborhood size cannot be held fixed when the population is rescaled: the radius that is most effective with a population of 100 is not the one that is most effective with a population of 900. We deliberately refrain from turning this observation into a general design rule. Two operating points are not enough to identify a functional relationship, and the two ratios we observe span no more than an order of magnitude. We therefore offer the following as a hypothesis to be tested in future work rather than as a finding: for this problem, the benefit of a cellular population appears to be governed by the ratio of neighborhood size to population size rather than by the radius in absolute terms, and to be obtained when that ratio is of the order of 10^{-1} . Establishing the form and the extent of this relationship would require a dedicated campaign sweeping population size, grid dimensionality, and




Fig. 7. Heat maps of individual fitness in a randomly chosen repetition for $n = 16$ and, for the cellular variants, $p = 0.5$.




Fig. 8. Distribution (population size 900, 111 generations), across repetitions, of the best final nonlinearity values. Cliff's Delta values annotate the cellular methods that are better than the non-cellular baseline.

radius jointly, which we identify as the principal direction for future work (Section 7).

7. Conclusion and Future Work

The main goal of this work was to investigate the impact of cellular structures on the behavior of a GA applied to the design of secure Boolean functions. We focused on how locally restricted selection and controlled diffusion of information influence the preservation of diversity and the effectiveness of the search on a highly discrete, multimodal landscape. By constraining all individuals to remain balanced and by using variation operators that preserve this property, we ensured that two key cryptographic requirements, balancedness and nonlinearity, are satisfied simultaneously. We emphasize that the contribution of this work lies in the empirical characterization of the cellular mechanism itself, and not in the introduction of new evolutionary operators or Boolean function representations. Consistently with this scope, the results should be read as evidence about the search and diversity dynamics of spatially structured populations on the nonlinearity landscape, and not as a claim of progress towards the construction of Boolean functions suitable for deployment in real ciphers. For that purpose, algebraic constructions and hybrid algebraic–evolutionary methods remain the state of the art, as documented in Table 3.

Our experimental results indicate that cellular GAs with small neighborhood radii ($r = 1$) and moderate selection pressure ($p = 0.5$) consistently outperform the standard GA as the dimension of the problem grows. This advantage is particularly evident for $n \geq 13$, where cellular configurations slow the propagation of dominant individuals, maintain higher semantic diversity during the early and intermediate stages of the run, and reduce premature convergence. These findings suggest that managing diversity is a critical challenge in Boolean function optimization, and that cellular selection may have an important role to play in addressing it.

The work highlights the relevance of cellular evolutionary dynamics for cryptographic design problems. As a preliminary study, it opens several directions for future research: a deeper analysis of cellular design choices such as neighborhood topology, adaptive radius, and selection pressure; the integration of complementary diversity-preserving mechanisms; and the extension to further cryptographic criteria. A follow-up study could examine how cellular grids affect other Boolean encodings, such as Walsh spectra or GP trees, and could include a principled comparison with GP, evolution strategies, particle swarm optimization [43], and hybrid algebraic–evolutionary approaches. Such a comparison would require the cellular selection and variation operators to be adapted to representations other than truth tables, and was therefore not attempted here. In particular, the relationship between the ratio of neighborhood size to population size and the benefit of the cellular structure, suggested by the two operating points reported in Section 6.3, should be established through a dedicated campaign sweeping population size, grid dimensionality, and neighborhood radius jointly; the present study is not sufficient to determine it.

This study is subject to several limitations. First, the population size, the number of generations, the mutation interval, and the dimensionality of the grid were fixed to values that are standard in the literature. Section 6.3 varies the population size and the generation budget at a constant evaluation budget, but a full sensitivity analysis covering the mutation interval and the grid dimensionality as well remains outside the scope of this study. Second, the nonlinearity values obtained remain below the best-known values reported in the literature (Table 2), which reflects both the untuned hyperparameters and the absence of any exploitation of algebraic structure in the truth-table representation adopted here. Third, the work is restricted to a single grid topology, a two-dimensional toroidal grid, and to a single representation of Boolean functions, namely truth tables. Extending the analysis to other grid topologies and dimensionalities, to other representations such as Walsh spectra or GP trees, and to comparisons with evolution strategies, particle swarm optimization [43], and hybrid algebraic–evolutionary approaches are all important directions for future work.

Acknowledgement

This research is partially supported by the PRIN 2022 PNRR project “Cellular Automata Synthesis for Cryptography Applications (CASCA)” (P2022MPFRT) financed by the European Union – Next Generation EU. This work was supported by national funds through FCT (Fundação para a Ciência e a Tecnologia), under the project - UID/04152/2025 - Centro de Investigação em Gestão de Informação (MagIC)/NOVA IMS - <https://doi.org/10.54499/UID/04152/2025> (2025–01–01 / 2028–12–31) and UID/PRR/04152/2025 <https://doi.org/10.54499/UID/PRR/04152/2025> (2025–01–01 / 2026–06–30). This work was funded by the European Union through the project 101084013 - DIGITAL4Business. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HADEA). Neither the European Union nor the granting authority can be held responsible for them.

Conflicts of Interest

The authors declare no conflicts of interest.

References

- [1] Katz, J. (2026). Introduction to modern cryptography, revised third edition: Student solutions manual. <https://doi.org/10.1201/9781003776994-1>
- [2] Carlet, C. (2010). Boolean functions for cryptography and error-correcting codes. In Y. Crama & P. L. Hammer (Eds.), *Boolean models and methods in mathematics, computer science, and engineering* (pp. 257–397). Cambridge University Press. <https://doi.org/10.1017/CBO9780511780448.011>
- [3] Carlet, C. (Ed.). (2020). *Boolean functions for cryptography and coding theory*. Cambridge University Press. <https://doi.org/10.1017/9781108606806>
- [4] Millan, W., Clark, A., & Dawson, E. (1998). Heuristic design of cryptographically strong balanced boolean functions. *International conference on the theory and applications of cryptographic techniques*, 489–499. <https://doi.org/10.1007/BFb0054148>
- [5] Millan, W., Clark, A., & Dawson, E. (1997). An effective genetic algorithm for finding highly nonlinear boolean functions. *International conference on information and communications security*, 149–158. <https://doi.org/10.1007/BFb0028471>
- [6] Picek, S., Carlet, C., Guilley, S., Miller, J. F., & Jakobovic, D. (2016). Evolutionary algorithms for boolean functions in diverse domains of cryptography. *Evolutionary computation*, 24(4), 667–694. https://doi.org/10.1162/EVCO_a_00190
- [7] Picek, S., Jakobovic, D., Miller, J. F., Batina, L., & Cupic, M. (2016). Cryptographic boolean functions: One output, many design criteria. *Applied Soft Computing*, 40, 635–653. <https://doi.org/10.1016/j.asoc.2015.10.066>
- [8] Bonin, L., Rovito, L., De Lorenzo, A., & Manzoni, L. (2024). Cellular geometric semantic genetic programming. *Genetic Programming and Evolvable Machines*, 25(1), 1–32. <https://doi.org/10.1007/s10710-024-09480-8>
- [9] Rovito, L., Bonin, L., Farinati, D., Vanneschi, L., Manzoni, L., Lorenzo, A. D., & Pietropolli, G. (2025). Exploring the integration of cellular structures in genetic programming-based methods. In B. Xue, L. Manzoni, & I. Bakurov (Eds.), *Genetic programming - 28th european conference, eurogp 2025, held as part of evostar 2025, trieste, italy, april 23-25, 2025, proceedings* (pp. 120–138, Vol. 15609). Springer. https://doi.org/10.1007/978-3-031-89991-1_8
- [10] Manzoni, L., Mariot, L., & Menara, G. (2026). Combinatorial designs and cellular automata: A survey. *Discret. Appl. Math.*, 379, 656–674. <https://doi.org/10.1016/J.DAM.2025.10.014>
- [11] Von Neumann, J., Burks, A. W., & Burks, A. W. (1966). Theory of self-reproducing automata. 1102024. <https://media-cdn.factba.se/pdf/theory-of-self-replicating-automata-john-von-neumann.pdf>
- [12] Zhang, M., Tian, N., Palade, V., Ji, Z., & Wang, Y. (2018). Cellular artificial bee colony algorithm with gaussian distribution. *Information Sciences*, 462, 374–401. <https://doi.org/10.1016/j.ins.2018.06.032>
- [13] Giacobini, M., Tomassini, M., Tettamanzi, A. G., & Alba, E. (2005). Selection intensity in cellular evolutionary algorithms for regular lattices. *IEEE Transactions on Evolutionary Computation*, 9(5), 489–505. <https://doi.org/10.1109/TEVC.2005.850298>
- [14] Pietropolli, G., Nichele, S., & Medvet, E. (2024). The role of the substrate in ca-based evolutionary algorithms. *Proceedings of the Genetic and Evolutionary Computation Conference*, 768–777. <https://doi.org/10.1145/3638529.3654112>
- [15] Salto, C., & Alba, E. (2019). Cellular genetic algorithms: Understanding the behavior of using neighborhoods. *Applied Artificial Intelligence*, 33(10), 863–880. <https://doi.org/10.1080/08839514.2019.1646005>
- [16] Moran, P. A. (1950). Notes on continuous stochastic phenomena. *Biometrika*, 37(1/2), 17–23. <https://doi.org/10.2307/2332142>
- [17] Dimovski, A., & Gligoroski, D. (2003). Generating highly nonlinear boolean functions using a genetic algorithm. *6th International Conference on Telecommunications in Modern Satellite, Cable and Broadcasting Service, 2003. TELSIKS 2003.*, 2, 604–607. <https://doi.org/10.1109/TELSKS.2003.1246297>
- [18] Asthana, R., Verma, N., & Ratan, R. (2014). Generation of boolean functions using genetic algorithm for cryptographic applications. *2014 IEEE International Advance Computing Conference (IACC)*, 1361–1366. <https://doi.org/10.1109/IADCC.2014.6779525>

- [19] Husa, J. (2019). Comparison of genetic programming methods on design of cryptographic boolean functions. In L. Sekanina, T. Hu, N. Lourenço, H. Richter, & P. García-Sánchez (Eds.), *Genetic programming* (pp. 228–244). Springer International Publishing. https://doi.org/10.1007/978-3-030-16670-0_15
- [20] Fuller, J., Dawson, E., & Millan, W. (2003). Evolutionary generation of bent functions for cryptography. *The 2003 Congress on Evolutionary Computation, 2003. CEC '03.*, 3, 1655–1661 Vol.3. <https://doi.org/10.1109/CEC.2003.1299871>
- [21] Dick, G., & Whigham, P. A. (2013). Controlling bloat through parsimonious elitist replacement and spatial structure. *European Conference on Genetic Programming*, 13–24. https://doi.org/10.1007/978-3-642-37207-0_2
- [22] Zheng, Y., & Zhang, X.-M. (1999). Plateaued functions. *International Conference on Information and Communications Security*, 284–300. https://doi.org/10.1007/978-3-540-47942-0_24
- [23] Cohen, G. D., & Litsyn, S. N. (1992). On the covering radius of reed-muller codes. *Discrete Mathematics*, 106, 147–155. [https://doi.org/10.1016/0012-365X\(92\)90542-N](https://doi.org/10.1016/0012-365X(92)90542-N)
- [24] Hou, X.-d. (1997). On the norm and covering radius of the first-order reed-muller codes. *IEEE Transactions on Information Theory*, 43(3), 1025–1027. <https://doi.org/10.1109/18.568715>
- [25] Rothaus, O. S. (1976). On “bent” functions. *Journal of Combinatorial Theory, Series A*, 20(3), 300–305. [https://doi.org/10.1016/0097-3165\(76\)90024-8](https://doi.org/10.1016/0097-3165(76)90024-8)
- [26] Mesnager, S. (2016). *Bent functions - fundamentals and results*. Springer. <https://doi.org/10.1007/978-3-319-32595-8>
- [27] Berlekamp, E. R. (2015). *Algebraic coding theory (revised edition)*. World Scientific. <https://doi.org/10.1142/9407>
- [28] Tarannikov, Y. V. (2000). On resilient boolean functions with maximal possible nonlinearity. In B. K. Roy & E. Okamoto (Eds.), *Progress in cryptology - INDOCRYPT 2000, first international conference in cryptology in india, calcutta, india, december 10-13, 2000, proceedings* (pp. 19–30, Vol. 1977). Springer. https://doi.org/10.1007/3-540-44495-5_3
- [29] Siegenthaler, T. (1984). Correlation-immunity of nonlinear combining functions for cryptographic applications (corresp.) *IEEE Transactions on Information Theory*, 30(5), 776–780. <https://doi.org/10.1109/TIT.1984.1056949>
- [30] Chen, Y., & Lu, P. (2011). Two classes of symmetric boolean functions with optimum algebraic immunity: Construction and analysis. *IEEE transactions on information theory*, 57(4), 2522–2538. <https://doi.org/10.1109/TIT.2011.2111810>
- [31] López-López, I., Sosa-Gómez, G., Segura, C., Oliva, D., & Rojas, O. (2020). Metaheuristics in the optimization of cryptographic boolean functions. *Entropy*, 22(9), 1052. <https://doi.org/10.3390/e22091052>
- [32] Clark, J. A., Jacob, J. L., Maitra, S., & Stănică, P. (2004). Almost boolean functions: The design of boolean functions by spectral inversion. *Computational Intelligence*, 20(3), 450–462. <https://doi.org/10.1111/j.0824-7935.2004.00245.x>
- [33] Djurasevic, M., Jakobovic, D., Mariot, L., & Picek, S. (2023). A survey of metaheuristic algorithms for the design of cryptographic boolean functions. *Cryptography and Communications*, 15(6), 1171–1197. <https://doi.org/10.1007/s12095-023-00662-2>
- [34] Carlet, C., Djurasevic, M., Jakobovic, D., Mariot, L., & Picek, S. (2022). Evolving constructions for balanced, highly nonlinear boolean functions. *Proceedings of the Genetic and Evolutionary Computation Conference*, 1147–1155. <https://doi.org/10.1145/3512290.3528871>
- [35] Koza, J. R. (1994). Genetic programming as a means for programming computers by natural selection. *Statistics and computing*, 4(2), 87–112. <https://doi.org/10.1007/BF00175355>
- [36] Carlet, C., Jakobovic, D., & Picek, S. (2021). Evolutionary algorithms-assisted construction of cryptographic boolean functions. *Proceedings of the Genetic and Evolutionary Computation Conference*, 565–573. <https://doi.org/10.1145/3449639.3459362>
- [37] Picek, S., & Jakobovic, D. (2016). Evolving algebraic constructions for designing bent boolean functions. *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, 781–788. <https://doi.org/10.1145/2908812.2908915>
- [38] Picek, S., Marchiori, E., Batina, L., & Jakobovic, D. (2014). Combining evolutionary computation and algebraic constructions to find cryptography-relevant boolean functions. *Parallel Problem Solving from Nature—PPSN XIII: 13th International Conference, Ljubljana, Slovenia, September 13-17, 2014. Proceedings 13*, 822–831. https://doi.org/10.1007/978-3-319-10762-2_81
- [39] Holland, J. H. (1973). Genetic algorithms and the optimal allocation of trials. *SIAM journal on computing*, 2(2), 88–105. <https://doi.org/10.1137/0202009>
- [40] Behera, P. K., & Gangopadhyay, S. (2022). An improved hybrid genetic algorithm to construct balanced boolean function with optimal cryptographic properties. *Evolutionary Intelligence*, 15(1), 639–653. <https://doi.org/10.1007/s12065-020-00538-x>
- [41] Manzoni, L., Mariot, L., & Tuba, E. (2020). Balanced crossover operators in genetic algorithms. *Swarm and Evolutionary Computation*, 54, 100646. <https://doi.org/10.1016/j.swevo.2020.100646>
- [42] Mariot, L., & Leporati, A. (2015a). A genetic algorithm for evolving plateaued cryptographic boolean functions. *International Conference on Theory and Practice of Natural Computing*, 33–45. https://doi.org/10.1007/978-3-319-26841-5_3
- [43] Mariot, L., & Leporati, A. (2015b). Heuristic search by particle swarm optimization of boolean functions for cryptographic applications. *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation*, 1425–1426. <https://doi.org/10.1145/2739482.2764674>
- [44] Miller, J. F. (1999). An empirical study of the efficiency of learning boolean functions using a cartesian genetic programming approach. *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation - Volume 2*, 1135–1142. <https://dl.acm.org/doi/abs/10.5555/2934046.2934074>
- [45] Hrbacek, R., & Dvorak, V. (2014). Bent function synthesis by means of cartesian genetic programming. In T. Bartz-Beielstein, J. Branke, B. Filipič, & J. Smith (Eds.), *Parallel problem solving from nature – ppsn xiii* (pp. 414–423). Springer International Publishing. https://doi.org/10.1007/978-3-319-10762-2_41

- [46] Picek, S., Jakobovic, D., Miller, J. F., Marchiori, E., & Batina, L. (2015). Evolutionary methods for the construction of cryptographic boolean functions. *European Conference on Genetic Programming*, 192–204. https://doi.org/10.1007/978-3-319-16501-1_16
- [47] Picek, S., Jakobovic, D., & Golub, M. (2013). Evolving cryptographically sound boolean functions. *Proceedings of the 15th Annual Conference Companion on Genetic and Evolutionary Computation*, 191–192. <https://doi.org/10.1145/2464576.2464671>
- [48] Mariot, L., Picek, S., Jakobovic, D., Djurasevic, M., & Leporati, A. (2022). Evolutionary construction of perfectly balanced boolean functions. *2022 IEEE Congress on Evolutionary Computation (CEC)*, 1–8. <https://doi.org/10.1109/CEC55065.2022.9870427>
- [49] Carlet, C., Djurasevic, M., Jakobovic, D., Mariot, L., & Picek, S. (2025). Degree is important: On evolving homogeneous boolean functions. *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 795–798. <https://doi.org/10.1145/3712255.3726779>
- [50] Carlet, C., Durasevic, M., Jakobovic, D., Mariot, L., & Picek, S. (2024). Look into the mirror: Evolving self-dual bent boolean functions. In M. Giacobini, B. Xue, & L. Manzoni (Eds.), *Genetic programming* (pp. 161–175). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-56957-9_10
- [51] Wang, Y., Gao, G., & Yuan, Q. (2022). Searching for cryptographically significant rotation symmetric boolean functions by designing heuristic algorithms. *Security and Communication Networks*, 2022(1), 8188533. <https://doi.org/10.1155/2022/8188533>
- [52] Carlet, C., Djurasevic, M., Gasperov, B., Jakobovic, D., Mariot, L., & Picek, S. (2024). A new angle: On evolving rotation symmetric boolean functions. In S. Smith, J. Correia, & C. Cintrano (Eds.), *Applications of evolutionary computation* (pp. 287–302). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-56852-7_19
- [53] Carlet, C., Đurasević, M., Jakobović, D., Picek, S., & Mariot, L. (2025). A systematic evaluation of evolving highly nonlinear boolean functions in odd sizes. In B. Xue, L. Manzoni, & I. Bakurov (Eds.), *Genetic programming* (pp. 18–34). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-89991-1_2
- [54] Carlet, C., Đurasević, M., Jakobovic, D., Mariot, L., & Picek, S. (2025). The more the merrier: On evolving five-valued spectra boolean functions. In P. García-Sánchez, E. Hart, & S. L. Thomson (Eds.), *Applications of evolutionary computation* (pp. 52–67). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-90062-4_4
- [55] Cussat-Blanc, S., Harrington, K., & Pollack, J. (2015). Gene regulatory network evolution through augmenting topologies. *IEEE Transactions on Evolutionary Computation*, 19(6), 823–837. <https://doi.org/10.1109/TEVC.2015.2396199>
- [56] Martins, T. M., & Neves, R. F. (2020). Applying genetic algorithms with speciation for optimization of grid template pattern detection in financial markets. *Expert Systems with Applications*, 147, 113191. <https://doi.org/10.1016/j.eswa.2020.113191>
- [57] Wickman, R., Poudel, B., Villarreal, T. M., Zhang, X., & Li, W. (2023). Efficient quality-diversity optimization through diverse quality species. *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*, 699–702. <https://doi.org/10.1145/3583133.3590581>
- [58] Xie, H., & Zhang, M. (2012). Impacts of sampling strategies in tournament selection for genetic programming. *Soft computing*, 16, 615–633. <https://doi.org/10.1007/s00500-011-0760-x>
- [59] Moraglio, A., Krawiec, K., & Johnson, C. G. (2012). Geometric semantic genetic programming. In C. A. C. Coello, V. Cutello, K. Deb, S. Forrest, G. Nicosia, & M. Pavone (Eds.), *Parallel problem solving from nature - ppsn xii* (pp. 21–31). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-32937-1_3
- [60] Pawlak, T. P., Wieloch, B., & Krawiec, K. (2015). Review and comparative analysis of geometric semantic crossovers. *Genetic Programming and Evolvable Machines*, 16, 351–386. <https://doi.org/10.1007/s10710-014-9239-8>
- [61] Vanneschi, L., Farinati, D., Rasteiro, D., Rosenfeld, L., Pietropolli, G., & Silva, S. (2025). Exploring non-bloating geometric semantic genetic programming. In *Genetic programming theory and practice xxi* (pp. 237–258). Springer Nature Singapore. https://doi.org/10.1007/978-981-96-0077-9_12
- [62] Mann, H. B., & Whitney, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, 18(1), 50–60. <https://doi.org/10.1214/aoms/1177730491>
- [63] Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, 65–70. <https://www.jstor.org/stable/4615733>
- [64] Garcia, S., & Herrera, F. (2008). An extension on statistical comparisons of classifiers over multiple data sets” for all pairwise comparisons. *Journal of machine learning research*, 9(12). <https://www.jmlr.org/papers/volume9/garcia08a/garcia08a.pdf>
- [65] Derrac, J., García, S., Molina, D., & Herrera, F. (2011). A practical tutorial on the use of nonparametric statistical tests as a methodology for comparing evolutionary and swarm intelligence algorithms. *Swarm and Evolutionary Computation*, 1(1), 3–18. <https://doi.org/10.1016/j.swevo.2011.02.002>
- [66] Meissel, K., & Yao, E. S. (2024). Using cliff’s delta as a non-parametric effect size measure: An accessible web app and r tutorial. *Practical Assessment, Research, and Evaluation*, 29(1). <https://doi.org/10.7275/pare.1977>