

# The Game of Life, Decision and Communication

Roland Mühlenbernd and Simon Schulz

## Abstract

The *game of life* represents a spatial environment of cells that live and die according to fixed *rules of nature*. In the basic variant of the game a cell's behavior can be described as reactive and deterministic since each cell's transition from an actual state to a subsequent state is straight-forwardly defined by the rules. Furthermore it can be shown that the alive cells' spatial *occupation share* of the environment decreases quickly and levels out at a really small value (around 3%), virtually independent of the initial number of alive cells. In this study we will show that this occupation share can be strongly increased if alive cells become active by playing *pre-games* before the rules of nature are accomplished. In such pre-games alive cells can make non-deterministic decisions according to their individual positions and on top of that have the ability to send advices to alive neighbor cells. This result stresses the assumption that individual behavior and local communication supports the optimization of resourcing and constitute important steps in the evolution of creature and man.

## 1 Conway's Game of Life

Conway's game of life [1] is a cellular automaton implemented on a  $m \times n$  grid of cells. A cell can be in two possible states: *dead* or *alive*. Each cell interacts with its eight horizontally, vertically, or diagonally adjacent neighbor cells. At each time step transitions occur, defined by four simple rules.<sup>1</sup> In this article these rules are considered as unalterably written in stone, and therefore called *rules of nature*, defined as follows:

1. *under-population*: any alive cell with fewer than two alive neighbor cells dies
2. *surviving*: any alive cell with two or three alive neighbor cells lives on to the next generation
3. *overcrowding*: any alive cell with more than three alive neighbor cells dies
4. *reproduction*: any dead cell with exactly three alive neighbors becomes an alive cell

By considering a finite grid size, the proportion of alive cells is called the *occupation share* since alive cells are considered as occupied, dead cells are considered as empty. By starting the game of life with a randomly chosen set of alive cells, it can be shown for sufficiently large grid sizes that after a while the occupation share stabilizes on a specific number.<sup>2</sup> This number is generally around 3% of all cells of the grid. Figure 1 shows the course of the number of alive cells over 3000 steps for 15 different simulation runs. Each run is an instance of the game of life applied on a  $70 \times 70$  grid (4900 cells) by starting with a 25% chance of each cell to be initially alive. Thus the expected initial *occupation share* is 25% (1225 alive cells). As a result, in each simulation run the number of live cells strongly decreased and finally leveled out at *occupation share* values between 1.9% (92 alive cells) and 4.1% (199 alive cells), in average 3.2% (157.6 alive cells).

---

<sup>1</sup>By only considering these rules, the evolution of the game is determined by the initial seed. This kind of game is called a *zero player game*.

<sup>2</sup>As we will show, because oscillating arrangements of alive cells may finally emerge, the number of stable cells can pulsate around a particular number.




Figure 1: The number of alive cells over 3000 simulation steps of the game of life for 15 different runs. The number of cells decreases from an occupation share of initially around 25% (1225 alive cells) to finally 3.2% (157.6 alive cells) on average over all runs.

Furthermore, it seems to be independent of the initial occupation share that the final occupation share is around 3%.<sup>3</sup> This fact is supported by further experiments with different initial occupation shares of 12%, 25% and 50%, 15 simulation runs for each setting. For each experiment the number of alive cells decreased during runtime and finally leveled out at an occupation share of around 3% on average.

Thus there is an interesting fact to observe: no matter how high the initial number of alive cells is, the rules of nature of the game of life cause a strong decrease of alive cells down to a very low level of around 3% until it stabilizes. If all cells are considered as resources and alive cells represent usage of resources, then the occupation share depicts a utilization value. As a consequence, with an occupation share of 3% of all possible resources utilization is fairly poor.

In this article we deal with the following question: is it possible to keep the occupation share and therefore the utilization value on a higher level?<sup>4</sup> If the alive cells have the opportunity to make a pre-'rules of nature' decision to sacrifice themselves (decide to die), is it then possible to keep the occupation share of a higher level? On the first view it sounds like a paradox that sacrifice might increase the occupation share, but it possibly restricts the overcrowding effect. In the next section we will show that the final occupation share can be increased by simple fixed pre-'rule of nature' decisions of sacrifice.

## 2 Sacrifice Decisions in Pre-Games

As mentioned in Section 1, an appropriate action that improves the occupation share of the population should not alter the basic game of life rules - they are fixed by nature. Furthermore the creation of new cells is not allowed and can only happen by the reproduction rule of nature. An action that can be added is the deletion of cells before the rules of nature appliance. Occupied resources would be freed. Since we want cells to make such decisions on their own, this could be seen as sacrifice.

<sup>3</sup>For the case, that the initial occupation share is not too low, thus at least above 10%.

<sup>4</sup>Note: under the requirement of complying with the rules of nature, since we consider them as unchangeable.

## 2.1 The Non-Deterministic $n$ -Die Game

For that purpose we integrate a pre-'rule of nature' game that an alive cell can play. In such a *pre-game* the cell can decide if to sacrifice itself and therefore to die, or to stay alive. In the following we introduce a simple variant for such a pre-game, called the *non-deterministic  $n$ -die game*.

**Definition 2.1** (The non-deterministic  $n$ -die game). *Given is the game of life as introduced.  $C_L$  denotes the set of alive cells in the actual round of play, where for all  $c_i \in C_L$  the neighborhood  $N_i$  is defined as follows:  $N_i = \{c_j \in C_L | c_j \text{ is a neighbor of } c_i\}$ . For a number  $n \in \mathbb{N}$ ,  $1 \leq n \leq 8$  the non-deterministic  $n$ -die game is defined by the following 'three phases' algorithm:*

*P1: Initialization*

1. Create a list  $A_L$  and include all alive cells  $c_i \in C_L$  in a random order

*P2: Sacrifice Decision*

2. For all  $j \in \mathbb{N}$ ,  $1 \leq j \leq |A_L|$ :
  - (a) Treat cell  $c_k$  at position  $j$  of list  $A_L$  as follows:
    - (b) If  $|N_k| = n$ :
      - i. Label  $c_k$  as dead
      - ii.  $\forall c_i \in C_L : N_i = N_i \setminus \{c_k\}$

3. Delete all with dead labeled cells

*P3: Rules of Nature*

4. Apply the rules of nature of the game of life

Note that the phases P1 and P2 constitute the pre-game: all cells with  $n$  neighbors sacrifice themselves. Further, this happens in a non-deterministic way: the cells are ordered in a random sequence and each sacrificing cell is also deleted in the neighborhood of all other cells (step 2.(b)ii.). Thus e.g. a cell that had initially more than  $n$  neighbors can probably have  $n$  neighbors when checked in step 2.(b). As a consequence, the fact that the list  $A_L$  is ordered randomly makes the the algorithm non-deterministic. All in all, the *non-deterministic  $n$ -die game* realizes a game of life with a pre-game, where cells are acting in a row and decide to die if they have exactly  $n$  neighbors at their turn of decision.

Note that the algorithm accomplishes a fixed decision rule for cells: *sacrifice yourself if you have exactly  $n$  neighbors*. Thus, to make a decision cells have only access to information of the direct neighborhood. And this is intended since even the rules of nature are exclusively based on neighborhood arrangements. In our opinion, the property of *access restriction to direct neighborhood information* is an important requirement for all following pre-games since this property reflects the spatial character of the rules of nature of the game of life. We denote this requirement as the *local information rule*.

## 2.2 Simulation Experiments & Results

To find out if the non-deterministic  $n$ -die game supports a better occupation share we started experiments for different  $n$  values. First, we ascertained from the basic tests that  $n$ -die games with  $n < 3$  have, if any, a detrimental effect on the game of life since these cells i) would die anyway by the under-population and ii) are highly possible in a shorthanded area

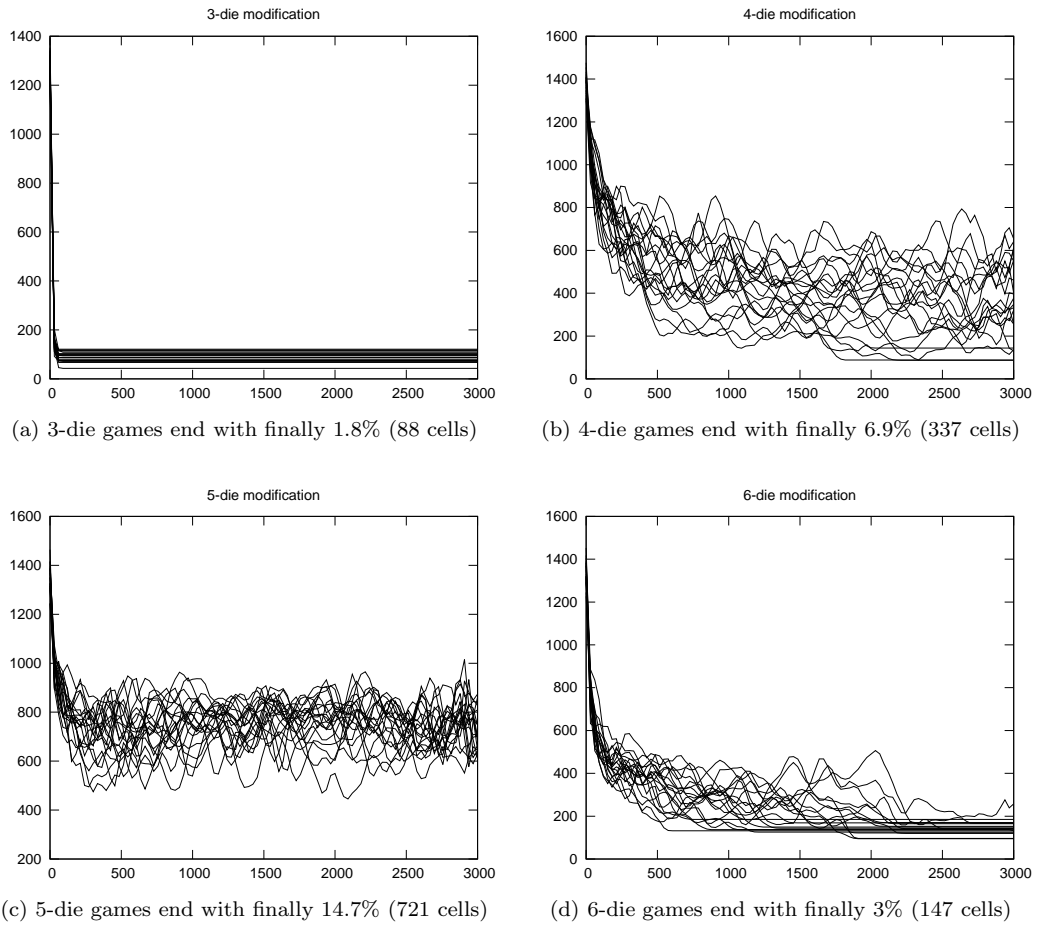


Figure 2: The number of alive cells over 3000 simulation steps for different  $n$ -die games. The number of cells decreases from an occupation share of initially 25% to average values of 1.8% for the 3-die game, 6.9% for the 4-die game, 14.7% for the 5-die game and 3% for the 6-die game.

and therefore probably important to support neighbor cells not to die by under-population. Furthermore  $n$ -die games with  $n > 6$  have (almost) no effect on the game of life, since they i) as good as never emerge during a simulation run and ii) would directly die by overcrowding. As a consequence, we started simulation experiments for  $n$ -die games with  $3 \leq n \leq 6$ : the 3-die game, the 4-die game, the 5-die game and the 6-die game. For each  $n$ -die game we performed 15 simulation runs over 3000 simulation steps. The resulting courses of the number of alive cells over time are depicted in Figure 2.

As shown in Figure 2, the different  $n$ -die games accuse quite different cells' behavior. While the 3-die game's performance of enlarging occupation share is below the performance of the basic game of life, the performance of the 4-die game is slightly better and the performance of the 5-die game is remarkable better. The 6-die game's performance is roughly as good as the basic game of life (as depicted in Figure 1).




Figure 3: Box plots of the final occupation share for the basic game and the 4  $n$ -die games

### 2.3 Explanatory Approach of Results

The simulation results of the four different  $n$ -die games reveal remarkable differences in performance. To highlight these differences and to compare them with the basic game, Figure 3 depicts the box plots of the final occupation shares for the basic game and all 4 experiments. The key to explain this variances is to analyze the way pre-game decisions might interact with the rules of nature of the game of life.

The 3-die game sacrifices cells that have 3 neighbors. These cells would survive because of the rules of nature, thus the sacrifice rule supports a faster dieing in comparison to the basic game of life. The speed of decreasing is not only faster, but also the final occupation share is smaller, with 1.8% almost only half of that of the basic game.

The 4-die game produces three effects: first, it sacrifices cells that would die anyway by overcrowding of the rules of nature, thus there is no acceleration of dieing speed like for the 3-die game. Second, in weakly crowded areas it causes neighbor cells to die by underpopulation, who would survive in the basic game. Thus it indirectly supports an accelerated dieing. Third, it clears crowded areas, thus it also rescues other cells that would otherwise die by overcrowding according to the rules of nature. The third effect seems to be much stronger than the second one, since the final average occupation share is with 6.9% more than twice as high as the final value of the basic game. In addition, the extremely high variation of these values reveals a strong competition of both effects.

The 5-die game also sacrifices cells that would die anyway by overcrowding of the rules of nature. And furthermore even the other effects, already described for the 4-die game, are expected. But since a cell with 5 neighbors is supposed to be in a really crowded area, the third effect described for the 4-die game is even much more helpful in the 5-die game: many cells in the neighborhood are rescued that would normally die by overcrowding of the rules of nature. This effect is remarkably strong, since the average occupation share after 3000

steps is with 14.7% almost 5 times as high as the one of the basic game.

The 6-die game sacrifices cells that would probably die because of overcrowding anyways. But it has to be taken into account, that cells with 6 neighbors are really rare during a simulation run, especially by reconsidering that each cell initially starts with two neighbors on average. Thus the effect on the population is minute and we observe with 3% a similar resulting occupation share as for the basic game of life.

In summary, we showed that a simple pre-game of cells that sacrifice themselves can strongly improve the performance of increasing occupation share. Especially the 5-die game shows a remarkable performance by revealing an occupation share that is with 14.7% almost 5 times as high as the one of the basic game. Nevertheless, we will show that further features like *learning* and *neighborhood communication* can improve the performance even more.

### 3 Neighborhood Situations & Learning

Within the  $n$ -die game cells make their decisions in dependence of their individual circumstances determined by the number of neighbors. But what would happen if they are able to gain more information? What if they can obtain neighbor cells' circumstances and take this additional information into account to make their decision? To answer these questions we applied a more elaborate pre-game that enables cells to request information about circumstances of neighbor cells and make their decision based on their own and a neighbor's situation, called the  $n \times m$ -die game. Furthermore, the cells do not have fixed rules like in the previous section, but learn rules by *learning dynamics*. Thus before we give the definition for the algorithm of the  $n \times m$ -die game, we first introduce the learning dynamics applied in our models, called *reinforcement learning*.

#### 3.1 Reinforcement learning

Reinforcement learning can be captured by a simple model based on urns, also known as *Pólya urns* [3]. An urn models a probabilistic choice in the sense that the probability of making a particular decision is proportional to the number of balls (in the urn) that correspond to that action choice. By adding or removing balls from an urn after each encounter, an agent's (here: cell's) behavior is gradually adjusted. In this work we apply reinforcement learning in a way that the cells of the game of life learn how to behave in the pre-game.

Let's formalize this: first of all we distinguish between different states of cells by means of the number of alive neighbor cells a particular alive cell has. Thus  $T = \{t_1, t_2, t_3, t_4, t_5, t_6, t_7, t_8\}$  is the set of states a cell can be in, where to be in state  $t_i$  means to have  $i$  alive neighbor cells.<sup>5</sup> Furthermore, there are two actions between a cell can choose: to sacrifice itself and therefore to die, or to stay alive. Thus  $A = \{a_{die}, a_{stay}\}$  is the set of possible actions. Finally the cell's basis for a decision also involves a neighbor's state; thus we define a *situation* as a set of state tuples:  $\Gamma = \{\gamma = \langle t_i, t_j \rangle | t_i \in T \text{ is the state of an alive cell } c, t_j \in T \text{ the state of an alive neighbor cell of } c\}$ .

A simple reinforcement learning account  $RL = \{\sigma, \Omega\}$  is defined by *response rule*  $\sigma$  and *update rule*  $\Omega$ . The response rule depicts a probabilistic action choice for a given situation, thus  $\sigma \in (\Gamma \rightarrow \Delta(A))$ . As already mentioned, such a response rule can be modeled by an urn model in the following way: for each situation  $\gamma$  there is an urn  $\mathcal{U}_\gamma$  filled with balls of any type  $a \in A$ . To make a probabilistic action choice in a given situation means to draw a ball from the appropriate urn. Thus the response rule  $\sigma(a|\gamma)$  is defined as follows:

---

<sup>5</sup>Notice that cells with 0 neighbors can be neglected, since we only reconsider situations where cells have at least one neighbor.

**Definition 3.1** (Reinforcement Learning Response Rule). *By defining  $a(\mathcal{U}_\gamma)$  as the number of balls of type  $a$  in urn  $\mathcal{U}_\gamma$  and  $|\mathcal{U}_\gamma|$  as the overall number of balls in this urn, the probability  $\sigma(a|\gamma)$  to make action choice  $a$  in situation  $\gamma$  is given as follows:*

$$\sigma(a|\gamma) = \frac{a(\mathcal{U}_\gamma)}{|\mathcal{U}_\gamma|} \quad (1)$$

Next, the *update rule* of a reinforcement learning account should reinforce successful behavior. But how can we measure that an action is successful for a given situation? In our model we call an action  $a$  successful, if it contributes in a positive way to the occupation share of the whole population. To be more concrete: let's say  $OS_a$  is the occupation share of the next round if action  $a$  is performed,  $OS_{-a}$  is the occupation share of the next round, if not. Now an action  $a$  is considered as successful, if and only if  $OS_a > OS_{-a}$ , i.o.w. if the occupation share of the next round is higher with performing action  $a$  than without performing action  $a$ . Consequently, the update rule  $\Omega$  is informally defined as follows: if action  $a$  is successful in situation  $\gamma$ , then increase the number of balls of type  $a$  in urn  $\mathcal{U}_\gamma$  by one ball. In this way successful behavior reinforces itself since it makes the appropriate choice more probable in subsequent rounds.

### 3.2 The $n \times m$ -Die Learning Game

The idea of the  $n \times m$ -die learning game is similar to the  $n$ -die game in the sense that cells can make a pre-'rules of nature' decision to sacrifice. But there are two crucial differences. First, for the  $n$ -die game the *basis of decision making* is an alive cell's state  $t$ , but for the  $n \times m$ -die learning game it is its situation  $\gamma$  (own state and an alive neighbor's state). Second, for the  $n \times m$ -die learning game the *decision-finding process* is modeled by reinforcement learning, while for the  $n$ -die game the rules are fixed.

Thus by taking the  $n$ -die game as template and incorporating reinforcement learning for decisions grounded on situations, the  $n \times m$ -die learning game can be defined as follows:

**Definition 3.2** (The  $n \times m$ -die learning game). *Given is the game of life as introduced.  $C_L$  denotes the set of alive cells in the actual round of play,  $N_i$  denotes the set of alive neighbor cells of cell  $c_i$ .  $RL = \{\sigma, \Omega\}$  is a reinforcement learning account with urns  $\mathcal{U}_\gamma \forall \gamma \in \Gamma$ . The  $n \times m$ -die learning game is defined by the following 'three phases' algorithm:*

*P1: Initialization*

1. For all  $\gamma \in \Gamma$ : fill urn  $\mathcal{U}_\gamma$  with 50 balls of type  $a_{die}$  and 50 balls of type  $a_{stay}$

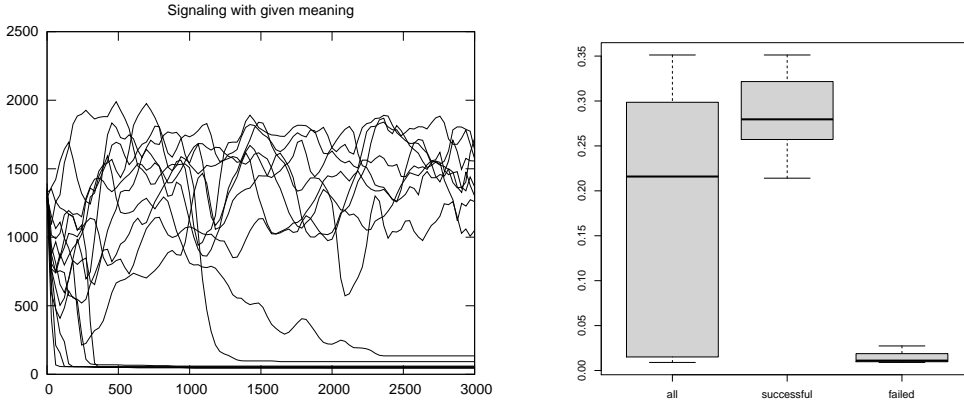
*P2: Sacrifice Decision*

2. For all  $c_i \in C$ :
  - (a) pick randomly a neighbor  $c_j \in N_i$  and request its state  $t_m$
  - (b) play action  $a$  via response rule  $\sigma(a|\langle t_n, t_m \rangle)$ , where  $t_n$  is the state of  $c_i$
  - (c) if  $a = a_{die}$ : label  $c_i$  with dead
  - (d) make an urn update of urn  $\mathcal{U}_{\langle t_n, t_m \rangle}$  via update rule  $\Omega$

3. Delete all with dead labeled cells

*P3: Rules of Nature*

4. Apply the rules of nature of the game of life



(a) 20 runs of the  $n \times m$ -die learning game

(b) Box plots of all, successful and failed runs

Figure 4: The number of alive cells over 3000 simulation steps for the  $n \times m$ -die learning game: the average occupation share over all runs is 17.6%. As evident in Figure 4a there is a clear separation of runs with successful and failed learning of an efficient strategy. For the former group the average occupation share is 28.4% (Figure 4b, middle box plot), for the latter group 1.4% (Figure 4b, right box plot).

In each round an alive cell requests a state of a random alive neighbor cell (step 2(a)) and for the state tuple of own and neighbor's state  $\langle t_n, t_m \rangle$  it chooses an action  $a$  according to the reinforcement learning response rule  $\sigma$  (step 2(b)). In case that action  $a_{die}$  is chosen, the cell gets the label *dead* and will be deleted after the loop (step 3). To complete the reinforcement learning process, the appropriate urn  $\mathcal{U}_{\langle t_n, t_m \rangle}$  will be updated in step 2(d) according to update rule  $\Omega$  of the reinforcement learning account.

### 3.3 Simulation Experiments & Results

To find out if the  $n \times m$ -die learning game supports a better occupation share than the basic game of life and probably outperforms the quite successful non-deterministic 5-die game, we started 20 simulation runs for the  $n \times m$ -die learning game, each over 3000 simulation steps. The resulting courses of the number of alive cells over time are depicted in Figure 4a.

While some of the simulations runs performed poorly and the number of alive cells decreases to 100, most of them performed remarkable well with values oscillating around 1500 alive cells after 3000 steps. This mismatch is caused by the fact that cells accomplish a successful learning strategy in most but not all simulation runs. In other words, the experiments yielded a partition of runs that (a) succeeded in developing a successful strategy (successful runs), and (b) runs that failed (failed runs). And as depicted in Figure 4b the difference is apparently sharp: failed runs have with an average occupation share of 1.4% a worse performance than the worst  $n$ -die game of the experiments in Section 2.1, the 3-die game. On the other hand, successful runs show with an average occupation share of 28.4% an almost twice as good performance as the 5-die game, the by far best  $n$ -die game of the experiments of Section 2.1.

There is a good case to belief that the cells seem to improve performance just by the fact that they have additional information, not only about their own state, but also about the neighbor state. And if they succeed to learn a successful strategy, they outperform the best fixed strategy, that only considered a cell's own state: the 5-die game. This result raises the



question of how such a successful strategy looks like: what kind of strategy must the cells learn to maintain a high occupation share?

As a basic result, all successful strategies that evolved in the simulation runs have particular properties in common: they all involve situations for which cells learn a *definite decision*. This can be illustrated by two sets: a set  $\Gamma_{die} = \{\gamma \in \Gamma \mid \sigma(a|\gamma) = 1.0\}$  of situations for which cells definitely sacrifice, and a set  $\Gamma_{stay} = \{\gamma \in \Gamma \mid \sigma(a|\gamma) = 0.0\}$  of situations for which cells do definitely not sacrifice.<sup>6</sup> In almost all successful simulation runs the cells' strategies contain the same definite decisions, depicted by the following two sets:

- $\Gamma_{die} = \{\langle t_i, t_4 \rangle \mid 1 \leq i \leq 8\}$
- $\Gamma_{stay} = \{\langle t_i, t_j \rangle \mid 1 \leq i \leq 8, 1 \leq j \leq 3, \}$

Note that both sets reveal the fact that definite decisions are only made in dependence of the cell's neighbor's state and completely independent of the cell's own state, since the own state  $t_i$  is defined for all possible numbers of alive neighbors  $1 \leq i \leq 8$  for both sets  $\Gamma_{die}$  and  $\Gamma_{stay}$ . Furthermore, two rules for a successful strategy can be derived from these results. These two rules are called *neighbor treatment rules*, given as follows:

**Definition 3.3** (Neighbor treatment rules). *For the  $n \times m$ -die learning game a successful strategy can be characterized by the following two rules:*

1. *Sacrifice if your neighbor has exactly 4 neighbors.*
2. *Never sacrifice if your neighbor has less than 4 neighbors.*

Further, that there weren't any other states or combinations that describe salient features that *all* successful strategies had in common. Thus, these results elicit two interesting conclusions. First, it is much more important for a cell's decision process to include its neighbor's state than its own: cells learn successful strategies by specific definite decisions independent of the own state, but completely dependent of the neighbor's state. Second, successful strategies that follow the two rules of Definition 3.3 achieve an average occupation share of 28.4% and therefore outperform the strategy of the 5-die game by almost factor 2.

## 4 Neighborhood Communication

In the  $n \times m$ -die game of Section 3 the cells are able to obtain the neighbor's state. They obtain the number of neighbors a neighbor has. But this precondition violates the *local information rule* we postulated at the end of Section 2.1 since cells gain information from observing facts beyond the direct neighborhood.

A possibility to comply with the *local information rule* and also give cells access to the states of direct neighbors is the following idea: instead of that a cell  $c_i$  can observe the state of a neighbor cell  $c_j$ , cell  $c_j$  can communicate its state to  $c_i$ , if  $c_i$  requests the information. Thus to gain information above the own direct neighborhood, cells have to communicate. This can be modeled accurately with a game-theoretic account, called the *signaling game*.

---

<sup>6</sup>Note that  $\Gamma_{rest} = \Gamma - (\Gamma_{die} \cup \Gamma_{stay})$  is the set of situations cells do not learn the same decisive behavior in successful simulation runs or they simply do not learn a decisive behavior at all.

## 4.1 The Signaling Game

A signaling game, first introduced by [2], is a dynamic game  $SG = \langle (S, R), T, M, A, U \rangle$  played between a sender  $S$  and a receiver  $R$ .  $S$  has private information: a state  $t \in T$ . To communicate the state,  $S$  sends a message  $m \in M$  to  $R$ , and  $R$  responds with a choice of action  $a \in A$ . For each round of play, players receive utilities depending on the communicative performance. Further, we will consider a variant of this game where the number of messages  $n = |M|$  is variable. We denote a signaling game with  $n$  messages as  $SG^n$ . In addition, the set of states  $T$  and the set of actions  $A$  is given as already defined in Section 3.1. The utility function is defined by the way the appropriate action improves the occupation share.<sup>7</sup> Finally, the cells have to learn how to assign messages to states and actions, thus we combine it with a reinforcement learning account: there are urns  $\mathcal{U}_t$  for states and urns  $\mathcal{U}_m$  for messages; and different response rules: the *sender response rule*  $\sigma(m|t) = \frac{a(\mathcal{U}_t)}{|\mathcal{U}_t|}$  and the *receiver response rule*  $\rho(a|m) = \frac{a(\mathcal{U}_m)}{|\mathcal{U}_m|}$  (see e.g. [4] for more a more detailed description).

## 4.2 The $n$ -Messages Signaling Game

As initially remarked, the innovation of the new game is the fact that a cell  $c_i$  cannot observe the state  $t_k$  of a neighbor cell  $c_j$ . Instead of that a cell  $c_i$  can request this information and  $c_j$  has to communicate that it is in state  $t_k$  in terms of sending a message  $m \in M$  that cell  $c_i$  has to construe. Furthermore, the decisions i) which message to send for a given state, and ii) how to construe a received message, is not initially given, but has to be learned by reinforcement learning. Thus the  *$n$ -messages signaling game* conforms to the  $n \times m$ -die learning game and extends it by communication via signaling games. It is defined as follows:

**Definition 4.1** (The  $n$ -messages signaling game). *Given the game of life as introduced.  $C_L$  denotes the set of alive cells in the actual round of play,  $N_i$  denotes the set of neighbor cells of a cell  $c_i$ .  $SG^n = \{(S, R), T, M, A, U\}$  is a signaling game as already introduced,  $RL = \{\sigma, \Omega\}$  is an reinforcement learning account with urns  $\mathcal{U}_t \forall t \in T$  and urns  $\mathcal{U}_m \forall m \in M$ . The  $n$ -messages signaling game is defined by the following 'three phases' algorithm:*

*P1: Initialization*

1.  $\forall t \in T \forall m \in M$ : fill urn  $\mathcal{U}_t$  with  $\frac{100}{|M|}$  balls of type  $m$
2.  $\forall m \in M$ : fill urn  $\mathcal{U}_m$  with 50 balls of type  $a_{die}$  and 50 balls of type  $a_{stay}$

*P2: Sacrifice Decision*

3. For all  $c_i \in C$ :
  - (a) pick randomly a neighbor  $c_j \in N_i$  and make a state request for its state  $t$
  - (b)  $c_j$  sends a message  $m \in M$  via response rule  $\sigma(m|t)$
  - (c)  $c_i$  plays action  $a \in A$  via response rule  $\rho(a|m)$
  - (d) if  $a = a_{die}$ : label  $c_i$  with dead
  - (e) make an urn update of urns  $\mathcal{U}_t$  and  $\mathcal{U}_m$  via update rule  $\Omega$

4. Delete all with dead labeled cells

*P3: Rules on Nature*

5. Apply the rules of nature of the game of life

---

<sup>7</sup>Given  $t$ :  $U(t, a) = \begin{cases} 1 & \text{if } OS_a > OS_{-a} \\ 0 & \text{else} \end{cases}$  with  $OS_a$  and  $OS_{-a}$  as introduced in Section 3.1.

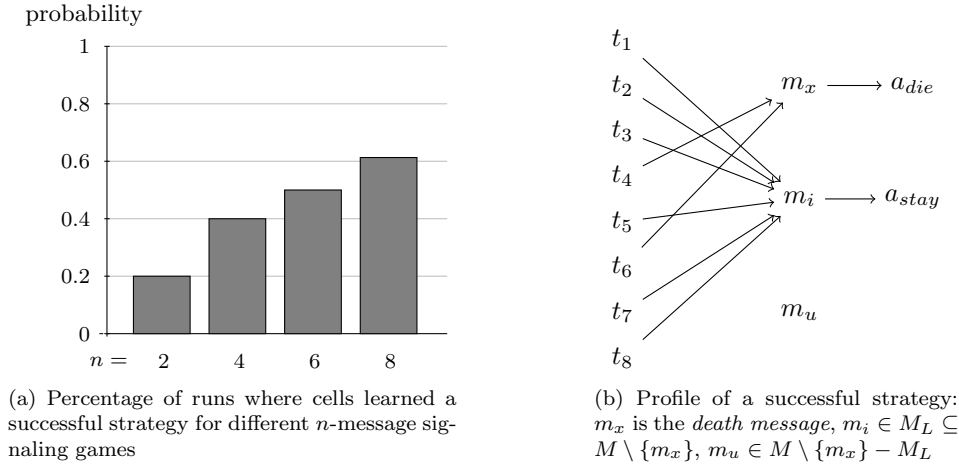


Figure 5: The number of messages  $n$  of the game supportively influences the probability of learning a successful strategy, as depicted in Figure 5a. The strategy that succeeds in the  $n$ -message signaling games needs at least two messages to compare between two actions. The death message  $m_x$  is generally sent in states  $s_4$  and  $s_6$ , as depicted in Figure 5b.

Notice that the game conforms to the  $n \times m$ -die learning game in almost all points, but instead of observing the neighbor's state, it is communicated via a signaling game (steps 2(a) - 2(c)), and communicative behavior is updated by reinforcement learning (step 2(e)).

### 4.3 Simulation Experiments & Results

To find out if the  $n$ -messages signaling game can at least uphold the performance of the  $n \times m$ -die learning game, we started experiments for different  $n$ -values for the  $n$ -messages signaling game: the 2-messages, the 4-messages, the 6-messages and the 8-messages signaling game, for each 20 simulation runs over 3000 simulation steps.

As a basic result, all different  $n$ -messages signaling games performed roughly as well as the  $n \times m$ -die learning game: if cells learned a successful communication strategy, the occupation share was around 25% on average after 3000 steps; if not, the occupation share was below 3% on average. Additionally, this results show that the number of messages  $n$  did not influence the performance of successful runs.

But the number of messages  $n$  influenced the percentage of simulation runs for which a successful strategy evolved. And that in the following way: the higher  $n$  and therefore the more messages are at disposal for communication, the more simulation runs were successful. This result is a good index for assuming that the probability that successful communication emerges increases with the number of messages  $n$ . The percentage of successful runs in dependence of  $n$  is depicted in Figure 5a.

In a next step we analyzed what kind of communication strategy turned out to lead to successful communication. As shown in the last section, the two *neighbor treatment rules* said 1. *sacrifice if your neighbor has 4 neighbors* and 2. *don't sacrifice if your neighbor has less than 4 neighbors*. And this rules are incorporated in almost all successful communication strategies. An exemplary successful strategy that evolved (at least in a quite similar fashion) in all successful runs is depicted in Figure 5b. The first rule is accomplished by the fact that cells learned exactly one message  $m_x$ , called the *death message*, to communicate that the sender cell is in state  $t_4$  and the receiver cell learned to construe it with  $a_{die}$ , thus to

sacrifice. Furthermore, this message was also quite often used to communicate  $t_6$ .<sup>8</sup>

All the other messages were either a) used to communicate that the sender cell is in one of the other states and the receiver cell learned to construe it with  $a_{stay}$  or b) not used at all. Thus in the former case all these messages  $m_i$  are member of the *live-on set*  $M_L \subseteq M \setminus \{m_x\}$ , that moreover accomplishes the second *neighbor treatment rule*. If  $M \setminus \{m_x\} - M_L \neq \emptyset$ , then there were also remaining messages not used for any state, called unused messages  $m_u \in M \setminus \{m_x\} - M_L$ .

## 5 Conclusion

As a basic result, we were able to show that the game of life's performance, in terms of occupation share of alive cells, can be strongly improved by integrating *pre-games* that give alive cells the opportunity to sacrifice before the *rules of nature* of the game of life are applied. In Section 2 and 3 we were able to show that such *pre-'rules of nature' decisions* were especially successful (in comparison with the basic game of life's occupation share of around 3%) by complying with the following strategies:

- *5-die rule*: Sacrifice, if you have exactly 5 neighbors (occupation share: 14.7%)
- *Neighbor treatment rules*: Sacrifice, if your neighbor has exactly 4 neighbors; and don't sacrifice, if your neighbor has less than 4 neighbors (occupation share: 28.4%)

Further, we argued that the state of the neighbor cannot be directly observed without violating the *local information rule* of direct neighbor access. Thus in Section 4 we integrated *signaling games* to obtain the neighbor's state by communication. We were able to show that i) two messages were sufficient to learn a successful communication strategy, but ii) the more messages are provided, the higher the probability that cells learn a successful communication strategy.

All in all, we were able to show that if cells have the possibility to make decisions according to their individual state and especially according to a neighbor's state, they can strongly increase the occupation share. Furthermore, by integrating communication via signaling games, we were able to detect specific strategy patterns that ensure successful communication. Thus, it would be quite interesting to analyze *how* such characteristics of strategic communication depend on the game of life's rules of nature. Experiments with altered rules of nature can possibly reveal such dependencies and are therefore interesting hints for the impact of environmental features on the evolution of communication.

## References

- [1] M. Gardner. Mathematical games: The fantastic combinations of John Conway's new solitaire game "life". *Scientific American*, 223(4):120–123, 1970.
- [2] D. Lewis. *Convention. A philosophical study*. Harvard University Press, 1969.
- [3] A. E. Roth, I. Erev. Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behaviour*, 8:164–212, 1995.
- [4] B. Skyrms. *Signals: Evolution, Learning & Information*. Oxford: Oxford University Press (2010)

---

<sup>8</sup>The reason for this way of usage remains to be analyzed and goes beyond the scope of this article.