

A Step Towards a Reinforcement Learning De Novo Genome Assembler

Kleber Padovani (✉ kleber.padovani@gmail.com)

University of the State of Amazonas

Roberto Xavier

Vale Technology Institute

André Carvalho

University of São Paulo

Anna Reali

University of São Paulo

Annie Chateau

Montpellier Laboratory of Informatics, Robotics and Microelectronics

Ronnie Alves

Vale Technology Institute

Research Article

Keywords: Genome assembly, Reinforcement Learning, De Novo Genome Assembler

Posted Date: July 29th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-753559/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

A step towards a reinforcement learning *de novo* genome assembler

Kleber Padovani^{1,2,*}, Roberto Xavier³, André Carvalho⁴, Anna Real⁵, Annie Chateau⁶, and Ronnie Alves^{1,3}

¹Federal University of Pará, Computer Science Graduate Program, Belém-PA, 66.075-110, Brazil

²University of the State of Amazonas, Itacoatiara-AM, 69.101-416, Brazil

³Vale Technology Institute, Sustainable Development, Belém-PA, 66.055-090, Brazil

⁴Institute of Mathematics and Computer Sciences, University of São Paulo, São Carlos-SP, 13566-590, Brazil

⁵Polytechnic School, University of São Paulo, São Paulo-SP, 05508-010, Brazil

⁶LIRMM, Univ Montpellier, CNRS, Montpellier, France

*kleber.padovani@gmail.com

ABSTRACT

Genome assembly is one of the most relevant and computationally complex tasks in genomics projects. It aims to reconstruct a genome through the analysis of several small textual fragments of such genome — named *reads*. Ideally, besides ignoring any errors contained in *reads*, the reconstructed genome should also optimally combine these *reads*, thus reaching the original genome. The quality of the genome assembly is relevant because the more reliable the genomes, the more accurate the understanding of the characteristics and functions of living beings, and it allows generating many positive impacts on society, including the prevention and treatment of diseases. The assembly becomes even more complex (and it is termed *de novo* in this case) when the assembler software is not supplied with a similar genome to be used as a reference. Current assemblers have predominantly used heuristic strategies on computational graphs. Despite being widely used in genomics projects, there is still no irrefutably best assembler for any genome, and the proper choice of these assemblers and their configurations depends on Bioinformatics experts. The use of reinforcement learning has proven to be very promising for solving complex activities without human supervision during their learning process. However, their successful applications are predominantly focused on fictional and entertainment problems - such as games. Based on the above, this work aims to shed light on the application of reinforcement learning to solve this relevant real-world problem, the genome assembly. By expanding the only approach found in the literature that addresses this problem, we carefully explored the aspects of intelligent agent learning, performed by the Q-learning algorithm, to understand its suitability to be applied in scenarios whose characteristics are more similar to those faced by real genome projects. The improvements proposed here include changing the previously proposed reward system and including state space exploration optimization strategies based on dynamic pruning and mutual collaboration with evolutionary computing. These investigations were tried on 23 new environments with larger inputs than those used previously. All these environments are freely available on the internet for the evolution of this research by the scientific community. The results suggest consistent performance progress using the proposed improvements, however, they also demonstrate the limitations of them, especially related to the high dimensionality of state and action spaces. We also present, later, the paths that can be traced to tackle genome assembly efficiently in real scenarios considering recent, successfully reinforcement learning applications — including deep reinforcement learning — from other domains dealing with high-dimensional inputs.

Introduction

Genome assembly is one of the most complex and computationally exhaustive tasks in genomics projects¹. Additionally, it is among the most important tasks which allow genome sequences analyses. The main goal of genome assembly is to reconstruct a genome using an assembler software, that does it by the analysis of several small fragments coming from a genome — commonly referred to as *target genome*². These fragments, named *reads*, are obtained from an equipment, called the DNA sequencer.

A high-performance assembler is something highly desired among researchers, as it will imply more accurate genomes, allowing researchers to reach a better understanding of the traits and functions of living organisms³. As a result, the knowledge acquired from these whole genomes produces positive impacts in several fields, such as medicine, biotechnology and biological sciences.

A DNA sequencer is a machine responsible for the initial, but fractional, reading of the genetic code of living organisms⁴. However, the genomes of most organisms — even microorganisms — are too long for being read in a single run in the

sequencer⁵. To surpass this limitation, a technique called Shotgun is applied, consisting of cutting up the genome into small pieces and producing small fragments of DNA which can be entirely (or mostly) read by the sequencer, representing the corresponding genetic information in text sequences (the *reads*)⁶.

Since DNA molecules are formed by sequential pairs of complementary nucleotides (each of them composed by Adenine-Thymine or Guanine-Cytosine), *reads* represent only a single nucleotide from each pair, sequentially written as a character in the text. In biological terms, we can say that nucleotides of only one (out of two) strand of each DNA fragment is read. The reading of each nucleotide is represented by its corresponding initial (A, C, G or T). Thus, the number of characters in each *read* is commonly referred to as base pairs (or *bp*).

The genome of an organism is the sequence of all nucleotides from its DNA molecules, represented by letters (A, C, T, G). Each nucleotide isolated does not represent any relevant biological information, however, when we put them all together, the corresponding sequences provide a deep knowledge about this organism. Within the organism's genome, for example, there are (among other information) the species genes. Genes are continuous fragments of the genome whose nucleotide sequences define species traits and behaviors (e.g. the human eye color)⁷. A single *read*, however, generally cannot represent the complete information from even a gene, thus genome assembly is commonly required to obtain the whole genome⁵.

Genome sequencing technology defines the maximum number of base pairs will be read from each DNA fragment⁵, which, in turn, defines the size of the produced *reads* and directly affects the number of *reads* produced during the sequencing process. As genome assembly is a computational task aiming to order *reads* in an attempt to reconstruct the original DNA sequence, the number of *reads* and their sizes directly impact the complexity of the assembly process — the more and larger the *reads*, the higher the complexity for assembling them.

From a computational perspective, given a set, $R = read_1, read_2, \dots, read_n$, with n small text sequences, called *reads*, the assembly problem is defined as the search for the superstring that originated the n reads of R by analysing the existing overlaps between them.

Genome assembly is usually carried out by assemblers adopting *de novo* strategy and/or the comparative strategy. The comparative approach is relatively simpler and computationally treatable, however it requires a previously assembled genome as reference (e.g. the genome of a similar species) to guide the assembly process by comparing the produced *reads* with the reference genome, meanwhile the *de novo* approach presents no such dependence⁸.

De novo strategy is particularly important given that only a small number of reference genomes are currently available, compared to the number of existing and non-sequenced genomes - it is estimated that the vast majority of microorganisms' genomes are still unknown⁹. However, although *de novo* assembly approach allows the assembly of new genomes without requiring a reference genome, it is considered a highly complex combinatorial problem, falling into the theoretically intractable class of computational problems, referred to as NP-hard¹⁰.

In computer science, the commonly applied strategies for *de novo* genome assembly process are based on heuristics and graphs, and they are known as *Greedy*, *Overlap-Layout-Consensus (OLC)*, and *De Bruijn graph*¹¹. For instance, in OLC strategy, each *read* is represented as a node in a graph (named overlap graph) and edges represent the overlap between *reads*. Thus, the reconstructed genome corresponds to the *reads* along the path traversing all the nodes. This algorithm corresponds to the Hamiltonian path algorithm. Another computational formulation for genome assembly is to find the shortest common superstring (SCS) formed by the *reads* — which can also be polynomially reduced to the Travelling Salesman Problem (TSP)¹².

Regardless the difficulties and limitations, current *de novo* assemblers are capable of producing acceptable solutions. However, the use and application of *de novo*-based assemblers normally require specific bioinformatics knowledge in order to correctly set configurations and parameters for the assembler. Nevertheless, optimal results are not always guaranteed — given the high complexity enclosed to this theme¹³.

Despite the great contributions of the current assemblers for developing scientific discoveries on the genomics analyses of organisms, genome assembly is not yet a fully solved problem. So, it is very important to continue the development of new and more robust assemblers, in order to assemble DNA sequences faster and with improved accuracy¹⁴. This challenge has been the aim of numerous currently ongoing researches, which apply computational techniques to genomics in the search for better solutions, including the use of machine learning (ML)¹.

Although machine learning is an alternative approach to heuristics for dealing with high complex computational problems (as it is the case of NP-hard problems), few approaches apply machine learning for dealing with the assembly problem. According to the literature review presented in Souza et al.¹, the few scientific investigations evaluating the application of machine learning techniques for genome assembly were recently published, and only one study dated to the end of the previous century while 12 others date from later. For comparison, in a similar review¹⁵, Henrique et al. reported 140 publications applying machine learning techniques into the problem of financial credit risk.

With the recent access to computational advances — including increased processing and storage power of computers, the investigation of machine learning application for complex and high-scale computational problems has started to increase in the scientific community and some good results have been reported¹⁶. This larger resources availability has also allowed the return

of reinforcement learning application for these problems¹⁷.

Reinforcement learning (RL), alongside supervised learning and unsupervised learning, is a basic machine learning paradigm that works with intelligent agents which take actions in an environment that represents a task. Ideally, this agent is supposed to solve this task if it is able to learn how to maximize the rewards received from its actions taken¹⁸.

Although scarcely applied in machine learning development, RL has shown some surprisingly positive results, especially for games¹⁷. Sutton et. al, by the way, recently suggested in 19 that RL, unlike supervised and unsupervised learning, may be the path to the development of general artificial intelligence. Two great (and commonly cited) examples of the application of reinforcement learning are the training of agents for playing the classic board game Go, as well as several Atari games, showing better results when compared to those from previous approaches and showing superior performances even when compared to trained humans²⁰⁻²².

However, RL successful applications are predominant in problems that rely on accurate environment simulators, such as games, where the rules and environments are known and allow the development of simulations for intensive training of intelligent agents²³. Despite the importance of games to both society and computer science, there is a growing expectation, and some initial efforts have been made, for extending the success of RL in games into real world problems — for which it is generally impossible or impracticable to create the required simulators that would provide appropriate training for the agents²⁴.

Such low use of RL in real world problems is also observed in the specific scenario of the genome assembly, where only a single study can be found, proposed by Bocicor et al.^{1,14}. This approach, which is also the object of this study and will be hence forward referred to as seminal approach, proposes the use of an episodic trained agent (i.e. whose training has been divided into episodes) applying the Q-learning reinforcement learning algorithm for learning the correct order of a set of *reads* and, consequently, for reaching the corresponding genome.

One of the most valued characteristics of the reinforcement learning is that it allows the agent's autonomy for learning. For example, in supervised learning, a great deal of intervention is required (usually by human specialists) during the learning process, given that all information provided for the machine is previously and properly labeled. In reinforcement learning, the agent learns through the consequences of its successive failures and successes.

This is particularly useful for solving tasks whose solutions extrapolate the human knowledge and capacity, such as genome assembly. The assembly complexity starts from the assembler's choice — as assemblers using similar strategies may produce different results^{25,26} — and extends until the assembler's configuration following the user's decisions. Obtaining intelligent and trained agents by reinforcement learning is important in this scenario as it could eliminate the need for human specialists.

Another relevant aspect for the application of reinforcement learning is the capacity of agents to deal with large volumes of data and extract new rules associated with the main task, which were not explicitly provided before and, in some cases, were not identified by humans. In games, as above mentioned^{20,21}, agents were able to play new games on their own - without human supervision - and in some cases, they were able to outperform the best known human players.

Considering that Q-learning algorithm requires a Markov decision process definition with established parameters of states and actions, together with a reward system to be achieved by the agent at each action in every state¹⁸, the problem was modeled by the authors¹⁴ through a space of states capable of representing all possible *read* permutations, so that only one initial state there existed and, in each state, there is one action to be taken by the agent for each *read* in the pool to be assembled.

Following these definitions, from the graph theory perspective, the proposed states space for n *reads* can be visualized as a complete n -ary tree, with height equal to n , as the set of states presents only one initial state and forms a connected and acyclic graph²⁷. Thus, we can reach the number of existing states in the states space, represented by Equation 1.

$$\text{number of states} = \frac{n^{n+1} - 1}{n - 1} \quad (1)$$

The proposed reward system depends, first, on the type of the state reached after taking each action, which can be an absorbing or a non-absorbing state. An absorbing state, once reached, does not allow any other state to be reached^{18,28}. The authors have defined that each state requiring n actions to be reached (being n the number of *reads* to be assembled) is an absorbing state.

A small and constant reward (e.g. 0.1) was assigned as reward for actions reaching non-absorbing states. This reward was also set to every action leading to absorbing states where repeated *reads* were used to achieve them. Finally, actions leading to other absorbing states (the final states) produce a reward corresponding to the sum of overlaps between all pairs of consecutive *reads* used to reach these states.

For your better understanding, Figure 1 presents a simple example of a space of states for a set of only 2 *reads*, identified as A and B. In this example, we can observe the existence of a single initial state, two actions associated with non-absorbing states and four absorbing states (highlighted in black), achieved after taking two any actions.

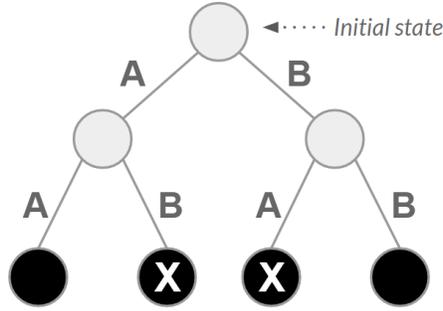


Figure 1. Example of state space for a set of two *reads*, here referred to by A and B.

In this figure two of the absorbing states are highlighted by the letter X. These states, unlike the previous ones, are final states, as they are the only ones in the space of states reached directly from the initial state without repeated actions — one is achieved from actions referred to as *read A* and *read B*, respectively, and the other from the actions of *read B* and, then, *read A*.

The Smith-Waterman algorithm (SW) was applied for obtaining the overlaps between pairs of reads, which were added for obtaining the rewards of actions that led to final states²⁹. The sum of overlaps when reaching a final state s , here referred to as Performance Measure (PM), is described in Equation 2, where $read_s$ corresponds to the sequence of reads associated with the actions taken for achieving s .

$$PM(s) = \sum_{i=1}^{n-1} sw(read_s[i], read_s[i+1]) \quad (2)$$

By using these definitions, the seminal approach produced attractive results, however, it has been evaluated by the authors against only two small sets of *reads*, one with 4 *reads* with less than 10bp and the second with 10 *reads* of 8bp each. These *reads* were obtained by simulating the sequencing process, assuming as the target genome only a small fragment (a microgenome) of the real genome of the bacterium *Escherichia coli*.

In order to perform a scalability analysis of the seminal approach, Xavier et al.³⁰ evaluated the performance of this approach against 18 datasets. These 18 datasets were produced following the same simulation methods. The first dataset corresponds to the set of 10 *reads* with 8bp from the seminal approach, originated from a 25bp microgenome. From this same microgenome and from a 50bp microgenome, a total of 17 new datasets were generated (8 from the minor microgenome and 9 from the major one) each containing 10, 20 or 30 *reads*, with 8bp, 10bp or 15bp.

Almost all definitions made by Bocicor et al. were replicated by Xavier et al. in this analysis, but they experimentally set α and β to 0.8 and 0.9, respectively, and the space of actions was slightly reduced, so that actions associated with *reads* that has already been taken previously were removed from the available actions. In the states space depicted in Figure 1, for example, the leftmost and rightmost leaves (i.e. absorbing states) would be removed after this change. Although this change reduces the number of states, the size of the space of states continues to grow exponentially, as we can observe in Equation 3.

$$\text{number of states} = \sum_{i=0}^n \frac{n!}{(n-i)!} \quad (3)$$

This study confirmed the previously positive results found in the seminal approach with the first dataset, however, as the dataset size increases the performance of the seminal approach decreases considerably, reaching the target microgenome in only 2 out of 17 major datasets. According to the authors, such bad results may be related to the high cost required by the agent to explore such a vast states space — as seen in Equation 1, this space grows exponentially — and also from the failures in the proposed reward system.

In order to continue the investigation of the application of reinforcement learning in the genome assembly problem and targeting the current challenge of applying the reinforcement learning into real-world problems²⁴, here we propose to analyze the limits of reinforcement learning into a real-world problem, which is also a key problem for the development of science. This analysis was carried out by evaluating the performance of strategies that are complementary to those previously studied, and that could be incorporated into the seminal approach for obtaining improved genome assembly.

Methods

In this study, 7 experiments were evaluated against the seminal approach — referred to as approaches 1.1, 1.2, 1.3, 1.4, 2, 3.1, and 3.2 — and their methodologies are described in the next 3 subsections. In each approach, as in the seminal approach, the main goal is to reach an agent trained by reinforcement learning capable of identifying the correct order of *reads* from a sequenced genome. Figure 2 illustrates this proposal, where the set of *reads* from the sequencing process is initially represented as the agent’s interaction environment; the agent learns the order of *reads* by observing the environment’s current state and the rewards produced from the successive actions taken — until it (ideally) reaches the correct order for the analyzed *reads*.

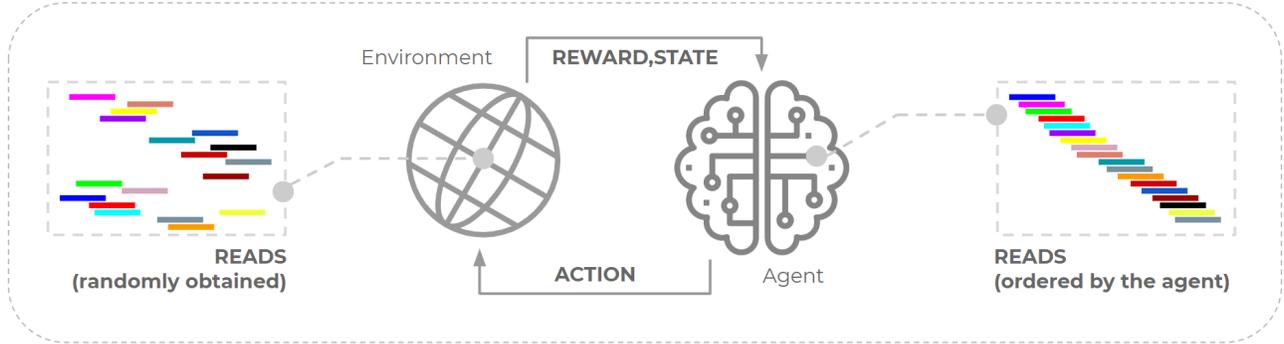


Figure 2. Illustration that demonstrates the application of reinforcement learning in the genome assembly problem. The set of *reads*, which are obtained in random order by the sequencer, is represented computationally by a reinforcement learning environment. Through successive interactions with the environment, caused by taking actions, the agent ideally learns the correct order of *reads* — thus allowing the target genome to be reached.

These approaches will consider the findings of the scalability analysis from the previously mentioned work of Xavier et al. Efforts were then made for improving the reward system adopted in the seminal approach — especially in the approaches described in the next subsection - and to optimize the agent’s exploration — in the approaches described in the last two subsections.

Approaches 1: Tackling sparse rewards

In approaches 1, the reward system, originally defined by Equation 3, has been improved for preventing that permutations of *reads* that are inconsistent, in terms of alignment, produce high reward values for the agent. This is an undesirable behavior as the agent’s learning process is based on maximizing the accumulated rewards. Therefore, ideally, high rewards are expected to be associated with high quality responses.

$$r(s, a, s') = \begin{cases} PM(s') & \text{if } s' \text{ is a final state,} \\ 0.1 & \text{otherwise} \end{cases} \quad (4)$$

This observed inconsistency stem from the *Smith-Waterman* (SW) algorithm, chosen for calculating the local alignment between two given sequences. With this algorithm, a numeric score is calculated to represent the major alignment size (even if partial) between two sequences. However, the SW algorithm has no constraint for the order between sequences. Therefore, it does not previously consider that a sequence (in our case, a *read*) must be aligned either left or right to the other one.

In the case of the genome assembly problem, where the alignment between subsequent *reads* is expected to follow an order, the overlap score obtained from the SW algorithm may induce the agent to find *reads* permutations with high overlap values in pairs, however, without presenting a consecutive alignment (suffix-prefix) between *reads* in the set.

In Figure 3 we can observe this type of inconsistency through an example, which presents two different permutations for a set of *reads*, identified by the letters ranging from A to J, obtained from a given microgenome. The first permutation is formed, in this order, by the *reads* A, B, C, D, E, F, G, H, I, and J, and according to SW the estimated overlap is 40.34. This permutation corresponds to the optimal permutation, because the union of subsequent *reads* through the overlap between the suffix of the previous read and the prefix of the posterior read produces the target microgenome (provided at the bottom of the figure). On the other hand, the second permutation, formed by the *reads* H, G, F, E, C, B, D, A, J, and I, respectively, is not optimal; however, for this permutation the SW presents an overlap of 43.02, which is greater than that obtained for the optimal permutation.

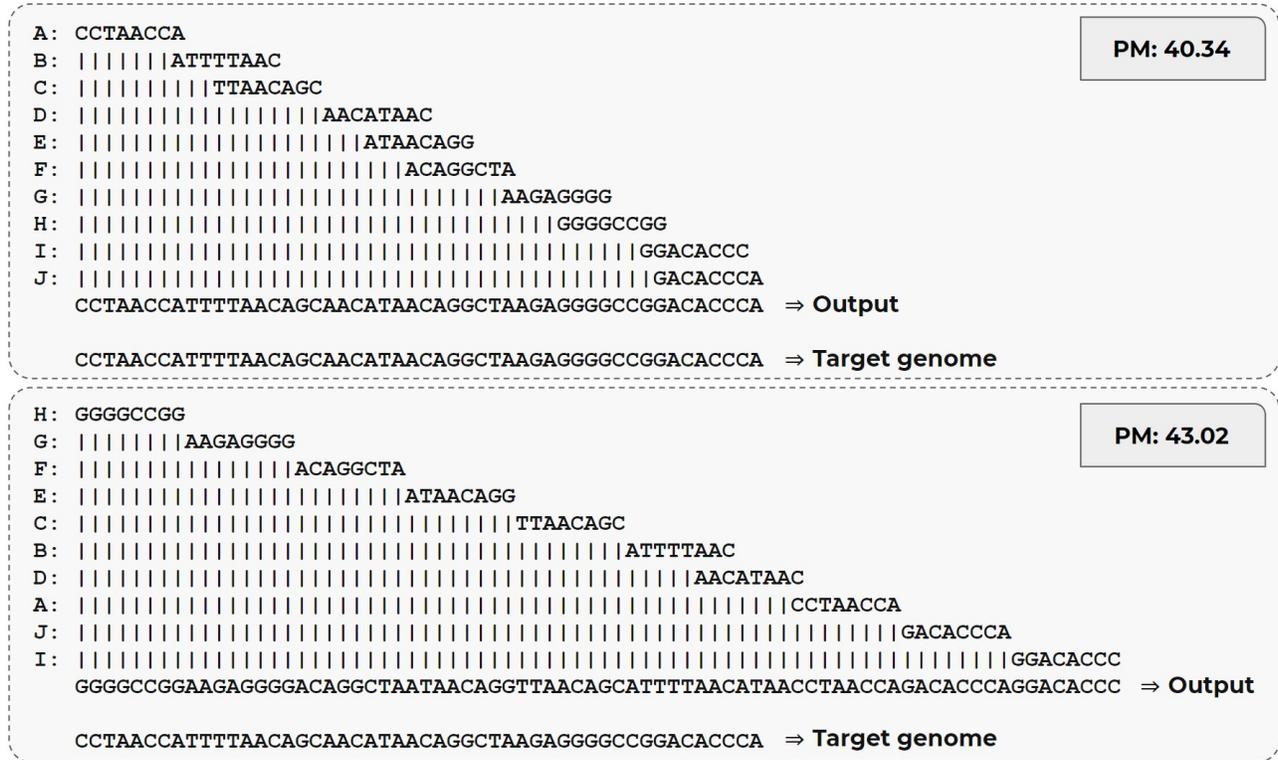


Figure 3. Illustration showing that the measure used as a reward to train agents does not produce maximum values for optimal outputs in some cases. Above, we have an optimal permutation of *reads*, for which the PM is 40.34 and whose corresponding genome is equal to the target genome itself; and, below, we have another permutation whose output differs from the target genome, but the corresponding PM is greater than the PM of the optimal permutation.

Another aspect to be considered in the reward system of the seminal approach is that the agent is able to receive high rewards only when it takes actions leading to sparse final states in the states space. That is, considering that the applied training is episodic, in each training episode, the agent receives a non-constant and high reward for only 1 of the n actions taken.

Given that, the reward system was adjusted in 4 different ways, in order to explore two aspects: (a) the use of overlap score that considers the relative order of *reads* and/or (b) the use of dense rewards. These new reward systems are referred here by approaches 1.1, 1.2, 1.3 and 1.4, as follows.

As proposed in the seminal approach, the reward system of the 1.1 approach defines that actions leading to final states produce a bonus reward (of 1.0) that is added to another numerical overlap score between all subsequent *reads* used since the initial state. Thus, these actions produces a reward corresponding to the sum of the normalized overlap score (ranging from 0 to 1) of each pair of *reads*, taking into account the relative order of them.

Still, every action that leads to a non-final state produces a constant and low reward (0.1). Equation 5 formalizes the reward system for Approach 1.1, with $PM_{norm}(s')$ representing the normalized overlap between the *reads* used to reach the s' (see more information on the normalized overlap calculation in Section 2 of the supplementary material).

$$r(s, a, s') = \begin{cases} PM_{norm}(s') + 1.0 & \text{if } s' \text{ is a final state,} \\ 0.1 & \text{otherwise} \end{cases} \quad (5)$$

Despite the use of overlap score that considers the order of the *reads* in approach 1.1, it is susceptible to the sparse rewards problem — as well as in the seminal approach. Although it often produces a small, constant and usually positive reward, and not a zero-value reward as traditionally applied by sparse reward systems, only few and sparse state-action pairs would produce higher rewards.

We can observe in both systems (from Equations 4 and 5) that there is no rewards provided during learning process to guide the agent towards its goal (since any *read* incorporated would produce a reward 0.1). Thus, the agent's learning process would depend exclusively on the sparse actions taken during the exploration of this vast space of states.

Considering that the agent’s learning process tends to take longer in environments suffering from the sparse reward problem³¹, the proposed changes in approaches 1.2, 1.3 and 1.4 mainly focused on improving this aspect and, for this, higher rewards, previously obtained only at the end of the episode, were distributed for each action taken in each episode.

Thus, these approaches — in addition to making the reward system dense, instead of sparse as originally proposed — focused on reducing or eliminating the occurrence of inconsistencies, which, from the genome assembly perspective, would allow permutations of unaligned *reads* to produce maximum accumulated rewards.

Equations 6, 7 and 8 represent, respectively, the reward systems for approaches 1.2, 1.3 and 1.4 — so that $ol_{norm}(s, s')$ represents the normalized overlap between the two subsequent *reads* that allowed to reach state s and then s' .

$$r(s, a, s') = PM_{norm}(s') \quad (6)$$

$$r(s, a, s') = \begin{cases} PM_{norm}(s') + 1.0 & \text{if } s' \text{ is a final state,} \\ ol_{norm}(s, s') & \text{otherwise} \end{cases} \quad (7)$$

$$r(s, a, s') = \begin{cases} ol_{norm}(s, s') + 1.0 & \text{if } s' \text{ is a final state,} \\ ol_{norm}(s, s') & \text{otherwise} \end{cases} \quad (8)$$

Approach 2: Pruning-based action elimination

One of the great challenges for applying reinforcement learning in real-world problems is the high dimensionality of the states space that must be explored by the agents^{32,33}. The Q-learning algorithm is especially susceptible to the dimensionality curse, as the number of states (as well as the number of actions) directly affects the data structure required for the agent’s learning³⁴. Given that the number of actions in each state directly affects the number of states in the seminal proposal, eliminating actions are a strategy for dealing with the high dimensionality³².

To reduce even more the states space proposed by the seminal approach, a heuristic procedure was applied to eliminate actions where states reached directly or indirectly had already been fully explored and where the maximum cumulative reward achieved is smaller than the cumulative reward obtained by taking any other action available in the state.

In Figure 4, we have an example of this action elimination in the given states space considering 3 reads. Looking again at the changed states space as a tree — so removing actions associated with *reads* that have already been used, we see 16 states in this illustration, 6 of them are absorbing states (in the base of the tree) and also correspond to final states following the proposed modeling approach.

For better understanding the pruning process, note that 3 out of 6 final states are highlighted in black, while the remaining states are in gray and white. Black states correspond to explored final states (i.e. already visited by the agent). Gray states, such as the one reached by taking action a in the initial state, represent the states in which all children states have been fully visited in the learning process. Finally, white states (final or not) are those not yet explored and/or that have unexplored children — e.g. the initial state, where one child is not explored and the other one is partially explored.

When reaching an unexplored final state, such as the rightmost final state in Figure 4, the accumulated rewards achieved since the initial state is maintained and propagated for its predecessors, maintaining only the highest value propagated for the children. Each reward is represented by integer numbers within the states in the figure. Thus, each non-final state stores the highest accumulated reward achieved from it during the training process.

Based on this information, it is possible to prune irrelevant actions, those actions that, after taken, do not produce the maximum accumulated reward. This type of action can be found in action a of the initial state in Figure 4. Note that all possible achievable states after taking this action have been explored and the maximum cumulative reward obtained is 6, obtained from consecutively taking the actions a, b, c . Also note that action c in the initial state, even that not fully explored, is capable of producing a higher reward, equals to 8, and obtained by taking the actions c, a and b , in that order.

Thus, when the agent first goes through the sequence of states corresponding to the actions c, a and b , the pruning mechanism propagates the maximum reward value up to the initial state and, at that moment, it cuts the action a from the initial state. The pseudo code presented in Algorithm 1 presents the procedure to update pruning process when the last explored final state (referred to as *state*) is reached obtaining the corresponding accumulated reward achieved (referred to as *newReward*).

Approaches 3: Evolutionary-based exploration

In this proposal, the potential for mutual collaboration between reinforcement learning and evolutionary computing was investigated — by applying the elitist selection of the genetic algorithm^{35,36} — to optimize the exploration of the states space. For assessing the individual contribution of the genetic algorithm in this hybrid proposal, this approach has been divided into two smaller approaches, referred to as approaches 3.1 and 3.2 and presented in the following subsections.

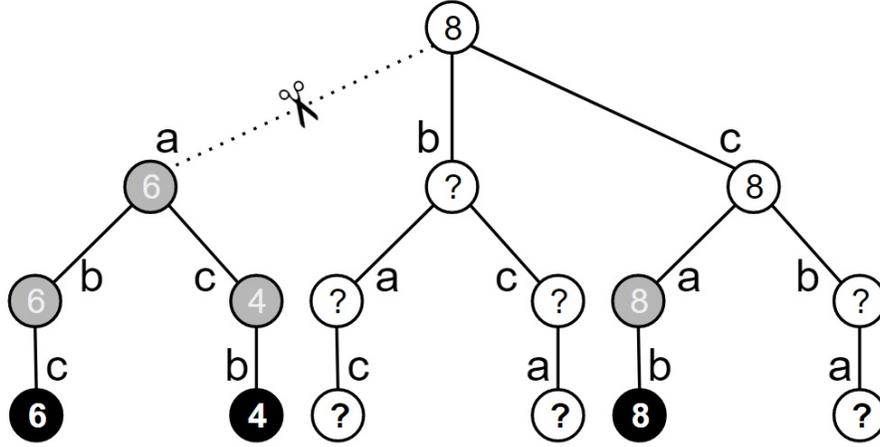


Figure 4. Illustration of the pruning procedure for a state space referring to the assembly of 3 reads, referred by a , b and c . The generic pruning procedure is defined in detail by Algorithm 1

Algorithm 1 Pruning's algorithm

```

1: procedure PRUNE( $state : treeNode, newReward : float$ )
2:   if  $state \neq null$  and ( $state.unseen$  or  $newReward > state.maxReward$ ) then  $\triangleright state.unseen$  starts  $true, \forall_{state}$ 
3:      $state.unseen \leftarrow false$ 
4:      $state.maxReward \leftarrow newReward$ 
5:     if  $state.final$  then
6:       PRUNEUSELESSCHILDREN( $state$ )  $\triangleright$  i.e. prune all fully explored children where  $maxReward < newReward$ 
7:       PRUNE( $state.parent, newReward$ )

```

Approach 3.1: Evolutionary-aided reinforcement learning assembly

The strategy of applying ϵ -greedy — and its variations — is a classic solution for expanding the exploration of agents trained by the Q-learning algorithm, as it allows a broader initial exploration, achieving the optimal policy once the states space has been sufficiently explored¹⁸. However, the existing trade-off between exploitation and exploration remains a major and challenging problem for reinforcement learning in high-dimensional environments^{37,38}.

Searching for a more efficient exploration process and also considering the good performance of genetic algorithms in a similar genome assembly approach carried out by Oliveira et al.³⁹, here, the interaction between reinforcement learning and evolutionary computation was introduced into the exploration process.

This approach is based on the traditional operation of the Q-learning algorithm. However, in each Q-learning episode, the sequence of actions is stored, and at the end of the episode it is transformed into a chromosome of an initial population, that later will evolve. This procedure is presented in Figure 5, where we can see the list of actions made in each episode being used for constructing a new individual.

New chromosomes are inserted into the initial population until the amount of chromosomes reaches the predefined and expected size for this population. At this point, agent training is interrupted and m genetic generations are carried out — being m also predetermined (for details see Section 4 of supplementary material) and applying the normalized sum of overlaps between reads as the adaptive function — the same applied in Equation 8 and detailed in Section 2 of the supplementary material.

After m generations, according to the objective function, the most fit individual is used for conducting the next episode in the agent's reinforcement learning training, hitherto interrupted. As each gene of the individual's chromosome corresponds to one possible action, the complete gene sequence will contain distinct successive actions to be taken by the agent in the current episode, producing then a mutual collaboration between reinforcement learning and the genetic algorithm — the initial populations of the genetic algorithm are produced by reinforcement learning and, as a counterpart, the results from the evolution of the genetic algorithm is introduced in an episode of reinforcement learning.

For your better understanding of all the aforementioned approaches, considering they share the common basis of the Q-learning algorithm, Figure 6 presents a flowchart where we can see the proposed updates for these approaches. The elements highlighted in gray represent the procedures and conditions of Q-learning algorithm — which, briefly, is based on the action

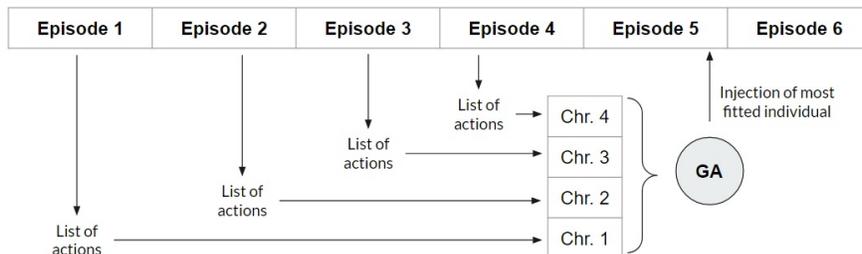


Figure 5. Illustration of the proposed interaction between reinforcement learning and genetic algorithm. At each reinforcement learning episode, the actions taken by the agent are converted into the chromosome (having each action as a gene) of an individual of the initial population of the genetic algorithm, whose size n is predefined. After n episodes (and thus n individuals in the initial population), this population evolves for an also predefined number of generations through the genetic algorithm. Then, the most adapted individual of the last generation is obtained. In the end, that individual’s chromosome genes are used as actions in the next episode of reinforcement learning.

choice, action taking and the update of Q-values in the Q-table during several episodes, all of them starting in an initial state and ending when an absorbing state is reached.

As a consequence, these elements (in gray) fully represent Approaches 1.1, 1.2, 1.3, and 1.4, as the proposal is focused on replacing the reward system proposed by the seminal approach. Approaches 2 and 3.1 emerge as complementary to this common basis of procedures with specific procedures, namely the insertion of the dynamic pruning mechanism in approach 2 — represented by the dashed element with double edges — and the introduction of mutual contribution with the genetic algorithm to improve the agent’s exploration performance — represented by the dashed elements with simple edges.

Approach 3.2: Evolutionary-based assembly

Finally, to estimate the contribution of the genetic algorithm in Approach 3.1, which applies a mutual collaboration between reinforcement learning and the genetic algorithm, the genetic algorithm assembling performance was evaluated separately, following the same configurations set of the previous approach, but this time, adopting as starting population a set of individuals whose chromosomes were built from random permutations without repetition of *reads*.

Datasets

To assess the performance of all approaches (including seminal approach), in addition to the 18 datasets proposed and made available by Xavier et al.³⁰, 5 novel datasets derived from other microgenomes extracted from the genome of the previous studies^{14,30} were complementarily created here. These microgenomes are not arbitrary genome fragments, as the previously used microgenomes (which had 25bp and 50bp), but represent larger fragments of previously annotated genes from the corresponding organism (i.e. *E. coli*). The experiments were then carried out on 23 datasets, whose microgenomes sizes, number of *reads* and sizes of *reads* are presented in Table 1 — the last 5 lines correspond to the 5 datasets derived from genes.

As each of these datasets correspond to a reinforcement learning environment, an environment for each of them was created in the OpenAI Gym toolkit⁴⁰, in order to share such reinforcement learning challenges in a simple way. These environments are provided in <http://github.com/kriowloo/gymnome-assembly> (see section 1 of the supplementary material for additional technical information) and use the reward system proposed in Approach 1.4. The identification names of each environment are presented in the last column of Table 1. The seminal reward system is also implemented and available — for running it, use the version 1, replacing v2 by v1 in the environment name field.

Two experiments were then carried out for evaluating the approaches. In each experiment, 20 successive runs of each evaluated approach was performed for all the 23 existing datasets; totalizing 460 runs per approach. Given that each approach has different levels of complexity, the real execution time for each approach was considered for comparing them. To reduce the interference of external factors in execution time, all experiments were individually and sequentially performed in the same station (with Ubuntu 16.04 in an AWS EC2 instance of the *r5a.large* type, dual core, with 16GB RAM and 30GB of storage).

In the first experiment, here referred to as *Experiment A*, the objective was to verify the impact of progressively including the previously described strategies. For this, the performance of the seminal approach was evaluated (according to 14) against approaches 1.1, 1.2, 1.3, 1.4 (which modify the reward system), 2 (which includes the pruning dynamic) and 3.1 (which uses the AG as an complementary factor).

In the second experiment, referred to as *Experiment B*, the main objective was to compare the performance of the new RL-based approaches against the performance of the AG alone. Therefore, in addition to Approaches 1.1, 1.2, 1.3, 1.4, 2 and

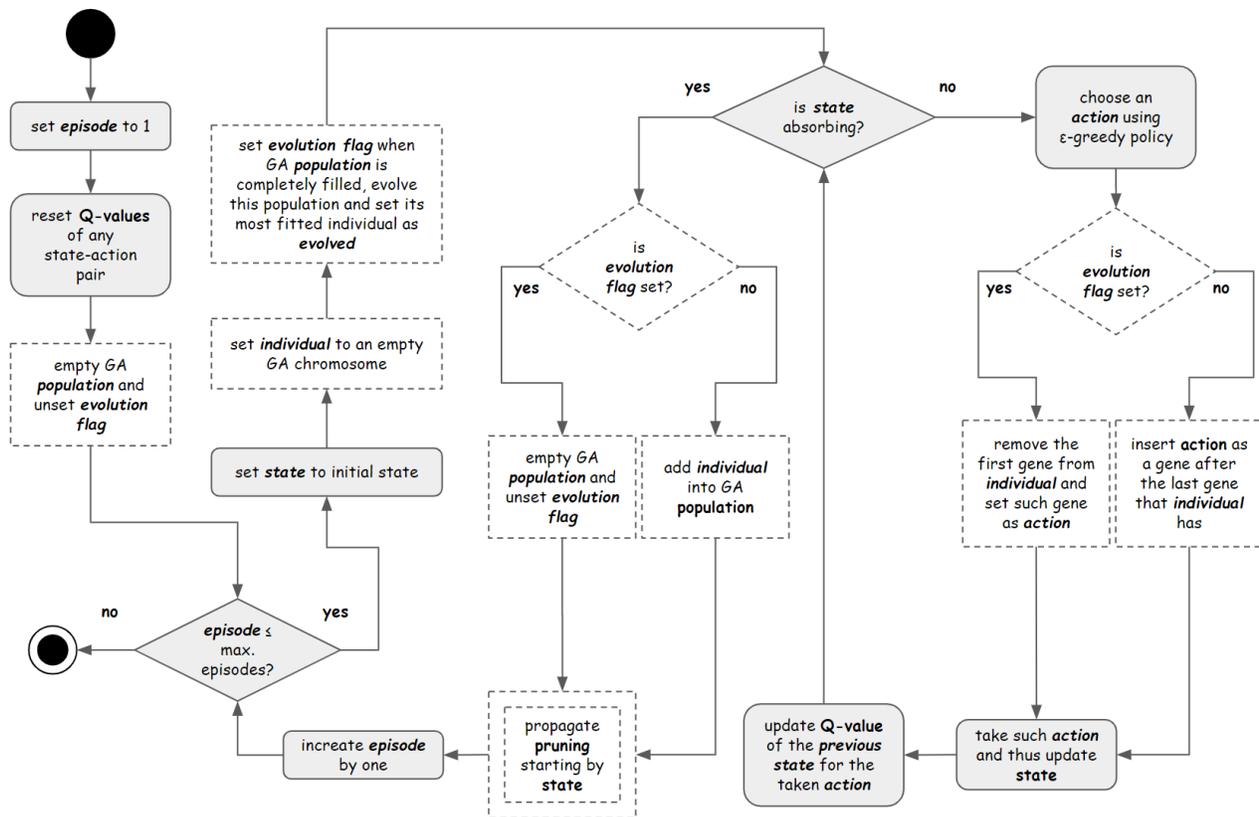


Figure 6. Flowchart representing Approaches 1.1, 1.2, 1.3, 1.4, 2 and 3.1, so that Approaches 1.1, 1.2, 1.3, and 1.4 are defined by the elements in gray, Approach 2 by the dashed element with double edges and Approach 3.1 by the dashed elements with single border.

3.1, the approach 3.2 (which explores GA alone) was performed in an equivalent time.

For performance measure in each experiment, two percentage measures were calculated, called distance-based measure (DM) and reward-based measure (RM). Evaluations of *de novo* assembly are commonly performed using proper metrics, such as the N50⁴¹. These metrics were created because, as previously indicated, *de novo* assemblies are not supported by a reference genome. This way, in some scenarios, it is not possible to accurately assess the results obtained from the assemblers — as the optimal output is unknown.

Here, although a *de novo* assembler is evaluated, its assessment environment is restricted and the target genomes are known, and this scenario allows the use of specific (and exact) evaluations, such as the DM and RM metrics.

DM considers that a run was successful when the consensus sequence resulting from the orders of *reads* produced in that run is exactly identical to the expected sequence. RM, however, consider any run as successful when the proposed order of *reads* presents the corresponding sum of PM_{norm} higher or equal to the sum of PM_{norm} from the optimal *reads* sequence (for details, see Section 3 of the supplementary material).

Results

The results obtained from *Experiment A* are presented in Table 2, where it is possible to observe that the seminal approach, although consuming the longest running time (23 hours and 34 minutes), also presented the lowest average performances, obtaining an optimal response in only 16.96% of the runs (i.e. 78 out of 460 executions) in terms of distance between the produced and the expected genome (DM) and 21.30% (98 out of 460) in terms of maximum reward (RM). This difference is based on the previously mentioned inconsistency of the proposed reward system, which allowed non-optimal permutations to produce maximum accumulated rewards.

Beyond that, following the updates in the rewards system, the DM and MR performances in Approaches 1.2, 1.3, and 1.4 surpassed those of the previous approach, and they also consumes about 4 hours less (19 hours and 38 minutes of execution time). In such experiments the part of the gains are due to the improved agent's performance in one of the sets where the sum of rewards for the optimal permutation of *reads* were not maximum in the previous reward system (as presented in Figure 3).

Microgenome size	Number of reads	Read size	Env. name
25	10	8	GymnomeAssembly_25_10_8-v2
25	10	10	GymnomeAssembly_25_10_10-v2
25	10	15	GymnomeAssembly_25_10_15-v2
50	10	8	GymnomeAssembly_50_10_8-v2
50	10	10	GymnomeAssembly_50_10_10-v2
50	10	15	GymnomeAssembly_50_10_15-v2
25	20	8	GymnomeAssembly_25_20_8-v2
25	20	10	GymnomeAssembly_25_20_10-v2
25	20	15	GymnomeAssembly_25_20_15-v2
50	20	8	GymnomeAssembly_50_20_8-v2
50	20	10	GymnomeAssembly_50_20_10-v2
50	20	15	GymnomeAssembly_50_20_15-v2
25	30	8	GymnomeAssembly_25_30_8-v2
25	30	10	GymnomeAssembly_25_30_10-v2
25	30	15	GymnomeAssembly_25_30_15-v2
50	30	8	GymnomeAssembly_50_30_8-v2
50	30	10	GymnomeAssembly_50_30_10-v2
50	30	15	GymnomeAssembly_50_30_15-v2
381	20	75	GymnomeAssembly_381_20_75-v2
567	30	75	GymnomeAssembly_567_30_75-v2
726	40	75	GymnomeAssembly_728_40_75-v2
930	50	75	GymnomeAssembly_930_50_75-v2
4224	230	75	GymnomeAssembly_4224_230_75-v2

Table 1. Data sets used in the experiments. The first column shows the size (in *bp*) of the microgenome used to generate the *reads* of each set; the second column shows the number of *reads* generated; the third column shows the size of the generated *reads*; and the fourth column shows the name of the environment built for each set in the OpenAI Gym toolkit (accessible in <http://github.com/kriowloo/gymnome-assembly>).

Despite the gains obtained from the updated reward system, it was possible to note, based on the results, that the previously mentioned inconsistencies were not completely resolved. In some of the datasets the agent reached and even surpassed the maximum expected accumulated rewards, however, without obtaining the target genome.

A minor improvement is observed after the application of the approach 2, with the incorporation of the agent’s dynamic pruning, presenting a performance slightly superior to that of the agent in Approaches 1, thus requiring approximately one hour less of processing. The great highlight for this comparison, however, is presented in Approach 3.1, which, while benefiting from the environment exploration improved by AG, has reached a significantly improved result in a reduced amount of time for execution.

Experiment A	Seminal	Appr. 1.1	Appr. 1.2	Appr. 1.3	Appr. 1.4	Appr. 2	Appr. 3.1
Average DM	16.96%	9.57%	18.48%	20.00%	20.43%	20.65%	73.91%
Average RM	21.30%	13.70%	21.30%	24.35%	24.78%	25.00%	80.87%
Total runtime	23h34m	19h38m	19h38m	19h38m	19h38m	18h41m	17h03m

Table 2. Results of Experiment A, which compares the performances of trained agents with different reinforcement learning strategies. The performances of each approach are expressed using distance-based (DM) and reward-based (RM) metrics (see *Methods* for details).

In order to analyze the contribution of each technique for the gain of performance in the approach 3.1, a second experiment was performed, here referred to as *Experiment B*, to compare the performance of RL-based approaches (approaches 1.4, 2 and 3) and the isolated use of AG, but using random initial populations (approach 3.2). The results of Experiment B are presented in Table 3.

Given the remarkable performance of Approach 3.2, Experiment B applied as reference the time taken by the AG to find an optimal solution in terms of RM for 22 out of the 23 datasets used (i.e. 95.65%), which corresponds to 1 hour and 34 minutes.

Experiment B	Appr. 1.4	Appr. 2	Appr. 3.1	Appr. 3.2
Average DM	13.91%	12.39%	14.78%	87.83%
Average RM	17.61%	16.30%	14.78%	95.65%
Total runtime	01h36m	01h36m	01h42m	01h34m

Table 3. Experimental performances considering similar running times (RT). Performances were expressed using Distance-based Measure (DM) and Reward-based Measure (RM) (see *Methods*)

Given this results, it is evident the superiority of the results obtained from Approach 3.2 when compared to Approaches 1.4, 2 and 3.1.

Finally, after verifying the dominance of Approach 3.2, an additional experiment was carried out, aiming to verify the limits of this approach. It considers the dataset for which no optimal response was achieved. This dataset corresponds to the reads from the largest gene (4Kbp) selected for the production of simulated datasets.

For that experiment, the running time was considerably increased, reaching approximately 38 hours of running — surpassing the total time allocated for running all datasets in the other experiments, considering about 30 minutes per run (against less than 2 minutes for the same dataset in the tests run for approach 3.2 of *Experiment B*).

Although with an expressive increase in execution time, no optimal solution has yet been reached for any of the 20 runs of this dataset. However, as presented Figure 7, through violinplots, it is possible to observe a consistent gain in performance, both in terms of DM (where longer runs had smaller distances than most distances reached by shorter runs) and RM (which was higher accumulated rewards in all longer time runs).

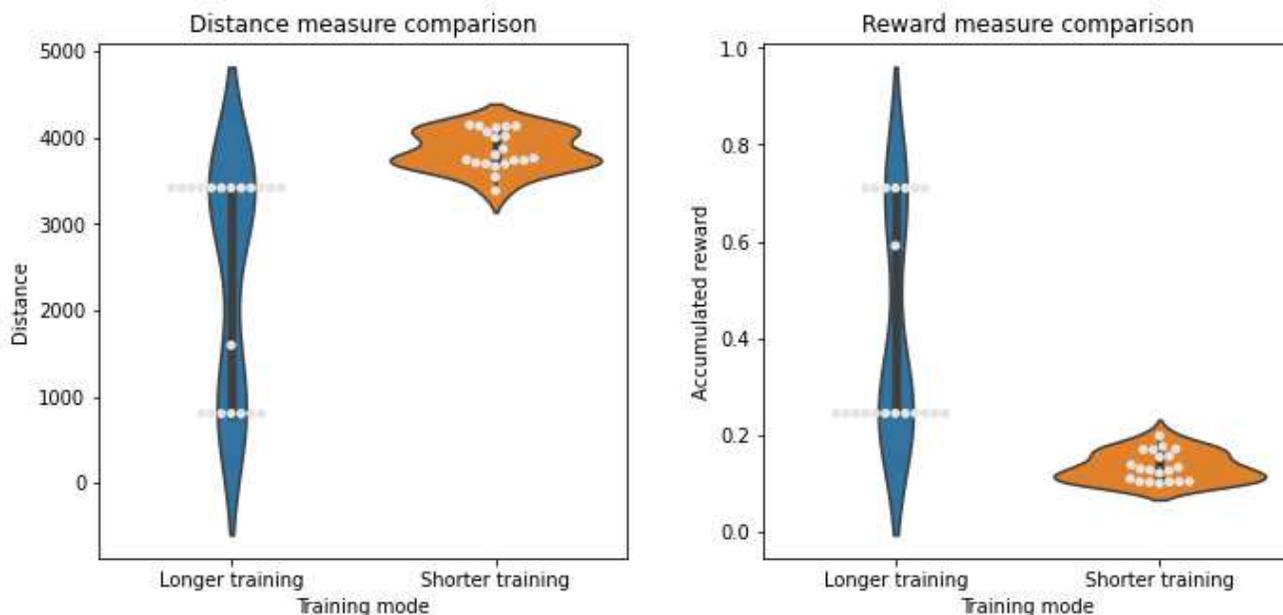


Figure 7. Violin plots demonstrating the performances obtained by the GA in experiments with short (1h34m) and long (37h58m) execution times in terms of sums of DM and RM. Gray dots represent the distance/reward obtained for all runs; the black line in the middle indicates interquartile range; and the violin curves shows the distribution density, where the wider the section, the greater the probability of the observations take the corresponding value.

Discussion

As mentioned, genome assembly is among the most complex problems confronted by computer scientists within the context of genomics projects, regardless the importance of its results for scientific development. This complexity, in computational terms, allocates the problem of finding an optimal permutation of sequenced *reads* and reaching the target genome into a class of problems called NP-hard which comprises the most difficult problems in computer science⁴².

This high complexity is particularly expressed in the vast space of states required for representing the assembly problem into the modeling proposed by the seminal approach and, consequently, into the approaches proposed here. To achieve the optimal solution in sets of *reads* of only 30 sequences in the seminal approach, for example, the agent should explore a states space composed of approximately $2e44$ states¹⁴ (this number exceeds the estimated number of stars in the universe and corresponds to the number of sand grains on Earth). It is also worth noting that, in real-world scenarios, it is common to sequence millions of *reads* per genome — which increases even more the corresponding states space.

The approaches proposed in this study aimed to expand the agent's learning based on two difficulties observed in the seminal approach for applying reinforcement learning into the genome assembly problem: (1) the reward system and (2) the agent's exploration strategies. Both the updates in the rewards system and the incorporation of new explorative strategies improved the agent's learning performance, as demonstrated by the results.

The definition of extrinsic rewards is one of the most challenging tasks for the construction of environments suitable for the agent's learning. The updates into the reward systems proposed here favored the agent's improved learning. However, it is not yet an optimal solution for the problem, as some non-optimal solutions are still present in datasets that produced maximum accumulated rewards. This would justify the occurrence of RM percentages higher than DM percentages in several experiments.

The dynamic pruning mechanism showed a discreet improvement for the learning process. However, the relationship between the additional processing cost resulting from the application of this mechanism and the benefit obtained from its implementation did not indicate a reasonable net gain from its use as bypass for the problem emerging from the high dimensionality of the space of states.

In this sense, the application of the RL strategy combined with the AG, in the hybrid approach, presented a much more attractive performance for supporting the exploration process in the long term. This combination was proved to be advantageous, probably given the dimensionality curse encountered by the Q-learning algorithm, as a strong AG support was observed for the agent while conduct the RL exploration.

Despite the performance improvements, it remains evident the insufficiency of the approaches when applied to real-world scenarios. This insufficiency is more evident in the experiments performed with the largest dataset analyzed, with a size that corresponds to a gene of approximately *4Kbp*. Despite being the largest dataset employed, this gene remains much smaller than the smallest genomes from living organisms. None of the applied approaches presented an optimal response in this scenario, not even when applying the isolated genetic algorithm for a time that was largely longer than the time applied for the tests with the other datasets.

This finding, together with the superior results obtained from the GA alone, allow us to conclude on the infeasibility of applying the Q-learning algorithm to solve the genomes assembly problem in search for an optimal *reads* permutation, as originally proposed in the seminal approach.

However, given the absence of approaches in the literature for tackling the problem through reinforcement learning and also considering the optimistic results obtained from RL, especially when combined to deep learning, further investigations on the applicability of reinforcement learning are required, including the use of different modeling approaches and algorithms.

Considering the importance of reproducing scientific studies and with the special intention of supporting future investigations, all experiments performed in this study, as well as the reinforcement learning environments applied to simulate the assembly problems, are available at https://osf.io/tp4zj/?view_only=8bd12290d9d4480f95bc2c6e49c73ff7, and are open for reuse (for details on how to reproduce the experiments see Section 5 of the supplementary material).

One of the major challenges for applying reinforcement learning to real-world problems is the low sample efficiency of the algorithms⁴³; and it is not different in the genome assembly problem treated here. From the time required by the agent trained by the Q-learning algorithm to reach an optimal solution, it is possible to perceive a high need for numerous interactions with the data. Considering that real-world inputs are even bigger than those experimentally applied here, obtaining a sample efficient algorithm for the problem would be required for supporting the development of solutions applicable into real-world problems.

One aspect that directly interferes with the agent's sample efficiency is optimizing the exploration of the space of states. In this sense, the application of techniques that remove duplicate *reads* — due to repeats — and the use of an intrinsic motivation could be alternatively investigated to bypass the exploration problem, given the high dimensionality of the proposed space of states^{43,44}.

As previously mentioned, it is important to continue the investigation for updating and replacing the Q-learning algorithm, as well as the use of distributed approaches and/or algorithms using eligibility traces. Although simple, Q-learning is a very powerful reinforcement learning method for solving tasks and remains applied for real-world applications^{45,46}. As is the case with other RL methods, however, its use in real-world problems where the space of states is too vast is not recommended. The use of approximating methods for the Q function may prove opportune in such cases, especially given the promising results obtained in this study through the combination of RL with deep learning⁴⁷.

Although it faces some obstacles for being implemented into commercial applications, it is reasonable to consider the recent achievements of deep reinforcement learning when applied to games, which presents an equivalent computational problem to

that of the genome assembly^{21,22,47}. One of the great challenges for production equivalent proposals for the assembly problem is that several promising works apply convolutional artificial networks, which use images as inputs.

The transformation of real problems into games is a possibility for reusing the promising technology developed focusing on games⁴⁸. In this sense, the modeling of the genome assembly problem through a game can work as an alternative representation of the problem and may reduce the space of states explored by the agent. Representing the assembly problem as a maze with multiple targets/objectives could be an example for such modeling approach.

One of the main benefits of representing the problem as a game is the reduction of the space of actions, which, in the proposed modeling approach, increases with the number of *reads*. Mnih et. al²⁰, for example, were capable of achieving a common deep reinforcement learning architecture capable of resolving several Atari games, being remarkable that the best performances were achieved in games with fewer actions, as in the game *Breakout*, which requires only two actions (move right or left). In this game the agent was able to learn, from its performance and without instruction, that producing a gap is a valid strategy for optimizing its results (illustrated in Figure 8).

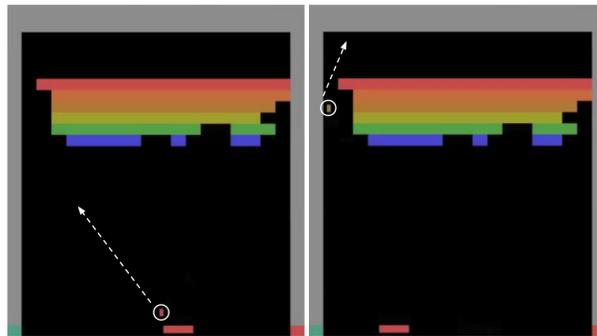


Figure 8. Example of application of deep reinforcement learning in games where the agent was able to find strategies to optimize accumulated rewards. Here, the agent was able to discover that caving a tunnel could maximize rewards gain (images adapted from a video demonstration of learning progress in ²⁰).

The use of *Graph Embedding* may act as an option of modeling approach allowing the use of deep reinforcement learning without requiring the conversion of the problem into an image — especially when considering that the genome assembly problem may be represented through a graph for optimizing the problem, in the shape of the Traveling Salesman Problem (TSP)^{49,50}. As presented by Vesselinova et. al, numerous studies investigated the application of deep reinforcement learning for solving graph problems, including TSP⁵¹.

Finally, one other aspect to be considered before the adoption of reinforcement learning into the genome assembly problem is the generalization of the agent's learning — a major challenge for the use of RL in real-world problems⁵². As designed into the RL environment for the genome assembly problem, the learning acquired by the agent when assembling a set of *reads* will hardly be applied for the assembly of a new set.

Data availability

All data generated or analysed during this study are included in this published article (and its Supplementary Information files).

References

1. de Souza, K. P. *et al.* Machine learning meets genome assembly. *Brief. Bioinforma.* DOI: [10.1093/bib/bby072](https://doi.org/10.1093/bib/bby072) (2018).
2. Li, Z. *et al.* Comparison of the two major classes of assembly algorithms: overlap-layout-consensus and de-bruijn-graph. *Briefings Funct. Genomics* **11**, 25–37, DOI: [10.1093/bfpg/elr035](https://doi.org/10.1093/bfpg/elr035) (2011).
3. Manzoni, C. *et al.* Genome, transcriptome and proteome: the rise of omics data and their integration in biomedical sciences. *Briefings Bioinforma.* **19**, 286–302, DOI: [10.1093/bib/bbw114](https://doi.org/10.1093/bib/bbw114) (2016).
4. Paszkiewicz, K. & Studholme, D. J. De novo assembly of short sequence reads. *Briefings Bioinforma.* **11**, 457–472, DOI: [10.1093/bib/bbq020](https://doi.org/10.1093/bib/bbq020) (2010).
5. Heather, J. M. & Chain, B. The sequence of sequencers: The history of sequencing DNA. *Genomics* **107**, 1–8, DOI: [10.1016/j.ygeno.2015.11.003](https://doi.org/10.1016/j.ygeno.2015.11.003) (2016).
6. Rodríguez-Ezpeleta, N., Hackenberg, M. & Aransay, A. M. (eds.) *Bioinformatics for High Throughput Sequencing* (Springer New York, 2012).

7. Portin, P. & Wilkins, A. The evolving definition of the term “gene”. *Genetics* **205**, 1353–1364, DOI: [10.1534/genetics.116.196956](https://doi.org/10.1534/genetics.116.196956) (2017).
8. Ji, P., Zhang, Y., Wang, J. & Zhao, F. MetaSort untangles metagenome assembly by reducing microbial community complexity. *Nat. Commun.* **8**, DOI: [10.1038/ncomms14306](https://doi.org/10.1038/ncomms14306) (2017).
9. Wong, H. L., MacLeod, F. I., White, R. A., Visscher, P. T. & Burns, B. P. Microbial dark matter filling the niche in hypersaline microbial mats. *Microbiome* **8**, DOI: [10.1186/s40168-020-00910-0](https://doi.org/10.1186/s40168-020-00910-0) (2020).
10. Medvedev, P., Georgiou, K., Myers, G. & Brudno, M. Computability of models for sequence assembly. In *Lecture Notes in Computer Science*, 289–301, DOI: [10.1007/978-3-540-74126-8_27](https://doi.org/10.1007/978-3-540-74126-8_27) (Springer Berlin Heidelberg, 2007).
11. Pop, M. Genome assembly reborn: recent computational challenges. *Briefings Bioinforma.* **10**, 354–366, DOI: [10.1093/bib/bbp026](https://doi.org/10.1093/bib/bbp026) (2009).
12. Böckenhauer, H.-J. & Bongartz, D. DNA sequencing. In *Algorithmic Aspects of Bioinformatics*, 171–209, DOI: [10.1007/978-3-540-71913-7_8](https://doi.org/10.1007/978-3-540-71913-7_8) (Springer Berlin Heidelberg, 2007).
13. Gurevich, A., Saveliev, V., Vyahhi, N. & Tesler, G. QUASt: quality assessment tool for genome assemblies. *Bioinformatics* **29**, 1072–1075, DOI: [10.1093/bioinformatics/btt086](https://doi.org/10.1093/bioinformatics/btt086) (2013).
14. Bocicor, M.-I., Czibula, G. & Czibula, I.-G. A reinforcement learning approach for solving the fragment assembly problem. In *2011 13th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*, DOI: [10.1109/synasc.2011.9](https://doi.org/10.1109/synasc.2011.9) (IEEE, 2011).
15. Henrique, B. M., Sobreiro, V. A. & Kimura, H. Literature review: Machine learning techniques applied to financial market prediction. *Expert. Syst. with Appl.* **124**, 226–251, DOI: [10.1016/j.eswa.2019.01.012](https://doi.org/10.1016/j.eswa.2019.01.012) (2019).
16. LeCun, Y. 1.1 deep learning hardware: Past, present, and future. In *2019 IEEE International Solid-State Circuits Conference - (ISSCC)*, DOI: [10.1109/isscc.2019.8662396](https://doi.org/10.1109/isscc.2019.8662396) (IEEE, 2019).
17. Botvinick, M. *et al.* Reinforcement learning, fast and slow. *Trends Cogn. Sci.* **23**, 408–422, DOI: [10.1016/j.tics.2019.02.006](https://doi.org/10.1016/j.tics.2019.02.006) (2019).
18. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (A Bradford Book, Cambridge, MA, USA, 2018).
19. Silver, D., Singh, S., Precup, D. & Sutton, R. S. Reward is enough. *Artif. Intell.* **299**, 103535, DOI: <https://doi.org/10.1016/j.artint.2021.103535> (2021).
20. Mnih, V. *et al.* Human-level control through deep reinforcement learning. *Nature* **518**, 529–533, DOI: [10.1038/nature14236](https://doi.org/10.1038/nature14236) (2015).
21. Silver, D. *et al.* Mastering the game of go without human knowledge. *Nature* **550**, 354–359, DOI: [10.1038/nature24270](https://doi.org/10.1038/nature24270) (2017).
22. Vinyals, O. *et al.* Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* **575**, 350–354, DOI: [10.1038/s41586-019-1724-z](https://doi.org/10.1038/s41586-019-1724-z) (2019).
23. Nian, R., Liu, J. & Huang, B. A review on reinforcement learning: Introduction and applications in industrial process control. *Comput. & Chem. Eng.* **139**, 106886, DOI: [10.1016/j.compchemeng.2020.106886](https://doi.org/10.1016/j.compchemeng.2020.106886) (2020).
24. Dulac-Arnold, G., Mankowitz, D. J. & Hester, T. Challenges of real-world reinforcement learning. In *ICML 2019 Workshop on Reinforcement Learning for Real Life (RLRL)* (2019).
25. Vollmers, J., Wiegand, S. & Kaster, A.-K. Comparing and evaluating metagenome assembly tools from a microbiologist’s perspective - not only size matters! *PLOS ONE* **12**, e0169662, DOI: [10.1371/journal.pone.0169662](https://doi.org/10.1371/journal.pone.0169662) (2017).
26. van der Walt, A. J. *et al.* Assembling metagenomes, one community at a time. *BMC Genomics* **18**, DOI: [10.1186/s12864-017-3918-9](https://doi.org/10.1186/s12864-017-3918-9) (2017).
27. Cormen, T. H., Leiserson, C. E., Rivest, R. L. & Stein, C. *Introduction to Algorithms, Third Edition* (The MIT Press, 2009), 3rd edn.
28. Grinstead, C. & Snell, J. *Introduction to Probability* (American Mathematical Society, 2012).
29. Smith, T. & Waterman, M. Identification of common molecular subsequences. *J. Mol. Biol.* **147**, 195–197, DOI: [10.1016/0022-2836\(81\)90087-5](https://doi.org/10.1016/0022-2836(81)90087-5) (1981).
30. Xavier, R., de Souza, K. P., Chateau, A. & Alves, R. Genome assembly using reinforcement learning. In Kowada, L. & de Oliveira, D. (eds.) *Advances in Bioinformatics and Computational Biology*, 16–28 (Springer International Publishing, 2020).

31. Trott, A., Zheng, S., Xiong, C. & Socher, R. Keeping your distance: Solving sparse reward tasks using self-balancing shaped rewards. In Wallach, H. M. *et al.* (eds.) *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, 10376–10386 (2019).
32. Zahavy, T., Haroush, M., Merlis, N., Mankowitz, D. J. & Mannor, S. Learn what not to learn: Action elimination with deep reinforcement learning. In Bengio, S. *et al.* (eds.) *Advances in Neural Information Processing Systems 31*, 3562–3573 (Curran Associates, Inc., 2018).
33. Dulac-Arnold, G. *et al.* An empirical investigation of the challenges of real-world reinforcement learning. *CoRR abs/2003.11881* (2020). [2003.11881](https://arxiv.org/abs/2003.11881).
34. de Wiele, T. V., Warde-Farley, D., Mnih, A. & Mnih, V. Q-learning in enormous action spaces via amortized approximate maximization. *CoRR abs/2001.08116* (2020). [2001.08116](https://arxiv.org/abs/2001.08116).
35. Baluja, S. & Caruana, R. Removing the genetics from the standard genetic algorithm. In *In Proceedings of ICML'95*, 38–46 (Morgan Kaufmann Publishers, 1995).
36. Konar, A. Evolutionary computing algorithms. In *Computational Intelligence*, 323–351, DOI: [10.1007/3-540-27335-2_12](https://doi.org/10.1007/3-540-27335-2_12) (Springer Berlin Heidelberg, 2005).
37. Gimelfarb, M., Sanner, S. & Lee, C.-G. Epsilon-bmc: A bayesian ensemble approach to epsilon-greedy exploration in model-free reinforcement learning. In Adams, R. P. & Gogate, V. (eds.) *Proceedings of Machine Learning Research*, vol. 115, 476–485 (PMLR, Tel Aviv, Israel, 2020).
38. Peterson, E. J. & Verstynen, T. D. A way around the exploration-exploitation dilemma. *bioRxiv* DOI: [10.1101/671362](https://doi.org/10.1101/671362) (2019).
39. Oliveira, R. R. M. *et al.* GAVGA: A genetic algorithm for viral genome assembly. In Oliveira, E. C., Gama, J., Vale, Z. A. & Cardoso, H. L. (eds.) *Progress in Artificial Intelligence - 18th EPIA Conference on Artificial Intelligence, EPIA 2017, Porto, Portugal, September 5-8, 2017, Proceedings*, vol. 10423 of *Lecture Notes in Computer Science*, 395–407, DOI: [10.1007/978-3-319-65340-2_33](https://doi.org/10.1007/978-3-319-65340-2_33) (Springer, 2017).
40. Brockman, G. *et al.* Openai gym (2016). [1606.01540](https://arxiv.org/abs/1606.01540).
41. Bradnam, K. R. *et al.* Assemblathon 2: evaluating de novo methods of genome assembly in three vertebrate species. *GigaScience* **2**, DOI: [10.1186/2047-217x-2-10](https://doi.org/10.1186/2047-217x-2-10) (2013).
42. Roughgarden, T. *Algorithms Illuminated (Part 4): Algorithms for NP-Hard Problems*. Algorithms Illuminated (Soundlikeyourself Publishing, LLC, 2020).
43. Yu, Y. Towards sample efficient reinforcement learning. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI'18*, 5739–5743 (AAAI Press, 2018).
44. Barto, A. G. Intrinsic motivation and reinforcement learning. In *Intrinsically Motivated Learning in Natural and Artificial Systems*, 17–47, DOI: [10.1007/978-3-642-32375-1_2](https://doi.org/10.1007/978-3-642-32375-1_2) (Springer Berlin Heidelberg, 2012).
45. Chakole, J., Kolhe, M., Mahapurush, G., Yadav, A. & Kurhekar, M. A q-learning agent for automated trading in equity stock markets. *Expert. Syst. with Appl.* **163**, DOI: [10.1016/j.eswa.2020.113761](https://doi.org/10.1016/j.eswa.2020.113761) (2021). Cited By 0.
46. Abdulhai, B., Pringle, R. & Karakoulas, G. J. Reinforcement learning for true adaptive traffic signal control. *J. Transp. Eng.* **129**, 278–285, DOI: [10.1061/\(asce\)0733-947x\(2003\)129:3\(278\)](https://doi.org/10.1061/(asce)0733-947x(2003)129:3(278)) (2003).
47. Fjelland, R. Why general artificial intelligence will not be realized. *Humanit. Soc. Sci. Commun.* **7**, DOI: [10.1057/s41599-020-0494-4](https://doi.org/10.1057/s41599-020-0494-4) (2020).
48. Reis, S., Reis, L. P. & Lau, N. Game adaptation by using reinforcement learning over meta games. *Group Decis. Negot.* DOI: [10.1007/s10726-020-09652-8](https://doi.org/10.1007/s10726-020-09652-8) (2020).
49. Cook, W. J. *Pushing the Limits*, 211–212 (Princeton University Press, 2012).
50. Li, Z. *et al.* Comparison of the two major classes of assembly algorithms: overlap–layout–consensus and de-bruijn-graph. *Briefings Funct. Genomics* **11**, 25–37, DOI: [10.1093/bfgp/elr035](https://doi.org/10.1093/bfgp/elr035) (2011).
51. Vesselinova, N., Steinert, R., Perez-Ramirez, D. F. & Boman, M. Learning combinatorial optimization on graphs: A survey with applications to networking. *IEEE Access* **8**, 120388–120416, DOI: [10.1109/ACCESS.2020.3004964](https://doi.org/10.1109/ACCESS.2020.3004964) (2020).
52. Ponsen, M., Taylor, M. E. & Tuyls, K. Abstraction and generalization in reinforcement learning: A summary and framework. In *Adaptive and Learning Agents*, 1–32, DOI: [10.1007/978-3-642-11814-2_1](https://doi.org/10.1007/978-3-642-11814-2_1) (Springer Berlin Heidelberg, 2010).

Acknowledgements

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

Author contributions statement

K.P., R.X and R.A. conceived the experiments, K.P. and R.X. conducted the experiments, K.P., R.A., A.C and A.R analysed the results. All authors reviewed the manuscript.

Additional information

Supplementary files: supplementary material and reproduction codes are available at https://osf.io/tp4zj/?view_only=8bd12290d9d4480f95bc2c6e49c73ff7; **Competing interests:** The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [srsupfinal.pdf](#)