Game Data to Algorithm Design: Tackling Combinatorial Problems through Citizen Science

Renata Mutalova

School of Computer Science
McGill University
Montreal, Quebec, Canada

April 16, 2025

A thesis submitted to McGill University in partial
fulfillment of the requirements of the degree of
Doctor of Philosophy

Abstract

Computation methods, while powerful, often struggle with complex problems that demand flexibility, intuition, and creative thinking. In contrast, human decision making excels, in scenarios that demand thinking, identifying patterns and making intuitive leaps—attributes that machines struggle to imitate. By encouraging individuals to apply these abilities through citizen science games, we not only boost engagement but also unleash their potential to contribute to real-world scientific problems.

By analyzing information gathered through community driven projects, like Borderlands Science and Project Discovery in EVE Online in this research paper shows how players' tactics frequently excel or work alongside methods — especially in problems, like Multiple Sequence Alignment and Cytometry Data Clustering. Despite the variety of individual approaches, players display a shared problem-solving logic that allows them to navigate potential solutions more efficiently than automated systems do. In my work, with Imitation Learning techniques I have created models that replicate the strategies used by players and capture the decision making processes that humans apply to tasks. This is a skill that often eludes current computational methods.

In this thesis, I will review algorithms that have merged instinct with a methodology to provide smarter and flexible solutions, therefore providing more efficient and adaptable solutions for complex scientific problems. By integrating human insights into these systems through imitation learning, I have created a powerful synergy where computation benefits from human flexibility, and human input is scaled through automation. This showcases opens up new possibilities for solving bioinformatics challenges and demonstrates the potential of combining human creativity with machine efficiency to push the boundaries of discovery.

Abrégé

Les méthodes informatiques, bien que puissantes, peuvent peiner à faire face à des problèmes complexes qui demandent de la flexibilité, de l'intuition et de la créativité. À l'inverse, les humains excellent dans ces scénarios qui nécessitent de la réflexion, l'identification de motifs et autres problèmes demandant de l'intuition et de la déduction. En encourageant des individus à appliquer ces habiletés à travers des jeux de science participative, il devient possible d'améliorer la capacité de ces jeux à motiver les joueurs tout en débloquant leur potentiel de contribution à la résolution de problèmes scientifiques concrets.

En analysant l'information accumulée par des projets communautaires de ce type comme Borderlands Science et Project Discovery, nous démontrons comment les stratégies employées par les joueurs excellent et se combinent aux méthodes traditionnelles, particulièrement dans le contexte de problèmes difficiles comme l'alignement de séquences biologiques et le clustering de données de cytométrie en flux. Malgré la l'hétérogénéité de ces approches individuelles, les joueurs montrent une approche commune de la résolution de ces problèmes, qui leur permet de naviguer parmi les solutions potentielles plus efficacement que les systèmes automatisés. Dans cette thèse, je présente mes travaux sur l'imitation de ces techniques, dans lesquels j'ai conçu des modèles qui répliquent les stratégies des joueurs et capturent le procédé décisionnel que les humains appliquent à ces tâches, une technique qui est encore très difficile pour les méthodes informatiques.

Dans cette thèse, je présente des algorithmes qui combinent l'instinct humain avec des méthodes d'intelligence artificielle afin d'atteindre des solutions plus intelligentes, efficaces et flexibles. En intégrant la perception humaine à ces systèmes à travers l'apprentissage par imitation, j'ai créé une synergie entre la puissance de l'informatique et la flexibilité humaine, qui ouvre de nouvelles possibilités pour résoudre des défis bioinformatiques et démontre le potentiel de la combinaison entre la créativité humaine et l'efficacité des

machines pour repousser les limites de la découverte scientifique.

Contribution

This thesis contains three research articles, each co-authored in collaboration with several contributors. My specific role, along with the contributions of each co-author, is outlined in the Preface of Chapters 1, 2 and 3. While the detailed scholarly contributions are provided in the prefaces of the respective chapters, they can be summarized as follows:

Part 1

- Developed a machine learning model trained on player data collected from the game Borderlands to solve small puzzles, which represent multiple sequence alignment problems. The model integrates human problem-solving strategies with biological constraints, resulting in improved alignment accuracy. By leveraging the cognitive patterns of human players, the model captures complex biological relationships more effectively than traditional alignment techniques (Chapter 1).

- Conducted a comprehensive comparison between the proposed algorithm and classical sequence alignment methods, evaluating both the accuracy and computational efficiency. The results demonstrate the superiority of my approach in handling real-world biological datasets (Chapter 1).

- Enhanced the model architecture by incorporating an advanced framework based on generative adversarial networks. This approach improves both the efficiency of the alignment process and the model's ability to generalize. By introducing adversarial training, the model adapts to a wider range of biological questions, further boosting alignment performance (Chapter 2).

Part 2

- Designed a novel model for clustering cytometry data, trained using player-generated solutions from the game EVE Online's Project Discovery. The model is tailored to identify complex clusters in biological data, leveraging the problem-solving approaches demonstrated by players. This human-in-the-loop method allows the model to recognize patterns in cytometry data that are often missed by traditional clustering techniques, leading to more accurate and interpretable biological insights (Chapter 3).

- Extracted and analyzed decision-making processes from player data to inform the clustering algorithm. By incorporating human intuition and strategy, the model identifies clusters with greater precision, enhancing the biological relevance of the results and providing new insights into cellular behaviors (Chapter 3).

Acknowledgements

I wish to express my heartfelt thanks to the following people, without placing importance on order:

- I would like to express my heartfelt gratitude to my supervisor, Jérôme Waldispühl, guidance, support, and encouragement were invaluable throughout this journey. I am deeply appreciative of the time, knowledge, and care he has generously offered. His mentorship has been essential to my growth as a researcher, and our collaboration has been a truly enriching experience.

- Audrey Durand and Jin Guo for being part of my committee. I appreciate your involvement and support in this important stage of my academic process.

- Thank you to Alexander Butyaev for always being there and understanding me in my native language. You were the person I could turn to and talk things through in Russian when I needed to open up and share my thoughts. Your support in those moments was truly invaluable.

- I am deeply grateful to my fellow lab members, each of whom has had a direct or indirect impact on my work: Roman Sarrazin-Gendron, Alexander Butyaev, Chris Drogaris, Julien Mounthanyvong, Jay White Bear, Parham Ghasemloo Gheidari, Tim Keding, Eddie Cai, and Elena Nazarova. Your support, insights, and camaraderie have contributed immensely to my journey.

- I would also like to extend my appreciation to the undergraduate student, Lucy, I had the privilege of co-supervising. Their incredible work

and dedication have greatly contributed to this thesis, and I am fortunate to have had the opportunity to collaborate with them. This work would not have been possible without their efforts.

- I would like to express my gratitude to the co-authors of the articles featured in this thesis: Amélie Brouillette, Gabriel Richard, Kornel Erhart, the Borderlands Science Gearbox Development Team, Randy Pitchford, Sébastien Caisse, Rob Knight, and Attila Szantner. Your contributions and collaboration have played a vital role in the success of this work.

- I would like to extend my sincere thanks to everyone who played Borderlands Science and Eve Online Project discovery. Your involvement has been crucial to the development of this thesis, and it would not have been possible without you.

- And I'm deeply grateful to my partner in crime, Aleksandr Berezutskii. He was there for me through the toughest moments, when I was ready to give up and just as present during the good times, celebrating every small win, from article submissions to conference talks and finally finishing this thesis. His support brought light to the hardest parts of this journey and made the good ones even brighter.

Contents

List of Figures

List of Tables

1

Introduction

Many challenges in bioinformatics belong to a category of problems known as NP problems that are extremely difficult to solve efficiently for large datasets due to computational constraints. At the same time, volumes of biological data are accumulating at an accelerating pace, spurring demand for computational efficiencies. Despite using the algorithms and computers available today, it remains difficult to manage the sheer complexity and volume of data, in a significant obstacle when processing and analyzing crucial biological information.

There is no denying how critical biological data is—it is the foundation of modern breakthroughs in health and science. However, the challenge is actually turning this sea of raw information into actionable knowledge. Unless analyzed and interpreted, this dataset is only a potential to uncover crucial information. This therefore underlines the importance of techniques to decipher data; they may change how we approach disease genetics and biomedical research in the future.

While machine learning and computational biology are strengths, there are tasks that are particularly challenging for machines yet relatively easy for humans. This is often because such tasks involve pattern recognition and decision making, areas where humans excel. Therefore, the main challenge lies in encouraging people to contribute time to help solve these problems for the advancement of science.

This is where Citizen Science Games step in. Public bioinformatics challenges allow the participation of hundreds or thousands of players to solve difficult puzzles on online platforms. Players generate data that can be captured and analyzed to extract valuable insights through machine learning.

These games have the potential to crowdsource scientific tasks that are currently beyond the capabilities of available algorithms.

In this thesis, I focus on the problem of aggregating and extracting insights from citizen science data in two key tasks: Multiple Sequence Alignment and Cytometry Data Clustering.

1. The Multiple Sequence Alignment (MSA) challenge is represented as small puzzles in the Borderlands 3 mini-game, Borderlands Science. This game leverages genome fragments of human gut microbes, sequenced by the Microsetta Initiative, and tasks players with aligning sequences. By engaging players in this way, we turn a computationally intense bioinformatics problem into a series of intuitive visual puzzles, enabling the generation of high-quality alignments that go beyond the limits of current algorithms.

2. Cytometry Data Clustering is explored in Eve Online through a minigame where players analyze graphical representations of cell populations. Players contributed to the analysis of protein structures related to SARS-CoV-2, helping accelerate scientific understanding During the COVID-19 pandemic. Players were presented with datasets and asked to classify cell populations, a task that algorithms struggle with due to the complexity of biological data, but one in which human pattern recognition excels.

The central hypothesis of this thesis is that the data generated from players playing these citizen science games can be aggregated to create algorithms capable of emulating human strategies in problem-solving. It is here that, by automating such processes, we create solutions not only capable of imitating human intuition but ones that go on to much more efficiently deal with these complex tasks in bioinformatics.

2

Background

Throughout the course of history many of the problems have been resolved not by individuals but through the combined efforts of numerous people. From ancient civilizations using shared knowledge to solve agricultural problems, to the impressive scientific collaborations of the modern days, the power of collaboration has been undeniable. As we entered into the era of technology and digital, this idea of collective problem-solving took on new forms, with the rise of crowdsourcing being a key development.

2.1 Crowdsourcing

`Crowdsourcing`, the method of engaging a crowd for views and ideas, is a popular approach in this Internet era. This is quite a popular approach in the digital age, offering a versatile framework for tackling a wide range of tasks.

`Amazon Mechanical Turk (AMT)` is an example that demonstrates the advantages of this method. The platform allows individuals, called "Turkers," to tackle tasks that remain difficult for automated systems to handle effectively, such as sentiment analysis, content moderation, and data annotation. By assigning these tasks to a broad participant base, AMT accelerates their completion while utilizing human skills that are hard to replicate through computational methods. The focus here is, on linking knowledge with computer processing power and highlighting the importance of crowdsourcing, in present day problem solving methods [RYHH10, SF08, PCI10, KCS08].

The difficulties associated with `AMT` have been widely acknowledged in research works. One major challenge is the cost involved in expanding projects to handle datasets [MW09]. Since workers are paid per task, costs can quickly escalate, especially when factoring in the need for quality control measures. Another concern is the absence of variety among participants [DFI18]. Many workers tend to be younger and tech-competent, often hailing from specific geographic regions, which can result in a more limited range of perspectives and contributions [IS13]. Furthermore since employees are mainly motivated by rewards, they often prioritize speed over accuracy, which can negatively impact the quality of the data [PCI10]. Finally, `AMT` faces the problem of low long-term engagement — once a task is completed, workers typically have little motivation to stay involved in the project, in the undertaking hindering efforts that necessitate continuous involvement.

2.2 Games and Gamified Systems

However, to overcome some of the limitations seen in traditional platforms like AMT, researchers have turned to gamified crowdsourcing as a more effective alternative [Bra08, MHK16, MHKM17]. Gaming offers motivation prompting participants to tackle problem solving tasks for pleasure challenge and rivalry than solely for monetary rewards. This leads to more thoughtful and high-quality contributions. Moreover, games intrinsically create long-term commitment: players become quite interested in the desire to achieve high scores, pass to new levels, or compete with their opponents and hence remain actively engaged for extended periods of time [CEO14]. It is such long-term engagement which forms one of the high valued resources in scientific projects that necessitates prolonged participation and massive datasets. Furthermore games appeal to a range of players, from backgrounds and demographics thus adding diverse perspectives to the collected dataset. [ESd$^+$17]. By turning complex scientific problems into game-like tasks, scientists can access the creative and problem solving skills of humans yielding more detailed and sophisticated data that aids in advancing scientific knowledge.

2.3 Examples of Success

A typical example of such gamification-based crowdsourcing involves the `Foldit` website [CKT$^+$08], which is basically an online video game-based puzzle where players can tinker with structures of proteins to find better folds. Indeed, protein folding is one of the very hard problems in molecular biology, although, while computational approaches can predict protein structures, many times with good accuracy. Results derived from `Foldit` let players solve such puzzles, and their cumulative efforts have yielded many significant scientific breakthroughs, including discoveries related to the structure of viral proteins, such as those in human immunodeficiency virus. By tapping the power of human intuition and three-dimensional spatial reasoning, `Foldit` has shown that non-experts can usefully contribute to scientific problems that require deep, creative problem-solving [CKT$^+$10, KCT$^+$11, KDG$^+$11].This achievement was recognized when biochemist David Baker, whose innovative approach harnessed gamification as a novel paradigm for solving complex scientific challenges paved the way for further advancements in the field. Baker's groundbreaking use of gamification has been instrumental in advancing protein design and prediction, which lies at the core of his Nobel Prize-winning work [Bak24].

Another successful example is `EyeWire`, which is a game of mapping neural connections within the brain by tracing 3D neuron structures [KGZ$^+$14, HBT$^+$13]. The work requires detailed visual pattern recognition-something that is difficult for machines to handle due to the complexity of the data. Large-scale contributions made by players in `EyeWire` are considered the main source in the development of neuroscience; they map out how the brain is wired, helping people to understand its structure and function. Development of better models mimicking functions of the brain remains indispensable with contributions.

The success of `Sea Hero Quest` has shown the potential of gamification for health awareness and behavioral change. This mobile game was developed to gather large-scale data on human spatial navigation, providing researchers with valuable insights into cognitive decline and early indicators of conditions such as Alzheimer's disease [CSM$^+$18, CMG$^+$19]. Players explore virtual worlds through a series of spatial tasks, generating data on large numbers of people at different ages and backgrounds that are otherwise difficult to obtain.

These projects demonstrate how the community could be involved in the solution of significant scientific challenges.

2.4 Human Abilities in Problem Solving

The human brain – especially visual perception and pattern recognition – is one of the reasons gamified crowdsourcing solves such tough scientific problems. Approximately 25% of the human cortex is dedicated to processing visual information, giving humans an exceptional **ability to recognize patterns, shapes, and colours** — abilities that machines often struggle to mimic accurately [HMT18, POLO23, GMR97]. This natural ability allows humans to quickly identify spatial relationships, detect anomalies, and make sense of visual information in ways that computers struggle with, especially when the data is noisy or incomplete.

For instance, research shows that people can easily tell the difference between groups of objects in two dimensions-something challenging for AI models without extensive preparation or data manipulation [AKDL22, dBTW08, RCZC19]. Such is the case in clustering problems, like the categorization of cells in microscopy images, where human intuition makes one notice subtle differences that might be passed over by algorithms. Research has demonstrated that humans can cluster visual data quite accurately in broad applications with very little training simply by using their innate pattern recognition abilities [BDTSW22].

It is also possible for a human being to recognize patterns and spatial structures that algorithms can't, particularly in cases where the data is unstructured or noisy [Lud23]. For example, in tasks like sequence alignment humans can naturally identify correlations and structures that might elude or be misconstrued by a computational approach because of their capacity to apply context specific comprehension and adjust their methods based on visual hints and changing details.

Humans play a role when engaging in the solving of challenging problems over and above simple calculations, where datasets include intricate or varied data that may not yield to fully automatic or algorithmically efficient processing. This is more so evident in tasks such as visual pattern recognition or sequence alignment, but it is equally so in the area of **combinatorial problems**, since solving such a problem needs something more than the application of mere computation, it needs insight, adaptability and innovation.

2.5 Combinatorial Problems

Data clustering, or the grouping of objects according to needs, is such a combinatorial problem. In general, combinatorial problems are concerned with the valuation of the most valued combination, arrangement, or selection of discrete objects that satisfy some constraints. While not all combinatorial problems are intractable, many combinatorial problems belong to NP-hard problems, whose computational complexity increases exponentially with the number of objects or constraints, hence intractable as scale increases [Kar75].

Common examples include optimization problems (like the *Traveling Salesman Problem*, where the objective is to find the shortest possible route between a set of cities), assignment problems (like the *Knapsack Problem*, where the goal is to maximize the value of items selected within a weight limit), and organizing objects, into clusters based on specific standards.

While for combinatorial problems, humans intuitively and heuristically can easily come up with a solution. For example, for such tasks as the *Traveling Salesman Problem or the Knapsack Problem*, instead of strict rules, humans use perception and flexible strategies to reveal patterns and shortcuts that possibly could not be seen by algorithms [MO96]. In the work of Acuna and Parada (2010), the authors investigated how individuals can solve combinatorial problems including clustering and optimization better than some algorithms. The research conducted by the authors proves that it is possible for human problem solving techniques to outperform algorithms in some instances [AP10].

Moreover, the wisdom of crowds theory by James Surowiecki where he explains how the collective intelligence of individuals can exceed that of experts as well as that of algorithms in solving a particular problem. In combinatorial problems like the *Traveling Salesman Problem* and *Minimum Spanning Tree*, solutions generated by groups of individuals often outperform both individual efforts and standard algorithms [Sur05, YSLD12]. This collective human insight, when applied through crowdsourcing platforms or gamified environments, shows that humans have a natural benefit in solving combinatorial challenges by using adaptive strategies that balance constraints and objectives.

2.5.1 Multiple Sequence Alignment

Bioinformatics is the area concerned with problems that can be approached by different configurations to solve them efficiently. Among those, one of the most important is the multiple sequence alignment (MSA) problem, a basic problem where the aim is to align more than two biological sequences with the purpose of finding regions of similarity that could indicate structural, functional, or evolutionary relationships.

Formally, given a set of sequences $S = \{s_1, s_2, \ldots, s_n\}$, where each sequence s_i is composed of characters from a finite alphabet, the goal of MSA is to insert gaps into these sequences so that they are all of equal length, maximizing a scoring function $f(A)$ over the alignment A, where A represents the aligned sequences. The scoring function typically reflects biological criteria such as matching characters, penalizing gaps, and handling substitutions.

Mathematically, this can be expressed as:

$$A^* = argmax_A f(A)$$

where A^* is the optimal alignment and $f(A)$ is the scoring function that sums matches, mismatches, and gap penalties across all aligned positions.

2.5.2 Computational Approaches for Multiple Sequence Alignment

The challenge with MSA becomes more intricate as the number of sequences and their lengths grow, leading to an increase in potential alignments to consider carefully. One of the foundational methods is *Dynamic Programming (DP)*, exemplified by algorithms like the *Needleman-Wunsch* [NW70] and *Smith-Waterman* [SW$^+$81] algorithms. These techniques use optimal pairwise alignments to generate globally optimal alignments between sequences. While dynamic programming offers solutions for problems at hand such solutions become costly in terms of computation as the quantity of sequences and their lengths grows because of its quadratic or cubic time complexity. For instance, aligning dozens of long DNA or protein sequences using dynamic programming becomes computationally prohibitive as the number of required calculations grows exponentially with the dataset size.

Besides this increasing difficulty, an exact solution of the MSA problem is viewed to be NP-hard. In other terms, this problem category represents problems in which computational effort drastically increases together with dataset growth. Accordingly, finding its exact solution has a low order of feasibility regarding time for large instances of the input datasets; generally speaking, scholars can only attain optimal or heuristic results within decent time limits using fast heuristics or approximating methods of alignments. These methods can only find near-optimal solutions, while possibly sacrificing some accuracy for highly divergent sequences.

A more general and profound problem in MSA lies in defining the most appropriate score, for evaluating the alignments, while defining the quality of an alignment has to be based on specific criteria, such as similarity of sequences or biological relevance. In many cases, the ideal scoring system, or at least a difficult-to-precise-define one, may not exist.

Therefore, in practice, proxy systems are often used, such as substitution matrices like PAM or BLOSUM [TN20], which approximate biological relationships between sequences. However, these proxies are not universally applicable and can introduce biases, especially when applied to diverse datasets. In the absence of a gold standard, the results of different alignment methods vary, and thus, comparing methods or validating results across different studies is difficult. These challenges have to be put into a wider context by realizing the inherent trade-offs within the methods of MSA. The computational resource issues preclude exact solutions in most cases, while the lack of a commonly accepted scoring system further complicates the evaluation of alignments. These are some of the important issues that need to be taken into consideration in developing robust and scalable algorithms which balance computational efficiency with biological relevance; at the same time, these also bring forth the need for improving proxy scoring methods or novel ways of developing an evaluation system.

To address the scalability issue, *approximation methods* such as *ClustalW* [THG94] and *MUSCLE* [Edg04] have been proposed. Indeed, these approaches are based on effective heuristics, either in the progressive alignment or iterative refinement directions, to compute MSA rapidly. Although these methods reduce computation time significantly, they do not ensure global optimality, while their accuracy might be inconsistent in the case of highly diverged sequences. Other heuristic methods that have been investigated are *simulated annealing* [KPC94], which relies on an iterative search for optimally aligned structures while progressively reducing the size of the search space; optimization-based approaches, including *particle swarm optimization* and *ant colony optimization* [RK03, XC09, MJ03].

Probabilistic Models, including *Hidden Markov Models (HMMs)* [Edd98] and profile-based alignments, use statistical frameworks to model sequence relationships and improve alignment accuracy by leveraging probabilities for matching, inserting, or deleting characters in sequences. These methods enhance alignment accuracy by considering probabilities for the match, insertion, or deletion of characters in sequences. These models require large amounts of training data and can be computationally intensive for larger datasets.

More modern approaches have also emerged in recent years, such as **MAFFT** *(Multiple Alignment using Fast Fourier Transforms)* [KMKM02] and *MUSCLE (MUltiple Sequence Comparison by Log-Expectation)* [Edg04]. MAFFT employs a progressive alignment method, using Fast Fourier Transform to cluster sequences quickly in the initial step, while MUSCLE begins with a rapid draft alignment followed by iterative refinement stages to optimize alignment accuracy. Specialized software like **PASTA** (Ultra-Large

Multiple Sequence Alignment) [NMKW15] has been designed for handling extremely large datasets by combining divide-and-conquer strategies with progressive alignment techniques.

Even though these approaches have improved, a concern over slow convergence and alignment accuracy makes it apparent that there's still much work to be done to make MSA better – especially with with large datasets.

Recently, reinforcement learning (RL) methods have been applied to tackle the MSA problem. RL is particularly appealing for its ability to explore large search spaces and learn policies that optimize the alignment process over time. In 2014, [MBD14] pioneered the use of *Q-learning*, a classic RL approach, to solve MSA. Their method involved representing ordinal numbers as RL states and learning to predict the optimal sequence alignment location as an action. Although this method showed promise, it faced limitations related to slow convergence due to a brute-force style search process that included many invalid operations.

To improve performance, researchers have employed *Deep Q-Networks (DQN)* for MSA, as seen in the work of [JJKR19]. By incorporating *Long Short-Term Memory (LSTM) [HS97]* networks, their DQN model could better process DNA sequences through sequential computation, allowing the RL agent to retain information over long sequences.

Another significant contribution came from [RSB18], who developed *RLALIGN*, an actor-critic architecture using **Asynchronous Advantage Actor-Critic (A3C)* and convolutional neural networks. Compared with the other existing RL approaches, RLALIGN would directly leverage the currently current alignment state in predicting the next action and representation for a nucleotide to move to a direction. Yet, when increasing the sequence number, there is a exponential increase of dimensions over the state space, creating convergence difficulties.

While progress is being made in the use of RL for MSA, a number of challenges do remain, particularly in the scaling of these models to tackle larger datasets and longer sequences. Indeed the paper [Joe21], has shown that RL models perform really well while aligning sequences of comparable lengths; however, they often fail while aligning longer or heavily diversified sequences, where the traditional algorithms, including Clustal and MUSCLE, are still unbeaten by the RL methods. However, techniques like *Negative Feedback Policy (NFP)*, introduced by proposed by [ZZL$^+$22] have been effective enough to improve the convergence of an RL model by incorporating sequence profiles into the alignment process.

Some researchers have been working on simplifying the complexity problem by creating RL techniques specifically designed for *pairwise sequence alignment*. For example, [SJS+21] introduced **DQNalign**, which relies on pre-processing techniques such as *Clustal* and *MUMmer* to break down MSA into iterative pairwise alignments, reducing computational time and memory requirements. Similarly, [LT23] developed **EdgeAligs**, an efficient RL model for pair-wise sequence alignment on embedded devices using a sliding window. While these methods help in understanding the issue from an easier-to-understand perspective, it is really important to consider the view for enhancing effectiveness and speed when dealing with extensive MSA assignments.

2.5.3 Flow Cytometry Clustering

Another problem in bioinformatics research is the flow cytometry data clustering task. Generally, flow cytometry is a technique employed in analysis, chemical properties of cells or particles in a liquid run them through a laser. The information coming from flow cytometry studies is multidimensional since every cell is represented as a vector in a space whose dimensions correspond to a given attribute, size, textural features, and identification of specific markers.

The objective in flow cytometry data analysis is to cluster cells into biologically meaningful groups based on their measured properties. Formally, given a set of cells $X = \{x_1, x_2, \dots, x_n\}$, where each cell x_i is a vector in R^d representing the d-dimensional feature space, the goal is to assign each cell x_i to one of k clusters $C = \{C_1, C_2, \dots, C_k\}$, where $C_j \subseteq X$ and $\cup_{j=1}^{k} C_j = X$, such that cells within each cluster are more similar to each other than to those in other clusters.

The clustering process can be represented as the following optimization problem:

$$\min_{C} \sum_{j=1}^{k} \sum_{x_i \in C_j} \|x_i - \mu_j\|^2$$

where μ_j is the centroid of cluster C_j, and $\|x_i - \mu_j\|^2$ is the Euclidean distance between cell x_i and the cluster centroid. The goal is to minimize the total within-cluster variance, leading to compact, well-separated clusters.

Clustering in flow cytometry is challenging because of the complexity of the data with its dimensionality and noise levels. Traditional clustering algorithms, such k-means [Llo82] or hierarchical clustering [Joh67], may not perform well on large and noisy data sets and may fail to identify the real biological clusters.

2.5.4 Computational Approaches for Flow Cytometry clustering

Different computational approaches have been employed for flow cytometry data analysis; however, all have, to a certain extent, surmounted the problems related to big, noisy, and complex cellular data in many ways. Conventional algorithms used for clustering, some graph-based techniques, machine learning, and statistical methods form parts of such attempts.

As classically clustering methods are very straightforward, a variety of classical methods were conventionally applied to the flow cytometry data. A two-stage clustering method couples Fuzzy-C-Means with Markov clustering identifies the cellular subtypes and detect anomalies effectively in various real-world datasets [PJBN16] with high F-measure value greater than 91%. But as it heavily relies on the fuzzy-C-Means, has limited capabilities handling complex structure of the data in high overlapping cases. Another potential defect the method may possess is the sensitivity to the initial parameters, which can eventually cause variations in the result regarding clustering.

In clinical settings, the MegaClust Algorithm is used for flow cytometry data, which enables effective recognition of immune cell subsets [BSH+21]. For detecting rare populations, it would be very effective since its tuning is dynamic to count the events. However, like for any other method, parameter tuning is needed-for instance, the minimum number of events to detect a cluster-so optimization of the findings across different datasets is hard.

It might be extended by several different advanced techniques and methods, including machine learning, in which deep learning approaches-Deep Convolutional Autoencoder-based Clustering (DCAEC) being one of them-have proven to be a very efficient tool. The results clearly indicated the impressive performance of the DCAEC model in differentiating between white blood cells of healthy ones and those suffering from leukemia with 97.9% accuracy [ZCT+23]. However, these models have a number of healthy and leukemic white disadvantages, including the requirement of big datasets for training and black-box approach, hence not the most wanted in clinics when understanding the process of decision-making is important.

Statistically, the *LAMBDA Framework* applies Bayesian data analysis for model-based clustering [AMM+20]. It is highly efficient in the detection of cell populations that are not yet known but may be computationally expensive. The framework requires carefully tuned parameters, and there may be some scalability limitations for large data sets.

Another recent work proposes a method of unsupervised clustering that aligns the histogram features of the two datasets and seeks to find maximum mutual information between the two [UOPS23]. This technique has improved accuracy in clustering by emphasizing features that are informative; thus, it is more suitable for subtle cell population differences. Similar to other feature-guided methods, its efficiency depends highly on appropriate selection of features, which can be a disadvantage in data with high variability.

These methods will always show the balance between accuracies, scalability, and computational complexities while solving a flow cytometry clustering problem. None of these techniques is devoid of limitations; the areas that need further developments are indicated.

Human intuition, especially, plays a key role in complex combinatorial problems like flow cytometry clustering or MSA, where the focus is on patterns, subtle groupings, conserved regions, and meaningful gaps in high-dimensional data. In each of these contexts, human problem-solving augments algorithmic approaches through the uncovering of biologically relevant patterns that algorithms may overlook, particularly when noisy or incomplete data are dealt with. With crowdsourcing or gamified platforms, for instance, researchers may more intuitively tap into human pattern recognition to guide the process of clustering or alignment. The combination of such a human insight with computational methods can help in dealing with intrinsic challenges presented by the combinatorial nature of such tasks in general, which would translate into higher accuracy, better alignment quality, and more meaningful interpretation of results in bioinformatics.

2.6 Citizen science games and bioinformatics

As mentioned in previous sections, bioinformatics has successfully applied gamified crowdsourcing to a wide variety of combinatorial problems. That provides an innovative way to utilize the human problem-solving capability for scientific discovery. Three of the most salient examples involve complex optimization tasks in molecular biology: `Foldit`, `EteRNA`, and `Phylo`. These provide very good illustrations of how such gamified platforms can leverage human problem-solving abilities to solve very intricate scientific problems that are difficult for algorithms to solve alone.

2.6.1 Borderlands Science

Besides these examples, the latest major citizen science applications that have emerged are video gaming-such as `Borderlands Science` and `EVE Online's Project Discovery` - where human players will create valuable scientific data during meaningful game play.

`Borderlands Science`, embedded in the popular video game `Borderlands 3`, has engaged over 5 million players to address the multiple sequence alignment problem. In the game, colored tiles are aligned that correspond to DNA sequences; these will help in solving a real-world genomics challenge. Since its launch in April 2020, it has generated above 150 million alignments, many of which provided valuable human-driven solutions with the potential to help researchers improve the accuracy of DNA sequence analysis. This effort uses the ability of players to recognize patterns and hence allows insight into bioinformatics problems that algorithms cannot capture, extending large-scale genomic research.

2.6.2 EVE Online's Project Discovery

Similarly, `EVE Online's Project Discovery` integrates citizen science into the massively multiplayer online game *EVE Online*, which boasts over 24 million registered users. In one iteration alone, more than 327,000 players classified cell populations from flow cytometry experiments and thus analyzed real-world scientific data. By identifying patterns and clusters in this complex, high-dimensional data, the players contributed to solving the challenge of grouping cells based on their physical and chemical characteristics. By 2021, over 40 million classifications had been submitted and provided researchers with data for a better understanding of the cell behaviour and population structure. This gamified approach solves a problem that current algorithms cannot by embracing human pattern recognition, which in turn produces serious scientific output.

Both games actually gamify complex scientific activity in an effective way, hence making bioinformatics problems very challenging yet accessible and agreeable to non-experts. Sequence alignment in the case of `Borderlands Science` or cellular data classification in `EVE Online's Project Discovery` comes as a fun puzzle while being a part of that game's universe. These examples give an idea of how scientific research is transformed into appealing experiences that can be used by the general public to contribute to critical research areas through well-designed games. By converting complex problems into approachable tasks, these platforms tap into the collective intelligence of non-experts, fostering widespread participation in scientific discovery through enjoyable gameplay.

First of all, the first time I started this project, I had already gathered so much data contributed by players, apparently very good to contribute to complex scientific problems. These are like little puzzle pieces embedded in a game, and producing them requires large-scale human effort. From here on, I began exploring how to take this crowdsourced, decentralized decision-making and turn it into a more unified, increasingly automated system. This goal means transforming many small, individual decisions made by players in scientific games into one system that can work automatically, without needing people to guide it every time. These player decisions are often useful but scattered and inconsistent. To make them work together, we need to collect and organize this data, and then use it to train computer models that can learn from how people solve problems. One way to do this is through imitation learning, where a system learns to copy human actions. By doing this, we can build models that not only give answers but also learn to "think" more like the human players who are good at spotting patterns in complex data. This kind of automation is important because it lets us apply human insight at a much larger scale - helping scientists analyze difficult biological data faster, more accurately, and without always needing a person in the loop.

The main possible solution to this problem can be in the area of *imitation learning*. With it, I would be able to learn models from this approach and imitate players' decisions. In general, the key behind imitation learning is that the model has to observe and imitate experts in order for it to be able to get closer to detailed strategies and solution-finding approaches. Unlike traditional approaches, the agent learns in a trial-and-error manner, interacting with an environment that provides rewards or penalties to the agent in; this is what's known as reinforcement learning. Imitation learning circumvents these challenges by skipping the need to design a reward function, instead directly duplicating human behavior. This is especially useful in complex environments wherein defining a reward function or exploration to learn optimal behavior is hard or computationally expensive.

Imitation learning draws its roots from both behavioral biology and some early works in robotics. From the perspective of behavioral biology, imitation has been said to be a basic mode of learning both in man and animals. Many forms of learning in both animals and humans at young, tender ages are normally recorded through observation and subsequent replica of such actions

[CF19]. This led to the inspiration of natural learning through observation and copy, which is known as imitation learning.

One of the very early and distinct findings of the use of imitation learning implemented on a robot came through: *Autonomous Land Vehicle in a Neural Network (ALVINN)*, developed by Pomerleau in 1989 [?]. ALVINN made use of learning through a neural means where it drove an unmanned ground-moving vehicle having experience exclusively as to how to drive a vehicle like itself yet mimic the same driving habits as reflected by the human. It had been trained on human-labelled data, i.e., pairs of visual inputs - road images - and corresponding steering actions; hence it learned a policy to steer the vehicle based on visual input. This pioneering work paved the way for what is known as *behavioral cloning*, a particular kind of imitation learning.

Other state-of-the-art applications for imitation learning, AlphaStar was trained, first and foremost, through the use of imitation of human games extracted from professional competitions [Aca23]. Indeed, this application has considered hundreds of thousands of human games where a general idea regarding strategies could be observed: micro-management or a decision-making profile. DeepMind has claimed the above two reasons to be responsible for a 95% win rate by AlphaStar against top-ranked players of the StarCraft II European server with regard to their paper: first, initial policy learned from human demonstrations gave an excellent baseline; second, after being further fine-tuned by self-play, surpassing human performance by an enormous margin seemed pretty much effortless [WSQ$^+$21, Dee19, Dee20].

Approached from another direction, various autonomous driving technologies have effectively drawn upon imitation learning, sometimes referred to as imitation with demonstrations. Waymo Self-Driving Car Technology used imitation learning. The work of Waymo is credited with one of the flagship milestones of training a model using more than 10 million miles of driving developed by human drivers [Way21]. This was possible by training first how the human drives, enabled from the collected dataset of approximately six years. Imitation learning yielded modeling driving behavior quite similar to that of humans, and brought a sharp improvement in road handling for the vehicle from diverse conditions and interactions on the road. Indeed, studies have shown that models trained on imitation learning attained a reduction of 30% regarding intervention rates when compared to traditional systems based on handcrafted rules. Imitation learning improves generalization in real-world, even complex urban environments; the performance of the autonomous system was shown to operate at a human level.

OpenAI applied imitation learning to robotic manipulation tasks such as solving a Rubik's Cube with a robot hand. In that respect, imitation learning allows foundational manipulation strategies to be endowed upon the system through structure human demonstrations [ZB24]. These strategies were further optimized using reinforcement learning and domain randomization. Under normal conditions, the system could solve the Rubik's cube with a success rate of 60%, and was even able to adapt to perturbations such as modification of the cube's physical properties or adding distractions in the environment. This result, considering only about 30 hours of training data from human demonstrations, constituted a demonstration of how imitation learning can effectively bootstrap complex motor control tasks in high-dimensional environments.

Imitation learning generally employs two main methodologies: *behavioral cloning*, and *inverse reinforcement learning*. Both methods aim at learning and replicating the expert's strategy and present different ways to approach tasks that require human expertise.

2.6.3 Behavioral cloning

In **behavioral cloning** [**?** BDTD$^+$16**?** , ZMJ$^+$18], the task of learning from human demonstrations is treated as a supervised learning problem. The goal is to train a model to imitate a human's policy π_H by minimizing the difference between the actions taken by the human and those predicted by the model. Let $D = \{(s_1, a_1), (s_2, a_2), \dots, (s_n, a_n)\}$ represent a set of demonstrations, where s_i is the state and a_i is the action taken by the human in that state. The objective of behavioral cloning is to learn a policy π_θ, parameterized by θ, that minimizes the error between the predicted action $\pi_\theta(s)$ and the demonstrated action a.

This can be formalized as:

$$L(\theta) = E_{(s,a)\sim D}[\ell(\pi_\theta(s), a)]$$

where $\ell(\cdot, \cdot)$ is a loss function depending on the nature of the task. The model is trained to minimize this loss by adjusting its parameters θ so that the predicted actions align as closely as possible with the human's demonstrated actions.

It has found effective applications in behavioral cloning for autonomous driving [BDTD$^+$16], robotics [ZMJ$^+$18], and video game AI, whereby machines can tap into expert human knowledge without having to worry about designing complex reward structures. This simplicity and efficiency are the reasons why it has found so many successful applications in domains where the environment provides the setting where human expertise can be modeled directly [**?** VBC$^+$19].

2.6.4 From Supervised Learning to Behavioral Cloning: The Shift in Objectives

What actually sets behavioral cloning apart from classic applications of supervised learning is not the learning mechanism itself, but rather the context in which it is applied and the type of behavior it aims to capture. At first glance, both rely on labeled data, yet their goals and operational dynamics diverge meaningfully. Supervised learning is traditionally designed to predict outputs by searching for patterns among static, independent samples—a methodology that serves very well in tasks that are, by their nature, classification or regression [GBC16]. But what if the task at hand requires reproducing human decision-making in dynamic, sequential environments? This is where BC becomes relevant — still supervised, yet operating under different assumptions and objectives.

BC seeks to replicate expert strategies by learning from trajectories, or sequences of state-action pairs indicative of decision-making over time [Pom91]. Unlike the standard i.i.d. assumption in supervised learning, BC must account for the fact that actions are interdependent, and their consequences propagate over time. For example, small mistakes made early in a trajectory can cascade due to prediction errors, gradually steering the system away from desirable behaviors. This well-known issue, referred to as covariate shift, illustrates how even a methodologically supervised process like BC can face challenges well beyond typical supervised settings [RGB11].

The distinction also deepens when considering the philosophical goals of these approaches. While supervised learning focuses on optimizing predictive accuracy, BC aims to approximate the often subtle, implicit logic of human expertise. Human demonstrations are more than just labeled data—they embed context, foresight, and the ability to adapt to changing conditions. BC excels in settings where these elements are critical, enabling the extraction of heuristics and strategies that defy purely statistical modeling.

In this light, behavioral cloning is best understood as a supervised learning technique applied within an imitation learning framework. Its methodology is familiar, but its function is tailored to capturing and reproducing expert behavior in high-dimensional, sequential environments. While supervised learning builds systems that learn from examples to generalize predictions, BC builds systems that learn from demonstrations to generalize strategic behavior—preserving the structure of human decision-making even under uncertainty. This difference does not negate BC's identity as supervised learning, but rather highlights how its goals stretch beyond standard predictive modeling into the realm of behavioral replication.

2.6.5 Generative Adversarial Imitation Learning

While generally effective, however, behavioral cloning can suffer from compounding errors if the encountered states shift from those glimpsed in training, the accumulation as the model acts in ever more unfamiliar ways leading to performance decay, commonly seen. Such limitations overcome the development of the following, more sophisticated ones: *Generative Adversarial Imitation Learning (GAIL)* that couples, in a consistent way, ideas of Imitation Learning with the adversarial training framework provided by *Generative Adversarial Networks (GANs)* [GPAM+20].

GANs, first proposed by Ian Goodfellow [GPAM+20], are a deep accomplishment in machine learning. They are composed of two neural networks: the generator and the discriminator. These networks compete against each other in a process. The generator tries to generate realistic data while the discriminator assesses the originality of such data against real data. This goes on until the generator learns to fool the discriminator, which in turn generates much better data. In an adversarial framework, this encourages the generator to come up with realistic outputs. Originally, GANs were first applied to image generation but are now applied to a wide variety of domains, including natural language processing and bioinformatics [MAP23, Tut23, ASF23]. The ability of GANs to model complex data distributions has made them very crucial in domains requiring data synthesis.

Building on the concept of GANs, GAIL emerged as a method for imitation learning, introduced by Jonathan Ho and colleagues in 2016 [HE16]: a generator G and a discriminator D . The generator G, representing the learned policy π_θ, aims to produce actions that mimic the human expert's actions. The discriminator D is trained to distinguish between actions taken by the expert and those generated by the policy π_θ. The goal of the generator is to fool the discriminator by producing actions that are indistinguishable from those of the expert.

Formally, the GAIL objective is to minimize the following adversarial loss:

$$\min_{\pi} \max_{D} E_{\pi_H}[\log D(s,a)] + E_{\pi_\theta}[\log(1 - D(s,a))]$$

In this formulation, π_H represents the expert policy, and π_θ is the learned policy. It trains the generator to produce actions that minimize the discriminator's capability of differentiating between human and machine-generated actions while training the discriminator to improve its capability for distinguishing between the two.

Imitation learning, especially behavioral cloning and some of its extensions like GAIL, provides a generic yet powerful paradigm for teaching machines to emulate human behaviors. By learning from expert demonstrations directly, these techniques avoid the pain points involved with reward design and explorations of the environment so that models often yield effective, scalable solutions to domains reliant on human intuition and problem-solving [HE16, OPN+18, ACVB09].

2.6.6 From Patterns to Policies: the Shift from GANs to GAIL

Having already explored GANs and their adaptation into generative adversarial imitation GAIL, it becomes essential to focus on the critical differences between the two. Both frameworks share a common foundation—a generator and a discriminator engaged in adversarial training—but their objectives and applications diverge profoundly. While GANs were designed to model static data distributions, such as images or sequences, GAIL shifts this paradigm to target the replication of human decision-making strategies in dynamic, sequential environments [GPAM+20, HE16].

The key difference is related to the nature of the generated output. GANs generate output that follows the distribution of the data they have been trained on, whether that is generating realistic images or samples of text. The fidelity of a dataset is the objective, and the discriminator works to make sure that generated outputs resemble examples labeled. On the other hand, GAIL does not seek to generate static data but rather attempts to mimic expert behavior. GAIL trains a generator to generate sequences of actions or policies while the discriminator evaluates if those policies are aligned with what was shown by the human expert and not necessarily similar data.

This difference in distinction changes all things in how the training and deployment of the models are carried out. GANs operate within a statically bound dataset where the generator iterates to refine their outputs against a fixed discriminator. Whereas in GAIL, it introduces one into a feedback loop with an environment: a generator must dynamically change its policy based on what happens in response to the actions. This embedding into a reinforcement learning framework implies that, instead of merely generating samples, GAIL is indeed learning policies in order for it to attain expert performance. Success in GAIL thus requires the replication not of patterns but of the intent and adaptability underlying human decisions. By targeting behavior instead of static outputs, GAIL overcomes challenges that GANs can't. It performs great in domains where the context and sequence matter; for example, demonstrations by a human encode not only what to do but also when and why to do it. For instance, GAIL is better suited for robotics, simulation tasks, or combinatorial problems where strategies evolve by considering the environment. That's a deep shift from generation of data to generation of intent, and one that speaks to a far greater possibility with general imitation learning in capturing rich detail in human expertise than is possible with GANs alone.

2.7 From Players to Algorithms

Open questions regarding the interaction of combinatorial problems and collective human problem-solving are still left untreated by literature. Whereas substantial advances have been achieved in computation methods, some combinatorial tasks still resist efficient solutions. Evidence from many practical and experimental settings does demonstrate that the very same problems can often be satisfactorily solved by players playing games. This points to a curious, under-explored dynamic between human decision-making and combinatorial optimization.

Although this connection between human problem-solving in games and combinatorial tasks has not been directly researched in depth, the deeper I delved into this area during my research, a variety of ideas and hypotheses came up, some of which were supported or challenged by the literature available.

- One of the most salient points of this discussion is that a large number of crowdsourcing platforms are already available, a significant number of which can be repurposed for solving some of the major bioinformatics challenges. These crowd-sourced initiatives utilize collective intelligence to bring forth new strategies for dealing with complex biological data, much of which cannot be meaningfully processed through traditional computational approaches. Section 3 Crowdsourcing expands on this idea.
- It soon became clear that there is a significant volume of current cit-

izen science projects focused on bioinformatics work, underlining the important contribution that can be made by the public toward solving some of the major challenges in this area. Such projects have proved exceptionally valuable, especially when there is a complex problem to solve which benefits from multiple contributors and insights. This will be discussed in more depth in Section 2.6, *Citizen Science Games and Bioinformatics*.

- Another important revelation here is that, intuitively, humans have this 'pattern-recognition' ability in data. A pattern recognition that quite often gives way to decision-making substantially different from, and sometimes superior to, the available traditional computational methods. The human players rely on intuitive thinking that allows them to find connections that algorithms do not. This points to a special link between citizen science games and bioinformatics: it indicates that human intuition and pattern recognition can be good tools for combinatorial tasks, and that such approaches can sometimes outperform the standard computational methods in a way.

- Yet another important notion is that humans have an innate aptitude for pattern recognition in data. This often leads to decision-making strategies quite different from, and sometimes superior to, those emerging from traditional computational approaches. Human players seem to rely on intuition that guides them toward connections which algorithms may not catch. This concept is discussed in Section 2.4, *Human Abilities in Problem Solving*, and further explored in Section 2.6, *Computational Approaches*.

These reflections have led me to ask new and important questions, which have shifted my approach toward the main issues in this research. By emphasizing the unique link between citizen science and bioinformatics, this work has deepened my understanding of how human problem-solving can complement computational biology. These can also be combined in a manner that leverages the strong points of both computational and intuitive decision-making as demonstrated through human players to devise hybrid approaches which marry human insight with algorithmic efficiency. The results from this work not only attempt to address some current challenges but also point to some very promising directions for the future in computational biology and human health.

Multiple Sequence Alignment

3

Playing the System: Can Puzzle Players Teach us How to Solve Hard Problems?

Preface

Background reflects upon points of intersection between combinatorial problems and their application within computational contexts; particular attention is paid to one of these - a Multiple Sequence Alignment problem - implemented within the game of Borderlands. In this paper, I introduce a behaviorally cloned algorithm on data from the game of Borderlands. As it is shown in this paper, this method outperforms the classical methods of sequence alignment in many aspects; hence, it brings a new perspective into the solving of alignment problems.

The core contribution of this work is in its novel application of behavioral cloning to multiple sequence alignment, using gaming data in ways not explored to date. This chapter, except for the present preface, was presented at the 2023 CHI Conference on Human Factors in Computing Systems.

The full list of authors is Renata Mutalova, Roman Sarrazin-Gendron, Eddie Cai, Gabriel Richard, Parham Ghasemloo, Sébastien Caisse, Rob Knight, Mathieu Blanchette, Attila Szantner, and Jérôme Waldispühl.

Eddie Cai created the post-analysis of the result and the first draft of the optimal solution searching algorithm. Roman Sarrazin-Gendron provided the expertise for the RNA analysis and made contributions to the writing.

Jérôme Waldispühl supervised the project, provided financial support, and offered guidance throughout its development. Gabriel Richard and Sébastien Caisse did the implementation of Borderlands Science.

Key findings of this research are given in **Figure 4.5**, showing the main methodology; **Figure 3.5** and **Figure 3.6** and show the results compared to classical approaches for solving the MSA problem.

The contribution of mine in this work involves dataset preparation, formatting for training, designing model architecture, and implementing all the benchmark methods. I also performed the post-processing of data and wrote most of the text. In addition, I ran all experiments and calculated the results presented in this section. This is the prepublication version of the following article: Published as: Playing the System: Can Puzzle Players Teach Us How to Solve Hard Problems? Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, pp. 1–15. The version you are reading is very close to the published version.

Introduction

In 2019, Americans spent an average of over 37 minutes a day playing video games and has only increased since then [Edd22]. The time invested in gaming is usually seen as pure entertainment, like watching a movie, but what if the efforts of the players could be harnessed to accomplish useful tasks? This is the objective of "games-with-a-purpose", which fill this gap and pair entertainment with a computational task such as image labelling [vAD04a, VA06] .

Nearly every game can be summarized as a series of tasks that the user has to solve. In many situations, players display inventiveness and creativity in solving these tasks [MCS17, ESB+12, ON20]. Typically, games are designed to optimize the enjoyment of the user [CSBG18, KGL+98]. However, when these tasks can be adapted to solve real-world problems [VAD04b], games become powerful vehicles to access the millions of hours of active problem-solving efforts conducted every week by the players [WSK+20].

Citizen science efforts, which aim to involve non-scientists in the solving of scientific problems, often struggle with finding enough participants to address large-scale problems. The potential contribution from the gaming community could bring a major paradigm change in the type and scale of problems that can be solved through this type of approach.

Enter Science Discovery Games (SDGs). These games are specifically designed to help with analysis of scientific data and assist research projects. They rely on the intuition of humans (i.e., the players) to solve computational problems that are challenging for computers because of their complexity [CKT+10, KRK+12] or because the solution is based on human perception and agreement [PAD+17, PKY+21, BDTSW22].

Among the types of important problems that classical algorithms struggle to solve are NP-hard problems, for which no polynomial time algorithms are known to exist. Many classical bioinformatics problems are NP-hard [WJ94], including the multiple sequence alignment problem. This problem, which has applications in fields ranging from biology to linguistics, is particularly difficult due to the lack of a well-defined ground truth, leading to many approaches being parameter-sensitive and difficult to generalize, or too inefficient for use on real-world problems. Many approaches have been proposed to tackle this problem [CMC+16], such as phylogeny-aware methods [LG08] or non-coding multiple sequence alignment [SS13], and new methods are regularly published, because there is still significant room for improvement. In particular, there is room for human contribution, because the lack of a well-established optimal target state is less of a challenge for citizen scientists than it would be for deterministic algorithms.

The first citizen science game to tap into this potential was Phylo (2010) [KRK$^+$12], a sequence alignment puzzle game in which players could improve computer-generated alignments in a lightly gamified environment. This project showed that humans could significantly improve alignments, but the limited gamification led to the game being heavily expert-dominated, which reduced the opportunity to learn how the average player plays an alignment game.

In order to bridge that gap by targeting a much wider group of potential players, Borderlands Science (see Figure 4.1) was released as a much more gamified mini-game inside Borderlands 3, a shooter-looter game played by millions of players [Sza16]. To attract a more general audience more used to fast-paced games and less familiar with puzzles, the methods from Phylo were adapted to simplify and speed up puzzle solving. With over 3 million participants and 100 million games played, and a relatively even contribution across individual players, we now have enough data to thoroughly investigate *how* human players solve these puzzles, and how this can be leveraged for the broader problem of sequence alignment.

3.0.1 Game design

In Borderlands Science (Figure 4.1), the player is shown 7 to 12 columns of bricks. Each column represents a homologous DNA sequence fragment (note that this representation is the transposed version of the typical multiple alignment representation where sequences constitute alignment rows). The number and length of sequences increases with the difficulty level. Each individual brick represents one of the four types of nucleotide bases in DNA that are encoded with a specific tile and color. Each puzzle in Borderlands Science is made of fragments of microbial 16S ribosomal RNA gene sequences provided by the American Gut Project [MHD$^+$18]. In the initial configuration of the puzzle, the bricks are piled up at the bottom of the screen (see Figure 4.1), as though they were under the effect of gravity.

The player is provided with a limited number of gap tokens, which they can insert between bricks to maximize the alignment of bricks against the guides on the left. These guides provide insight about the overall context of the alignment region the sequence fragments come from [WSK$^+$20]. The cost of adding a gap is conveyed to the player as the resistance to the gravity effect.

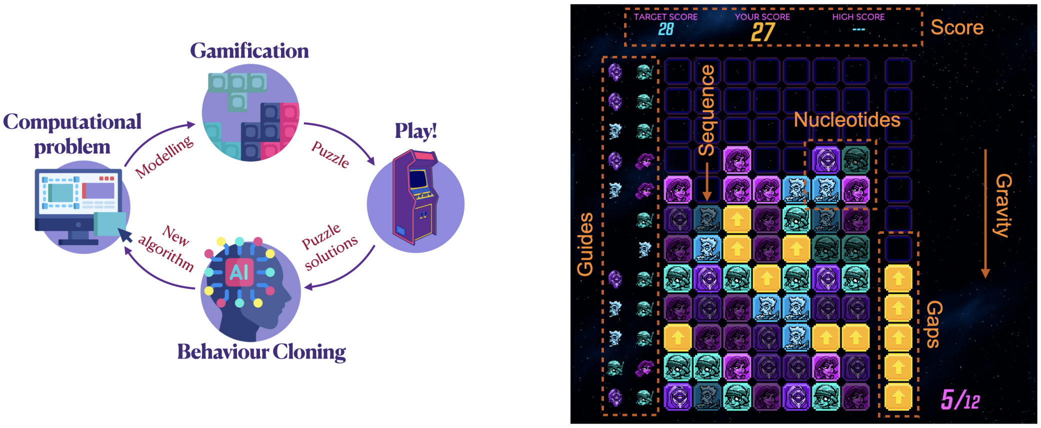

Figure 3.2: Borderlands Science interface: The sequences to align are presented vertically with tiles of four colors representing the four nucleotides A,C,G,T. On the left, the guides represent the tiles to match in each row to collect points. The user must insert gaps (i.e., yellow tiles) to maximize the reward. The number of gaps is limited and full rows get an extra bonus.

The main challenge of the game is to reach a target score which is the integer that is above the nominal score set by the naive greedy player in the face of limited number of gap tokens and move on to the next puzzle. We also display the highest score previously submitted to push the participants to optimize that score. (See official trailer at `https://youtu.be/L_mH6Ak_Ny0`.)

3.0.2 Hypotheses

The strategies employed by the players to achieve this goal appears to be heterogeneous without obvious pattern. Are these strategies efficient (i.e., the solution simultaneously maximize the score and minimize the number of gaps used)? Can we learn from the the collected solutions new heuristics to solve the puzzles?

This article explores the potential of crowdsourcing human intuition to tackle NP-hard problems, focusing on the typical problem of multiple sequence alignment. We attempt to formalize, extract, and reproduce the players' puzzle-solving strategies to show they match an effective and reproducible strategy. By establishing the meaningfulness of the solutions submitted by the players, we open an avenue for further work in human-inspired multiple sequence alignment algorithms, especially as new developments in reinforcement learning and Transformers-based methods unlock new possibilities for exploiting crowdsourcing results. We hypothesized that the wisdom of crowds could be an effective solution for the multiple sequence alignment task in particular, and for other types of NP-hard problems more generally. Specifically:

H1: Human solutions use complex strategies that are not replicating basic heuristics. H1 aims to tests whether the strategies used by the participants follow simple and intuitive rules. To address this question, we device several bots simulating players that are using simple explicit rules to solve the puzzles. We compare the solutions returned by the human participants to those from the bots to identify similarities and differences.

H2: Human solutions perform at least as well as standard algorithms for the MSA problem. This hypothesis compares several well-established and non-trivial algorithms designed to solve the multiple sequence alignment problem embedded in Borderlands Science. Again, we compare the performance and similarities between solutions obtained from these algorithms and humans.

H3: We can learn and reproduce player strategies using behavioural cloning. This part of the work considers whether the decisions of the players are reproducible using behaviour cloning. We hypothesize that being able to reproduce the puzzle solutions obtained by the players is an indirect confirmation that the players' strategies have features in common.

We test these hypotheses and aim to establish the relevance of small, fast-paced puzzles, when previous work in this field typically relied on very complex puzzles, such as those in Phylo [KRK$^+$12].

3.1 Related work

Understanding human abilities related to combinatorics and optimization problems has been of interest for the past few decades. In 1990, MacGregor et al. found that human solutions to the Travelling Salesman Problem significantly outperformed solutions obtained using Nearest Neighbor [MO96], Largest Interior Angle or Convex Hull algorithms [MO96]. They noticed that human-based decisions were not related to how state-of-the-art algorithms solve these problems, but were based on human perception. Acuña similarly observed that human performance outperformed not only a random approach, but also many common heuristics, and that humans can improve on the best existing solutions [AP10]. Hidalgo-Herrero provides an analysis of humans solving the Knapsack and Vertex Cover problems [HHRRR13], concluding that human solutions to these problems outperform genetic algorithms. They also observed that human performance decreases with increasing task complexity, and that children develop more diverse and interesting strategies than adults.

Problem-solving tasks in games are good environments to mobilize the human knack for tackling complex problems. Gamification of scientific tasks has gained traction since the release of Foldit in 2008 [CKT$^+$10], with many science discovery games reaching hundreds of thousands of players, such as EteRNA [LKL$^+$14], Galaxy Zoo, Eyewire, Project Discovery and, in particular, Phylo (2010) [KRK$^+$12], the first multiple sequence alignment SDG, which showed humans could improve existing alignments through participating in a game.

James Surowiecki's wisdom of crowds theory [Sur05] suggests that a combination of solutions from different sources may be better than an individual solution. This theory has been described and applied in projects in diverse fields including business, economics and sociology [Sur05], but has only recently been applied to computer science. In 2010, Yi defined classic combinatorics problems such as the Minimum Spanning Tree and Traveling Salesman as puzzles divided into fragments solved by human participants [YSLD12]. These fragments were then assembled into a global solution that outperformed the results of standard algorithms.

The main limitation of this class of crowdsourcing initiatives is the reliance on a large number of participants to solve a single problem. There is thus a natural synergy with artificial intelligence methods that augment crowdsourced data or apply the knowledge gained from humans to other problems. This led to the idea of the wisdom of artificial crowds (WoAC), a strategy successfully applied to general computational problems such as the Travelling Salesman Problem, as well as to real-world games such as Sudoku [YEB11, RSS$^+$15].

The multiple sequence alignment problem, due to its importance in biology, has been studied for over 50 years, and hundreds of algorithms have been proposed to tackle it [THG94, BKR$^+$04, PHB$^+$08, DMBB05, MWP$^+$05]. While these methods are still very popular, an increasing number of new approaches involving machine learning have been published recently [AN16, KLG20], with a small minority involving Reinforcement Learning (RL) [RSB18].

In Reinforcement Learning, an agent interacts with an environment by following a policy in order to maximize its reward function [SB18]. An example of this class of method is the application of Q-learning to the MSA problem [JJKR19]. Their results were later significantly improved with Deep

Q-networks (DQN) [JTKRJ20]. Another application of reinforcement learning to MSA is RLALIGN (2018), which is based on Asynchronous Advantage Actor Critic (A3C) [RSB18].

The main disadvantage of the RL approach is the explicit definition of the Reward Function, which can be quite complex [SG15]. Alternatively, Behavioural Cloning (BC) can be used as an approach focused on the capture and reproduction of human abilities and sequences of actions [Sam10]. This approach has been successfully applied to various video game tasks [HVP$^+$18, GHT$^+$19, ZWL$^+$20]. A recent benchmarking paper reviewing 10 modern games showed that, despite low productivity compared to humans, agents can learn the basic strategy and rules and emphasize that data quality rather than quantity matters for RL [KPH20].

Game-related tasks are well suited for Reinforcement Learning because they typically involve a well-defined action space and reward system: these are key features of game design. Reinforcement Learning has important limitations: it requires a designed environment in which an agent can interact with a game, which can be costly to implement and maintain, and once the environment is in place, agent training can take thousands of years of game time.

However, we know that humans can learn to perform tasks through imitation, and can leverage this approach for automated learning. Applying this learning method to an autonomous agent is referred to as learning from demonstration (LfD) or Imitation Learning (IL) [HGEJ17]. The commonly used approaches within this paradigm can be divided into two broad categories: Behaviour cloning (BC) and Inverse Reinforcement Learning (IRL). AI Players based on the players behavior can solve a large number of different tasks [HGEJ17]. For example, using player experience as observational data, we can train a generative model to play Atari games, which is a benchmark for many RL approaches [CTG$^+$21]. Pfau used this approach to ensure the balancing of in-game parameters and classes to ensure the success of the Aion game [PLV$^+$20]. These methods can also be successfully applied in the real world with systems in autonomous vehicles [LMYDM22].

In this article, we focus on a Behaviour Cloning approach, because it tends to be both simple and effective in solving policy search problems [Far19, CSLG19], and our main goal is to assess whether the player solutions can be mimicked, a task that does not require a complex model.

Behavioural cloning [**?** BS95, RGB11, DBH16] is one of the main approaches for imitation learning. Rather than learning an optimal policy that maximizes the long-term cumulative rewards like in traditional reinforcement learning, a set of demonstrations from an expert is used as a base for learning.

The notation is similar to that used in RL. Expert demonstrations are divided into state-action pairs, and supervised learning is applied as a classification or regression model. The loss function depends entirely on the application and the nature of the data. More formally, it can be written as:

- Collect demonstrations (τ^* trajectories) from expert
- Treat the demonstrations as i.i.d. state-action pairs: $(s_0^*, a_0^*), (s_1^*, a_1^*), \ldots$
- Learn π_θ policy using supervised learning by minimizing the loss function $L(a^*, \pi_\theta(s))$

This approach has many attractive properties. Specialized game modifications and environment creation are not needed, lengthy training is avoided, and the same method can be adapted and reused in different games. Until now, this class of approaches struggles to outperform human experts, and requires a significant amount of expert examples in order to achieve similar proficiency, despite recent advancements combining simulation and optimization yielding promising results [CZW$^+$20]. However, when expert imitation is achieved, the positive outcomes can be significant because expert-level strategies can be computed much faster than human solutions, and scales by adding computational resources rather than through the lengthy process of training new human experts.

The current state of the art in both human solving of hard computational problems and reinforcement learning methods creates a remarkably favourable context to explore the synergy between these two fields. With Borderlands Science, we are given a unique opportunity to explore this synergy and fully unlock the potential of human contributions.

3.2 Methods

In this section, we describe the multiple sequence alignment problem, and the machine learning techniques we used to capture strategies from players.

3.2.1 Data filtering

The Borderlands Science game data consists of human gut microbe genome fragments, sequenced by the Microsetta Initiative, which were pre-aligned with PASTA [MNG+15]. The puzzles submitted to players are carved out of this pre-alignment. In the game, player optimize a bi-objective function: reaching the maximal number of matches (the score), with a limited number of gaps they can insert. This is the same bi-objective function traditional sequence alignment algorithm optimize, but powered by human intuition rather than a scoring matrix.

For this paper, we focus on a sample of 1,000,000 randomly selected puzzle solutions, played between April 2020 and June 2021.

To determine the value of a submitted solution, approximate the Pareto front for each puzzle from all solutions collected for that puzzle (on average, we collected 45 solutions per puzzle). The Pareto front is the set of solutions to a puzzle that optimize our bi-objective function. In other words, a solution is Pareto-optimal if there does not exist another solution that simultaneously matches/increases the score and matches/reduces the number of gaps. We estimate the quality of a solution from its distance from the Pareto front on the x axis (i.e., the score difference between the two).

We aggregated puzzle solutions and extracted those close to the Pareto front, in order to only include the best human solutions. We define a proportional horizontal Pareto distance of a single solution to a puzzle as the score improvement of that solution over the worst human solution, divided by the improvement of the best score for this number of gaps (pareto-optimal) over the worst human solution. Solutions associated with a distance over 0.7 were excluded from the data. This threshold was established through visual assessment of clustering outcomes by human experts. This resulted in 53.4 percent of solutions were filtered out.

3.2.2 Multiple sequence alignment

Definition

A multiple sequence alignment can be described as a set of n sequences $S_1, S_2, \dots S_n$, represented as vectors of the matrix A where each element $a_{i,j}$ comes from the set of nucleotides or gap $(A, C, G, T, -)$. The input to the multiple sequence alignment problem is typically a set of ungapped sequences, and the output sequences contain gaps. This formulation of the problem allows us to consider several sequences as a matrix. This input matrix A is dense, and the output matrix A' contains gaps $-$, inserted in a specific arrangement to maximize some score, such as a phylogenetically-aware scoring scheme [BKR+04] or a simple sum-of-pairs scoring scheme [SMC73]. The mathematical form of the MSA problem is NP-hard for both classes of scoring scheme [WJ94].

The many methods presented to tackle this problem, [THG94, BKR+04, PHB+08, DMBB05, MWP+05] tend to have different performance characteristics that depend on the type, length and number of input sequences. In particular, PASTA [MNG+15] is specifically designed to handle large alignments and was used as the scaffold for the Borderlands Science puzzles. It should be noted all these algorithms are heuristics with no guarantees of optimality, and that multiple sequence alignment remains an open problem.

Algorithms used for comparison

The methods described above for solving the problem have different focuses and approaches. We chose several algorithms for comparison to our methods in two categories corresponding to our hypotheses. The first category, which corresponds to hypothesis 2, contains traditional computational methods developed to tackle the MSA problem, included for the purpose of benchmarking the performance of our methods. It includes: Dynamic Programming (Needleman-Wunsch) [Nee70], PASTA [NMKW15], profile alignment using HMMER [Edd92], and Greedy algorithm. The second category, corresponding to hypothesis 1, aims at providing context to interpret the particularity of player solutions. It includes methods built from basic heuristics that can be applied by a human to solve the game. The Greedy algorithm, a Progressive Profile Strategy (PPS) algorithm, and a Random algorithm are chosen for this category. We designed the greedy and PPS algorithms ourselves. A summary of our methods can be seen in Table 3.1.

Our dynamic programming algorithm is inspired by the Needleman–Wunsch algorithm, a classic sequence alignment algorithm introduced in 1970 [Nee70]. Unfortunately, its complexity is $O(n^k)$, where n is the number of nucleotides in the sequences, and k is the number of sequences. Therefore, we use an adapted method that aligns each sequence to the guide while applying a maximum total width to the alignment. This type of algorithm has many equally optimal solutions (per each sequence). In order to satisfy the gap limit constraint of the Borderlands Science game, we take a 2-step approach: first, we invert the sequences and the guide and remove the gap insertion penalty at the start (that corresponds to the top of the puzzle now), effectively accounting for the "gravity" effect of the puzzle. This is required because NW always fills the entire grid. After obtaining optimal solutions, we try different combinations of these optimal solutions and count the number of gaps

Name	Short Description	Reference	Hypothesis
Random	Randomly places gaps	–	H1
Progressive Profile Strategy (PPS)	Progressive profile alignment	–	H1
Greedy	Chooses best place to insert a gap iteratively	–	H1 & H2
Needleman-Wunsch (NW)	Dynamic programming	[Nee70]	H2
PASTA	Uses alignment and tree estimation, and HMMs	[NMKW15]	H2
HMMER	Uses HMMs	[Edd92]	H2

Table 3.1: Summary of reference algorithms used for comparison.

until we find one that respects the gap limit. If none exist, we incrementally increase the gap limit until such a solution exists. However, due to running time constraints of the dynamic programming algorithm for complex puzzle, a proportion of solutions are skipped.

PASTA [NMKW15] (Practical Alignments Using SAT'e and TrAnsitivity) allows for the computation of MSA alignments for very large nucleotide datasets. Prior to aligning the sequences, PASTA estimates an alignment and a guide tree from a subset of the sequences using a very simple profile Hidden Markov Model (HMM)-based method. A set of HMMs is created from the alignment and the tree, and the rest of the sequences are aligned to each HMM and the best one is used to update the alignment with the sequence. The PASTA alignment is taken directly from the uncollapsed alignment used to generate the puzzles, so does not respect the gap limit given to the player solutions and other algorithms.

In order to use HMMer to solve our puzzles, we used pyhmmer package. We first build a Hidden Markov Model (HMM) profile from our guides with DNA Alphabet, then use the HMMer Trace Aligner method. The main challenge was to fit the solution of the algorithm into the same grid that a player sees in the BLS and use fewer gaps than the maximum that players are allowed. To do this, we varied the consensus impact scale, between 0.8 and 0.95 to yield solutions that use fewer gaps, however, many solutions still do not respect the gap limit.

Greedy algorithms, which optimize for immediate reward, can be effective approaches for solving the MSA problem [ZSWM00] despite their simplicity. Borderlands Science uses a naive greedy solver to set score targets for the players. Because all solutions submitted by players must outperform the greedy solution, that algorithm is intentionally simplistic. For the purpose of this study, we added stochasticity to explore the limits of its reach in the solution space, and simulate players using a strategy aiming for maximum reward after each move.

The Progressive Profile Strategy (PPS) algorithm was inspired by Progressive MSA, combined with an effort to reproduce how players are usually observed to solve puzzles. In this method, first, a profile is created from the guide. We then start from the leftmost sequence and align it to the profile by trying all possible combinations of gaps and choosing the best one. Then we add the solution found for the first sequence to the profile, and move on to aligning the second sequence with it. We continue this until all sequences are aligned with the profile.

We also compared the trained models to a random player who randomly selects a gap location based solely on the size of the puzzle and the number of gaps added by the human players to solve the same problem. We are using this algorithm to simulate a player without strategy.

We tested whether the strategy used by the players is distinct from these methods, and outperforms them on any relevant criteria.

3.2.3 Imitation learning

Data representation

We consider two ways of representing our data (See Figure 3.3). First, we can represent the puzzles as a matrix, where each of the four nucleotides is associated with a channel (e.g., red, green and blue are three distinct channels for traditional computer representation of images). The input is thus a $4 \times L \times N$ matrix, where, L is the number of rows, and N the number of columns. And the output will be a two-dimensional $L \times N$ matrix storing gap positions. We call it a image2image task.

The second option consists in representing each individual sequence as a string. In this case, the output will be the same string in which we eventually

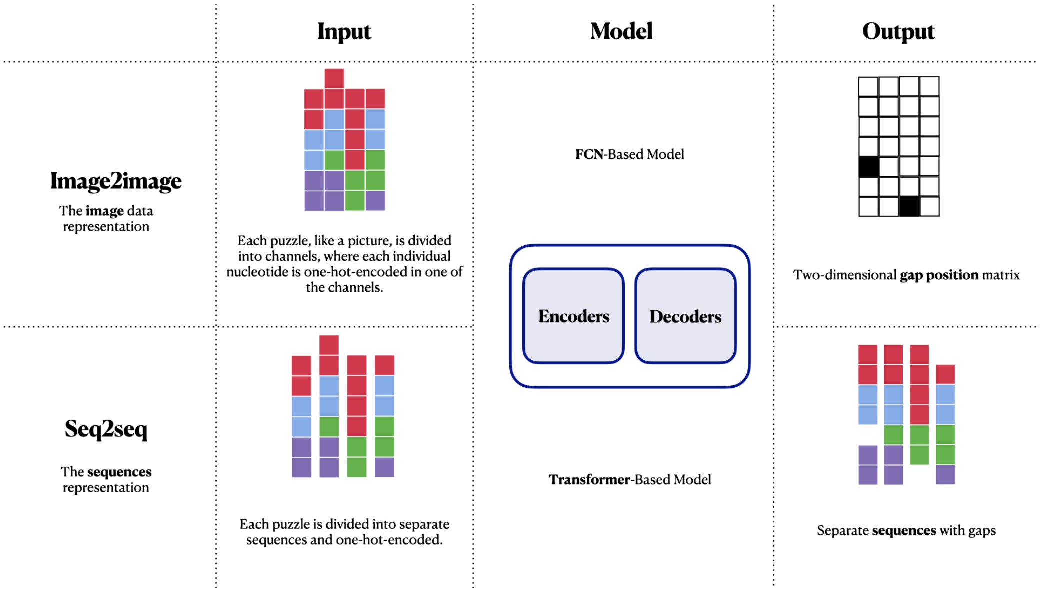

Figure 3.3: Different representations of the data for image2image and seq2seq tasks. The main differences in terms of input, output data and architecture of the model are presented.

insert gaps gaps. We call it a seq2seq task.

Predictive models

- **Fully Convolutional Network.** FCNs were originally presented as a solution to the segmentation problem [LSD15]. Unlike classical CNNs, which rely on a fully connected layer to obtain a fixed-length feature vector for classification after the convolutional layer, FCNs can accept input images of any size because of the set of deconvolutional layers. This architecture has been shown to outperform the state-of-the-art without further machinery in several settings [LLY$^+$19, LZX16, PCMY15].

- **Transformers.** Transformers are a state-of-the-art model with self-attention [VSP$^+$17]. Recent research has shown that self-attention is an effective way to model text sequences [VSP$^+$17, DCLT18]. It consists of an encoder and a decoder. An encoder converts the input information into one or more vectors, and a decoder generates output information from these vectors. The input data passes through the layers of the

> encoder: some of them are standard fully-connected layers, and others are residual connections similar to ResNet [HZRS16]. The most novel component of the encoder section is the Multi-head attention layer, a special layer that allows each input vector to interact with other words through the attention mechanism, instead of passing through a hidden state like in an RNN [She20] or CNN [LBD$^+$89]. Its inputs are *Query* vectors and several *Key-Value* pairs that are responsible for positional information. This embedding will then be used by the decoder. Internal attention layers in a decoder work a little differently than layers in an encoder. In the decoder, the internal attention layer can only focus on previous positions in the output sentence.

We implemented our Transformer model with 8 layers, 64 embedding sizes, and 8 attention heads using a batch size of 100 puzzles, and them trained for 100 epochs. To optimize the transformers-based model, we used Adam optimization with a learning rate of 0.004 and a weight decay of 0.004. We also set attention dropout to 0.1. In order to get the most out of the self-attention mechanism, we selected "axial attention" architecture proposed in [HKWS19]. This allows embedding of a 2D tensor with an autoregressive model, while remaining economical in computation.

By contrast, our FCN model consists of stacked blocks of 2D convolution layers, Dropout and BatchNormalization regularization and a RELU activation function. We have three layers in the encoder and three layers in the decoder with the following parameters: filters = 4, kernel size=1, strides=1. Each position in output data can be associated with one of two classes, indicating whether this position corresponds to a gap. In experimenting with this model, we focused on the more suitable image2image approach, leveraging the encoder-decoder approach as presented by Noh [NHH15]. The input is a puzzle of variable size, and the output is a label assigned to each pixel from one of two classes, determining whether the pixel moved or stayed in place. To optimize the FCN model, we kept the same learning rate of 0.0001, 100 epochs. Optimization was performed with cross-entropy loss minimization as the metric and Adam optimization. The cross-entropy loss is computed against the flattened real player solution image.

3.2.4 Evaluation

The choice of a metric that measure the similarity between outputs is challenging, because puzzles have different sizes and numbers of gaps. To obtain a complete picture of the similarity between solutions, we selected three metrics: Cosine Similarity, Hamming distance, and Jensen-Shannon distance.

The first metric, Cosine Similarity [NB10, XZL15], measures the angle between two vectors in multidimensional space, and is useful for evaluating sparse data such as ours because cosine it ignores 0-0 matches. Cosine similarity is a widely used quality metric that is used in a wide variety of ML areas, such as text classification [LH13] or images analysis [NB10], also finds high use in RL tasks [RLZ17]. We also use cosine similarity to obtain the distance between each solution from algorithms and humans to the consensus of gaps between humans.

We also present the Hamming distance [NFS12] between predicted sequences and player solutions. The metric shows the number of positions in which the characters corresponding to them are different. The Hamming distance is normalized to the string length to generalize to any puzzle size. This metric is extremely useful as it is well interpreted by humans - visually it is very easy to identify mismatched elements in two lines of the test. Hamming distance is often used for language models [LPL21, RMS$^+$20]. These approaches can also be used in work and genetic chains, presenting language model information [LSB22, PPS12].

We report the Jensen-Shannon distance [FT04], which measures the differences between probability distributions. We compare the distribution of gaps between predictions and real solutions. If two distributions are similar, the Jensen-Shannon distribution between them is 0. This metric is more complex as it considers the real and predicted distributions. Jensen-Shannon divergence has also been used by many authors to analyze their research. For example, [WDJ04] use this metric for a scoring scheme for Sequence alignment profiles task, also the metric is used for DNA analysis [GBGC$^+$02].

We also report a similarity to the centroid of player solutions. The centroid of player solutions is a representative solution that is calculated by identifying the player solution that has the highest cosine similarity to the consensus of gaps between all player solutions for a specific puzzle. In the context of our analysis, we use the Levenshtein distance to compare the predicted solutions to the centroid of all player solutions for that puzzle. The Levenshtein distance is a measure of the similarity between two sequences by measuring the amount of insertions, deletions and substitutions needed to transform one sequence to another, and is often used in natural language processing tasks to compare the similarity between two pieces of text. By using the Levenshtein distance on the centroid, we can determine the similarity of the predictions to the most popular player solution, without being heavily influenced by outliers. As this information is not at all used in the training of RL agents, it also serves as a sanity check.

For the methods that do not adhere to the gap limit provided to the player such as HMMer and PASTA, we can compare their performances by looking at how many incompatibilities there are with the player solutions. Incompatibilities are described as the number of gap deletions required to transform a player's solution into the algorithm's solution using only gap insertions. This allows us to see how similar the player solutions are to a particular algorithm's strategy, regardless of the number of gaps the solutions use. The assumption is that if there are few incompatibilities, it means that the player solutions are likely incorporating part of the strategy, but may not have enough gap insertions to fully solve the puzzle. Figure 3.9 shows the mean of these solutions across all puzzle difficulties and highlights the similarities between the various algorithms.

We also measure the mean number of gaps per puzzle and the distribution of gaps across the columns of each puzzle. These values are split into three groups of puzzle difficulties (which correspond to the size and complexity of the puzzle) to allow us to compare the performance of the different methods across different levels of difficulty. Analyzing the distribution of gaps can help us understand the strategies that lead to the results we observe.

All metrics are measured between the flattened forms of the algorithm or agent-produced and player-solved puzzle matrices.

3.3 Results

3.3.1 Data from Borderlands Science

The Borderlands Science data was processed as described in section Methods 4.2. Table 3.2 presents a detailed description of this dataset. It contains a total of 25,000 puzzles and 1,145,001 solutions.

3.3.2 Testing Hypothesis 1

To test whether the player strategies were significantly different from simple heuristics, we solved the puzzles with a greedy player, a Progressive Profile Strategy (PPS) algorithm, and a random player. We compared these algorithmic solutions to human solutions with the methods listed in section 3.4.

First, when comparing solutions, we observe that the PPS, greedy and random players achieve lower game scores than humans (Figure 4.5). They

Difficulty	Mean used gaps	Number of puzzles
1	5.68	5443
2	5.62	3627
3	5.03	2108
4	6.60	2001
5	7.95	2292
6	9.76	3753
7	10.34	2833
8	11.37	2851

Table 3.2: Level-by-level description of the Borderlands Science data the agents were trained on, with the average number of gaps inserted per puzzle for each level

	CosSimilarity	HammingDist	JensenShannonDist
Transformer	0.79	0.07	0.40
FCN	0.77	0.08	0.41
NW	0.77	0.04	0.37
Greedy	0.64	0.04	0.30
PPS	0.58	0.11	0.44
Random	0.55	0.07	0.44

Table 3.3: The three similarity metrics computed on the final state were obtained with seven reference algorithms. We report the similarity/distance between an algorithm's solution to a puzzle and the Pareto-optimality-filtered player solutions for this puzzle. The optimal result is 1 for cosine similarity, and 0 for the two distance metrics.

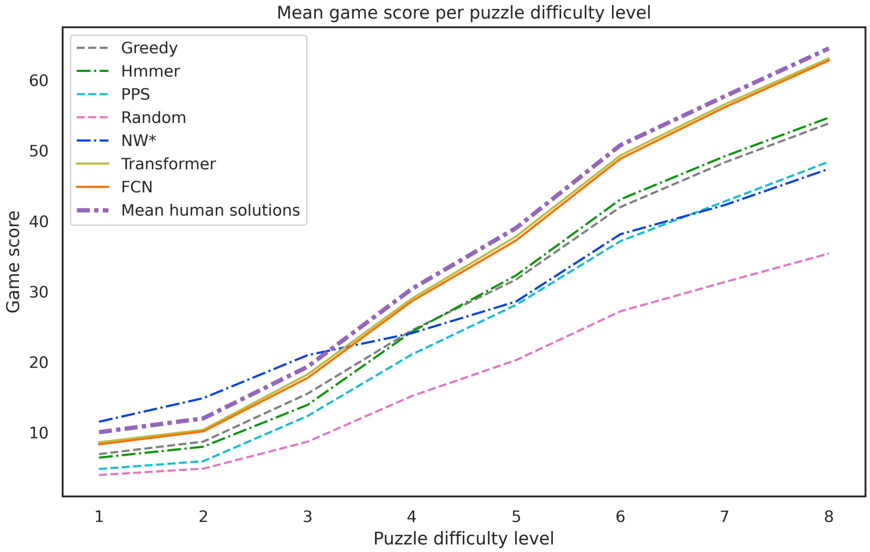

Figure 3.4: The average game score obtained by the players (real) and models (i.e., reference algorithms) for puzzles at increasing difficulty levels. (*) Results for the Needleman-Wunsch (NW) algorithm are computed for 10,000 puzzles on which our implementation could return a result.

also generate final states that are less similar to the averaged human solutions (Figure 3.5) and further from the human solution centroid (Figure 3.6) than the Needleman-Wunsch algorithm. Overall, this appears to indicate the similarity between human solutions and PPS, greedy and random solutions is low.

Second, we compare strategies. We observe that PPS places a similar number of gaps per puzzle as the human players (near the maximum allowed) for all difficulty levels (Figure 3.8). However, the column placement patterns differ: humans build more columns with one gap than columns with zero, whereas PPS builds more columns with zero gaps than columns with one (Figure 3.9). The Greedy gap placement behaviour also significantly differs from humans. An example result of the difference in gap placement patterns can be seen in Figure 4.8.

Conclusion: Our results indicate that both in terms of strategy and performance, there is a low similarity between human players and basic heuris-

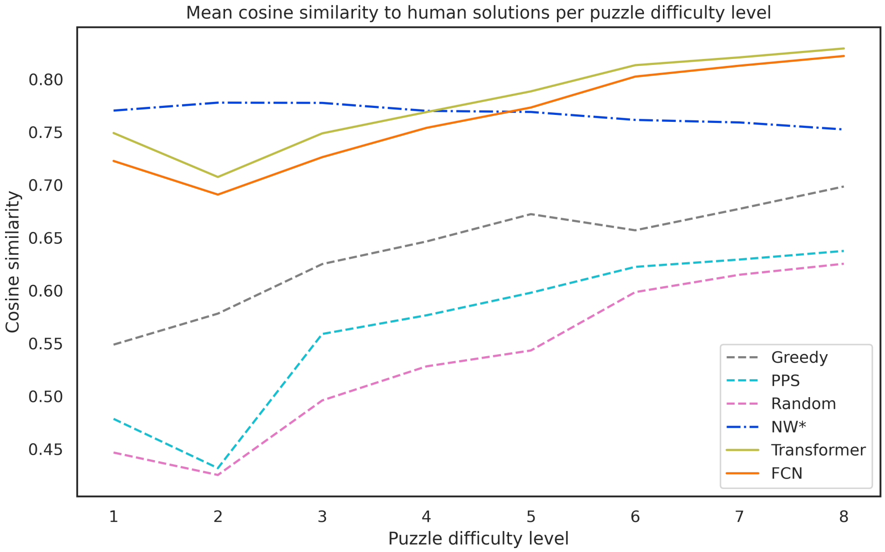

Figure 3.5: Cosine similarity to averaged player solutions. (*) Results for the Needleman-Wunsch (NW) algorithm are computed for 10,000 puzzles on which our implementation could return a result.

tics.

3.3.3 Testing Hypothesis 2

To test whether the player strategies provide a satisfying solution to the multiple sequence alignment problem, we also included standard algorithms to our comparison set, such as a modified Needleman-Wunsch and a greedy algorithm.

Since these two methods could be modified to accommodate the gap limit in the game (see methods), we can compare them to the human performance in terms of game score, a score obtained primarily from correctly matching nucleotides to guides. We observe that Needleman-Wunsch outperforms the average player for the first three difficulty levels (smallest puzzles), but that as the problem gets more complex it falls behind human solutions (Figure 4.5). This phenomenon suggests that the players use a strategy that allows them to quickly identify a near-optimal solution. When the combinatorial complexity of the puzzle is growing, an exhaustive search is no longer feasible and players clearly outperform this optimization method. The greedy algorithm is overall

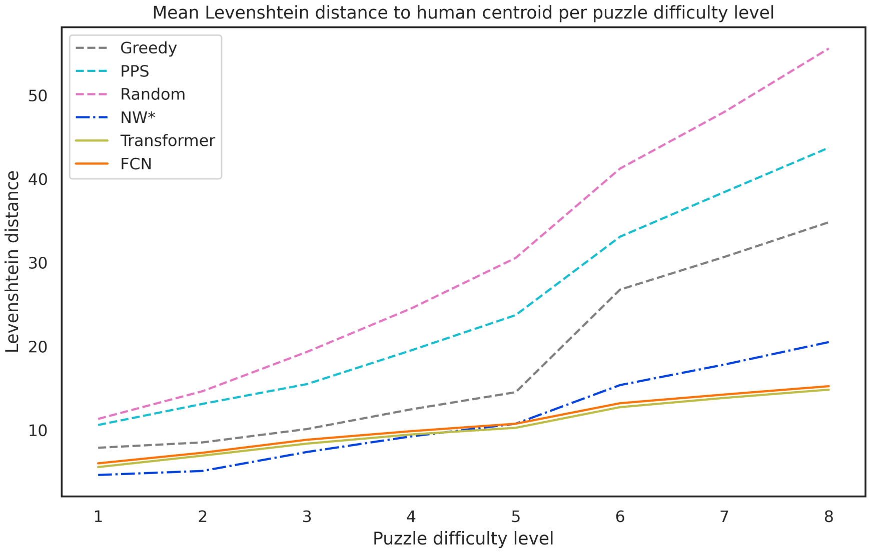

Figure 3.6: Levenshtein distance from the player solution centroid. (*) Results for the Needleman-Wunsch (NW) algorithm are computed for 10,000 puzzles on which our implementation could return a result.

outperformed by the average human.

Because some other standard algorithms such as HMMer and PASTA cannot accommodate the gap limit imposed on the players and cannot easily be modified to accommodate it (as we did for Needleman-Wunsch), we also performed a separate investigation on the compatibility of human solutions with the outputs of advanced and widely-used algorithms HMMer [Edd92] and PASTA [NMKW15] (see section 3.5).

In this investigation (Figure 3.9), we estimated the average number of gaps placed by players that are not compatible with the software solution. We observed that the behavior of HMMer is highly compatible with that of players. In other words, the gaps inserted by players are overwhelmingly gaps inserted by HMMer. By contrast, PASTA showed significantly more disagreement with player solutions, in large part because it is very gap-adverse (i.e., PASTA tends to avoid opening new gaps in sequences that has none). This is consistent with the role of PASTA in the Borderlands Science pipeline; it builds a very tight scaffold and then players add missing gaps to improve the

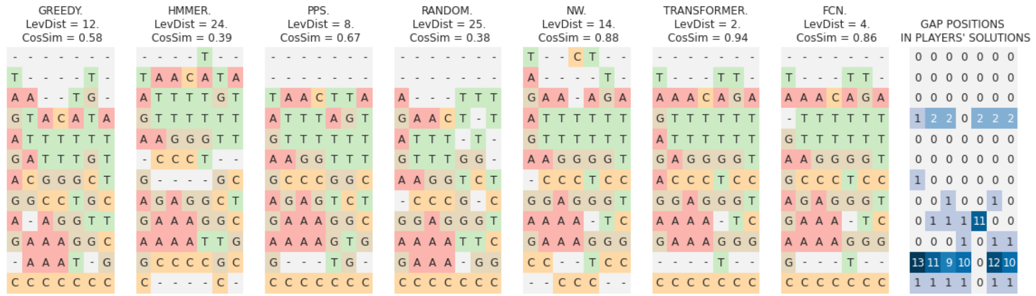

Figure 3.7: Solution Example: Greedy, Hmmer, PPS, Random, NW, Transformer, FCN, Average Human Solution. Average Human Solution includes a number of how often each position has been chosen as a move.

alignment. This however comes with PASTA obtaining a much poorer score.

Conclusion: In the context of the game, humans appear to outperform traditional algorithms for puzzles above the easiest difficulty levels, and the moves made by human players are compatible with gap-friendly methods such as HMMer and complementary to gap-adverse methods such as PASTA. Interestingly, solutions of the Needleman-Wunsch optimization algorithm bear similarities with those of humans. Yet, the player strategy quickly becomes more efficient at identifying a good solution when the difficulty increases.

3.3.4 Testing Hypothesis 3

To test whether it was possible to accurately learn from the human solutions, we leveraged two well-researched behaviour cloning approaches.

We observed that Transformers and FCN achieve high cosine similarity to human consensus (Figure 3.5) and low Levenshtein distance against human centroid (Figure 3.6). RL methods obtain game scores (Figure 4.5) comparable to humans, but slightly worse, notably because they place fewer gaps (Figure 3.8), limiting itself to high-consensus gaps. In other words, after having added the gaps that most humans agree with, adding more gaps increases the distance against the majority of humans. This can be a strength in the context of sequence alignments where gap insertion and extensions are usually limited by some form of penalty to avoid having too many gaps. The very small difference in score, in light of the significant difference in gap insertions, indicates that the additional gaps inserted by humans provide a diminishing returns, suggesting the FCN and Transformer's gap-efficiency is a strength.

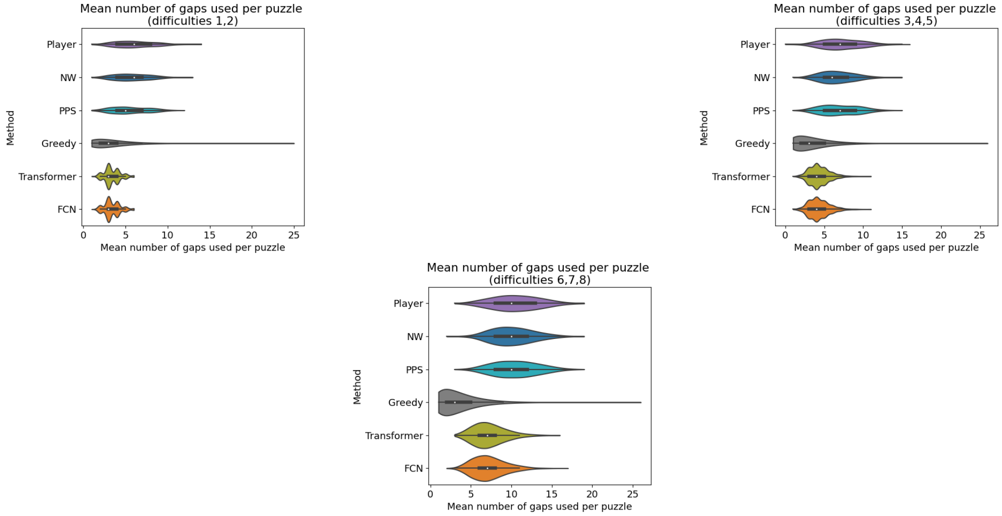

Figure 3.8: Comparison of the number of gaps per puzzle difficulty category (1-2, 3-4-5, 6-7-8) to identify solution strategies between the player solutions, ML models and naive algorithm.

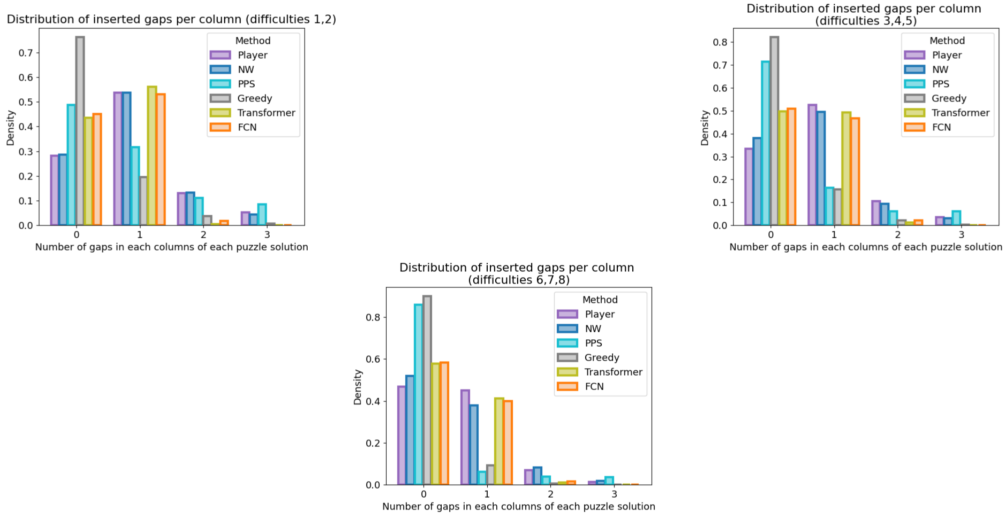

Figure 3.9: Comparison of the number of inserted gaps per column per puzzle for each puzzle difficulty category (1-2, 3-4-5, 6-7-8) to identify solution strategies between the player solutions, ML models and naive algorithm.

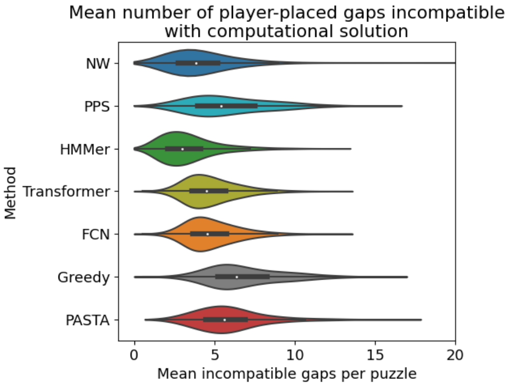

Figure 3.10: The similarity of solutions from reference algorithms to human solutions (i.e., mean player solution) using a customized edit distance. Our distance reports the average number of incompatible gaps in the solutions of the reference algorithms that must be removed to reach the human solutions. We report the statistics for all difficulty levels together.

It should also be noted that this high game score was obtained by the agents despite the fact that it does not perfectly copy player behavior. This indicates that even when the agent fails to correctly imitate the players, it still provides a solution that obtains a decent score suggesting it presents a valid solution to the general problem of multiple sequence alignment.

In terms of strategy, RL methods produce nearly exactly as many columns with one gap as humans over the three difficulty categories, and its gap placement patterns become more similar to humans as the complexity of the problem increases (and as human solutions progressively outperform Needleman-Wunsch).

Conclusion: We can successfully capture strategies employed by humans using behaviour cloning techniques. Both behaviour cloning methods we present accurately mimic the strategies of human players, with Transformers slightly out-performing FCN.

3.4 Discussion

We have shown that all three hypotheses presented in this paper are corroborated by experimental results: the player solutions are original, high-quality and mimickable.

By simultaneously demonstrating that the player strategies are original

(H1) and that they produce outputs comparable in quality to state of the art algorithms (H2), we established the relevance of training an agent to learn from these players on Borderlands Science. This confirms our hypothesis based on previous successes in the solving of computational problems by humans [MO96, AP10, YEB11, RSS$^+$15]. Notably, humans out-performed greedy algorithms on Borderlands Science as they did in [AP10], and we confirmed the presence of a perceptible regular pattern in how humans solve these tasks, as in [MO96].

We also demonstrated that accurately learning from players was achievable (H3). Our results show that behavioral cloning can be an effective approach for a puzzle game such as Borderlands Science, confirming our hypothesis based on its performance on Atari [ZWL$^+$20]. In particular, we discovered that learning from the consensus of players leads to a gap-efficient agent able to achieve scores comparable to humans while excluding gaps that provide diminishing returns, a particularity that is highly relevant for tasks such as multiple sequence alignment.

A preliminary analysis of the moves played by our model reveals that a strategy captured from the players aims at finding a specific trade-off between the number and length of gaps (i.e., the length of gaps is defined as the number of adjacent gap tiles). It turns out that the determination of a gap penalty scheme is one of the most sensitive parameters in biological sequence alignment [VW94]. Hence, this work offers a piece of promising information to design new sequence alignment algorithms.

There are two main application paths for the methods we have presented. The most obvious one is to leverage the solutions from the players by applying the strategy they taught us to new sequences, effectively generalizing the results of the Borderlands Science initiative and taking full advantage of the data, limiting the impact of the significant overhead cost of the project. Another promising avenue is to use the reinforcement learning agents as a partner for the players, for example by providing hints, thus creating a fully-fledged Human-in-the-loop (HITL) system to leverage both human and machine intelligence for solving combinatorial problems.

From a technical perspective, both machine learning frameworks yielded comparable and positive results. Yet, it appears that transformers trained with the seq2seq framing outperform FCNs trained on the full image on two of our three metrics. This could be significant in terms of indicating a better framing of the problem, but it could also be due to the high performance of state-of-the-art transformer-based methods. We leave a full investigation of this question to future work.

3.4.1 Limitations

First, the algorithmic methods presented here as benchmarks are limited. Typical multiple sequence alignment methods do not usually have a hard gap limit and integrating that constraint can be difficult: our modified Needleman-Wunsch was only able to terminate for 40% of the puzzles. We only reported results for the problems it was able to terminate on. Greedy algorithms work well with gap limits since they optimize step by step but they are rarely optimal for MSAs. We were unable to integrate a gap limit to PASTA and HMMer and had to limit ourselves to assessing their compatibility to human solutions. The bottom line here is that humans and algorithms solve the MSA problem with different approaches and a perfect benchmarking of one against the other is not possible.

Second, while the purpose of the approach we present is to be able to apply the strategies learned from human players to any dataset, its performance on a dataset of non-microbial sequences has not yet been tested. However, we have evidence that at least in the context this approach has been tested, it presents a reasonable solution to the MSA problem, which seems to indicate the agents learn general strategies that should apply to any nucleic acid sequence.

Third, a limitation of the Borderlands Science game is that the players are not explicitly told to optimize a bi-objective function. They have a limited number of gaps but they have no incentive not to use all the gaps they are given. We alleviate this limitation by showing players different numbers of gap tokens for the same puzzle, and by filtering player output for solutions near the pareto front, which results in a large enough set of nearly optimal solutions.

Also, pareto-optimality, which we use here as a proxy to identify good solutions to the MSA problem, is not guaranteed to provide the best solution to this problem. It is not impossible that there exists a better way to filter this data and extract solutions that constitute the best strategies to tackle the MSA problem. However, this limitation is damped by the logical relevance of pareto-optimality to the problem at hand; while we cannot guarantee the player solutions are globally optimal, we have shown strong evidence that they generally outperform computational methods on at least one important evaluation criteria, which is sufficient to support the claim that these solutions provide valuable information that is worth learning to reproduce (see

H1).

Another typical limitation of a large-scale participative science approach is the potential presence of disruptive behavior and its handling [PM19]. In Borderlands Science, several steps were taken to limit the impact of such behavior: the requirement of reaching the par score forces players to at least attempt to solve the puzzle, and, solutions that are too distant from the consensus are removed. We are inclined to trust this consensus between players because the main feedback we received from our interaction with the player-base was excitement about contributing to science. Finally, the main rewards offered to player, in-game cosmetics, can only be received once per level, which reduces the long-term value of playing the game and thus de-incentivizes botting.

Additionally, the reinforcement learning methods presented in this article do not cover the entirety of the state of the art for imitation learning. We intentionally selected simpler models that are known to work well for this type of problem rather than more complex models such as DQN to assess whether player solutions could be mimicked with a simple model. It is possible alternate methods outperform the ones presented here. Nevertheless, we reserve the investigation of alternative strategies for future work, as recent work has confirmed it is possible to train an agent to play at an experienced level in a game with behavioural cloning [VBC$^+$19], and this is the extent of the evidence we needed to accumulate for the claims presented in this work.

Finally, the data presented in this study is difficult to acquire. Assembling a million solutions from players requires selecting the right problem, adding an adequate level of gamification, and finding a way to massively distribute the game. The overhead cost is significant. However, the methods presented in this paper represent a promising avenue to make full use of the output by widening its range of applications, further justifying that overhead cost.

3.5 Conclusion

In this work, we showed that Borderlands Science players leveraged heterogeneous and efficient strategies and that these strategies could be mimicked with machine learning. This contribution further supports the relevance of scientific discovery games to tackle scientific data analysis challenges by offering an avenue to apply these results to new data. Yet, the main novelty of this contribution is the successful application of this strategy to an NP-hard computational problem, charting a promising course for solving complex problems with the wisdom of the crowd through massive-scale citizen science games.

3.6 Acknowledgements

The authors want to recognize the contribution of the entire Borderlands Science (BLS) consortium* to this project, which includes (but is not limited to) the design, implementation, production, promotion of the BLS game and preparation of the data sets. The team is composed of members from Gearbox Studio Québec and Gearbox Software (Gearbox), Massively Multiplayer Online Science (MMOS), The Microsetta Initiative at the University of California at San Diego (UCSD), and the School of Computer Science at McGill University (McGill).

The BLS consortium is made of (by alphabetical order): David Bélanger (Gearbox), Alexis Belley Duffault (Gearbox), Mathieu Blanchette (McGill), Michael Bouffard (Gearbox), Amélie Brouillette (Gearbox), Alexander Butyaev (McGill), Eddie Cai (McGill), Sébastien Caisse (Gearbox), Joshua Davidson (Gearbox), Chrisostomos Drogaris (McGill), Kornel Erhart (MMOS), Mathieu Falaise (Gearbox), Fanny Pelletier (Gearbox), Vincent Fiset (Gearbox), Olivier Gagné Houle (Gearbox), Parham Ghasemloo (McGill), Steven Hebert (Gearbox), Dan Hewitt (Gearbox), Jonathan Huot (Gearbox), Raphaël Jean (Gearbox), Timothy Keding (McGill), Seung Kim (Gearbox), Rob Knight (UCSD), Daniel McDonald (UCSD), Jonathan Moreau-Genest (Gearbox), Julien Mounthanyvong (McGill), Renata Mutalova (McGill), David Najjab (Gearbox), Elena Nazarova (McGill), Randy Pitchford (Gearbox), Steve Prince (Gearbox), Gabriel Richard (Gearbox), Ludger Saintélien (Gearbox), Roman Sarrazin-Gendron (McGill), Hugues Thibodeau (Gearbox), Attila Szantner (MMOS), Jérôme Waldispühl (McGill), Jiayue Zheng (McGill), and Yuxue Zhu (McGill).

Additionally, this project would not have been possible without the participation of the millions of Borderlands Science players; the results presented in this paper are the direct product of their contributions.

Finally, the authors want to thank Genome Canada and Génome Québec for their financial support to this project through the Genomic Applications Partnership Program (GAPP).

Bibliography

[AN16] Massimo Andreatta and Morten Nielsen. Gapped sequence alignment using artificial neural networks: application to the mhc class i system. *Bioinformatics*, 32(4):511–517, 2016.

[AP10] Daniel E Acuña and Víctor Parada. People efficiently explore the solution space of the computationally intractable traveling salesman problem to find near-optimal tours. *PloS one*, 5(7):e11685, 2010.

[BDTSW22] Alexander Butyaev, Chrisostomos Drogaris, Olivier Tremblay-Savard, and Jérôme Waldispühl. Human-supervised clustering of multidimensional data using crowdsourcing. *Royal Society open science*, 9(5):211189, 2022.

[BKR$^+$04] Mathieu Blanchette, W James Kent, Cathy Riemer, Laura Elnitski, Arian FA Smit, Krishna M Roskin, Robert Baertsch, Kate Rosenbloom, Hiram Clawson, Eric D Green, et al. Aligning multiple genomic sequences with the threaded blockset aligner. *Genome research*, 14(4):708–715, 2004.

[BS95] Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Machine Intelligence 15*, pages 103–129, 1995.

[CKT$^+$10] Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker, Zoran Popović, et al. Predicting protein structures with a multiplayer online game. *Nature*, 466(7307):756–760, 2010.

[CMC$^+$16] Maria Chatzou, Cedrik Magis, Jia-Ming Chang, Carsten Kemena, Giovanni Bussotti, Ionas Erb, and Cedric Notredame. Multiple sequence alignment modeling: methods and applications. *Briefings in bioinformatics*, 17(6):1009–1023, 2016.

[CSBG18] Irene Cristofori, Carola Salvi, Mark Beeman, and Jordan Grafman. The effects of expected reward on creative problem solving. *Cognitive, Affective, & Behavioral Neuroscience*, 18(5):925–931, 2018.

[CSLG19] Felipe Codevilla, Eder Santana, Antonio M López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9329–9338, 2019.

[CTG+21] Brian Chen, Siddhant Tandon, David Gorsich, Alex Gorodetsky, and Shravan Veerapaneni. Behavioral cloning in atari games using a combined variational autoencoder and predictor model. In *2021 IEEE Congress on Evolutionary Computation (CEC)*, pages 2077–2084. IEEE, 2021.

[CZW+20] Xinyue Chen, Zijian Zhou, Zheng Wang, Che Wang, Yanqiu Wu, and Keith Ross. Bail: Best-action imitation learning for batch deep reinforcement learning. *Advances in Neural Information Processing Systems*, 33:18353–18363, 2020.

[DBH16] Shreyansh Daftry, J Andrew Bagnell, and Martial Hebert. Learning transferable policies for monocular reactive mav control. In *International Symposium on Experimental Robotics*, pages 3–11. Springer, 2016.

[DCLT18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[DMBB05] Chuong B Do, Mahathi SP Mahabhashyam, Michael Brudno, and Serafim Batzoglou. Probcons: Probabilistic consistency-based multiple sequence alignment. *Genome research*, 15(2):330–340, 2005.

[Edd92] Sean Eddy. Hmmer user's guide. *Department of Genetics, Washington University School of Medicine*, 2(1):13, 1992.

[Edd22] Sean Eddy. 2022 essential facts about the video game industry. *Entertainment Software Association*, 2022.

[ESB+12] Christopher B Eiben, Justin B Siegel, Jacob B Bale, Seth Cooper, Firas Khatib, Betty W Shen, Barry L Stoddard, Zoran Popovic, and David Baker. Increased diels-alderase activity through backbone remodeling guided by foldit players. *Nature biotechnology*, 30(2):190–192, 2012.

[Far19] Wael Farag. Cloning safe driving behavior for self-driving cars using convolutional neural networks. *Recent Patents on Computer Science*, 12(2):120–127, 2019.

[FT04] Bent Fuglede and Flemming Topsoe. Jensen-shannon divergence and hilbert space embedding. In *International Symposium onInformation Theory, 2004. ISIT 2004. Proceedings.*, page 31. IEEE, 2004.

[GBGC$^+$02] Ivo Grosse, Pedro Bernaola-Galván, Pedro Carpena, Ramón Román-Roldán, Jose Oliver, and H Eugene Stanley. Analysis of symbolic sequences using the jensen-shannon divergence. *Physical Review E*, 65(4):041905, 2002.

[GHT$^+$19] William H Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, and Ruslan Salakhutdinov. Minerl: A large-scale dataset of minecraft demonstrations. *arXiv preprint arXiv:1907.13440*, 2019.

[HGEJ17] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.

[HHRRR13] Mercedes Hidalgo-Herrero, Pablo Rabanal, Ismael Rodriguez, and Fernando Rubio. Comparing problem solving strategies for np-hard optimization problems. *Fundamenta Informaticae*, 124(1-2):1–25, 2013.

[HKWS19] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.

[HVP$^+$18] Todd Hester, Matej Vecerik, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Dan Horgan, John Quan, Andrew Sendonaris, Ian Osband, et al. Deep q-learning from demonstrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[HZRS16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[JJKR19] Reza Jafari, Mohammad Masoud Javidi, and Marjan Kuchaki Rafsanjani. Using deep reinforcement learning approach for solving the multiple sequence alignment problem. *SN Applied Sciences*, 1(6):1–12, 2019.

[JTKRJ20] Reza Jafari, Hamed Tabrizchi, Marjan Kuchaki Rafsanjani, and M.M. Javidi. A distributed deep q-learning based model for solving the multiple sequence alignment problem. 04 2020.

[KGL$^+$98] Matthias J Koepp, Roger N Gunn, Andrew D Lawrence, Vincent J Cunningham, Alain Dagher, Tasmin Jones, David J Brooks, Christopher J Bench, and PM Grasby. Evidence for striatal dopamine release during a video game. *Nature*, 393(6682):266–268, 1998.

[KLG20] Mengmeng Kuang, Yong Liu, and Lufei Gao. Dlpalign: A deep learning based progressive alignment method for multiple protein sequences. In *CSBio'20: Proceedings of the Eleventh International Conference on Computational Systems-Biology and Bioinformatics*, pages 83–92, 2020.

[KPH20] Anssi Kanervisto, Joonas Pussinen, and Ville Hautamäki. Benchmarking end-to-end behavioural cloning on video games. In *2020 IEEE conference on games (CoG)*, pages 558–565. IEEE, 2020.

[KRK$^+$12] Alexander Kawrykow, Gary Roumanis, Alfred Kam, Daniel Kwak, Clarence Leung, Chu Wu, Eleyine Zarour, Phylo players, Luis Sarmenta, Mathieu Blanchette, et al. Phylo: a citizen science approach for improving multiple sequence alignment. *PloS one*, 7(3):e31362, 2012.

[LBD$^+$89] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.

[LG08] Ari Löytynoja and Nick Goldman. Phylogeny-aware gap placement prevents errors in sequence alignment and evolutionary analysis. *science*, 320(5883):1632–1635, 2008.

[LH13] Baoli Li and Liping Han. Distance weighted cosine similarity measure for text classification. In *International conference on intelligent data engineering and automated learning*, pages 611–618. Springer, 2013.

[LKL$^+$14] Jeehyung Lee, Wipapat Kladwang, Minjae Lee, Daniel Cantu, Martin Azizyan, Hanjoo Kim, Alex Limpaecher, Snehal Gaikwad, Sungroh Yoon, Adrien Treuille, et al. Rna design rules from a massive open laboratory. *Proceedings of the National Academy of Sciences*, 111(6):2122–2127, 2014.

[LLY$^+$19] Xin Lu, Buyu Li, Yuxin Yue, Quanquan Li, and Junjie Yan. Grid r-cnn. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7363–7372, 2019.

[LMYDM22] Luc Le Mero, Dewei Yi, Mehrdad Dianati, and Alexandros Mouzakitis. A survey on imitation learning techniques for end-to-end autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[LPL21] Hong-Liang Li, Yi-He Pang, and Bin Liu. Bioseq-blm: a platform for analyzing dna, rna and protein sequences based on biological language models. *Nucleic acids research*, 49(22):e129–e129, 2021.

[LSB22] Umberto Lupo, Damiano Sgarbossa, and Anne-Florence Bitbol. Protein language models trained on multiple sequence alignments learn phylogenetic relationships. *bioRxiv*, 2022.

[LSD15] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.

[LZX16] Bo Li, Tianlei Zhang, and Tian Xia. Vehicle detection from 3d lidar using fully convolutional network. *arXiv preprint arXiv:1608.07916*, 2016.

[MCS17] David C Moffat, William Crombie, and Olga Shabalina. Some video games can increase the player's creativity. *International Journal of Game-Based Learning (IJGBL)*, 7(2):35–46, 2017.

[MHD$^+$18] Daniel McDonald, Embriette Hyde, Justine W Debelius, James T Morton, Antonio Gonzalez, Gail Ackermann, Alexander A Aksenov, Bahar Behsaz, Caitriona Brennan, Yingfeng Chen, et al. American gut: an open platform for citizen science microbiome research. *Msystems*, 3(3):e00031–18, 2018.

[MNG$^+$15] Siavash Mirarab, Nam Nguyen, Sheng Guo, Li-San Wang, Junhyong Kim, and Tandy Warnow. Pasta: ultra-large multiple sequence alignment for nucleotide and amino-acid sequences. *Journal of Computational Biology*, 22(5):377–386, 2015.

[MO96] James N MacGregor and Tom Ormerod. Human performance on the traveling salesman problem. *Perception & psychophysics*, 58(4):527–539, 1996.

[MWP$^+$05] Burkhard Morgenstern, Nadine Werner, Sonja J Prohaska, Rasmus Steinkamp, Isabelle Schneider, Amarendran R Subramanian, Peter F Stadler, and Jan Weyer-Menkhoff. Multiple sequence alignment with user-defined constraints at gobics. *Bioinformatics*, 21(7):1271–1273, 2005.

[NB10] Hieu V Nguyen and Li Bai. Cosine similarity metric learning for face verification. In *Asian conference on computer vision*, pages 709–720. Springer, 2010.

[Nee70] SB Needleman. Needleman-wunsch algorithm for sequence similarity searches. *J Mol Biol*, 48:443–453, 1970.

[NFS12] Mohammad Norouzi, David J Fleet, and Russ R Salakhutdinov. Hamming distance metric learning. In *Advances in neural information processing systems*, pages 1061–1069, 2012.

[NHH15] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1520–1528, 2015.

[NMKW15] Nam-phuong D Nguyen, Siavash Mirarab, Keerthana Kumar, and Tandy Warnow. Ultra-large alignments using phylogeny-aware profiles. *Genome biology*, 16(1):1–15, 2015.

[ON20] Csaba Osvath and Erica Newport. The creative gamer: Creativity and creative problem solving in video games. 2020.

[PAD⁺17] Vineet Pandey, Amnon Amir, Justine Debelius, Embriette R Hyde, Tomasz Kosciolek, Rob Knight, and Scott Klemmer. Gut instinct: Creating scientific theories with online learners. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pages 6825–6836, 2017.

[PCMY15] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015.

[PHB⁺08] Benedict Paten, Javier Herrero, Kathryn Beal, Stephen Fitzgerald, and Ewan Birney. Enredo and pecan: genome-wide mammalian consistency-based multiple alignment with paralogs. *Genome research*, 18(11):1814–1828, 2008.

[PKY⁺21] Vineet Pandey, Tushar Koul, Chen Yang, Daniel McDonald, Mad Price Ball, Bastian Greshake Tzovaras, Rob Knight, and Scott Klemmer. Galileo: Citizen-led experimentation using a social computing system. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–14, 2021.

[PLV⁺20] Johannes Pfau, Antonios Liapis, Georg Volkmar, Georgios N Yannakakis, and Rainer Malaka. Dungeons & replicants: automated game balancing via deep player behavior modeling. In *2020 IEEE Conference on Games (CoG)*, pages 431–438. IEEE, 2020.

[PM19] Johannes Pfau and Rainer Malaka. Can you rely on human computation? a large-scale analysis of disruptive behavior in games with a purpose. In *Extended Abstracts of the Annual*

Symposium on Computer-Human Interaction in Play Companion Extended Abstracts, pages 605–610, 2019.

[PPS12] Aluísio Pinheiro, Hildete Prisco Pinheiro, and Pranab Kumar Sen. The use of hamming distance in bioinformatics. In *Handbook of Statistics*, volume 28, pages 129–162. Elsevier, 2012.

[RGB11] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.

[RLZ17] Yongming Rao, Jiwen Lu, and Jie Zhou. Attention-aware deep reinforcement learning for video face recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 3931–3940, 2017.

[RMS$^+$20] Roshan Rao, Joshua Meier, Tom Sercu, Sergey Ovchinnikov, and Alexander Rives. Transformer protein language models are unsupervised structure learners. *Biorxiv*, 2020.

[RSB18] Ramchalam Kinattinkara Ramakrishnan, Jaspal Singh, and Mathieu Blanchette. Rlalign: a reinforcement learning approach for multiple sequence alignment. In *2018 IEEE 18th International Conference on Bioinformatics and Bioengineering (BIBE)*, pages 61–66. IEEE, 2018.

[RSS$^+$15] J Redding, Jacob Schreiver, C Shrum, A Lauf, and R Yampolskiy. Solving np-hard number matrix games with wisdom of artificial crowds. In *2015 Computer Games: AI, Animation, Mobile, Multimedia, Educational and Serious Games (CGAMES)*, pages 38–43. IEEE, 2015.

[Sam10] Claude Sammut. Behavioral cloning., 2010.

[SB18] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction.* MIT press, 2018.

[SG15] Seungmoon Song and Hartmut Geyer. A neural circuitry that emphasizes spinal feedback generates diverse behaviours of human locomotion. *The Journal of physiology*, 593(16):3493–3511, 2015.

[She20] Alex Sherstinsky. Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena*, 404:132306, 2020.

[SMC73] David Sankoff, Cristiane Morel, and Robert J Cedergren. Evolution of 5s rna and the non-randomness of base replacement. *Nature New Biology*, 245(147):232–234, 1973.

[SS13] Rafik A Salama and Dov J Stekel. A non-independent energy-based multiple sequence alignment improves prediction of transcription factor binding sites. *Bioinformatics*, 29(21):2699–2704, 2013.

[Sur05] James Surowiecki. *The wisdom of crowds*. Anchor, 2005.

[Sza16] Attila Szantner. Massively multiplayer online science. In *Levelling Up: The Cultural Impact of Contemporary Videogames*, pages 103–110. Brill, 2016.

[THG94] Julie D Thompson, Desmond G Higgins, and Toby J Gibson. Clustal w: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice. *Nucleic acids research*, 22(22):4673–4680, 1994.

[VA06] Luis Von Ahn. Games with a purpose. *Computer*, 39(6):92–94, 2006.

[vAD04a] Luis von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’04, page 319–326, New York, NY, USA, 2004. Association for Computing Machinery.

[VAD04b] Luis Von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326, 2004.

[VBC$^+$19] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

[VSP$^+$17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

[VW94] Martin Vingron and Michael S Waterman. Sequence alignment and penalty choice: Review of concepts, case studies and implications. *Journal of molecular biology*, 235(1):1–12, 1994.

[WDJ04] Guoli Wang and Roland L Dunbrack Jr. Scoring profile-to-profile sequence alignments. *Protein Science*, 13(6):1612–1626, 2004.

[WJ94] Lusheng Wang and Tao Jiang. On the complexity of multiple sequence alignment. *Journal of computational biology*, 1(4):337–348, 1994.

[WSK$^+$20] Jérôme Waldispühl, Attila Szantner, Rob Knight, Sébastien Caisse, and Randy Pitchford. Leveling up citizen science. *Nature Biotechnology*, 38(10):1124–1126, 2020.

[XZL15] Peipei Xia, Li Zhang, and Fanzhang Li. Learning similarity with cosine similarity ensemble. *Information Sciences*, 307:39–52, 2015.

[YEB11] Roman V Yampolskiy and Ahmed El-Barkouky. Wisdom of artificial crowds algorithm for solving np-hard problems. *International Journal of Bio-inspired computation*, 3(6):358–369, 2011.

[YSLD12] Sheng Kung Michael Yi, Mark Steyvers, Michael D Lee, and Matthew J Dry. The wisdom of the crowd in combinatorial problems. *Cognitive science*, 36(3):452–470, 2012.

[ZSWM00] Zheng Zhang, Scott Schwartz, Lukas Wagner, and Webb Miller. A greedy algorithm for aligning dna sequences. *Journal of Computational biology*, 7(1-2):203–214, 2000.

[ZWL$^+$20] Ruohan Zhang, Calen Walshe, Zhuode Liu, Lin Guan, Karl Muller, Jake Whritner, Luxin Zhang, Mary Hayhoe, and Dana Ballard. Atari-head: Atari human eye-tracking and demonstration dataset. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 6811–6820, 2020.

4

Player-Guided AI Outperforms Standard AI in Sequence Alignment Puzzles

Preface

After the publication of our first paper, we realized a few limitations concerning the determinism of the model presented. We wanted to conserve the stochastic nature of the solutions generated by players and also investigate comparisons with models based on reinforcement learning.

These and a number of other enhancements over the previous architecture have deepened our view on how player-generated data can be used to inform a strong algorithm, which not only closely emulates player behavior but also exceeds classical methods.

In this chapter, I present a novel architecture based on Generative Adversarial Imitation Learning, applied to the dataset provided by Borderlands Science, that allows maintaining the stochastic nature of human decision-making while achieving superior alignment results with respect to state-of-the-art methods.

The main novelty of the research is that this model, to the best of our knowledge, for the first time has been applied to Borderlands data in a way that grasps quite nicely the randomness of human decision-making processes.

This chapter, bar the present preface, was presented at the 2023 CM Collective Intelligence Conference.

The full list of authors for this work includes Renata Mutalova, Ro-

---

man Sarrazin-Gendron, Parham Ghasemloo, Eddie Cai, Gabriel Richard, Sébastien Caisse, Rob Knight, Mathieu Blanchette, Attila Szantner, and Jérôme Waldispühl.

Jérôme Waldispühl guided and supported the biological aspects of the project. Roman Sarrazin-Gendron contributed to the writing and performed background analysis. Gabriel Richard and Sébastien Caisse played a crucial role in the realization of Borderlands Science, an important part of this work.

My contributions towards them were the development and actual implementation of almost all coding activities related to the pre- and post-processing. I rewrote the algorithm that exists to find the optimal solution, making it run up to 100 times faster. I also implemented the GAIL architecture and a Deep Q-Network serving as a comparative baseline algorithm. I created the figures and tables presented in this section. This work was originally published as Mutalova, R., Sarrazin-Gendron, R., Ghasemloo Gheidari, P., Cai, E., Richard, G., Caisse, S., ... & Waldispühl, J. (2023, November). Player-guided AI outperforms standard AI in sequence alignment puzzles. In Proceedings of The ACM Collective Intelligence Conference (pp. 53-62).

4.1 Introduction

Citizen science is an approach that engages non-scientists in scientific projects, where they actively contribute their time, efforts, and expertise. It harnesses the collective power of people to assist in data collection, analysis, and interpretation. Citizen science projects can take the form of a game (Citizen science games, or CSGs), in which players become active contributors to real scientific research while enjoying gameplay. This approach has been leveraged to tackle complex problems such as Multiple Sequence Alignment (MSA) [KRK$^+$12].

The purpose of a MSA is to identify homologous regions of DNA, RNA, or protein sequences. In this process, gaps are inserted into the sequence to accurately represent insertions or deletions of nucleotides, reflecting evolutionary changes. Aligning the sequences facilitates their comparison and downstream analysis by revealing conserved motifs and structural insights. An MSA allows the identification of similarities and distinctions within multiple sequences, offering valuable insights into evolutionary connections and conserved regions among the sequences. Similar to many classical bioinformatics problems, MSA is known to be NP-hard [WJ94], which means that its computational complexity increases exponentially with the number of sequences to be aligned [WJ94].

Several approaches have been proposed to tackle this challenging problem. Dynamic programming serves as a prevalent approach in sequence alignment algorithms [NW70, SW$^+$81]. To address the computational challenge of the MSA problem, researchers have proposed alternative approaches and techniques as a potential solution. Examples include simulated annealing [KPC94], optimization-based algorithms [RK03, XC09, MJ03], heuristic methods [MNG$^+$15, WF93] and even quantum algorithms [HSX10]. Additionally, researchers have made multiple attempts to apply Artificial Intelligence (AI) techniques to this problem [MBD14, JJKR19, RSB18, Joe21, ZZL$^+$22, MSGC$^+$23]. However, these attempts have encountered issues related to convergence rate and alignment accuracy [MBD14, RSB18, ZZL$^+$22].

Crowd computing, and more particularly citizen science games, have emerged as a promising approach for tackling the MSA problem. One notable implementation of this concept is Phylo, released in 2010 [KRK$^+$12]. Phylo utilizes crowd computing principles to engage players in gameplay while simultaneously contributing to scientific research and genetic analysis. In Phylo, players are presented with a set of puzzle pieces representing homologous segments of DNA from various species. By aligning these sequences, players help researchers in the analysis of genetic data and the study of evolutionary relationships. The success of Phylo highlighted the potential of crowd computing in harnessing the power of human intuition and pattern recognition to address complex scientific problems like MSA.

Building on the success of Phylo was Borderlands Science (BLS), a minigame integrated into the popular video game Borderlands 3, released in 2020. The integration of BLS into the actual game has attracted a significant number of players compared to Phylo. Also a sequence alignment game but focused on smaller tasks, BLS offers the advantage of coming with tens of millions of player solutions, unlocking a unique potential to explore novel strategies not discovered by other algorithms, including AI. This raises the question of how we can effectively leverage the knowledge gained from players' participation in BLS and transfer these valuable ideas to tackle new alignment problems, ultimately advancing our ability to solve a wide range of alignment challenges beyond the game's scope.

One approach to amplify the analysis strategies in Borderlands Science is to incorporate Reinforcement Learning (RL), a prominent field within the realm of Artificial Intelligence (AI). RL algorithms empower AI agents to learn from gathered data, continually refining their decision-making processes through feedback and rewards.

DQN, or Deep Q-Network, is a fundamental reinforcement learning algorithm that utilizes deep neural networks to learn optimal policies. By iteratively updating the network based on interactions with the environment, DQN enables AI agents to improve their decision-making abilities and solve complex problems effectively. Nevertheless, the DQN model exhibits limitations, including high computational demands that restrict its scalability for larger-scale problems, and a tendency to overestimate action values, resulting in suboptimal performance [MKS$^+$15].

To overcome the mentioned drawbacks, improve the performance of AI models, and mitigate the complexities of data aggregation in crowd computing systems, we propose the utilization of a hybrid approach, Generative Adversarial Imitation Learning (GAIL) [HE16]. GAIL enhances the performance of the basic AI algorithm by incorporating data from players, allowing for the integration of human knowledge and experiences. This integration enables the algorithm to leverage extra insights and enhance its ability to make informed decisions. Consequently, this approach effectively integrates the strengths of AI and crowd computing.

4.1.1 Hypotheses

This paper presents a comprehensive analysis of the potential benefits of leveraging crowdsourcing to enhance the efficiency of RL approaches for solving challenging puzzles from Borderlands Science.

Specifically, we aim to investigate three hypotheses:

H1: Does incorporating data from players through a human-in-the-loop approach improve the performance of the basic AI model in solving the BLS puzzles?

The primary objective of H1 - is to investigate the potential performance enhancement of the DQN algorithm by incorporating optimal strategy trajectories provided by players. To examine this, we employ a GAIL approach that combines player solutions with the DQN model. This comparative analysis aims to identify potential similarities and distinctions among these methodologies.

H2: Can the proposed human-in-the-loop algorithm generate solutions that are more optimal and exhibit reduced randomness in terms of the score-gap trade-off when compared to the basic AI model?

This Hypothesis aims to compare the solutions provided by DQN and GAIL in terms of game score, gaps used, and uniqueness. In this part of the research, we analyze which of the generated solutions are scientifically significant and involve a balance between the quantity of used gaps and the game score achieved. Our hypothesis suggests that the GAIL-based algorithm has the potential to generate more optimal solutions compared to DQN.

H3: To what extent does the proposed human-in-the-loop algorithm compare to the performance level of players?

In this phase of the research, we conduct an analysis to determine the similarity between the generated solutions and the solutions provided by the players. Additionally, we examine the common characteristics shared by these solutions and assess the extent to which they align with each other. This analysis allows us to gain insights into the level of agreement and consistency between the generated solutions and the player-provided solutions.

4.1.2 Game Design

Borderlands Science features a puzzle game where each column within the game corresponds to a distinct fragment of DNA sequence. These specific DNA sequences originate from the American Gut Project [MHD$^+$18], a component of The Microsetta Initiative. Subsequently, these sequences were

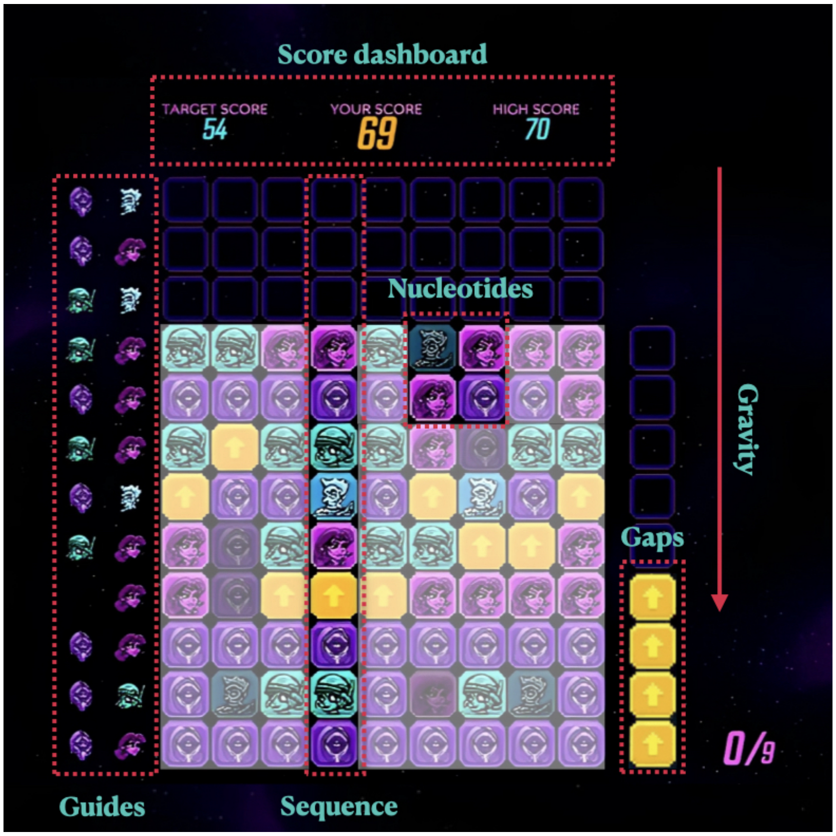

Figure 4.1: Borderlands Science interface: the alignment sequences are displayed vertically, represented by tiles of four different colors corresponding to the nucleotides A, C, G, and T. On the left side, there are guides indicating which tiles should be matched in each row to earn points. To maximize the reward, the user needs to strategically insert gaps (represented by yellow tiles) within the alignment. However, there is a restriction on the number of gaps allowed, and completing full rows grants an additional bonus.

aligned using the PASTA algorithm [NMKW15]. This alignment procedure generated a tight alignment, which then served as the scaffold for the creation of the puzzles.

The puzzles in Borderlands Science feature a grid filled with coloured blocks, with each column representing a sequence. The game typically consists of 5 to 12 columns, and their difficulty level determines the number and length of the DNA sequences. The objective for players is to align the bricks to the guides located on the left side of the screen in **Figure 4.1**. These guides offer information about the consensus nucleotides of the corresponding column of the PASTA alignment. To assist in the alignment process, players are given a limited number of yellow tokens representing gaps in MSA. These tokens can be inserted between the bricks to improve the alignment according to the guides.

The main challenge of the game is to achieve a target score that exceeds the score achieved by the default greedy player. Furthermore, the game showcases the highest score previously attained to encourage players to enhance their performance.

4.1.3 Pareto Optimality

In Borderlands Science puzzles, we collect multiple solutions per puzzle, enabling us to explore the solution space and identify Pareto optimal solutions. A solution is considered Pareto optimal if it is not strictly worse than another solution. These are the solutions that maximize trade-offs between used gaps and game score.

4.2 Related work

4.2.1 Reinforcement learning

Reinforcement learning, a widely employed method in the field of artificial intelligence, enables the development of intelligent agents capable of decision-making and learning within various domains, including games [Szi12, GEFL18, RSdW$^+$18, MKS$^+$13, VBC$^+$19, BBC$^+$19]. RL can be employed to train game-playing agents to accomplish specific tasks, such as winning a game or achieving a particular objective. These agents can learn from their own experiences and interactions with the game environment, leading to an improvement in their performance over time. RL has been successfully utilized across a range of game genres, from traditional board games like chess and Go [SHM$^+$16] to contemporary video games like Dota 2 [BBC$^+$19] and StarCraft II [VBC$^+$19].

RL algorithms, such as the Deep Q-Network, have been applied to Atari games which share with BLS a common aspect of utilizing gaming mechanics for problem-solving, yielding impressive results but failing to surpass human-level performance [MKS$^+$13]. One of the key reasons for this discrepancy is the complexity and diversity of the real world. RL algorithms heavily rely on extensive training in a simulated environment to learn optimal strategies, but they struggle when faced with the vast range of uncertainties, dynamics,

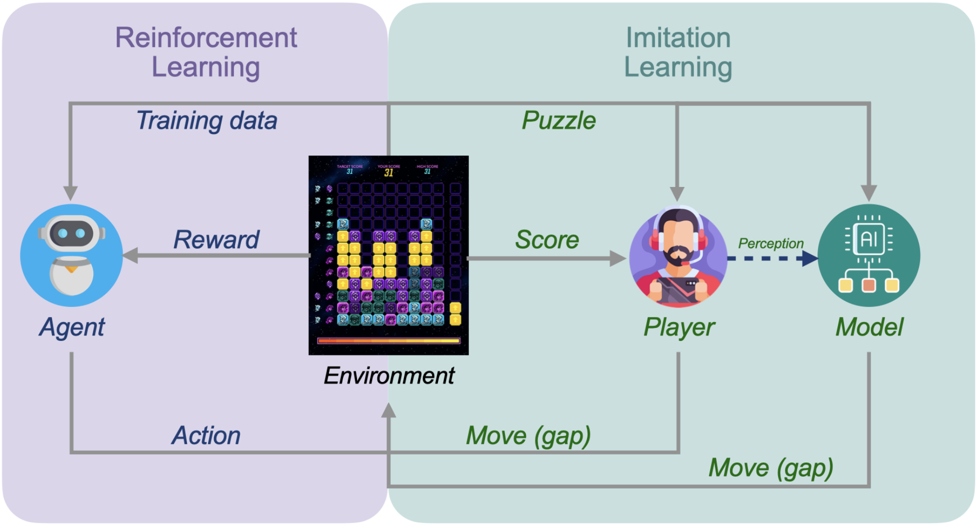

Figure 4.2: The illustration depicts two approaches in the field of machine learning: Reinforcement Learning and Imitation Learning. **On the left-hand side**, RL is illustrated as a process where an agent interacts with an environment to learn a policy that maximizes a reward signal. The agent takes actions based on the current state of the environment and receives a reward based on the outcome of the action. The policy is updated using the rewards obtained and the next state of the environment. **On the right-hand side**, imitation learning is illustrated as a process where an agent learns from expert demonstrations. The expert provides a set of examples of actions taken in different states of the environment. The agent learns to mimic the expert's behavior by learning a mapping between the current state and the action taken by the expert in that state. This mapping is learned using techniques such as supervised learning.

and unstructured situations that humans effortlessly navigate. Humans have natural cognitive abilities, intuition, and common sense reasoning, which allow them to adapt and apply knowledge in various situations. In contrast, RL algorithms often need large amounts of data and computational resources to achieve similar levels of performance.

4.2.2 Reinforcement learning for MSA

Multiple attempts have been made by researchers to utilize RL to address the challenge of MSA problem [MBD14, JJKR19, RSB18, Joe21, ZZL^{+}22]. Many of these papers leverage the DQN algorithm as the primary approach for addressing MSA challenges [MBD14, JJKR19, Joe21]. However, researchers often encounter common challenges associated with employing the DQN algorithm for MSA, including issues such as suboptimal convergence and the constraint of handling a restricted quantity of sequences with limited lengths [MBD14, RSB18]. Dealing with a large number of sequences can be a major challenge for models due to the exponential growth of states and actions. As a result, many studies have focused on pairwise sequence alignment to avoid this complexity [SJS^{+}21, SC21, LT23]. Although there have been many proposed RL methods for multiple sequence alignment (MSA), there is still significant room for improving their accuracy.

Can the integration of human intelligence and the diversity of machines overcome the limitations faced by RL algorithms and mitigate the aforementioned challenges?

4.2.3 Deep Q-Network

Deep Q-Network (DQN) is a reinforcement learning algorithm that combines the Q-learning algorithm with deep neural networks to learn policies from raw sensory inputs, allowing agents to learn to play games or perform other tasks through trial and error [MKS^{+}15]. DQN has been successful in learning to play a wide range of video games at a superhuman level, demonstrating its ability to learn complex behaviors from high-dimensional sensory inputs [MKS^{+}15]. The algorithm has also been extended to other domains, such as robotics [LFDA16] and finance [DBK^{+}16], where it has shown promise in solving complex problems. Additionally, DQN has been combined with other reinforcement learning algorithms, such as policy gradient methods [WSH^{+}16] and actor-critic methods [LHP^{+}15], to achieve even better performance.

4.2.4 Generative Adversarial Imitation Learning

Generative Adversarial Imitation Learning (GAIL) and Reinforcement Learning (RL) [HE16, WVO12] are machine learning algorithms that can be applied for diverse tasks. Despite some similarities, there exist fundamental differences between the two approaches.

GAIL is an imitation learning method that involves training a generative model to create sequences similar to input sequences. This approach is particularly effective for scenarios where expert behavior is available and can be imitated. In the context of decision-making tasks, GAIL employs a generator-discriminator setup to guide the policy network towards replicating the expert's behavior. This is reminiscent of the broader Generative Adversarial Networks (GANs) framework, where GAIL's focus is on imitating expert actions rather than generating creative samples. This approach has found notable application in reinforcement learning tasks such as robotic control and autonomous driving [FDSF18, BWP$^+$22, KMWK17], enabling the learning of complex behaviors from expert demonstrations.

On the other hand, RL is a learning technique that involves training an agent to perform actions in an environment to maximize a reward signal. In contrast to GAIL's supervised learning nature, RL is a quintessential form of reinforcement learning where an agent interacts with an environment, takes actions, and learns through trial and error. The agent explores different actions and receives feedback in the form of rewards, adapting its policy to optimize its decision-making strategy over time. This optimization, performed by adapting the policy, is directly linked to the rewards received from the environment, as the agent aims to enhance its action-selection strategy to achieve higher cumulative rewards.

The primary difference between GAIL and RL is that GAIL is a type of supervised learning, whereas RL is a form of reinforcement learning. In GAIL, the model is trained using examples of the desired outcome/behavior, whereas, in RL, the model learns by exploring and receiving feedback from the environment shown in **Figure 4.2**.

Although GAIL is a relatively new method, there have been some recent studies exploring its application in various fields. For example, GAIL has been used to model human vehicle driving behavior [BWP$^+$22] to improve simulations for vehicle safety validation processes, and to discover diverse strategies for human-robot interaction by observing human-human behavior [WPDX$^+$22]. However, to the best of our knowledge, no previous studies have investigated the use of GAIL for tackling the MSA problem or for game-related tasks.

4.3 Data

The Borderlands Science game uses human gut microbe genome fragments that have been pre-aligned with PASTA[MNW14] to create puzzles for players to solve. Players aim to optimize a bi-objective function by maximizing the number of matches while minimizing the number of gaps. In this paper, the focus is on a sample of 25,000 puzzles with over 1 million solutions. The quality of each solution is estimated based on its distance from the Pareto front, which is the set of solutions that optimize the bi-objective function between score and used gaps.

We selected the top-performing human solutions by aggregating puzzle solutions and identifying those close to the Pareto front. To determine their inclusion, we calculated the proportional horizontal Pareto distance of each solution to a puzzle. This distance represents the improvement in score compared to the worst human solution, divided by the improvement achieved by the best score for the same number of gaps (considered Pareto-optimal) over the worst human solution. Solutions with a distance exceeding 0.7 were excluded based on visual clustering assessment by human experts. As a result, 53.4% of solutions were filtered out.

4.4 Methods

In this section, we will discuss the implementation details of the algorithms.

4.4.1 DQN Implementation

The DQN algorithm focuses on training a neural network to estimate the Q-values corresponding to various actions in a given state. The agent selects the action with the highest predicted Q-value, enabling the acquisition of a policy that maximizes the expected cumulative reward over time. Additionally, the incorporation of Memory Replay proves significant, wherein the agent's experiences are stored in a memory buffer, facilitating learning from past interactions. Instead of relying solely on the most recent experience, the agent randomly samples and learns from a batch of experiences from the memory buffer. This mechanism enhances learning efficiency and decision-making capabilities, ultimately leading to a stabilized learning process shown in **Figure 4.3**.

In our experiment, the architecture consists of a transformer attention mechanism, which is comprised of several layers of self-attention and feed-forward neural networks with Sinusoid Positional Encoding. To improve training stability and prevent overfitting, the model utilizes several techniques, including layer normalization, residual connections, and dropout. Additionally, a replay memory is employed to store and sample past moves sequences during training (memory size = 1000). We limited the number

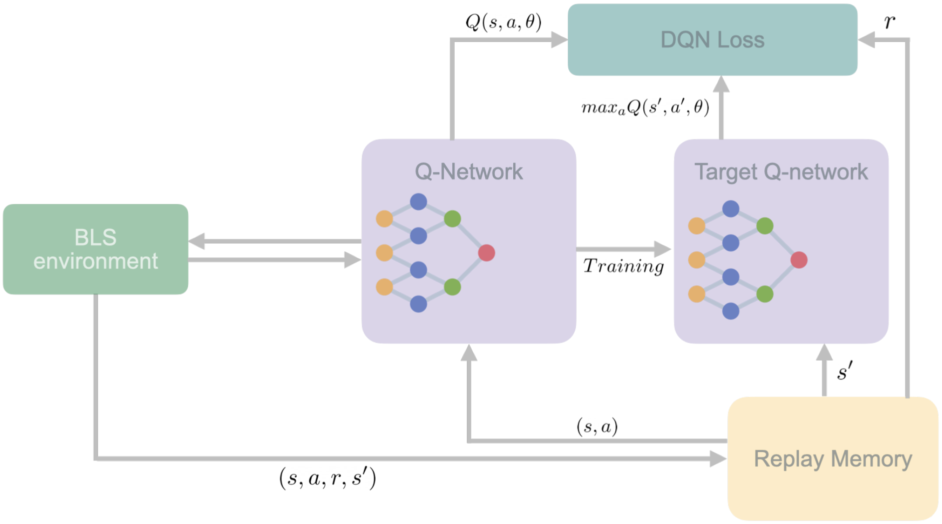

Figure 4.3: Deep Q-Network Architecture. The neural network is trained to approximate the Q-values, allowing the agent to estimate the expected future rewards for different actions in a given state, enabling informed decision-making. Replay buffer: The replay buffer in the DQN stores the agent's experiences, creating a dataset from which random samples are drawn during training, allowing the agent to learn from a diverse set of experiences. In the DQN, the iterative process involves repeatedly interacting with the environment, collecting experiences, storing them in the replay buffer, and periodically sampling from the buffer to train the neural network, updating its weights to minimize the discrepancy between predicted and actual Q-values, resulting in an improved policy over time.

of possible actions according to the game mechanism. DQN has methods for selecting an action based on an epsilon-greedy policy, predicting the best action for a given state, and updating the network weights using the replay memory and a target network. The main hyperparameters used in this implementation are the learning rate (0.001), batch size (128), discount factor (gamma = 1), the exploration-exploitation tradeoff parameter (epsilon = 0.8), and decrement iteration (5).

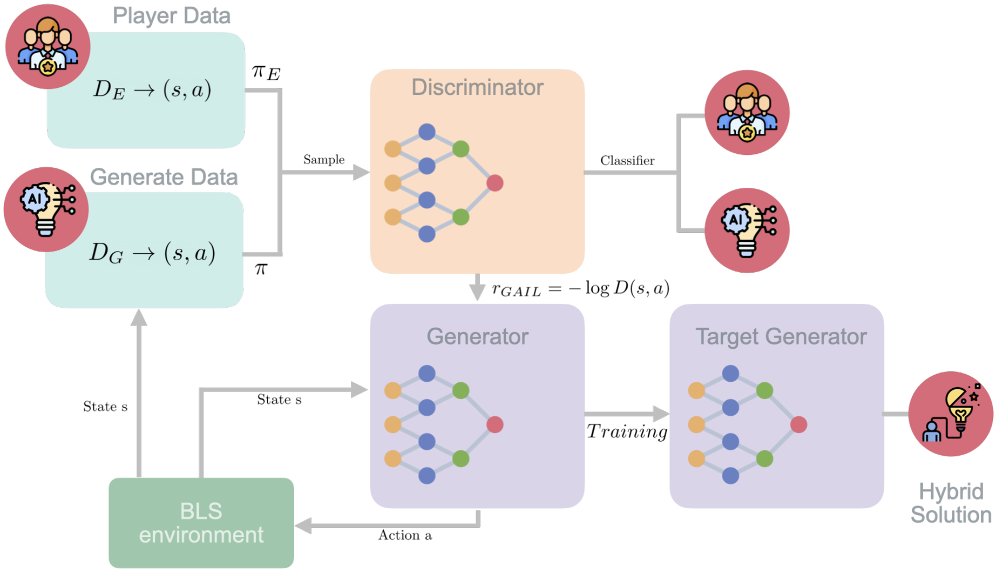

Figure 4.4: Generative Adversarial Imitation Learning Architecture: a novel approach to reinforcement learning with human expertise integration. Generator: the generator learns to generate player strategies, that closely resemble the demonstrations or desired output. Discriminator: the discriminator is responsible for distinguishing between the generated data from the generator and the real players demonstrations, helping to provide feedback on the quality of the generated data. The iterative process involves training the generator and discriminator networks in an adversarial manner, where the generator aims to fool the discriminator while the discriminator continuously improves its ability to differentiate between real and generated data, resulting in a refined generator that produces increasingly realistic outputs.

4.4.2 GAIL Implementation

In our experiment, we extend the architecture shown in Figure 4.3 to the circuit in Figure 4.4 and GAIL involves the following steps. The GAIL algorithm involves training a generative model to mimic the behavior of players, enabling the acquisition of policies through adversarial learning. First, data from players demonstrations of the task the agent will be performing needs to be collected. This data is used to train the discriminator network in the Generative adversarial network (GAN) [WGD$^+$17, GPAM$^+$20]. Second, the generator (base DQN agent) and discriminator networks in the GAN are trained using expert data. The agent interacts with the BLS environment by choosing actions based on the DQN agent's policy. The resulting state, reward, and done flag are observed and stored in the Replay Buffer. The DQN agent is updated by sampling a batch of experiences from the replay buffer. The Q-network weights are updated using the DQN loss and optimizer (Adam with a learning rate of 0.001). The discriminator is updated by sampling a batch of expert demonstrations and a batch of agent trajectories generated using the current DQN agent. The DQN agent can leverage its ability to learn from trial and error to improve its policy, while GAIL provides a framework for imitation learning by utilizing the discriminator network to provide feedback and guide the generator's training. The discriminator's weights are updated using the GAIL loss (Binary Cross-Entropy loss) and optimizer (Adam, with a learning rate of 0.0001). The DQN agent is further updated using the GAIL reward with Discount Factor = 0.99. The general scheme is shown in **Figure 4.4**.

4.5 Results and Discussion

In this section, we share the findings from our examination of the complete dataset gathered from the Borderlands Science game. We present the solutions provided by all players, separated the Pareto optimal solutions contributed by the players, as well as the solutions generated by the DQN in GAIL. Notably, GAIL is a player-guided model, while DQN solely relies on an AI-driven approach.

4.5.1 Testing Hypothesis 1

In order to evaluate the impact of integrating player data on the effectiveness of the basic DQN model, we compared two strategies, DQN and GAIL. In **Figure 4.5** we can compare algorithms performance in terms of game score, a score obtained primarily from correctly matching bricks to guides. We show that classical Reinforcement learning methods such as DQN are unable to achieve human-level performance, even with precise parameter tuning. In the context of the game score, the performance of this model barely passes the minimum required score. By contrast, the results obtained with the GAIL method are at least close to the average human performance.

Conclusion: Our findings demonstrate that the performance of the basic AI model in solving the BLS puzzles is enhanced by incorporating data from players through a human-in-the-loop approach.

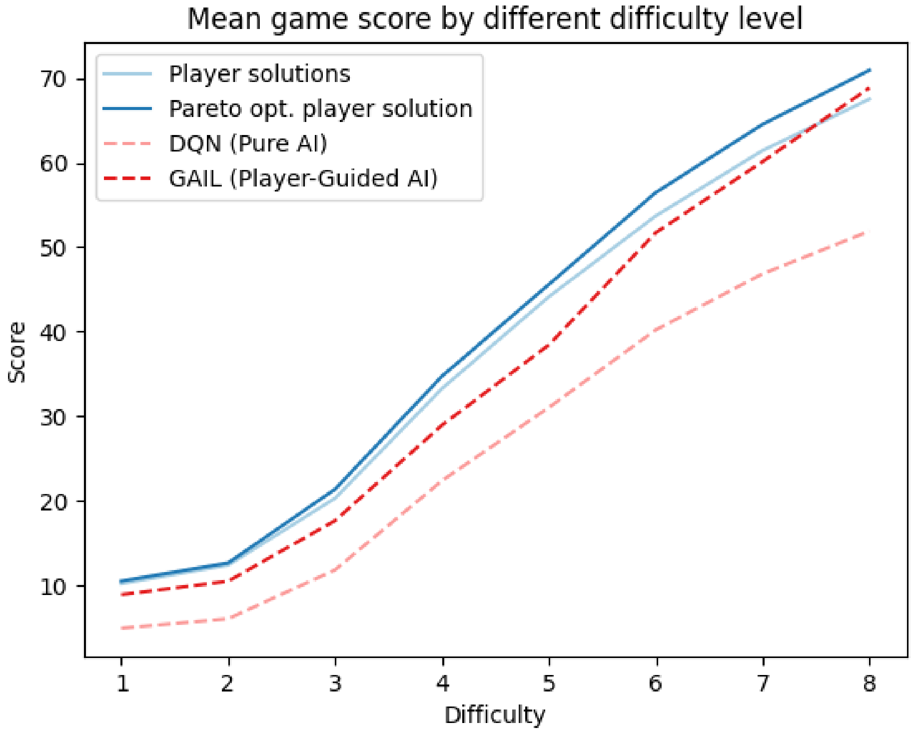

Figure 4.5: Comparison of algorithms performance in terms of game score: the average game score obtained by the players (All players solution and Pareto optimal solution) and models (DQN and GAIL) for puzzles at increasing difficulty levels. The game score is primarily determined by the correct matching of bricks to guides within the puzzles. Blue and solid line style - Players solutions, red and dashed line style - proposed player-guided AI and uninformed AI

4.5.2 Testing Hypothesis 2

To evaluate the effectiveness of the proposed GAIL algorithm in generating solutions that are more optimal and display reduced randomness in terms of the score-gap trade-off, we computed the percent of Pareto optimal solutions per original solutions and generated. In **Figure 4.6** we observe that GAIL is capable of generating a higher proportion of optimal solutions which is define in the section **Data**. However, it is worth noting that as the puzzle complexity increases, this percentage tends to decrease. Nonetheless, even with the decrease, the GAIL algorithm still produces a larger proportion of solutions that closely resemble real solutions.

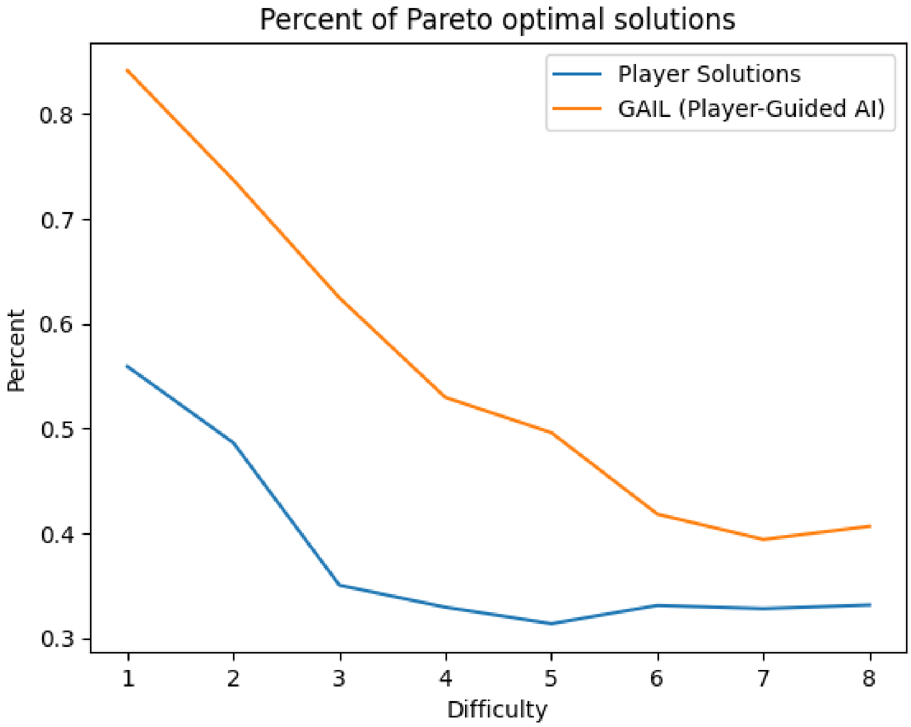

Figure 4.6: Comparison of the percentage of solutions that are Pareto-optimal between player solutions and generated solutions

Second, we conducted an evaluation of the score-gap trade-off for all algorithms. **Figure 4.7** illustrates that the DQN algorithm tends to utilize fewer gaps compared to the other algorithms. On the other hand, GAIL demonstrates the lowest gap usage among all the described algorithms. However, **Figure 4.5** highlights that, on average, DQN achieves significantly lower scores compared to GAIL.

Third, to evaluate the reduction of noise in solutions, we have given a detailed analysis of one puzzle on **Figures 4.8, 4.9, 4.10**. The results on 4.10 indicate that the GAIL algorithm has the ability to generate solutions that closely to the Pareto optimal front, while effectively minimizing the number of solutions with noise in comparison to the DQN algorithm.

In **Figure 4.8**, we depict solutions or an integral representation of a solution for a single puzzle. It is apparent that the average of the DQN

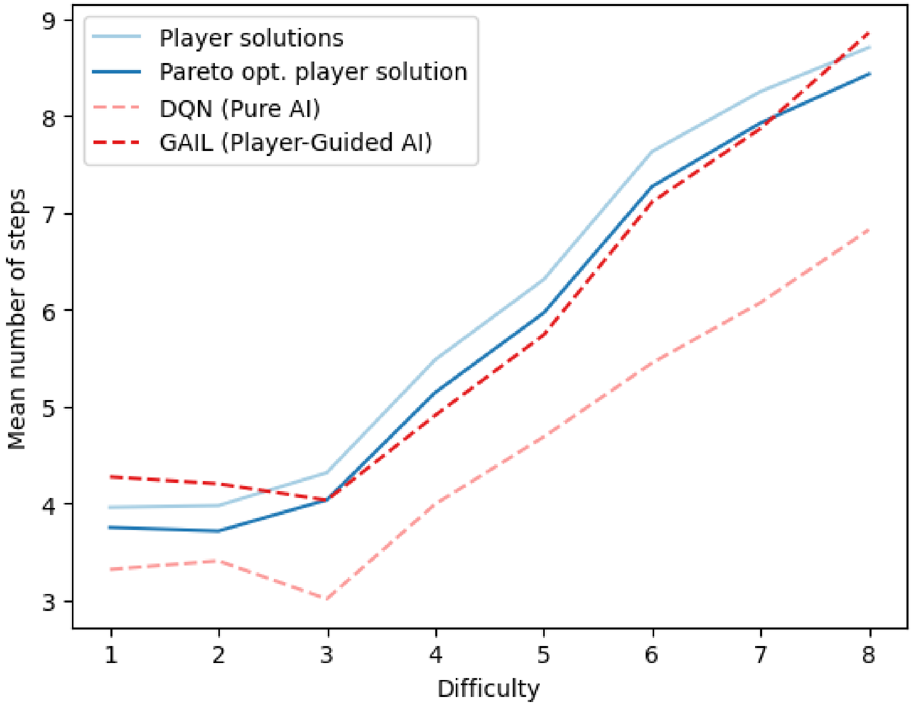

Figure 4.7: Score-Gap trade-off evaluation for all algorithms: the average number of gaps used by the players and models for puzzles of increasing difficulty levels. Blue and solid line style - players solutions, red and dashed line style - proposed player-guided AI and uninformed AI

solution exhibits a significant amount of noise and does not align well with the other solutions. The GAIL solutions demonstrate good agreement with the overall solutions and display significantly less noise compared to the DQN solution.

Figure 4.9 demonstrates that the score distribution for DQN has a wide range and exhibits a distribution that is close to normal, centered around approximately score equals 30 points. On the other hand, GAIL showcases a significant variation in solutions with a range extending beyond the score of 30 points.

Conclusion: By incorporating player solutions, we can effectively enhance the RL model, resulting in reduced noise and the generation of compelling options that strike a balance between the number of used gaps and the game score.

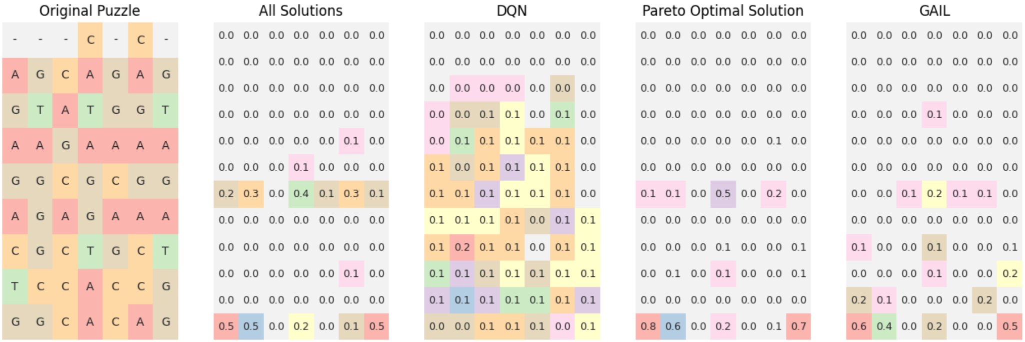

Original Puzzle

-	-	-	C	-	C	-
A	G	C	A	G	A	G
G	T	A	T	G	G	T
A	A	G	A	A	A	A
G	G	C	G	C	G	G
A	G	A	G	A	A	A
C	G	C	T	G	C	T
T	C	C	A	C	C	G
G	G	C	A	C	A	G

All Solutions

0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.1	0.0
0.0	0.0	0.0	0.1	0.0	0.0	0.0
0.2	0.3	0.0	0.4	0.1	0.3	0.1
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.1	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.5	0.5	0.0	0.2	0.0	0.1	0.5

DQN

0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.1	0.1	0.0	0.1	0.0
0.0	0.1	0.1	0.1	0.1	0.1	0.0
0.1	0.0	0.1	0.1	0.1	0.1	0.0
0.1	0.1	0.1	0.1	0.1	0.1	0.0
0.1	0.1	0.1	0.1	0.0	0.1	0.1
0.1	0.2	0.1	0.1	0.0	0.1	0.1
0.1	0.1	0.1	0.1	0.1	0.1	0.1
0.1	0.1	0.1	0.1	0.1	0.1	0.1
0.0	0.0	0.1	0.1	0.1	0.0	0.1

Pareto Optimal Solution

0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.1	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.1	0.1	0.0	0.5	0.0	0.2	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.1	0.0	0.0	0.1
0.0	0.1	0.0	0.1	0.0	0.0	0.1
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.8	0.6	0.0	0.2	0.0	0.1	0.7

GAIL

0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.1	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.1	0.2	0.1	0.1	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.1	0.0	0.0	0.1	0.0	0.0	0.1
0.0	0.0	0.0	0.1	0.0	0.0	0.2
0.2	0.1	0.0	0.0	0.0	0.2	0.0
0.6	0.4	0.0	0.2	0.0	0.0	0.5

Figure 4.8: Solutions comparison for a single puzzle. The percentage distribution of steps for all predicted solutions. The layout is similar to the game: the sequences are presented vertically. The GAIL solutions demonstrate good agreement with the all-players and Pareto-optimal solutions and display significantly less noise in comparison to the DQN solution.

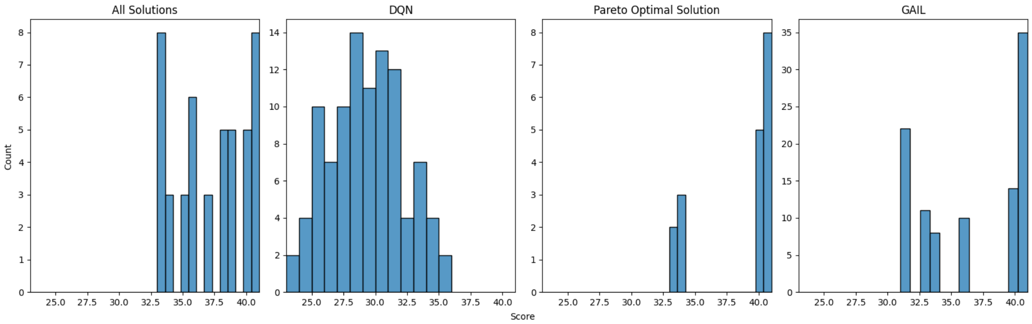

Figure 4.9: Score distribution for a single puzzle. The x-axis represents the number of gaps used. The y-axis represents the Game Score, which measures the overall quality or effectiveness of each solution. The score distribution provides insights into the range and variability of the achieved scores for each algorithm and players solutions.

4.5.3 Testing Hypothesis 3

To explore if GAIL solutions match the performance level of players we can compare them in terms of the game score. **Figure 4.5** indicates that the highest scores achieved mostly by players. Also, for puzzles with high difficulty, the performance of our GAIL approach is comparable to All-players solutions, and on the highest difficulty even surpasses. However, **Figure 4.8**,

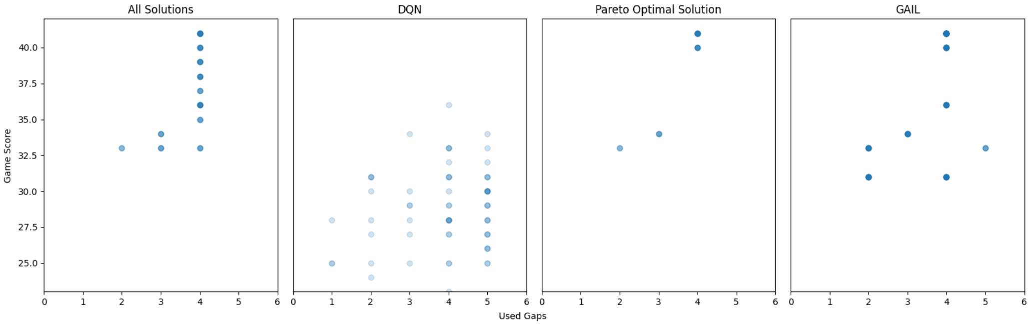

Figure 4.10: Trade-off between gaps used and game score for a single puzzle. The x-axis represents the number of gaps used. The y-axis represents the game score, which measures the overall quality or effectiveness of each solution. The findings demonstrate that the GAIL algorithm can generate solutions that are close to Pareto optimal solutions and significantly reduce the number of noisy solutions compared to DQN. The opacity of the points represents the frequency of the solution. The varying number of points reflects the finite number of solutions contributed by players (approximately 25 for each puzzle) while we conducted 100 runs of the generated solutions by DQN and GAIL to assess prediction variability.

we can observe the similarities between GAIL solutions and the decisions made by the players. The GAIL solutions form a distribution that closely resembles the overall solutions, with only a few distinct solutions deviating from the pattern.

In **Figure 4.9**, we can observe that both the players and GAIL exhibit a general tendency to achieve scores above 30 for this puzzle. Furthermore, the distribution of the players' actions in **Figure 4.9** closely aligns with the distribution of GAIL solutions.

Conclusion: Our findings indicate that a model trained using this approach can demonstrate performance levels that are close to those of humans, exhibiting a similar degree of variability and quality in decision-making.

4.6 Discussion

The outcomes of our study answer the three hypotheses we formulated, providing further evidence to support the claim that the performance of a basic AI model in solving BLS puzzles can be enhanced by incorporating data from players through a human-in-the-loop approach. Our findings indicate that this approach significantly enhances the performance of the AI model.

By incorporating insights and strategies from human players, we observed a substantial increase in the game score achieved by the AI model. This suggests that the human-in-the-loop approach allows the AI model to benefit from the expertise of experienced players, enabling it to adopt a distinct problem-solving strategy that leads to higher scores compared to relying solely on the basic AI model.

Furthermore, we investigated whether the proposed human-in-the-loop algorithm could generate solutions that are more optimal and exhibit reduced randomness in terms of the score-gap trade-off, when compared to the basic AI model. Our results demonstrated that incorporating player solutions effectively enhanced the RL model, resulting in reduced noise and the generation of compelling options. The algorithm struck a better balance between utilizing gaps and maximizing the game score.

Additionally, we explored the extent to which the proposed human-in-the-loop algorithm matched the performance level of players. Compared to humans, the player-guided AI model exhibited respectable variability and quality in decision-making, suggesting that it successfully captured the gameplay characteristics observed in human players. This approach demonstrates it is possible to extend the usefulness of player contributions to a citizen science project beyond the limits of the original problem. Unlocking this potential is a key step in charting a course for next-generation citizen science, as collecting solutions from millions of participants for a single problem could be overkill, but having the ability to leverage these contributions into training robust AI systems to multiply their impact will allow scientists to fully take advantage of the heterogeneity of human contributions to analyze large scale data sets.

4.7 Limitations

There are however a few limitations to the methods and results presented here, both in terms of how the system is trained and what it is trained on.

The limitations of the Borderlands Science game design for AI stem from its mechanics and puzzle specifics. AI agents trained on Borderlands Science may face challenges in adapting to larger state and action spaces, which can impede performance and generalization. The game may not fully capture the complexity and scalability required for demanding AI tasks, making it difficult to transfer learned policies to different game environments with diverse rules, actions, and dynamics.

Another drawback is related game it self Borderlands Science game is that it doesn't encourage players to optimize the bi-objective function for score-gaps trade-off. Although they have a limited number of spaces, there is no incentive for them to leave any spaces unused. To address this, we introduce a change: players are presented with varying numbers of gap tokens for the same puzzle. We then filter their solutions to find those that are close to the Pareto front, resulting in a significant set of nearly optimal solutions. Furthermore, it is important to note that the use of Pareto optimality as a proxy for identifying good solutions to the Multiple Sequence Alignment (MSA) problem does not guarantee the discovery of the absolute best solution for that problem.

Another drawback of the Borderlands Science game is that it does not include a direct incentive for players to optimize the score-gaps trade-off. Despite having a limited number of gaps, there is no motivation for players to leave any gaps unused. To address this issue, players are presented with varying numbers of gap tokens for the same puzzle. Subsequently, we filter their solutions to identify those that are close to the Pareto front, resulting in a significant set of nearly optimal solutions. However, it is important to note that the use of Pareto optimality as a proxy for identifying good solutions to the Multiple Sequence Alignment (MSA) problem does not guarantee the discovery of the absolute best solution for that problem.

Several important limitations are associated with the models. While DQN is effective in game-based reinforcement learning, it possesses certain limitations. It can be unstable, overestimate action values, and result in suboptimal policies and poor convergence. Achieving a balance between exploration and exploitation is challenging for DQN, as it tends to prioritize exploitation, potentially missing out on optimal strategies. Nevertheless, researchers are actively working to enhance DQN's performance in gaming applications. Additionally, GAIL may encounter difficulties in generalizing to unseen or complex scenarios, as it primarily learns from provided expert demonstrations, leading to limited diversity and suboptimal behaviour when generating new samples.

4.8 Conclusion

In conclusion, our research provides compelling evidence that integrating player data through a human-in-the-loop approach significantly improves the performance of AI models in solving BLS puzzles. The incorporation of player solutions effectively enhances the AI standard model, reducing noise and generating compelling options that achieve a harmonious balance between used gaps and game score. Our findings underscore the remarkable potential of using games as an effective method to collect valuable information for training AI models. Moreover, our study aligns with a wealth of existing evidence demonstrating the superiority of human + AI collaboration over AI alone. By leveraging the collective intelligence and decision-making capabilities of both humans and AI, we can unlock new frontiers of problem-solving and achieve new performance levels.

Bibliography

[BBC+19] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

[BWP+22] Raunak Bhattacharyya, Blake Wulfe, Derek J Phillips, Alex Kuefler, Jeremy Morton, Ransalu Senanayake, and Mykel J Kochenderfer. Modeling human driving behavior through generative adversarial imitation learning. *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[DBK+16] Yue Deng, Wei Bao, Hongwei Kong, Yufei Ren, Yong Dai, and Ruoyu Wang. Deep direct reinforcement learning for financial signal representation and trading. In *Proceedings of the IEEE Conference on Computational Intelligence for Financial Engineering & Economics*, pages 1–7. IEEE, 2016.

[FDSF18] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Learning temporal strategic relationships using generative adversarial imitation learning. *arXiv preprint arXiv:1805.04969*, 2018.

[GEFL18] Abhishek Gupta, Benjamin Eysenbach, Chelsea Finn, and Sergey Levine. Continuous adaptation via meta-learning in nonstationary and competitive environments. pages 1882–1891, 2018.

[GPAM+20] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and

Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

[HE16] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29:4565–4573, 2016.

[HSX10] Hong-Wei Huo, Vojislav Stojkovic, and Qiao-Luan Xie. A quantum-inspired genetic algorithm based on probabilistic coding for multiple sequence alignment. *Journal of bioinformatics and computational biology*, 8(01):59–75, 2010.

[JJKR19] Reza Jafari, Mohammad Masoud Javidi, and Marjan Kuchaki Rafsanjani. Using deep reinforcement learning approach for solving the multiple sequence alignment problem. *SN Applied Sciences*, 1(6):1–12, 2019.

[Joe21] Roman Joeres. Multiple sequence alignment using deep reinforcement learning. *SKILL 2021*, 2021.

[KMWK17] Alex Kuefler, Jeremy Morton, Tim Wheeler, and Mykel Kochenderfer. Imitating driver behavior with generative adversarial networks. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 204–211. IEEE, 2017.

[KPC94] Jin Kim, Sakti Pramanik, and Moon Jung Chung. Multiple sequence alignment using simulated annealing. *Bioinformatics*, 10(4):419–426, 1994.

[KRK^{+}12] Alexander Kawrykow, Gary Roumanis, Alfred Kam, Daniel Kwak, Clarence Leung, Chu Wu, Eleyine Zarour, Phylo players, Luis Sarmenta, Mathieu Blanchette, et al. Phylo: a citizen science approach for improving multiple sequence alignment. *PloS one*, 7(3):e31362, 2012.

[LFDA16] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. In *Proceedings of the 29th International Conference on Neural Information Processing Systems*, pages 3572–3580, 2016.

[LHP+15] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *Proceedings of the International Conference on Learning Representations*, 2015.

[LT23] Aryan Lall and Siddharth Tallur. Deep reinforcement learning-based pairwise dna sequence alignment method compatible with embedded edge devices. *Scientific Reports*, 13(1):2773, 2023.

[MBD14] Ioan-Gabriel Mircea, Maria-Iuliana Bocicor, and Andra. Dîincu. On reinforcement learning based multiple sequence alignment. *Studia Universitatis Babes-Bolyai, Informatica*, 59(2), 2014.

[MHD+18] Daniel McDonald, Embriette Hyde, Justine W Debelius, James T Morton, Antonio Gonzalez, Gail Ackermann, Alexander A Aksenov, Bahar Behsaz, Caitriona Brennan, Yingfeng Chen, et al. American gut: an open platform for citizen science microbiome research. *Msystems*, 3(3):e00031–18, 2018.

[MJ03] Jonathan Moss and Colin G Johnson. An ant colony algorithm for multiple sequence alignment in bioinformatics. In *Artificial Neural Nets and Genetic Algorithms: Proceedings of the International Conference in Roanne, France, 2003*, pages 182–186. Springer, 2003.

[MKS+13] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[MKS+15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

[MNG+15] Siavash Mirarab, Nam Nguyen, Sheng Guo, Li-San Wang, Junhyong Kim, and Tandy Warnow. Pasta: ultra-large multiple sequence alignment for nucleotide and amino-acid sequences. *Journal of Computational Biology*, 22(5):377–386, 2015.

[MNW14] Siavash Mirarab, Nam Nguyen, and Tandy Warnow. Pasta: ultra-large multiple sequence alignment. In *Research in Computational Molecular Biology: 18th Annual International Conference, RECOMB 2014, Pittsburgh, PA, USA, April 2-5, 2014, Proceedings 18*, pages 177–191. Springer, 2014.

[MSGC$^+$23] Renata Mutalova, Roman Sarrazin-Gendron, Eddie Cai, Gabriel Richard, Parham Ghasemloo Gheidari, Sébastien Caisse, Rob Knight, Mathieu Blanchette, Attila Szantner, and Jérôme Waldispühl. Playing the system: Can puzzle players teach us how to solve hard problems? In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2023.

[NMKW15] Nam-phuong D Nguyen, Siavash Mirarab, Keerthana Kumar, and Tandy Warnow. Ultra-large alignments using phylogeny-aware profiles. *Genome biology*, 16(1):1–15, 2015.

[NW70] Saul B Needleman and Christian D Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of molecular biology*, 48(3):443–453, 1970.

[RK03] Thomas Kiel Rasmussen and Thiemo Krink. Improved hidden markov model training for multiple sequence alignment by a particle swarm optimization—evolutionary algorithm hybrid. *Biosystems*, 72(1-2):5–17, 2003.

[RSB18] Ramchalam Kinattinkara Ramakrishnan, Jaspal Singh, and Mathieu Blanchette. Rlalign: a reinforcement learning approach for multiple sequence alignment. In *2018 IEEE 18th International Conference on Bioinformatics and Bioengineering (BIBE)*, pages 61–66. IEEE, 2018.

[RSdW$^+$18] Tabish Rashid, Mikayel Samvelyan, Cedric de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, pages 4295–4304. PMLR, 2018.

[SC21] Yong-Joon Song and Dong-Ho Cho. Local alignment of dna sequence based on deep reinforcement learning. *IEEE Open Journal of Engineering in Medicine and Biology*, 2:170–178, 2021.

[SHM+16] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

[SJS+21] Yong-Joon Song, Dong Jin Ji, Hyein Seo, Gyu Bum Han, and Dong-Ho Cho. Pairwise heuristic sequence alignment algorithm based on deep reinforcement learning. *IEEE Open Journal of Engineering in Medicine and Biology*, 2:36–43, 2021.

[SW+81] Temple F Smith, Michael S Waterman, et al. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197, 1981.

[Szi12] István Szita. Reinforcement learning in games. In *Reinforcement learning*, pages 539–577. Springer, 2012.

[VBC+19] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

[WF93] John C Wootton and Scott Federhen. Statistics of local complexity in amino acid sequences and sequence databases. *Computers & chemistry*, 17(2):149–163, 1993.

[WGD+17] Kunfeng Wang, Chao Gou, Yanjie Duan, Yilun Lin, Xinhu Zheng, and Fei-Yue Wang. Generative adversarial networks: introduction and outlook. *IEEE/CAA Journal of Automatica Sinica*, 4(4):588–598, 2017.

[WJ94] Lusheng Wang and Tao Jiang. On the complexity of multiple sequence alignment. *Journal of computational biology*, 1(4):337–348, 1994.

[WPDX$^+$22] Chen Wang, Claudia Pérez-D'Arpino, Danfei Xu, Li Fei-Fei, Karen Liu, and Silvio Savarese. Co-gail: Learning diverse strategies for human-robot collaboration. In *Conference on Robot Learning*, pages 1279–1290. PMLR, 2022.

[WSH$^+$16] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando Freitas. Dueling network architectures for deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1995–2003, 2016.

[WVO12] Marco A Wiering and Martijn Van Otterlo. Reinforcement learning. *Adaptation, learning, and optimization*, 12(3):729, 2012.

[XC09] Fasheng Xu and Yuehui Chen. A method for multiple sequence alignment based on particle swarm optimization. In *Emerging Intelligent Computing Technology and Applications. With Aspects of Artificial Intelligence: 5th International Conference on Intelligent Computing, ICIC 2009 Ulsan, South Korea, September 16-19, 2009 Proceedings 5*, pages 965–973. Springer, 2009.

[ZZL$^+$22] Yongqing Zhang, Qiang Zhang, Yuhang Liu, Meng Lin, and Chunli Ding. Multiple sequence alignment based on deep q network with negative feedback policy. *Computational Biology and Chemistry*, 101:107780, 2022.

Cytometry Data Clustering

5

How EVE Online Players Uncover Patterns in Cytometry Data

Preface

In the first part of my thesis, we explored multiple sequence alignment using data provided by the Borderlands Science. In this chapter, we generalize this approach to other combinatorial problems by extending the dataset with data coming from EVE Online, an online game providing annotated cytometry data marked by players.

Since this task has a different structure, and the data type is very different from what we've been used to, the methods had to be adapted as well. Central to our approach, even with the changes, has been the use of player-generated data and behavioural cloning. In this chapter, I design and experiment with a model that can help extract deeper insights from player-provided cytometry data annotations. That would fall under visual data, and perhaps better dealt with by a technique known as behavioral cloning. Here, we upgrade it using the U-Net architecture, which is perfect for image-based data processing.

I contributed to the design of the research framework, implementation of the project, code development, data processing, model architecture design, fine-tuning, validation, testing, and analysis of results.

I prepared a draft of the manuscript, figured out the artwork, and formatted tables used here. Alexander Butyav played an important role in data preparation - actually, crucial work for such kinds of projects. Jérôme Waldispühl guided the general design of the research, reviewed the analysis, critically reviewed, and edited the draft of this manuscript.

At present, the article is a manuscript awaiting submission, so this work is still going on. We intend to submit this paper either next year.

5.1 Introduction

Within the spaces of online gaming, there are very few rivaling environments for level and multi-layered nature as the virtual universes found in the Massively Multiplayer Online category. Probably the most complex example would be that of EVE Online, a game that has earned a reputation for an expansive and highly detailed virtual world in which players can run economic systems, forge and maintain political alliances, and conduct large-scale battles. EVE Online is one of a kind regarding the player-driven economy and open-ended gameplay, while enabling players to solve complex problems, strategize on an unprecedented scale, and cooperate [Wik21, Gam24, Onl21, Gam21]. Besides conquering and exploring space, players also need to work in a team toward attaining complex goals; these are decisions that may lead to widespread ramifications inside the game. This environment has also been used to solve real-world scientific problems by crowdsourcing data analysis tasks where collective human decisions can strongly benefit.

In the last couple of years alone, the symbiosis between gaming and academic research has given way to new ways of crowdsourcing in problem-solving. Already, various games such as Foldit, EyeWire, and Borderlands Science have successfully shown how players can enable scientific discovery in domains ranging from molecular biology to neuroscience. For example, Foldit challenged players to fold proteins in ways that led to novel insights in biochemistry [CKT+10, KCT+11], and EyeWire challenged players to map neural connections in the brain [TLRSH17]. Borderlands Science called on gamers to help align genetic data and helped move the study of microbiomes forward [WSK+20, SGGGB+24]. These projects really show how even the most abstract scientific challenges can be turned into interactive puzzles, which players, in some cases with great success, solve in game environments.

Similarly, EVE Online has entered the field of scientific contribution through various projects, including Project Discovery [Gam03], which lets players take part in the analysis of real-world datasets, such as flow cytometry data. Flow cytometry is a key technique in biological and clinical research that allows for the measurement of various physical and chemical properties of cells, including size, shape, and protein expression [McK18, GHP+14].

Despite the value of this data, many large datasets are difficult to interpret using flow cytometry; this is because the process of identifying clusters or patterns in these data requires a high degree of expertise [McK18]. One of the major problems in clustering flow cytometry data is the high dimensionality, hence overfitting and computation-intensive processes. This is driving a related interest in the development of more intuitive and scalable analysis tools to handle increasingly complex and higher-volume cytometric data from modern experimental setups [AFH$^+$13, PHW$^+$09].

By engaging EVE Online players in this process, Project Discovery attempts to turn such challenges by leveraging the intuitive pattern recognition capability of the human brain in a novel way for meaningful engagement of non-experts in data analysis [Mat14]. It will also contribute to crowdsourcing better, scalable methods to handle complex biological data and provide insight into how human decision-making can improve automated clustering techniques.

The manner with which non-experts interact and analyze such complex scientific data forms a vital clue for how human intuition can be exploited to solve scientific challenges. Non-experts are in a better position to come up with novel approaches for pattern recognition that complement the traditional algorithms and might offer insights into the better design of more user-friendly and scalable data analysis tools. We base our study on four major hypotheses:

5.1.1 Hypotheses

Do Players Reach Consensus in Cytometry Data Clustering? The hypothesis is an investigation as to whether or not, when players clustering cytometry data, these intuitive procedures can reach agreement, find patterns, and even find signals in highly complex high-dimensional datasets. In support of this we have evaluated the various clustering choices made by the gamers to establish for sure how consistently non-expert intuition catches meaningful structures with regard to human intuition, unravelling possibly congruent or even better signals according to traditional algorithmic approach.

Which Clustering Models Resemble Human Decision-Making? The second hypothesis tests the proposition of which clustering models result closest to human decision-making: Are any of these classical clustering algorithms producing results similar to the intuitive patterns players identified in the cytometry data? Elucidation of this relationship may provide insight into the similarities between algorithmic and human approaches to data clustering, highlighting the models that reflect the strategies utilized by non-experts in reaching decisions.

How Well Do Imitation Algorithms Replicate Player Strategies? The third hypothesis probes how well imitation learning algorithms can mimic strategies that players utilize to cluster cytometry data. In other words, it is about whether the machine learning models learn from the way non-expert players decide and hence reflect their particular ways of looking at data.

What Insights Can We Extract from Human Decision-Making in Clustering? The fourth hypothesis is an investigation into what can be learnt from human decision-making in clustering, by means of indirect understanding of strategies, patterns, and intuitive approaches through modeling. Analyses of how human solutions relate to computational models will show deeper cognitive processes and strategies that players use when approaching the clustering of complex cytometry data.

5.1.2 Game Design

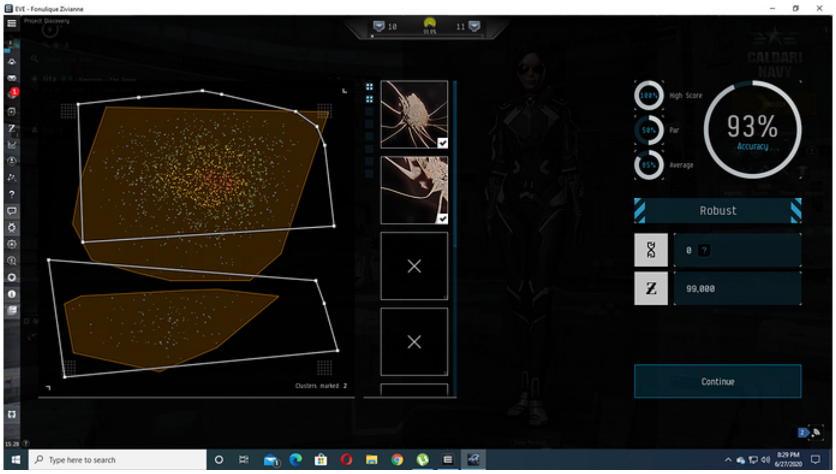

Figure 5.1: Clustering in EVE Online: Left side shows a screenshot from EVE Online. Overlaid on the left is a clustering visualization of player activities and resources. The right side presents metrics such as scores and resource collection rates through graphs and charts, illustrating player dynamics.

In EVE Online, Project Discovery makes for a unique crossroads between gaming and science, with the creation of excellently designed mini-games that drive the players to accomplish real-world scientific tasks. Most recently, it has focused on the analysis of flow cytometry data. The player is presented with a graphical representation of cell populations to analyze in a mini-game fashion. What will result on each player's screen is that a flow cytometry chart results with scattered plots representing another cell type. This does involve the drawing of a polygon around clusters of the cells so that the drawn polygon will act as scientific means of categorizing it in the process of "gating" in old-fashioned flow cytometry.

The game is set within the universe of EVE Online. Thus, players' contributions go a long way in helping actual scientific research processes- yet without having to leave the comfort zone of the game itself. Its intuitive interface facilitates a scientifically sound environment wherein its view and plot details are easily adjustable, so the data analysis accuracy turns out to be very high. It's a tutorial introducing a new player to the mini-game mechanics in such a way that even a person not educated in much science is able to contribute.

Project Discovery Performance: Success in Project Discovery is quantitatively measured through a scoring system which leverages an evaluation against a 'gold standard' set by scientific experts in order to judge the correctness of player cell categorizations. Immediately after each and every performance, feedback is provided because scores impact both experience points and in-game currency rewarded, thus working as incentive systems to do more analyses with more precision.

Indeed, devised by scientists and developers under consultation, Project Discovery typified how to successfully gamify scientific research. They enlist the millions in EVE Online to crowdsolve data analysis that enables such research to scale up exponentially in speed. Interesting idea, it then was praised as ingenious, extremely viable, serving data of true value to the scientific world. However, it was engaging for them and educational for the players in this joint effort.

5.2 Related work

The intersection of science and gaming, especially in games like EVE Online's Project Discovery, can probably be considered the most fantastic crossing of disciplines in scientific theory. It is a meeting place for many: computational biology, crowd psychology, data science-all flow into new platforms whose ambition it is to capitalize on public engagement for the furtherance of frontiers in science.

Scientific task gamification is borrowed from high-publicity projects that first deployed the public use of computing resources for scientific research. This is inclusive of the seminal project Folding@home, which, starting in 2000, applied distributed computing power to simulate protein folding through the idle processing of tens of thousands of volunteer-Contributed computers across the world. This project set a precedent for how non-scientists might contribute to complex scientific problems by providing a format that is both accessible and useful to the scientific community [CKT$^+$08]. We build upon this work from Folding@home and introduce Phylo to advance the multiple sequence alignment problem from both outside in, by its outward appeal, and inside out, by its inward contents. Launched in 2011, Phylo asked players to arrange sequences of DNA in an effort to contribute to the scientific study of genetics and evolutionary biology [KRK$^+$11].

Borderlands Science took the success of, and framework from, these previously mentioned projects and introduced them into the video game series Borderlands. This newer program not only challenged players with solving puzzles that would lead to mapping out the human microbiome but also further evolved the gamification of science by showing how complex scientific tasks could be effectively embedded in commercially available video games to reach a wider audience [WSK$^+$20].

This connection between Phylo and Borderlands Science demonstrates how ideas from initial projects ignite a spark, evolve through development, and become much more sophisticated in subsequent initiatives. Better still, these projects undertake the coming together of scientific inquiry with gameplay, within which each entrant actually garners some new lesson and successes its ancestors might have learned. While this approach accelerates scientific data processing, it also extends public engagement in education, bridging the two strata: scientists and citizens.

Flow cytometry is the name of advanced techniques found in both immunology and oncology, mainly focused on the analysis of physical and chemical characteristics of fluid particles by passing them through a laser. It gets even more complicated when it comes to the need for identification and quantification appropriately of various cell populations, mostly similar in nature, often with small differences in fluorescence and light scatter. Here we describe in detail some of the computational techniques most commonly in use for the effective management of large data generated by flow cytometry and for its analysis.

With the successful integration of games and scientific research, this model applied to other complex tasks is a very promising frontier; for example, the analysis of flow cytometry. It is a technique that analyses the physical and chemical characteristics of particles in a fluid as they pass through at least one laser. Flow cytometry requires advanced analysis because the data are high-dimensional. It has a number of uses in various machine learning model studies for enhancing the existing flow cytometry data analyses. On one side, unsupervised machine learning methods, such as clustering algorithms, identify cell populations in a non-prelabeled way. On the other hand, supervised methods-like decision trees or neural networks-classify into specific cell types based on the known labels [QM90, AFH^{+}13].

Considering the technical difficulty and expert knowledge necessary to analyze flow cytometry data, this opens the potential for numerous implications that the introduction of this task into a gamified environment may have in widening participation and improving the technique of analysis. Gamification taps into the efforts of large audiences in processing data, and in effect, creates novel insights in ways that no other public strategy or intuition has realized. While currently used in a range of scientific contexts for projects involving protein folding with Foldit, it could do the same in democratizing flow cytometry analysis to the masses, increasing public understanding and participation in the process of science to address the most complex biomedical challenges. [CKT^{+}10].

Although the methodologies of machine learning represent state-of-the-art, their application within the analysis of complex data, such as flow cytometric ones, has been rather inhomogeneous concerning human analytics. An interesting critical study by Kim et al. showed how tasks in which careful image segmentation had to be performed saw much higher precisions and recall by experienced players than a state-of-the-art deep convolutional neural network algorithm designed for this task. In this respect, human input is still invaluable, even when pitted against the most advanced artificial intelligence; sometimes, intuition and human experience can prove even more vital than automated systems [KGZ^{+}14].

The paper ends with a discussion on the general capabilities of the human brain in cluster discrimination of flow cytometry as related to its natural tendencies to handle visual information appropriately. A variety of perceptual cues emanating from the visual cortex, in particular the inferotemporal cortex and the lateral fusiform gyrus, has been employed by humans in distinguishing objects. The ability to make sense of and classify objects extends further to shapes, colors, textures, and spatial relations. For example, in the recognition of various fruits there is the recognition of typical colors and shapes where apples are round and red, while bananas are yellow and elongated [MIT19].

Color sensitivity in humans is a significant factor in object recognition. It contributes to the decision about specific attributes of objects, like maturity or potential danger, and thus allows one to group similar objects [Wil02]. Besides, humans make use of contextual information, such as the standard environment of an object or its average size, to still further increase the accuracy of recognition. The integration of the visual cue with the cognitive processes enhances recognition accuracy by orders of magnitude [ar520]. These capacities then create an argument for the potential gain that could be had from gleaning human perceptual capabilities in analyzing very complex data sets, which flow cytometry generally churns out. In effect, it gamifies such natural human skills and carves a new direction toward enhancing the efficiency and precision of tasks on data categorization, which typically has been done with so much less versatility in computation models. This approach not only plays into the intrinsic visual and cognitive strengths of human beings but can also be productive in fusing together human ingenuity with machine efficiency to lead perhaps to superior analytical outcomes in scientific research.

5.3 Data

The third phase of Project Discovery in EVE Online used a dataset of protein structures of SARS-CoV-2, where players participated to speed up scientific analysis against the backdrop of the ongoing pandemic scenario caused by COVID-19. Participants in this round received a dataset of high-resolution images on these clusters of the protein structure of the virus, which is important for understanding the mechanics of the virus at a molecular level.

The EVE online players will be involved in classifying and labeling of huge datasets of such protein structures. Thus, scientific data validation is achieved by a very large community in its aim. The players would be tasked with a various kind of different protein clusters. The aim shall always be to single out that feature of interest, in which biologists want the interaction between viruses and the host. This is so important because it locates potential targets for therapy or even drug development.

Technically, the dataset was reworked for a gamified environment in such a way that accessibility and engagement would be maintained but not at the

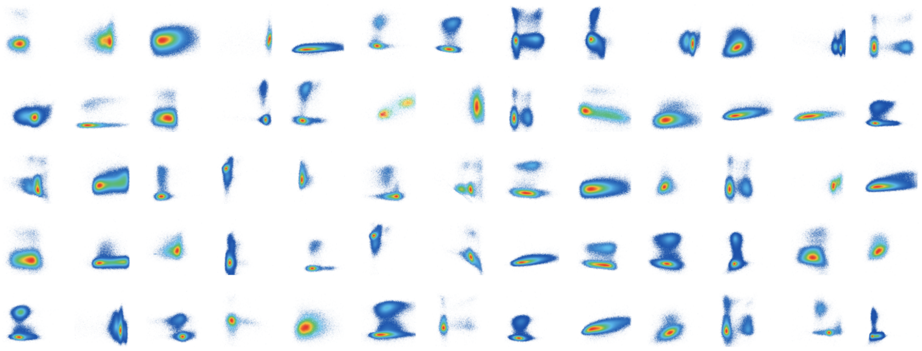

Figure 5.2: Example of flow cytometry data plots, illustrating cell population densities. Higher-density regions are represented in red, indicating a greater concentration of cells, while lighter areas are shown in blue, highlighting regions with lower cell densities

cost of scientific integrity 5.2. To each protein image, a set of predefined possible features had been added that players had to choose according to shape, distribution, and apparent interactions depicted in the photos.

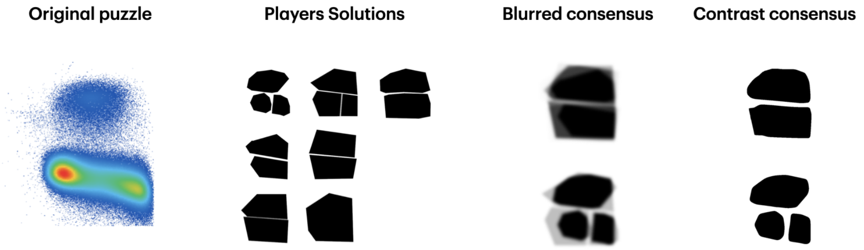

Figure 5.3: Example of flow cytometry data plots illustrating different stages of problem-solving: From left to right: (1) Initial puzzle, representing the raw, unprocessed data; (2) All players' solutions, showcasing individual attempts to solve the problem; (3) Overlapping solutions, highlighting areas of agreement among players; (4) Cleaned consensus solution, depicting the refined outcome derived from combining and refining the overlapping solutions.

Data visualization in black and white pixel segments happens in post-processing, long after the gamers have played the game. Solutions submitted are filtered out for noise, leaving the dataset with the real relevant number of solutions. For example, filtering out solutions which pick out only one cluster, that normally doesn't carry much scientific insight.

Next, there is a need to synthesize all other solutions that may point towards this number of clusters. The very nature of the task being flexible makes aggregation of these multiple solutions to a single analytical output a critical task. This is achieved by laying one solution over the other and searching for commonalities and disparities in how the different players decide about the clustering output.

This makes the clustered output less jagged and hence more amenable to scientific evaluation. We apply a Gaussian blur with boundaries smoothed, using sigma of 0.5; this would blur edges of clusters into others, decreasing the sharpness of the boundary lines, hence allowing visual clarity of the cluster results.

Enhancement is done after the application of Gaussian blur, where the remaining gray areas are removed and sharpened to give a better distinction between clusters. It is in this processing of enhancement with high contrast and well-defined delineation that the real scientific analysis and interpretation become immensely important 5.3.

5.4 Methods

The dataset consisted of RGB images and their corresponding binary segmentation masks, each resized to 256×256 pixels. These images represent microscopy views or density plots, and each mask indicates the location of distinct black clusters. The input image set is denoted as $X \in R^{N\times256\times256\times3}$ and the corresponding masks as $Y \in R^{N\times256\times256\times1}$, where N is the total number of samples.

As part of preprocessing, several features were extracted to better capture structural and density-related characteristics of the images. These included: (1) pixel intensity values, used directly to represent density; (2) color gradient information, reflecting transitions in density levels typically color-coded from blue (low) to red (high); (3) histogram of oriented gradients, capturing local edge and shape patterns; (4) contour-based edge features, highlighting boundaries using gradient detection methods; (5) geometric shape descriptors such as area, elongation, compactness, and convexity to characterize regions; and (6) statistical features including mean, variance, skewness, and kurtosis of pixel intensity distributions.

To obtain ground truth cluster counts from the binary masks, each image was converted to grayscale, inverted, and thresholded to create a binary mask suitable for contour detection. The number of outer contours was taken as the estimated number of clusters. These counts were normalized by the maximum observed count across the dataset and used as targets for regression.

A convolutional neural network was trained to predict the number of clusters the input images. The network consisted of multiple convolutional and pooling layers followed by fully connected layers, producing a single continuous output. The model was trained using mean squared error loss over 50 epochs with a batch size of 8. After training, predicted counts were rescaled to the original range and spatially expanded to match the image dimensions. This additional feature was then concatenated to the original 3-channel input with extracted features.

For the segmentation task, a U-Net architecture was used, modified to accept the 4-channel input. The model included an encoder-decoder structure with skip connections, and applied convolutional layers with non-linear activations, followed by a sigmoid activation in the final layer to produce a binary mask. Training was performed for 50 epochs using binary cross-entropy loss, and performance was evaluated based on pixel-wise accuracy.

Predicted masks were refined through post-processing. Masks were first binarized and cleaned using morphological operations to eliminate small artifacts and noise. Connected components were then labeled to identify distinct clusters, and the total count of labeled regions was recorded. For visualization, convex hulls were computed around each cluster to highlight the boundaries and provide a clearer structural interpretation.

To assess the model qualitatively, five test samples were randomly selected. For each, the original input image, ground truth mask, predicted mask, and convex hull overlays were displayed. The actual and predicted cluster counts were also compared to evaluate both segmentation quality and counting accuracy.

5.5 Results and Discussion

5.5.1 Testing Hypothesis 1: Do Players Reach Consensus in Cytometry Data Clustering?

We then collected all clustering solutions provided by players and analyzed them, noting the number of clusters each player identified and the degree of overlap between them. We then created a consensus by overlaying

Model Type	IoU	Dice Coefficient	Accuracy
Agglomerative	0.1236	0.0039	0.4157
Birch	0.0721	0.0473	0.2316
GaussianMixture	0.0687	0.0460	0.2232
KMeans	0.0692	0.0463	0.2235

Table 5.1: Mean Metrics for Clustering Models. (1) Intersection over Union (IoU): measures overlap between two masks as the ratio of the intersection area to the union area. (2) Dice Coefficient (F1 Score): evaluates similarity between masks by calculating 2 times the intersection area divided by the sum of the areas of both masks. (3) Pixel-wise Accuracy: quantifies correctly classified pixels in both masks, providing a measure of model alignment with human solutions.

similar solutions, highlighting areas of agreement among players.

We further purified this consensus map by applying contrast enhancement, which isolated the most representative clustering solution that balanced the number of clusters and their overlapping regions. We then compared this enhanced-contrast solution with all player solutions exhibiting similar overlap and cluster counts **Figure: 5.3** Preprocessing and analysis of player clustering solutions.

Our results show that players independently identify and determine the number of clusters in the data. When players perceive a certain number of clusters, they consistently agree on their positions. This suggests that a strong signal is being detected from the cytometry data, as player clustering solutions are converging both in cluster count and spatial arrangement.

5.5.2 Testing Hypothesis 2: Which Clustering Models Resemble Human Decision-Making?

In the next section, we will discuss how close to human decision-making various clustering models are by comparing human-generated solutions with the results of four of the most popular clustering algorithms: KMeans [Llo82], GaussianMixture [DLR77], Birch [ZRL96] and Agglomerative [LW67], **Figure: 5.4**.

Three major metrics are used for the evaluation of the likeness between human and model solutions. The reasons for choosing this are that the

Figure 5.4: Mean metrics for clustering models applied to flow cytometry data. The table compares the performance of Agglomerative, Birch, Gaussian Mixture, and KMeans models based on Intersection over Union (IoU), Dice Coefficient, and Accuracy

data is constituted of diverse features, and we wanted to consider the topic from different perspectives to ensure comprehensiveness from the following viewpoints:

- **Intersection over Union (IoU)**: This metric measures the overlap between two masks as the ratio of the intersection area (shared space between human and model clusters) to the union area (total space covered by both). It provides a direct indication of how well the model's clusters align with human-identified clusters.

$$IoU = \frac{\text{Intersection Area}}{\text{Union Area}}$$

- **Dice Coefficient (F1 Score)**: The Dice coefficient evaluates the similarity between two masks. It is calculated as twice the intersection area divided by the sum of the areas of both masks. This score balances precision and recall, focusing on the degree of overlap between clusters.

$$Dice = \frac{2 \times \text{Intersection Area}}{\text{Area of Mask 1} + \text{Area of Mask 2}}$$

- **Pixel-wise Accuracy**: This metric measures how many pixels are classified correctly in both masks (human and model), giving a direct measure of how accurately the model clusters the data compared to human solutions.

$$Accuracy = \frac{\text{Correctly classified pixels}}{\text{Total pixels}}$$

These results make it obvious that no model perfectly captures human decision-making; however, Agglomerative clustering is the best of the compared models across all metrics: IoU of 0.1236, Dice coefficient of 0.0039, and pixel-wise accuracy of 0.4157 **Table: 5.1**. From this, we can see that a higher IoU and accuracy mean that Agglomerative clustering fits human-generated clusters more than other models. However, the far lower Dice coefficient shows great variability in the shape and size of clusters between human solution and model solutions.

While Birch, GaussianMixture, and KMeans results were closer to each other, lower on IoU and accuracy, they tended to align less with the human clustering solutions. That is, these models totally failed to recognize the much finer patterns that human judgment produced, with poor overlap into the image and poorly defined clusters.

The discrepancy between human clustering and algorithmic models suggests that human players may be employing more intuitive, nonlinear strategies that these models fail to capture and hints at the possible value of further refinement of clustering algorithms to better align with human intuition in complex datasets such as cytometry.

5.5.3 Testing Hypothesis 3: How Well Do Imitation Algorithms Replicate Player Strategies?

We adopted an imitation learning model, as described in the "Models" section, and trained it using player-generated data. For evaluating how well this model mimics human decision-making, we used the same three metrics previously used for comparing the classical clustering methods against human solutions: IoU, Dice Coefficient, and Pixel-wise Accuracy.

The model indeed demonstrated a significant IoU alignment with human-generated results, reaching up to 0.5636, **Figure: 5.1**. These values indicate that the model not only captures the salient features of human decisions but also mimics human performance quite well both in terms of cluster overlap and accuracy.

This result underlines that imitation learning can be very proficient in the implementation of intuitive strategies devised by players and points towards very strong potential for the future applications of the model in clustering tasks with human-like analyses of complex data.

5.5.4 Testing Hypothesis 4: What Insights Can We Extract from Human Decision-Making in Clustering?

To test this hypothesis, we first generated an artificial dataset varying the attributes between clusters (**Figure 5.5**).

First experiment we are varying the distances between two clusters. The results indicated that the human-inspired predictive model reliably identified separate clusters when inter-cluster distances were sufficiently large (e.g., distances of [6, 6] or [8, 8]), clearly delineating distinct groups. However, as inter-cluster distances decreased (e.g., distances of [1, 1] or [2, 2]), clusters significantly overlapped, leading the model to predict a single cluster. Intermediate distances ([4, 4]) represented a critical threshold, where the model consistently differentiated two separate clusters. The experiment demonstrated that human-inspired clustering models are sensitive to the spatial distance between clusters. Specifically, the model reliably distinguishes clusters that are well-separated spatially, yet merges clusters when distances decrease and overlap occurs. Thus, spatial separation significantly influences human decision-making in clustering tasks, emphasizing the need to incorporate distance-based criteria in clustering strategies.

In the second artificial dataset, we investigated how oval-shaped covariance parameters, influencing cluster elongation, impacted predictions (**Figure 5.6**). Clusters with minimal elongation ([0.2, 0.2]) were reliably identified individually, resulting in accurate cluster separation. However, increased elongation ([4.0, 2.0] or [4.0, 4.0]) produced significant overlap, prompting predictions of single clusters due to unclear boundaries. Asymmetric covariance ([1.0, 0.5] or [2.0, 0.5]) introduced uncertainty, with predictions varying between one and two clusters. This experiment showed that the shape and elongation of clusters strongly affect how clusters are predicted. The model reliably separated clusters when they were compact and symmetrical, but elongated or unevenly shaped clusters often led to overlapping and unclear boundaries, making predictions uncertain. These results suggest human clustering decisions depend significantly on cluster shapes, highlighting the importance of considering shape when developing clustering methods.

Next, we explored scenarios involving three overlapping clusters of varying densities and proximities (**Figure 5.7**). The model demonstrated robust performance when clusters had distinct densities and adequate separation, accurately predicting all three clusters. Conversely, substantial overlaps or density variations caused the model to merge sparse clusters into denser ones, predicting fewer clusters. In this experiment, we saw that the model clearly identified three clusters when they were separated enough and had similar densities. However, when clusters overlapped significantly or had very different densities, the model combined sparse clusters with denser ones, predicting fewer clusters. These results show that humans naturally focus on denser, more compact groups, suggesting future clustering methods should

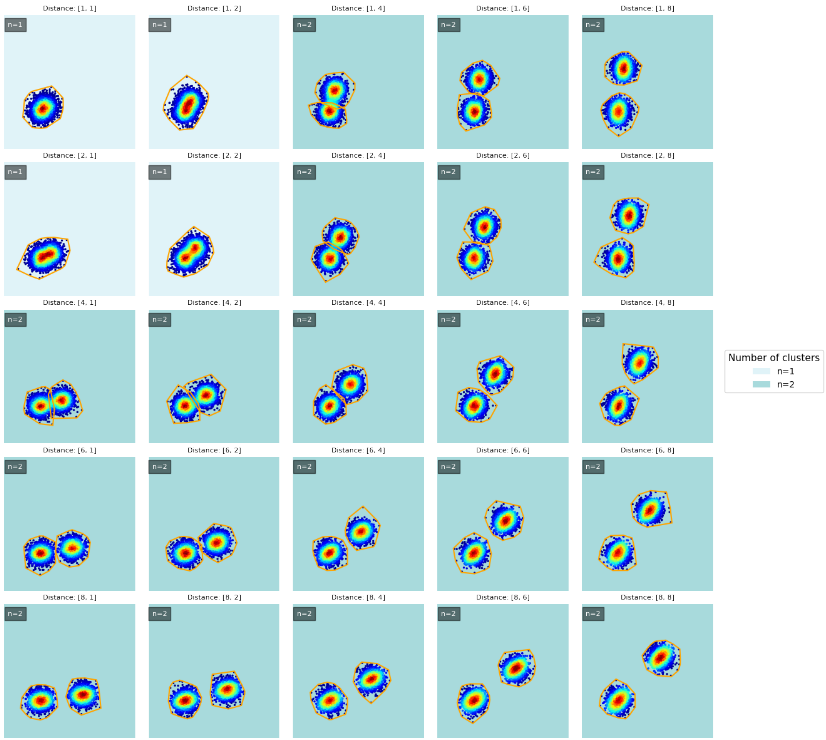

Figure 5.5: Model predictions for cluster number across varying inter-cluster distances. Larger distances resulted in clear separation and accurate cluster prediction, whereas smaller distances produced significant overlap and ambiguous predictions.

account for differences in cluster density.

In an additional dataset, the impact of asymmetric oval covariance parameters on three-cluster scenarios was analyzed (Figure 5.8). While clusters with low covariance ([0.2, 0.2]) were accurately identified, increasing covariance asymmetry led to elongated and irregular shapes, causing ambiguity and resulting in fewer predicted clusters. In this experiment, we found that clusters with low asymmetry were easily recognized, but as asymmetry increased, clusters became elongated and irregular, causing the model to combine them into fewer clusters. This highlights that people are sensitive to unusual or stretched cluster shapes, indicating that clustering methods should better handle irregular shapes to match human perceptions.

Finally, we examined how varying cluster rotation angles influenced the

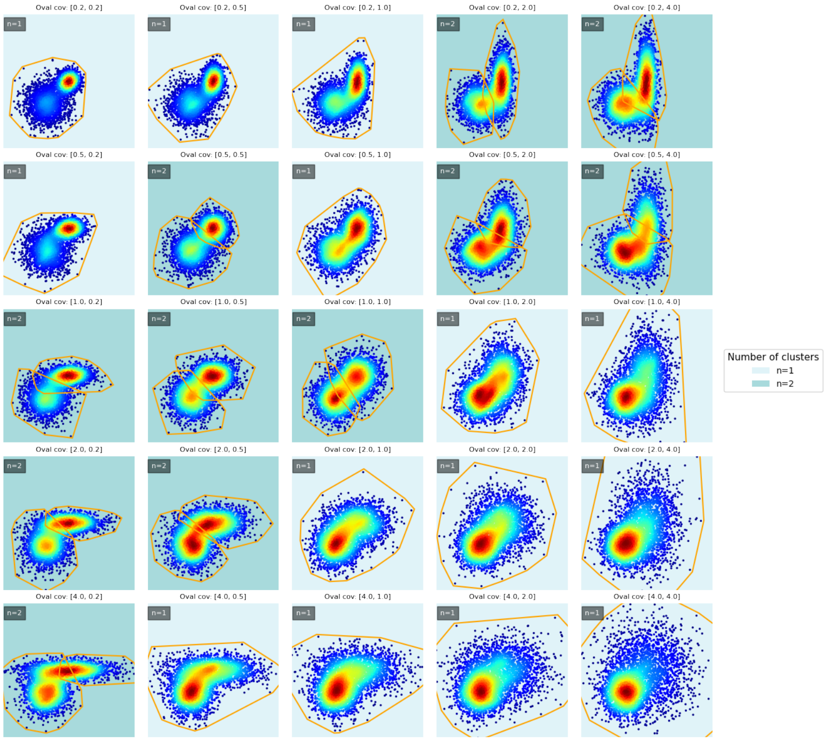

Figure 5.6: Effect of oval covariance parameters on cluster predictions. Higher elongation led to ambiguous boundaries and fewer predicted clusters, indicating shape significantly influences human decision-making in clustering tasks.

model's ability to identify distinct clusters (**Figure 5.9**). Results indicated that small rotation angles (20°–40°) often resulted in overlapping clusters, with fewer clusters predicted. As angles increased (90°–120°), clusters became spatially distinct along rotated axes, leading to improved detection of multiple clusters. In this experiment, we saw that small rotation angles often made clusters overlap, causing the model to identify fewer clusters. Larger rotation angles helped separate clusters clearly, improving detection. This shows that people's clustering decisions depend not only on distance and shape but also strongly on how clusters are oriented. Future clustering methods could benefit from considering cluster orientation more carefully.

Our experiments showed that human-inspired clustering decisions strongly depend on cluster features such as distance, shape, density, and orientation.

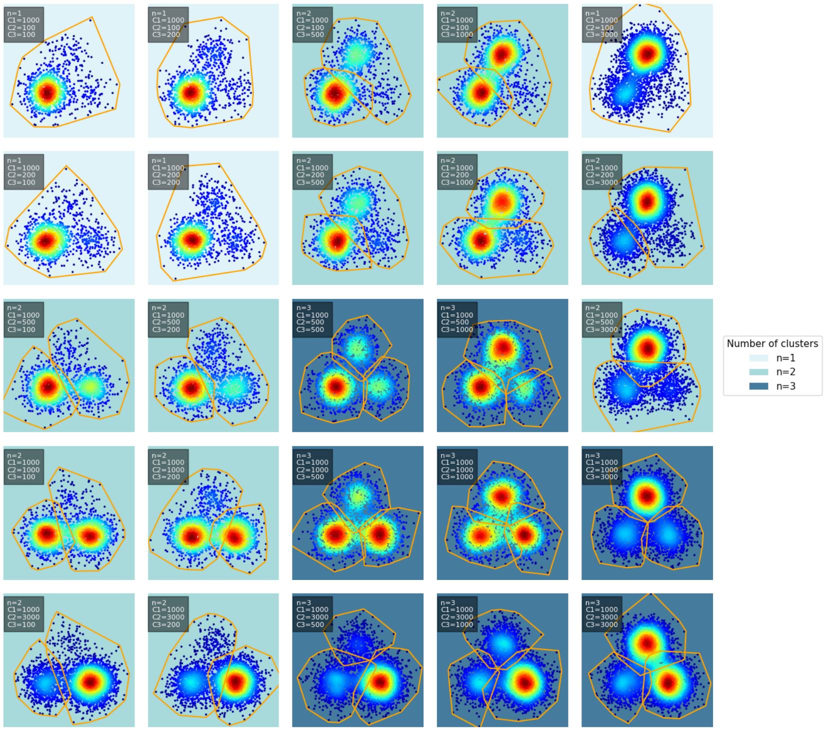

Figure 5.7: Model predictions across scenarios involving three overlapping clusters with varying densities and spatial arrangements. Predictions revealed human biases toward recognizing denser and compact regions as distinct clusters, merging sparse or heavily overlapping areas into fewer clusters.

People clearly identify clusters when they're well separated, compact, and similarly dense. However, when clusters overlap, become elongated, irregularly shaped, or differ significantly in density, people tend to merge them into fewer groups. We also saw that cluster orientation matters, with better clustering accuracy when clusters are rotated apart at larger angles. Overall, these results suggest that clustering methods inspired by human decision-making should carefully consider distance, shape, density, and orientation to effectively match how people naturally group data.

5.6 Limitations

While our study provides important insights into how players can contribute to the clustering of cytometry data and how their strategies compare

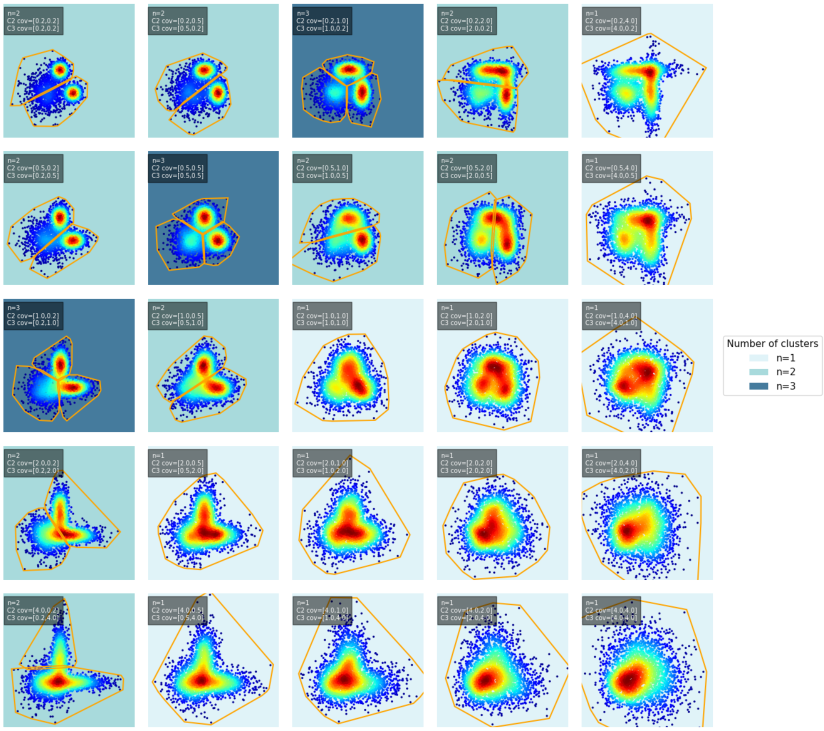

Figure 5.8: Impact of asymmetric covariance parameters on cluster predictions in three-cluster scenarios. Increasing asymmetry and elongation led to ambiguous boundaries and fewer clusters predicted, illustrating human perceptual sensitivity to irregular cluster shapes.

with machine learning models, there are several limitations that have to be considered. First, the consensus generated from player solutions may not fully capture the most accurate clustering of the data, since players' decisions were based on intuition rather than domain-specific expertise in cytometry. This introduces a bias in how the tendencies of clusters become perceived and grouped. Secondly, the general, generic models tested here to describe clustering performance, like KMeans, GaussianMixture, Birch, Agglomerative, have intrinsic limits in general and more importantly when high-dimensionality of cytometry data deals the cards, hence their results were limited to be completely successful in emulating a human decision. The imitation learning model has also achieved success in reproducing strategies of players; however, it may be difficult to generalize to more other datasets

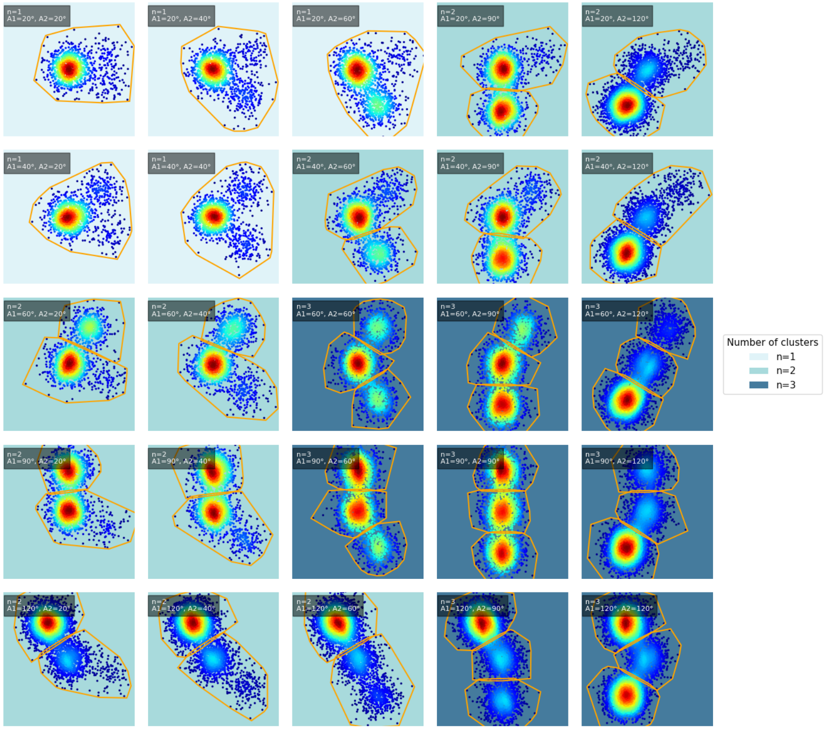

Figure 5.9: Influence of cluster rotation angles on model predictions. Smaller angles resulted in overlapping clusters and fewer predicted clusters, whereas larger angles improved spatial distinction and clustering accuracy, emphasizing the role of cluster orientation in human decision-making.

or tasks than those of the player data on which it was originally trained in this study. Lastly, the metrics for evaluation, such as IoU, Dice Coefficient, and pixel-wise accuracy, are biased toward spatial, whereas this may not encompass other more qualitative aspects of human intuition-for example, why humans chose their specific clustering. Future work shall move into more advanced modeling and consider a wider array of evaluative criteria toward a deepened understanding of how intuition can be best integrated within data clustering approaches.

5.7 Conclusion

In this work, we have explored how players of EVE Online can contribute to clustering complex cytometry data and the ways in which their decision-making strategies compared to both classical and imitation learning models. Our results showed that players were able to discern coherent patterns in the data, reaching a very rapid consensus in their solutions regarding clustering, which suggests the capability of human intuition to successfully detect high-dimensional data signals. Among all classical clustering models, Agglomerative clustering is closest to human strategies. However, imitation learning algorithms worked out much better in the case of replicating human decisions and proved that machine learning might not only master but also develop further in those cases where human solutions seem irretrievable. By finally looking at human decision-making, we got a good insight into what factors can potentially influence cluster distance, orientation, and shape, and this can help us improve clustering methods for future applications.

Bibliography

[AFH+13] Nima Aghaeepour, Greg Finak, Holger Hoos, Tim R Mosmann, Ryan Brinkman, Raphael Gottardo, and Richard H Scheuermann. Critical assessment of automated flow cytometry data analysis techniques. *Nature Methods*, 10(3):228–238, 2013.

[ar520] ar5iv Labs. Computational models mimicking human visual recognition processes, 2020.

[CKT+08] S. Cooper, F. Khatib, A. Treuille, J. Barbero, J. Lee, M. Beenen, A. Leaver-Fay, D. Baker, Z. Popović, and Foldit Players. Foldit, 2008. An online puzzle video game about protein folding. It is part of an experimental research project developed by the University of Washington, Center for Game Science, in collaboration with the UW Department of Biochemistry.

[CKT+10] Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker, Zoran Popović, et al. Predicting protein structures with a multiplayer online game. *Nature*, 466(7307):756–760, 2010.

[DLR77] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm.

Journal of the Royal Statistical Society: Series B (Methodological), 39(1):1–22, 1977.

[Gam03] CCP Games. Eve online. *World Wide Web. http://www. eve-online. com*, 2003.

[Gam21] CCP Games. Eve online's "project discovery" wins webby people's voice award for public service, activism, and social impact, 2021. Accessed: 2023-06-23.

[Gam24] CCP Games. Phase four of eve online's project discovery revealed, expanding to cancer research and a companion mobile app launching in 2024, 2024. Accessed: 2023-06-23.

[GHP$^+$14] Craig A Gedye, Ali Hussain, Joshua Paterson, Alannah Smrke, Harleen Saini, Danylo Sirskyj, Keira Pereira, Nazleen Lobo, Jocelyn Stewart, Christopher Go, et al. Cell surface profiling using high-throughput flow cytometry: a platform for biomarker discovery and analysis of cellular heterogeneity. *PloS one*, 9(8):e105602, 2014.

[KCT$^+$11] Firas Khatib, Seth Cooper, Michael D Tyka, Kefan Xu, Ilya Makedon, Zoran Popović, David Baker, and Foldit Players. Algorithm discovery by protein folding game players. *Proceedings of the National Academy of Sciences*, 108(47):18949–18953, 2011.

[KGZ$^+$14] Jinseop S Kim, Matthew J Greene, Aleksandar Zlateski, Kisuk Lee, Mark Richardson, Srinivas C Turaga, Michael Purcaro, Matthew Balkam, Amy Robinson, Bardia F Behabadi, et al. Space–time wiring specificity supports direction selectivity in the retina. *Nature*, 509(7500):331–336, 2014.

[KRK$^+$11] A. Kawrykow, G. Roumanis, A. Kam, D. Kwak, C. Leung, C. Wu, E. Zarour, L. Sarmenta, M. Blanchette, and J. Waldispühl. Phylo: A citizen science approach for improving multiple sequence alignment, 2011. A citizen science platform that converts multiple sequence alignment problems into a game, engaging the public in solving complex genetic puzzles.

[Llo82] Stuart Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.

[LW67] Godfrey N Lance and William Thomas Williams. A general theory of classificatory sorting strategies: 1. hierarchical systems. *The computer journal*, 9(4):373–380, 1967.

[Mat14] Mark P Mattson. Superior pattern processing is the essence of the evolved human brain. *Frontiers in neuroscience*, 8:265, 2014.

[McK18] Katherine M McKinnon. Flow cytometry: an overview. *Current protocols in immunology*, 120(1):5–1, 2018.

[MIT19] MIT News. The brain's inferotemporal cortex plays a crucial role in object recognition, 2019.

[Onl21] EVE Online. Project discovery - level up!, 2021. Accessed: 2023-06-23.

[PHW⁺09] Saumyadipta Pyne, Xiaoxu Hu, Kui Wang, Erica Rossin, Ting-I Lin, Lisa M Maier, Clare Baecher-Allan, Geoffrey J McLachlan, Pablo Tamayo, David A Hafler, et al. Automated high-dimensional flow cytometric data analysis. *Proceedings of the National Academy of Sciences*, 106(21):8519–8524, 2009.

[QM90] Yingdong Qian and Robert F Murphy. Automated identification of subpopulations in flow cytometric list mode data using cluster analysis. *Cytometry*, 11(5):583–589, 1990.

[SGGGB⁺24] Roman Sarrazin-Gendron, Parham Ghasemloo Gheidari, Alexander Butyaev, Timothy Keding, Eddie Cai, Jiayue Zheng, Renata Mutalova, Julien Mounthanyvong, Yuxue Zhu, Elena Nazarova, et al. Improving microbial phylogeny with citizen science within a mass-market video game. *Nature Biotechnology*, pages 1–9, 2024.

[TLRSH17] Ramine Tinati, Markus Luczak-Roesch, Elena Simperl, and Wendy Hall. An investigation of player motivations in eyewire, a gamified citizen science project. *Computers in Human Behavior*, 73:527–540, 2017.

[Wik21] EVE University Wiki. Project discovery: Flow cytometry, 2021. Accessed: 2023-06-23.

[Wil02] Wiley Online Library. The significance of color in object recognition, 2002.

[WSK$^+$20] Jérôme Waldispühl, Attila Szantner, Rob Knight, Sébastien Caisse, and Randy Pitchford. Leveling up citizen science. *Nature Biotechnology*, 38(10):1124–1126, 2020.

[ZRL96] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. Birch: An efficient data clustering method for very large databases. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pages 103–114, 1996.

Discussion and Conclusions

6

Discussion

In this work, two basic combinatorial problems arising in bioinformatics where computational approaches are not yet effective are identified. Most often, these challenges refer either to the time needed to yield a solution or to the completeness of the solution itself. In the present thesis, the alternative of Citizen Science-a method that marshals crowds as a means to apply the collective intelligence to solve such open problems-is suggested. This framework opened new opportunities in those areas where computational methods still faced severe limitations.

The core of this thesis is the question of how the solutions contributed in citizen science can be integrated into automated systems by imitation learning methods such that the findings and strategies devised by human contributors are exploited for improving the performance of computational models. That is, enabling such models to solve not only the original problems but also similar or even more challenging problems.

As discussed in the **Section Background**, multiple sequence alignment is one of the most fundamental tasks in bioinformatics, and cytometry also presents its own unique challenges. Citizen projects such as `Borderlands Science` or `EVE Online's Project Discovery` highlighted the potential to engage thousands of players in complex problems but also created valuable data based on crowdsourcing for active solution findings in complex scientific problem situations. This work extends that approach by applying player-generated data to an even more challenging task-namely, the interrelated challenges of MSA and cytometry. Specifically, two behavioral cloning models are used here in order to mimic and extend human strategies for those tasks.

In the **Playing the System** part of this thesis, I introduced two behavioral cloning models. I trained these models on strategies players used to solve alignment tasks, allowing the models to learn from human problem-solving.

The results were very positive. Not only did the model match the traditional methods, but at times, it even outperformed them. That is all the more important because it outperforms the standard approaches across different conditions. While these traditional methods are based on rigid heuristics, human-driven models appear to capture an intuitive sense that enables them to better handle the complexity involved in sequence alignment.

Encouraged by these findings, in the **Player-Guided AI** chapter, I extended this approach further by introducing a more advanced model in which the previous one would be integrated. This model was bench-marked to a traditional reinforcement learning approach on its efficacy. The results showed clear advantages for the human-inspired model. This pointed to embedding human strategies into computational models as key to better and more scalable solving of bioinformatics challenges.

In the **Cytomerty** part of this thesis, I extend our approach to a new kind of data and focus on the flow cytometry clustering problem. Training the Behavioral Cloning model to do the challenging tasks in the clustering problems was done. Although the focus shifted from sequence alignment to cytometry, the results remained consistent. The approach adapted effectively to this new form of data. What stands out in this work is the demonstration of the flexibility and applicability of this approach to a wide range of bioinformatics problems. There is significant potential here, with many problems that could benefit from the combination of human creativity and computational power.

6.0.1 Data Limitations

While this work highlights the success of imitation learning models trained on data from citizen science games, the strategies used to select this data also introduce important limitations that may affect the models' generalization and broader applicability.

In the case of Borderlands Science, training data was filtered to include only those player solutions that lay close to the Pareto front, optimizing the trade-off between game score and number of gaps. This approach ensures high-quality demonstrations and promotes efficient strategy learning. However, it also narrows the diversity of training examples by excluding suboptimal yet informative solutions. As a result, the models may struggle to recover from imperfect states, since they were not exposed to a wider spectrum of human strategies, including corrections or explorations off the optimal path. Furthermore, Pareto-based filtering is inherently tied to the game's internal scoring logic. While effective within that system, this may inadvertently encode game-specific biases that do not reflect broader biological relevance or utility beyond the game context.

For Project Discovery in EVE Online, clustering labels were derived from consensus-based aggregation of multiple players' solutions. Although consensus is a robust way to reduce noise and achieve high label quality, it also has its drawbacks. The aggregation process removes individual variation, potentially masking less common but valid strategies. It also assumes that the majority interpretation is always correct, which can reinforce systematic biases or errors shared among the player base. Most importantly, this approach captures only the final results of clustering, not the human decision-making process leading up to them. As a consequence, imitation learning models trained on consensus data may learn to reproduce the output but not the strategy, reducing their ability to extrapolate to novel or ambiguous data.

These limitations highlight a trade-off: while filtering and aggregation improve data consistency and label quality, they may also limit the model's exposure to the full range of human reasoning and flexibility. This can lead to inflated performance on curated in-distribution data while hindering generalization to real-world or less-structured scenarios. Future work may benefit from integrating a wider range of player behaviors and from capturing not just what players decided, but how they reached those decisions. Doing so could further strengthen the connection between human insight and scalable AI systems trained through gameplay.

6.0.2 Gameplay Insights

The potential to combine human intuition with computational power is an exciting prospect in dealing with complex problems. However, effective human problem-solving is much more than just fancy recombinations of how to present the problem for solution.

In particular, the challenge in `Borderlands Science` was to improve microbial sequence alignments by making use of human pattern recognition. While the question being asked is more fundamental, it is all the same abstract and mostly algorithm-driven.

It took some abstraction to set it up as a game, translating nucleotide sequences with colored tiles bearing the face of main characters that are moved by players to achieve maximum alignment while minimizing gaps. This was a simplified graphical representation that maintained many key biological constraints of the problem while allowing players to leverage intuition rather than formal training in making.

Interestingly, these alignments from players were actually often better than those from standard algorithms. For example, players might perceive patterns or overlaps which algorithms could not, such as subtle compensatory changes indicative of conserved secondary structures. Of course, this success did not come without its challenges. While the essence of alignment was preserved in this game, some of the biological subtleties-like how RNA secondary structure informs sequence conservation-were lost in abstraction. This required rigorous validation to make sure that player-generated alignments were scientifically valid. Of course, not all of the challenges of cytometry clustering lent themselves quite so easily to gamification. Data produced in flow cytometry are complex, high-dimensional captures of fluorescence intensities across many markers for millions of cells. For `EVE Online's Project Discovery`, this was reimagined for the players by reducing the data down to two-dimensional scatterplots, which represent different projections in multi-dimensional space. These provided a visual approximation of structure in the data, with players able to sketch over clusters by drawing boundaries around points. While there has ever been a set of exciting opportunities, citizen science games still have to deal with all those challenges typical of gaming in general: addiction, poor sleeping, and many other social well-being concerns.

For that reason, the design of games need to be approached with such possible downsides in order to assure that they might not overshadow the gains of the games. While these concerns are attached to larger issues of problem gaming, it is not yet clear if the effects have their full impact in these games compared to traditional forms. Further research is needed to understand properly how these impact players, in order to develop ways of maximizing engagement while minimizing harm. While there is certainly a risk of it, gamification has mostly been in use for the facilitation of bridging between human intuition and computational challenges.

This methodology raised concerns, however, about the fidelity of the results, since the players might misinterpret distortions in data that had been simplified. In this vein, tasks were made redundant, so that multiple players analyzed the same data for consistency and reliability. Expert curation then further validated these most contentious results, integrating human intuition with biological expertise. It is this iterative relationship-in which algorithms are informed by human judgments, and human insight is refined through algorithmic output-that best illustrates the full potential of hybrid systems based on gamification.

6.0.3 Implications for Game Design

These gameplay insights also provide fertile ground for shaping the next generation of citizen science platforms. Building on the lessons from `Borderlands Science` and `Project Discovery`, we can draw clear implications for designing future games that not only attract and retain users but also produce useful data for scientific research. This includes well-designed user interfaces, smart reward systems, and strong technical setups that support large-scale participation and easy integration with machine learning. As future problems become more abstract and complex, design choices will play a key role in keeping tasks both accessible to players and scientifically meaningful.

On the gameplay side, both platforms showed how difficult scientific problems can be turned into familiar and enjoyable experiences. These examples highlight the importance of reducing complexity through simple interactions and visual clarity, while still keeping the underlying task challenging and scientifically valid. Features like limited action options, clear scoring, and progression through increasing difficulty helped guide players toward better solutions while keeping them engaged.

Motivation was another key element. While rewards such as in-game currency, items, or rankings helped keep players active, many were also driven by the feeling that they were contributing to real science. Both games made efforts to communicate the scientific purpose of the tasks, using in-game messages, status updates, or community goals. Players were more likely to stay involved when they saw that their actions had real value. Future games can build on this by showing how individual contributions improve models or advance research, and by offering more personal or social feedback—such as progress bars, community milestones or comparison to peers.

From a technical perspective, these games were carefully designed to produce data that was immediately useful. Each move a player made represented a specific action in the scientific task, whether aligning sequences or outlining cell populations. The Borderlands system also tracked not just the final results, but the entire decision process, which is especially useful for training imitation learning models that aim to learn human-like strategies. Aggregating multiple player solutions through consensus methods improved the reliability of the data, turning variation among users into a strength. This suggests that future platforms should include built-in ways to gather and combine multiple views on the same problem, and to identify which solutions are most trustworthy.

As citizen science expands to other areas, some problems may not have such obvious visual or game-like representations. More abstract tasks—such as modeling networks, solving equations or exploring theoretical spaces—will require new ways of thinking about game design. One approach could be to split big problems into smaller, easier ones, or to build simulations where player actions follow scientific rules. Storytelling, tutorials, and adaptive difficulty could also help teach players the skills they need as they progress. These ideas will be important for keeping players engaged while still collecting valid and useful data.

Finally, future systems could use a combination of human and machine input. Models trained on earlier player data could suggest starting points or complete easier parts of puzzles, leaving more difficult cases for humans. This hybrid approach can save time while still taking advantage of human intuition. To make this work well, systems should continue to record how decisions are made, not just what the answers are. This richer data opens the door for training models that better reflect how humans solve problems.

In summary, these insights point to a design strategy that brings together user-friendly interaction, strong motivational structure, and solid technical foundations. When these parts are aligned, citizen science games can do more than just engage the public—they can become powerful tools for tackling scientific challenges that are hard for computers alone to solve. With careful design, these platforms can expand the role of players in research and create new paths for collaboration between humans and AI.

6.0.4 Healthcare Applications

Gamification and citizen science are new frontiers representing great promises for bioinformatics, with a very important impact on healthcare diagnostics and drug development. This approach leverages human pattern recognition and computational scalability for solving problems which, if purely algorithmic methods were employed, are traditionally far beyond the limits.

Arguably, one outstanding success in this direction involves gamification in malaria diagnostics. Players analyzing digitized blood smear images reach an accuracy of over 99% in diagnoses [Waz18], thereby often outperforming many algorithmic approaches. Such a platform leverages human spatial reasoning to extend and reinforce diagnostic workflows, particularly for resource-constrained settings where automated solutions may not as yet perform quite so well.

There have been a number of broad applications of gamification in managing chronic diseases, such as diabetes, cancer, and hypertension. These gamification techniques aim to improve the level of patient activity, rehabilitation training, symptom management, and medication adherence. Examples are interactive applications and games that engage one in consistent therapeutic routines in return for rewards, leaderboards, and feedback. However, most studies lack a uniform standard whereby gamification effectiveness is reviewed and hence limits comparability of personal interventions. Interventions should be customized to fit individual patient needs for outcomes to be optimized [HXL+23].

Gamified interventions also proved to work in education and cancer prevention. "Tapamole" amongst other virtual reality games serve in educating melanoma patients, the signs of the disease, and use sun protection. These will eventually help in improving self-efficacy in patients and self-screen on a habitual basis. Games such as "Re-Mission involves young adults in gaining the awareness of the seriousness or severity of cancer and thus its prevention measures. These interventions help in cultivating attitudes that are preventive while encouraging active health behaviors among the clients [PGV+23].

Interventions based on gamification also seem to be of promise in promoting physical activities among cancer survivors. Positive effects on quality of life and general health have indeed been reported for web-based interventions focused on the motivational aspects of moderate physical activities. However, meta-analysis has pointed out that results for those are inconsistent, which widely has been attributed to profound methodological heterogeneity within studies. Larger and strictly designed trials are needed but confirmatory findings are partly established by strong evidence within this context [PGV+23].

My thesis shows how human ingenuity and machine learning can change healthcare. Think of algorithms that are trained on gamers' strategies and adapt to find anomalies in blood smears, diagnosing with precision in remote clinics. Think of how gamified insights might optimize drug discovery, uncovering molecular patterns that would have remained hidden to traditional methods. Melding human intuition with computational power, this approach revolutionizes diagnostics, therapy personalization, and disease prevention-heralding smarter, more accessible healthcare solutions.

6.0.5 Drugs Development

Another domain in which gamified platforms are very likely to make a transformation is that of drug development.

The most noticeable application of gamification in drug development is education and training. The "SCREENER" game is a very good example of how such gamified tools can simulate the whole process of drug discovery and development for educational purposes.

SCREENER was the first hybrid board and card game designed to take them through each stage of a DDD pipeline, including target validation and through into regulatory approval [NXM$^+$21]. The much-needed accessible education tool, SCREENER, needed in this arena, is now available; pharma companies' previous attempts to gamify the process either stalled entirely or were shut behind proprietary firewalls, accessible within companies only. Other evidence of the educational impact of gamification comes from a systematic quantitative review of gamification in pharmacy education. It indicated that gamified interventions are widely adopted across the globe, with 12 countries represented, involving over 8,000 students and health professionals. The interventions ranged from board games to immersive simulations, with the most popular format involving escape rooms [HGRK23, HGRK23].

While these initiatives have also shown promise in increasing engagement and retaining knowledge, the review highlighted the need for more robust reporting and closer alignment between education objectives and outcomes.

Other ingenious uses have come in the form of "MedChem Game," which is an Android-based teaching aid for medicinal chemistry [DŁPP24]. Gamification merged with artificial intelligence in one platform for teaching medicinal chemistry not only provided knowledge at a foundational level to the users but also interactively involved them in designing new small-molecule drugs. AI in MedChem Game navigates a user through a solution of complex tasks in medicinal chemistry in an easy and interactive way. Beyond education, citizen science with gamification has also had remarkable value in the realm of contributing directly to drug discovery.

While gamification holds immense potential in health, there are also many challenges on the way to implementation. Privacy and security of subject data hold prime importance for health care. Regulations in GDPR and HIPAA play an important role in participant privacy. More importantly, such gamified platforms need to be accessible across a wide population in order to eliminate bias and problems of representativeness in the collected data. Accomplishment of these goals can take place by ensuring that multilingual interfaces with adaptive difficulty levels are present to finally ensure inclusivity.

Other challenges include player-generated data validation. Through expert curation and introduction of redundancy in tasks, one can ensure that results of high fidelity are indeed realized; curated data may be used to train more advanced machine learning models further.

In a nutshell, gamification of drug development goes beyond education in rewriting the way drugs are found and optimized. Melding human creativity with computational power, these platforms have created a promising pathway to accelerate therapeutic innovation and develop the next generation of researchers.

My thesis redefines the process of drug development by transforming gamified platforms into tools of education, innovation, and discovery. By leveraging player-generated data and baking inclusivity into practice, your work accelerates drug discovery and democratizes access to state-of-the-art pharmaceutical innovation. The combination of human creativity and AI advances the pace of therapeutic breakthroughs while nurturing the next generation of scientific leaders.

6.0.6 Societal Impact

The impact that citizen science games have is much wider in societal impacts compared to the participation of individuals and scientific data collection.

There are several motivations for people participating in citizen science games. Most studies have shown that the major reasons for citizens participating in citizen science are based on a background interest in science and helping to advance scientific research [MGGC22, PHKK18, RBG$^+$09, FCS$^+$17]. However, in most cases this is followed by a power-law distribution of work pattern with only a few highly active participants [ACL$^+$22].

The use of citizen science games has attracted much interest, especially when obtaining relevant scientific results and in the development of more public involvement in science. On the one hand, it creates an inherent contradiction between the motivational aspects of gaming and the generation of relevant scientific results. Secondly, there is ongoing debate regarding the tensions between game elements and scientific rigour, and for gamification approaches to be tailored to specific tasks and domains [Cur14a, Sch24].

However, the challenges identified in the work do not impede either the evolution or the reach of citizen science games. Recent developments include the introduction of citizen science activities within popular commercial video games that could widen the audience for scientific participation [SGGGB$^+$24]. As such, while this field is still evolving, further research will be required into the experiences of players, motivations, and engagement to realize the impacts and effectiveness within the contexts of citizen science game development for the advance of scientific knowledge and public engagement in science.

I believe, this approach brings us closer to a world where citizen science games become an entry point for society actively to engage in scientific discoveries. By gamifying science, we turn complex problems into engaging puzzles. Therefore, we empower the masses to make remarkable contributions. This work also envisions such platforms as hubs for inclusivity and education, where anyone - regardless of age or background - can learn, play, and leave a real-world impact. Imagine these games as platforms that evolve to become ecosystems, fund education, raise awareness, and drive the next generation of citizen scientists. Beyond data collection, these games have a ripple effect on driving public curiosity about science, fostering a global community of problem-solvers, and redefining what it means for society to "play for change."

6.0.7 The Power of Collaboration: From Citizen Science Games to Scalable AI

This integration of classical bioinformatics, citizen science games, and imitation learning forms a revolutionary paradigm to address complex scientific problems. It combines the flexibility and creativity of human intuition with the precision and scalability of AI, thus achieving unparalleled progress in solving bioinformatics problems.

Citizen science games like `Borderlands Science` and `EVE Online's Project Discovery` have already shown that non-experts can be a great help in solving complex scientific tasks. These platforms convert these abstract bioinformatic challenges into interesting video games, such as aligning microbial sequences or performing cytometry clustering. Since such platforms challenge players in the relevant tasks, they produce insights that uniquely have their origins in the use of human cognitive capacities of pattern recognition and intuition based on tacit thinking. For instance, when alignment of microbial sequences has taken place, players repeatedly bring out conserved motifs along with subtle patterns which existing algorithms had failed to demonstrate.

These insights bring out the importance of human problem-solving in bioinformatics workflows. First, crowdsourcing complex problems through these games allowed for breakthroughs that had been hard to achieve conventionally [Cur14b]. As I mentioned, Foldit has made remarkable breakthroughs in protein structure prediction by demonstrating the power of human collective intelligence in solving computationally hard problems. It complements the computational analyses through human intuition and pattern recognition capabilities, thereby enhancing the data quality and showing relationships or anomalies not easily depicted by an algorithm [Kat19]. This is true in the case of cytometry clustering, among other areas that usually involve high-dimensional data and need subtle interpretation. Above all, these games allow outreach and education: they constitute places of informal science learning, creating the opportunity for the general public to be meaningfully engaged in research, hence encouraging further appreciation for scientific research [Cur14b].

Second, scalability: the models trained by combining AI and human strategy are scalable to bigger datasets and completely new problems [MSGC$^+$23]. This makes the solutions extendable beyond what was envisioned in the scope of the citizen science games themselves and enables breakthroughs in genomics, proteomics, and immunology, among other diverse bioinformatics areas. Third, the iterative nature of this framework supports continuous improvement. As more and more people start using these platforms and hence feeding data into them, the AI models start to evolve and get refined. This continuous cycle of human input and machine learning ensures that the models remain robust and accurate, even as the complexity and scale of the tasks increase over time.

This synergy eventually leads to new research directions in both bioinformatics and artificial intelligence [YA23, LYY$^+$24]. This marriage of strengths brings forth completely new directions of interdisciplinary research, such as using human strategies to help focus machine learning, and may result, for example, in next-generation advances in personalized medicine in which nuanced human insights complement data-driven approaches. Using this framework sets up the platform to further investigate the gamification of AI for domains such as environmental monitoring and healthcare diagnostics. Imitation learning has become a critical bridge between human problem-solving and AI. For instance, behavioral cloning can enable AI systems to closely mimic player strategies by distilling decision-making processes underlying human expertise. These models can go beyond simple replication to generalize human strategies toward much larger and more complex problems.

For instance, AI models trained on player-generated data for microbial sequence alignment have performed not only comparably to traditional algorithms but have even outperformed them in cases with noisy or highly variable datasets.

Central to this is the "humans-in-the-loop" paradigm: the integration of human insights directly into computational workflows. The players provide initial data by solving gamified tasks, while their solutions inform the training of machine learning models. In return, the strategies are fine-tuned and scaled using these models, creating a feedback loop that improves the contributions by humans and machines. This two-way relationship enriches the interaction in finding solutions not only with high accuracy and efficiency but also in understanding deeper biological insights from the problem at hand.

In conclusion, citizen science games in combination with classical bioinformatics and imitation learning are the power of interdisciplinary collaboration, enhancing problem-solving capabilities, enabling scalability and continuous improvement, and new research directions. This is the new paradigm for how scientific challenges are approached. As this synergy develops further, it might turn out to change the panorama of bioinformatics and unleash a similar approach in a wide range of scientific and technological domains.

6.1 Conclusion

This work demonstrates the potential transformation possible through the integration of citizen science, bioinformatics, and imitation learning—a bold step toward solving some of the most intractable problems in genomics and systems biology. The successes reported herein stand as testimony to the power of collaboration between human intuition and computational precision, offering new avenues by which to explore limitations and opportunities within these crossroads. But the way ahead needs critical reflection, with an eye to what has been accomplished and remaining complexity.

Basically, this research underlines the power of leveraging human problem-solving strategies in computational models to solve bioinformatics challenges. The successes within MSA and cytometry clustering are exemplary of how insights guided by humans can lead to the improvement in the performance of algorithms, especially while working with variability and noise that confound traditional methods. Yet, these successes should not be seen as panaceas but rather as proofs-of-concept for what can be achieved when human creativity and machine learning harmonize. They represent the beginnings, a scaffolding upon which more comprehensive, developed methodologies can be built.

Yet, the ramifications of this work go far beyond the direct applications of MSA and cytometry. They question the very borders of the traditional areas of computational methods, underlining the role of human intuition as an integral part of the problem-solving pipeline. Bioinformatics as a discipline is increasingly having to come to grips with data that is not only vast but also multi-dimensional and deeply interconnected. This, in a sense, points to a view whereby the future of bioinformatics will be based on whether it can afford hybrid systems where human insight is regularly captured, enhanced, and embedded into scalable computation flows. This is how such a redefinition of workflow with human-computer collaboration holds the potential for change: in the ways of our thinking about scalability, adaptability, and even creativity in scientific problem-solving.

This work also highlights the fragility of such a schema. Particularly when done through gamification, depending on human-derived strategies brings in complexities that cannot be ignored. This abstraction used to translate the bioinformatics problems into game mechanics may result in a loss of critical nuances and, in turn, can compromise the biological validity of solutions generated by players. Similarly, while redundancy and expert curation go some way to mitigate these risks, they also suggest the importance of ensuring that human contributions make sense in intricate, sometimes hidden, relationships between biological data. It is this tension between simplification for accessibility and scientific rigour where future research will need to tread with a great deal of care.

Another powerful dimension, however, is the intersection of innovation with inclusiveness. This study democratizes participation in bioinformatics, allowing non-experts to meaningfully contribute to scientific discovery through engaging the public in citizen science games.

While it enhances not only the pool of data with diverse perspectives but also fosters shared ownership of science, participation itself is simply not enough for democratization. This has to be sustained with the education and empowerment of contributors toward ensuring equitably distributed and far-reaching benefits of the system. This becomes even more important in a global context where access to such technologies and opportunities might be distributed very unevenly. The implications are huge for the future. The methodologies pioneered here extend an invitation toward the exploration of even more complex challenges in bioinformatics, like systems-level interactions in proteomics, epigenomics, and microbiomics. These areas involve approaches able to explore intricate networks of biological relations while adapting to the scale and diversity of the underlying data. At the same time, this work points toward a general philosophical question: *how is one to design the interplay between human intuition and machine precision so as to maximize discovery without losing the particular advantages of each?*

In all probability, the explanation is rooted in deeper interdisciplinary collaboration. As this work illustrates, the intersection of domain knowledge in bioinformatics, machine learning engineering, and game design is not merely helpful but rather crucial. The future of hybrid systems in science depends on how well these disciplines can come together, playing to their respective strengths to develop tools that are rigorous yet accessible. Moreover, incorporation of recent advances in areas like explainable AI and multimodal data integration will be critical in making the developed systems interpretable and robust. The work, therefore, is not an end but a beginning; a foundation upon which more sophisticated, inclusive, and impactful systems can be built. The combination of citizen science, bioinformatics, and imitation learning is not a methodological advance but rather a paradigm shift in the way we approach some of the most complex scientific problems. This research provides a kind of roadmap to the future in recognizing the limits of these first steps while celebrating their promise a vision of science collaborative, adaptive, and profoundly human-centered. It represents an invitation to further innovation grounded in the conviction that the most important discoveries lie ahead where creativity, rigor, and inclusivity intersect disciplines.

Bibliography

[ACL+22] Bradley C Allf, Caren B Cooper, Lincoln R Larson, Robert R Dunn, Sara E Futch, Maria Sharova, and Darlene Cavalier. Citizen science as an ecosystem of engagement: Implications for learning and broadening participation. *BioScience*, 72(7):651–663, 2022.

[Cur14a] Vickie Curtis. Online citizen science games: Opportunities for the biological sciences. *Applied & translational genomics*, 3(4):90–94, 2014.

[Cur14b] Vickie Curtis. Online citizen science games: Opportunities for the biological sciences. *Applied Translational Genomics*, 3(4):90–94, 2014.

[DŁPP24] Tomasz Danel, Jan Łski, Sabina Podlewska, and Igor T Podolak. Medchem game: gamification of drug design, 2024.

[FCS+17] T Frensley, Alycia Crall, Marc Stern, Rebecca Jordan, Steven Gray, Michelle Prysby, Greg Newman, Cindy Hmelo-Silver, David Mellor, and Joey Huang. Bridging the benefits of online and community supported citizen science: A case study on motivation and retention with conservation-oriented volunteers. *Citizen Science: Theory and Practice*, 2(1):4–4, 2017.

[HGRK23] Denise L Hope, Gary D Grant, Gary D Rogers, and Michelle A King. Gamification in pharmacy education: a systematic quantitative literature review. *International Journal of Pharmacy Practice*, 31(1):15–31, 2023.

[HXL+23] Xiting Huang, Xinyue Xiang, Yang Liu, Zhiqian Wang, Zhili Jiang, and Lihua Huang. The use of gamification in the self-management of patients with chronic diseases: scoping review. *JMIR Serious Games*, 11:e39019, 2023.

[Kat19] Tarun Reddy Katapally. The smart framework: integration of citizen science, community-based participatory research, and systems science for population health science in the digital age. *JMIR mHealth and uHealth*, 7(8):e14056, 2019.

[LYY$^+$24] Jiajia Liu, Mengyuan Yang, Yankai Yu, Haixia Xu, Kang Li, and Xiaobo Zhou. Large language models in bioinformatics: applications and perspectives. *ArXiv*, 2024.

[MGGC22] Josh Aaron Miller, Kutub Gandhi, Anna Gander, and Seth Cooper. A survey of citizen science gaming experiences. *Citizen Science: Theory and Practice*, 7(1):1–12, 2022.

[MSGC$^+$23] Renata Mutalova, Roman Sarrazin-Gendron, Eddie Cai, Gabriel Richard, Parham Ghasemloo Gheidari, Sébastien Caisse, Rob Knight, Mathieu Blanchette, Attila Szantner, and Jérôme Waldispühl. Playing the system: Can puzzle players teach us how to solve hard problems? In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2023.

[NXM$^+$21] F Noël, G Xexéo, E Mangeli, A Mothé, P Marques, J Kritz, F Blanchard, H Vermelho, and B de Paiva. Screener, an educational game for teaching the drug discovery and development process. *Brazilian Journal of Medical and Biological Research*, 54:e11786, 2021.

[PGV$^+$23] Andrea Poliani, Silvia Gnecchi, Giulia Villa, Debora Rosa, and Duilio F Manara. Gamification as an educational approach for oncological patients: a systematic scoping review. In *Healthcare*, volume 11, page 3116. MDPI, 2023.

[PHKK18] Marisa Ponti, Thomas Hillman, Christopher Kullenberg, and Dick Kasperowski. Getting it right or being top rank: Games in citizen science. *Citizen Science: Theory and Practice*, 3(1):1–1, 2018.

[RBG$^+$09] M Jordan Raddick, Georgia Bracey, Pamela L Gay, Chris J Lintott, Phil Murray, Kevin Schawinski, Alexander S Szalay, and Jan Vandenberg. Galaxy zoo: Exploring the motivations of citizen science volunteers. *arXiv preprint arXiv:0909.2925*, 2009.

[Sch24] Karen Schrier. Leadership through citizen science games. *Journal of Leadership Studies*, 17(4):32–37, 2024.

[SGGGB+24] Roman Sarrazin-Gendron, Parham Ghasemloo Gheidari, Alexander Butyaev, Timothy Keding, Eddie Cai, Jiayue Zheng, Renata Mutalova, Julien Mounthanyvong, Yuxue Zhu, Elena Nazarova, et al. Improving microbial phylogeny with citizen science within a mass-market video game. *Nature Biotechnology*, pages 1–9, 2024.

[Waz18] Kerri Wazny. Applications of crowdsourcing in health: an overview. *Journal of global health*, 8(1), 2018.

[YA23] Malik Yousef and Jens Allmer. Deep learning in bioinformatics. *Turkish Journal of Biology*, 47(6):366–382, 2023.

Bibliography

[Aca23] EITCA Academy. How does alphastar's use of imitation learning from human gameplay data differ from its reinforcement learning through self-play, and what are the benefits of combining these approaches?, 2023. Accessed: 2024-10-17.

[ACL+22] Bradley C Allf, Caren B Cooper, Lincoln R Larson, Robert R Dunn, Sara E Futch, Maria Sharova, and Darlene Cavalier. Citizen science as an ecosystem of engagement: Implications for learning and broadening participation. *BioScience*, 72(7):651–663, 2022.

[ACVB09] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.

[AFH+13] Nima Aghaeepour, Greg Finak, Holger Hoos, Tim R Mosmann, Ryan Brinkman, Raphael Gottardo, and Richard H Scheuermann. Critical assessment of automated flow cytometry data analysis techniques. *Nature Methods*, 10(3):228–238, 2013.

[AKDL22] Vladislav Ayzenberg, Frederik S Kamps, Daniel D Dilks, and Stella F Lourenco. Skeletal representations of shape in the human visual cortex. *Neuropsychologia*, 164:108092, 2022.

[AMM+20] Ko Abe, Kodai Minoura, Yuka Maeda, Hiroyoshi Nishikawa, and Teppei Shimamura. Model-based clustering for flow and mass cytometry data with clinical information. *BMC bioinformatics*, 21:1–15, 2020.

[AN16] Massimo Andreatta and Morten Nielsen. Gapped sequence alignment using artificial neural networks: application to the mhc class i system. *Bioinformatics*, 32(4):511–517, 2016.

[AP10] Daniel E Acuña and Víctor Parada. People efficiently explore the solution space of the computationally intractable traveling salesman problem to find near-optimal tours. *PloS one*, 5(7):e11685, 2010.

[ar520] ar5iv Labs. Computational models mimicking human visual recognition processes, 2020.

[ASF23] Xusheng Ai, Melissa C Smith, and Frank Alex Feltus. Generative adversarial networks applied to gene expression analysis: an interdisciplinary perspective. *Computational and Systems Oncology*, 3(3):e1050, 2023.

[Bak24] David Baker. The nobel prize in chemistry 2024, 2024. Accessed: 2025-01-10.

[BBC+19] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

[BDTD+16] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016. Describes the application of convolutional neural networks for training autonomous vehicles by imitating human driving behaviors.

[BDTSW22] Alexander Butyaev, Chrisostomos Drogaris, Olivier Tremblay-Savard, and Jérôme Waldispühl. Human-supervised clustering of multidimensional data using crowdsourcing. *Royal Society open science*, 9(5):211189, 2022.

[BKR$^+$04] Mathieu Blanchette, W James Kent, Cathy Riemer, Laura Elnitski, Arian FA Smit, Krishna M Roskin, Robert Baertsch, Kate Rosenbloom, Hiram Clawson, Eric D Green, et al. Aligning multiple genomic sequences with the threaded blockset aligner. *Genome research*, 14(4):708–715, 2004.

[Bra08] Daren C Brabham. Moving the crowd at istockphoto: The composition of the crowd and motivations for participation in a crowdsourcing application. *First monday*, 2008.

[BS95] Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Machine Intelligence 15*, pages 103–129, 1995.

[BSH$^+$21] Petra Baumgaertner, Martial Sankar, Fernanda Herrera, Fabrizio Benedetti, David Barras, Anne-Christine Thierry, Denarda Dangaj, Lana E Kandalaft, George Coukos, Ioannis Xenarios, et al. Unsupervised analysis of flow cytometry data in a clinical setting captures cell diversity and allows population discovery. *Frontiers in immunology*, 12:633910, 2021.

[BWP$^+$22] Raunak Bhattacharyya, Blake Wulfe, Derek J Phillips, Alex Kuefler, Jeremy Morton, Ransalu Senanayake, and Mykel J Kochenderfer. Modeling human driving behavior through generative adversarial imitation learning. *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[CEO14] Ilaria Caponetto, Jeffrey Earp, and Michela Ott. Gamification and education: A literature review. In *European conference on games based learning*, volume 1, page 50. Academic Conferences International Limited, 2014.

[CF19] Ioana Carcea and Robert C Froemke. Biological mechanisms for observational learning. *Current opinion in neurobiology*, 54:178–185, 2019.

[CKT$^+$08] S. Cooper, F. Khatib, A. Treuille, J. Barbero, J. Lee, M. Beenen, A. Leaver-Fay, D. Baker, Z. Popović, and Foldit Players.

Foldit, 2008. An online puzzle video game about protein folding. It is part of an experimental research project developed by the University of Washington, Center for Game Science, in collaboration with the UW Department of Biochemistry.

[CKT+10] Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker, Zoran Popović, et al. Predicting protein structures with a multiplayer online game. *Nature*, 466(7307):756–760, 2010.

[CMC+16] Maria Chatzou, Cedrik Magis, Jia-Ming Chang, Carsten Kemena, Giovanni Bussotti, Ionas Erb, and Cedric Notredame. Multiple sequence alignment modeling: methods and applications. *Briefings in bioinformatics*, 17(6):1009–1023, 2016.

[CMG+19] Antoine Coutrot, Ed Manley, Sarah Goodroe, Christopher Gahnstrom, Joseph Filomena, Jan M Wiener, Ruth C Dalton, Michael Hornberger, and Hugo J Spiers. The impact of mobile-based navigation training on real-world navigation skills. *Scientific Reports*, 9:4692, 2019.

[CSBG18] Irene Cristofori, Carola Salvi, Mark Beeman, and Jordan Grafman. The effects of expected reward on creative problem solving. *Cognitive, Affective, & Behavioral Neuroscience*, 18(5):925–931, 2018.

[CSLG19] Felipe Codevilla, Eder Santana, Antonio M López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9329–9338, 2019.

[CSM+18] Antoine Coutrot, Ricardo Silva, Ed Manley, Will de Cothi, Saber Sami, Véronique D Bohbot, Jan M Wiener, Christoph Hölscher, Ruth C Dalton, Michael Hornberger, et al. Global determinants of navigation ability. *Current Biology*, 28(17):2861–2866, 2018.

[CTG⁺21] Brian Chen, Siddhant Tandon, David Gorsich, Alex Gorodetsky, and Shravan Veerapaneni. Behavioral cloning in atari games using a combined variational autoencoder and predictor model. In *2021 IEEE Congress on Evolutionary Computation (CEC)*, pages 2077–2084. IEEE, 2021.

[Cur14a] Vickie Curtis. Online citizen science games: Opportunities for the biological sciences. *Applied & translational genomics*, 3(4):90–94, 2014.

[Cur14b] Vickie Curtis. Online citizen science games: Opportunities for the biological sciences. *Applied Translational Genomics*, 3(4):90–94, 2014.

[CZW⁺20] Xinyue Chen, Zijian Zhou, Zheng Wang, Che Wang, Yanqiu Wu, and Keith Ross. Bail: Best-action imitation learning for batch deep reinforcement learning. *Advances in Neural Information Processing Systems*, 33:18353–18363, 2020.

[DBH16] Shreyansh Daftry, J Andrew Bagnell, and Martial Hebert. Learning transferable policies for monocular reactive mav control. In *International Symposium on Experimental Robotics*, pages 3–11. Springer, 2016.

[DBK⁺16] Yue Deng, Wei Bao, Hongwei Kong, Yufei Ren, Yong Dai, and Ruoyu Wang. Deep direct reinforcement learning for financial signal representation and trading. In *Proceedings of the IEEE Conference on Computational Intelligence for Financial Engineering & Economics*, pages 1–7. IEEE, 2016.

[dBTW08] Hans P Op de Beeck, Katrien Torfs, and Johan Wagemans. Perceived shape similarity among unfamiliar objects and the organization of the human object vision pathway. *Journal of Neuroscience*, 28(40):10111–10123, 2008.

[DCLT18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[Dee19] DeepMind. AlphaStar: Grandmaster level in StarCraft II using multi-agent reinforcement learning. `https://deepmind.google/discover/blog/alphastar-grandmaster-level-in-starcraft-ii-using-multi-agent-reinfor` 2019. Accessed: 2025-01-10.

[Dee20] DeepMind. AlphaStar: Mastering the real-time strategy game StarCraft II. `https://deepmind.google/discover/blog/alphastar-mastering-the-real-time-strategy-game-starcraft-ii/`, 2020. Accessed: 2025-01-10.

[DFI18] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. Demographics and dynamics of mechanical turk workers. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 135–143, 2018.

[DŁPP24] Tomasz Danel, Jan Łski, Sabina Podlewska, and Igor T Podolak. Medchem game: gamification of drug design, 2024.

[DLR77] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.

[DMBB05] Chuong B Do, Mahathi SP Mahabhashyam, Michael Brudno, and Serafim Batzoglou. Probcons: Probabilistic consistency-based multiple sequence alignment. *Genome research*, 15(2):330–340, 2005.

[Edd92] Sean Eddy. Hmmer user's guide. *Department of Genetics, Washington University School of Medicine*, 2(1):13, 1992.

[Edd98] Sean R. Eddy. Profile hidden markov models. *Bioinformatics (Oxford, England)*, 14(9):755–763, 1998.

[Edd22] Sean Eddy. 2022 essential facts about the video game industry. *Entertainment Software Association*, 2022.

[Edg04] Robert C Edgar. Muscle: multiple sequence alignment with high accuracy and high throughput. *Nucleic acids research*, 32(5):1792–1797, 2004.

[ESB+12] Christopher B Eiben, Justin B Siegel, Jacob B Bale, Seth Cooper, Firas Khatib, Betty W Shen, Barry L Stoddard, Zoran Popovic, and David Baker. Increased diels-alderase activity through backbone remodeling guided by foldit players. *Nature biotechnology*, 30(2):190–192, 2012.

[ESd+17] Anna Everett, Braxton Soderman, Jennifer deWinter, Carly Kocurek, Nina B Huntemann, Gabrielle Trepanier-Jobin, Irene Chien, Soraya Murray, Rachael Hutchinson, Lisa Patti, et al. *Gaming representation: Race, gender, and sexuality in video games.* Indiana University Press, 2017.

[Far19] Wael Farag. Cloning safe driving behavior for self-driving cars using convolutional neural networks. *Recent Patents on Computer Science*, 12(2):120–127, 2019.

[FCS+17] T Frensley, Alycia Crall, Marc Stern, Rebecca Jordan, Steven Gray, Michelle Prysby, Greg Newman, Cindy Hmelo-Silver, David Mellor, and Joey Huang. Bridging the benefits of online and community supported citizen science: A case study on motivation and retention with conservation-oriented volunteers. *Citizen Science: Theory and Practice*, 2(1):4–4, 2017.

[FDSF18] Tharindu Fernando, Simon Denman, Sridha Sridharan, and Clinton Fookes. Learning temporal strategic relationships using generative adversarial imitation learning. *arXiv preprint arXiv:1805.04969*, 2018.

[FT04] Bent Fuglede and Flemming Topsoe. Jensen-shannon divergence and hilbert space embedding. In *International Symposium onInformation Theory, 2004. ISIT 2004. Proceedings.*, page 31. IEEE, 2004.

[Gam03] CCP Games. Eve online. *World Wide Web. http://www. eve-online. com*, 2003.

[Gam21] CCP Games. Eve online's "project discovery" wins webby people's voice award for public service, activism, and social impact, 2021. Accessed: 2023-06-23.

[Gam24] CCP Games. Phase four of eve online's project discovery revealed, expanding to cancer research and a companion mobile app launching in 2024, 2024. Accessed: 2023-06-23.

[GBC16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep learning (adaptive computation and machine learning series), 2016.

[GBGC+02] Ivo Grosse, Pedro Bernaola-Galván, Pedro Carpena, Ramón Román-Roldán, Jose Oliver, and H Eugene Stanley. Analysis of symbolic sequences using the jensen-shannon divergence. *Physical Review E*, 65(4):041905, 2002.

[GEFL18] Abhishek Gupta, Benjamin Eysenbach, Chelsea Finn, and Sergey Levine. Continuous adaptation via meta-learning in nonstationary and competitive environments. pages 1882–1891, 2018.

[GHP+14] Craig A Gedye, Ali Hussain, Joshua Paterson, Alannah Smrke, Harleen Saini, Danylo Sirskyj, Keira Pereira, Nazleen Lobo, Jocelyn Stewart, Christopher Go, et al. Cell surface profiling using high-throughput flow cytometry: a platform for biomarker discovery and analysis of cellular heterogeneity. *PloS one*, 9(8):e105602, 2014.

[GHT+19] William H Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, and Ruslan Salakhutdinov. Minerl: A large-scale dataset of minecraft demonstrations. *arXiv preprint arXiv:1907.13440*, 2019.

[GMR97] Stephen Grossberg, Ennio Mingolla, and William D Ross. Visual brain and visual perception: how does the cortex do perceptual grouping? *Trends in neurosciences*, 20(3):106–111, 1997.

[GPAM+20] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

[HBT$^+$13] Moritz Helmstaedter, Kevin L Briggman, Srinivas C Turaga, Viren Jain, H Sebastian Seung, and Winfried Denk. Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature*, 500(7461):168–174, 2013.

[HE16] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29:4565–4573, 2016.

[HGEJ17] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.

[HGRK23] Denise L Hope, Gary D Grant, Gary D Rogers, and Michelle A King. Gamification in pharmacy education: a systematic quantitative literature review. *International Journal of Pharmacy Practice*, 31(1):15–31, 2023.

[HHRRR13] Mercedes Hidalgo-Herrero, Pablo Rabanal, Ismael Rodriguez, and Fernando Rubio. Comparing problem solving strategies for np-hard optimization problems. *Fundamenta Informaticae*, 124(1-2):1–25, 2013.

[HKWS19] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.

[HMT18] Trevor Huff, Navid Mahabadi, and Prasanna Tadi. Neuroanatomy, visual cortex. 2018.

[HS97] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[HSX10] Hong-Wei Huo, Vojislav Stojkovic, and Qiao-Luan Xie. A quantum-inspired genetic algorithm based on probabilistic coding for multiple sequence alignment. *Journal of bioinformatics and computational biology*, 8(01):59–75, 2010.

[HVP$^+$18] Todd Hester, Matej Vecerik, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Dan Horgan, John Quan, Andrew

Sendonaris, Ian Osband, et al. Deep q-learning from demonstrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[HXL$^+$23] Xiting Huang, Xinyue Xiang, Yang Liu, Zhiqian Wang, Zhili Jiang, and Lihua Huang. The use of gamification in the self-management of patients with chronic diseases: scoping review. *JMIR Serious Games*, 11:e39019, 2023.

[HZRS16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[IS13] Lilly C Irani and M Six Silberman. Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 611–620, 2013.

[JJKR19] Reza Jafari, Mohammad Masoud Javidi, and Marjan Kuchaki Rafsanjani. Using deep reinforcement learning approach for solving the multiple sequence alignment problem. *SN Applied Sciences*, 1(6):1–12, 2019.

[Joe21] Roman Joeres. Multiple sequence alignment using deep reinforcement learning. *SKILL 2021*, 2021.

[Joh67] Stephen C Johnson. Hierarchical clustering schemes. *Psychometrika*, 32(3):241–254, 1967.

[JTKRJ20] Reza Jafari, Hamed Tabrizchi, Marjan Kuchaki Rafsanjani, and M.M. Javidi. A distributed deep q-learning based model for solving the multiple sequence alignment problem. 04 2020.

[Kar75] Richard M Karp. On the computational complexity of combinatorial problems. *Networks*, 5(1):45–68, 1975.

[Kat19] Tarun Reddy Katapally. The smart framework: integration of citizen science, community-based participatory research, and systems science for population health science in the digital age. *JMIR mHealth and uHealth*, 7(8):e14056, 2019.

[KCS08] Aniket Kittur, Ed H Chi, and Bongwon Suh. Crowdsourcing user studies with mechanical turk. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 453–456, 2008.

[KCT$^+$11] Firas Khatib, Seth Cooper, Michael D Tyka, Kefan Xu, Ilya Makedon, Zoran Popović, David Baker, and Foldit Players. Algorithm discovery by protein folding game players. *Proceedings of the National Academy of Sciences*, 108(47):18949–18953, 2011.

[KDG$^+$11] Firas Khatib, Frank DiMaio, Foldit Contenders Group, Foldit Void Crushers Group, Seth Cooper, Maciej Kazmierczyk, Miroslaw Gilski, Szymon Krzywda, Helena Zabranska, Iva Pichova, et al. Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nature structural & molecular biology*, 18(10):1175–1177, 2011.

[KGL$^+$98] Matthias J Koepp, Roger N Gunn, Andrew D Lawrence, Vincent J Cunningham, Alain Dagher, Tasmin Jones, David J Brooks, Christopher J Bench, and PM Grasby. Evidence for striatal dopamine release during a video game. *Nature*, 393(6682):266–268, 1998.

[KGZ$^+$14] Jinseop S Kim, Matthew J Greene, Aleksandar Zlateski, Kisuk Lee, Mark Richardson, Srinivas C Turaga, Michael Purcaro, Matthew Balkam, Amy Robinson, Bardia F Behabadi, et al. Space–time wiring specificity supports direction selectivity in the retina. *Nature*, 509(7500):331–336, 2014.

[KLG20] Mengmeng Kuang, Yong Liu, and Lufei Gao. Dlpalign: A deep learning based progressive alignment method for multiple protein sequences. In *CSBio'20: Proceedings of the Eleventh International Conference on Computational Systems-Biology and Bioinformatics*, pages 83–92, 2020.

[KMKM02] Kazutaka Katoh, Kazuharu Misawa, Kei-ichi Kuma, and Takashi Miyata. Mafft: a novel method for rapid multiple sequence alignment based on fast fourier transform. *Nucleic acids research*, 30(14):3059–3066, 2002.

[KMWK17] Alex Kuefler, Jeremy Morton, Tim Wheeler, and Mykel Kochenderfer. Imitating driver behavior with generative adversarial networks. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 204–211. IEEE, 2017.

[KPC94] Jin Kim, Sakti Pramanik, and Moon Jung Chung. Multiple sequence alignment using simulated annealing. *Bioinformatics*, 10(4):419–426, 1994.

[KPH20] Anssi Kanervisto, Joonas Pussinen, and Ville Hautamäki. Benchmarking end-to-end behavioural cloning on video games. In *2020 IEEE conference on games (CoG)*, pages 558–565. IEEE, 2020.

[KRK$^+$11] A. Kawrykow, G. Roumanis, A. Kam, D. Kwak, C. Leung, C. Wu, E. Zarour, L. Sarmenta, M. Blanchette, and J. Waldispühl. Phylo: A citizen science approach for improving multiple sequence alignment, 2011. A citizen science platform that converts multiple sequence alignment problems into a game, engaging the public in solving complex genetic puzzles.

[KRK$^+$12] Alexander Kawrykow, Gary Roumanis, Alfred Kam, Daniel Kwak, Clarence Leung, Chu Wu, Eleyine Zarour, Phylo players, Luis Sarmenta, Mathieu Blanchette, et al. Phylo: a citizen science approach for improving multiple sequence alignment. *PloS one*, 7(3):e31362, 2012.

[LBD$^+$89] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.

[LFDA16] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. In *Proceedings of the 29th International Conference on Neural Information Processing Systems*, pages 3572–3580, 2016.

[LG08] Ari Löytynoja and Nick Goldman. Phylogeny-aware gap placement prevents errors in sequence alignment and evolutionary analysis. *science*, 320(5883):1632–1635, 2008.

[LH13] Baoli Li and Liping Han. Distance weighted cosine similarity measure for text classification. In *International conference on intelligent data engineering and automated learning*, pages 611–618. Springer, 2013.

[LHP$^+$15] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *Proceedings of the International Conference on Learning Representations*, 2015.

[LKL$^+$14] Jeehyung Lee, Wipapat Kladwang, Minjae Lee, Daniel Cantu, Martin Azizyan, Hanjoo Kim, Alex Limpaecher, Snehal Gaikwad, Sungroh Yoon, Adrien Treuille, et al. Rna design rules from a massive open laboratory. *Proceedings of the National Academy of Sciences*, 111(6):2122–2127, 2014.

[Llo82] Stuart Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.

[LLY$^+$19] Xin Lu, Buyu Li, Yuxin Yue, Quanquan Li, and Junjie Yan. Grid r-cnn. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7363–7372, 2019.

[LMYDM22] Luc Le Mero, Dewei Yi, Mehrdad Dianati, and Alexandros Mouzakitis. A survey on imitation learning techniques for end-to-end autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[LPL21] Hong-Liang Li, Yi-He Pang, and Bin Liu. Bioseq-blm: a platform for analyzing dna, rna and protein sequences based on biological language models. *Nucleic acids research*, 49(22):e129–e129, 2021.

[LSB22] Umberto Lupo, Damiano Sgarbossa, and Anne-Florence Bitbol. Protein language models trained on multiple sequence alignments learn phylogenetic relationships. *bioRxiv*, 2022.

[LSD15] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.

[LT23] Aryan Lall and Siddharth Tallur. Deep reinforcement learning-based pairwise dna sequence alignment method compatible with embedded edge devices. *Scientific Reports*, 13(1):2773, 2023.

[Lud23] Dylan Ludwig. The functions of consciousness in visual processing. *Neuroscience of Consciousness*, 2023(1):niac018, 2023.

[LW67] Godfrey N Lance and William Thomas Williams. A general theory of classificatory sorting strategies: 1. hierarchical systems. *The computer journal*, 9(4):373–380, 1967.

[LYY$^+$24] Jiajia Liu, Mengyuan Yang, Yankai Yu, Haixia Xu, Kang Li, and Xiaobo Zhou. Large language models in bioinformatics: applications and perspectives. *ArXiv*, 2024.

[LZX16] Bo Li, Tianlei Zhang, and Tian Xia. Vehicle detection from 3d lidar using fully convolutional network. *arXiv preprint arXiv:1608.07916*, 2016.

[MAP23] Taslim Murad, Sarwan Ali, and Murray Patterson. Exploring the potential of gans in biological sequence analysis. *Biology*, 12(6):854, 2023.

[Mat14] Mark P Mattson. Superior pattern processing is the essence of the evolved human brain. *Frontiers in neuroscience*, 8:265, 2014.

[MBD14] Ioan-Gabriel Mircea, Maria-Iuliana Bocicor, and Andra. Dîincu. On reinforcement learning based multiple sequence alignment. *Studia Universitatis Babes-Bolyai, Informatica*, 59(2), 2014.

[McK18] Katherine M McKinnon. Flow cytometry: an overview. *Current protocols in immunology*, 120(1):5–1, 2018.

[MCS17] David C Moffat, William Crombie, and Olga Shabalina. Some video games can increase the player's creativity. *International Journal of Game-Based Learning (IJGBL)*, 7(2):35–46, 2017.

[MGGC22] Josh Aaron Miller, Kutub Gandhi, Anna Gander, and Seth Cooper. A survey of citizen science gaming experiences. *Citizen Science: Theory and Practice*, 7(1):1–12, 2022.

[MHD$^+$18] Daniel McDonald, Embriette Hyde, Justine W Debelius, James T Morton, Antonio Gonzalez, Gail Ackermann, Alexander A Aksenov, Bahar Behsaz, Caitriona Brennan, Yingfeng Chen, et al. American gut: an open platform for citizen science microbiome research. *Msystems*, 3(3):e00031–18, 2018.

[MHK16] Benedikt Morschheuser, Juho Hamari, and Jonna Koivisto. Gamification in crowdsourcing: a review. In *2016 49th Hawaii international conference on system sciences (HICSS)*, pages 4375–4384. IEEE, 2016.

[MHKM17] Benedikt Morschheuser, Juho Hamari, Jonna Koivisto, and Alexander Maedche. Gamified crowdsourcing: Conceptualization, literature review, and future agenda. *International Journal of Human-Computer Studies*, 106:26–43, 2017.

[MIT19] MIT News. The brain's inferotemporal cortex plays a crucial role in object recognition, 2019.

[MJ03] Jonathan Moss and Colin G Johnson. An ant colony algorithm for multiple sequence alignment in bioinformatics. In *Artificial Neural Nets and Genetic Algorithms: Proceedings of the International Conference in Roanne, France, 2003*, pages 182–186. Springer, 2003.

[MKS$^+$13] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[MKS$^+$15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al.

Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

[MNG$^+$15] Siavash Mirarab, Nam Nguyen, Sheng Guo, Li-San Wang, Junhyong Kim, and Tandy Warnow. Pasta: ultra-large multiple sequence alignment for nucleotide and amino-acid sequences. *Journal of Computational Biology*, 22(5):377–386, 2015.

[MNW14] Siavash Mirarab, Nam Nguyen, and Tandy Warnow. Pasta: ultra-large multiple sequence alignment. In *Research in Computational Molecular Biology: 18th Annual International Conference, RECOMB 2014, Pittsburgh, PA, USA, April 2-5, 2014, Proceedings 18*, pages 177–191. Springer, 2014.

[MO96] James N MacGregor and Tom Ormerod. Human performance on the traveling salesman problem. *Perception & psychophysics*, 58(4):527–539, 1996.

[MSGC$^+$23] Renata Mutalova, Roman Sarrazin-Gendron, Eddie Cai, Gabriel Richard, Parham Ghasemloo Gheidari, Sébastien Caisse, Rob Knight, Mathieu Blanchette, Attila Szantner, and Jérôme Waldispühl. Playing the system: Can puzzle players teach us how to solve hard problems? In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2023.

[MW09] Winter Mason and Duncan J Watts. Financial incentives and the" performance of crowds". In *Proceedings of the ACM SIGKDD workshop on human computation*, pages 77–85, 2009.

[MWP$^+$05] Burkhard Morgenstern, Nadine Werner, Sonja J Prohaska, Rasmus Steinkamp, Isabelle Schneider, Amarendran R Subramanian, Peter F Stadler, and Jan Weyer-Menkhoff. Multiple sequence alignment with user-defined constraints at gobics. *Bioinformatics*, 21(7):1271–1273, 2005.

[NB10] Hieu V Nguyen and Li Bai. Cosine similarity metric learning for face verification. In *Asian conference on computer vision*, pages 709–720. Springer, 2010.

[Nee70] SB Needleman. Needleman-wunsch algorithm for sequence similarity searches. *J Mol Biol*, 48:443–453, 1970.

[NFS12] Mohammad Norouzi, David J Fleet, and Russ R Salakhutdinov. Hamming distance metric learning. In *Advances in neural information processing systems*, pages 1061–1069, 2012.

[NHH15] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1520–1528, 2015.

[NMKW15] Nam-phuong D Nguyen, Siavash Mirarab, Keerthana Kumar, and Tandy Warnow. Ultra-large alignments using phylogeny-aware profiles. *Genome biology*, 16(1):1–15, 2015.

[NW70] Saul B Needleman and Christian D Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of molecular biology*, 48(3):443–453, 1970.

[NXM$^+$21] F Noël, G Xexéo, E Mangeli, A Mothé, P Marques, J Kritz, F Blanchard, H Vermelho, and B de Paiva. Screener, an educational game for teaching the drug discovery and development process. *Brazilian Journal of Medical and Biological Research*, 54:e11786, 2021.

[ON20] Csaba Osvath and Erica Newport. The creative gamer: Creativity and creative problem solving in video games. 2020.

[Onl21] EVE Online. Project discovery - level up!, 2021. Accessed: 2023-06-23.

[OPN$^+$18] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, Jan Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018.

[PAD$^+$17] Vineet Pandey, Amnon Amir, Justine Debelius, Embriette R Hyde, Tomasz Kosciolek, Rob Knight, and Scott Klemmer. Gut instinct: Creating scientific theories with online learners.

In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pages 6825–6836, 2017.

[PCI10] Gabriele Paolacci, Jesse Chandler, and Panagiotis G Ipeirotis. Running experiments on amazon mechanical turk. *Judgment and Decision making*, 5(5):411–419, 2010.

[PCMY15] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015.

[PGV$^+$23] Andrea Poliani, Silvia Gnecchi, Giulia Villa, Debora Rosa, and Duilio F Manara. Gamification as an educational approach for oncological patients: a systematic scoping review. In *Healthcare*, volume 11, page 3116. MDPI, 2023.

[PHB$^+$08] Benedict Paten, Javier Herrero, Kathryn Beal, Stephen Fitzgerald, and Ewan Birney. Enredo and pecan: genome-wide mammalian consistency-based multiple alignment with paralogs. *Genome research*, 18(11):1814–1828, 2008.

[PHKK18] Marisa Ponti, Thomas Hillman, Christopher Kullenberg, and Dick Kasperowski. Getting it right or being top rank: Games in citizen science. *Citizen Science: Theory and Practice*, 3(1):1–1, 2018.

[PHW$^+$09] Saumyadipta Pyne, Xiaoxu Hu, Kui Wang, Erica Rossin, Ting-I Lin, Lisa M Maier, Clare Baecher-Allan, Geoffrey J McLachlan, Pablo Tamayo, David A Hafler, et al. Automated high-dimensional flow cytometric data analysis. *Proceedings of the National Academy of Sciences*, 106(21):8519–8524, 2009.

[PJBN16] Maziyar Baran Pouyan, Vasu Jindal, Javad Birjandtalab, and Mehrdad Nourani. Single and multi-subject clustering of flow cytometry data for cell-type identification and anomaly detection. *BMC Medical Genomics*, 9:99–110, 2016.

[PKY$^+$21] Vineet Pandey, Tushar Koul, Chen Yang, Daniel McDonald, Mad Price Ball, Bastian Greshake Tzovaras, Rob Knight, and

Scott Klemmer. Galileo: Citizen-led experimentation using a social computing system. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–14, 2021.

[PLV$^+$20] Johannes Pfau, Antonios Liapis, Georg Volkmar, Georgios N Yannakakis, and Rainer Malaka. Dungeons & replicants: automated game balancing via deep player behavior modeling. In *2020 IEEE Conference on Games (CoG)*, pages 431–438. IEEE, 2020.

[PM19] Johannes Pfau and Rainer Malaka. Can you rely on human computation? a large-scale analysis of disruptive behavior in games with a purpose. In *Extended Abstracts of the Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts*, pages 605–610, 2019.

[POLO23] Cyriel MA Pennartz, Matthijs N Oude Lohuis, and Umberto Olcese. How 'visual'is the visual cortex? the interactions between the visual cortex and other sensory, motivational and motor systems as enabling factors for visual perception. *Philosophical Transactions of the Royal Society B*, 378(1886):20220336, 2023.

[Pom91] Dean A Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural computation*, 3(1):88–97, 1991.

[PPS12] Aluísio Pinheiro, Hildete Prisco Pinheiro, and Pranab Kumar Sen. The use of hamming distance in bioinformatics. In *Handbook of Statistics*, volume 28, pages 129–162. Elsevier, 2012.

[QM90] Yingdong Qian and Robert F Murphy. Automated identification of subpopulations in flow cytometric list mode data using cluster analysis. *Cytometry*, 11(5):583–589, 1990.

[RBG$^+$09] M Jordan Raddick, Georgia Bracey, Pamela L Gay, Chris J Lintott, Phil Murray, Kevin Schawinski, Alexander S Szalay, and Jan Vandenberg. Galaxy zoo: Exploring the motivations of citizen science volunteers. *arXiv preprint arXiv:0909.2925*, 2009.

[RCZC19] Alma Rodríguez, Erik Cuevas, Daniel Zaldivar, and Luis Castañeda. Clustering with biological visual models. *Physica A: Statistical Mechanics and its Applications*, 528:121505, 2019.

[RGB11] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.

[RK03] Thomas Kiel Rasmussen and Thiemo Krink. Improved hidden markov model training for multiple sequence alignment by a particle swarm optimization—evolutionary algorithm hybrid. *Biosystems*, 72(1-2):5–17, 2003.

[RLZ17] Yongming Rao, Jiwen Lu, and Jie Zhou. Attention-aware deep reinforcement learning for video face recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 3931–3940, 2017.

[RMS+20] Roshan Rao, Joshua Meier, Tom Sercu, Sergey Ovchinnikov, and Alexander Rives. Transformer protein language models are unsupervised structure learners. *Biorxiv*, 2020.

[RSB18] Ramchalam Kinattinkara Ramakrishnan, Jaspal Singh, and Mathieu Blanchette. Rlalign: a reinforcement learning approach for multiple sequence alignment. In *2018 IEEE 18th International Conference on Bioinformatics and Bioengineering (BIBE)*, pages 61–66. IEEE, 2018.

[RSdW+18] Tabish Rashid, Mikayel Samvelyan, Cedric de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, pages 4295–4304. PMLR, 2018.

[RSS$^+$15] J Redding, Jacob Schreiver, C Shrum, A Lauf, and R Yampolskiy. Solving np-hard number matrix games with wisdom of artificial crowds. In *2015 Computer Games: AI, Animation, Mobile, Multimedia, Educational and Serious Games (CGAMES)*, pages 38–43. IEEE, 2015.

[RYHH10] Cyrus Rashtchian, Peter Young, Micah Hodosh, and Julia Hockenmaier. Collecting image annotations using amazon's mechanical turk. In *Proceedings of the NAACL HLT 2010 workshop on creating speech and language data with Amazon's Mechanical Turk*, pages 139–147, 2010.

[Sam10] Claude Sammut. Behavioral cloning., 2010.

[SB18] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction.* MIT press, 2018.

[SC21] Yong-Joon Song and Dong-Ho Cho. Local alignment of dna sequence based on deep reinforcement learning. *IEEE Open Journal of Engineering in Medicine and Biology*, 2:170–178, 2021.

[Sch24] Karen Schrier. Leadership through citizen science games. *Journal of Leadership Studies*, 17(4):32–37, 2024.

[SF08] Alexander Sorokin and David Forsyth. Utility data annotation with amazon mechanical turk. In *2008 IEEE computer society conference on computer vision and pattern recognition workshops*, pages 1–8. IEEE, 2008.

[SG15] Seungmoon Song and Hartmut Geyer. A neural circuitry that emphasizes spinal feedback generates diverse behaviours of human locomotion. *The Journal of physiology*, 593(16):3493–3511, 2015.

[SGGGB$^+$24] Roman Sarrazin-Gendron, Parham Ghasemloo Gheidari, Alexander Butyaev, Timothy Keding, Eddie Cai, Jiayue Zheng, Renata Mutalova, Julien Mounthanyvong, Yuxue Zhu, Elena Nazarova, et al. Improving microbial phylogeny with citizen science within a mass-market video game. *Nature Biotechnology*, pages 1–9, 2024.

[She20] Alex Sherstinsky. Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena*, 404:132306, 2020.

[SHM$^+$16] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

[SJS$^+$21] Yong-Joon Song, Dong Jin Ji, Hyein Seo, Gyu Bum Han, and Dong-Ho Cho. Pairwise heuristic sequence alignment algorithm based on deep reinforcement learning. *IEEE Open Journal of Engineering in Medicine and Biology*, 2:36–43, 2021.

[SMC73] David Sankoff, Cristiane Morel, and Robert J Cedergren. Evolution of 5s rna and the non-randomness of base replacement. *Nature New Biology*, 245(147):232–234, 1973.

[SS13] Rafik A Salama and Dov J Stekel. A non-independent energy-based multiple sequence alignment improves prediction of transcription factor binding sites. *Bioinformatics*, 29(21):2699–2704, 2013.

[Sur05] James Surowiecki. *The wisdom of crowds*. Anchor, 2005.

[SW$^+$81] Temple F Smith, Michael S Waterman, et al. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197, 1981.

[Sza16] Attila Szantner. Massively multiplayer online science. In *Levelling Up: The Cultural Impact of Contemporary Videogames*, pages 103–110. Brill, 2016.

[Szi12] István Szita. Reinforcement learning in games. In *Reinforcement learning*, pages 539–577. Springer, 2012.

[THG94] Julie D Thompson, Desmond G Higgins, and Toby J Gibson. Clustal w: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice. *Nucleic acids research*, 22(22):4673–4680, 1994.

[TLRSH17] Ramine Tinati, Markus Luczak-Roesch, Elena Simperl, and Wendy Hall. An investigation of player motivations in eyewire, a gamified citizen science project. *Computers in Human Behavior*, 73:527–540, 2017.

[TN20] Rakesh Trivedi and Hampapathalu Adimurthy Nagarajaram. Substitution scoring matrices for proteins-an overview. *Protein Science*, 29(11):2150–2163, 2020.

[Tut23] Omics Tutorials. Generative adversarial networks (gans) in science and biology: Advancing image manipulation and data augmentation, 2023. Accessed: 2024-10-17.

[UOPS23] Thomas A Ujas, Veronica Obregon-Perko, and Ann M Stowe. A guide on analyzing flow cytometry data using clustering methods and nonlinear dimensionality reduction (tsne or umap). In *Neural Repair: Methods and Protocols*, pages 231–249. Springer, 2023.

[VA06] Luis Von Ahn. Games with a purpose. *Computer*, 39(6):92–94, 2006.

[vAD04a] Luis von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, page 319–326, New York, NY, USA, 2004. Association for Computing Machinery.

[VAD04b] Luis Von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326, 2004.

[VBC$^+$19] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

[VSP$^+$17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia

Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

[VW94] Martin Vingron and Michael S Waterman. Sequence alignment and penalty choice: Review of concepts, case studies and implications. *Journal of molecular biology*, 235(1):1–12, 1994.

[Way21] Waymo. Waymo safety report. Technical report, 2021. Accessed: 2024-10-17.

[Waz18] Kerri Wazny. Applications of crowdsourcing in health: an overview. *Journal of global health*, 8(1), 2018.

[WDJ04] Guoli Wang and Roland L Dunbrack Jr. Scoring profile-to-profile sequence alignments. *Protein Science*, 13(6):1612–1626, 2004.

[WF93] John C Wootton and Scott Federhen. Statistics of local complexity in amino acid sequences and sequence databases. *Computers & chemistry*, 17(2):149–163, 1993.

[WGD$^+$17] Kunfeng Wang, Chao Gou, Yanjie Duan, Yilun Lin, Xinhu Zheng, and Fei-Yue Wang. Generative adversarial networks: introduction and outlook. *IEEE/CAA Journal of Automatica Sinica*, 4(4):588–598, 2017.

[Wik21] EVE University Wiki. Project discovery: Flow cytometry, 2021. Accessed: 2023-06-23.

[Wil02] Wiley Online Library. The significance of color in object recognition, 2002.

[WJ94] Lusheng Wang and Tao Jiang. On the complexity of multiple sequence alignment. *Journal of computational biology*, 1(4):337–348, 1994.

[WPDX$^+$22] Chen Wang, Claudia Pérez-D'Arpino, Danfei Xu, Li Fei-Fei, Karen Liu, and Silvio Savarese. Co-gail: Learning diverse strategies for human-robot collaboration. In *Conference on Robot Learning*, pages 1279–1290. PMLR, 2022.

[WSH$^+$16] Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando Freitas. Dueling network architectures for deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1995–2003, 2016.

[WSK$^+$20] Jérôme Waldispühl, Attila Szantner, Rob Knight, Sébastien Caisse, and Randy Pitchford. Leveling up citizen science. *Nature Biotechnology*, 38(10):1124–1126, 2020.

[WSQ$^+$21] Xiangjun Wang, Junxiao Song, Penghui Qi, Peng Peng, Zhenkun Tang, Wei Zhang, Weimin Li, Xiongjun Pi, Jujie He, Chao Gao, et al. Scc: An efficient deep reinforcement learning agent mastering the game of starcraft ii. In *International conference on machine learning*, pages 10905–10915. PMLR, 2021.

[WVO12] Marco A Wiering and Martijn Van Otterlo. Reinforcement learning. *Adaptation, learning, and optimization*, 12(3):729, 2012.

[XC09] Fasheng Xu and Yuehui Chen. A method for multiple sequence alignment based on particle swarm optimization. In *Emerging Intelligent Computing Technology and Applications. With Aspects of Artificial Intelligence: 5th International Conference on Intelligent Computing, ICIC 2009 Ulsan, South Korea, September 16-19, 2009 Proceedings 5*, pages 965–973. Springer, 2009.

[XZL15] Peipei Xia, Li Zhang, and Fanzhang Li. Learning similarity with cosine similarity ensemble. *Information Sciences*, 307:39–52, 2015.

[YA23] Malik Yousef and Jens Allmer. Deep learning in bioinformatics. *Turkish Journal of Biology*, 47(6):366–382, 2023.

[YEB11] Roman V Yampolskiy and Ahmed El-Barkouky. Wisdom of artificial crowds algorithm for solving np-hard problems. *International Journal of Bio-inspired computation*, 3(6):358–369, 2011.

[YSLD12] Sheng Kung Michael Yi, Mark Steyvers, Michael D Lee, and Matthew J Dry. The wisdom of the crowd in combinatorial problems. *Cognitive science*, 36(3):452–470, 2012.

[ZB24] Xinyu Zhang and Abdeslam Boularias. One-shot imitation learning with invariance matching for robotic manipulation. *arXiv preprint arXiv:2405.13178*, 2024.

[ZCT$^+$23] Zunming Zhang, Xinyu Chen, Rui Tang, Yuxuan Zhu, Han Guo, Yunjia Qu, Pengtao Xie, Ian Y Lian, Yingxiao Wang, and Yu-Hwa Lo. Interpretable unsupervised learning enables accurate clustering with high-throughput imaging flow cytometry. *Scientific Reports*, 13(1):20533, 2023.

[ZMJ$^+$18] Tianhao Zhang, Zoe McCarthy, Owen Jow, Dennis Lee, Xi Chen, Ken Goldberg, and Pieter Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 5628–5635. Ieee, 2018.

[ZRL96] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. Birch: An efficient data clustering method for very large databases. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pages 103–114, 1996.

[ZSWM00] Zheng Zhang, Scott Schwartz, Lukas Wagner, and Webb Miller. A greedy algorithm for aligning dna sequences. *Journal of Computational biology*, 7(1-2):203–214, 2000.

[ZWL$^+$20] Ruohan Zhang, Calen Walshe, Zhuode Liu, Lin Guan, Karl Muller, Jake Whritner, Luxin Zhang, Mary Hayhoe, and Dana Ballard. Atari-head: Atari human eye-tracking and demonstration dataset. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 6811–6820, 2020.

[ZZL$^+$22] Yongqing Zhang, Qiang Zhang, Yuhang Liu, Meng Lin, and Chunli Ding. Multiple sequence alignment based on deep q network with negative feedback policy. *Computational Biology and Chemistry*, 101:107780, 2022.