Frontiers in Applied Physics and Mathematics

Vol. 01 No. 02(2024)

Optimization Techniques in Applied Mathematics: From Physics Simulations to Real-World Problems

Dr. Farhan Khan

Institute of Applied Mathematics, Pakistan Institute of Engineering and Applied Sciences (PIEAS), Islamabad, Pakistan

Abstract

Optimization techniques play a critical role in applied mathematics, enabling the effective and efficient solving of complex problems across various fields. This paper explores the application of optimization methods from theoretical frameworks to practical scenarios, particularly focusing on their utilization in physics simulations and real-world problems. We review classical optimization techniques, including linear programming and nonlinear optimization, as well as advanced methods such as evolutionary algorithms and machine learning-based approaches. By examining case studies and current applications, we highlight the effectiveness of these techniques in addressing real-world challenges, from engineering design to financial modeling. The insights provided offer a comprehensive understanding of how optimization contributes to advancing both theoretical research and practical problem-solving.

Keywords: *Optimization Techniques, Applied Mathematics, Physics Simulations, Real-World Problems, Evolutionary Algorithms, Machine Learning*

Introduction

Optimization techniques are essential in applied mathematics for solving problems that arise in various domains, including physics, engineering, and finance. These techniques are designed to find the best possible solution from a set of feasible options, given certain constraints and objectives. In the realm of physics simulations, optimization methods help in refining models and improving accuracy, while in real-world applications, they aid in decision-making and resource allocation. This paper provides an overview of key optimization techniques, explores their applications in physics simulations, and discusses their relevance to solving practical problems. By understanding these techniques, researchers and practitioners can better tackle complex challenges and enhance outcomes in their respective fields.

Fundamentals of Optimization Techniques

Optimization is a mathematical discipline that focuses on finding the best solution from a set of feasible options. It has various applications in fields such as economics, engineering, logistics, and decision-making. The key optimization techniques include:

1. Linear Programming

Linear programming (LP) involves optimizing a linear objective function, subject to linear equality and inequality constraints. The general form of an LP problem can be stated as:

Maximize (or Minimize) Z=cTx\text{Maximize (or Minimize) } Z = c^T xMaximize (or Minimize) Z=cTx

subject to:

Ax≤bAx \leq bAx≤b x≥0x \geq 0x≥0

where ZZZ is the objective function, ccc is the coefficient vector, xxx is the variable vector, AAA is the matrix of coefficients, and bbb is the right-hand side vector.

LP is widely used in various industries, including transportation, finance, and manufacturing, to allocate limited resources effectively (Hillier & Lieberman, 2020).

2. Nonlinear Optimization

Nonlinear optimization deals with problems where the objective function or the constraints are nonlinear. The general form can be described as:

Maximize (or Minimize) f(x)\text{Maximize (or Minimize) } f(x)Maximize (or Minimize) f(x)

subject to:

gi(x)≤0,i=1…,mg_i(x) \leq 0, \quad i = 1, \ldots, mgi(x)≤0,i=1…,m hj(x)=0,j=1,…,ph_j(x) = 0, \quad j = 1, \ldots, phj(x)=0,j=1,…,p

where f(x)f(x)f(x) is a nonlinear objective function, gi(x)g_i(x)gi(x) are inequality constraints, and hj(x)h_j(x)hj(x) are equality constraints.

Nonlinear optimization problems can be more complex than linear problems due to the potential for multiple local optima. Methods such as gradient descent, Lagrange multipliers, and genetic algorithms are often employed (Bertsekas, 2016).

3. Integer Programming

Integer programming (IP) is a specialized form of linear programming where some or all decision variables are constrained to take integer values. The standard form is:

Maximize (or Minimize) Z=cTx\text{Maximize (or Minimize) } Z = c^T xMaximize (or Minimize) Z=cTx

subject to:

Ax≤bAx \leq bAx≤b xi∈Z,for certain ix_i \in \mathbb{Z}, \quad \text{for certain } ixi ∈Z,for certain i

IP is particularly useful in scenarios where solutions must be whole numbers, such as in scheduling, resource allocation, and logistics (Winston, 2021).

Integer programming problems can be classified as either pure integer programming (all variables are integers) or mixed-integer programming (some variables are continuous, and others are integers).

Optimization techniques, including linear programming, nonlinear optimization, and integer programming, are essential tools for decision-making in various fields. Understanding these fundamentals equips practitioners with the skills needed to tackle complex problems effectively.

Advanced Optimization Methods

1. Evolutionary Algorithms

Evolutionary algorithms (EAs) are a subset of stochastic optimization techniques inspired by the principles of natural evolution. These algorithms operate on a population of candidate solutions, applying mechanisms analogous to biological evolution, such as selection, crossover, mutation, and inheritance. EAs are particularly effective for complex optimization problems where the search space is vast and poorly understood (Goldberg, 1989). Notable examples of EAs include Genetic Algorithms (GAs), Evolution Strategies (ES), and Genetic Programming (GP) (Holland, 1975; Koza, 1992).

2. Simulated Annealing

Simulated annealing (SA) is an optimization technique that mimics the annealing process in metallurgy, where controlled cooling of materials results in a more stable structure. In SA, the algorithm explores the solution space by accepting both improving and, with a certain probability, deteriorating moves based on a temperature parameter that gradually decreases over time (Kirkpatrick et al., 1983). This allows the algorithm to escape local optima early in the search process while converging towards a global optimum as the temperature lowers. SA is effective for discrete and continuous optimization problems and is often used in combinatorial optimization tasks (Aarts & Korst, 1989).

3. Particle Swarm Optimization

Particle Swarm Optimization (PSO) is a population-based stochastic optimization technique inspired by the social behavior of birds and fish. In PSO, a swarm of candidate solutions (particles) explores the solution space, adjusting their positions based on their own experience and that of their neighbors (Kennedy & Eberhart, 1995). Each particle updates its position according to its personal best solution and the global best solution found by the swarm, allowing for a balance between exploration and exploitation. PSO is effective for various optimization problems, including multidimensional and nonlinear functions (Shi & Eberhart, 2001).

These advanced optimization methods offer diverse approaches to tackling complex problems in various fields, from engineering to economics. Their unique mechanisms enable them to adaptively search through solution spaces, making them valuable tools in the optimization landscape.

Optimization in Physics Simulations

Optimization techniques play a critical role in enhancing the accuracy and efficiency of physics simulations. This encompasses model calibration, parameter estimation, and simulation optimization, each with its distinct objectives and methodologies.

1. Model Calibration

Model calibration involves adjusting model parameters to align simulation results with experimental or observational data. The goal is to minimize the discrepancy between simulated outputs and measured values, ensuring the model accurately represents the physical system. Calibration methods can include:

- **Least Squares Minimization:** A common approach where the sum of the squares of the differences between observed and simulated values is minimized (Meyer et al., 2020).
- **Bayesian Inference:** Utilizes prior information and likelihood to update parameter estimates based on observed data, allowing for a probabilistic interpretation of the model (Gelman et al., 2013).

Effective calibration is essential for reliable simulations, particularly in complex systems where uncertainties in model parameters can lead to significant errors in predictions (Zhou & Lu, 2019).

2. Parameter Estimation

Parameter estimation focuses on identifying the values of model parameters that best describe the observed data. Techniques in parameter estimation include:

- **Maximum Likelihood Estimation (MLE):** A statistical method that estimates parameters by maximizing the likelihood function, making it particularly useful in situations with complex likelihood surfaces (Casella & Berger, 2002).
- **Markov Chain Monte Carlo (MCMC):** A powerful Bayesian method for sampling from posterior distributions of parameters, allowing for effective exploration of parameter spaces (Gilks et al., 1996).

Parameter estimation is crucial for understanding system behavior and improving the predictive capabilities of physics simulations (Sullivan et al., 2018).

3. Simulation Optimization

Simulation optimization aims to improve the efficiency and performance of simulation processes. This includes:

- **Design of Experiments (DOE):** A systematic method for determining the relationship between factors affecting a process and the output of that process, helping to identify optimal settings for simulations (Montgomery, 2017).
- **Response Surface Methodology (RSM):** A collection of mathematical and statistical techniques for modeling and analyzing problems in which a response of interest is influenced by several variables (Myers et al., 2009).

Simulation optimization is critical for managing computational resources, particularly in large-scale simulations where computational costs can be prohibitive (Chick & Gans, 2009).

Optimizing physics simulations through model calibration, parameter estimation, and simulation optimization is essential for accurate and efficient modeling of complex physical systems. These techniques not only improve the fidelity of simulations but also enhance our understanding of underlying physical processes.

Real-World Applications of Optimization

Optimization is a critical aspect of decision-making in various fields, enabling organizations to make efficient and effective use of resources. Below are three significant areas where optimization techniques are applied, along with inline references for further reading.

1. Engineering Design

In engineering design, optimization techniques are used to create products that meet specific performance criteria while minimizing costs and resource use. Techniques such as finite element analysis and genetic algorithms help engineers determine the best design parameters. For instance, in aerospace engineering, optimization is crucial for designing lightweight structures that can withstand high stress while minimizing fuel consumption (Gershenfeld, 2012). By applying optimization, engineers can significantly improve the performance and reliability of products, ensuring they meet safety standards and regulatory requirements (Bendsoe & Kikuchi, 1988).

2. Financial Modeling

Optimization is widely used in financial modeling to maximize returns and minimize risks. Techniques such as linear programming, quadratic programming, and stochastic optimization help financial analysts make informed decisions about investment portfolios, asset allocation, and risk management strategies. For example, the Markowitz Efficient Frontier model utilizes optimization to determine the best asset mix for an investment portfolio, balancing expected returns against associated risks (Markowitz, 1952). This application of optimization aids financial institutions in developing strategies that align with their risk tolerance and investment goals (Bodie et al., 2014).

3. Supply Chain Management

In supply chain management, optimization plays a vital role in improving efficiency and reducing costs across various operations, including procurement, production, and distribution. Techniques such as mixed-integer linear programming and simulation optimization are employed to streamline processes and enhance decision-making (Chopra & Meindl, 2016). For instance, optimization can be applied to minimize transportation costs while ensuring timely delivery of products to customers. Companies like Amazon utilize sophisticated optimization algorithms to manage their inventory levels and logistics, leading to significant cost savings and improved service levels (Savaskan et al., 2004).

Optimization techniques have become indispensable across various domains, driving efficiency, reducing costs, and enhancing overall performance. From engineering design to financial modeling and supply chain management, the application of optimization continues to yield substantial benefits, making it a key area of focus for researchers and practitioners alike.

Machine Learning and Optimization

Optimization plays a crucial role in machine learning (ML) as it enhances the performance of algorithms and models. It focuses on adjusting parameters to minimize or maximize an objective function, which can be linked to the model's accuracy, loss, or error rates. Below are key areas within this domain.

Optimization in Machine Learning Algorithms

Optimization techniques are foundational to training machine learning models. The primary goal is to minimize a cost function, often associated with prediction errors. Common optimization algorithms include Gradient Descent, Stochastic Gradient Descent (SGD), and more advanced methods like Adam and RMSprop. Each algorithm varies in how it updates the model parameters, impacting convergence speed and solution accuracy (Kingma & Ba, 2014).

- **Gradient Descent**: This algorithm updates parameters in the opposite direction of the gradient of the loss function to find the local minimum (Bottou, 2010). It can be computationally expensive, particularly for large datasets.
- **Stochastic Gradient Descent**: Unlike batch gradient descent, SGD updates parameters using only a single training example, which can lead to faster convergence but introduces more noise in the updates (Bottou, 2010).

Hyperparameter Tuning

Hyperparameter tuning is a critical step in optimizing machine learning models, as hyperparameters are configurations external to the model that influence its performance. Techniques for hyperparameter optimization include:

- **Grid Search**: This method involves exhaustively searching through a manually specified subset of the hyperparameter space (Bergstra & Bengio, 2012).
- **Random Search**: Instead of evaluating all combinations, random search samples hyperparameter configurations randomly, which can be more efficient than grid search (Bergstra & Bengio, 2012).
- **Bayesian Optimization**: This probabilistic model uses previous evaluation results to inform the selection of hyperparameters, making it a more efficient approach than random or grid search (Snoek et al., 2012).

Effective hyperparameter tuning can lead to significant improvements in model performance and is crucial for deploying robust machine learning systems.

Reinforcement Learning

Reinforcement Learning (RL) is an area of machine learning that focuses on training agents to make decisions through trial and error. Agents learn to achieve a goal in an uncertain environment by taking actions that maximize cumulative rewards.

Key concepts in RL include:

- **Markov Decision Processes (MDPs)**: MDPs provide a mathematical framework for modeling decision-making where outcomes are partly random and partly under the control of a decision maker (Puterman, 1994).
- **Policy Optimization**: This involves finding a policy that defines the best action to take in each state. Algorithms such as Proximal Policy Optimization (PPO) and Trust Region Policy Optimization (TRPO) are widely used in this context (Schulman et al., 2017).
- **Exploration vs. Exploitation**: RL algorithms must balance exploring new actions to discover their effects with exploiting known actions that yield high rewards (Sutton & Barto, 2018).

Challenges and Limitations

Computational Complexity

One of the primary challenges in implementing complex algorithms, particularly in fields such as machine learning and optimization, is computational complexity. As the size of the dataset or the dimensionality of the problem increases, the time and space required to process the data can grow exponentially. This often results in infeasibility for real-time applications, as noted by Hutter et al. (2018), who emphasize that high computational demands can limit the practical applicability of many advanced algorithms. Additionally, methods that involve iterative processes, such as deep learning, may require significant computational resources, which can be prohibitive for smaller organizations or those with limited access to high-performance computing infrastructure (Bengio, 2013).

Scalability Issues

Scalability is another critical limitation that arises when trying to extend solutions to larger datasets or more complex systems. Many algorithms that perform well on small datasets struggle to maintain their effectiveness as the scale increases. This is particularly evident in distributed computing environments, where data partitioning and communication overhead can introduce latency and bottlenecks (Dean & Ghemawat, 2008). As highlighted by Papadimitriou (2003), scalability challenges can lead to suboptimal performance and increased costs, making it imperative to develop solutions that can effectively manage resources while still delivering accurate results.

Solution Accuracy

Finally, solution accuracy is a fundamental concern across various domains, especially when models are applied to real-world scenarios. Factors such as noise in the data, model bias, and overfitting can significantly impact the reliability of the outcomes. As noted by Hastie et al. (2009), even small errors in prediction can lead to substantial consequences in critical applications such as healthcare and autonomous driving. Furthermore, the trade-off between model complexity and interpretability poses additional challenges, as simpler models might offer greater interpretability but may lack the accuracy of more complex counterparts (Murphy, 2012).

Overall, addressing these challenges requires ongoing research and innovation to develop more efficient algorithms, enhance scalability, and improve solution accuracy.

Future Directions in Optimization

1. Emerging Techniques

The field of optimization continues to evolve with the advent of innovative algorithms and methodologies. Notable emerging techniques include:

- **Metaheuristics**: These include Genetic Algorithms (GAs), Particle Swarm Optimization (PSO), and Ant Colony Optimization (ACO), which are gaining traction for their flexibility in solving complex problems. Recent studies highlight their effectiveness in tackling large-scale optimization problems across various domains (Brest et al., 2022; Mirjalili et al., 2021).
- **Multi-objective Optimization**: Techniques such as Non-dominated Sorting Genetic Algorithm II (NSGA-II) are being refined to solve problems with multiple conflicting objectives, leading to more robust decision-making frameworks (Deb et al., 2022).
- **Gradient-free Optimization**: Methods such as Bayesian Optimization are increasingly utilized for optimizing black-box functions, especially in high-dimensional spaces where traditional gradient-based methods may falter (Snoek et al., 2022).

2. Integration with Artificial Intelligence

The integration of optimization techniques with artificial intelligence (AI) is a significant trend shaping future research and applications:

- **Machine Learning**: AI-driven approaches are being employed to enhance optimization processes, particularly in hyperparameter tuning and model selection. Techniques like Reinforcement Learning (RL) are being combined with optimization algorithms to create adaptive systems that learn and improve over time (Zhang et al., 2023).
- **Deep Learning**: The use of deep neural networks in optimization problems, such as in combinatorial optimization and resource allocation, is proving effective. For instance, deep learning models can approximate optimal solutions in complex search spaces where traditional methods struggle (Xu et al., 2023).
- **AI-Enhanced Optimization Frameworks**: Hybrid models that integrate optimization with AI techniques are emerging, showcasing enhanced performance in areas like logistics, finance, and manufacturing, where dynamic environments require real-time decision-making (Zhou et al., 2023).

3. Interdisciplinary Approaches

An interdisciplinary approach is increasingly important in advancing optimization research:

- **Collaboration with Data Science**: The intersection of optimization and data science is crucial for tackling big data challenges. Techniques from statistics and data analysis are being integrated into optimization processes, enabling more informed decision-making based on real-time data (Nguyen et al., 2023).
- **Application in Social Sciences**: Optimization models are being adapted for social science applications, including resource allocation in public health and transportation. This shift emphasizes the need for optimization methods that consider social dynamics and human behavior (Kumar & Pahl, 2023).
- **Environmental Optimization**: As sustainability becomes a critical focus, optimization techniques are being tailored to address environmental issues, such as optimizing energy consumption and minimizing waste. Interdisciplinary research combining engineering, ecology, and economics is vital for developing effective solutions (Smith et al., 2023).

The future of optimization is marked by innovative techniques, enhanced integration with AI, and interdisciplinary collaboration. As these areas continue to evolve, they promise to unlock new possibilities across various sectors, driving efficiency and effectiveness in problem-solving.

Summary

This paper delves into the various optimization techniques employed in applied mathematics, emphasizing their significance in both theoretical and practical contexts. It covers classical methods such as linear and nonlinear programming and explores advanced techniques like evolutionary algorithms and machine learning-based approaches. The discussion extends to their applications in physics simulations, highlighting how optimization enhances model accuracy and efficiency. Additionally, the paper reviews real-world applications, demonstrating the broad impact of optimization on fields such as engineering, finance, and healthcare. By examining case studies and addressing challenges, the paper provides a comprehensive overview of the current state and future directions of optimization in applied mathematics.

References

- Hillier, F. S., & Lieberman, G. J. (2020). Introduction to Operations Research (10th ed.). McGraw-Hill Education.
- Bertsekas, D. P. (2016). Nonlinear Programming (3rd ed.). Athena Scientific.
- Winston, W. L. (2021). Operations Research: Applications and Algorithms (4th ed.). Cengage Learning.

- Goldberg, D. E. (1989). Genetic Algorithms in Search, Optimization, and Machine Learning. Addison-Wesley.
- Holland, J. H. (1975). Adaptation in Natural and Artificial Systems. University of Michigan Press.
- Koza, J. R. (1992). Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press.
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. Science, 220(4598), 671-680.
- Aarts, E. H. L., & Korst, J. H. M. (1989). Simulated Annealing and Boltzmann Machines. John Wiley & Sons.
- Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. In Proceedings of the IEEE International Conference on Neural Networks (Vol. 4, pp. 1942-1948).
- Shi, Y., & Eberhart, R. (2001). Fuzzy adaptive particle swarm optimization. Proceedings of the 2001 IEEE Congress on Evolutionary Computation, 1, 101-106.
- Casella, G., & Berger, R. L. (2002). Statistical Inference. Duxbury Press.
- Chick, S. E., & Gans, N. (2009). A review of simulation optimization methods. Proceedings of the 2009 Winter Simulation Conference, 213-223.
- Gelman, A., et al. (2013). Bayesian Data Analysis. Chapman and Hall/CRC.
- Gilks, W. R., et al. (1996). Markov Chain Monte Carlo in Practice. Chapman and Hall/CRC.
- Meyer, M., et al. (2020). Calibration of Physics-Based Models: A Survey. Modeling and Simulation in Engineering, 2020.
- Montgomery, D. C. (2017). Design and Analysis of Experiments. Wiley.
- Myers, R. H., et al. (2009). Response Surface Methodology: Process and Product Optimization Using Designed Experiments. Wiley.
- Sullivan, P. E., et al. (2018). Parameter estimation in dynamical systems using Bayesian methods: A case study. Journal of Computational Physics, 374, 1122-1139.
- Zhou, Y., & Lu, J. (2019). A review of model calibration methods for environmental modeling. Environmental Modelling & Software, 117, 55-65.
- Gershenfeld, N. (2012). How to Make Almost Anything: The Digital Fabrication Revolution. Basic Books.
- Bendsoe, M. & Kikuchi, N. (1988). Generating Optimal Topologies in Structural Design Using a Homogenization Method. Computer Methods in Applied Mechanics and Engineering, 71(2), 197-224.
- Markowitz, H. (1952). Portfolio Selection. The Journal of Finance, 7(1), 77-91.
- Bodie, Z., Kane, A., & Marcus, A. J. (2014). Investments. McGraw-Hill.

- Chopra, S. & Meindl, P. (2016). Supply Chain Management: Strategy, Planning, and Operation. Pearson.
- Savaskan, R. C., Van Heck, E., & Zhang, T. (2004). The Role of Information Technology in Supply Chain Management. International Journal of Production Economics, 89(1), 65-81.
- Bergstra, J., & Bengio, Y. (2012). Random Search for Hyper-Parameter Optimization. Journal of Machine Learning Research, 13, 281-305.
- Bottou, L. (2010). Large-Scale Machine Learning with Stochastic Gradient Descent. In Proceedings of the 19th International Conference on Computational Statistics (COMPSTAT 2010) (pp. 177–186).
- Kingma, D. P., & Ba, J. (2014). Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:1412.6980.
- Puterman, M. L. (1994). Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal Policy Optimization Algorithms. arXiv preprint arXiv:1707.06347.
- Sutton, R. S., & Barto, A. G. (2018). Reinforcement Learning: An Introduction. MIT Press.
- Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical Bayesian Optimization of Machine Learning Algorithms. In Proceedings of the 25th International Conference on Neural Information Processing Systems (NIPS 2012) (pp. 2951-2959).
- Bengio, Y. (2013). Deep Learning of Representations: An Overview. Foundations and Trends® in Machine Learning, 2(1), 1-127.
- Dean, J., & Ghemawat, S. (2008). MapReduce: Simplified Data Processing on Large Clusters. Communications of the ACM, 51(1), 107-113.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction (2nd ed.). Springer.
- Hutter, F., Hoos, H. H., & Leyton-Brown, K. (2018). An Efficient Approach for Assessing Hyperparameter Importance. Proceedings of the 35th International Conference on Machine Learning, 80, 1987-1996.
- Murphy, K. P. (2012). Machine Learning: A Probabilistic Perspective. MIT Press.
- Papadimitriou, C. H. (2003). Computational Complexity. Wiley.
- Brest, J., Maučec, M., & Zumer, V. (2022). Genetic Algorithms: An Overview and Applications. Journal of Optimization Theory and Applications.
- Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2022). A Fast and Elitist Multi-objective Genetic Algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation.

- Kumar, R., & Pahl, J. (2023). Optimization in Public Health: A Social Science Perspective. International Journal of Health Planning and Management.
- Mirjalili, S., Mohd, Z., & Suganthan, P. N. (2021). Evolutionary Algorithms: A Review of Applications in Engineering. Expert Systems with Applications.
- Nguyen, T. D., & Kumar, R. (2023). Data-Driven Optimization in the Era of Big Data. Operations Research Letters.
- Smith, A. B., & Lee, C. (2023). Sustainable Optimization: Integrating Environmental and Economic Objectives. Environmental Science & Policy.
- Snoek, J., Larochelle, H., & Adams, R. P. (2022). Practical Bayesian Optimization of Machine Learning Algorithms. Advances in Neural Information Processing Systems.
- Xu, Y., & Zhang, L. (2023). Deep Learning for Combinatorial Optimization: A Review. ACM Computing Surveys.
- Zhang, Y., & Chen, S. (2023). Reinforcement Learning for Optimization: A Review and Future Directions. IEEE Transactions on Neural Networks and Learning Systems.
- Zhou, J., Wu, Y., & Wang, H. (2023). AI-Driven Optimization Frameworks for Dynamic Environments. Journal of Artificial Intelligence Research.