Cost-based policies for the narrative observation of unpredictable events

Tiernan Lindauer
Department of Computer Science
Texas A&M University
College Station, United States
tiernanlind@tamu.edu

Dylan A. Shell
Department of Computer Science
Texas A&M University
College Station, United States
dshell@tamu.edu

***Abstract*— *Passive observation of a series of unpredictable but consistent events unfolding in real-time presents a unique challenge for various applications, such as home robotics, robot videography, sensor platforms, and other real-time systems. These systems require important decision-making capabilities to effectively capture data about significant events as they occur. The captured data is then composed into a narrative based on the temporal sequence of the events, allowing for the crafting of the narrative according to a pre-determined specification. This paper builds upon previous research that models a stochastic process using a Markov model by containerizing the source code and introducing a cost-based approach. Unlike the original approach, which minimizes the expected number of steps, the proposed method focuses on minimizing the cost between nodes. This is particularly relevant for autonomous systems, where minimizing the distance traveled can help reduce wear on components and improve overall efficiency. The paper presents experiments that demonstrate the feasibility of calculating an optimal policy for a system using this cost-based modeling approach. The results show that it is indeed possible to compute an optimal policy that takes into account the specific costs associated with transitioning between different states in the system. This research has important implications for the development of autonomous systems that need to make real-time decisions while optimizing for various cost factors, ultimately leading to more efficient and effective operation in dynamic environments.*

Keywords—Planning, cost, passive, robotics, Markov

I. Motivation and Introduction

This paper involves cost-based planning problems in which the goals are expressed as time-extended sequences of discrete events whose occurrence the actor cannot causally influence, and which transitions from state to state exact some cost from the agent. This improves the previous model by accounting for real-world factors like distance traveled, time spent, resources expended, or other cost metrics.

II. Cost Matrix Addition

The addition of a cost matrix is an addition to the previous approach [3]. A cost transition function *C* is provided below,

$$\mathbb{C}(s_0, s_1) = K \tag{1}$$

where the constant cost *K* is evaluated given the current state s_0 and the state the model is transitioning to, s_1. To optimize cost using this function, the following equation can be used iteratively to determine the best transition to take, given each transition T_n contains a source state s_0 and destination state s_1.

$$min_{T_n \in T}\left(G(T_{n_{s_1}}) \cdot p(T_n) + \frac{\mathbb{C}(T_{n_{s_0}}, T_{n_{s_1}})}{\mathbf{max}(\mathbb{C})}\right) \tag{2}$$

where *T* is the set of all transitions to take, *G* stores values for the value-iteration algorithm, $p(T_n)$ is the transition probability, and *C* is the cost matrix from equation 1.

III. Wedding FOM Simulation and Analysis

Using this alteration, the Wedding FOM scenario was loaded. An example cost function for this scenario is provided below, representing the cost from the source state (rows) to destination state (columns).

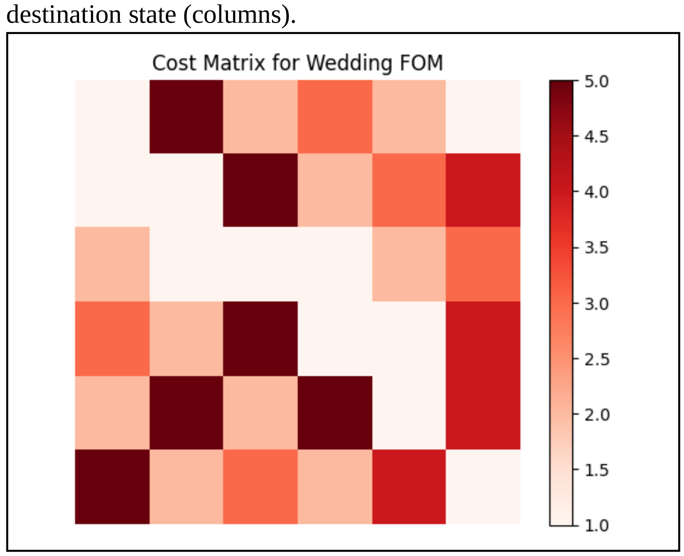

Fig. 1. The provided cost matrix for the wedding FOM example. Note that the minimum cost is one (1) to incentivise movement across the states.

During the training process, the value difference normalized from zero to one for both the original step-based optimization and cost-based optimization (equation 2) was collected and the results are provided below.

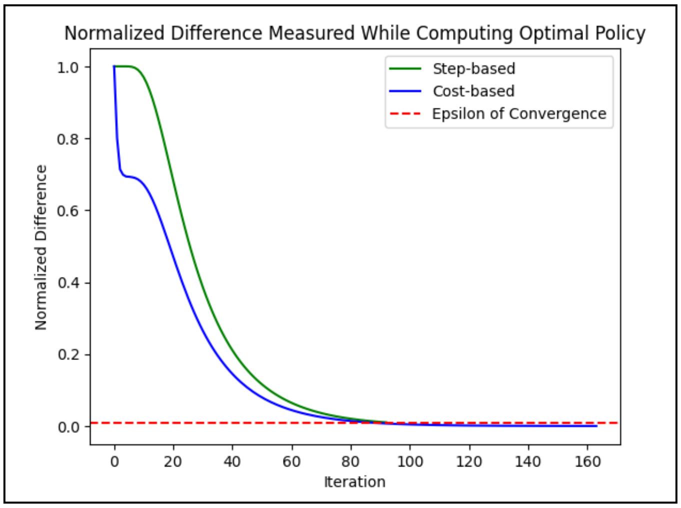

Fig. 2. The training differences for the step and cost based algorithms. Note the epsilon of convergence in this instance is set to 0.01. The cost-based algorithm appears to dip below it because of the scaling applied.

After training concluded, the two general algorithms were compared across 1,000 simulation trials, and the number of steps and costs accrued for each trial was recorded. The distribution of steps is provided below.

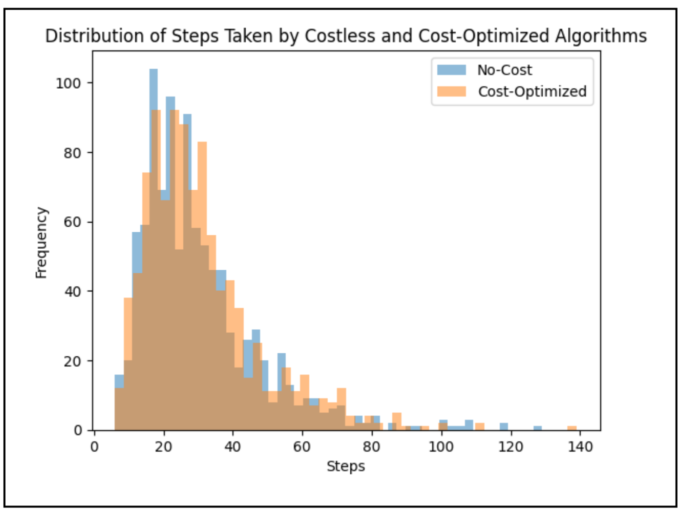

Fig. 3. The distribution of the number of steps for the no-cost and cost-optimized algorithm. As expected, there is not see a significant leftwards shift of the distribution, as the number of steps is not what the cost-optimization algorithm is primarily optimizing.

Perhaps of more interest, the authors constructed a histogram for the distribution of costs incurred by both general algorithms, with the results provided below.

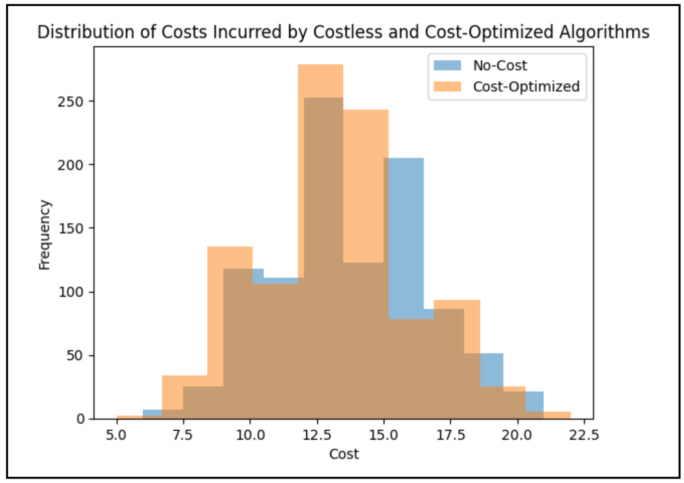

Fig. 4. The evaluation of 1000 trials on the costs incurred by the costless and cost-optimized general algorithms.

Summaries of both general algorithms across the 1,000 trials were then generated and are provided below. The step summary is provided first, with the cost summary provided after.

TABLE I. STEP SUMMARY ACROSS GENERAL ALGORITHMS

Algorithm	**Step Summary**	
	Measurement	***Value***
No-Cost	Minimum	6
	Q1	18
	Median	27
	Mean	30.313
	Q3	37
	Maximum	129
Cost-Optimized (Ours)	Minimum	6
	Q1	19
	Median	27
	Mean	30.725
	Q3	37
	Maximum	139

As we can see, the mean number of steps in each simulation for the cost-optimized algorithm is higher. This is likely due to the nature of the cost function, as optimizing for a lower cost may result in a higher number of steps. The table for the cost summaries across the no-cost and cost-optimized algorithms is provided below.

TABLE II. COST SUMMARY ACROSS GENERAL ALGORITHMS

Algorithm	**Cost Summary**	
	Measurement	***Value***
No-Cost	Minimum	6
	Q1	11
	Median	13
	Mean	**13.48**
	Q3	16
	Maximum	21
Cost-Optimized (Ours)	Minimum	5
	Q1	11
	Median	13
	Mean	**13.192**
	Q3	15
	Maximum	22

The authors then conducted a two-sample t-test with alpha of 0.05 to determine if the cost-optimized algorithm reduced either the number of steps or mean cost across the 1,000 trials.

TABLE III. T-TEST RESULTS FOR A MEANINGFUL STEP/COST DIFFERENCE

Test Conducted	**T-Test Results**	
	Measurement	***Value***
Step Test	T-Statistic	0.5417
	P-Value	0.7059
Cost Test	T-Statistic	-2.3171
	P-Value	0.0103

So, the authors fail to reject the null hypothesis for the step t-test: the cost-optimized algorithm is not significantly lower in steps than the no-cost general algorithm.

The authors rejected the null hypothesis for the cost t-test: the cost-optimized algorithm is significantly lower in cost than the no-cost algorithm. This implies that the cost-optimization is effective, though could be made more effective.

IV. ADDITIONAL IMPROVEMENTS

In addition to the cost-based algorithm itself, the authors have developed a comprehensive Python library to facilitate the implementation of FOM-based systems. This library provides a clean and accessible framework for researchers and practitioners to integrate the cost-based approach into their own projects. The containerization of the source code ensures cross-platform compatibility, making it easier to deploy and utilize the system across different environments.

To further enhance the usability and accessibility of the system, the authors have designed an extensible JSON configuration file. This allows users to configure and customize the system without requiring in-depth knowledge of the underlying technical implementation or programming language. By providing a user-friendly interface for system setup and configuration, the authors aim to lower the barrier to entry and encourage wider adoption of the cost-based approach.

V. FUTURE IMPROVEMENTS

Looking ahead, there are several potential avenues for additional improvements and extensions to the current work. One area of interest is the integration of real-time learning and adaptation capabilities into the system. By incorporating feedback loops and online learning mechanisms, the system could dynamically adjust its cost matrix and optimization strategies based on the observed outcomes and changing environmental conditions. This would enable the system to continuously refine its decision-making process and adapt to evolving circumstances.

Another promising direction is the exploration of multi-objective optimization techniques within the cost-based framework. In many real-world scenarios, there may be multiple competing objectives or constraints that need to be balanced, such as minimizing cost while maximizing event coverage or ensuring a certain level of narrative coherence. Developing algorithms that can effectively handle these multi-objective optimization problems would greatly enhance the flexibility and applicability of the cost-based approach.

Finally, the integration of the cost-based algorithm with other complementary technologies, such as computer vision, natural language processing, and human-robot interaction, could unlock new possibilities for intelligent and adaptive narrative observation systems. By leveraging the strengths of these different domains, researchers can develop more sophisticated and context-aware systems that can better understand and respond to the nuances of real-world events and human behaviors.

In conclusion, the cost-based algorithm presented in this paper and its experimental results, coupled with the development of accessible tools and frameworks, lay a solid foundation for further research and practical applications. As the field continues to evolve, it is essential to explore new directions, incorporate emerging technologies, and collaborate across disciplines to unlock the full potential of cost-based planning of this nature in real-world scenarios.

VI. Conclusion

The cost-based algorithm presented in this paper demonstrates a significant improvement over traditional value-based algorithms for planning in systems that involve passive observation of unpredictable events. By incorporating a cost matrix and optimizing for cost minimization, the proposed approach takes into account real-world factors such as distance traveled, time spent, and resources expended. The experimental results, based on a large sample size and a statistically significant difference in means, validate the effectiveness of the cost-based algorithm in reducing overall costs compared to the no-cost alternative.

However, it is important to acknowledge that there may be further opportunities for optimization and refinement of the cost-based approach. Future research should explore variations and extensions of the proposed algorithm to identify even more efficient solutions. This could involve investigating alternative cost functions, incorporating additional constraints, or exploring hybrid approaches that combine cost-based optimization with other techniques such as heuristics or machine learning.

Moreover, the applicability of the cost-based algorithm to a wider range of real-world scenarios should be examined. While the Wedding FOM simulation provides a valuable proof-of-concept, it is crucial to assess the performance and scalability of the approach in more complex and diverse environments. This may require adapting the algorithm to handle larger state spaces, more intricate event sequences, and dynamic cost structures.

VII.

Acknowledgment

The authors would like to thank the Texas A&M Department of Computer Science.

References

[1] Rahmani H and O’Kane JM (2019) Optimal temporal logic planning with cascading soft constraints. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Macau, China, 3-8 Nov. 2019, IEEE, pp. 2524–2531.

[2] Rahmani H, Shell D.A. and O’Kane JM (2020) Planning to chronicle. In: Algorithmic Foundations of Robotics (WAFR XIV), 21-23 June 2021.

[3] Rahmani, H., Shell, D. A., & O’Kane, J. M. (2023). Planning to chronicle: Optimal policies for narrative observation of unpredictable events. The International Journal of Robotics Research, 42(6), 412-432.