AdaCred: Adaptive Causal Decision Transformers with Feature Crediting

Hemant Kumawat Georgia Institute of Technology Atlanta, United States of America hkumawat6@gatech.edu

ABSTRACT

Reinforcement learning (RL) can be formulated as a sequence modeling problem, where models predict future actions based on historical state-action-reward sequences. Current approaches typically require long trajectory sequences to model the environment in offline RL settings. However, these models tend to over-rely on memorizing long-term representations, which impairs their ability to effectively attribute importance to trajectories and learned representations based on task-specific relevance. In this work, we introduce AdaCred, a novel approach that represents trajectories as causal graphs built from short-term action-reward-state sequences. Our model adaptively learns control policy by crediting and pruning low-importance representations, retaining only those most relevant for the downstream task. Our experiments demonstrate that AdaCred-based policies require shorter trajectory sequences and consistently outperform conventional methods in both offline reinforcement learning and imitation learning environments.

KEYWORDS

Offline Learning, Decision Transformer, Feature Crediting

ACM Reference Format:

Hemant Kumawat and Saibal Mukhopadhyay . 2025. AdaCred: Adaptive Causal Decision Transformers with Feature Crediting. In Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

1 INTRODUCTION

Transformer-based approaches, such as Decision Transformer (DT) [3] and its variants [8, 22, 26], have redefined offline reinforcement learning (RL) as a supervised learning problem by leveraging extended historical contexts to relax the strict Markov assumption. These models use self-attention mechanisms to process entire trajectory sequences, enabling the capture of dependencies across long temporal horizons. Despite their advantages, these methods face three primary challenges: (a) *Long Sequence Memory:* The requirement for long trajectory sequences to predict individual actions [22, 31], (b) *Suboptimal Trajectories:* Limited adaptability in processing datasets containing suboptimal trajectories [2, 6, 27], and (c) *Credit Assignment:* Lack of an explicit mechanism to assess and weigh the importance of trajectories and representations according to task-specific relevance [7, 19]. These challenges arise primarily

This work is licensed under a Creative Commons Attribution International 4.0 License. Saibal Mukhopadhyay Georgia Institute of Technology Atlanta, United States of America saibal.mukhopadhyay@ece.gatech.edu



b.) Causal View of the Representations for POMDPs

Figure 1: POMDP view of RL

from the inherent limitations of Transformers in RL, which, while proficient at memorizing sequences, struggle to assign appropriate importance to task-relevant trajectories [19]. Consequently, DT and its variants often require extended sequences-spanning 30 to 50 past trajectories-just to generate a single action. Moreover, these methods exhibit poor adaptability in scenarios involving suboptimal trajectories. By attending to entire sequences indiscriminately, they risk learning policies based on suboptimal behavior, thereby limiting overall performance. Recent advancements, such as Elastic Decision Transformer [26], aim to address these shortcomings by selectively "stitching" optimal segments from suboptimal trajectories. These methods estimate the expected reward for each segment and selectively attend to recent observations when the estimated reward is low. However, the exclusion of earlier observations when a trajectory is deemed suboptimal risks diluting essential relational information or overemphasizing redundant recent data points, which can ultimately hinder performance.

Unlike prior works that focus on learning and retaining all representations, we recast the problem as learning a **parsimonious causal graphical representation**. Each representation learned from the observations is treated as a node in a causal graph, and we reformulate the original task as identifying and pruning causal variables that are redundant for policy learning, based on the dynamically evolving causal structure. This adaptive approach allows the model to capture essential dependencies while discarding irrelevant data. Figure 1a illustrates the MDP perspective of reinforcement learning when the full state is available, whereas Figure 1b depicts the causal structure of representations when the full state is

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).



Figure 2: Our model design

not accessible. This representation explicitly encodes the structural relationships among different observations across the temporal domain. The core hypothesis of our approach is that, by structuring the learned representations as a causal graph, not all representations contribute to the future reward. Consequently, low-credit representations-those that do not influence future rewards-can be pruned without negatively impacting the agent's performance. This pruning mechanism selectively filters out irrelevant state components, allowing the model to focus on the most informative parts of the trajectory. Our approach significantly reduces the effective memory length of the sequence, allowing the model to allocate computational resources more efficiently by prioritizing high-credit representations. By selectively pruning low-credit representations, our framework not only mitigates the negative impact of suboptimal trajectories but also reinforces critical relationships within the sequence, thus enhancing decision-making quality. Figure 2 illustrates the overall design of spatial and temporal credit assignment for decision transformer models.

We support our framework with rigorous theoretical results, where we prove that the proposed credit assignment mechanism leads to the identification of a minimal set of latent state components that are sufficient for effective policy learning. We formally demonstrate that the pruning process does not negatively impact the agent's ability to maximize future rewards, as irrelevant features are effectively removed while preserving the necessary structural relationships in the causal graph. We empirically validate our method on well-established benchmarks, including the Atari and Gym environments, where our model demonstrates superior decision-making and sequence generation capabilities compared to Decision Transformer and its variants. The experimental results reveal that our credit assignment and pruning mechanisms lead to more effective policy learning while maintaining computational efficiency. Our method introduces minimal overhead during training, yet yields significant improvements in performance. By surpassing state-of-the-art methods in both offline and imitation learning tasks, our approach opens new avenues for future research in reinforcement learning models that emphasize structured and efficient representation learning.

In this paper, we make the following key contributions:

- We introduce **AdaCred**: Adaptive Causal Decision Transformers, a novel approach that addresses the challenge of suboptimal sequence representations in Decision Transformer and related models by leveraging causal latent graph structures with feature crediting.
- The causal latent graph representation allows AdaCred to assign credit to individual sequence representations and prune low-importance components, identifying optimal sequences even from suboptimal trajectories, thereby improving decision-making efficiency.
- We provide a rigorous theoretical framework for our model and our experimental results on standard benchmarks demonstrate the superior performance of AdaCred in both offline and imitation learning settings.

2 RELATED WORK

Transformer-based RL Methods. One of the most notable Transformer based models for offline RL is the Decision Transformer (DT) [3], which recasts RL as a supervised learning task, treating trajectory sequences as input and predicting the optimal actions. The DT and its variants [8, 22, 26] have demonstrated strong performance across a range of benchmarks, thanks to their ability to model long-term dependencies via self-attention mechanisms. However, these methods face challenges when applied to suboptimal trajectories or when dealing with long sequences, often requiring the model to attend to up to 50 past trajectories to generate a single action. Recent efforts, such as Elastic Decision Transformer (EDT) [26], address these limitations by selectively attending to optimal trajectory segments and dynamically stitching together segments with high estimated rewards. However, such approaches may lose essential information by prematurely discarding earlier observations, particularly in long-horizon tasks. Our approach extends the Transformer-based RL framework by introducing adaptive credit assignment to mitigate the reliance on long sequences and suboptimal data, allowing the model to focus on task-relevant representations.

Reinforcement Learning with Suboptimal Data Incorporating suboptimal trajectories into RL models has proven to be a challenging task, as suboptimal data can significantly degrade the quality of learned policies. Methods such as Conservative Q-Learning (CQL) [12] have attempted to mitigate the effects of suboptimal data by introducing conservative constraints during policy optimization. Similarly, Data-Driven Offline RL [18] aims to learn policies from imperfect datasets while minimizing the impact of suboptimal actions and states. Recent works have also explored sequence stitching techniques [26] to combine high-reward segments from suboptimal trajectories, improving overall performance. However, these methods often fail to adapt dynamically to the changing relevance of different parts of the sequence.

3 ADAPTIVE CAUSAL DECISION TRANSFORMER

3.1 Theoretical Discussion

In this section, we present a formal theoretical framework that underpins our AdaCred model. The motivation for this analysis is to rigorously define the representation of latent states in reinforcement learning, introduce a feature credit assignment mechanism, and study the effects of pruning latent representations on policy learning.

POMDPs in Offline Reinforcement Learning. We consider a partially observable Markov decision process $M = (S, O, A, P, Z, R, \gamma, T)$, where at each time step $t \in \{1, ..., T\}$, an agent is in a state $s_t \in S$, receives an observation $o_t \in O$, selects an action, $a_t \in A$ and receives a reward r_t . The reward is distributed as $r_t \sim R(s_t, a_t)$, and the environment transitions to the next state according to $s_{t+1} \sim P(\cdot | s_t, a_t)$. The next observation o_{t+1} is sampled from the observation model $Z(\cdot|s_{t+1})$. The initial state s_1 is sampled from the distribution $P(s_1)$, and the initial observation o_1 is sampled from $Z(o_1|s_1)$. The problem is characterized by a time horizon $T \in \mathbb{N}^+ \cup \{+\infty\}$ and a discount factor $\gamma \in [0, 1]$. In the offline setting, the agent has access to a trajectory $\tau = \{o_1, a_1, r_1, \dots, o_t, a_t, r_t\},\$ which reflects the agent's experience of the environment. For sequence modeling in offline RL, we employ Transformer architectures, treating observations, actions, and rewards as input tokens to approximate the optimal action $P(a_t | o_{1:t}, a_{1:t-1}, r_{1:t-1})$.

Graphical Representation of Latent State Transitions. We denote the underlying latent states for o_t as $\mathbf{g}_t = (g_{1,t}, \dots, g_{d,t})^{\mathsf{T}}$, where *d* is the dimensionality of the latent states. The environment's generative process can be described by the following transition, observation, and reward functions for each dimension of **g**:

$$g_{i,t} = f_i(c_i^{g \to g} \odot \mathbf{g}_{t-1}, c_i^a \cdot \mathbf{a}_{t-1}, c_i^r \cdot \mathbf{r}_{t-1}, \epsilon_{i,t}^g), \quad \text{for } i = 1, \dots, d,$$

$$o_t = g(c^{g \to o} \odot \mathbf{g}_t, \epsilon_t^o),$$

$$r_t = h(c_r^{g \to r} \odot \mathbf{g}_{t-1}, c^{a \to r} \cdot \mathbf{a}_{t-1}, \epsilon_t^r),$$

where \odot represents element-wise multiplication, and ϵ_t^g , ϵ_t^o , ϵ_t^r are independent noise terms. The terms $c \rightarrow \cdot$ are binary masks indicating causal relationships between variables. The latent states g_{t+1} form a Markov decision process (MDP): given g_t and a_t , the next latent state g_{t+1} is independent of previous states and actions. The observed signals o_t are generated from the latent states g_t , and rewards are determined by both the latent states and actions.

Definition 1 (Graphical Representation of Latent State Transitions). Given a causal graph G = (V, E), where the nodes Vrepresent the latent state components $\mathbf{g}_t = (g_{1,t}, \dots, g_{d,t})^{\top}$, and edges $E \subseteq V \times V$ represent causal relationships between them, a binary mask $\mathbf{c}^{g \to g} \in \{0, 1\}^{d \times d}$ encodes the presence of edges. The *i*, *j*-th entry of $\mathbf{c}^{g \to g}$ is 1 if $g_{i,t+1}$ is influenced by $g_{j,t}$, and 0 otherwise. Similarly, binary masks $\mathbf{c}^{a \to g}, \mathbf{c}^{r \to g} \in \{0, 1\}^d$ represent causal effects of actions and rewards on latent states.

For each latent state component $g_{i,t}$, the crediting mechanism outputs a binary mask $c_{i,t} \in \{0, 1\}$, where $c_{i,t} = 1$ if $g_{i,t}$ contributes to the future reward r_{t+k} for some k > 0, and 0 otherwise. The binary mask is learned through a credit assignment process that leverages causal relationships encoded in c, as well as the temporal dependence of state features on rewards. We now present key theoretical results, starting with the minimal sufficient representations [10, 25] required for policy learning.

Theorem 1 (Minimal Sufficient Representations for Policy Learning). Under the assumption that the causal graph G is Markov and faithful to the observed data, the set of minimal latent states $\mathbf{g}_t^{\min} \subseteq \mathbf{g}_t$ is defined as: $g_{i,t} \in \mathbf{g}_t^{\min}$ if $c_i^{g \to r} = 1$ or $g_{i,t}$ has a directed path to a future reward through other latent states This set forms a minimal and sufficient representation for policy learning, ensuring that only these latent state components are necessary to maximize the expected future rewards. (For full proof, please refer to Appendix B.1 [17])

PROOF SKETCH. The proof follows from the Markov property, which ensures that future rewards depend only on the current latent states and actions. By d-separation, any state $g_{i,t}$ that is not part of \mathbf{g}_t^{\min} is conditionally independent of future rewards given \mathbf{g}_t^{\min} , and thus, pruning these states does not affect the policy's ability to predict future rewards. The remaining states in \mathbf{g}_t^{\min} are those that directly or indirectly contribute to future rewards, making them both necessary and sufficient for policy learning.

Theorem 2 (Structural Identifiability). Suppose the underlying latent states g_t are observed, i.e., the environment follows a Markov Decision Process (MDP). Then, under the Markov condition and faithfulness assumption, the structural matrices $c^{g \rightarrow g}$, $c^{a \rightarrow g}$, $c^{g \rightarrow r}$, $c^{a \rightarrow r}$ are identifiable. (For full proof, please refer to Appendix B.2 [17])

PROOF SKETCH. The identifiability of the structural matrices follows from conditional independence relations in the data. By the faithfulness assumption, the observed conditional independencies reflect the structure of the causal graph G. Since the dynamic Bayesian network (DBN) is time-invariant and adheres to the Markov property, the causal dependencies encoded in the matrices c can be identified by analyzing the conditional distributions of the latent states, actions, and rewards. This guarantees the identifiability of the structural matrices.

Causal Graph Pruning As previously discussed, not all latent states contribute directly to future actions and rewards. During learning, we can identify minimal sufficient state representations that are essential for task performance. Based on the theoretical results above, we can classify the state representations and prune the causal graph accordingly. This pruning process aims to reduce redundancy in the latent space, thus improving both the efficiency and generalization capability of the model.

Following the structural analysis of latent states, we define two classes of state representations, adapted from [29, 30]:

1. **Compact State Representation** (\mathbf{g}_t^c) : This class includes latent state variables that either influence the observation o_t , contribute to the reward r_{t+1} , or affect other latent states $g_{j,t+1}$ at





b.) Temporal Causal Transformer

Figure 3: a.) Spatial Transformer with Crediting. Output of spatial transformer is sent to the next spatial transformer. b.) Temporal Causal Transformer with Crediting

the next time step. Formally, a state variable $g_{i,t}$ is part of this compact set if any of the following structural dependencies hold: $c^{g \to o} = 1$ or $c^{g \to r} = 1$ or $c^{g \to g_j} = 1$, $\forall j \neq i$. These variables form the core of the latent state space that is essential for policy learning and prediction of future rewards.

2. Non-Compact State Representation (\mathbf{g}_t^c) : This class contains state variables that do not satisfy the criteria for compact representations. These variables can be pruned without loss of policy learning capability, as they neither influence the immediate observations nor the future rewards.

By distinguishing between compact and non-compact state representations, we can simplify the causal graph by removing nonessential latent states. To further facilitate the pruning process, we introduce a regularization term that promotes sparsity in the structural masks $c^{\cdot \rightarrow \cdot}$ by encouraging non-essential dependencies to be pruned. Inspired by the edge-minimality principle [30], we define the regularization term as follows:

$$J_{\text{reg}} = -\lambda_{\text{reg}} \left(\| c^{g \to o} \|_1 + \| c^{g \to r} \|_1 + \| c^{g \to g} \|_1 + \| c^{a \to g} \|_1 + \| \theta \|_1 \right),$$

where λ_{reg} is a regularization coefficient, and $\|\cdot\|_1$ represents the ℓ_1 norm, which encourages sparsity. Incorporating this regularization term into the objective function *J* enables simultaneous pruning and model training. The presence of this regularization term naturally leads to the selective removal of non-compact state variables, as entries in $c^{-\rightarrow}$ transition from 1 to 0, promoting a sparse, efficient representation of the causal relationships. By reducing the number of latent states and their interactions, the model becomes more interpretable and generalizes better to new tasks.

3.2 AdaCred Algorithm

To effectively separate the learning of spatial and temporal representations, we design our causal decision transformer with two distinct components: (a) a Spatial Transformer and (b) a Causal Temporal Transformer. The Spatial Transformer learns spatial representations from the state, action, and reward at each individual time step by performing self-attention on state-action-reward tokens within a single time-step window. These spatial representations are then passed to the Causal Temporal Transformer, which combines them with the pure state representations to model the entire sequence of states. The final action predictions are generated by a prediction decoder. We address the challenge of temporal and spatial feature credit assignment by introducing a feature crediting mechanism that quantifies the importance of latent representations over time. At each time step t, the agent observes a partial state and produces a latent representation $g_t \in \mathbb{R}^d$. We introduce a crediting function $\sigma: \mathbb{R}^d \to [0,1]^d$ to compute token selection probabilities for each latent feature, which represent the probability that a given feature contributes significantly to future rewards. The feature crediting mechanism outputs a binary mask, $m_t \in \{0, 1\}^d$, The resultant mask m_t serves as a gating mechanism, assigning high values to those features deemed influential in maximizing future rewards and low values to those with minimal impact. Each of the spatial and causal temporal transformer block has this crediting function. In the following sections, we will discuss the design and implementation of these components in detail.

Algorithm 1 Feature Crediting and Pruning for Single Spatial and Temporal Transformer

- 1. **Input:** Observations $\{o_t\}$, Actions $\{a_t\}$, Rewards $\{r_t\}$, Time Horizon *T*
- 2. Initialize parameters θ for Transformers
- 3. **For** each time step t = 1 to T:
 - 3.1. $g_t \leftarrow \text{SpatialTransformer}(o_t, a_t, r_t; \theta)$
 - 3.2. $\sigma(q_t) \leftarrow \text{CreditFunction}(q_t; \theta)$
 - 3.3. $m_t \leftarrow \text{GumbelSigmoid}(\sigma(g_t))$
 - 3.4. $g_t^{\text{selected}} \leftarrow m_t * g_t$
 - 3.5. $g_t^{\text{processed}} \leftarrow \text{TemporalTransformer}(g_t^{\text{selected}}; \theta)$
 - 3.6. $\sigma_{\text{temporal}}(g_t^{\text{processed}}) \leftarrow \text{CreditFunction}(g_t^{\text{processed}})$
 - 3.7. $m_t^{\text{temporal}} \leftarrow \text{GumbelSigmoid}(\sigma_{\text{temporal}}(g_t^{\text{processed}}))$
 - 3.8. $a_{\text{final}}^{\text{final}} \leftarrow m_{\text{final}}^{\text{temporal}} * a^{\text{processed}}$

5.8.
$$g_t \leftarrow m_t \qquad * g_t$$

- 3.9. $a_t^{\text{pred}} \leftarrow \text{ActionDecoder}(g_t^{\text{final}}; \theta)$
- 3.10. Observe reward r_t and next observation o_{t+1}

4. End For

5. **Output:** Predicted actions $\{a_t^{\text{pred}}\}$, optimized policy π

Spatial Transformer We adopt an Observation-Action-Reward (OAR) grouping mechanism, following a similar approach as in [22]. Each trajectory is partitioned into segments comprising the previous action, the corresponding reward, and the current observation, highlighting the causal interrelationships among these elements. To obtain a fine-grained representation of the state, we partition each input state into non-overlapping spatial patches, thereby enabling localized interactions between actions, rewards, and regions of the state. The action and reward are linearly embedded to produce corresponding representations, which are then combined with the observation token embeddings. This results in an initial representation set of observation, action, and reward tokens, defined as: $S_0^t = \{s_{a_{t-1}}, s_{r_t}, s_{o_{1,t}}, s_{o_{2,t}}, \dots, s_{o_{n,t}}\}$, where $s_{a_{t-1}}$, s_{r_t} , and $s_{o_{t,t}}$ represent the embeddings for actions, rewards, and

observation patches, respectively. This is shown in figure 3a where spatial crediting function adaptively computes the tokens to be processed by next layer and prunes the low credit ones .

To process these tokens, we employ a Transformer-based architecture. Specifically, the Spatial Transformer layer maps token representations from the previous layer (S_t^{l-1}) to the current layer (S_t^l) using the transformation: $S_t^l = F_{step}^l(S_t^{selected,l-1})$, where $S_t^{selected,l-1}$ represents the subset of features selected by a spatial token policy, which determines the relevant spatial tokens to be processed by the Spatial Transformer block. This token policy serves as a crediting mechanism, selectively pruning unnecessary spatial features before passing them to the subsequent layer, thereby enhancing computational efficiency. For each Spatial Transformer layer l, an aggregated Observation-Action-Reward representation $g_t^l \in \mathbb{R}^D$ is generated by pooling the output tokens $(S_t^l \in \mathbb{R}^{n \times d})$, computed as: $g_t^l = FC([concat(S_t^l)]) + e_t^{temporal}$, where concat(·) denotes the concatenation of the tokens within each group, and $e_t^{temporal} \in \mathbb{R}^D$ is a temporal positional embedding. The resulting OAR representation g_t^l is then passed to a Temporal Causal Transformer layer for long-term temporal modeling, effectively capturing both local and global dependencies across the trajectory.

Temporal Causal Transformer: To effectively model long-term sequences, we introduce the Temporal Causal Transformer, which utilizes the learned spatial representations from the Spatial Transformer, along with pure state representations embedded without any explicit action or review, processed by an encoder.

The pure state representations are defined as follows: $h_t^0 =$ $Conv(o_t) + e_t^{temporal}$, where o_t represents the image input at time t. The output of the Spatial Transformer and the learned convolutional state embedding are combined into a sequence, preserving the temporal order in which they were received: Y_{in}^{l} = $\{g_1^l, h_1^{l-1}, g_2^l, h_2^{l-1}, \dots, g_T^l, h_T^{l-1}\}$. This Transformer layer takes input from each intermediate spatial layer, rather than only the final spatial representation after all layers. Similar to the Spatial Transformer, we employ a token policy that removes certain tokens based on the token policy. The intermediate output of the Temporal Causal Transformer is given by: $Y_{out}^l = F_{sequence}^l(Y_{in}^l)$ where similar to the spatial transformer the credit function determines which representations to prune adaptively. This is shown in the figure 3b. From Y_{out}^l , we select tokens at even indices (starting from 1) to be the pure state tokens $h_i^l := y_{\text{out};2i}^l$. These tokens are then passed to the next Temporal Causal Transformer layer. The final output of the Temporal Causal Transformer layer is used to predict actions using a linear head, denoted as \hat{a} . Algorithm 1 shows the overall algorithm for our model for single layer networks.

AdaCred Training Method The training process for our Adaptive Causal Decision Transformer framework consists of two main stages: training a static model and co-training the model along with the crediting mechanism, ensuring that both spatial and temporal feature assignments are learned effectively.

Stage 1: Static Model Training: The initial phase involves training the static components of the AdaCred model, comprising the Spatial Transformer and Temporal Causal Transformer. In this stage, we focus on learning the basic spatial and temporal representations that underpin the entire sequence modeling. The training objective

for this stage is to minimize the action prediction loss, which is defined as: $\mathcal{L}_{stage1} = \sum_{t=1}^{T} \mathcal{L}_{action}(\hat{a}_t, a_t)$ where \hat{a}_t represents the predicted action at time t, a_t is the ground-truth action, and T is the total number of time steps. The action prediction loss, \mathcal{L}_{action} , measures the discrepancy between the predicted and ground-truth actions.

Stage 2: Policy Network with Feature crediting: In this stage, we focus on jointly training the model with the crediting mechanism, optimizing the binary masks $m_t \in \{0, 1\}^d$ that govern feature selection for both Spatial and Temporal Causal Transformers. At each time step t, the Spatial Transformer and the Temporal Causal Transformer use a feature crediting mechanism $\sigma : \mathbb{R}^d \to [0, 1]^d$ to compute token selection probabilities for each latent feature representation q_t . The crediting mechanism assigns importance to each feature, generating binary masks that act as gates, allowing only the most influential features to pass through. During co-training, we optimize both the Spatial Transformer and Temporal Causal Transformer layers to ensure that the selected features (via binary masks) maximize the future rewards while maintaining computational efficiency. The training objective for this stage incorporates both the task prediction loss and an efficiency loss, which accounts for the number of features selected: $\mathcal{L}_{stage2} = \mathcal{L}_{action} + \alpha \mathcal{L}_{eff}$ where \mathcal{L}_{action} is the total loss across all action predictions, and \mathcal{L}_{eff} measures the efficiency of feature selection. The efficiency loss is calculated as: $\mathcal{L}_{\text{eff}} = \text{MSE}\left(\omega_d \cdot \frac{T_{\text{activ}}}{T_{\text{total}}}, \omega_d \cdot \frac{T_{\text{target}}}{T_{\text{total}}}\right)$ where T_{activ} represent the number of activated tokens, while T_{target} are the desired activation percentages. The term ω_d accounts for the different embedding dimensions at different layers, ensuring that the efficiency loss reflects both the number and the size of the selected features, adequately balancing the computational burden.

4 EXPERIMENTS

Our experiments aim to address the following key research questions, each corresponding to specific aspects of our study:

- **Performance Comparison**: Does AdaCred outperform Decision Transformers (DT) and its variants in Atari and Gym benchmarks? (See *Section 4.2.1*)
- Effectiveness in Feature Pruning: How effective is Ada-Cred in identifying and pruning both spatial and temporal features? (See *Section 4.2.2*)
- Effect of Pruning percentage on performance: How much can we prune spatially and temporally in Adacred (See *Section 4.2.3*)
- **Sparse Task Performance**: How well does AdaCred perform in tasks that require sparse representations? (See *Section 4.2.4*)

4.1 Experimental setting

4.1.1 Dataset. Our experiments encompass both offline reinforcement learning (RL) [9] and imitation learning via behavior cloning. For offline RL, we use a fixed dataset of sub-optimal trajectory rollouts, which presents challenges such as distribution shift, a key difficulty in offline RL compared to standard online RL settings [13, 23, 24]. In the imitation learning scenario, the agent is



Performance vs. Sequence Length for Games for Imitation Learning

Figure 4: Performance for 75% Spatial and Temporal Crediting

Dataset	DT	CQL	QR-DQN	REM	Baseline (100% S & 100% T*)	Ours (75% S & 75% T*)
Assault	462 ± 139	432	142	350	458 ± 166	658 ± 66
Boxing	78.3 ± 4.6	56.2	14.3	12.7	78 ± 14	78.4 ± 14.5
Breakout	76.9 ± 17.1	55.8	4.5	2.4	52 ± 14	77 ± 12
Pong	12.8 ± 3.2	13.5	2.2	0.0	14 ± 7.68	18 ± 2.54
Qbert	3488 ± 631	14012	0.0	0.0	1612 ± 1860	4487 ± 5099
Seaquest	1131 ± 168	685	161	282	890 ± 454	968 ± 505

Table 1: Performance Comparison of Various Methods in Offline RL Setting. (X % S and Y% T means X % of Spatial and Y% of Temporal tokens have been used for learning the model)

trained without reward signals or real-time environment interaction, which significantly increases the learning difficulty. Unlike traditional imitation learning, where additional data collection and Inverse Reinforcement Learning are feasible, we work solely with sub-optimal trajectories. Following the approach in [22], we derive the imitation learning dataset by removing rewards from the offline dataset. We evaluate our model in both discrete and continuous action space environments. For discrete environments, we use the Atari benchmark, selecting six games: Assault, Boxing, Breakout, Pong, Qbert, and Seaquest. We employ 1% (500,000 transitions) of the DQN replay buffer to facilitate thorough and fair comparisons. For continuous control tasks, we use the D4RL benchmark [5], which includes HalfCheetah, Hopper, Walker, and a 2D reacher task with sparse rewards to explore beyond the typical locomotion benchmarks. The datasets for these experiments feature three settings: (1) 'Medium,' which consists of 1 million timesteps generated by a policy with moderate performance (about one-third of an expert); (2) 'Medium-Replay,' which contains the replay buffer data collected during training of the medium policy; and (3) 'Medium-Expert,' combining 1 million timesteps from both medium and expert policies. We report episodic returns as absolute values, averaged over ten random seeds for the Atari games and gym tasks. Each seed is evaluated across ten episodes, with random initial conditions to provide a robust performance assessment.

4.1.2 Baselines. For baseline comparisons, we include our method using 100% of spatial and temporal tokens, following the approach in [22]. We evaluate performance by progressively reducing the proportion of spatial and temporal tokens for action prediction. Additionally, we compare against the Decision Transformer [3], its variants with variable-length trajectories [26], a Q-learningbased method for suboptimal trajectories [28] and EDT [26]. We further benchmark against state-of-the-art non-Transformer offline RL methods, including CQL[12], QR-DQN[4], REM[1], and BEAR[11].

4.2 Results

4.2.1 Analysis of Atari Games for Offline Reinforcement Learning and Imitation Learning. We evaluate the effectiveness of our credit assignment and pruning mechanism on Atari games, focusing on both the imitation learning and offline reinforcement learning (RL) settings. To establish a baseline, we first train a model without any credit assignment or pruning mechanisms. We then compare it to a version of the model that incorporates 75% spatial and temporal pruning. For both models, we use sequence lengths of 10, 20, and 30 steps, consistent with the experimental setup of the Decision Transformer (DT). The results for four Atari environments are illustrated in Figure 4, while Figure 5 presents the results for the offline RL setting. In both settings, the baseline model exhibits the highest performance when sequence trajectories of 30 time steps are provided, with a noticeable decline in performance as the sequence length is reduced. This indicates that the baseline model relies on longer sequences to capture relevant information for decision-making. Conversely, our model demonstrates a marked improvement in efficiency, achieving optimal performance with sequences as short as 10 time steps and maintaining consistent performance across all tested sequence lengths. Remarkably, even with just 10 steps, our model consistently outperforms the baseline model's best results with 30-step sequences. This clearly demonstrates that our approach effectively learns to assign credit and prune unnecessary



Figure 5: Performance for 75% Spatial and Temporal Crediting for Offline RL

Dataset	Environment	DT Sparse	QDT Sparse	BEAR Dense	EDT Dense	AWR	Baseline Dense	Ours Dense	Baseline Sparse	Ours Sparse
Medium-Expert	HalfCheetah	-	-	53.4	-	52.7	77.4 ± 1.2	77.8 ± 1.3	67 ± 1.2	81 ± 1.6
Medium-Expert	Hopper	-	-	96.3	-	27.1	97.9 ± 4.6	102.7 ± 5.6	106.2 ± 2.6	97.33 ± 3.6
Medium-Expert	Walker	-	-	40.1	-	53.8	108.5 ± 7.6	108.8 ± 5.6	108 ± 3.6	109 ± 4.6
Medium	HalfCheetah	42.2 ± 0.2	42.4 ± 0.5	41.7	42.5 ± 0.9	37.4	42.4 ± 1.6	42.8 ± 2.6	42.1 ± 0.5	42.7 ± 0.7
Medium	Hopper	57.3 ± 2.4	50.7 ± 5.0	52.1	63.5 ± 5.8	35.9	51.2 ± 2.1	55.6 ± 1.2	52.11 ± 0.6	54.47 ± 2.1
Medium	Walker	69.9 ± 2.0	63.7 ± 6.4	59.1	72.8 ± 6.2	17.4	72.4 ± 1.8	74 ± 3.9	38.46 ± 3.8	75.15 ± 4.2

Table 2: Performance comparison across datasets and environments in Medium and Medium-Expert settings.



Figure 6: Visualization of our spatial masks with attention

outputs

representations, thereby retaining only the most relevant portions of the trajectory data. This result is particularly significant for resource-constrained environments, as the model's performance remains robust while processing only 33% of the sequence data compared to the baseline.

In Table 1, we provide a comprehensive performance comparison between our proposed model with 75% spatial and temporal pruning and various baseline methods, including Decision Transformer (DT), Conservative Q-Learning (CQL), QR-DQN, and REM, across six different Atari games. The metrics presented reflect the average rewards obtained, along with the standard deviation where applicable. Our approach demonstrates superior performance in nearly all tasks, particularly excelling in games that typically require longer-term planning and spatial awareness, such as Assault and Seaquest. For Qbert, which involves complex decision-making and long-term strategy, our model exhibits an impressive performance of 4487 \pm 5099, significantly improving upon the baseline (1612 \pm 1860) and outperforming DT (3488 \pm 631). The results indicate that our model effectively manages credit assignment over longer sequences, even in environments with extensive trajectory dependencies. Overall, Table 1 highlights the competitive advantage of our approach across various environments, outperforming or matching traditional methods, including DT and CQL. Notably, while methods like CQL are specifically designed to optimize policies from suboptimal data, they struggle to generalize effectively in these environments, achieving suboptimal rewards.

4.2.2 Effectiveness of Spatial and Temporal Crediting. To demonstrate the model's capability in assigning credit and pruning representations that are not relevant to the task, we present a visualization of the spatial token masks in Figure 6. The figure showcases the Breakout game scenario, where the model was trained for 30 timesteps, highlighting key moments and decision-making processes during gameplay. The top row of Figure 6 displays the output of the spatial mask from the first layer of the spatial transformer for the frames with highest temporal importance based on temporal token mask. In this visualization, the black regions correspond to areas of the image that are disregarded by the model, indicating that they provide little or no useful information for the task at hand. From the spatial mask, it is evident that the model learns to focus on the relevant regions containing dynamic elements, such as the paddle and the ball, while ignoring other static parts of the environment. Specifically, the central region of the game screen, which primarily contains unchanging elements, is discarded, as it does not contribute to the decision-making process. The spatial transformer efficiently identifies the regions where important interactions occur, such as the movement of the ball and paddle, which are assigned

higher attention scores. This is clearly reflected in the attention maps, where the ball and paddle are shown to receive the highest attention, indicating their critical role in the task. Moreover, the temporal token mask reveals that the model focuses on key events, which can be categorized into three main interactions: a.) Collision with the paddle (E1): The ball's interaction with the paddle is crucial for determining the next trajectory and is therefore heavily weighted. b.) Collision with the side wall (E2): The side wall collision alters the ball's path and is a significant event that affects the gameplay. c.) Hitting a game tile (E3): This event results in a reward, making it an essential moment for credit assignment. The model successfully learns to emphasize these key interactions, which are crucial for effective decision-making. By focusing on these spatial and temporal aspects, the network is able to prune irrelevant representations and learn efficient action policies. This targeted credit assignment mechanism enables our model to outperform others in similar tasks by prioritizing relevant events and discarding extraneous information.

4.2.3 Effect of pruning percentage on performance. In Figure 7, we examine the effect of varying degrees of spatial and temporal pruning on the performance of our model in the Breakout game. The pruning ratios are defined as (X%, Y%), where X% represents the percentage of spatial features retained by the model, and Y% represents the percentage of temporal features retained. We test the following combinations: (50%, 50%), (50%, 100%), (40%, 80%), and (100%, 50%), with the baseline model using 100% of both spatial and temporal features. The model achieves its highest performance with the (50%, 100%) pruning ratio, where 50% of the spatial features are retained while all temporal features are used. This result supports our hypothesis that, among the 30 sequences analyzed, the majority of redundancies are spatial rather than temporal. In the context of Breakout, the spatial redundancies primarily come from the images, where large portions of the background contain redundant or irrelevant information that does not contribute to decision-making. By pruning these less critical spatial features, the model can generalize better, focusing on the most relevant spatial information, such as the paddle and ball, while retaining the full temporal context to capture important trajectory-based decisions. The combination of spatial pruning and full temporal retention leads to a significant improvement in performance, with the model achieving approximately 2.8 times the baseline performance. On the other hand, the lowest performance is observed with the (100%, 50%) pruning ratio, where all spatial features are used but 50% of the temporal features are pruned. In this case, retaining all spatial features introduces redundant spatial information, which in turn establishes misleading causal relationships in the temporal domain. The absence of sufficient temporal information exacerbates the issue, leading the model to infer incorrect temporal dependencies. If proper temporal pruning had been applied, these erroneous temporal relationships could have been removed, allowing for better decision-making. This pruning ratio significantly underperforms compared to the baseline, emphasizing the critical role of temporal features in tasks requiring long-term planning and time-sensitive decisions. Overall, with the exception of the (100%, 50%) case, all other pruning strategies outperform the baseline.



Figure 7: Effect of Spatial and Temporal Pruning

4.2.4 Analysis in Sparse Reward Settings. In sparse reward environments, traditional reinforcement learning methods face significant challenges in credit assignment due to the delayed nature of rewards and the need to track which past actions contributed to success. Our proposed model addresses this by implementing a feature crediting mechanism that quantifies the importance of latent features at each decision step, enabling more precise credit assignment without heavy reliance on memory. The results in Table 2 demonstrate the effectiveness of this approach across Medium and Medium-Expert environments. Our model consistently outperforms baselines, such as Decision Transformer (DT) and EDT, particularly in environments like HalfCheetah, Hopper, and Walker, where our method excels in sparse reward settings. For example, in the Medium-Expert HalfCheetah environment, our model achieves 81 ± 1.6 , significantly surpassing the baseline (67 ± 1.2). Similarly, strong performance is observed in Walker and Hopper environments, demonstrating the robustness of our crediting mechanism across different tasks.

5 CONCLUSION

In this paper, we introduced AdaCred, a novel framework that addresses the challenges of suboptimal sequence representations in reinforcement learning through adaptive credit assignment and pruning of low-importance latent representations. By leveraging a causal latent graph structure, our approach efficiently identifies and retains critical components of the trajectory, improving both decision-making and computational efficiency. Theoretical analysis and empirical results on standard benchmarks demonstrate the effectiveness of AdaCred, surpassing state-of-the-art methods in both offline and imitation learning settings. In future, we would like to extend this to transformer based multiagent control [14, 15], as well as computer vision tasks such as detection and classification[16, 20, 21].

ACKNOWLEDGMENTS

This work is based in part on research supported by SRC JUMP2.0 (CogniSense Center, Grant: 2023-JU-3133). The views, opinions, findings, and conclusions expressed in this material are those of the author(s) and do not necessarily reflect the official policies or positions of SRC.

REFERENCES

- Rishabh Agarwal, Dale Schuurmans, and Mohammad Norouzi. 2020. An optimistic perspective on offline reinforcement learning. In *International conference* on machine learning. PMLR, 104–114.
- [2] Ian Char, Viraj Mehta, Adam Villaflor, John M. Dolan, and Jeff Schneider. 2022. BATS: Best Action Trajectory Stitching. arXiv:2204.12026 [cs.LG] https://arxiv. org/abs/2204.12026
- [3] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. 2021. Decision transformer: Reinforcement learning via sequence modeling. Advances in neural information processing systems 34 (2021), 15084–15097.
- [4] Will Dabney, Mark Rowland, Marc Bellemare, and Rémi Munos. 2018. Distributional reinforcement learning with quantile regression. In Proceedings of the AAAI conference on artificial intelligence, Vol. 32.
- [5] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. 2020. D4rl: Datasets for deep data-driven reinforcement learning. arXiv preprint arXiv:2004.07219 (2020).
- [6] Charles A. Hepburn and Giovanni Montana. 2022. Model-based Trajectory Stitching for Improved Offline Reinforcement Learning. arXiv:2211.11603 [cs.LG] https://arxiv.org/abs/2211.11603
- [7] Chia-Chun Hung, Timothy Lillicrap, Josh Abramson, Yan Wu, Mehdi Mirza, Federico Carnevale, Arun Ahuja, and Greg Wayne. 2018. Optimizing Agent Behavior over Long Time Scales by Transporting Value. arXiv:1810.06721 [cs.AI] https://arxiv.org/abs/1810.06721
- [8] Michael Janner, Qiyang Li, and Sergey Levine. 2021. Offline reinforcement learning as one big sequence modeling problem. Advances in neural information processing systems 34 (2021), 1273-1286.
- Michael Janner, Qiyang Li, and Sergey Levine. 2021. Offline Reinforcement Learning as One Big Sequence Modeling Problem. arXiv:2106.02039 [cs.LG] https://arxiv.org/abs/2106.02039
- [10] Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick van der Smagt. 2017. Deep Variational Bayes Filters: Unsupervised Learning of State Space Models from Raw Data. arXiv:1605.06432 [stat.ML] https://arxiv.org/abs/1605.06432
- [11] Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. 2019. Stabilizing off-policy q-learning via bootstrapping error reduction. Advances in neural information processing systems 32 (2019).
- [12] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative q-learning for offline reinforcement learning. Advances in Neural Information Processing Systems 33 (2020), 1179–1191.
- [13] Hemant Kumawat, Biswadeep Chakraborty, and Saibal Mukhopadhyay. 2024. RoboKoop: Efficient Control Conditioned Representations from Visual Input in Robotics using Koopman Operator. In 8th Annual Conference on Robot Learning. https://openreview.net/forum?id=NiA8hVdDS7
- [14] Hemant Kumawat, Biswadeep Chakraborty, and Saibal Mukhopadhyay. 2024. STAGE Net: Spatio-Temporal Attention-based Graph Encoding for Learning Multi-Agent Interactions in the presence of Hidden Agents. https://openreview. net/forum?id=tsj6rDzI0V
- [15] Hemant Kumawat, Biswadeep Chakraborty, and Saibal Mukhopadhyay. 2024. STEMFold: Stochastic Temporal Manifold for Multi-Agent Interactions in the Presence of Hidden Agents. arXiv:2401.14522 [cs.MA] https://arxiv.org/abs/ 2401.14522

- [16] Hemant Kumawat and Saibal Mukhopadhyay. 2022. Radar Guided Dynamic Visual Attention for Resource-Efficient RGB Object Detection. In 2022 International Joint Conference on Neural Networks (IJCNN). 1–8. https://doi.org/10.1109/ IJCNN55064.2022.9892184
- [17] Hemant Kumawat and Saibal Mukhopadhyay. 2024. AdaCred: Adaptive Causal Decision Transformers with Feature Crediting. arXiv:2412.15427 [cs.LG] https: //arxiv.org/abs/2412.15427
- [18] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems. arXiv:2005.01643 [cs.LG] https://arxiv.org/abs/2005.01643
- [19] Tianwei Ni, Michel Ma, Benjamin Eysenbach, and Pierre-Luc Bacon. 2023. When Do Transformers Shine in RL? Decoupling Memory from Credit Assignment. In Thirty-seventh Conference on Neural Information Processing Systems. https: //openreview.net/forum?id=APGXBNkt6h
- [20] Kruttidipta Samal, Hemant Kumawat, Priyabrata Saha, Marilyn Wolf, and Saibal Mukhopadhyay. 2022. Task-Driven RGB-Lidar Fusion for Object Tracking in Resource-Efficient Autonomous System. *IEEE Transactions on Intelligent Vehicles* 7, 1 (2022), 102–112. https://doi.org/10.1109/TIV.2021.3087664
- [21] Kruttidipta Samal, Hemant Kumawat, Marilyn Wolf, and Saibal Mukhopadhyay. 2022. A methodology for understanding the origins of false negatives in dnn based object detectors. In 2022 International Joint Conference on Neural Networks (IJCNN). IEEE, 1–8.
- [22] Jinghuan Shang, Kumara Kahatapitiya, Xiang Li, and Michael S Ryoo. 2022. Starformer: Transformer with state-action-reward representations for visual reinforcement learning. In *European conference on computer vision*. Springer, 462–479.
- [23] Richard S. Sutton and Andrew G. Barto. 2018. Reinforcement learning: An introduction. In *MIT press*.
- [24] Amit Ranjan Trivedi, Sina Tayebati, Hemant Kumawat, Nastaran Darabi, Divake Kumar, Adarsh Kumar Kosta, Yeshwanth Venkatesha, Dinithi Jayasuriya, Nethmi Jayasinghe, Priyadarshini Panda, Saibal Mukhopadhyay, and Kaushik Roy. 2025. Intelligent Sensing-to-Action for Robust Autonomy at the Edge: Opportunities and Challenges. arXiv:2502.02692 [cs.RO] https://arxiv.org/abs/2502.02692
- [25] Haoqing Wang, Xun Guo, Zhi-Hong Deng, and Yan Lu. 2022. Rethinking Minimal Sufficient Representation in Contrastive Learning. arXiv:2203.07004 [cs.CV] https://arxiv.org/abs/2203.07004
- [26] Yueh-Hua Wu, Xiaolong Wang, and Masashi Hamaya. 2024. Elastic decision transformer. Advances in Neural Information Processing Systems 36 (2024).
- [27] Taku Yamagata, Ahmed Khalil, and Raul Santos-Rodriguez. 2023. Q-learning Decision Transformer: Leveraging Dynamic Programming for Conditional Sequence Modelling in Offline RL. arXiv:2209.03993 [cs.LG] https://arxiv.org/abs/ 2209.03993
- [28] Taku Yamagata, Ahmed Khalil, and Raul Santos-Rodriguez. 2023. Q-learning decision transformer: Leveraging dynamic programming for conditional sequence modelling in offline rl. In *International Conference on Machine Learning*. PMLR, 38989–39007.
- [29] Yupei Yang, Biwei Huang, Fan Feng, Xinyue Wang, Shikui Tu, and Lei Xu. 2024. Towards Generalizable Reinforcement Learning via Causality-Guided Self-Adaptive Representations. arXiv preprint arXiv:2407.20651 (2024).
- [30] Kun Zhang and Peter Spirtes. 2011. Causal inference and causal explanation with background knowledge. Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence (2011), 581–588.
- [31] Qinqing Zheng, Amy Zhang, and Aditya Grover. 2022. Online Decision Transformer. arXiv:2202.05607 [cs.LG] https://arxiv.org/abs/2202.05607