



# Model-based trajectory stitching for improved behavioural cloning and its applications

Charles A. Hepburn<sup>1</sup> · Giovanni Montana<sup>2,3,4</sup>

Received: 1 December 2022 / Revised: 12 May 2023 / Accepted: 16 August 2023 /  
Published online: 27 September 2023  
© The Author(s) 2023

## Abstract

Behavioural cloning (BC) is a commonly used imitation learning method to infer a sequential decision-making policy from expert demonstrations. However, when the quality of the data is not optimal, the resulting behavioural policy also performs sub-optimally once deployed. Recently, there has been a surge in offline reinforcement learning methods that hold the promise to extract high-quality policies from sub-optimal historical data. A common approach is to perform regularisation during training, encouraging updates during policy evaluation and/or policy improvement to stay close to the underlying data. In this work, we investigate whether an offline approach to improving the quality of the existing data can lead to improved behavioural policies without any changes in the BC algorithm. The proposed data improvement approach - Model-Based Trajectory Stitching (MBTS) - generates new trajectories (sequences of states and actions) by ‘stitching’ pairs of states that were disconnected in the original data and generating their connecting new action. By construction, these new transitions are guaranteed to be highly plausible according to probabilistic models of the environment, and to improve a state-value function. We demonstrate that the iterative process of replacing old trajectories with new ones incrementally improves the underlying behavioural policy. Extensive experimental results show that significant performance gains can be achieved using MBTS over BC policies extracted from the original data. Furthermore, using the D4RL benchmarking suite, we demonstrate that state-of-the-art results are obtained by combining MBTS with two existing offline learning methodologies reliant on BC, model-based offline planning (MBOP) and policy constraint (TD3+BC).

---

Editors: Fabio Vitale, Tania Cerquitelli, Marcello Restelli, Charalampos Tsourakakis.

---

✉ Giovanni Montana  
g.montana@warwick.ac.uk

Charles A. Hepburn  
charlie.hepburn@warwick.ac.uk

<sup>1</sup> Mathematics Institute, University of Warwick, Coventry, England

<sup>2</sup> Department of Statistics, University of Warwick, Coventry, England

<sup>3</sup> WMG, University of Warwick, Coventry, England


<sup>4</sup> Alan Turing Institute, London, England

**Keywords** Offline reinforcement learning · Behavioural cloning · Imitation learning

## 1 Introduction

Behavioural cloning (BC) (Pomerleau, 1988, 1991) is one of the simplest imitation learning methods to obtain a decision-making policy from expert demonstrations. BC frames the imitation learning problem as a supervised learning one. Given expert trajectories - the expert's paths through the state space - a policy network is trained to reproduce the expert behaviour: for a given observation, the action taken by the policy must closely approximate the one taken by the expert. Although a simple method, BC has shown to be very effective across many application domains (Kadous et al., 2005; Pearce & Zhu, 2022; Pomerleau, 1988; Sammut et al., 1992), and has been particularly successful in cases where the dataset is large and has wide coverage (Codevilla et al., 2019). An appealing aspect of BC is that it is applied in an offline setting, using only the historical data. Unlike reinforcement learning (RL) methods, BC does not require further interactions with the environment. Offline policy learning can be advantageous in many circumstances, especially when collecting new data through interactions is expensive, time-consuming or dangerous; or in cases where deploying a partially trained, sub-optimal policy in the real-world may be unethical, e.g. in autonomous driving and medical applications.

BC extracts the behaviour policy which created the dataset. Consequently, when applied to sub-optimal data (i.e. when some or all trajectories have been generated by non-expert demonstrators), the resulting behavioural policy is also expected to be sub-optimal. This is due to the fact that BC has no mechanism to infer the importance of each state-action pair. Other drawbacks of BC are its tendency to overfit when giving a small number of demonstrations and the state distributional shift between training and test distributions (Codevilla et al., 2019; Ross et al., 2011). In the area of imitation learning, significant efforts have been made to overcome such limitations, however the available methodologies generally rely on interacting with the environment (Finn et al., 2016; Ho & Ermon, 2016; Le et al., 2018; Ross et al., 2011). So, a question arises: can we help BC infer a superior



**Fig. 1** Simplified illustration of Model-Based Trajectory Stitching. Each original trajectory (a sequence of states and actions) in the dataset  $\mathcal{D}$  is indicated as  $\mathcal{T}_i$  with  $i = 1, \dots, 3$ . A first stitching event is seen in trajectory  $\mathcal{T}_1$  whereby a transition to a state originally visited in  $\mathcal{T}_2$  takes place. A second stitching event involves a jump to a state originally visited in  $\mathcal{T}_3$ . At each event, jumping to a new state increases the current trajectory's future expected returns. The resulting trajectory (in bold) consists of a sequence of states, all originally visited in  $\mathcal{D}$ , but connected by imagined actions; it replaces  $\mathcal{T}_1$  in the new dataset

policy only from available sub-optimal data without the need to collect additional expert demonstrations?

Our investigation is related to the emerging body of work on offline RL, which is motivated by the aim of inferring expert policies with only a fixed set of sub-optimal data (Lange et al., 2012; Levine et al., 2020). A major obstacle towards this aim is posed by the notion of *action distributional shift* (Fujimoto et al., 2019; Kumar et al., 2019; Levine et al., 2020). This is introduced when the policy being optimised deviates from the behaviour policy, and is caused by the action-value function overestimating out-of-distribution (OOD) actions. A number of existing methods address the issue by constraining the actions that can be taken. In some cases, this is achieved by constraining the policy to actions close to those in the dataset (Fujimoto et al., 2019; Fujimoto & Gu, 2021; Jaques et al., 2019; Kumar et al., 2019; Wu et al., 2019; Zhou et al., 2020), or by manipulating the action-value function to penalise OOD actions (Kumar et al., 2020; An et al., 2021; Kostrikov et al., 2021; Yu et al., 2021). In situations where the data is sub-optimal, offline RL has been shown to recover a superior policy to BC (Fujimoto et al., 2019; Kumar et al., 2022). Improving BC will in turn improve many offline RL policies that rely on an explicit behaviour policy of the dataset (Argenson & Dulac-Arnold, 2020; Fujimoto & Gu, 2021; Zhan et al., 2021).

In contrast to existing offline learning approaches, we turn the problem on its head: rather than trying to regularise or constrain the policy somehow, we investigate whether the data quality itself can be improved using only the available demonstrations. To explore this avenue, we propose a model-based data improvement method called Model-Based Trajectory Stitching (MBTS). Our ultimate aim is to develop a procedure that identifies sub-optimal trajectories and replaces them with better ones. New trajectories are obtained by stitching existing ones together, without the need to generate unseen states. The proposed strategy consists of replaying each existing trajectory in the dataset: for each state-action pair leading to a particular next state along a trajectory, we ask whether a different action could have been taken instead, which would have landed at a different seen state from a different trajectory. An actual jump to the new state only occurs when generating such an action is plausible and it is expected to improve the quality of the original trajectory - in which case we have a *stitching event*.

An illustrative representation of this procedure can be seen in Fig. 1, where we assume to have at our disposal only three historical trajectories. In this example, a trajectory has been improved through two stitching events. To determine the stitching points, MBTS uses a probabilistic view of state-reachability that depends on learned dynamics models of the environment. These models are evaluated only on in-distribution states enabling accurate prediction. In order to assess the expected future improvement introduced by a potential stitching event, we utilise a state-value function and a reward model. Thus, MBTS can be thought of as a data-driven, automated procedure yielding highly plausible and higher-quality demonstrations to facilitate supervised learning; at the same time, sub-optimal demonstrations are removed altogether whilst keeping the diverse set of seen states.

Our experimental results show that MBTS produces higher-quality data, with BC-derived policies always superior than those inferred on the original data. Remarkably, we demonstrate that MBTS-augmented data allow BC to compete with state-of-the-art offline RL algorithms on highly complex continuous control openAI gym tasks implemented in MuJoCo using the D4RL offline benchmarking suite (Fu et al., 2020). Furthermore, we show that integrating MBTS with existing offline learning methods that explicitly use BC such as model-based planning (Argenson & Dulac-Arnold, 2020) and TD3+BC (Fujimoto & Gu, 2021) can significantly boost their performance.

## 2 Related work

### 2.1 Imitation learning

Imitation learning aims to emulate a policy from expert demonstrations (Hussein et al., 2017). BC is the simplest of such category of methods and uses supervised learning to clone the actions in the dataset. BC is a powerful method and has been used successfully in many applications such as learning a quadrotter to fly (Giusti et al., 2015), self-driving cars (Bojarski et al., 2016; Farag & Saleh, 2018) and games (Pearce & Zhu, 2022). These application are highly complex and shows accurate policy estimation from high-quality offline data.

One drawback from using BC is the state distributional shift between training and test distributions. Improved imitation learning methods have been introduced to reduce this distributional shift, however they usually require online exploration. For instance, DAgger (Ross et al., 2011) is an online learning approach that iteratively updates a deterministic policy; it addresses the state distributional shift problem of BC through an on-policy method for data collection; similarly to MBTS, the original dataset is augmented, but this involves online interactions. Another algorithm, GAIL (Ho & Ermon, 2016), iteratively updates a generative adversarial network (Goodfellow et al., 2014) to determine whether a state-action pair can be deemed as expert; a policy is then inferred using a trust region policy optimisation step (Schulman et al., 2015). MBTS also uses generative modelling, but this is to create data points likely to have come from the data that connect high-value regions.

While expert demonstrations are crucial for imitation learning, our MBTS approach generates higher quality datasets from existing, potentially sub-optimal data, thereby enhancing offline policy learning. Furthermore, MBTS leverages a reward function to learn an improved policy, which distinguishes it from the imitation learning setting where access to rewards may not always be available. This key difference enables MBTS to deliver better performance in certain scenarios compared to traditional imitation learning methods.

### 2.2 Offline reinforcement learning

Offline RL aims to learn an optimal policy from sub-optimal datasets without further interactions with the environment (Lange et al., 2012; Levine et al., 2020). Similarly to BC, offline RL suffers from distributional shift. However this shift comes from the policy selecting OOD actions leading to overestimation of the value function (Fujimoto et al., 2019; Kumar et al., 2019). In the online setting, this overestimation encourages the agent to explore, but offline this leads to a compounding of errors where the agent believes OOD actions lead to high returns. Many offline RL algorithms bias the learned policy towards the behaviour-cloned one (Argenson & Dulac-Arnold, 2020; Fujimoto & Gu, 2021; Zhan et al., 2021) to ensure the policy does not deviate too far from the behaviour policy. Many of these offline methods are therefore expected to directly benefit from enhanced datasets yielding higher-achieving behavioural policies.

### 2.2.1 Model-free methods

Many model-free offline RL methods typically deal with distributional shift either by regularising the policy to stay close to actions given in the dataset (Fujimoto & Gu, 2021; Fujimoto et al., 2019; Jaques et al., 2019; Kumar et al., 2019; Wu et al., 2019; Zhou et al., 2020) or by pessimistically evaluating the Q-value to penalise OOD actions (An et al., 2021; Kostrikov et al., 2021; Kumar et al., 2020). Both options involve explicitly or implicitly capturing information about the unknown underlying behaviour policy. This behaviour policy can be fully captured using BC. For instance, batch-constrained Q-learning (BCQ) (Fujimoto et al., 2019) is a policy constraint method which uses a variational autoencoder to generate likely actions in order to constrain the policy. The TD3+BC algorithm (Fujimoto & Gu, 2021) offers a simplified policy constraint approach; it adds a behavioural cloning regularisation term to the policy update biasing actions towards those in the dataset. Alternatively, conservative Q-learning (CQL) (Kumar et al., 2020) adjusts the value of the state-action pairs to “push down” on OOD actions and “push up” on in-distribution actions. CQL manipulates the value function so that OOD actions are discouraged and in-distribution actions are encouraged. Implicit Q-learning (IQL) (Kostrikov et al., 2021) avoids querying OOD actions altogether by manipulating the Q-value to have a state-value function in the SARSA-style update. All the above methods try to directly deal with OOD actions, either by avoiding them or safely handling them in either the policy improvement or evaluation step. In contrast, our method rethinks the problem of learning from sub-optimal data. Rather than using RL to learn a policy, instead we use RL-based approaches to enrich the data enabling BC to extract an improved policy. Our method generates unseen actions between in-distribution states; by doing so, we avoid distributional shift by evaluating a state-value function only on seen states.

### 2.2.2 Model-based methods

Model-based algorithms rely on an approximation of the environment’s dynamics (Janner et al., 2019; Sutton, 1991), that is probability distributions where the next state and reward are predicted from a current state and action. In the online setting, model-based methods tend to improve sample efficiency (Buckman et al., 2018; Chua et al., 2018; Feinberg et al., 2018; Janner et al., 2019; Kalweit & Boedecker, 2017). In an offline learning context, the learned dynamics have been exploited in various ways.

One approach consists of using the models to improve the policy learning. For instance, Model-based offline RL (MOREL) (Kidambi et al., 2020) is an algorithm which constructs a pessimistic Markov Decision Model (P-MDP), based off a learned forward dynamics model and a state-action detector. The P-MDP is given an additional absorbing state, which gives large negative reward for unknown state-actions. Model-based Offline policy Optimization (MOPO) (Yu et al., 2020) augments the dataset by performing rollouts using a learned, uncertainty-penalised, MDP. Unlike MOPO, MBTS does not introduce imagined states, but only actions between reachable unconnected states.

Another opportunity to exploit learnt models of the environment is in decision-time planning. Model-based offline planning (MBOP) (Argenson & Dulac-Arnold, 2020) uses the learnt environment dynamics and a BC policy to roll-out a trajectory from the current state, one transition at a time. The best trajectory from the current state is found where the

trajectory horizon has been extended using a value function and the first action is selected. This process is repeated for each new state. Model-based offline planning with trajectory pruning (MOPP) (Zhan et al., 2021) extends the MBOP idea, but prunes the trajectory roll-outs based on an uncertainty measure, safely handling the problem of distributional shift. Diffuser (Janner et al., 2022) uses a diffusion probabilistic model to predict a whole trajectory in one step. Rather than using a model to predict a single next state at decision-time, diffuser can generate unseen trajectories that have high likelihood under the data and maximise the cumulative rewards of a trajectory ensuring long-horizon accuracy. However, diffuser’s individual plans are very slow which limits its use case for real-world applications. Our MBTS method can be used in direct conjunction with planning, especially with MBOP and MOPP, which both use a BC policy to guide the trajectory sampling.

### 2.3 State similarity metrics

A central aspect of the proposed MBTS approach consists of a *stitching event*, which uses a notion of state similarity to determine whether two states are “close” together. Relying on only geometric distances would often be inappropriate; e.g. two states may be close in Euclidean distance, yet reaching one from another may be impossible (e.g. in navigation task environments where walls or other obstacles preclude reaching a nearby state). Bisimulation metrics (Ferns et al., 2004) capture state similarity based on the dynamics of the environment. These have been used in RL mainly for system state aggregation (Ferns et al., 2012; Kemertas et al., 2021; Zhang et al., 2020); they are expensive to compute (Chen et al., 2012) and usually require full-state enumeration (Bacci et al., 2013a, b; Dadashi et al., 2021). A scalable approach for state-similarity has recently been introduced by using a pseudometric (Castro, 2020) which facilitates the calculation of state-similarity in offline RL. PLOFF (Dadashi et al., 2021) is an offline RL algorithm that uses a state-action pseudometric to bias the policy evaluation and improvement steps. Whereas PLOFF uses a pseudometric to stay close to the dataset, we bypass this notion altogether by only using states in the dataset and generating unseen actions connecting them. Our stitching event is based from the decomposition of the trajectory distribution which allows us to pick unseen actions, but with high likelihood, determined by the future state.

### 2.4 Data re-sampling and augmentation approaches

In offline RL, data re-sampling strategies aim to only learn from high-performing transitions. For instance, best-action imitation learning (BAIL) (Chen et al., 2020) imitates state-action pairs based from the upper-envelope of the dataset. Monotonic Advantage Re-Weighted Imitation Learning (MARWIL) (Wang et al., 2018) weights state-action pairs from an exponentially-weighted advantage function during policy learning by BC. Return-based data re-balance (ReD) (Yue et al., 2022) re-samples the data based from the trajectory returns and then applies offline reinforcement learning methods. The proposed MBTS differs from BAIL, MARWIL and ReD as we increase the dataset by adding impactful stitching transitions as well as removing the low-quality transitions. MBTS has the effect of re-sampling high-value transitions in the trajectory as well supplementing the dataset with stitched transitions, connecting high-value regions.

Best action trajectory stitching (BATS) (Char et al., 2022) is a related trajectory stitching method: it augments the dataset by adding transitions through model-based planning. MBTS differs from BATS in a number of fundamental ways. First, BATS takes a geometric

approach to defining state similarity; state-actions are rolled-out using the dynamics model until a state is found that is within a short distance of a state in the dataset. Relying exclusively on geometric distances may result in poor results; as such, our stitching events are based on the dynamics of the environment and are only assessed between two in-distribution states. Second, BATS generates new states that are not in the dataset. Due to compounding model error, resulting in unlikely rollouts, the rewards are penalised for the generated transitions which favours state-action pairs within the dataset. In contrast, we only allow one-step stitching between in-distribution states and use the value function to extend the horizon rather than a learned model. Finally, BATS adds all stitched actions to the original dataset, then create a new dataset by running value iteration, which is eventually used to learn a policy through BC. In contrast, our MBTS method has been designed to be more directly suited to policy learning through BC: since the lower-value experiences have been removed through stitching events, the resulting dataset contains only high-quality trajectories to learn from.

### 3 Methods

#### 3.1 Problem setup

We consider the offline RL problem setting, which consists of finding an optimal decision-making policy from a fixed dataset. The policy is a mapping from states to actions,  $\pi : \mathcal{S} \rightarrow \mathcal{A}$ , whereby  $\mathcal{S}$  and  $\mathcal{A}$  are the state and action spaces, respectively. The dataset is made up of transitions  $\mathcal{D} = \{(s_t, a_t, r_t, s_{t+1})\}$  that include the current state,  $s_t$ , the action performed in that state,  $a_t$ , the next state after the action has been taken,  $s_{t+1}$ , and the reward resulting for transitioning,  $r_t$ . The actions are assumed to follow an unknown behavioural policy,  $\pi_\beta$ , acting in a Markov decision process (MDP). The MDP is defined as  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ , where  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the transition probability function which defines the dynamics of the environment,  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  is the reward function and  $\gamma \in (0, 1]$  is a scalar discount factor (Sutton & Barto, 1998).

In offline RL, the agent must learn a policy,  $\pi(a_t | s_t)$ , that maximises the returns defined as the expected sum of discounted rewards,  $\mathbb{E}_\pi [\sum_{t=0}^{\infty} r_t \gamma^t]$ , without ever having access to  $\pi_\beta$ . Here we are interested in performing imitation learning through BC, which mimics  $\pi_\beta$  by performing supervised learning on the state-action pairs in  $\mathcal{D}$  (Pomerleau, 1988, 1991). More specifically, BC finds a deterministic policy,

$$\pi^{\text{BC}}(s) = \arg \min_{\pi} \mathbb{E}_{s_t, a_t \sim \mathcal{D}} [(\pi(s_t) - a_t)^2].$$

This solution is known to minimise the KL-divergence between  $\pi_\beta$  and the trajectory distributions of the learned policy Ke et al. (2020). Our objective is to enhance the dataset, such that it has the effect of being collected by an improved behaviour policy. Thus, training a policy by BC on the improved dataset will lead to higher returns than  $\pi_\beta$ .

#### 3.2 Model-based trajectory stitching

Under our modelling assumptions, the probability distribution of any given trajectory  $T = (s_0, a_0, s_1, a_1, s_2, \dots, s_H)$  in  $\mathcal{D}$  can be expressed as

$$p(\mathcal{T}) = p(s_0) \prod_{t=1}^H p(a_t | s_t) p(s_{t+1} | s_t, a_t). \quad (1)$$

where  $p(a_t | s_t)$  is the policy and  $p(s_{t+1} | s_t, a_t)$  is the environment's dynamics. First, we note that, in the offline case, Eq. (1) can be re-written in an alternative, but equivalent form as

$$p(\mathcal{T}) = p(s_0) \prod_{t=1}^H p(s_{t+1} | s_t) p(a_t | s_t, s_{t+1}), \quad (2)$$

which now depends on two different conditional distributions:  $p(s_{t+1} | s_t)$ , the environment's forward dynamics, and  $p(a_t | s_t, s_{t+1})$ , its inverse dynamics. Both distributions can be approximated using the available data,  $\mathcal{D}$  (see Section 3.3). We also pre-train a state-value function  $V_{\pi_\beta}$  to estimate the future expected sum of rewards for being in a state  $s$  following the behaviour policy  $\pi_\beta$  as well as a reward function (see Section 3.4), which will be used to predict  $r(s_t, \hat{a}_t, s_{t+1})$  for any action  $\hat{a}_t$  not in  $\mathcal{D}$ .

Equation (2) informs our data-improvement strategy, as follows. For a given transition,  $(s_t, a_t, s_{t+1}) \in \mathcal{D}$ , our aim is to replace  $s_{t+1}$  with  $\hat{s}_{t+1} \in \mathcal{D}$  using a synthetic connecting action  $\hat{a}_t$ . A necessary condition for such a state swap to occur is that  $\hat{s}_{t+1}$  should be plausible, conditional on  $s_t$ , according to the learnt forward dynamic model,  $p(s_{t+1} | s_t)$ . Furthermore, such a state swap should only happen when landing on  $\hat{s}_{t+1}$  leads to higher expected returns. Accordingly, two criteria need to be satisfied in order to allow swapping states:  $p(\hat{s}_{t+1} | s_t) \geq p(s_{t+1} | s_t)$  and  $V_{\pi_\beta}(\hat{s}_{t+1}) > V_{\pi_\beta}(s_{t+1})$ . The first criterion ensures that the new next state must be at least as likely to have been observed as the candidate state under the learnt dynamic model. Furthermore, to be beneficial, the candidate next state must not only be likely to be reached from  $s_t$  under the environment dynamics, but must also lead to higher expected returns compared to the current  $s_{t+1}$ . This requirement is captured by the second criterion using the pre-trained value function. In practice, finding a suitable candidate  $\hat{s}_{t+1}$  involves a search for candidate next states amongst all the states that has been visited by any trajectory in  $\mathcal{D}$  (see Section 3.3). Where the two criteria above are satisfied, a plausible action connecting  $s_t$  and the newly found  $\hat{s}_{t+1}$  is obtained by generating an action that maximises the learnt inverse dynamics model. In summary, we have:

**Definition 1** A candidate *stitching event* consists of a transition  $(s_t, \hat{a}_t, \hat{s}_{t+1}, r(s_t, \hat{a}_t, \hat{s}_{t+1}))$  that replaces  $(s_t, a_t, s_{t+1}, r(s_t, a_t, s_{t+1}))$  and it is such that, starting from  $s_t$ , the new state satisfies

$$\hat{s}_{t+1} = \arg \max_{s_{t+1} \in \mathcal{D}} V_{\pi_\beta}(s_{t+1}) \quad \text{s.t. } p(\hat{s}_{t+1} | s_t) > p(s_{t+1} | s_t)$$

and the new action is generated by

$$\hat{a}_t = \arg \max_{\hat{a}} p(\hat{a} | s_t, \hat{s}_{t+1}).$$

For every trajectory in the dataset, starting from the initial state, we sequentially identify candidate stitching events. For instance, in Fig. 1, two such events have been identified along the  $\mathcal{T}_1$  trajectory and eventually they yield a new trajectory,  $\hat{\mathcal{T}}_1$ . When the cumulative sum of rewards along the newly formed trajectory are higher than those



observed in the original trajectory, the old trajectory is replaced by the new one in  $\mathcal{D}$ . This is captured by the following definition.

**Definition 2** A *trajectory replacement event* is such that, if a new trajectory  $\hat{\mathcal{T}}$  started at the initial state  $s_0$  of  $\mathcal{T}$  has been compiled after a sequence of candidate stitching events, then  $\hat{\mathcal{T}}$  replaces  $\mathcal{T}$  in  $\mathcal{D}$  when the following condition is satisfied:

$$(1 + \tilde{p}) \sum_{t \in \mathcal{T}} r_t < \sum_{u \in \hat{\mathcal{T}}} r_u.$$

In this definition,  $\tilde{p}$  is a small positive constant and the  $(1 + \tilde{p})$  terms ensures that the cumulative sum of returns in the new trajectory improves upon the old one by a given margin. This conservative approach takes into account potential prediction errors incurred by using the learnt reward model when assessing the rewards for  $\hat{\mathcal{T}}$ .

The procedure above is repeated for all the trajectories in the current dataset. When any of the original trajectories are replaced by new ones, a new and improved dataset is formed. The new dataset can then be thought of as being collected by a different, and improved, behaviour policy. Using the new data, the value function is trained again, and a search for trajectory replacement events is started again. This iterative procedure is summarised below.

**Definition 3** *Trajectory Stitching* is an iterative process whereby every trajectory in a dataset  $\mathcal{D}$  may be entirely replaced by a new one formed through trajectory replacement events. When such replacements take place, resulting in a new dataset, an updated value function is inferred and the process is repeated again.

The trajectory stitching method enforces a greedy next state selection policy (Definition 1) and guarantees that the trajectories produced by this policy have higher returns than under the previous policy (Definition 2). Therefore, we obtain a new dataset (Definition 3) collected under a new behaviour policy for which a new value function can be learned and the trajectory stitching process can be repeated. This iterative data improvement process is terminated when no more trajectory replacements are possible, or earlier (see Section 4).

The MBTS approach is highly versatile and can be implemented in a variety of ways. In the remainder of this section, we describe our chosen methods for modeling the two probability distributions featured in Eq. (1), as well as our techniques for estimating the state-value function and predicting the environment's rewards. Our discussion is primarily motivated by continuous control problems, focusing on continuous state and action space domains. For applications in discrete spaces, alternative models should be employed.

### 3.3 Candidate next state search

The search for a candidate next state requires a learned forward dynamics model, i.e.  $p(s_{t+1} | s_t)$ . Model-based RL approaches typically use such dynamics' models conditioned on the action as well as the state to make predictions (Janner et al., 2019; Yu et al., 2020; Kidambi et al., 2020; Argenson & Dulac-Arnold, 2020). Here, we use the model differently, only to guide the search process and identify a suitable next state to transition to. Specifically, conditional on  $s_t$ , the dynamics model is used to assess the relative likelihood of observing any other  $s_{t+1}$  in the dataset compared to the observed one. The environment

dynamics is assumed to follow a Gaussian distribution whose mean vector and covariance matrix are approximated by a neural network, i.e.

$$\hat{p}_{\xi}(s_{t+1} | s_t) = \mathcal{N}(\mu_{\xi_1}(s_t), \Sigma_{\xi_2}(s_t))$$

where  $\xi = (\xi_1, \xi_2)$  indicate the parameters of the neural network. This modelling assumption is fairly common in applications involving continuous state spaces (Janner et al., 2019; Kidambi et al., 2020; Yu et al., 2020, 2021).

In our implementation, we take an ensemble of  $N$  Gaussian models,  $\mathcal{E}$ ; each component of  $\mathcal{E}$  is characterised by its own parameter set,  $(\mu_{\xi_1}, \Sigma_{\xi_2})$ . This approach has been shown to take into account epistemic uncertainty, i.e. the uncertainty in the model parameters (Argenson & Dulac-Arnold, 2020; Buckman et al., 2018; Chua et al., 2018; Yu et al., 2021). Each individual model's parameter vector is estimated via maximum likelihood by optimising

$$\mathcal{L}_{\hat{p}}(\xi) = \mathbb{E}_{s_t, s_{t+1} \sim \mathcal{D}} [(\mu_{\xi_1}(s_t) - s_{t+1})^T \Sigma_{\xi_2}^{-1}(s_t) (\mu_{\xi_1}(s_t) - s_{t+1}) + \log |\Sigma_{\xi_2}(s_t)|],$$

where  $|\cdot|$  is the determinant of a matrix, and each model's parameter set is initialised differently prior to estimation. Upon fitting the models, a state  $s_{t+1}$  is replaced by  $\hat{s}_{t+1}$  only when

$$\min_{i \in \mathcal{E}} \hat{p}_{\xi_i}(\hat{s}_{t+1} | s_t) > \text{mean}_{i \in \mathcal{E}} \hat{p}_{\xi_i}(s_{t+1} | s_t).$$

In this context, we adopt a conservative approach, as we have greater confidence in the likelihood prediction of observed state-next state pairs,  $\hat{p}_{\xi_i}(s_{t+1} | s_t)$ , compared to unseen state-next state pairs,  $\hat{p}_{\xi_i}(\hat{s}_{t+1} | s_t)$ .

Performing this search over all next states in the dataset can be computationally inefficient. To address this, we assume that states that are far apart in Euclidean distance are not reachable, and therefore, we only evaluate these models on nearby states. We employ a nearest neighbors search organized by a KD-tree, as shown in line 7 of Algorithm 1, with the complete procedure detailed in the Appendix.

### 3.4 Value and reward function estimation

Value functions are widely used in reinforcement learning to determine the quality of an agent's current position (Sutton & Barto, 1998). In our context, we use a state-value function to assess whether a candidate next state offers a potential improvement over the original next state. To accurately estimate the future returns given the current state, we calculate a state-value function dependent on the behaviour policy of the dataset. The function  $V_{\theta}(s)$  is approximated by a MLP neural network parameterised by  $\theta$ . The parameters are learned by minimising the squared Bellman error (Sutton & Barto, 1998),

$$\mathcal{L}_V(\theta) = \mathbb{E}_{s_t, r_t, s_{t+1} \sim \mathcal{D}} [(r_t + \gamma V_{\theta}(s_{t+1}) - V_{\theta}(s_t))^2]. \quad (3)$$

In our context,  $V_{\theta}$  is only used to observe the value of in-distribution states, thus avoiding the OOD issue when evaluating value functions which occurs in offline RL. The value function will only be queried once to determine whether a candidate stitching event has been found (Definition 1).

Value functions require rewards for training, therefore a reward must be estimated for unseen tuples  $(s_t, \hat{a}_t, \hat{s}_{t+1})$ . There are many different modelling choices available; e.g., under

a Gaussian model, the mean and variance of the reward can be estimated allowing uncertainty quantification. Other alternatives include a Wasserstein-GAN, a VAE, and a standard multilayer neural network. In practice, the impact of the specific reward model and its effects when used for MBTS appears negligible (e.g. see Section 4.4.1). In the remainder of this section, we provide further details for one such model, based on Wasserstein-GAN (Arjovsky et al., 2017; Goodfellow et al., 2014), which we have extensively used in all our experiments (Section 4) and in our early investigations (Hepburn & Montana, 2022).

Wasserstein-GANs consist of a generator,  $G_\phi$  and a discriminator  $D_\psi$ , with parameters of the neural networks  $\phi$  and  $\psi$  respectively. The discriminator takes in the state, action, reward, next state and determines whether this transition is from the dataset. The generator loss function is:

$$\mathcal{L}_G(\phi) = \mathbb{E}_{\substack{z \sim p(z) \\ s_t, a_t, s_{t+1} \sim \mathcal{D} \\ \tilde{r}_t \sim G_\phi(z, s_t, a_t, s_{t+1})}} [D_\psi(s_t, a_t, s_{t+1}, \tilde{r}_t)].$$

Here  $z \sim p(z)$  is a noise vector sampled independently from  $\mathcal{N}(0, 1)$ , the standard normal. The discriminator loss function is:

$$\mathcal{L}_D(\psi) = \mathbb{E}_{s_t, a_t, r_t, s_{t+1} \sim \mathcal{D}} [D_\psi(s_t, a_t, s_{t+1}, r_t)] - \mathbb{E}_{\substack{z \sim p(z) \\ s_t, a_t, s_{t+1} \sim \mathcal{D} \\ \tilde{r}_t \sim G_\phi(z, s_t, a_t, s_{t+1})}} [D_\psi(s_t, a_t, s_{t+1}, \tilde{r}_t)].$$

Once trained, a reward will be predicted for the stitching event when a new action has been generated between two previously disconnected states.

### 3.5 Action generation

Sampling a suitable action that leads from  $s_t$  to the newly found state  $\hat{s}_{t+1}$  requires an inverse dynamics model. Specifically, we require that a synthetic action must maximise the estimated conditional density,  $p(a_t | s_t, \hat{s}_{t+1})$ . Given our requirement of sampling synthetic actions, a conditional variational autoencoder (CVAE) (Kingma & Welling, 2013; Sohn et al., 2015) provides a suitable approximation for the inverse dynamics model. The CVAE consists of an encoder  $q_{\omega_1}$  and a decoder  $p_{\omega_2}$  where  $\omega_1$  and  $\omega_2$  are the respective parameters of the neural networks.

The encoder maps the input data onto a lower-dimensional latent representation  $z$  whereas the decoder generates data from the latent space. We train a CVAE to maximise the conditional marginal log-likelihood,  $\log p(a_t | s_t, \hat{s}_{t+1})$ . While intractable in nature, the CVAE objective enables us to maximize the variational lower bound instead,

$$\begin{aligned} \max_{\omega_1, \omega_2} \log p(a_t | s_t, \hat{s}_{t+1}, z) &\geq \max_{\omega_1, \omega_2} \mathbb{E}_{z \sim q_{\omega_1}} [\log p_{\omega_2}(a_t | s_t, \hat{s}_{t+1}, z)] \\ &\quad - D_{\text{KL}}[q_{\omega_1}(z | a_t, s_t, \hat{s}_{t+1}) || P(z | s_t, \hat{s}_{t+1})], \end{aligned}$$

where  $z \sim \mathcal{N}(0, 1)$  is the prior for the latent variable  $z$ , and  $D_{\text{KL}}$  represents the KL-divergence (Kullback & Leibler, 1951; Kullback, 1997). To generate an action between two unconnected states,  $s_t$  and  $\hat{s}_{t+1}$ , we use the decoder  $p_\omega$  to sample from  $p(a_t | s_t, \hat{s}_{t+1})$ . This process ensures that the most plausible action is generated conditional on  $s_t$  and  $\hat{s}_{t+1}$ .

**Algorithm 1** Model-based Trajectory Stitching

**Initialise:** An action generator  $p_{\omega_1}$ , a reward generator  $G_\phi$ , an ensemble of dynamics models  $\{\hat{p}_{\xi^i}(s' | s)\}_{i=1}^N$ , an acceptance threshold  $\tilde{p}$ , and a dataset  $\mathcal{D}_0$  made up of  $T$  trajectories  $(\mathcal{T}_1, \dots, \mathcal{T}_T)$

```

1: for  $k = 0, \dots, K$  do
2:   Train state-value function,  $V$  on  $\mathcal{D}_k$  by minimising Eq. (3).
3:   for  $t = 1, \dots, T$  do
4:     Select  $s, s' = s_0, s'_0 \in \mathcal{T}_t$ 
5:     Initialise new trajectory,  $\hat{\mathcal{T}}_t$ 
6:     while not done do
7:       Create set of candidate next states from dataset,
            $\{\hat{s}'_j\}_{j=1}^N \sim \mathcal{D}_k$ 
8:       Evaluate dynamics models for new set of states and take
           minimum,  $\min_i \hat{p}_{\xi^i}(\hat{s}'_j | s)$ 
9:       if  $\min_i \hat{p}_{\xi^i}(\hat{s}'_j | s) > \text{mean}_i \hat{p}_{\xi^i}(s' | s)$ ,  $V(\hat{s}'_j) = \max_i V(\hat{s}'_i)$  and
            $V(\hat{s}'_j) > V(s')$  then
10:        Generate a new action and reward,
             $\tilde{a} \sim p_{\omega_1}(z, s, \hat{s}'_j)$ ,  $\tilde{r} \sim G_\phi(z, s, \tilde{a}, \hat{s}'_j)$ 
11:        Add  $(s, \tilde{a}, \tilde{r}, \hat{s}'_j)$  to new trajectory  $\hat{\mathcal{T}}_t$ 
12:        Set  $s = \hat{s}'_j$ 
13:       else
14:        Add original transition,  $(s, a, r, s')$  to the new trajectory  $\hat{\mathcal{T}}_t$ 
15:        Set  $s = s'$ 
16:       end if
17:     end while
18:     if  $\sum_{i \in \hat{\mathcal{T}}_t} r_i > (1 + \tilde{p}) * \sum_{j \in \mathcal{T}_t} r_j$  then
19:        $\hat{\mathcal{T}}_t = \hat{\mathcal{T}}_t$ 
20:     else
21:        $\hat{\mathcal{T}}_t = \mathcal{T}_t$ 
22:     end if
23:   end for
24:   Collect trajectories into dataset,  $\mathcal{D}_{k+1} = (\hat{\mathcal{T}}_1, \dots, \hat{\mathcal{T}}_T)$ 
25: end for

```

## 4 Experimental results

In this section we first investigate whether MBTS can improve the quality of existing datasets for the purpose of inferring decision-making policies through BC in an offline fashion, without collecting any more data from the environment. Furthermore, we show that MBTS can help existing methods that explicitly use a BC term for offline learning to achieve higher performance. Specifically, we explore the use of MBTS in combination

with two algorithms: model-based offline planning (MBOP) (Argenson & Dulac-Arnold, 2020), which uses an explicit BC policy to select new actions, and TD3+BC (Fujimoto & Gu, 2021), which has an explicit BC policy constraint. Our experiments rely on the D4RL datasets, a collection of commonly used benchmarking tasks, and include comparisons with selected offline RL methods. These comparisons provide an insight into the potential gains that can be achieved when MBTS is combined with BC-based algorithms, which often reach or even improve upon current state-of-the-art performance levels in offline RL. In Section 4.2, we show empirically that even with a small amount of expert data, the MBTS+BC policies become closer to the expert policy, in KL divergence. In all experiments, we run MBTS for five iterations; these have been found to be sufficient to increase the quality of the data without being overly computationally expensive (Section 4.3). Finally we provide ablation studies into the choice of reward model, as well as alternative extraction policies to BC.

#### 4.1 Performance assessment on D4RL data

We compare our MBTS method on the D4RL (Fu et al., 2020) benchmarking datasets of the openAI gym MuJoCo tasks. Three complex continuous environments are tested - Hopper, Halfcheetah and Walker2d - each with different levels of difficulty. The “medium” datasets were gathered by the original authors using a single policy produced from the early-stopping of an agent trained by soft actor-critic (SAC) (Haarnoja et al., 2018a, b). The “medium-replay” datasets are the replay buffers from the training of the “medium” policies. The “expert” datasets were obtained from a policy trained to an expert level, and the “medium-expert” datasets are the combination of both the “medium” and “expert” datasets. A BC-cloned policy that used a MBTS dataset is denoted by MBTS+BC. All results and comparisons are summarised in Table 1 and detailed explanations of our methods are in order. We run MBTS for 3 different seeds, giving 3 datasets, we then train BC over 5 seeds for each new dataset giving 15 MBTS +BC policies.

##### 4.1.1 Behaviour cloning: MBTS+BC

The first method we investigate using MBTS with on the D4RL datasets is BC. Enriching the dataset with more high-value transitions and removing low quality ones leaves the dataset with closer-to-expert trajectories making BC the most suitable policy extraction algorithm. From Table 1 we can see that MBTS+BC improves over BC in all cases, showing that MBTS creates a higher quality dataset as claimed.

##### 4.1.2 Model-based offline planning: MBTS+MBOP


Given previously presented evidence that MBTS improves over BC, a natural next step is to investigate whether MBTS can also improve on other methods that are reliant on BC. Model-based offline planning (MBOP) (Argenson & Dulac-Arnold, 2020) is an offline model-based planning method that uses a BC policy to rollout multiple trajectories picking the action that leads to the trajectory with highest returns. For this study, we alter MBOP slightly to obtain MBTS+MBOP: in this version, actions are selected using our MBTS extracted policy and we use our trained value function.

As can be observed in Table 1, MBTS+MBOP improves over the MBOP baseline in all cases. We also compare MBTS+MBOP to state-of-the-art model-based algorithms such as a MOPO (Yu et al., 2020), MOREL (Kidambi et al., 2020) and Diffuser (Janner et al., 2022); in these comparisons, MBTS+MBOP achieves higher performance in 5 out of the 9 comparable tasks. Only in the hopper medium and medium-replay tasks does another model-based method outperform MBTS+MBOP.

#### 4.1.3 Model-free offline RL: TD3+ MBTS+BC

We also investigate the benefits of using MBTS in conjunction with a model-free offline RL algorithm. TD3+BC (Fujimoto & Gu, 2021) explicitly using BC in the policy improvement step to regularise the policy to take actions close to the dataset. As MBTS removes low-quality data, the learned Q-values will be inaccurate when trained solely on the new MBTS data. To counter this, we warm start TD3+BC on the original dataset, then use the new MBTS data to fine-tune both the critic and actor after the Q-values have been sufficiently trained. To keep this a fair comparison, we train the policy over the same number of iterations as reported in Fujimoto and Gu (2021). We make one small amendment to the Walker2d medium-replay dataset where we train the critic only using the original data, and use the MBTS data only to fine-tune the policy. We run TD3+ MBTS+BC on the same 5 seeds as reported in the original dataset.

As reported in Table 1, we find that, in all cases, TD3+ MBTS+BC outperforms the baseline method thus solidifying the positive effect of MBTS in offline RL. For this comparison, we also consider two additional state-of-the-art model-free offline RL algorithms: IQL (Kostrikov et al., 2021) and CQL (Kumar et al., 2020). In 6 out of the 9 comparable tasks, TD3+ MBTS+BC significantly improves over the model-free baselines. In the hopper medium-replay task, we find that TD3+ MBTS+BC under-performs compared to other model-free methods (IQL and CQL).

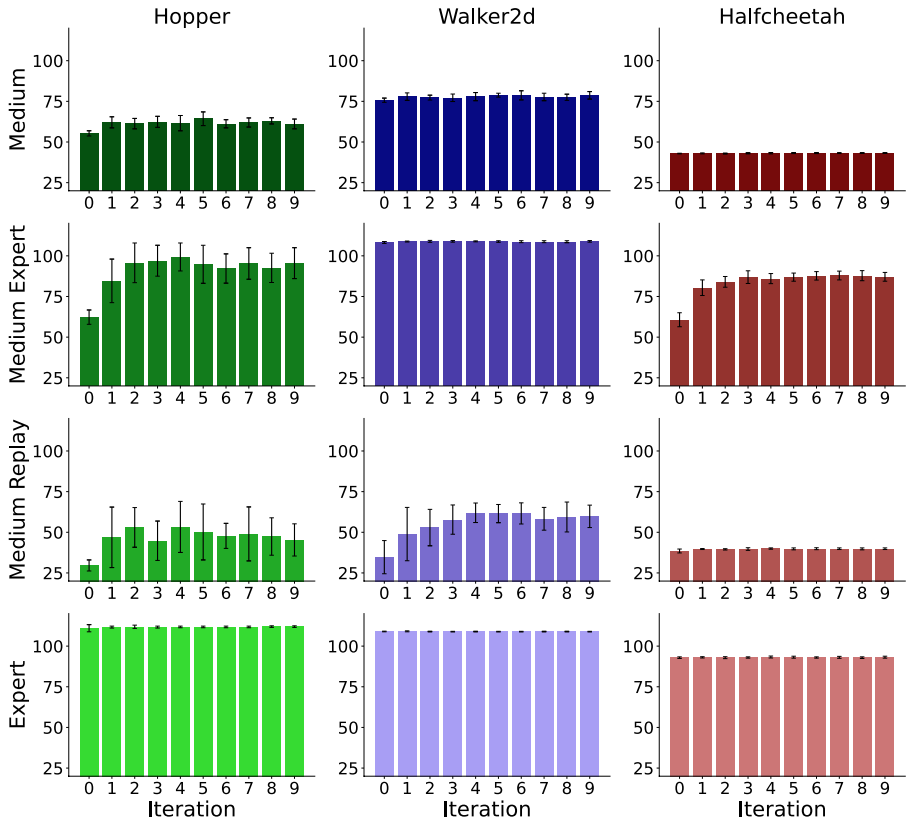


**Fig. 2** Comparative performance of BC and MBTS+BC as the fraction of expert trajectories increases up to 40%. For two environments, Hopper (left) and Walked2D (right), we report the average return of 10 trajectory evaluations of the best checkpoint during BC training. BC has been trained over 5 random seeds and MBTS has produced 3 datasets over different random seeds


**Table 1** Average normalised scores of state-of-the-art offline RL methods achieved on three locomotion tasks (Hopper, Halfcheetah and Walker2d) using the D4RL v2 data sets

Dataset	BC	TD3+BC	IQL	CQL	MOPO	MOReL	Diffuser	MBOP	MBTS+BC (ours)	TD3+ MBTS+BC (ours)	MBTS+MBOP (ours)
Medium	Hopper	55.3	59.3	66.3	58.5	<b>95.4</b>	58.5	56.9	64.3 ± 4.2 (+16.3%)	64.1 ± 4.4 (+8.1%)	66.5 ± 5.5 (+16.9%)
	Halfcheetah	42.9	48.3	47.4	44.0	42.1	44.2	51.2	43.2 ± 0.3 (+0.7%)	48.4 ± 0.4 (+0.2%)	<b>51.3 ± 0.3</b> (+0.2%)
	Walker2d	75.6	83.7	78.3	72.5	77.8	79.7	73.5	78.8 ± 1.2 (+4.2%)	<b>84.2 ± 1.4</b> (+0.6%)	77.3 ± 2.8 (+5.2%)
MedExp	Hopper	62.3	98.0	91.5	105.4	23.7	107.2	70.7	94.8 ± 11.7 (+52.2%)	109.1 ± 3.9 (+11.9%)	<b>110.4 ± 1.2</b> (+56.2%)
	Halfcheetah	60.7	90.7	86.7	91.6	63.3	79.8	63.5	86.9 ± 2.5 (+43.2%)	93.8 ± 3.4 (+3.4%)	<b>94.1 ± 1.0</b> (+48.2%)
	Walker2d	108.2	110.1	109.6	108.8	44.6	108.4	111.0	108.8 ± 5.5 (+16.9%)	110.3 ± 0.4 (+0.2%)	<b>111.1 ± 0.2</b> (+0.1%)
MedRep	Hopper	29.6	60.9	94.7	95.0	67.5	<b>96.8</b>	40.5	50.2 ± 17.2 (+69.6%)	77.4 ± 17.0 (+27.1%)	68.2 ± 9.5 (+68.4%)
	Halfcheetah	38.5	44.6	44.2	45.5	39.0	42.2	45.4	39.8 ± 0.6 (+3.4%)	44.7 ± 0.6 (+0.2%)	<b>46.7 ± 1.0</b> (+2.9%)
	Walker2d	34.7	81.8	73.9	77.2	53.1	61.2	53.8	61.5 ± 5.6 (+77.2%)	<b>82.8 ± 3.4</b> (+1.2%)	71.9 ± 5.6 (+33.6%)
Expert	Hopper	111.0	108.8	–	–	–	–	111.3	<b>111.8 ± 0.5</b> (+0.7%)	110.9 ± 2.7 (+2.9%)	111.3 ± 1.1 (±0.0%)
	Halfcheetah	92.9	96.7	–	–	–	–	98.2	93.2 ± 0.6 (+0.3%)	97.6 ± 0.6 (+0.9%)	<b>98.7 ± 1.0</b> (+0.5%)
	Walker2d	109.0	110.2	–	–	–	–	109.0	108.9 ± 0.2 (−0.1%)	<b>110.3 ± 0.3</b> (+0.1%)	109.5 ± 0.1 (+0.5%)

The results for competing methods have been gathered from the original publications. Bold scores represent the highest scores per task. MBTS+BC, TD3+ MBTS+BC, MBTS+MBOP: In brackets we report the percentage improvement achieved by MBTS relative to their respective baselines



**Fig. 4** Returns of BC extracted policies as the number of iterations of MBTS is increased. Iteration 0 are the BC scores on the original D4RL datasets. The errors bars represent the standard deviation of the average returns of 10 trajectory evaluations over 5 random seeds of BC and 3 random seeds of MBTS



**Fig. 3** Estimated KL-divergence and MSE of the BC and MBTS+BC policies on the Hopper and Walker2d environments as the fraction of expert trajectories increases. (Left) Relative difference between the KL-divergence of the BC policy and the expert and the KL-divergence of the MBTS+BC policy and the expert. Larger values represent the MBTS+BC policy being closer to the expert than the BC policy. MSE between actions evaluated from the expert policy and the learned policy on states from the Hopper (Middle) and Walker2d (Right) environments. The y-axes (Middle and Right) are on a log-scale. All policies were collected by training BC over 5 random seeds, with MBTS being evaluated over 3 different random seeds. All KL-divergences were scaled between 0 and 1, depending on the minimum and maximum values per task, before the difference was taken




## 4.2 Expected performance on sub-optimal data

It is well known that BC minimises the KL-divergence of trajectory distributions between the learned policy and  $\pi_\beta$  (Ke et al., 2020). As MBTS has the effect of improving  $\pi_\beta$ , this suggests that the KL-divergence between the trajectory distributions of the learned policy and the expert policy would be smaller post MBTS. To investigate this hypothesis, we use two complex locomotion tasks, Hopper and Walker2D, in OpenAI's gym (Brockman et al., 2016). Independently for each task, we first train an expert policy,  $\pi^*$ , with TD3 (Fujimoto et al., 2018), and use this policy to generate a baseline noisy dataset by sampling the expert policy in the environment and adding white noise to the actions, i.e.  $a = \pi^*(s) + \epsilon$ , which gives us a stochastic behaviour policy. A range of different, sub-optimal datasets are created by adding a certain amount of expert trajectories to the noisy dataset so that they make up  $x\%$  of the total trajectories. Using this procedure, we create eight different datasets by controlling  $x$ , which takes values in the set  $\{0, 0.1, 2.5, 5, 10, 20, 30, 40\}$ . BC is run on each dataset for 5 random seeds. We run MBTS (for five iterations) on each dataset over three different random seeds and then create BC policies over the 5 random seeds, giving 15 MBTS+BC policies. Random seeds cause different MBTS trajectories as they affect the latent variables sampled for the reward function and inverse dynamics model. Also, the initialisation of weights is randomised for the value function and BC policies hence the robustness of the methods is tested over multiple seeds. The KL divergences are calculated following (Ke et al., 2020) as

$$D_{KL}(p_{\pi^*}(\mathcal{T}), p_{\pi}(\mathcal{T})) = \mathbb{E}_{s \sim p_{\pi^*}, a \sim \pi^*(s)} [\log \pi^*(a | s) - \log \pi(a | s)].$$

Figure 2 shows the scores as average returns from 10 trajectory evaluations of the learned policies. MBTS+BC consistently improves on BC across all levels of expertise for both the Hopper and Walker2d environments. As the percentage of expert data increases, MBTS is available to leverage more high-value transitions, consistently improving over the BC baseline. Fig. 3 (left) shows the average difference in KL-divergences of the BC and MBTS+BC policies against the expert policy. Precisely, the y-axis represents  $D_{KL}(p_{\pi^*}(\mathcal{T}), p_{\pi^{BC}}(\mathcal{T})) - D_{KL}(p_{\pi^*}(\mathcal{T}), p_{\pi^{MBTS+BC}}(\mathcal{T}))$ , where  $p_{\pi}(\mathcal{T})$  is the trajectory distribution for policy  $\pi$ , Eq. (1). A positive value represents the MBTS+BC policy being closer to the expert, and a negative value represents the BC policy being closer to the expert, with the absolute value representing the degree to which this is the case. We also scale the average KL-divergence between 0 and 1, where 0 is the smallest KL-divergence and 1 is the largest KL-divergence, per task. This makes the scale comparable between Hopper and Walker2d. The figure shows that BC can extract a behaviour policy closer to the expert after performing MBTS on the dataset, except in the 0% case for Walker2D, however the difference is not significant. MBTS seems to work particularly well with a minimum of 2.5% expert data for Hopper and 0.1% for Walker2d.

Furthermore, Fig. 3 (middle and right) shows the mean square error (MSE) between actions from the expert policy and the learned policy for the Hopper (middle) and Walker2d (right) tasks. Actions are selected by collecting 10 trajectory evaluations of an expert policy. As we expect, the MBTS+BC policies produce actions closer to the experts on most levels of dataset expertise. A surprising result is that for 0% expert data on the Walker2d environment the BC policy produces actions closer to the expert than the MBTS+BC policy. This is likely due to MBTS not having any expert data to leverage. However, even in this case, MBTS still produces a higher-quality dataset than previous as shown by the increased performance on the average returns. Overall, these results offer empirical



**Fig. 5** Assessment of different types of models to predict reward on the hopper-medium D4RL dataset. The MSE between predicted and true rewards are assessed during training on a test set and training set of the same size

**Table 2** MSE between true and predicted rewards from the reward functions evaluated on the other D4RL hopper datasets

Networks	Hopper-random	Hopper-expert	Hopper-medium replay
GAN	$0.013 \pm 0.059$	$0.00019 \pm 0.0037$	$0.0039 \pm 0.050$
VAE	$0.011 \pm 0.055$	$0.000021 \pm 0.00011$	$0.0019 \pm 0.032$
MLP	$0.011 \pm 0.061$	$0.000024 \pm 0.00014$	$0.0022 \pm 0.047$
Gaussian	$5.18 \pm 2.05$	$0.60 \pm 0.68$	$1.59 \pm 1.79$

This table shows the performance of the reward models when evaluated on unseen data. The standard deviation is over the whole dataset

confirmation that MBTS does have the effect of improving the underlying behaviour policy of the dataset.

### 4.3 On the number of MBTS iterations

We investigate empirically how the quality of the dataset improves after each iteration; see Definition 3. We repeat MBTS on each D4RL dataset, each time using a newly estimated value function to take into account the newly generated transitions. In all our experiments, we choose 5 iterations. Figure 4 shows the scores of the D4RL environments on the different iterations, with the standard deviation across seeds shown as the error bar. With iteration 0 we indicate the BC score as obtained on the original D4RL datasets. For all datasets, we observe that the average scores of BC increase initially over a few iterations, then remain stable with only some minor random fluctuations. We see less improvement in the expert datasets as there are fewer trajectory improvements to be made. Conversely, for the medium expert datasets more iterations are required to reach an improved performance.

For Hopper and Walker2d medium-replay, there is a higher degree of standard deviation across the seeds, which gives a less stable average as the number of iterations increases.

## 4.4 Ablation studies

In this Section we perform ablation studies to assess the impact of the reward model on MBTS performance and the effect of value-weighted BC.

### 4.4.1 Choice of reward model

MBTS requires a predictive model for rewards associated to the stitched transitions enabling a value function to be learned on the new dataset. Unlike some online methods (Chua et al., 2018; Nagabandi et al., 2018) we do not have access to the true reward function during training time and so a model must be trained to predict rewards. There are many choices of models. For example, MBPO (Janner et al., 2019), MOPO (Yu et al., 2020) and MBOP (Argenson & Dulac-Arnold, 2020) use a neural network that outputs the parameters of a Gaussian distribution, to predict the next state and reward. These models are coupled with the next state as well as reward. We solely want to predict the reward and consider the following options: a Gaussian distribution whose parameters are modelled by a neural network, a Wasserstein-GAN, a VAE and multilayer neural network that minimizes the mean square error between true and predicted reward.

We evaluate the reward models on the D4RL hopper-medium dataset and perform a 95 : 5 training and test split. To make it a fair test all models are trained on the same training data and each model has two hidden layers with dimension size 512. Fig. 5 shows the mean-square error (MSE) between predicted and true rewards during training on the test and train set. From this clearly the VAE model and MLP model perform the best by attaining the smallest error, getting training and test error to  $10^{-5}$ . The average reward for a

**Table 3** Comparison of BC, MBTS(WGAN)+BC and MBTS(MLP)+BC on the D4RL locomotion tasks

Dataset	BC	MBTS(WGAN)+BC	MBTS(MLP) +BC
Hopper-medium	55.3	64.3 ± 4.2	63.7 ± 3.3
Halfcheetah-medium	42.9	43.2 ± 0.3	43.2 ± 0.2
Walker2d-medium	75.6	78.8 ± 1.2	77.6 ± 2.4
Hopper-mediumexpert	62.3	94.8 ± 11.7	97.7 ± 11.0
Halfcheetah-mediumexpert	60.7	86.9 ± 2.5	86.7 ± 2.8
Walker2d-mediumexpert	108.2	108.8 ± 0.5	109.0 ± 0.5
Hopper-mediumreplay	29.6	50.2 ± 17.2	51.9 ± 10.9
Halfcheetah-mediumreplay	38.5	39.8 ± 0.6	40.0 ± 0.4
Walker2d-mediumreplay	34.7	61.5 ± 5.6	58.8 ± 8.9
Hopper-expert	111.0	111.8 ± 0.5	111.5 ± 0.9
Halfcheetah-expert	92.9	93.2 ± 0.6	92.9 ± 0.7
Walker2d-expert	109.0	108.9 ± 0.2	108.8 ± 0.1

For the MBTS methods, the mean performance is provided over 3 datasets of MBTS and 5 seeds of BC and the standard deviation is given over the total of 15 policies

transition in the hopper-medium dataset is 3.11, so in fact the GAN also performs very well by attaining a training and test error of order  $10^{-4}$ .

In MBTS we want to predict a reward for an unseen transition, where  $s$  and  $s'$  are in the dataset but have never been connected by an observed action. Therefore, we evaluate the trained reward models on unseen data to test their OOD performance. Table 2 shows the MSE between predicted and true rewards of the models on the rest of the D4RL hopper datasets: random, expert and medium replay. The GAN, VAE and MLP perform very similarly achieving accurate predictions on all three datasets. The VAE and MLP outperform the GAN in predicting rewards of the expert dataset. The Gaussian model performed very poorly on these datasets.

Finally we compare MBTS(WGAN)+BC with MBTS(MLP)+BC on the D4RL datasets; here, either a WGAN or MLP is used to predict the reward. Table 3 shows that the decision between using a WGAN or MLP is insignificant as they are both accurate enough at predicting rewards.

#### 4.4.2 Value-weighted BC

MBTS uses a value function to estimate the future returns from any given state. Therefore MBTS+BC has a natural advantage over just BC which uses only the states and actions. To ensure that using a value function is only sufficient to improve the performance of BC, we investigate a weighted version of the BC loss function whereby the weights are given by the estimated value function, i.e.

$$\pi^{\text{BC}}(s) = \arg \min_{\pi} \mathbb{E}_{s,a \sim \mathcal{D}} [V_{\theta}(s)(\pi(s) - a)^2].$$

This weighted-BC method gives larger weight to the high-value states and lower weight to the low-value states during training.

On the Hopper medium and medium-expert datasets, training this weighted-BC method only gives a slight improvement over the original BC-cloned policy. For Hopper-medium, weighted-BC achieves an average score of 59.21 (with standard deviation 3.4); this is an improvement over BC (55.3), but lower than MBTS+BC (64.3). Weighted-BC on hopper-medexp achieves an average score of 66.02 (with standard deviation 6.9); again, this is a slight improvement over BC (62.3), but significantly lower than MBTS+BC (94.8). The experiments indicate that using a value function to weight the relative importance of seen states when optimising the BC objective function is not sufficient to achieve the performance gains introduced by MBTS.

## 5 Discussion

The proposed method, MBTS, has been presented for learning an optimal policy in continuous state and action spaces. For other domains, alternative modeling techniques for dynamic models should be considered. As demonstrated in Section 4.2, MBTS is expected to be beneficial in settings with sub-optimal data, even with a small percentage of additional expert data. Notably, MBTS does not require expert transitions within the initial dataset, as evidenced by the nearly 80% improvement in Walker2d medium-replay, which contains no expert data.

Empirically, MBTS does not damage performance; however, this does not guarantee that performance will not decrease under all circumstances. We believe that the primary risk of MBTS harming performance lies in the use of imperfect models. To mitigate model error in the forward models, we employ an ensemble and take a conservative approach to determining the next state, as shown in Eq. . To further account for model error, we only replace existing trajectories with new ones if they significantly increase returns, as per Definition 2. Due to Definition 2, we believe that the MBTS procedure is consistently safe when replacing existing trajectories, as it guarantees performance improvement. However, this may result in the loss of potentially useful information that MBTS can no longer use in future iterations.

Our method might be perceived as computationally intensive due to the number of models and the need to iterate over the entire dataset. However, we use our models with their limitations in mind, avoiding extrapolation of unseen states and generating actions only between in-distribution reachable states. The primary computational burden comes from searching for reachable next states. We have reduced this burden by evaluating the forward model only on “reasonably close” states using nearest neighbors organized by a KD-tree. KD-trees are well-studied and have a worst-case time complexity of  $\mathcal{O}(k \cdot n^{1-\frac{1}{k}})$  for  $k$ -dimensional trees and  $n$  data points (Lee & Wong, 1977). Consequently, we assume that states far apart in Euclidean distance are not reachable and do not require evaluation. Our approach is detailed in the Appendix; however, alternative methods could be employed to reduce complexity, and the technique we used is not strictly integral to the MBTS framework. Moreover, our method converges in a remarkably few number of iterations (max 5), as shown in Fig. 4, significantly reducing the computational cost.

## 6 Conclusion

In this paper, we have proposed an iterative data improvement strategy, Model-Based Trajectory Stitching, which can be applied to historical datasets containing demonstrations of sequential decisions taken to solve a complex task. At each iteration, MBTS performs one-step stitching between reachable states within the dataset that lead to higher future expected returns. We have demonstrated that, without further interactions with the environment, MBTS improves the quality of the historical demonstrations, which in turn has the effect of boosting the performance of BC-extracted policies significantly. Extensive experimental results using the D4RL benchmarking data have demonstrated that MBTS always improves the underlying behaviour policy. We have also demonstrated that MBTS is beneficial beyond BC, when combined with existing offline reinforcement learning methods. In particular, MBTS can be used to extract an improved explicit BC-based regulariser for TD3+BC, as well as an improved BC prior for offline model-based planning (MBOP). MBTS-based methods achieve state-of-the-art results in 10 out of the 12 D4RL datasets considered.

We believe that this work opens up a number of directions for future investigation. For example, MBTS could be extended to multi-agent offline policy learning by reformulating Eq. 2 to actions taken by multiple agents. Besides the realm of offline RL, MBTS may also be useful for learning with sub-optimal demonstrations, e.g. by inferring a reward function through inverse RL. Historical demonstrations can also be used to guide RL and improve the data efficiency of online RL (Hester et al., 2018). In these cases, BC can be used to initialise or regularise the training policy (Rajeswaran et al., 2017; Nair et al., 2018).

## Appendix A: Further implementation details

In this Appendix we report on all the hyperparameters required for MBTS as used on the D4RL datasets. All hyperparameters have been kept the same for every dataset, notable the acceptance threshold of  $\bar{p} = 0.1$ . MBTS consists of four components: a forward dynamics model, an inverse dynamics model, a reward function and a value function. Table 4 provides an overview of the implementation details and hyperparameters for each MBTS component. As our default optimiser we have used Adam (Kingma & Ba, 2014) with default hyperparameters, unless stated otherwise. Our code implementation is provided at <https://github.com/CharlesHepburn1/Model-Based-Trajectory-Stitching>.

### Forward dynamics model

Each forward dynamics model in the ensemble consists of a neural network with three hidden layers of size 200 with ReLU activation. The network takes a state  $s$  as input and outputs a mean  $\mu$  and standard deviation  $\sigma$  of a Gaussian distribution  $\mathcal{N}(\mu, \sigma^2)$ . For all experiments, an ensemble size of 7 is used with the best 5 being chosen.

### Inverse dynamics model

To sample actions from the inverse dynamics model of the environment, we have implemented a CVAE with two hidden layers with ReLU activation. The size of the hidden layer depends on the size of the dataset (Zhou et al., 2020): when the dataset has less than 900, 000 transitions (e.g. the medium-replay datasets) the layer has 256 nodes; when larger, it has 750 nodes. The encoder  $q_{\omega_1}$  takes in a tuple consisting of state, action and next state; it encodes it into a mean  $\mu_q$  and standard deviation  $\sigma_q$  of a Gaussian distribution  $\mathcal{N}(\mu_q, \sigma_q)$ . The latent variable  $z$  is then sampled from this distribution and used as input for the decoder along with the state,  $s$ , and next state,  $s'$ . The decoder outputs an action that is likely to connect  $s$  and  $s'$ . The CVAE is trained for 400, 000 gradient steps with hyperparameters given in Table 4.

### Reward function

The reward function is used to predict reward signals associated with new transitions,  $s, a, s'$ . For this model, we use a conditional-WGAN with two hidden layers of size 512. The generator,  $G_\phi$ , takes in a state  $s$ , action  $a$ , next state  $s'$  and latent variable  $z$ ; it outputs a reward  $r$  for that that transition. The decoder takes a full transition of  $(s, a, r, s')$  as input to determine whether this transition is likely to have come from the dataset or not. In the reward ablation study all models use the same number of hidden layers and dimension size and are trained for 500k iterations.

### Value function

Similarly to previous methods (Fujimoto et al., 2019), our value function  $V_\theta$  takes the minimum of two value functions,  $\{V_{\theta_1}, V_{\theta_2}\}$ . Each value function is a neural network with two hidden layers of size 256 and a ReLU activation. The value function takes in a state  $s$  and

determines the sum of future rewards of being in that state and following the policy (of the dataset) thereon.

### KL-divergence experiment

As the KL-divergence requires a continuous policy, the BC policy network is a 2-layer MLP of size 256 with ReLU activation, but with the final layer outputting the parameters of a Gaussian,  $\mu_s$  and  $\sigma_s$ . We carry out maximum likelihood estimation using a batch size of 256. For the Walker2d experiments, MBTS was slightly adapted to only accept new trajectories if they made less than ten changes. For each level of difficulty, MBTS is run 3 times and the scores are the average of the mean returns over 10 evaluation trajectories of 5 random seeds of BC. To compute the KL-divergence, a continuous expert policy is also required, but TD3 gives a deterministic one. To overcome this, a continuous expert policy is created by assuming a state-dependent normal distribution centred around  $\pi^*(s)$  with a standard deviation of 0.01.


### Search procedure for candidate next states

Calculating  $p(s' | s)$  for all  $s' \in \mathcal{D}$  may be computationally inefficient. To speed this up in the MuJoCo environments, we initially select a smaller set of candidate next states by thresholding the Euclidean distance. Although on its own a geometric distance would not be sufficient to identify stitching events, we found that in our environments it can help reduce the set of candidate next states thus alleviating the computational workload. To pre-select a smaller set of candidate next states, we use two criteria. Firstly, from a transition  $(s, a, r, s') \in \mathcal{D}$ , a neighbourhood of states around  $s$  is taken and the following state in the trajectory is collected. Secondly, all the states in a neighbourhood around  $s'$  are collected. This process ensures all candidate next states are geometrically-similar to  $s'$  or are preceded by geometrically-similar states. The neighbourhood of a state is an  $\epsilon$  – ball around the state. When  $\epsilon$  is large enough, we can retain all feasible candidate next states for evaluation with the forward dynamic model. Fig. 6 illustrates this procedure.

### D4RL experiments

For the D4RL experiments, we run MBTS 3 times for each dataset and average the mean returns over 10 evaluation trajectories of 5 random seeds of BC, to attain the results for MBTS+BC. For the BC results, we average the mean returns over 10 evaluation trajectories of 5 random seeds. The BC policy network is a 2-layer MLP of size 256 with ReLU activation, the final layer has tanh activation multiplied by the action dimension. We use the Adam optimiser with a learning rate of  $1e - 3$  and a batch size of 256.

The hyperparameters we use for MBOP are given in Table 5. TD3+BC is trained for 1000k iterations we train TD3+ MBTS+BC also for 1000k iterations with the actor and critic dimensions the same as the original implementation. For TD3+ MBTS+BC we warm start the algorithm on the original data and train for 800k iterations and then carry on training for the remaining 200k iterations on the new MBTS data. As the MBTS dataset contains many duplicate transitions we remove all duplicates from the dataset when training with TD3+BC. For the hopper datasets (except medium-expert)



**Fig. 6** Visualisation of our two definitions of a neighbourhood. For a transition  $(s_t, a_t, s_{t+1}) \in \mathcal{D}$ , the neighbourhoods are used to reduce the size of the set of candidate next states. (Left) All states within an  $\epsilon$ -ball of the current state,  $s_t$ , are taken and the next state in their respective trajectories (joined by an action shown as an arrow) are added to the set of candidate next states. (Right) All states within an  $\epsilon$ -ball of the next state,  $s_{t+1}$  are added to the set of candidate next states. The full set of candidate next states are highlighted in yellow (Color figure online)

**Table 4** Hyperparameters and values for models used in MBTS

Hyperparameter	Value
<i>Forward dynamics model</i>	
Optimiser	Adam
Learning rate	3e-4
Batch size	256
Ensemble size	7
<i>Inverse dynamics model</i>	
Optimiser	Adam
Learning rate	1e-4
Batch size	100
Latent dim	2*action dim
<i>Reward function</i>	
Optimiser	Adam
	$\beta = (0.5, 0.999)$
Learning rate	1e-4
Batch size	256
Latent dim	2
L2 regularisation	1e-4
<i>Value function</i>	
Optimiser	Adam
Learning rate	3e-4
Batch size	256

the policy is improved if the data is swapped to the MBTS dataset at 600k iterations. Also the critic is fixed and training on the MBTS dataset starts at 900k iterations for the walker2d medium-replay dataset.



**Table 5** Hyperparameters used for the MBOP method across the D4RL datasets

	Dataset	Horizon	# Samples	Kappa	Sigma	Beta
Medium	Hopper	2	100	1	0.2	0.0
	Halfcheetah	2	100	3	0.2	0.0
	Walker2d	4	1000	3	0.01	0.0
MedExp	Hopper	2	100	1	0.05	0.0
	Halfcheetah	2	100	1	0.01	0.0
	Walker2d	2	1000	3	0.1	0.0
MedRep	Hopper	8	100	1	0.01	0.0
	Halfcheetah	2	100	0.3	0.2	0.0
	Walker2d	2	1000	0.3	0.2	0.0
Expert	Hopper	2	100	0.3	0.01	0.0
	Halfcheetah	4	100	0.3	0.05	0.0
	Walker2d	2	1000	3	0.05	0.0

**Acknowledgements** CH acknowledges support from the Engineering and Physical Sciences Research Council through the Mathematics of Systems Centre for Doctoral Training at the University of Warwick (EP/S022244/1). GM acknowledges support from a UKRI Turing AI Acceleration Fellowship (EPSRC EP/V024868/1).

**Author Contributions** CAH contributed to the idea, wrote the code, performed the experiments, generated figures and tables, and co-wrote the paper. GM contributed to the idea, advised on experiments, and co-wrote the paper.

**Funding** CH acknowledges support from the Engineering and Physical Sciences Research Council through the Mathematics of Systems Centre for Doctoral Training at the University of Warwick (EP/S022244/1). GM acknowledges support from a UKRI Turing AI Acceleration Fellowship (EPSRC EP/V024868/1).

**Data availability** The data was obtained from the public D4RL offline RL benchmarking datasets.

**Code availability** All code will be made public on Github upon release of the paper.

## Declarations

**Conflict of interest** No conflicts of interest.

**Ethical approval** No ethics approval required- the work presented in this manuscript relies on simulated data.

**Consent to participate** Not applicable

**Consent for publication** Not applicable

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- An, G., Moon, S., Kim, J.-H., & Song, H.O. (2021). Uncertainty-based offline reinforcement learning with diversified q-ensemble. In: *Advances in Neural Information Processing Systems* **34**
- Argenson, A., & Dulac-Arnold, G.: (2020). Model-based offline planning. arXiv preprint [arXiv:2008.05556](https://arxiv.org/abs/2008.05556)
- Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. In: *International Conference on Machine Learning*, pp. 214–223 PMLR
- Bacci, G., Bacci, G., Larsen, K.G., & Mardare, R. (2013). Computing behavioral distances, compositionally. In: *International Symposium on Mathematical Foundations of Computer Science*, pp. 74–85 . Springer
- Bacci, G., Bacci, G., Larsen, K.G., & Mardare, R. (2013). On-the-fly exact computation of bisimilarity distances. In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pp. 1–15 . Springer
- Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L.D., Monfort, M., Muller, U., & Zhang, J. et al.: (2016). End to end learning for self-driving cars. arXiv preprint [arXiv:1604.07316](https://arxiv.org/abs/1604.07316)
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). Openai gym. arXiv preprint [arXiv:1606.01540](https://arxiv.org/abs/1606.01540)
- Buckman, J., Hafner, D., Tucker, G., Brevdo, E., & Lee, H. (2018). Sample-efficient reinforcement learning with stochastic ensemble value expansion. In: *Advances in neural information processing systems* **31**
- Castro, P.S. (2020). Scalable methods for computing state similarity in deterministic markov decision processes. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 10069–10076
- Char, I., Mehta, V., Villafior, A., Dolan, J.M., & Schneider, J. (2022). Bats: Best action trajectory stitching. arXiv preprint [arXiv:2204.12026](https://arxiv.org/abs/2204.12026)
- Chen, D., Breugel, F.v., & Worrell, J. (2012). On the complexity of computing probabilistic bisimilarity. In: *International Conference on Foundations of Software Science and Computational Structures*, pp. 437–451 . Springer
- Chen, X., Zhou, Z., Wang, Z., Wang, C., Wu, Y., & Ross, K. (2020). Bail: Best-action imitation learning for batch deep reinforcement learning. *Advances in Neural Information Processing Systems*, *33*, 18353–18363.
- Chua, K., Calandra, R., McAllister, R., & Levine, S. (2018). Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In: *Advances in neural information processing systems* **31**
- Chua, K., Calandra, R., McAllister, R., & Levine, S. (2018). Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In: *Advances in neural information processing systems* **31**
- Codevilla, F., Santana, E., López, A.M., & Gaidon, A. (2019). Exploring the limitations of behavior cloning for autonomous driving. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9329–9338
- Dadashi, R., Rezaeifar, S., Vieillard, N., Hussenot, L., Pietquin, O., & Geist, M. (2021). Offline reinforcement learning with pseudometric learning. In: *International Conference on Machine Learning*, pp. 2307–2318 . PMLR
- Farag, W., & Saleh, Z. (2018). Behavior cloning for autonomous driving using convolutional neural networks. In: *2018 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT)*, pp. 1–7 . IEEE
- Feinberg, V., Wan, A., Stoica, I., Jordan, M.I., Gonzalez, J.E., & Levine, S. (2018). Model-based value estimation for efficient model-free reinforcement learning. arXiv preprint [arXiv:1803.00101](https://arxiv.org/abs/1803.00101)
- Ferns, N., Castro, P.S., Precup, D., & Panangaden, P. (2012). Methods for computing state similarity in markov decision processes. arXiv preprint [arXiv:1206.6836](https://arxiv.org/abs/1206.6836)
- Ferns, N., Panangaden, P., & Precup, D. (2004). Metrics for finite markov decision processes. In: *UAI*, vol. 4, pp. 162–169
- Finn, C., Levine, S., & Abbeel, P. (2016). Guided cost learning: Deep inverse optimal control via policy optimization. In: *International Conference on Machine Learning*, pp. 49–58 PMLR
- Fu, J., Kumar, A., Nachum, O., Tucker, G., & Levine, S. (2020). D4rl: Datasets for deep data-driven reinforcement learning. arXiv preprint [arXiv:2004.07219](https://arxiv.org/abs/2004.07219)
- Fujimoto, S., & Gu, S.S. (2021). A minimalist approach to offline reinforcement learning. In: *Advances in Neural Information Processing Systems* **34**
- Fujimoto, S., Hoof, H., & Meger, D. (2018). Addressing function approximation error in actor-critic methods. In: *International Conference on Machine Learning*, pp. 1587–1596 PMLR
- Fujimoto, S., Meger, D., & Precup, D. (2019). Off-policy deep reinforcement learning without exploration. In: *International Conference on Machine Learning*, pp. 2052–2062. PMLR

- Giusti, A., Guzzi, J., Cireşan, D. C., He, F.-L., Rodríguez, J. P., Fontana, F., Faessler, M., Forster, C., Schmidhuber, J., Di Caro, G., et al. (2015). A machine learning approach to visual perception of forest trails for mobile robots. *IEEE Robotics and Automation Letters*, *1*(2), 661–667.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems* **27**
- Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: *International Conference on Machine Learning*, pp. 1861–1870. PMLR
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., & Abbeel, P. (2018). et al.: Soft actor-critic algorithms and applications. arXiv preprint [arXiv:1812.05905](https://arxiv.org/abs/1812.05905)
- Hepburn, C.A., & Montana, G. (2022). Model-based trajectory stitching for improved offline reinforcement learning. arXiv preprint [arXiv:2211.11603](https://arxiv.org/abs/2211.11603)
- Hester, T., Vecerik, M., Pietquin, O., Lanctot, M., Schaul, T., Piot, B., Horgan, D., Quan, J., Sendonaris, A., & Osband, I. et al: (2018). Deep q-learning from demonstrations. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32
- Ho, J., Ermon, S. (2016). Generative adversarial imitation learning. In: *Advances in neural information processing systems* **29**
- Hussein, A., Gaber, M. M., Elyan, E., & Jayne, C. (2017). Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, *50*(2), 1–35.
- Janner, M., Du, Y., Tenenbaum, J.B., & Levine, S. (2022). Planning with diffusion for flexible behavior synthesis. arXiv preprint [arXiv:2205.09991](https://arxiv.org/abs/2205.09991)
- Janner, M., Fu, J., Zhang, M., & Levine, S. (2019). When to trust your model: Model-based policy optimization. In: *Advances in Neural Information Processing Systems* **32**
- Jaques, N., Ghandeharioun, A., Shen, J.H., Ferguson, C., Lapedriza, A., Jones, N., Gu, S., & Picard, R. (2019). Way off-policy batch deep reinforcement learning of implicit human preferences in dialog. arXiv preprint [arXiv:1907.00456](https://arxiv.org/abs/1907.00456)
- Kadous, M.W., Sammut, C., & Sheh, R. (2005). Behavioural cloning for robots in unstructured environments. In: *Advances in Neural Information Processing Systems Workshop*
- Kalweit, G., & Boedecker, J. (2017). Uncertainty-driven imagination for continuous deep reinforcement learning. In: *Conference on Robot Learning*, pp. 195–206. PMLR
- Ke, L., Choudhury, S., Barnes, M., Sun, W., Lee, G., & Srinivasa, S. (2020). Imitation learning as f-divergence minimization. In: *International Workshop on the Algorithmic Foundations of Robotics*, pp. 313–329 Springer
- Kemertas, M., & Aumentado-Armstrong, T. (2021). Towards robust bisimulation metric learning. In: *Advances in Neural Information Processing Systems* **34**
- Kidambi, R., Rajeswaran, A., Netrapalli, P., & Joachims, T. (2020). Morel: Model-based offline reinforcement learning. *Advances in neural information processing systems*, *33*, 21810–21823.
- Kingma, D.P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
- Kingma, D.P., & Welling, M. (2013). Auto-encoding variational bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114)
- Kostrikov, I., Fergus, R., Tompson, J., & Nachum, O. (2021). Offline reinforcement learning with fisher divergence critic regularization. In: *International Conference on Machine Learning*, pp. 5774–5783 PMLR
- Kostrikov, I., Nair, A., & Levine, S. (2021). Offline reinforcement learning with implicit q-learning. arXiv preprint [arXiv:2110.06169](https://arxiv.org/abs/2110.06169)
- Kullback, S. (1997). Information theory and statistics. Courier Corporation
- Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, *22*(1), 79–86.
- Kumar, A., Fu, J., Soh, M., Tucker, G., & Levine, S. (2019). Stabilizing off-policy q-learning via bootstrapping error reduction. In: *Advances in Neural Information Processing Systems* **32**
- Kumar, A., Hong, J., Singh, A., & Levine, S. (2022). When should we prefer offline reinforcement learning over behavioral cloning? arXiv preprint [arXiv:2204.05618](https://arxiv.org/abs/2204.05618)
- Kumar, A., Zhou, A., Tucker, G., & Levine, S. (2020). Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, *33*, 1179–1191.
- Lange, S., Gabel, T., & Riedmiller, M. (2012). Batch reinforcement learning. *Reinforcement Learning* (pp. 45–73). Berlin: Springer.
- Le, H., Jiang, N., Agarwal, A., Dudik, M., Yue, Y., & Daumé III, H. (2018). Hierarchical imitation and reinforcement learning. In: *International Conference on Machine Learning*, pp. 2917–2926. PMLR

- Lee, D.-T., & Wong, C.-K. (1977). Worst-case analysis for region and partial region searches in multidimensional binary search trees and balanced quad trees. *Acta Informatica*, 9(1), 23–29.
- Levine, S., Kumar, A., Tucker, G., & Fu, J. (2020). Offline reinforcement learning: Tutorial, review, and perspectives on open problems. arXiv preprint [arXiv:2005.01643](https://arxiv.org/abs/2005.01643)
- Nagabandi, A., Kahn, G., Fearing, R.S., & Levine, S. (2018). Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 7559–7566 IEEE
- Nair, A., McGrew, B., Andrychowicz, M., Zaremba, W., & Abbeel, P. (2018). Overcoming exploration in reinforcement learning with demonstrations. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 6292–6299 . IEEE
- Pearce, T., & Zhu, J. (2022). Counter-strike deathmatch with large-scale behavioural cloning. In: 2022 IEEE Conference on Games (CoG), pp. 104–111. IEEE
- Pomerleau, D. A. (1988). Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1.
- Pomerleau, D. A. (1991). Efficient training of artificial neural networks for autonomous navigation. *Neural computation*, 3(1), 88–97.
- Rajeswaran, A., Kumar, V., Gupta, A., Vezzani, G., Schulman, J., Todorov, E., & Levine, S. (2017). Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. arXiv preprint [arXiv:1709.10087](https://arxiv.org/abs/1709.10087)
- Ross, S., Gordon, G., & Bagnell, D. (2011). A reduction of imitation learning and structured prediction to no-regret online learning. In: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, pp. 627–635 . JMLR Workshop and Conference Proceedings
- Sammut, C., Hurst, S., Kedzier, D., & Michie, D. (1992). Learning to fly. In: Machine Learning Proceedings 1992, pp. 385–393. Elsevier
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., & Moritz, P. (2015). Trust region policy optimization. In: International Conference on Machine Learning, pp. 1889–1897 PMLR
- Sohn, K., Lee, H., & Yan, X. (2015). Learning structured output representation using deep conditional generative models. In: Advances in neural information processing systems 28
- Sutton, R.S., & Barto, A.G. (1998). Reinforcement Learning: An Introduction, MIT press
- Sutton, R. S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4), 160–163.
- Wang, Q., Xiong, J., Han, L., Liu, H., & Zhang, T. et al.: (2018). Exponentially weighted imitation learning for batched historical data. In: Advances in Neural Information Processing Systems 31
- Wu, Y., Tucker, G., & Nachum, O. (2019). Behavior regularized offline reinforcement learning. arXiv preprint [arXiv:1911.11361](https://arxiv.org/abs/1911.11361)
- Yu, T., Kumar, A., Rafailov, R., Rajeswaran, A., Levine, S., & Finn, C. (2021). Combo: Conservative offline model-based policy optimization. In: Advances in Neural Information Processing Systems 34
- Yue, Y., Kang, B., Ma, X., Xu, Z., Huang, G., & Yan, S. (2022). Boosting offline reinforcement learning via data rebalancing. arXiv preprint [arXiv:2210.09241](https://arxiv.org/abs/2210.09241)
- Yu, T., Thomas, G., Yu, L., Ermon, S., Zou, J. Y., Levine, S., Finn, C., & Ma, T. (2020). Mopo: Model-based offline policy optimization. *Advances in Neural Information Processing Systems*, 33, 14129–14142.
- Zhan, X., Zhu, X., & Xu, H (2021). Model-based offline planning with trajectory pruning. arXiv preprint [arXiv:2105.07351](https://arxiv.org/abs/2105.07351)
- Zhang, A., McAllister, R., Calandra, R., Gal, Y., & Levine, S. (2020). Learning invariant representations for reinforcement learning without reconstruction. arXiv preprint [arXiv:2006.10742](https://arxiv.org/abs/2006.10742)
- Zhou, W., Bajracharya, S., & Held, D. (2020). Plas: Latent action space for offline reinforcement learning. arXiv preprint [arXiv:2011.07213](https://arxiv.org/abs/2011.07213)