## A. BMPO Performance Guarantee



Figure 6. Bidirectional rollout.

Lemma A.1. (Bidirectional Branched Rollout Returns Bound). Let $\eta_{1}, \eta_{2}$ be the expected returns of two bidirectional branched rollouts. Out of the branch, we assume that the expected total variation distance between these two dynamics at each timestep $t$ is bounded as $\max _{t} E_{(s, a) \sim p_{1}^{t}(s, a)} D_{T V}\left(p_{1}^{\mathrm{pre}}\left(s^{\prime} \mid s, a\right) \| p_{2}^{\text {pre }}\left(s^{\prime} \mid s, a\right)\right) \leq \epsilon_{m}^{\text {pre }}$, similarly, the forward branch dynamic bounded as $\max _{t} E_{(s, a) \sim p_{1}^{t}(s, a)} D_{T V}\left(p_{1}^{\text {for }}\left(s^{\prime} \mid s, a\right) \| p_{2}^{\text {for }}\left(s^{\prime} \mid s, a\right)\right) \leq \epsilon_{m}^{\text {for }}$, and the backward branch dynamic bounded as $\max _{t} E_{\left(s^{\prime}, a\right) \sim p_{1}^{t}\left(s^{\prime}, a\right)} D_{T V}\left(p_{1}^{\text {back }}\left(s \mid s^{\prime}, a\right) \| p_{2}^{\text {back }}\left(s \mid s^{\prime}, a\right)\right) \leq \epsilon_{m}^{\text {back }}$. Likewise, the total variation distance of


$$
\begin{array}{r}
\left|\eta_{1}-\eta_{2}\right| \leq 2 r_{\max }\left[\frac{\gamma^{k_{1}+k_{2}+1}}{(1-\gamma)^{2}}\left(\epsilon_{m}^{\mathrm{pre}}+\epsilon_{\pi}^{\mathrm{pre}}\right)+\frac{\gamma^{k_{1}+k_{2}}}{1-\gamma} \epsilon_{\pi}^{\mathrm{pre}}+\frac{1-\gamma^{k_{1}}}{1-\gamma}\left(k_{1}\left(\epsilon_{m}^{\mathrm{back}}+\epsilon_{\pi}^{\mathrm{back}}\right)+\epsilon_{\pi}^{\mathrm{back}}\right)\right.  \tag{11}\\
\left.+\frac{\gamma^{k_{1}}}{1-\gamma}\left(k_{2}\left(\epsilon_{m}^{\mathrm{for}}+\epsilon_{\pi}^{\mathrm{for}}\right)+\epsilon_{\pi}^{\mathrm{for}}\right)\right]
\end{array}
$$

Proof. Lemma B. 1 and Lemma B. 2 imply that state marginal error at each timestep can be bounded by the divergence at the current timestep plus the state marginal error at the next (Lemma B.1), or previous (Lemma B.2) timestep. And by employing Lemma B.3, we can convert the ( $\mathrm{s}, \mathrm{a}$ ) joint distribution to marginal distributions. Thus, letting $d_{1}(s, a)$ and $d_{2}(s, a)$ denote the state-action marginals, we can write:
For $t \leq k_{1}$ :

$$
\begin{align*}
D_{T V}\left(d_{1}^{t}(s, a) \| d_{2}^{t}(s, a)\right) & \leq D_{T V}\left(d_{1}^{t}(s) \| d_{2}^{t}(s)\right)+\max _{s^{\prime}} D_{T V}\left(\pi_{1}\left(a \mid s^{\prime}\right) \| \pi_{2}\left(a \mid s^{\prime}\right)\right)  \tag{12}\\
& \leq\left(k_{1}-t\right)\left(\epsilon_{m}^{\text {back }}+\epsilon_{\pi}^{\text {back }}\right)+\epsilon_{\pi}^{\text {back }} \leq k_{1}\left(\epsilon_{m}^{\text {back }}+\epsilon_{\pi}^{\text {back }}\right)+\epsilon_{\pi}^{\text {back }}
\end{align*}
$$

Similarly, for $k_{1}<t \leq k_{1}+k_{2}$ :

$$
\begin{equation*}
D_{T V}\left(d_{1}^{t}(s, a) \| d_{2}^{t}(s, a)\right) \leq\left(t-k_{1}\right)\left(\epsilon_{m}^{\mathrm{for}}+\epsilon_{\pi}^{\mathrm{for}}\right)+\epsilon_{\pi}^{\mathrm{for}} \leq k_{2}\left(\epsilon_{m}^{\mathrm{for}}+\epsilon_{\pi}^{\mathrm{for}}\right)+\epsilon_{\pi}^{\mathrm{for}} \tag{13}
\end{equation*}
$$

And for $t>k_{1}+k_{2}$ :

$$
\begin{equation*}
D_{T V}\left(d_{1}^{t}(s, a) \| d_{2}^{t}(s, a)\right) \leq\left(t-k_{1}-k_{2}\right)\left(\epsilon_{m}^{\mathrm{pre}}+\epsilon_{\pi}^{\mathrm{pre}}\right)+k_{2}\left(\epsilon_{m}^{\mathrm{for}}+\epsilon_{\pi}^{\mathrm{for}}\right)+\epsilon_{\pi}^{\mathrm{pre}}+\epsilon_{\pi}^{\mathrm{for}} \tag{14}
\end{equation*}
$$

We can now bound the difference in occupancy measures by averaging the state marginal error over time, weighted by the discount:

$$
\begin{aligned}
D_{T V}\left(d_{1}(s, a) \| d_{2}(s, a)\right) \leq & (1-\gamma) \sum_{t=0}^{\infty} \gamma^{t} D_{T V}\left(d_{1}^{t}(s, a) \| d_{2}^{t}(s, a)\right) \\
\leq & (1-\gamma) \sum_{t=0}^{k_{1}} \gamma^{t}\left(k_{1}\left(\epsilon_{m}^{\text {back }}+\epsilon_{\pi}^{\text {back }}\right)+\epsilon_{\pi}^{\text {back }}\right) \\
& +(1-\gamma) \sum_{t=k_{1}}^{k_{1}+k_{2}} \gamma^{t}\left(k_{2}\left(\epsilon_{m}^{\text {for }}+\epsilon_{\pi}^{\text {for }}\right)+\epsilon_{\pi}^{\text {for }}\right) \\
& +(1-\gamma) \sum_{t=k_{1}+k_{2}}^{\infty} \gamma^{t}\left(\left(t-k_{1}-k_{2}\right)\left(\epsilon_{m}^{\text {pre }}+\epsilon_{\pi}^{\text {pre }}\right)+k_{2}\left(\epsilon_{m}^{\text {for }}+\epsilon_{\pi}^{\text {for }}\right)+\epsilon_{\pi}^{\mathrm{pre}}+\epsilon_{\pi}^{\text {for }}\right) \\
= & \left(k _ { 1 } \left(\epsilon_{m}^{\text {back }}+\epsilon_{\pi}^{\left.\left.\operatorname{back}^{\text {ack }}\right)+\epsilon_{\pi}^{\text {back }}\right)\left(1-\gamma^{k_{1}}\right)+\left(k_{2}\left(\epsilon_{m}^{\text {for }}+\epsilon_{\pi}^{\text {for }}\right)+\epsilon_{\pi}^{\text {for }}\right)\left(\gamma^{k_{1}}\right)} \begin{array}{rl} 
& +\frac{\gamma^{k_{1}+k_{2}+1}}{1-\gamma}\left(\epsilon_{m}^{\text {pre }}+\epsilon_{\pi}^{\text {pre }}\right)+\gamma^{k_{1}+k_{2}} \epsilon_{\pi}^{\mathrm{pre}}
\end{array}\right.\right.
\end{aligned}
$$

Multiplying this bound by $\frac{2 r_{\text {max }}}{1-\gamma}$ to convert the occupancy measure difference into a returns bound completes the proof.
Theorem A.1. (BMPO Return Discrepancy Upper Bound) Assume that the expected total variation distance between the learned forward model $\hat{p}$ and the true dynamics $p$ at each timestep $t$ is bounded as $\max _{t} E_{(s, a) \sim \pi_{t}}\left[D_{T V}\left(p\left(s^{\prime} \mid s, a\right) \| \hat{p}\left(s^{\prime} \mid s, a\right)\right)\right] \leq \epsilon_{m}^{\text {for }}$. Similarly, the error of backward model $\hat{q}$ is bounded as $\max _{t} E_{\left(s^{\prime}, a\right) \sim \pi_{t}}\left[D_{T V}\left(q\left(s \mid s^{\prime}, a\right) \| \hat{q}\left(s \mid s^{\prime}, a\right)\right)\right] \leq \epsilon_{m}^{\text {back }}$ and the variation between current policy and the behavioral policy is bounded as $\max _{s} D_{T V}\left(\pi_{D}(a \mid s) \| \pi(a \mid s)\right) \leq \epsilon_{\pi}$. Assume $\epsilon_{m}^{\text {for }} \approx \epsilon_{m}^{\text {back }}=\epsilon_{m}$ and $\epsilon_{\pi}^{\text {back }}=0$, then under a branched rollouts scheme with a backward branch length of $k_{1}$ and a forward branch length of $k_{2}$, the returns are bounded as:

$$
\begin{equation*}
\left|\eta[\pi]-\eta^{\text {branch }}[\pi]\right| \leq 2 r_{\max }\left[\frac{\gamma^{k_{1}+k_{2}+1} \epsilon_{\pi}}{(1-\gamma)^{2}}+\frac{\gamma^{k_{1}+k_{2}} \epsilon_{\pi}}{(1-\gamma)}+\frac{\max \left(k_{1}, k_{2}\right)}{1-\gamma}\left(\epsilon_{m}\right)\right] . \tag{15}
\end{equation*}
$$

Proof. Using Lemma A.1, out of the branch, we only suffer from error of executing old policy $\pi_{D}$, so, set $\epsilon_{\pi}^{\text {pre }}=\epsilon_{\pi}$ and $\epsilon_{m}^{\text {pre }}=0$. Then in the branched rollout, we execute current policy, so the only error comes from using the learned model to simulate. Set $\epsilon_{\pi}^{\text {for }}=\epsilon_{\pi}^{\text {back }}=0$ and $\epsilon_{m}^{\text {for }}=\epsilon_{m}^{\text {back }}=\epsilon_{m}$. Plugging these in Lemma B. 1 we can get:

$$
\begin{align*}
\left|\eta[\pi]-\eta^{\text {branch }}[\pi]\right| & \leq 2 r_{\max }\left[\frac{\gamma^{k_{1}+k_{2}+1} \epsilon_{\pi}}{(1-\gamma)^{2}}+\frac{\gamma^{k_{1}+k_{2}} \epsilon_{\pi}}{(1-\gamma)}+\frac{k_{1}\left(1-\gamma^{k_{1}}\right)+k_{2}\left(\gamma^{k_{1}}\right)}{1-\gamma}\left(\epsilon_{m}\right)\right] \\
& \leq 2 r_{\max }\left[\frac{\gamma^{k_{1}+k_{2}+1} \epsilon_{\pi}}{(1-\gamma)^{2}}+\frac{\gamma^{k_{1}+k_{2}} \epsilon_{\pi}}{(1-\gamma)}+\frac{\left.\max \left(k_{1}, k_{2}\right)\left(1-\gamma^{k_{1}}+\gamma^{k_{1}}\right)\right)}{1-\gamma}\left(\epsilon_{m}\right)\right]  \tag{16}\\
& \leq 2 r_{\max }\left[\frac{\gamma^{k_{1}+k_{2}+1} \epsilon_{\pi}}{(1-\gamma)^{2}}+\frac{\gamma^{k_{1}+k_{2}} \epsilon_{\pi}}{(1-\gamma)}+\frac{\max \left(k_{1}, k_{2}\right)}{1-\gamma}\left(\epsilon_{m}\right)\right]
\end{align*}
$$

## B. Useful Lemmas

In this section, we give proofs of the lemmas used before.
Lemma B.1. (Backward State Marginal Distance Bound). Suppose the expected total variation distance between two backward dynamics is bounded as $\max _{t} E_{\left(s^{\prime}, a\right) \sim p_{1}^{t}}\left[D_{T V}\left(p_{1}\left(s \mid s^{\prime}, a\right) \| p_{2}\left(s \mid s^{\prime}, a\right)\right)\right] \leq \epsilon_{m}^{\text {back }}$ and the backward policy divergences are bounded as $\max _{s^{\prime}} D_{T V}\left(\pi_{1}\left(a \mid s^{\prime}\right) \| \pi_{2}\left(a \mid s^{\prime}\right)\right) \leq \epsilon_{\pi}^{\text {back. Then the state marginal distance at timestep } t \text { can be }}$ bounded as:

$$
\begin{equation*}
D_{T V}\left(p_{1}^{t}(s) \| p_{2}^{t}(s)\right) \leq \epsilon_{m}^{\text {back }}+\epsilon_{\pi}^{\text {back }}+D_{T V}\left(p_{1}^{t+1}(s) \| p_{2}^{t+1}(s)\right) . \tag{17}
\end{equation*}
$$

Proof. Let the total variation distance of state at time $t$ be denoted as $\epsilon_{t}=D_{T V}\left(p_{1}^{t}(s) \| p_{2}^{t}(s)\right)$.

$$
\begin{aligned}
& \left|p_{1}^{t}(s)-p_{2}^{t}(s)\right|=\left|\sum_{s^{\prime}, a} p_{1}\left(s_{t}=s \mid s^{\prime}, a\right) p_{1}^{t+1}\left(s^{\prime}, a\right)-p_{2}\left(s_{t}=s \mid s^{\prime}, a\right) p_{2}^{t+1}\left(s^{\prime}, a\right)\right| \\
& \leq \\
& =\sum_{s^{\prime}, a}\left|p_{1}\left(s_{t}=s \mid s^{\prime}, a\right) p_{1}^{t+1}\left(s^{\prime}, a\right)-p_{2}\left(s_{t}=s \mid s^{\prime}, a\right) p_{2}^{t+1}\left(s^{\prime}, a\right)\right| \\
& =\sum_{s^{\prime}, a} \mid p_{1}\left(s_{t}=s \mid s^{\prime}, a\right) p_{1}^{t+1}\left(s^{\prime}, a\right)-p_{2}\left(s_{t}=s \mid s^{\prime}, a\right) p_{1}^{t+1}\left(s^{\prime}, a\right) \\
& \\
& \quad+p_{2}\left(s_{t}=s \mid s^{\prime}, a\right) p_{1}^{t+1}\left(s^{\prime}, a\right)-p_{2}\left(s_{t}=s \mid s^{\prime}, a\right) p_{2}^{t+1}\left(s^{\prime}, a\right) \mid \\
& \leq \sum_{s^{\prime}, a} p_{1}^{t+1}\left(s^{\prime}, a\right)\left|p_{1}\left(s \mid s^{\prime}, a\right)-p_{2}\left(s \mid s^{\prime}, a\right)\right|+p_{2}\left(s \mid s^{\prime}, a\right)\left|p_{1}^{t+1}\left(s^{\prime}, a\right)-p_{2}^{t+1}\left(s^{\prime}, a\right)\right| \\
& =E_{s^{\prime}, a \sim p_{1}^{t+1}}\left[\left|p_{1}\left(s \mid s^{\prime}, a\right)-p_{2}\left(s \mid s^{\prime}, a\right)\right|\right]+\sum_{s^{\prime}, a} p_{2}\left(s \mid s^{\prime}, a\right)\left|p_{1}^{t+1}\left(s^{\prime}, a\right)-p_{2}^{t+1}\left(s^{\prime}, a\right)\right| \\
& \epsilon_{t}=D_{T V}\left(p_{1}^{t}(s) \| p_{2}^{t}(s)\right) \\
& =
\end{aligned} \begin{aligned}
2 & \sum_{s}\left|p_{1}^{t}(s)-p_{2}^{t}(s)\right| \\
& \leq \frac{1}{2} \sum_{s}\left(E_{s^{\prime}, a \sim p_{1}^{t+1}}\left[\left|p_{1}\left(s \mid s^{\prime}, a\right)-p_{2}\left(s \mid s^{\prime}, a\right)\right|\right]+\sum_{s^{\prime}, a} p_{2}\left(s \mid s^{\prime}, a\right)\left|p_{1}^{t+1}\left(s^{\prime}, a\right)-p_{2}^{t+1}\left(s^{\prime}, a\right)\right|\right) \\
& =E_{s^{\prime}, a \sim p_{1}^{t+1}}\left[D_{T V}\left(p_{1}\left(s \mid s^{\prime}, a\right) \| p_{2}\left(s \mid s^{\prime}, a\right)\right)\right]+D_{T V}\left(p_{1}^{t+1}\left(s^{\prime}, a\right) \| p_{2}^{t+1}\left(s^{\prime}, a\right)\right) \\
& \leq \epsilon_{m}^{\text {back }}+D_{T V}\left(p_{1}^{t+1}\left(s^{\prime}\right) \| p_{2}^{t+1}\left(s^{\prime}\right)\right)+\max _{s^{\prime}} D_{T V}\left(p_{1}\left(a \mid s^{\prime}\right) \| p_{2}\left(a \mid s^{\prime}\right)\right) \\
& =\epsilon_{m}^{\text {back }}+\epsilon_{\pi}^{\text {back }}+D_{T V}\left(p_{1}^{t+1}(s) \| p_{2}^{t+1}(s)\right)
\end{aligned}
$$

Lemma B.2. (Forward State Marginal Distance Bound) ((Janner et al., 2019), Lemma B.2, B.3). Suppose the expected TVD between two forward dynamics is bounded as $\max _{t} E_{(s, a) \sim p_{1}^{t}}\left[D_{T V}\left(p_{1}\left(s^{\prime} \mid s, a\right) \| p_{2}\left(s^{\prime} \mid s, a\right)\right)\right] \leq \epsilon_{m}^{\text {for }}$ and the forward policy divergences are bounded as $\max _{s^{\prime}} D_{T V}\left(\pi_{1}(a \mid s) \| \pi_{2}(a \mid s)\right) \leq \epsilon_{\pi}^{\text {for }}$. Then the state marginal distance at timestep $t$ can be bounded as:

$$
\begin{equation*}
D_{T V}\left(p_{1}^{t}(s) \| p_{2}^{t}(s)\right) \leq \epsilon_{m}^{\text {for }}+\epsilon_{\pi}^{\text {for }}+D_{T V}\left(p_{1}^{t-1}(s) \| p_{2}^{t-1}(s)\right) . \tag{18}
\end{equation*}
$$

Lemma B.3. (TVD Of Joint Distributions) ((Janner et al., 2019), Lemma B.1). Suppose we have two distributions $p_{1}(x, y)=p_{1}(x) p_{1}(y \mid x)$ and $p_{2}(x, y)=p_{2}(x) p_{2}(y \mid x)$. We can bound the total variation distance of the joint distributions as:

$$
\begin{equation*}
D_{T V}\left(p_{1}(x, y) \| p_{2}(x, y)\right) \leq D_{T V}\left(p_{1}(x) \| p_{2}(x)\right)+\max _{x} D_{T V}\left(p_{1}(y \mid x) \| p_{2}(y \mid x)\right) . \tag{19}
\end{equation*}
$$

## C. Environment Settings

In this section, we provide a comparison of the environment settings used in our experiments. Among them, 'Hopper-NT' and 'Walker2d-NT' refer to the settings in Langlois et al. (2019) and others are the standard version.

Table 1. Observation and action dimension, and task horizon of the environments used in our experiments.

| Environment Name | Observation Space Dimension | Action Space Dimension | Steps Per Epoch |
| :---: | :---: | :---: | :---: |
| Pendulum | 3 | 1 | 200 |
| Hopper | 11 | 3 | 1000 |
| Hopper-NT | 11 | 3 | 1000 |
| Walker2d | 17 | 6 | 1000 |
| Walker2d-NT | 17 | 6 | 1000 |
| Ant | 27 | 8 | 1000 |

Table 2. Reward function and termination states condition of the environments used in our experiments. $\theta_{t}$ denotes the joint angle, $x_{t}$ denotes the position in x direction, $a_{t}$ denotes the action control input, and $z_{t}$ denotes the height.

| Environment Name | Reward Function | Termination States Condition |
| :---: | :---: | :---: |
| Pendulum | $-\theta_{t}^{2}-0.1 \dot{\theta}_{t}^{2}-0.001\left\\|a_{t}\right\\|_{2}^{2}$ | None |
| Hopper | $\dot{x}_{t}-0.001\left\\|a_{t}\right\\|_{2}^{2}+1$ | $z_{t} \leq 0.7$ or $\theta_{t} \geq 0.2$ |
| Hopper-NT | $\dot{x}_{t}-0.1\left\\|a_{t}\right\\|_{2}^{2}-3.0 \times\left(z_{t}-1.3\right)^{2}+1$ | None |
| Walker2d | $\dot{x}_{t}-0.001\left\\|a_{t}\right\\|_{2}^{2}+1$ | $z_{t} \leq 0.8$ or $z_{t} \geq 2.0$ or $\left\|\theta_{t}\right\| \geq 1.0$ |
| Walker2d-NT | $\dot{x}_{t}-0.1\left\\|a_{t}\right\\|_{2}^{2}-3.0 \times\left(z_{t}-1.3\right)^{2}+1$ | None |
| Ant | $\dot{x}_{t}-0.5\left\\|a_{t}\right\\|_{2}^{2}+1$ | $z_{t} \leq 0.2$ or $z_{t} \geq 1.0$ |

## D. Hyperparameters

Table 3. Hyperparameter settings for BMPO. $x \rightarrow y$ over epochs $a \rightarrow b$ means clipped linear function, i.e. for epoch e, $f(e)=$ $\operatorname{clip}\left(\left(x+\frac{e-a}{b-a} \cdot(x-y)\right), x, y\right)$. Other hyperparameters not listed here are the same as those in MBPO (Janner et al., 2019).

| Environment Name | $k_{1}$ | $k_{2}$ | $\beta$ | MPC Horizon | Epochs |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Pendulum | $1 \rightarrow 5$ over <br> epochs $1 \rightarrow 5$ | $1 \rightarrow 5$ over <br> epochs $1 \rightarrow 5$ | $0.01 \rightarrow 0$ over <br> epochs $0 \rightarrow 10$ | 6 | 20 |
| Hopper | $1 \rightarrow 15$ over <br> epochs $20 \rightarrow 150$ | $1 \rightarrow 15$ over <br> epochs $20 \rightarrow 150$ | $0.004 \rightarrow 0.003$ over <br> epochs $20 \rightarrow 30$ | 6 | 100 |
| Hopper-NT | $1 \rightarrow 15$ over <br> epochs $20 \rightarrow 150$ | $1 \rightarrow 15$ over <br> epochs $20 \rightarrow 150$ | 0.01 | 6 | 100 |
| Walker2d | 1 | 1 | $0.01 \rightarrow 0$ over <br> epochs $0 \rightarrow 100$ | 1 | 200 |
| Walker2d-NT | 1 | 1 | 0.01 | 0 | 200 |
| Ant | 1 | $1 \rightarrow 25$ over <br> epochs $20 \rightarrow 100$ | 0.003 | 0 | 300 |

## E. Computing Infrastructure

In this section, we provide a description of the computing infrastructure used to run all the experiments in Table 4. We also show the computation time comparison between our algorithm and the MBPO baseline in Table 5.

Table 4. Computing infrastructure.

| CPU | GPU | Memory |
| :---: | :---: | :---: |
| AMD2990WX | RTX2080TI $\times 4$ | 256 GB |

Table 5. Computation time in hours for one experiment.

|  | Pendulum | Hopper | Hopper-NT | Walker2d | Walker2d-NT | Ant |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| BMPO | 0.49 | 16.34 | 17.98 | 27.24 | 27.34 | 71.51 |
| MBPO | 0.41 | 10.33 | 11.12 | 22.26 | 21.32 | 57.42 |

