

A PROOF OF MAIN RESULTS

A.1 PROOF OF THEOREM 1

Theorem 1. *Let the joint policy be the product of the individual policy $\{\pi_i\}_{i=1}^k$, where π_i with respect to the individual flow function $F_i(o_i, a_i)$, i.e.,*

$$\pi_i(a_i|o_i) = \frac{F_i(o_i, a_i)}{F_i(o_i)}, \forall i = 1, \dots, k. \quad (16)$$

Assume that the individual flow $F_i(o_i, a_i)$ satisfies the condition in Definition 2. Define a flow function \hat{F} , if all agents generate trajectories using independent policies π_i , $i = 1, \dots, k$ and the matching conditions

$$\forall s' > s_0, \hat{F}(s') = \sum_{s \in \mathcal{P}(s')} \hat{F}(s \rightarrow s') \text{ and } \forall s' < s_f, \hat{F}(s') = \sum_{s'' \in \mathcal{C}(s')} \hat{F}(s' \rightarrow s'') \quad (17)$$

are satisfied. Then, we have:

- 1) $\pi(s_f) \propto R(s_f)$;
- 2) \hat{F} uniquely defines a Markovian flow F matching \hat{F} such that

$$F(\tau) = \frac{\prod_{t=1}^{n+1} \hat{F}(s_{t-1} \rightarrow s_t)}{\prod_{t=1}^n \hat{F}(s_t)}. \quad (18)$$

Proof: We first prove the part 1). Since

$$F(s_t, \mathbf{a}_t) = \prod_i F_i(o_t^i, a_t^i),$$

then we have the global state flow as

$$F(s_t) = \sum_{\mathbf{a}_t \in \mathcal{A}} F(s_t, \mathbf{a}_t) = \sum_{\mathbf{a}_t \in \mathcal{A}} \prod_i F_i(o_t^i, a_t^i). \quad (19)$$

According to the flow definitions, the observation flow $F_i(o_t^i)$ and individual observation flows have the relationship:

$$F_i(o_t^i) = \sum_{a_t^i \in \mathcal{A}^i} F_i(o_t^i, a_t^i). \quad (20)$$

Hence, we have

$$\prod_{i=1}^k F_i(o_t^i) = \prod_{i=1}^k \left\{ \sum_{a_t^i \in \mathcal{A}^i} F_i(o_t^i, a_t^i) \right\} \quad (21)$$

$$= \sum_{a_t^1 \in \mathcal{A}^1} F_i(o_t^1, a_t^1) \cdots \sum_{a_t^k \in \mathcal{A}^k} F_i(o_t^k, a_t^k) \quad (22)$$

$$= \sum_{a_t^1, \dots, a_t^k \in \mathcal{A}^1 \times \dots \times \mathcal{A}^k} F_i(o_t^1, a_t^1) \cdots F_i(o_t^k, a_t^k) \quad (23)$$

$$= \sum_{\mathbf{a}_t \in \mathcal{A}} \prod_{i=1}^k F_i(o_t^i, a_t^i), \quad (24)$$

yielding $F(s_t) = \prod_i F_i(o_t^i)$. Therefore, the joint policy

$$\begin{aligned} \pi(\mathbf{a}|s) &= \frac{F(s_t, \mathbf{a}_t)}{F(s_t)} = \frac{\prod_i F_i(o_t^i, a_t^i)}{F(s_t)} \\ &= \frac{\prod_i F_i(o_t^i, a_t^i)}{\prod_i F_i(o_t^i)} = \prod_i \pi_i(a_i|o_i). \end{aligned} \quad (25)$$

Equation 25 indicates that if the conditions in Definition 2 is satisfied, we can establish the consistency of joint and individual policies. Based on Lemma 1, we can conclude that the reward of the generated state satisfies $\pi(s_f) \propto R(s_f)$ using the individual policy $\pi_i(a_i|o_i)$ of each agent.

Next, we prove the part 2). We first prove the necessity part. According to Definition 2 and Bengio et al. (2021b) we have

$$\begin{aligned} F(s') &= \prod_i F_i(o^{i'}) = \prod_i \sum_{o^i \in \mathcal{P}(o^{i'})} F_i(o^i \rightarrow o^{i'}) = \sum_{o \in \mathcal{P}(o')} \prod_i F_i(o^i \rightarrow o^{i'}), \\ F(s') &= \prod_i F_i(o^{i'}) = \prod_i \sum_{o^{i''} \in \mathcal{C}(o^{i'})} F_i(o^{i'} \rightarrow o^{i''}) = \sum_{o'' \in \mathcal{C}(o')} \prod_i F_i(o^{i'} \rightarrow o^{i''}). \end{aligned}$$

Then we prove the sufficiency part. We first present Lemma 3, which shows that

$$\sum_{\tau \in \mathcal{T}_{0,s}} P_B(\tau) = \sum_{\tau \in \mathcal{T}_{0,s}} \prod_{s_t \rightarrow s_{t+1} \in \tau} P_B(s_t | s_{t+1}) = 1.$$

Lemma 3 (Independent Transition Probability) *Define the independent forward and backward transition respectively as*

$$P_F(o_{t+1}^i | o_t^i) := P_i(o_t^i \rightarrow o_{t+1}^i | o_t^i) = \frac{F_i(o_t^i \rightarrow o_{t+1}^i)}{F_i(o_t^i)}, \quad (26)$$

and

$$P_B(o_t^i | o_{t+1}^i) := P_i(o_{t+1}^i \rightarrow o_t^i | o_{t+1}^i) = \frac{F_i(o_{t+1}^i \rightarrow o_t^i)}{F_i(o_{t+1}^i)}. \quad (27)$$

Then we have

$$\begin{aligned} \sum_{\tau \in \mathcal{T}_{s,f}} P_F(\tau) &= 1, \forall s \in \mathcal{S} \setminus \{s_f\}, \\ \sum_{\tau \in \mathcal{T}_{0,s}} P_B(\tau) &= 1, \forall s \in \mathcal{S} \setminus \{s_0\}, \end{aligned} \quad (28)$$

where $\mathcal{T}_{s,f}$ is the set of trajectories starting in s and ending in s_f and $\mathcal{T}_{0,s}$ is the set of trajectories starting in s_0 and ending in s .

Define $\hat{Z} = \hat{F}(s_0)$ as the partition function and \hat{P}_F as the forward probability function. Then, according to Proposition 18 in Bengio et al. (2021b), we have there exists a unique Markovian flow F with forward transition probability function $P_F = \hat{P}_F$ and partition function Z , and such that

$$F(\tau) = \hat{Z} \prod_{t=1}^{n+1} \hat{P}_F(s_t | s_{t-1}) = \frac{\prod_{t=1}^{n+1} \hat{F}(s_{t-1} \rightarrow s_t)}{\prod_{t=1}^n \hat{F}(s_t)}, \quad (29)$$

where $s_{n+1} = s_f$. Thus, we have for $s' \neq s_0$:

$$\begin{aligned} F(s') &= \hat{Z} \sum_{\tau \in \mathcal{T}_{0,s'}} \prod_{(s_t \rightarrow s_{t+1}) \in \tau} \hat{P}_F(s_{t+1} | s_t) \\ &= \hat{Z} \frac{\hat{F}(s')}{\hat{F}(s_0)} \sum_{\tau \in \mathcal{T}_{0,s'}} \prod_{(s_t \rightarrow s_{t+1}) \in \tau} \hat{P}_B(s_t | s_{t+1}) = \hat{F}(s'). \end{aligned} \quad (30)$$

Combining equation 30 with $P_F = \hat{P}_F$, we have $\forall s \rightarrow s' \in \mathcal{A}, F(s \rightarrow s')$. Finally, for any Markovian flow F' matching \hat{F} on states and edges, we have $F'(\tau) = F(\tau)$ according to Proposition 16 in Bengio et al. (2021b), which shows the uniqueness property. Then we complete the proof.

A.2 PROOF OF LEMMA 2

Lemma 2. *Suppose Assumption 1 holds and the environment has a tree structure, based on the IGC and IGM conditions we have:*

- 1) $Q_{\text{tot}}^\mu(s, \mathbf{a}) = F(s, \mathbf{a})f(s)$;
- 2) $(\arg \max_{a_i} Q_i(o_i, a_i))_{i=1}^k = (\arg \max_{a_i} F_i(o_i, a_i))_{i=1}^k$.

Proof: The proof is an extension of that of Proposition 4 in Bengio et al. (2021a). For any (s, \mathbf{a}) satisfies $s_f = T(s, \mathbf{a})$, we have $Q_{\text{tot}}^\mu(s, \mathbf{a}) = R(s_f)f(s)$ and $F(s, \mathbf{a}) = R(s_f)$. Therefore, we have $Q_{\text{tot}}^\mu(s, \mathbf{a}) = F(s, \mathbf{a})f(s)$. Then, for each non-final node s' , based on the action-value function in terms of the action-value at the next step, we have by induction:

$$\begin{aligned} Q_{\text{tot}}^\mu(s, \mathbf{a}) &= \hat{R}(s') + \mu(\mathbf{a}|s') \sum_{\mathbf{a}' \in \mathcal{A}(s')} Q_{\text{tot}}^\mu(s', \mathbf{a}'; \hat{R}) \\ &\stackrel{(a)}{=} 0 + \mu(\mathbf{a}|s') \sum_{\mathbf{a}' \in \mathcal{A}(s')} F(s', \mathbf{a}'; R)f(s'), \end{aligned} \quad (31)$$

where $\hat{R}(s')$ is the reward of $Q_{\text{tot}}^\mu(s, \mathbf{a})$ and (a) is due to that $\hat{R}(s') = 0$ if s' is not a final state. Since the environment has a tree structure, we have

$$F(s, \mathbf{a}) = \sum_{\mathbf{a}' \in \mathcal{A}(s')} F(s', \mathbf{a}'), \quad (32)$$

which yields

$$Q_{\text{tot}}^\mu(s, \mathbf{a}) = \mu(\mathbf{a}|s')F(s, \mathbf{a})f(s') = \mu(\mathbf{a}|s')F(s, \mathbf{a})f(s) \frac{1}{\mu(\mathbf{a}|s')} = F(s, \mathbf{a})f(s).$$

According to the IGC condition we have $F(s_t, \mathbf{a}_t) = \prod_i F_i(o_t^i, a_t^i)$, yielding

$$\begin{aligned} \arg \max_{\mathbf{a}} Q_{\text{tot}}(s, \mathbf{a}) &\stackrel{(a)}{=} \arg \max_{\mathbf{a}} \log F(s, \mathbf{a})f(s) \\ &\stackrel{(b)}{=} \arg \max_{\mathbf{a}} \sum_{i=1}^k \log F_i(o_i, a_i) \\ &\stackrel{(c)}{=} \left(\arg \max_{a_1 \in \mathcal{A}_1} F_1(o_1, a_1), \dots, \arg \max_{a_k \in \mathcal{A}_k} F_k(o_k, a_k) \right), \end{aligned} \quad (33)$$

where (a) is based on the fact F and $f(s)$ are positive, (b) is due to the IGC condition. Combining with the IGM condition

$$\arg \max_{\mathbf{a} \in \mathcal{A}} Q_{\text{tot}}(s, \mathbf{a}) = \left(\arg \max_{a_1 \in \mathcal{A}_1} Q_1(o_1, a_1), \dots, \arg \max_{a_k \in \mathcal{A}_k} Q_k(o_k, a_k) \right), \forall s \in \mathcal{S}. \quad (34)$$

we can conclude that

$$\left(\arg \max_{a_i \in \mathcal{A}_i} F_i(o_i, a_i) \right)_{i=1}^k = \left(\arg \max_{a_i \in \mathcal{A}_i} Q_i(o_i, a_i) \right)_{i=1}^k.$$

Then we complete the proof.

A.3 PROOF OF LEMMA 3

Lemma 3 [Independent Transition Probability]. *Define the independent forward and backward transition respectively as*

$$P_F(o_{t+1}^i | o_t^i) := P_i(o_t^i \rightarrow o_{t+1}^i | o_t^i) = \frac{F_i(o_t^i \rightarrow o_{t+1}^i)}{F_i(o_t^i)}, \quad (35)$$

and

$$P_B(o_t^i | o_{t+1}^i) := P_i(o_{t+1}^i \rightarrow o_t^i | o_{t+1}^i) = \frac{F_i(o_{t+1}^i \rightarrow o_t^i)}{F_i(o_{t+1}^i)}. \quad (36)$$

Then we have

$$\begin{aligned} \sum_{\tau \in \mathcal{T}_{s,f}} P_F(\tau) &= 1, \forall s \in \mathcal{S} \setminus \{s_f\}, \\ \sum_{\tau \in \mathcal{T}_{0,s}} P_B(\tau) &= 1, \forall s \in \mathcal{S} \setminus \{s_0\}, \end{aligned} \quad (37)$$

where $\mathcal{T}_{s,f}$ is the set of trajectories starting in s and ending in s_f and $\mathcal{T}_{0,s}$ is the set of trajectories starting in s_0 and ending in s .

Proof: When the maximum length of trajectories is not more than 1, we have

$$\sum_{\tau \in \mathcal{T}_{s,f}} P_F(\tau) = 1. \quad (38)$$

Then we have the following results by induction:

$$\begin{aligned} \sum_{\tau \in \mathcal{T}_{s,f}} P_F(\tau) &= \sum_{s' \in \mathcal{C}(s)} \sum_{\tau \in \mathcal{T}_{s' \rightarrow s',f}} P_F(\tau) = \sum_{\mathbf{o}' \in \mathcal{C}(\mathbf{o})} P_F(\mathbf{o}'|\mathbf{o}) \sum_{\tau \in \mathcal{T}_{s',f}} P_F(\tau) \\ &= \sum_k \sum_{\mathbf{o}'_i \in \mathcal{C}(\mathbf{o}_i)} P_F(\mathbf{o}'_i|\mathbf{o}_i) \sum_{\tau \in \mathcal{T}_{s',f}} P_F(\tau) = 1, \end{aligned} \quad (39)$$

where $\mathcal{C}(\cdot)$ is the children set of the current state or observation and the last equation is based on the fact $\sum_{\mathbf{o}'_i \in \mathcal{C}(\mathbf{o}_i)} P_F(\mathbf{o}'_i|\mathbf{o}_i) = 1$. Since the proof process of P_B is similar to that of P_F , it is omitted here.

B EXPERIMENTAL DETAILS

B.1 HYPER-GRID ENVIRONMENT

Here we present the experimental details on the Hyper-Grid environments. Figure 5 shows the curve of the flow matching loss function with the number of training steps. The loss of our proposed algorithm gradually decreases, ensuring the stability of the learning process. For some RL algorithms based on the state-action value function estimation, the loss usually oscillates. This may be because RL-based methods use experience replay buffer and the transition data distribution is not stable enough. The method we propose uses an on-policy based optimization method, and the data distribution changes with the current sampling policy, hence the loss function is relatively stable. We set the same number of training steps for all algorithms for a fair comparison. Moreover, we list the key hyperparameters of the different algorithms in Tables 2 3 4 5.

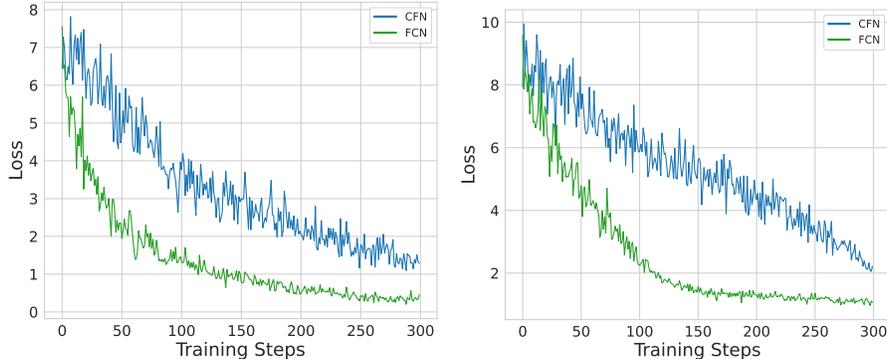


Figure 5: The flow matching loss of different algorithm.

We study the effect of different reward in Figure 6. In particular, we set $R_0 = \{10^{-1}, 10^{-2}, 10^{-4}\}$ for different task challenge. A smaller value of R_0 makes the reward function distribution more

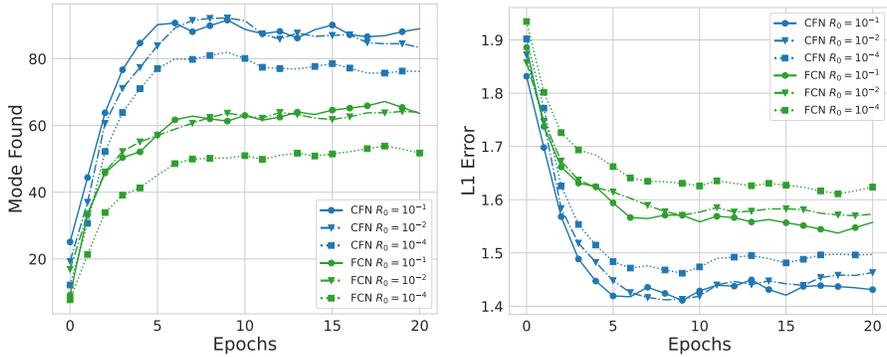


Figure 6: The effect of different reward R_0 on different algorithm according to L1 error and mode found.

sparse, which makes policy optimization more difficult Bengio et al. (2021a); Riedmiller et al. (2018); Trott et al. (2019). As shown in Figure 6, we found that our proposed method is robust with the cases $R_0 = 10^{-1}$ and $R_0 = 10^{-2}$. When the reward distribution becomes sparse, the performance of the proposed algorithm degrades slightly.

Table 2: Hyper-parameter of MAPPO under different environments

	Hyper-Grid-v1	Hyper-Grid-v2	Hyper-Grid-v3
Train Steps	20000	20000	20000
Agent	2	2	3
Grid Dim	2	3	3
Grid Size	[8,8]	[8,8]	[8,8]
Actor Network Hidden Layers	[256,256]	[256,256]	[256,256]
Optimizer	Adam	Adam	Adam
Learning Rate	0.0001	0.0001	0.0001
Batchsize	64	64	64
Discount Factor	0.99	0.99	0.99
PPO Entropy	1e-1	1e-1	1e-1

Table 3: Hyper-parameter of MASAC under different environments

	Hyper-Grid-v1	Hyper-Grid-v2	Hyper-Grid-v3
Train Steps	20000	20000	20000
Grid Dim	2	3	3
Grid Size	[8,8]	[8,8]	[8,8]
Actor Network Hidden Layers	[256,256]	[256,256]	[256,256]
Critic Network Hidden Layers	[256,256]	[256,256]	[256,256]
Optimizer	Adam	Adam	Adam
Learning Rate	0.0001	0.0001	0.0001
Batchsize	64	64	64
Discount Factor	0.99	0.99	0.99
SAC Alpha	0.98	0.98	0.98
Target Network Update	0.001	0.001	0.001

Table 4: Hyper-parameter of FCN under different environments

	Hyper-Grid-v1	Hyper-Grid-v2	Hyper-Grid-v3
Train Steps	20000	20000	20000
R_2	2	2	2
R_1	0.5	0.5	0.5
Grid Dim	2	3	3
Grid Size	[8,8]	[8,8]	[8,8]
Trajectories per steps	16	16	16
Flow Network Hidden Layers	[256,256]	[256,256]	[256,256]
Optimizer	Adam	Adam	Adam
Learning Rate	0.0001	0.0001	0.0001
ϵ	0.0005	0.0005	0.0005

Table 5: Hyper-parameter of CFN under different environments

	Hyper-Grid-v1	Hyper-Grid-v2	Hyper-Grid-v3
Train Steps	20000	20000	20000
Trajectories per steps	16	16	16
R_2	2	2	2
R_1	0.5	0.5	0.5
Grid Dim	2	3	3
Grid Size	[8,8]	[8,8]	[8,8]
Flow Network Hidden Layers	[256,256]	[256,256]	[256,256]
Optimizer	Adam	Adam	Adam
Learning Rate	0.0001	0.0001	0.0001
ϵ	0.0005	0.0005	0.0005