A Proofs for §5

A.1 Proofs of Theorem 1

Proof of Theorem 1. Assume that MF-MI-Greedy terminates within k episodes. Let us use $\mathcal{E}^{(1)}, \ldots, \mathcal{E}^{(k)}$ to denote the sequence of actions selected by MF-MI-Greedy, where $\mathcal{E}^{(j)} := \mathcal{L}^{(j)} \cup \{\langle x, m\rangle^{(j)}\}$ denotes the sequence of actions selected at the j^{th} episode. Further, let $\Lambda_{\mathcal{E}}^{(j)}$ be the cost of the j^{th} episode, and $\Lambda_{\mathcal{L}}^{(j)} = \Lambda_{\mathcal{E}}^{(j)} - \lambda_m$ the cost of lower fidelity actions of the j^{th} episode. The budget allocated for the target fidelity is $k\lambda_m = \Lambda - \sum_{j=1}^{k} \Lambda_{\mathcal{L}}^{(j)}$. By definition of the cumulative regret (Eq. (5.2)), we get

$$\begin{aligned}
R(\pi, \Lambda) &= \frac{\Lambda}{\lambda_m} f_m^* - \sum_{j=1}^{k} q(\mathcal{L}^{(j)} \cup \{\langle x, m\rangle^{(j)}\}) \\
&= \frac{\Lambda}{\lambda_m} f_m^* - \left(\sum_{j=1}^{k} \cancelto{0}{q(\mathcal{L}^{(j)})} + \sum_{j=1}^{k} q(\langle x, m\rangle^{(j)}) \right) \\
&= \frac{\Lambda}{\lambda_m} f_m^* - \sum_{j=1}^{k} f_m\left(x^{(j)}\right) \\
&= \left(\frac{\Lambda}{\lambda_m} - k\right) f_m^* + \sum_{j=1}^{k} \left(f_m^* - f_m\left(x^{(j)}\right)\right) \\
&= \frac{f_m^*}{\lambda_m} \cdot \sum_{j=1}^{k} \Lambda_{\mathcal{L}}^{(j)} + \sum_{j=1}^{k} \left(f_m^* - f_m\left(x^{(j)}\right)\right)
\end{aligned} \tag{A.1}$$

The first term on the R.H.S. of Eq. (A.1) represents the regret incurred from exploring the lower fidelity actions, while the second term represents the regret from the target fidelity actions (chosen by SF-GP-OPT).

According to the stopping condition of Algorithm 2 at Line 8, we know that when Explore-LF terminates at episode j, the selected low fidelity actions $\mathcal{L}$ satisfy

$$\frac{\mathbb{I}\left(\mathbf{y}_{\mathcal{L}}^{(j)}; f_m \,\middle|\, \mathbf{y}_{\mathcal{E}}^{(1:j-1)}\right)}{\Lambda_{\mathcal{L}}^{(j)}} \geq \beta_j,$$

where $\mathbf{y}_{\mathcal{E}}^{(1:j)} := \bigcup_{u=1}^{j} \mathbf{y}_{\mathcal{E}}^{(u)}$ denotes the observations obtained up to episode j, and β_j specifies the stopping condition of Explore-LF at episode j. Therefore

$$\begin{aligned}
\sum_{j=1}^{k} \Lambda_{\mathcal{L}}^{(j)} &\leq \sum_{j=1}^{k} \frac{\mathbb{I}\left(\mathbf{y}_{\mathcal{L}}^{(j)}; f_m \,\middle|\, \mathbf{y}_{\mathcal{E}}^{(1:j-1)}\right)}{\beta_j} \\
&\overset{(a)}{\leq} \alpha_\Lambda \sum_{j=1}^{k} \mathbb{I}\left(\mathbf{y}_{\mathcal{L}}^{(j)}; f_m \,\middle|\, \mathbf{y}_{\mathcal{E}}^{(1:j-1)}\right) \\
&= \alpha_\Lambda \gamma_{\mathcal{L}}
\end{aligned} \tag{A.2}$$

where step (a) is because $\alpha_\Lambda = \max_B \alpha(B) > \frac{1}{\beta_j}$ for $j \in [k]$. Recall that $\beta_1 = \frac{1}{o(\sqrt{\Lambda})}$. Therefore,

$$\frac{f_m^*}{\lambda_m} \cdot \sum_{j=1}^{k} \Lambda_{\mathcal{L}}^{(j)} \leq \frac{f_m^*}{\lambda_m} \cdot \alpha_\Lambda \gamma_{\mathcal{L}}. \tag{A.3}$$

Note that the second term of Eq. (A.1) is the regret of MF-MI-Greedy on the target fidelity. Since all the target fidelity actions are selected by the subroutine SF-GP-OPT, by assumption, we know $\sum_{j=1}^{k} \left(f_m^* - f_m\left(x^{(j)}\right)\right) = \sqrt{C\gamma_m k \lambda_m} \leq \sqrt{C\gamma_m \Lambda}$. Combining this with Eq. (A.3) completes the proof. □

A.2 Proof of Corollary 2

To show that running MF-MI-Greedy with subroutine GP-UCB (Srinivas et al., 2010), EST (Wang et al., 2016), or GP-MI (Contal et al., 2014) in the optimization phase is no-regret, it suffices to show that the candidate subroutines GP-UCB, EST, and GP-MI satisfy the assumption on SF-GP-OPT as provided in Theorem 1. From the references above we know that the statement is true.