Learning from diverse feedback types

- demonstrations: ξ_D
- comparisons (preferences): ξ_A or ξ_B ?
- physical corrections $\tilde{c}$
- turning the robot ξ_I ①
- saying something λ = "don't step on the carpet"
- specifying a reward? $\tilde{\theta}$
- the current state of the world s_0

$$b'(\theta) \propto b(\theta) \underbrace{P(\text{feedback} \mid \theta)}_{\text{human model}}$$

choices $c \in C$

choose based on reward

$R_\theta(c)$? ×

$R_\theta(\psi(c))$

↑ grounding of c into traj

$\psi : c \mapsto \xi$

or $\psi : c \mapsto \Delta$ $\mathbb{E}[R_\theta(\xi) \mid \xi \sim \psi(c)]$

$$P(C^* \mid \theta) = \frac{e^{R_\theta(\psi(C^*))}}{\sum_{C \in \mathcal{C}} e^{R_\theta(\psi(C))}}$$

demonstrations

$\mathcal{C} = \{\xi\}$ $\psi(\xi) = \xi$ $P(\xi_D \mid \theta) = \frac{e^{R_\theta(\xi_D)}}{\sum_\xi e^{R_\theta(\xi)}}$

comparisons

$\mathcal{C} = \{\xi_A, \xi_B\}$ $\psi(\xi) = \xi$ $P(\xi_A \mid \theta) = \frac{e^{R_\theta(\xi_A)}}{e^{R_\theta(\xi_A)} + e^{R_\theta(\xi_B)}}$

$= \frac{1}{1 + e^{R_\theta(\xi_B) - R_\theta(\xi_A)}}$

$\mathcal{Q}$ switch

$\mathcal{C} = \{0, 1\}$ $\psi(-) = \xi_R$ $\psi(0) = \xi_D^{0:t} \xi_D^{t} \cdots \xi_R^{t}$

corrections

$\mathcal{C} = \{\Delta q\}$ $\psi(\Delta q) = \xi_R + A^{-1} \Delta q$

(proxy) specified reward

$\mathcal{C} = \{\theta\}$ $\psi(\theta) \sim P(\xi \mid \theta, \underline{M_{device}})$

current state

$$C = \{ s \} \quad \psi(s) \sim \mathrm{Unif}\left(S_{t+1}^{-T:0} \mid S_t^0 = s \right)$$