

377 8 Supplementary Material

378 8.1 Background on Linearly Solvable MDP

379 Since the Reference-Based POMDP expands the Linearly Solvable (fully observed) MDPs [18, 19, 20]
380 to POMDPs, for completeness, here, we summarise Linearly Solvable MDPs.

381 A standard infinite horizon MDP is specified by tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, R, \gamma \rangle$, where \mathcal{S} and \mathcal{A} are the state
382 and action spaces, $\mathcal{T}(s, a, s')$ is the conditional probability function $P(s' | s, a)$ that specifies the
383 probability the agent arrives at state $s' \in \mathcal{S}$ after performing action $a \in \mathcal{A}$ at state $s \in \mathcal{S}$, R is the
384 reward function, and $\gamma \in (0, 1)$ is the discount factor. The solution to an MDP problem is a an
385 optimal policy $\pi^* : \mathcal{S} \rightarrow \mathcal{A}$ that maximises the value function:

$$V^*(s) = \max_{a \in \mathcal{A}} \left[R(s, a) + \gamma \sum_{s' \in \mathcal{S}} \mathcal{T}(s, a, s') V^*(s') \right] \quad (30)$$

386 The works in [18, 19, 20] consider a class of MDPs where, the state space \mathcal{S} is finite and for any states
387 $s, s' \in \mathcal{S}$, there exists a one-step (not necessarily time-homogeneous) transition probability $p(s' | s)$
388 representing the *passive dynamics* of the system. They propose a new formulation of MDPs, called
389 Linearly Solvable MDPs, to be specified by $\langle \mathcal{S}, p, r, \gamma \rangle$, where $r : \mathcal{S} \rightarrow \mathbb{R}$ is the reward function. A
390 solution to the Linearly Solvable MDP is a stochastic state-to-state transition probability $u(\cdot | s)$ that
391 maximises:

$$v(s) = \sup_{u(\cdot | s) \in \mathcal{U}_p(s)} \left(r(s) - \text{KL}(u(\cdot | s) \| p(\cdot | s)) \right) + \gamma \sum_{s' \in \mathcal{S}} u(s' | s) v(s'). \quad (31)$$

392 where $\mathcal{U}_p(s)$ is the set of admissible controls. An admissible control $u(\cdot | s)$ is one that prohibits state
393 transitions that are not feasible under the passive dynamics $p(\cdot | s)$.

394 Now, suppose $w(s) := e^{v(s)}$ for any $s \in \mathcal{S}$, then (31) is equivalent to

$$w(s) = e^{r(s)} \sum_{s' \in \mathcal{S}} p(s' | s) w^\gamma(s'). \quad (32)$$

395 Moreover, the solution w^* to the above equation exists and is unique. The optimal stochastic transition
396 to the equation (31) is given by

$$u^*(\cdot | s) = \frac{p(\cdot | s) w^{*\gamma}(\cdot)}{D[w^{*\gamma}](s)}. \quad (33)$$

397 where $D[w^{*\gamma}](s) := \sum_{s' \in \mathcal{S}} p(s' | s) w^{*\gamma}(s')$ is a normaliser. Intuitively, one can view w^* as the
398 desirability score, so that (33) represents distorting the passive dynamics to transition dynamics that
399 favour transitioning to states with higher desirability scores. Of course, w^* is not known a priori
400 but it can be determined by iterating the Bellman backup operator given by (32). This computation
401 essentially reduces to taking expectations under the reference dynamics, which can be computed
402 faster than searching for the optimal value function in (30) directly.

403 A standard MDP can be *embedded* in a linearly solvable MDP. This implies that, for a given standard
404 MDP problem $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, R, \gamma \rangle$, one can embed it as an instance of a linearly solvable MDP, use
405 the above efficient machinery to determine the solution to the linearly solvable MDP $u^*(\cdot | s)$, and
406 then choose the symbolic action $a^* \in \mathcal{A}$ such that $\mathcal{T}(s' | s, a^*)$ is as close as possible to $u^*(\cdot | s)$.
407 Empirical results in [19] indicate that there is a close correspondence between the optimal value of
408 the embedded standard MDP $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, R, \gamma \rangle$ and the optimal value of the linearly solvable MDP.

409 8.2 Proof of Lemma 3.1

410 *Step 1.* We first need to verify that a maximiser to the supremum in (10) exists. To this end, define
411 $\mathcal{W}(b) := e^{\mathcal{V}(b)}$ for any $b \in \mathcal{B}$ and notice that the terms inside the supremum in the RHS of equation

412 (10) can be rewritten as

$$\begin{aligned}
& \sum_{a,o} \mathbb{U}(a,o|b) \left[\mathcal{R}(b,a) - \log \left\{ \frac{\mathbb{U}(a,o|b)}{\mathbb{U}^{\mathbb{P}}(a,o|b)} \right\} \right. \\
& \quad \left. + \gamma \sum_{a,o} \mathcal{V}(\tau(b,a,o)) \right] \\
&= - \sum_{a,o} \mathbb{U}(a,o|b) \left[\log \left\{ \frac{\mathbb{U}(a,o|b)}{\mathbb{U}^{\mathbb{P}}(a,o|b) e^{\mathcal{R}(b,a)} \mathcal{W}^{\gamma}(\tau(b,a,o))} \right\} \right] \\
&= - \sum_{a,o} \mathbb{U}(a,o|b) \left[\log \left\{ \frac{\mathbb{U}(a,o|b) \mathcal{D}[\mathcal{W}^{\gamma}](b)}{\mathbb{U}^{\mathbb{P}}(a,o|b) e^{\mathcal{R}(b,a)} \mathcal{W}^{\gamma}(\tau(b,a,o))} \right\} \right. \\
& \quad \left. - \log \{ \mathcal{D}[\mathcal{W}^{\gamma}](b) \} \right] \\
&= - \text{KL} \left(\mathbb{U}(\cdot, \cdot | b) \parallel \frac{\mathbb{U}^{\mathbb{P}}(\cdot, \cdot | b) e^{\mathcal{R}(b,a)} \mathcal{W}^{\gamma}(\tau(b,a,o))}{\mathcal{D}[\mathcal{W}^{\gamma}](b)} \right) \\
& \quad + \log \{ \mathcal{D}[\mathcal{W}^{\gamma}](b) \} \quad (34)
\end{aligned}$$

413 where $\mathcal{D}[\mathcal{W}^{\gamma}](b) := \sum_{a,o} \mathbb{U}^{\mathbb{P}}(a,o|b) e^{\mathcal{R}(b,a)} \mathcal{W}^{\gamma}(\tau(b,a,o))$ is a normalising factor. Only the KL
414 divergence term in the last line above depends on \mathbb{U} . We know that the KL divergence is minimised
415 when its two component distributions are identical. That is, when

$$\mathbb{U}^*(a,o|b) = \frac{\mathbb{U}^{\mathbb{P}}(a,o|b) e^{\mathcal{R}(b,a)} \mathcal{W}^{\gamma}(\tau(b,a,o))}{\mathcal{G}(b)}. \quad (35)$$

416 It is clear that \mathbb{U}^* belongs to the space $\mathcal{Z}_{\mathbb{P}}(b)$ since $\mathbb{U}^{\mathbb{P}}(a,o|b) = 0$ implies that $\mathbb{U}^*(a,o|b) = 0$ too.
417 Therefore, we conclude that the supremum is attained and that \mathbb{U}^* is the maximiser.

418 *Step 2.* Now, we can essentially repeat the classical argument from Ross [15] (see e.g. Theorem 6.5).
419 Namely, let $\Phi : \mathbb{B}(\mathcal{B}) \rightarrow \mathbb{B}(\mathcal{B})$ be the Bellman backup operator

$$\begin{aligned}
\Phi \mathcal{V}(b) := & \sup_{\mathbb{U} \in \mathcal{Z}_{\mathbb{P}}(b)} \left(\mathcal{R}(b, \mathbb{U}) - \text{KL}(\mathbb{U} \parallel \mathbb{U}^{\mathbb{P}}) \right. \\
& \left. + \gamma \mathbb{E}_{\mathbb{U}}[\mathcal{V}(\tau, \cdot, \cdot)] \right) \quad \forall b \in \mathcal{B} \quad (36)
\end{aligned}$$

420 where, for brevity, we write

$$\text{KL}(\mathbb{U} \parallel \mathbb{U}^{\mathbb{P}}) := \text{KL}(\mathbb{U}(\cdot, \cdot | b) \parallel \mathbb{U}^{\mathbb{P}}(\cdot, \cdot | b)) \quad (37)$$

421 and

$$\mathbb{E}_{\mathbb{U}}[\mathcal{V}(\tau, \cdot, \cdot)] := \sum_{a,o} \mathbb{U}(a,o|b) \mathcal{V}(\tau(b,a,o)). \quad (38)$$

422 We want to show that Φ is a contraction. For any $b \in \mathcal{B}$ and any $\mathcal{V}_1, \mathcal{V}_2 \in \mathbb{B}(\mathcal{B})$,

$$\begin{aligned}
& (\Phi \mathcal{V}_1)(b) - (\Phi \mathcal{V}_2)(b) \\
&= \sup_{\mathbb{U} \in \mathcal{Z}_{\mathbb{P}}(b)} \left(\mathcal{R}(b, \mathbb{U}) - \text{KL}(\mathbb{U} \parallel \mathbb{U}^{\mathbb{P}}) + \gamma \mathbb{E}_{\mathbb{U}}[\mathcal{V}_1(\tau, \cdot, \cdot)] \right) \\
& \quad - \sup_{\tilde{\mathbb{U}} \in \mathcal{Z}_{\mathbb{P}}(b)} \left(\mathcal{R}(b, \tilde{\mathbb{U}}) - \text{KL}(\tilde{\mathbb{U}} \parallel \mathbb{U}^{\mathbb{P}}) + \gamma \mathbb{E}_{\tilde{\mathbb{U}}}[\mathcal{V}_2(\tau, \cdot, \cdot)] \right) \\
& \leq \left(\mathcal{R}(b, \mathbb{U}^*) - \text{KL}(\mathbb{U}^* \parallel \mathbb{U}^{\mathbb{P}}) + \gamma \mathbb{E}_{\mathbb{U}^*}[\mathcal{V}_1(\tau, \cdot, \cdot)] \right) \\
& \quad - \left(\mathcal{R}(b, \mathbb{U}^*) - \text{KL}(\mathbb{U}^* \parallel \mathbb{U}^{\mathbb{P}}) + \gamma \mathbb{E}_{\mathbb{U}^*}[\mathcal{V}_2(\tau, \cdot, \cdot)] \right) \\
&= \gamma \sum_{a,o} \mathbb{U}^*(a,o|b) \left[\mathcal{V}_1(\tau(b,a,o)) - \mathcal{V}_2(\tau(b,a,o)) \right] \\
& \leq \gamma \|\mathcal{V}_1 - \mathcal{V}_2\|_{\infty} \quad (39)
\end{aligned}$$

423 where \mathbb{U}^* is the maximiser of

$$\mathcal{R}(b, \mathbb{U}) - \text{KL}(\mathbb{U} \parallel \mathbb{U}^{\mathfrak{p}}) + \gamma \mathbb{E}_{\mathbb{U}}[\mathcal{V}_1(\tau, \cdot, \cdot)]. \quad (40)$$

424 Reversing the roles of \mathcal{V}_1 and \mathcal{V}_2 and using the fact that $b \in \mathcal{B}$ is arbitrary, we conclude that

$$\|\Phi \mathcal{V}_1 - \Phi \mathcal{V}_2\|_{\infty} \leq \gamma \|\mathcal{V}_1 - \mathcal{V}_2\|_{\infty}. \quad (41)$$

425 Since we assumed that $\gamma \in (0, 1)$, we conclude that Φ is a contraction.

426 8.3 Proof of Theorem 3.1

427 Repeating the argument in Step 1 of 8.2, we see that the Bellman equation (10) reduces to

$$\begin{aligned} \mathcal{V}(b) &= \log [\mathcal{D}[w^{\gamma}](b)] \\ &= \log \left[\sum_{a,o} \mathbb{U}^{\mathfrak{p}}(a, o | b) e^{\mathcal{R}(b,a)} \mathcal{W}^{\gamma}(\tau(b, a, o)) \right] \end{aligned} \quad (42)$$

428 which, after taking exponents, justifies the equivalence to (12). Given this equivalence and Lemma
 429 3.1, it is clear that (12) has a unique solution. To be more explicit, suppose for a contradiction that
 430 (12) does not have exactly one solution (up to $\|\cdot\|$ -equivalence of solutions). Then by the equivalence
 431 between the two Bellman equations, (10) would either have no solutions or more than one solution
 432 which contradicts the existence and uniqueness guaranteed by Lemma 3.1. Finally, (13) follows from
 433 the form of the maximiser at each Bellman step.

434 8.4 Proof of Proposition 3.1

435 For brevity, we will fix a $b \in \mathcal{B}$ and drop it from our notation. Also write $u = u(\cdot | b)$ and $u_a = u(a | b)$.
 436 The Lagrangian for the constrained problem is

$$\begin{aligned} \mathcal{L}(u, \lambda) &= \sum_{a,o} P(o | a) u_a \log \left[\frac{P(o | a) u_a}{\mathbb{U}^*(a, o)} \right] \\ &\quad + \lambda \left(\sum_a u_a - 1 \right). \end{aligned} \quad (43)$$

437 We require, in addition, that the minimiser u^* (which exists due to the Weierstrass extreme value
 438 theorem) is such that $u_a^* \geq 0$ for each $a \in \mathcal{A}$. The first order necessary conditions gives

$$u_a = e^{-(1+\lambda)} \exp[-\Pi(a)] \quad \forall a \in \mathcal{A} \quad (44)$$

439 and the constraint equation gives

$$1 = \sum_a u_a = e^{-(1+\lambda)} \sum_a \exp[-\Pi(a)]. \quad (45)$$

440 Hence the only candidate for the minimiser is u^* such that

$$u_a^* = \frac{\exp[-\Pi(a)]}{\sum_{\hat{a} \in \mathcal{A}} \exp[-\Pi(\hat{a})]} \quad \forall a \in \mathcal{A}. \quad (46)$$

441 The Hessian of \mathcal{L} is positive definite for any λ and $u \in \Delta(\mathcal{A})$, so we conclude that u^* is a minimiser.
 442 Finally, that $u_a^* \geq 0$ for every $a \in \mathcal{A}$ is clear from (46).

443 8.5 Proof of Proposition 4.1

444 *Proof.* Fix an $\hat{a} \in \mathcal{A}$ and $b \in \mathcal{B}$. Clearly \mathfrak{p} as defined in (24) has full support on \mathcal{A} . Thus, if we set

$$u^{\hat{a}}(a | b) := \begin{cases} 1, & a = \hat{a} \\ 0, & \text{otherwise} \end{cases} \quad (47)$$

445 and

$$\mathbb{U}^{\hat{a}}(a, o | b) := P(o | \hat{a}, b) u^{\hat{a}}(a | b) \in \mathcal{U}_{\mathfrak{p}} \quad (48)$$

446 for any $u^{\hat{a}} \in \Delta(\mathcal{A})$ then the constraint (21) is satisfied trivially. Straightforward computations show
 447 that the constraint (22) is satisfied by $\mathbb{U}^{\hat{a}}(a, o | b)$ with (ρ, \mathfrak{p}) as defined in (23) and (24). \square

448 **8.6 Proof of Proposition 4.2**

449 *Proof.* For a fixed $\delta > 0$, the matrix equation (25) has a solution \mathbf{x}_δ if and only if

$$\sum_{s \in \mathcal{S}} b(s) \rho_\delta(s) = \ell(b) \quad \forall b \in C_\delta. \quad (49)$$

450 As $\delta \downarrow 0$ the δ -cover converges to the set \mathcal{R} which proves that the pair (ρ, \mathbf{p}) is an embedding. \square

451 **8.7 Algorithm REFSOLVER**

Algorithm 1 REFSOLVER

parameters:
 $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma \rangle$
max-depth
max-rollout-depth
 $\alpha \triangleright \text{expl const} = 1 - \alpha$
require: $\gamma \in (0, 1), \alpha \in [0, 1]$

PRE-PROCESS (OFFLINE)

1: $\pi^{\text{FO}} \leftarrow \text{GENERATE-FO-POLICY}(\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma \rangle)$

2: **RUNTIME (ONLINE)**

3: **procedure** PLAN-AND-EXECUTE(h)
4: **repeat**
5: **if** $h = \emptyset$ **then**
6: $s \sim \mathcal{I}$
7: **else**
8: $s \sim \mathcal{B}(h)$
9: **end if**
10: SIMULATE($s, h, 0$)
11: **until** TIMEOUT()
12: **return** OPTIMAL-STOCHASTIC-POLICY(h)
13: **end procedure**

14: **procedure** ROLLOUT(s, h, depth)
15: $a \leftarrow \pi^{\text{FO}}(s)$
16: **if** $s \in \mathcal{G}$ or $\text{depth} > \text{max-rollout-depth}$ **then**
17: **return** $\mathcal{R}(s, a)$
18: **end if**
19: $(s', o, \mathcal{R}) \sim \mathcal{G}(s, a) \triangleright \text{generative model}$
20: **return** $\mathcal{R}(s, a) + \text{ROLLOUT}(s', hao, \text{depth} - 1)$
21: **end procedure**

22: **procedure** SIMULATE(s, h, depth)
23: **if** $s \in \mathcal{G}$ or $\text{depth} > \text{max-depth}$ **then**
24: **return** $\text{exp}(\text{ROLLOUT}(s, h, \text{max-depth}))$
25: **end if**
26: $\mathcal{B}(h) \leftarrow \mathcal{B}(h) \cup \{s\}$
27: $N(h) \leftarrow N(h) + 1$
28: $X \sim \text{Bernoulli}(\alpha)$
29: $a \leftarrow \pi^{\text{FO}}(s) I_{\{X=1\}} + (1 - \alpha) \times I_{\{X=0\}}$
30: $(s', o, \mathcal{R}) \sim \mathcal{G}(s, a)$
31: $N(ha) \leftarrow N(ha) + 1$
32: $\widehat{\mathcal{R}}(ha) \leftarrow \widehat{\mathcal{R}}(ha) + \frac{\mathcal{R}(s,a) - \widehat{\mathcal{R}}(ha)}{N(ha)}$
33: $\widehat{\mathcal{W}} \leftarrow \widehat{\mathcal{W}} + \frac{e^{\widehat{\mathcal{R}}(ha)} \text{SIMULATE}(s', hao, \text{depth} + 1) - \widehat{\mathcal{W}}(h)}{N(h)}$
34: **return** $\widehat{\mathcal{W}}(h)^\gamma$
35: **end procedure**

36: **procedure** OPTIMAL-STOCHASTIC-POLICY(h)
37: $\mathcal{D} \leftarrow 0 \triangleright \text{Normaliser}$
38: **for** $a \in \mathcal{A}$ and $o \in \mathcal{O}$ **do**
39: **if** $hao \notin \mathcal{T}$ **then**
40: $\widehat{\mathcal{U}}^*(hao) = 0$
41: **else**
42: $\widehat{\mathcal{U}}^*(hao) \leftarrow \frac{N(hao)}{N(h)} e^{\widehat{\mathcal{R}}(hao)} \mathcal{W}^\gamma(hao)$
43: **end if**
44: $\mathcal{D} \leftarrow \mathcal{D} + \widehat{\mathcal{U}}^*(hao)$
45: **end for**
46: **for** $a \in \mathcal{A}$ **do**
47: $\Pi(a) \leftarrow \frac{N(hao)}{N(ha)} \log \left[\frac{N(hao) \mathcal{D}}{N(ha) \widehat{\mathcal{U}}^*(hao)} \right]$
48: **end for**
49: $\mathbf{u}^* \leftarrow \{a : \exp[-\Pi(a)] / \mathcal{D}^\Pi\}$
50: **return** RANDOM-SAMPLE(\mathbf{u}^*)
51: **end procedure**

452 **8.8 Details of Navigation1 Scenario**

453 The robot can move in the four cardinal directions with 0.1 probability of actuator failure. If the
454 realised movement leads to a collision with an obstacle or the edge of the map, no movement occurs
455 and the robot remains in its current position. If the robot's true state is a landmark, the robot receives
456 a position reading uniformly in the 9×9 grid around the robot's true state. Outside the landmarks,
457 the robot receives no observation. The robot receives a penalty of -100 for entering a danger zone

458 and a reward of +300 for entering a goal state. In both cases, the problem terminates. Every other
459 state incurs a reward of -1. The discount parameter was 0.99. The robot's initial belief was equally
460 distributed between two initial positions that were uniformly sampled from the southern-most row of
461 the map.

462 **8.9 Details of Navigation2 Scenario**

463 Similar to Navigation1, the robot's action space consists of moves anywhere in the four cardinal
464 directions NORTH, SOUTH, EAST, WEST. To simulate noise in the robot's actuator's, actions fail
465 with 0.1 probability, and if this occurs, the robot moves randomly in a direction orthogonal to the
466 one specified. If the realised movement leads to a collision with an obstacle or the edge of the
467 map, no movement occurs and the robot remains in its current position. If the robot's true state is a
468 landmark, the robot receives a position reading uniformly in the 9×9 grid around the robot's true
469 state. Otherwise, the robot receives no observation. The robot receives a reward of +600 for being in
470 a goal state, and -3 for being in any other state. The discount parameter was $\gamma = 0.99$. The robot's
471 initial belief was equally distributed between two initial positions that were uniformly sampled from
472 the southern-most row of the map.

473 **8.10 Source code**

474 We also include the source codes for REF Solver, which is developed on top of pomdp_py.