

---

# Supplementary Material to Control What You Can: Intrinsically Motivated Task-Planning Agent

---

Anonymous Author(s)

Affiliation

Address

email

1 The supplementary information is structured as follows. We start with the algorithmic details in the  
2 next section and elaborate on the goal proposal networks. We provide the pseudocode in Sec. B and  
3 further information on the environment in Sec. C. In Sec. D we discuss the oracle baselines followed  
4 by Sec. E providing further analysis of the ablation studies. Finally we report the parameters for the  
5 architectures and the training in Sec. F.

## 6 A Details of the method

### 7 A.1 Final task selector

8 The task selector  $\mathcal{T}$  [3] models the learning progress when attempting to solve a task and is im-  
9 plemented as a multi-armed bandit. The reward is given in Eq. 1. We use the absolute value of  
10 the learning progress  $|\rho|$  because the system should both learn when it can improve, but also if  
11 performance degrades [2]. Initially, the surprise term dominates the quantity. As soon as actual  
12 progress can be made  $\rho$  takes the leading role. The reward is non-stationary and the action-value is  
13 updated according to

$$Q^{\mathcal{T}}(i) = Q^{\mathcal{T}}(i) + \alpha^{\mathcal{T}}(r^{\mathcal{T}}(i) - Q^{\mathcal{T}}(i)) \quad (\text{S1})$$

14 with learning rate  $\alpha^{\mathcal{T}}$ . The task selector is to choose the (final) task for each rollout relative to their  
15 value accordingly. We want to maintain exploration, such that we opt for a stochastic policy with  
16  $p^{\mathcal{T}}(\tau = i) = Q^{\mathcal{T}}(i) / \sum_j Q^{\mathcal{T}}(j)$ .

### 17 A.2 Low-level control

18 Each task  $i$  has its own policy  $\pi_i$  which is trained separately using an off-policy deep RL algorithm.  
19 We use soft actor critic (SAC) [4] in the synthetic environment and DDPG+Her [1] in the robotics  
20 environment. Policies and the critic networks are parametrized by the goal (UVFA [6]).

### 21 A.3 Subgoal sampling

22 For each subtask the goal is selected with the maximal value in the attention map. However,  
23 coordinates of tasks that are still to be solved in the task-chain are fixed, because they can likely not  
24 be controlled by the current policy. Formally:

$$s^* = \arg \max_{s'} G_{i,j}(s') \quad (\text{S2})$$

subject to  $s'_{m_\tau} = s_{m_\tau}, \forall \tau \in \kappa(i+)$

25 where  $\kappa$  is the task-chain and  $\kappa(i+)$  denotes all tasks after  $i$  and including  $i$ .  $s_{m_\tau}$  selects the  
26 coordinates belonging to task  $\tau$ , see Sec. 3.1. The goal for subtask  $j$  is then  $g_j(s) = s_{n_j}^*$ . This is a  
27 convex program and its solution can be computed analytically.

28 **A.4 Intrinsic motivations**

29 For computing the success rate we use a running mean of the last  $Z = 10$  attempts of the particular  
30 task:

$$sr_i = 1/Z \sum_{z=0}^Z succ_i(-z) \tag{S3}$$

31 where  $succ_i(-z) \in [0, 1]$  denotes the success in the  $z$ -th last rollout where task  $i$  was attempted to be  
32 solved.

33 The learning progress  $\rho_i$  is then given as the finite difference of  $sr_i$  between subsequent attempts of  
34 task  $i$ .

35 To compute the surprise signal  $surprise_i$ , we compute the statistics of the prediction error over all the  
36 collected experience, i. e. we assume

$$(e_i(t) - e_i(t - 1)) \sim \mathcal{N}(\mu_i, \sigma_i^2) \tag{S4}$$

37 and compute the empirical  $\mu$  and  $\sigma$ . Denoting the finite difference by  $\dot{e}_i$ , surprise within one rollout  
38 is then defined as

$$surprise_i(t) = \begin{cases} 1 & \text{if } |\dot{e}_i(t)| > \mu_i + \theta \cdot \sigma_i \\ 0 & \text{otherwise.} \end{cases} \tag{S5}$$

39 where  $\theta$  is a hyperparameter that needs to be choose.

40 **A.5 Training details of the goal proposal network**

41 In an ever-changing environment as the ones presented in this paper, the goal proposal networks  
42 are a critical component of our framework that aim to learn relations between entities in the world.  
43 Transitions observed in the environment are labeled by the agent in interesting and undetermined  
44 transitions. Interesting transitions are those, in which a surprising event (high prediction error) occurs  
45 or which lead to an success in task  $i$  given some other task  $j$  was solved before, see Eq. 4. All other  
46 transitions are labeled as undetermined, since they might contain transition which are similar to those  
47 that are labeled interesting but didn't spark high interest. Coming back to our running example:  
48 bumping into, hence suddenly moving, the tool might spark interest in the tool because of a suddenly  
49 jump in prediction error. In general, the behaviour of an object after the surprising event is unknown  
50 and label for these transitions is not clear. Conclusively, we discard all undetermined transition within  
51 a rollout that come after a transition with positive label.

52 After removing all data that might prevent the goal proposal networks from learning the right relations  
53 it remains the problem that positive events are rare compared to the massive body of undetermined  
54 data. Hence, we balance the training data in each batch during training.

55 To make efficient use of the few positive samples we collect in the beginning of the training we impose  
56 a structural prior on the goal proposal network given by Eq. 3. The weight matrices are depicted  
57 in Fig. S1. This particular structure restricts the hypothesis space of the component to positional  
58 relations between components in the observation space that contains entities in the environment. In  
59 the main text, Figure 7 shows a compact representation of the initial and final weight matrices for  
60 different tasks that are computed by taking the minimum over  $|w^1|$  (left column) and  $|w^2|$  (middle  
61 column) in Fig. S1.

62 To understand the parametrization, consider to model that two components  $(k, l)$  of  $s$  should have the  
63 same value for a possitive signal, then  $w_{kl}^1 \approx -w_{kl}^2$  should be nonzero and  $w_k^3 = 0$ . In this case the  
64 corresponding term in the exponent of Eq. 4 is zero if  $s_l = s_k$ . We see that in the case of the learned  
65  $G$  in Fig. S1 this relationship is true for the relevant components (position of agent, tool and object).

66 **A.6 Training / overall procedure**

67 All components of CWYC start in a complete uninformed state. A rollout starts by randomly scramble  
68 the environment. The (final) task is chosen by the task selector. The task planner constructs the  
69 task chain  $\kappa$ . Every 5 steps in the environment, the goal proposal networks computes a goal for the

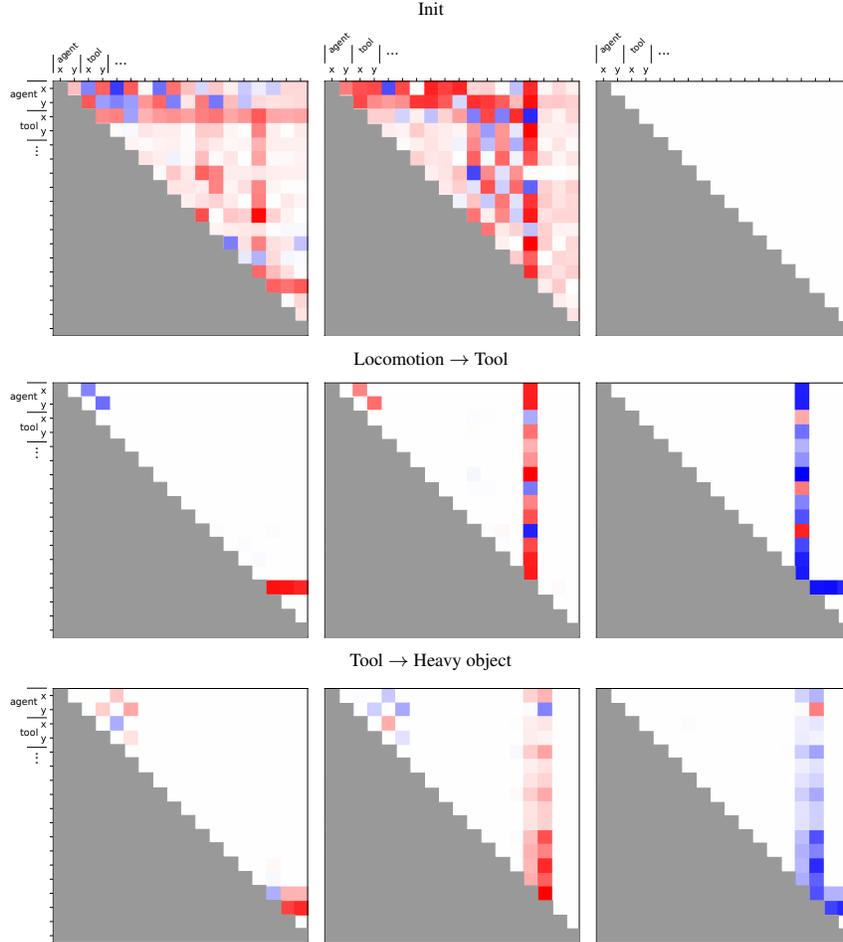


Figure S1: Weights learned by goal proposal networks for different task transitions. The left column shows the weights of  $w^1$ , the middle column of  $w^2$  and the right column of  $w^3$  (see Eq. 3).

70 current task. Given the subgoal the goal-parametric policy of that task is used. Whenever the goal is  
 71 reached (up to a certain precision) a switch to the next task occurs. Again the goal proposal network  
 72 is employed to select a goal in this task, unless it is the final task where the final goal is obviously  
 73 used. If a goal cannot be reached the task ends after  $T_T$  steps. In practice we run 5 rollouts in parallel.  
 74 Then all components are trained using the collected data. For the task selector and task planner we  
 75 use Eq. S1 and Eq. 2, respectively. Forward model and  $G$ s are trained using square-loss and Adam [5].  
 76 The policies are trained according to SAC/DDPG+HER. Pseudo-code and implementation details  
 77 can be found in Sections (B, F).

## 78 B Pseudocode

79 The pseudocode for the method is given in Algorithm 1.

## 80 C Environments

### 81 C.1 Synthetic environment

82 The synthetic environment is depicted in Fig. 2a and is simulated by the physics engine MuJoCo.  
 83 The agent is modeled by a ball that is controlled by applying force in the  $x$  and  $y$  axis, so the agent's  
 84 action corresponds to a 2-dimensional vector:

$$a = (F_x, F_y) \tag{S6}$$

---

**Algorithm 1** CWYC

---

```
1: for episode in episodes do
2:   sample main task  $\tau_{\text{final}} \sim \mathcal{T}$ 
3:   sample main goal  $g_{\tau_{\text{final}}}$  from environment
4:   compute task chain  $\kappa$  using  $B$  starting from  $\tau_{\text{final}}$ 
5:                                     //  $\kappa$  contains list task indices
6:    $i = 1$ 
7:   while  $t < T^{\text{max}}$  and no success in  $\tau_{\text{final}}$  do
8:      $\tau = \tau_{\kappa[i]}$ 
9:     if  $\tau \neq \tau_{\text{final}}$  then
10:      sample goal  $g_{\tau}$  from  $G_{\kappa[i], \kappa[i+1]}$  // Eq. S2
11:      end if
12:      try to reach  $g_{\tau}$  with policy  $\tau$ 
13:      if  $\text{succ}_{\kappa[i]}$  then
14:         $i = i + 1$  // next task in task chain
15:      end if
16:    end while
17:    store episode in history buffer
18:    calculate statistics based on history
19:    train policies for each task
20:    train B // Sec. 3.2
21:    train all G // Sec. 3.2
22:    train  $\mathcal{T}$  // Sec. 3.2
23: end for
```

---

85 The motion of the agent is subject to the laws of motion with the application of friction from the  
86 environment which makes it non-trivial to control. Other than the agent, the environment contains  
87 objects with different dynamics. The positions of the objects are part of the observation space of the  
88 agent along with a flag that specifies if the object has been picked up by the agent. We are dealing  
89 with a fully observable environment.

90 We define the goal spaces of the tasks as corresponding to the position of the individual objects. Some  
91 objects are harder to move than others and have other objects as dependencies. This means that the  
92 agent has to find this relation between them in order to successfully master the environment.

93 The types of objects that are used in the experiments are the following:

- 94 • *Static objects* cannot be moved
- 95 • *Random objects* move randomly in the environment, but cannot be moved by the agent
- 96 • *50% light objects* can be moved in 50% of the rollouts
- 97 • *Tool* can be moved and used to move the heavy object
- 98 • *Heavy objects* can be moved when using the tool

99 The observation vector for  $d$  objects is structured as follows  $(x, y, o_x^1, o_y^1, \dots, o_x^d, o_y^d, \dot{x}, \dot{y}, p^1, \dots, p^d)$ ,  
100 where  $(x, y)$  is the position of the agent,  $(o_x^i, o_y^i)$  is the position of the  $i$ -th object and  $p^i$  indicates  
101 whether the agent is in possession of the  $i$ -th object. The goal spaces are the coordinates of the agent  
102  $(x, y)$  and the coordinates of each object  $(o_x^i, o_y^i)$ .

## 103 C.2 Robotic environment

104 The robotic environment is depicted in Fig. 2c. The state space is 40 dimensional. It consists of the  
105 agent position and velocity, the gripper state, the absolute and relative positions of the box and the  
106 hook, respectively, as well as their velocities and rotations.

107 The environment is based on the OpenAI gym [3] PickAndPlace-v1 environment.

108 Final goals for the reach and tool task are sampled close to the initial gripper and tool location,  
109 respectively. Final goals for the object task are spawned in close proximity to the initial box position

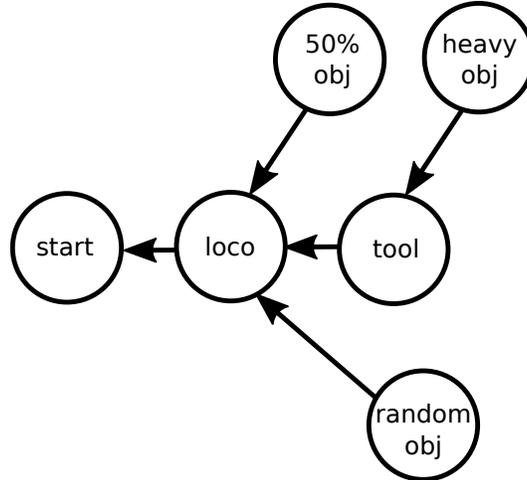


Figure S2: Oracle dependency graph.

110 such that the box needs to be pulled closer to the robot but never pushed away. The box is spawned in  
 111 close proximity to (closer to the robot) the upper end of the hook.

## 112 D Oracle baselines

### 113 D.1 CWYC with oracle goals

114 To assess the maximum performance of CWYC in the described settings, we crafted an upper baseline  
 115 in which all learned high-level components, except for the final task selector  $\mathcal{T}$ , are fixed and set to  
 116 their optimal value.

117 In the distractor setting, every task is solved by first doing the locomotion task. The goal proposal  
 118 network  $G_{i,j}(s)$  returns always the state value  $s_{m_i}$ , reflecting the ground truth relation we try to  
 119 learn.

120 In the synthetic tool-use setting, the task graph depicted in Figure S2 is used. The goal proposal  
 121 network  $G_{i,j}(s)$  returns always the state value  $s_{m_i}$ , reflecting the ground truth relation we try to  
 122 learn.

### 123 D.2 HIRO/SAC with oracle reward

124 To see if HIRO manages to solve the synthetic environment at all, we constructed a oracle version  
 125 of HIRO. The oracle receives as input not only the distance from, e.g., tool to target position but  
 126 additionally the distance from agent to tool. This signal is rich enough to allow HIRO to solve the  
 127 tool manipulation task as shown in Fig. 8(d) in the main text, although it still takes a lot of time  
 128 compared to CWYC. We trained the SAC baseline on the same hybrid reward as well.

## 129 E Additional analysis of the ablation studies

130 Without the surprise signal CWYC\* neither learns a meaningful resource allocation schedule, see  
 131 Fig. S3(a), nor a task dependency graph, see Fig. S3(b). This highlights again the critical role of the  
 132 surprise signal.

## 133 F Training Details and Parameters

### 134 F.1 Synthetic environment

- 135 • Training:
- 136 # parallel rollout workers: 5

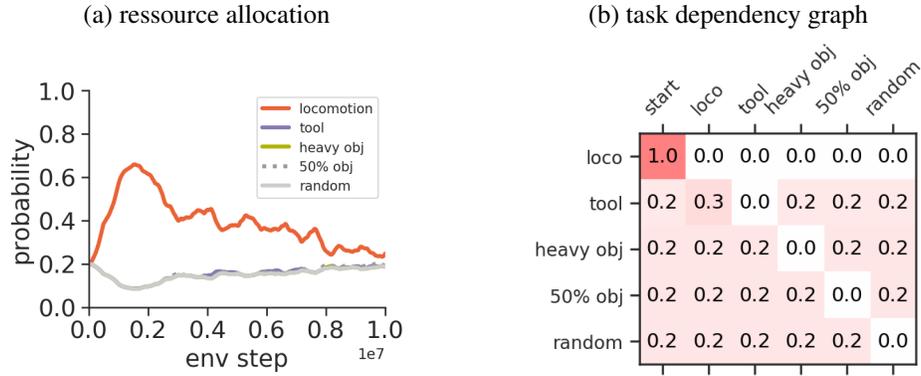


Figure S3: (a) Ressource allocation and (b) task dependency graph for the ablated version CWYC\*. In (a) all tasks except locomotion behave identically because no progress is made.

- 137     • **Environment:**  
arena size:     $20 \times 20$   
138      $T^{\max}$ :        1600  
 $\delta$ :            1.0
- 139     • **SAC:**  
lr:                 $3 \times 10^{-4}$   
batch size:        64  
policy type:       gaussian  
discount:          0.99  
reward scale:      5  
target update interval: 1  
140     tau (soft update)    $5 \times 10^{-3}$   
action prior:      uniform  
reg:                $1 \times 10^{-3}$   
layer size ( $\pi, q, v$ ): 256  
# layers ( $\pi, q, v$ ): 2  
# train iterations: 200  
buffer size:        $1 \times 10^6$
- 141     • **Forward model:**  
lr:                 $10^{-4}$   
batch size:        64  
input:              $(o_{t-1}, u_{t-1})$   
confidence interval: 5  
142     network type:   MLP  
layer size:        100  
 $\theta$ :             5  
# layers:          9  
# train iterations: 100
- 143     • **Final task selector:**  
 $\beta^T$ :             $10^{-1}$   
lr:                 $10^{-1}$   
144     random\_eps:     0.05  
surprise history weighting: 0.99
- 145     • **Task planner:**  
 $\beta^B$ :             $10^{-3}$   
avg. window size: 100  
146     surprise history weighting: 0.99  
sampling\_eps:     0.05
- 147     • **Goal proposal network:**

lr:  $10^{-4}$   
 batch size: 64  
 L1 reg.: 0.0  
 L2 reg.: 0.0  
 $\gamma$  init: 1.0  
 $\gamma$  trainable: True  
 # train iterations: 100

## 149 F.2 Robotic environment

### 150 • Training:

151 # parallel rollout workers: 5

### 152 • Environment:

153  $T^{\max}$ : 150  
 $\delta$ : 0.05

### 154 • DDPG+HER:

Q\_lr:  $10^{-3}$   
 pi\_lr:  $10^{-3}$   
 batch size: 256  
 polyak: 0.95  
 layer size ( $\pi, q$ ): 256  
 # layers ( $\pi, q$ ): 3  
 # train iterations: 80  
 155 buffer size:  $1 \times 10^6$   
 action\_l2: 1.0  
 relative goals: false  
 replay strategy: future  
 replay\_k: 4  
 random\_eps: 0.3  
 noise\_eps: 0.2

### 156 • Forward model:

lr:  $10^{-4}$   
 batch size: 64  
 input:  $(o_{t-1}, u_{t-1})$   
 confidence interval: 3  
 157 network type: MLP  
 layer size: 100  
 $\theta$ : 3  
 # layers: 9  
 # train iterations: 100

### 158 • Final task selector:

$\beta^{\mathcal{T}}$ :  $10^{-1}$   
 lr:  $10^{-1}$   
 159 random\_eps: 0.05  
 surprise history weighting: 0.99

### 160 • Task planner:

$\beta^B$ :  $10^{-3}$   
 avg. window size: 100  
 161 surprise history weighting: 0.99  
 sampling\_eps: 0.05

### 162 • Goal proposal network:

lr:  $10^{-4}$   
batch size: 64  
L1 reg.: 0.0  
163 L2 reg.: 0.0  
 $\gamma$  init: 1.0  
 $\gamma$  trainable: True  
# train iterations: 30

## 164 **References**

- 165 [1] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder,  
166 Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience  
167 replay. In *Advances in Neural Information Processing Systems*, pages 5048–5058, 2017.
- 168 [2] Adrien Baranes and Pierre-Yves Oudeyer. Active Learning of Inverse Models with Intrinsically  
169 Motivated Goal Exploration in Robots. *Robotics and Autonomous Systems*, 61(1):69–73, January  
170 2013.
- 171 [3] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang,  
172 and Wojciech Zaremba. Openai gym, 2016.
- 173 [4] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy  
174 maximum entropy deep reinforcement learning with a stochastic actor. In Jennifer Dy and  
175 Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*,  
176 volume 80 of *PMLR*, pages 1861–1870. PMLR, 10–15 Jul 2018.
- 177 [5] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *in*  
178 *Proceedings of ICLR*, 2015. arXiv preprint <https://arxiv.org/abs/1412.6980>.
- 179 [6] Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approxi-  
180 mators. In *International Conference on Machine Learning*, pages 1312–1320, 2015.