

OMNI: OPEN-ENDEDNESS VIA MODELS OF HUMAN NOTIONS OF INTERESTINGNESS

Jenny Zhang

Department of Computer Science
University of British Columbia
Vector Institute
jennyzzt@cs.ubc.ca

Joel Lehman

Stochastic Labs
lehman.154@gmail.com

Kenneth Stanley

Maven
kennethostanley@gmail.com

Jeff Clune

Department of Computer Science
University of British Columbia
Vector Institute
Canada CIFAR AI Chair
jeffclune@ubc.ca

ABSTRACT

Open-ended algorithms aim to learn new, interesting behaviors forever. That requires a vast environment search space, but there are thus infinitely many possible tasks. Even after filtering for tasks the current agent can learn (i.e., learning progress), countless learnable yet uninteresting tasks remain (e.g., minor variations of previously learned tasks). An Achilles Heel of open-endedness research is the inability to quantify (and thus prioritize) tasks that are not just learnable, but also *interesting* (e.g., worthwhile and novel). We propose solving this problem by *Open-endedness via Models of human Notions of Interestingness* (OMNI). The insight is that we can utilize large (language) models (LMs) as a model of interestingness (MoI), because they *already* internalize human concepts of interestingness from training on vast amounts of human-generated data, where humans naturally write about what they find interesting or boring. We show that LM-based MoIs improve open-ended learning by focusing on tasks that are both learnable *and interesting*, outperforming baselines based on uniform task sampling or learning progress alone. This approach has the potential to dramatically advance the ability to intelligently select which tasks to focus on next (i.e., auto-curricula), and could be seen as AI selecting its own next task to learn, facilitating self-improving AI and AI-Generating Algorithms.¹

1 INTRODUCTION

Provided that the real, significant challenges of AI safety and existential risk can be solved (Critch & Krueger, 2020; Bostrom, 2002; Turchin & Denkenberger, 2020; Ecoffet et al., 2020), there are tremendous gains to be had by creating more powerful AI or even AGI. A great hope for AI is that one day it can produce breakthroughs that fundamentally improve the human condition. These so-far uniquely human advancements and discoveries are the hallmark of civilization, from the invention of the wheel, to farming, vaccines, computers, and even rock and roll. Perhaps someday, unlimited energy and a cure for cancer can be added to the list. What does AI need to possess to discover such new paradigms, as only humans have until now?

Much discussed in open-endedness research (Stanley et al., 2023), the ephemeral fuel behind civilization’s prodigious output is the human intuition for *interestingness*. Drawing upon eons of human experience, we can sense potential even when we don’t precisely know where it leads. Conventional Reinforcement Learning (RL) tools (e.g., intrinsic motivation (Aubret et al., 2019; Pathak et al., 2017;

¹View website at <https://www.jennyzhangzt.com/omni/>.

Osband et al., 2018; Colas et al., 2022b; Oudeyer et al., 2007) and learning progress (Kanitscheider et al., 2021; Matiisen et al., 2019; Portelas et al., 2020; Graves et al., 2017; Kovač et al., 2022; Baranes & Oudeyer, 2013)) are so far only shadows of what such a human sense could do. However, with the rise of large models (LMs) (a.k.a. foundation models (Bommasani et al., 2021)), an intriguing prospect has arisen – trained on vast troves of human experience, perhaps LMs have the potential to grapple for the first time with the critical question of what is actually interesting to explore.

Open-ended learning algorithms, which could leverage such a notion of interestingness, seek to create AI agents that, like humans, continuously learn a variety of different skills within a vast, complex, ever-changing environment. The challenge addressed by interestingness is that, in such environments, there are an infinite number of possible tasks, requiring some method to choose which tasks to try to learn next at every point in training. Handcrafting curricula for training agents in open-ended environments can be extremely challenging due to the sheer number of tasks and the need to adapt to the agent’s skill level and learning progress. In pursuit of an algorithm that is applicable in any domain and enables perpetual learning, handcrafting curricula proves to be an impractical solution. Learning progress methods are a type of auto-curriculum approach that estimates which tasks are at appropriate difficulty levels for the agent to learn from (Kanitscheider et al., 2021; Matiisen et al., 2019; Portelas et al., 2020; Graves et al., 2017; Kovač et al., 2022; Baranes & Oudeyer, 2013). However, such methods can be distracted by learnable yet uninteresting tasks. For example, an agent could be bogged down indefinitely with rearranging silverware in slightly new configurations, hindering it from trying other interesting tasks. Even after filtering for tasks that the current agent can learn, countless learnable yet uninteresting tasks may persist (e.g., slight variations of previously learned tasks). A key challenge in open-endedness research is the inability to quantify and thus focus on tasks that are not only learnable but also interesting. There have been many attempts to quantify interestingness, but, as we detail in Section 2, such simple, hand-crafted formulas consistently fall short of truly capturing the essence of interestingness, creating crippling pathologies. This paper proposes a different path forward.

To borrow from Newton, modern AI sees further by standing on the shoulders of giant human datasets. Training on vast amounts of human-generated data has proven powerful in many cases, such as text generation (e.g., GPT-3 (Brown et al., 2020)), image generation (e.g., DALL-E (Ramesh et al., 2021)), and representation learning (e.g., CLIP (Radford et al., 2021)). We propose *Open-endedness via Models of human Notions of Interestingness* (OMNI). OMNI leverages the power of LMs that have already been trained on extensive human-generated data and have an inherent understanding of human notions of interestingness (Brown et al., 2020; OpenAI, 2023). OMNI utilizes LMs as a model of interestingness (MoI) to focus on tasks that are: (1) learnable, at appropriate difficulty levels for the agents to learn from, and (2) interesting, roughly meaning worthwhile to learn and sufficiently novel. The concepts of “interestingness”, “worthwhile”, and “novelty” are challenging to explicitly define, let alone quantify, which is precisely what OMNI addresses. Humans can intuitively assess these qualities despite their elusive and abstract nature, echoing Justice Potter Stewart’s sentiment of “I know it when I see it” (Stewart, 1964). The goal of OMNI is to emulate this human capacity for nuanced interestingness judgement in open-ended learning. We evaluate OMNI on two challenging domains, *Crafter* (Hafner, 2021) (a 2D version of Minecraft) and *BabyAI* (Chevalier-Boisvert et al., 2018) (a 2D grid world for grounded language learning). OMNI outperforms baselines based on uniform task sampling or learning progress alone. Overall, OMNI has the potential to significantly enhance the ability of AI to intelligently select which tasks to concentrate on next for endless learning and could pave the way for self-improving AI and AI-Generating Algorithms (Clune, 2019).

2 RELATED WORK

2.1 AUTO-CURRICULUM LEARNING

Training neural networks with a curriculum has been extensively studied (Bengio et al., 2009). Auto-curriculum learning has emerged as a promising research area in RL (Kanitscheider et al., 2021; Matiisen et al., 2019; Portelas et al., 2020; Graves et al., 2017; Kovač et al., 2022; Baranes & Oudeyer, 2013; Lehman & Stanley, 2011a; Eysenbach et al., 2018; Wang et al., 2019; 2020; Akkaya et al., 2019; Florensa et al., 2018; Zhang et al., 2020; Campero et al., 2020; OpenAI et al., 2021; Dennis et al., 2020; Gur et al., 2021; Jiang et al., 2021; Dharna et al., 2022), with approaches based on success probabilities and reward thresholds (Wang et al., 2019; 2020; Akkaya et al., 2019;




Figure 1: **Overview of OMNI.** OMNI enables open-ended learning in *vast* environment search spaces by ensuring that tasks trained on not only have high learning progress, but are also *interesting* (harnessing large AI models to make such a heretofore impossible judgement).

Campero et al., 2020; Tan et al., 2023), regret (Dennis et al., 2020; Gur et al., 2021; Jiang et al., 2021), or learning progress (Kanitscheider et al., 2021; Matiisen et al., 2019; Portelas et al., 2020; Graves et al., 2017; Kovač et al., 2022; Baranes & Oudeyer, 2013). Static threshold-based approaches provide a straightforward method for curriculum design. These approaches involve setting fixed criteria for tasks based on their difficulty or complexity. An agent progresses to the subsequent task in a predefined order only after mastering a simpler one. To handcraft an effective curriculum, one would have to understand the relative difficulty of each task and identify tasks of suitable difficulty corresponding to each phase of the agent’s learning trajectory. Doing this in a vast task space is extremely difficult or even impossible. Regret-based methods compute per-task regret by taking the difference between the maximum known return and the average return over multiple rollouts. Regret-based methods typically select tasks with high regret, under the assumption that these tasks still offer substantial learning opportunities (Dennis et al., 2020; Gur et al., 2021; Jiang et al., 2021). However, in stochastic environments, this approach may favor more stochastic and less learnable tasks instead of less stochastic and more learnable ones (Kanitscheider et al., 2021). Learning-progress-based curricula have the potential to mitigate these issues by monitoring the agent’s progress and adapting the task selection accordingly (Kanitscheider et al., 2021; Matiisen et al., 2019; Portelas et al., 2020; Graves et al., 2017; Kovač et al., 2022; Baranes & Oudeyer, 2013). Kanitscheider et al. (2021) demonstrated that learning progress can be measured reliably and that learning-progress-based curricula can be applied to hard RL problems at scale. Our work extends the learning-progress-based curriculum proposed by Kanitscheider et al. (2021). A notable limitation of existing auto-curricula approaches is their inability to distinguish between interesting and uninteresting tasks. Despite filtering for learnable tasks, open-ended environments may still contain infinite learnable but uninteresting tasks. This paper proposes a novel method for identifying and filtering interesting tasks and integrates it with a learning-progress-based auto-curriculum.

2.2 ATTEMPTS TO QUANTIFY INTERESTINGNESS

Many prior research papers have tried to encourage a predefined metric of novelty, diversity, exploration, or open-endedness, but doing so requires quantifying these ineffable qualities. The problem is that optimizing these quantitative measures often leads to undesirable or pathological outcomes, resulting in an output that conforms to the defined metrics, rather than achieving the intended goal (Aubret et al., 2019; Pathak et al., 2017; Osband et al., 2018; Colas et al., 2022b; Oudeyer et al., 2007; Etcheverry et al., 2020; Lehman & Stanley, 2011a; Mouret, 2011; Mouret & Clune, 2015; Lehman & Stanley, 2011b; Eysenbach et al., 2018; Bellemare et al., 2016; Ecoffet et al., 2019; 2021; Mendonca et al., 2023; Lehman & Stanley, 2012; Lehman et al., 2020; Nguyen et al., 2015; Auerbach & Bongard, 2010). As Goodhart’s law posits, “when a measure becomes a target, it ceases to be a good measure” (Strathern, 1997). For example, an agent might exploit a novelty measure by generating many superficially different but ultimately trivial solutions, thus undermining the goal of discovering genuinely interesting outcomes (Lehman & Stanley, 2011a). Similarly, based on how intrinsic motivation is measured, an agent could be biased towards certain types of solutions, leading to a narrow exploration of the problem space rather than developing diverse and valuable insights and innovations (Aubret et al., 2019). Attempting to manually specify a criteria for what constitutes an interesting learning challenge is unlikely to yield satisfactory results. Instead, this paper proposes harnessing LMs to model ineffable human notions of interestingness, gleaned from large text corpora of existing human-generated data (e.g. training on the Internet).

2.3 PRE-TRAINED LARGE MODELS IN OPEN-ENDEDNESS

Large (language) models have recently shown a remarkable ability to capture rich knowledge on an extensive array of subjects from large-scale text corpora. They achieve impressive performance across a wide range of natural language processing tasks (Brown et al., 2020; OpenAI, 2023; Kenton & Toutanova, 2019; Liu et al., 2019; Min et al., 2021; Li et al., 2022; Colas et al., 2023) and display profound understanding of complex concepts such as physics. Consequently, they are utilized in many robotics domains (Huang et al., 2022b;a; Ahn et al., 2022; Yang et al., 2023; Lynch & Sermanet, 2020; Sharma et al., 2021; Kant et al., 2022; Kwon et al., 2023; Du et al., 2023; Driess et al., 2023). There has been growing interest in using them for task selection or generation. Some studies have investigated the application of LMs in breaking down high-level instructions into a sequence of sub-goals, which can be executed by an agent in a zero-shot manner (Huang et al., 2022b;a; Ahn et al., 2022; Yang et al., 2023; Colas et al., 2023) or used to train modular sub-policies (Lynch & Sermanet, 2020; Sharma et al., 2021). Kant et al. (2022) queries LMs for zero-shot commonsense priors and apply them to a planning task. Other studies have utilized LMs to estimate success rates for a given task or desired behavior (Kwon et al., 2023; Du et al., 2023; Colas et al., 2023; Wang et al., 2023). Moreover, LMs have been employed to generate tasks, enabling structured exploration in various environments (Du et al., 2023; Colas et al., 2023; Wang et al., 2023). OMNI differs from Du et al. (2023) by considering an agent’s past successes and employing LMs’ commonsense knowledge for adaptive task selection. Unlike Wang et al. (2023), which employs a code API generated by LLMs, OMNI promotes direct action learning via environment interaction, demanding potentially higher computational resources but bypassing the need for and, critically, limitations of, domain-specific code APIs. While Colas et al. (2023) use deterministic environments and binary reward signals for trajectory success, OMNI adopts a more nuanced approach in stochastic settings, recognizing that agents often improve over time and may not always achieve consistent success rates.

3 METHODS

3.1 PROBLEM FORMULATION

We train task-conditioned agents, and formulate the RL problem as a partially observed Markov decision process (Kaelbling et al., 1998) defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \mathcal{O}, \Omega, \gamma)$. Observations $o \in \Omega$ depend on the new environment states $s \in \mathcal{S}$ and actions taken $a \in \mathcal{A}$ via $\mathcal{O}(o|s, a)$. The task which the agent is conditioned on is part of the environment state s . $\mathcal{T}(s'|s, a)$ describes the dynamics of the environment. $\mathcal{R}(s, a)$ is the environment’s reward function. γ is a discount factor. OMNI focuses on generating learnable and interesting tasks to condition the RL agent on.

3.2 LEARNING PROGRESS CURRICULUM

The task pool in open-ended environments can be very large and diverse, making it challenging for an agent to learn effectively through uniform sampling. Most randomly sampled tasks are likely to be impossible (or at least currently too hard for the agent to learn). To automatically identify tasks at the frontier of the agent’s capabilities, we extend the learning-progress-based curriculum (without the dynamic exploration bonus) from Kanitscheider et al. (2021).

The curriculum predominantly samples tasks with high learning progress, defined as an agent’s recent change in task success probability. During training, the agent is periodically evaluated, and a recent success probability estimate p_{recent} is calculated by applying an exponential moving average (EMA) to the evaluated task success rates. p_{recent} is smoothed with a second, identical EMA to obtain a slower-to-change reflection $p_{gradual}$ of the success probability. Since tasks with low success probabilities are more likely to be novel and are harder to learn because the agent observes fewer successes, p_{recent} and $p_{gradual}$ are reweighted to magnify the learning progress in tasks with low success probabilities and reduce the learning progress in tasks with high success probabilities. This reweighting also compensates for the temporal delay caused by the EMA (Figure 5). Bidirectional learning progress, the absolute difference between the reweighted p_{recent} and $p_{gradual}$, is used to also focus learning on tasks where performance is degrading due to forgetting. Sampling of training tasks is biased towards those that score the highest on this bidirectional learning progress measure.

The reweighting mechanism magnifies differences in low probabilities, putting additional focus on tasks that have high learning progress and low success rates. However, it causes a bias against tasks with initially high success rates achieved by chance even though the agent has not yet learned them. We do not remove the reweighting mechanism since it remains useful in the later stages of training. Instead, we propose an extension to the approach from Kanitscheider et al. (2021), normalizing the task success rates with the success rates achieved by a random action policy. Normalizing the success rates reduces the bias against tasks with initially high success rates, increasing their sample frequency during the early stages of the learning progress curriculum. Appendix A has more details on the learning progress (LP) curriculum and the performance difference without our extension.

3.3 MODELING WHAT HUMANS FIND INTERESTING

An LP curriculum can be distracted by endless variations of uninteresting tasks. To address this challenge, a Model of Interestingness (MoI) is used to focus on the selection of interesting tasks that offer substantial learning value. Humans often intuitively know what might be useful for learning new skills or achieving goals much later (Stanley & Lehman, 2015). This is evident in children playing to unknowingly acquire skills, or scientists exploring new areas to uncover unexpected and beneficial knowledge for future endeavors. This paper presents one specific instance of the OMNI principle, but there are many other possible instantiations to explore in future work (Section 5). First, this section outlines the process of using an LM to determine which tasks are interesting. Then, the section describes how the interestingness predictions are utilized to obtain task sampling weights.

Determining Interesting Tasks. This paper capitalizes on the capabilities of autoregressive LMs to emulate human notions of interestingness. LMs are pretrained on vast and diverse text corpora, enabling them to amass a significant amount of world knowledge. We prompt the LM in a few-shot manner by providing it with examples of choosing which tasks are interesting. It takes into account the agent’s existing proficiency on a given set of tasks and suggests what humans would typically find interesting to learn next. Davinci GPT-3 (Brown et al., 2020) was utilized for the Crafter experiments because it was the state-of-the-art language model available when the experiments were run. GPT-4 (OpenAI, 2023) was used for the BabyAI experiments, which were conducted later. Appendices B and C show the full prompts. The input prompt consists of several components:

1. **Directives encouraging interestingly different behaviors**, such as “The ultimate goal that [the agent] would like your help with is to learn as many interestingly different skills as possible ...”
2. **Environment description**, including the possible objects in the environment, and how a task in the environment is specified
3. **Tasks that the agent has done well and tasks to predict the interestingness of**

Sampling Weights. OMNI aims to improve open-ended learning by focusing on tasks that are both learnable and interesting (Figure 1). Task sampling rates are first assigned based on the LP curriculum, with higher rates for tasks with higher learning progress (Section 3.2). Then, an LM-based MoI predicts which tasks are interesting (Appendix D). Boring tasks have their sampling weights reduced by multiplying by 0.001. Finally, task sampling rates are normalized to probabilities that sum to 1.

4 EXPERIMENTS

4.1 CRAFTER ENVIRONMENT

We evaluate OMNI on Crafter (Hafner, 2021), a 2D version of Minecraft that enables collecting and creating a set of artifacts organized along a technology tree. This means that certain tasks need to be completed, often multiple times, as prerequisites for other more challenging tasks (Figure 2). Agents receive RGB pixel observations (64 x 64 resolution) of a 9 x 9 grid area surrounding their position within a 64 x 64 grid landscape that varies with each episode, offering a complex and engaging testing ground. The agent is provided a target task (represented with a bag-of-words encoding) as part of its observation and rewarded +1 upon successful completion of the conditioned task. We modify the game to focus on gathering and crafting skills by eliminating the survival component. This removes the need for the agent to learn and continually apply survival tactics against enemies or for food

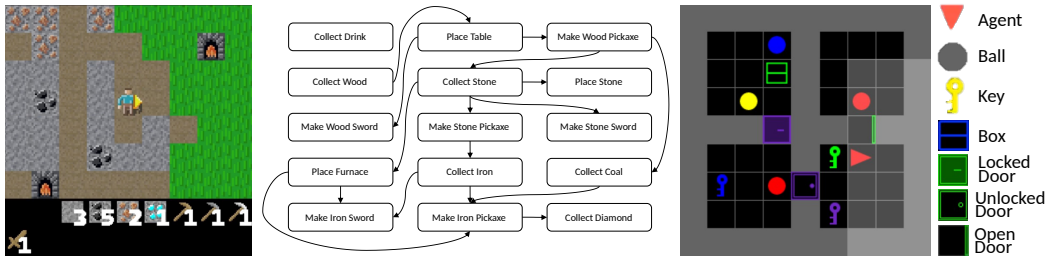


Figure 2: **Crafter and BabyAI environments.** (Left) Agent view in a procedurally generated Crafter world, showing terrain types, resources, and the agent’s inventory. (Middle) The 15 tasks considered interesting for Crafter analyses. Arrows indicate which tasks in the technology tree must be completed, often multiple times, along the way to perform more challenging tasks. (Right) Bird’s-eye view of a randomly generated BabyAI environment, showing different object types, colors, locations, and states. The agent is the red triangle and its view (sometimes occluded) is highlighted in light grey. In this example, the agent starts from the bottom right room, and is tasked to “go to a red ball”. To succeed, the agent must open the green door (sometimes locked) to reach the red ball.

gathering. The “sleep” and “place plant” actions are important for survival in the original game and have been omitted due to their reduced relevance in our modified context, which excludes the survival aspect. The original game consists of 22 tasks, of which, the 15 tasks unrelated to survival are selected and considered interesting.

To investigate our hypothesis that focusing on interesting tasks with high learning progress will improve performance, we dilute the 15 interesting tasks with 90 “boring” tasks and 1023 “extremely challenging” tasks that serve as potential distractors for learning-progress-based approaches. Boring tasks are generated as numerical repeats of interesting tasks, e.g., “collect N wood” where $N \in [2, 10]$, analogous to how minor numerical variations of real-world tasks are less interesting than tasks that differ qualitatively. See Appendix E for the full list of boring tasks. Extremely challenging tasks represent tasks that are too difficult for the agent to complete at its current state of learning, serving as tasks that uniform sampling will waste time on, but that learning-progress-based methods should successfully ignore. The agent is assumed to always fail at these extremely challenging tasks and hence is always assigned a success rate of 0 for them. By analogy, consider the futility of attempting to cook a 5-course meal before learning the basic skill of cutting a vegetable.

4.2 BABYAI ENVIRONMENT

We also evaluate OMNI on BabyAI (Chevalier-Boisvert et al., 2018), a readily available benchmark domain characterized by its partially observable 2D grid world environment (Figure 2). We test on the *MiniBossLevel*. While *BossLevel* is the most challenging level in BabyAI, we choose *MiniBossLevel* as it has the same features as *BossLevel* but with a smaller room and lower probability of locked rooms, speeding up training. For each episode, the room layout and item configuration are randomly generated (using off-the-shelf configurations from Chevalier-Boisvert et al. (2018)). The grid world can have objects in six colors (*red, green, blue, purple, yellow, grey*), and of four types (*key, ball, box, door*). The agent is randomly spawned at a location in the 9 x 9 grid world, containing four 3 x 3 rooms. The agent’s observation includes one-hot encodings of each of the 7 x 7 grid cells in front of the agent (observations are set to a special symbol if occluded), and a description of the task in natural language (embedded with a look-up table and GRU, see Appendix G for more details). The agent receives a reward, proportional to the number of steps it took to finish, only when it has successfully completed the given task. While the Baby Language grammar (Chevalier-Boisvert et al., 2018) is limited to sequential tasks with a maximum of 2 instructions, we expanded this by permitting tasks with up to 5 instructions, resulting in 1364 unique tasks. Each task is a sequence of instructions (*GoTo, PickUp, OpenDoor, PutNextTo*), linked by the ordering constraint *then*. Object placements are randomized each episode. Tasks with the same sequence of instructions but different object instances are considered the same when sampling (e.g., “go to a blue ball” and “go to a red key” are considered the same task “go to <object>”).




Figure 3: **Results in Crafter.** (Left) Conditional success probabilities of all tasks in Crafter. Tasks are organized from simple to complex based on the prerequisite tasks that must be accomplished before completing the target task. Task names (left of each row) are readable in a digital format with zoom. (Right) Performance in Crafter on all tasks. While OMNI biases training towards *interesting tasks*, it achieves higher average task success rates and learns more tasks than uniform sampling or choosing tasks based on learning progress alone, even across *all tasks*.




Figure 4: **Results in BabyAI.** (Left) Conditional success probabilities of a subset of tasks in BabyAI. These plots only show tasks with a success rate of at least 0.05 by any method at any timestep. Tasks are organized from simple to complex based on the instruction length. (Right) Performance in BabyAI on all tasks. The average task success rate scale for BabyAI is low because it is averaged over the entire task set, which includes many tasks that are difficult to learn. This captures in microcosm the real world, where there can be infinitely many difficult or even impossible tasks. OMNI achieves much higher average task success rates and learns more tasks than uniform sampling or choosing tasks based on learning progress alone.

4.3 RESULTS

Both Crafter and BabyAI RL agents are trained with PPO (Schulman et al., 2017), a standard RL algorithm. Policy details and hyperparameters for the Crafter and BabyAI settings are in Appendices F and G. We compare the performance of agents trained with: (1) **Uniform** sampling, (2) Learning Progress (**LP**) only, and (3) OMNI: Learning Progress with additional filtering by a Model of Interestingness (**OMNI: LP + MoI**). Uniform sampling, the control, samples all tasks with equal probabilities. Uniform sampling is the most naive and samples tasks that are too easy or too difficult for the agent most of the time. LP samples tasks based on the calculated learning progress weights (Section 3.2), but is distracted by the many boring tasks. OMNI: LP + MoI focuses on the subset of tasks with high learning progress that are also interesting (Section 3.3). All experiments are run for

100 million time steps and are repeated 10 times with different random seeds. Each experiment takes about 33 hrs for Crafter and 60 hrs for BabyAI on a 24GB NVIDIA A10 GPU with 30 virtual CPUs.

We evaluate our methods with two metrics: (1) the average task success rate, and (2) the number of tasks with success rates exceeding a predetermined threshold α . This study sets $\alpha = 0.2$, consistent with the selections made in related literature (Kanitscheider et al., 2021; Team et al., 2023). The first metric reflects the agent’s average performance across all tasks, while the second metric captures the extent to which the agent is a generalist that has decent competency on many different tasks. These metrics are calculated on the full task set (Figures 3 and 4). Metrics calculated on interesting tasks only are shown in Appendix H. All confidence intervals given are 95% median bootstrap confidence intervals obtained by resampling 1000 times. Confidence intervals are reported with the following notation: *stat* (CI: *lower* – *upper*) where *stat* is the median across runs. Shaded areas in graphs also indicate the 95% median bootstrap confidence interval obtained by resampling 1000 times.

Uniform sampling. As expected, the results with uniform sampling are poor. Worse, the agents did not improve over time as most tasks sampled are too difficult or too easy for the agent and successes are extremely sparse (Figures 3 and 4). The agent is considered to have learned a task if its conditional success probability on that task is at least 0.2. In Crafter, the agent learns 4 (CI: 4 – 6) tasks (interesting or boring) and only 3 (CI: 2 – 3) interesting tasks. The agent achieves an average task success rate of 0.030 (CI: 0.026 – 0.033) on interesting and boring tasks, and 0.103 (CI: 0.087 – 0.120) on interesting tasks only. In BabyAI, the agent learns only 1 (CI: 0 – 1) task and achieves an average task success rate of $4.7e-3$ (CI: $4.6e-3$ – $5.0e-3$) on all tasks.

Learning Progress Curriculum. By focusing on tasks with suitable difficulty, the agent learns to do a lot more tasks with higher success rates than uniform sampling. In Crafter, the agent learns 55 (CI: 54 – 56) tasks (interesting or boring) and 9 (CI: 9 – 11) interesting tasks. The agent achieves an average task success rate of 0.42 (CI: 0.41 – 0.43) on interesting and boring tasks, and 0.52 (CI: 0.50 – 0.56) on interesting tasks only. In BabyAI, the agent learns 4 (CI: 4 – 6) tasks and achieves an average task success rate of $5.9e-3$ (CI: $5.5e-3$ – $6.2e-3$) on all tasks. Across all metrics and in both domains, the differences in performance between LP and Uniform at 25%, 50%, 75%, and 100% of the way through training are statistically significant (all $p < 1e-3$, Mann Whitney U test), showing that LP significantly outperforms uniform sampling (Figures 3 and 4). LP samples tasks that are at the frontier of the agent’s capabilities (Figures 3, 4, 10, 12). When a task’s conditional success probability changes, LP focuses more on it. Hence, there will be more rollouts where the task is the given goal and thus more positive examples from which the agent can learn to solve the conditioned task. However, LP is distracted by boring tasks (Figures 10 and 12). When the conditional success probabilities of boring tasks change, LP allocates higher sampling weights to them even though they are similar to other sampled tasks and might not expand the agent’s range of skills.

OMNI: Learning Progress + a Model of Interestingness. To automatically select and focus on interesting tasks, an LM is prompted in a few-shot manner to predict which tasks are interesting. By combining LP with an MoI, OMNI focuses on the subset of high learning progress tasks that are interesting. In Crafter, the agent learns 82 (CI: 80 – 87) tasks (interesting or boring) and 14 (CI: 14 – 14) interesting tasks. The agent achieves an average task success rate of 0.56 (CI: 0.54 – 0.58) on interesting and boring tasks, and 0.78 (CI: 0.76 – 0.80) on interesting tasks only. In BabyAI, the agent learns 8 (CI: 7 – 10) tasks and achieves an average task success rate of $7.5e-3$ (CI: $7.3e-3$ – $7.7e-3$) on all tasks. Across all metrics and in both domains, the differences in performance between OMNI and LP at 25%, 50%, 75%, and 100% of the way through training are statistically significant (all $p < 1e-3$, Mann Whitney U test), showing that OMNI significantly outperforms an LP-only curriculum (Figures 3 and 4). OMNI is not distracted by uninteresting yet learnable tasks, and focuses on the interesting tasks only (Figures 10 and 12). The trained agent not only achieves higher average task success rates, but also learns more challenging tasks faster (Figures 3 and 4).

We thus know OMNI performs better than LP alone, but how good is it at predicting interesting tasks? To address this, we created an oracle for the MoI, termed the Oracle Model of Interestingness (OMOI). Impressively, the performance of the LM-based MoI is nearly on par with the oracle, suggesting that OMNI is highly effective in identifying interesting tasks for the agent to learn on. Across all metrics and in both domains, the performance of OMNI is almost indistinguishable from that of LP + OMOI (Figures 13 and 16), and most of the differences in performance between them at 25%, 50%, 75%, and 100% of the way through training are not statistically significant (all but one $p > 0.05$, Mann Whitney U test). Appendix I presents more details on the OMOI’s design, and provides plots of task

sample rates and task success rates for LP + OMoI. These findings underscore the effectiveness of OMNI in harnessing the power of LMs to guide AI agents towards interesting and learnable tasks.

In the main Crafter setting, we consider repetitive tasks as boring tasks. One may wonder how OMNI performs on different types of boring tasks. To address this, we explored two additional settings in Crafter, detailed in Appendices J and K, where boring tasks are generated as compounds and synonyms respectively. Compound tasks are combinations of interesting tasks (e.g., “collect wood and make stone pickaxe”). Synonymous tasks are tasks with different names but the same success condition (e.g., “collect wood” vs. “gather wood”). The MoI also captures what humans find interesting in these additional settings, and as a result, OMNI significantly outperforms the LP-only and Uniform curricula. These findings suggest that OMNI is robust and adaptable to various tasks, including those that deviate from the repetitive tasks initially considered. By demonstrating its effectiveness across diverse task settings, OMNI further strengthens its position as a valuable approach for guiding AI agents towards interesting and learnable tasks.

5 DISCUSSION, FUTURE WORK, AND CONCLUSION

This study explores the vision of using human notions of interestingness to accelerate open-ended learning, an approach we term *Open-endedness via Modeling human Notions of Interestingness* (OMNI). OMNI has several advantages over other methods by leveraging human concepts of interestingness to guide task selection in open-ended learning (Section 4.3). In this first version, we estimate learning progress through statistical methods and utilize a Model of Interestingness (MoI) based on human data distilled into LMs. A different version of OMNI could have the LM judge both learning progress and interestingness by giving it a history of task selection and task performance. That could allow for a more flexible notion of learning progress, as it can potentially recognize not just success or failure, but also important stepping stones and patterns leading to a solution, echoing the complexity of human learning progression. Consider an agent learning object manipulation, eventually progressing to more complex tasks like peeling an egg. Despite prolonged lack of reward, the task may not be “too difficult” as statistical learning progress might suggest. Given the agent’s proven ability to handle intricate items, it may be primed for egg peeling, simply requiring more attempts. Preliminary results suggest that LMs can successfully integrate these aspects (Appendix L). Notably, as LMs continue to advance (Kaplan et al., 2020), all versions of OMNI are expected to improve correspondingly. Ultimately, OMNI offers a general recipe for accelerated learning in open-ended environments with potentially infinite tasks, and steers open-ended learning towards meaningful and interesting progress, instead of meandering aimlessly amidst endless possibilities.

Looking ahead, there are several promising avenues for future work. First, pushing beyond the reliance on a predetermined set of tasks, systems that automatically generate tasks could enable learning in a broader range of tasks (Wang et al., 2019; 2020). One way of representing a diverse and vast task space is through language. Its advantages include ease of specification by humans, the ability to encapsulate more abstract concepts than standard state-based goals (Colas et al., 2022a), and the potential to enhance the trained agent’s generalizability due to its partial compositionality (Colas et al., 2022a; 2020; Hermann et al., 2017). Second, instead of using handcrafted reward signals, we propose training foundation models to automatically provide reward signals, exemplified by previous approaches (Huang et al., 2022b; Fan et al., 2022). One version could be a video-language foundation model that evaluates the degree to which an agent completes natural language tasks like “build a house” or “herd sheep to a hilltop”. It is worth noting that these two points are closely related, as the automatic estimation of reward signals for arbitrary tasks specified in natural language is essential in order to use RL with automatically generated natural language tasks.

Another possibility is to incorporate multi-modal models, such as vision-language models and other modalities into the *MoI*. This could give the MoI richer representations and a better comprehension of the agent’s capabilities, facilitating a more accurate assessment of the agent’s learning progress and task diversity. For example, a vision-language MoI might see that the agent is making progress on or very close to solving a task for which it is getting no reward, such as peeling an egg. Another idea is to allow the MoI to autonomously analyze quantitative performance measures, make its own assessment of learning progress, and incorporate that into its notion of interestingness. Our preliminary investigations indicate that LMs, such as GPT-3 and GPT-4, can automatically analyze numerical results and adjust their understanding of interestingness (Figures 24 and 27).

Critically, by Goodhart’s law, we expect that any model will have pathologies uncovered once it is a target metric being optimized against. For instance, fine-tuning a task generator to produce tasks that the MoI finds interesting might eventually result in the MoI not being a good indicator of what is interesting. Hence, refining and updating the MoI with additional human feedback could lead to more effective learning systems, an algorithm within the OMNI paradigm we call *Open-Endedness with Human Feedback* (OEHF). Similar to Reinforcement Learning with Human Feedback (RLHF) (Christiano et al., 2017), the objective is to train a model that can effectively capture an ineffable property that, although challenging to quantitatively measure, is readily identifiable upon observation (e.g., a backflip for RLHF, or whether a task is interesting for OEHF). One version would be to build upon the insights of this paper and start with a model *already* well-versed with human concepts of interestingness (via unsupervised pre-training on internet-scale human data), and further fine-tune that MoI with additional human evaluation of its output (as often as necessary). Such fine-tuning can help minimize suboptimal interestingness judgements, enhance the MoI’s understanding of skills not initially present in its zero-shot repertoire, or tailor the MoI to a specific domain.

In conclusion, our work demonstrates the potential of using an MoI to significantly enhance auto-curricula and the quest for open-ended learning algorithms by intelligently focusing on learnable and interesting tasks. OMNI addresses the Achilles Heel of open-ended systems, which lies in defining and quantifying interestingness, as previous attempts have resulted in pathologies when optimizing against such definitions and quantifications. OMNI mitigates this problem by leveraging human notions of interestingness to guide AI systems. There are numerous ways to implement the principles of this new paradigm, and exploring different versions presents an exciting avenue for future research. The generality and applicability of OMNI to other open-ended domains with large task spaces further underscores its significance. In the long run, it hints at a synergy between LMs and open-endedness that simultaneously addresses looming challenges for both: how will LMs ultimately rise to the level of creativity seen in the best of human innovation, and how will open-endedness overcome the trap of diverging into a vast space of uninspiring mediocrity? By playing off each other’s strengths, LMs can perhaps someday become essential engines of open-ended discovery and begin to participate in the creative dance that has defined civilization since its inception.

ACKNOWLEDGMENTS

This work was supported by the Vector Institute, a grant from Schmidt Futures, an NSERC Discovery Grant, and a generous donation from Rafael Cosman. We also thank Andrew Dai, Cédric Colas, and members in our lab at the University of British Columbia, namely Aaron Dharna, Ben Norman, and Shengran Hu, for insightful discussions and feedback.

REFERENCES

- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- Arthur Aubret, Laetitia Matignon, and Salima Hassas. A survey on intrinsic motivation in reinforcement learning. *arXiv preprint arXiv:1908.06976*, 2019.
- Joshua E Auerbach and Josh C Bongard. Evolving CPPNs to grow three-dimensional physical structures. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pp. 627–634, 2010.
- Adrien Baranes and Pierre-Yves Oudeyer. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61(1):49–73, 2013.
- Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29, 2016.

- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Nick Bostrom. Existential risks: Analyzing human extinction scenarios and related hazards. *Journal of Evolution and technology*, 9, 2002.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Andres Campero, Roberta Raileanu, Heinrich Küttler, Joshua B Tenenbaum, Tim Rocktäschel, and Edward Grefenstette. Learning with amigo: Adversarially motivated intrinsic goals. *arXiv preprint arXiv:2006.12122*, 2020.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. Babyai: A platform to study the sample efficiency of grounded language learning. *arXiv preprint arXiv:1810.08272*, 2018.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- Jeff Clune. AI-GAs: AI-generating algorithms, an alternate paradigm for producing general artificial intelligence. *arXiv preprint arXiv:1905.10985*, 2019.
- Cédric Colas, Tristan Karch, Nicolas Lair, Jean-Michel Dussoux, Clément Moulin-Frier, Peter Dominey, and Pierre-Yves Oudeyer. Language as a cognitive tool to imagine goals in curiosity driven exploration. *Advances in Neural Information Processing Systems*, 33:3761–3774, 2020.
- Cédric Colas, Tristan Karch, Clément Moulin-Frier, and Pierre-Yves Oudeyer. Language and culture internalization for human-like autotelic AI. *Nature Machine Intelligence*, 4(12):1068–1076, 2022a.
- Cédric Colas, Tristan Karch, Olivier Sigaud, and Pierre-Yves Oudeyer. Autotelic agents with intrinsically motivated goal-conditioned reinforcement learning: a short survey. *Journal of Artificial Intelligence Research*, 74:1159–1199, 2022b.
- Cédric Colas, Laetitia Teodorescu, Pierre-Yves Oudeyer, Xingdi Yuan, and Marc-Alexandre Côté. Augmenting autotelic agents with large language models. *arXiv preprint arXiv:2305.12487*, 2023.
- Andrew Critch and David Krueger. AI research considerations for human existential safety (ARCHES). *arXiv preprint arXiv:2006.04948*, 2020.
- Michael Dennis, Natasha Jaques, Eugene Vinitsky, Alexandre Bayen, Stuart Russell, Andrew Critch, and Sergey Levine. Emergent complexity and zero-shot transfer via unsupervised environment design. *Advances in neural information processing systems*, 33:13049–13061, 2020.
- Rahul Dey and Fathi M Salem. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*, pp. 1597–1600. IEEE, 2017.
- Aaron Dharna, Amy K Hoover, Julian Togelius, and Lisa Soros. Transfer dynamics in emergent evolutionary curricula. *IEEE Transactions on Games*, 2022.
- Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.
- Yuqing Du, Olivia Watkins, Zihan Wang, Cédric Colas, Trevor Darrell, Pieter Abbeel, Abhishek Gupta, and Jacob Andreas. Guiding Pretraining in Reinforcement Learning with Large Language Models. *arXiv preprint arXiv:2302.06692*, 2023.

- Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.
- Adrien Ecoffet, Jeff Clune, and Joel Lehman. Open Questions in Creating Safe Open-ended AI: Tensions Between Control and Creativity. In *ALIFE 2020: The 2020 Conference on Artificial Life*, pp. 27–35. MIT Press, 2020.
- Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. First return, then explore. *Nature*, 590(7847):580–586, 2021.
- Mayalen Etcheverry, Clément Moulin-Frier, and Pierre-Yves Oudeyer. Hierarchically organized latent modules for exploratory search in morphogenetic systems. *Advances in Neural Information Processing Systems*, 33:4846–4859, 2020.
- Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- Linxi Fan, Guanzhi Wang, Yunfan Jiang, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu, and Anima Anandkumar. Minedojo: Building open-ended embodied agents with internet-scale knowledge. *arXiv preprint arXiv:2206.08853*, 2022.
- Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *International conference on machine learning*, pp. 1515–1528. PMLR, 2018.
- Alex Graves, Marc G Bellemare, Jacob Menick, Remi Munos, and Koray Kavukcuoglu. Automated curriculum learning for neural networks. In *international conference on machine learning*, pp. 1311–1320. PMLR, 2017.
- Izzeddin Gur, Natasha Jaques, Kevin Malta, Manoj Tiwari, Honglak Lee, and Aleksandra Faust. Adversarial environment generation for learning to navigate the web. *arXiv preprint arXiv:2103.01991*, 2021.
- Danijar Hafner. Benchmarking the spectrum of agent capabilities. *arXiv preprint arXiv:2109.06780*, 2021.
- Karl Moritz Hermann, Felix Hill, Simon Green, Fumin Wang, Ryan Faulkner, Hubert Soyer, David Szepesvari, Wojciech Marian Czarnecki, Max Jaderberg, Denis Teplyashin, et al. Grounded language learning in a simulated 3d world. *arXiv preprint arXiv:1706.06551*, 2017.
- Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pp. 9118–9147. PMLR, 2022a.
- Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*, 2022b.
- Minqi Jiang, Edward Grefenstette, and Tim Rocktäschel. Prioritized level replay. In *International Conference on Machine Learning*, pp. 4940–4950. PMLR, 2021.
- Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- Ingmar Kanitscheider, Joost Huizinga, David Farhi, William Hebguss, Brandon Houghton, Raul Sampedro, Peter Zhokhov, Bowen Baker, Adrien Ecoffet, Jie Tang, et al. Multi-task curriculum learning in a complex, visual, hard-exploration domain: Minecraft. *arXiv preprint arXiv:2106.14876*, 2021.
- Yash Kant, Arun Ramachandran, Sriram Yenamandra, Igor Gilitschenski, Dhruv Batra, Andrew Szot, and Harsh Agrawal. Housekeep: Tidying virtual households using commonsense reasoning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIX*, pp. 355–373. Springer, 2022.

- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, volume 1, pp. 2, 2019.
- Grgur Kovač, Adrien Laversanne-Finot, and Pierre-Yves Oudeyer. Grimgep: learning progress for robust goal sampling in visual deep reinforcement learning. *IEEE Transactions on Cognitive and Developmental Systems*, 99:1–1, 2022.
- Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. Reward design with language models. *arXiv preprint arXiv:2303.00001*, 2023.
- Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011a.
- Joel Lehman and Kenneth O Stanley. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pp. 211–218, 2011b.
- Joel Lehman and Kenneth O Stanley. Beyond open-endedness: Quantifying impressiveness. In *ALIFE 2012: The Thirteenth International Conference on the Synthesis and Simulation of Living Systems*, pp. 75–82. MIT Press, 2012.
- Joel Lehman, Jeff Clune, Dusan Misevic, Christoph Adami, Lee Altenberg, Julie Beaulieu, Peter J Bentley, Samuel Bernard, Guillaume Beslon, David M Bryson, et al. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial life*, 26(2):274–306, 2020.
- Xiang Lorraine Li, Adhiguna Kuncoro, Jordan Hoffmann, Cyprien de Masson d’Autume, Phil Blunsom, and Aida Nematzadeh. A systematic investigation of commonsense knowledge in large language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 11838–11855, 2022.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Corey Lynch and Pierre Sermanet. Language conditioned imitation learning over unstructured data. *arXiv preprint arXiv:2005.07648*, 2020.
- Tambet Matiisen, Avital Oliver, Taco Cohen, and John Schulman. Teacher–student curriculum learning. *IEEE transactions on neural networks and learning systems*, 31(9):3732–3740, 2019.
- Russell Mendonca, Shikhar Bahl, and Deepak Pathak. Alan: Autonomously exploring robotic agents in the real world. *arXiv preprint arXiv:2302.06604*, 2023.
- Bonan Min, Hayley Ross, Elinor Sulem, Amir Pouran Ben Veyseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heinz, and Dan Roth. Recent advances in natural language processing via large pre-trained language models: A survey. *arXiv preprint arXiv:2111.01243*, 2021.
- Jean-Baptiste Mouret. Novelty-based multiobjectivization. In *New Horizons in Evolutionary Robotics: Extended Contributions from the 2009 EvoDeRob Workshop*, pp. 139–154. Springer, 2011.
- Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.
- Anh Mai Nguyen, Jason Yosinski, and Jeff Clune. Innovation engines: Automated creativity and improved stochastic optimization via deep learning. In *Proceedings of the 2015 annual conference on genetic and evolutionary computation*, pp. 959–966, 2015.
- OpenAI. GPT-4 Technical Report. *ArXiv*, abs/2303.08774, 2023.

- OpenAI OpenAI, Matthias Plappert, Raul Sampedro, Tao Xu, Ilge Akkaya, Vineet Kosaraju, Peter Welinder, Ruben D'Sa, Arthur Petron, Henrique P d O Pinto, et al. Asymmetric self-play for automatic goal discovery in robotic manipulation. *arXiv preprint arXiv:2101.04882*, 2021.
- Ian Osband, John Aslanides, and Albin Cassirer. Randomized prior functions for deep reinforcement learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286, 2007.
- Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pp. 2778–2787. PMLR, 2017.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- Rémy Portelas, Cédric Colas, Katja Hofmann, and Pierre-Yves Oudeyer. Teacher algorithms for curriculum learning of deep rl in continuously parameterized environments. In *Conference on Robot Learning*, pp. 835–853. PMLR, 2020.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831. PMLR, 2021.
- John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Pratyusha Sharma, Antonio Torralba, and Jacob Andreas. Skill induction and planning with latent language. *arXiv preprint arXiv:2110.01517*, 2021.
- Kenneth O. Stanley and Joel Lehman. *Why Greatness Cannot Be Planned: The Myth of the Objective*. Springer Publishing Company, Incorporated, 2015. ISBN 3319155237.
- Kenneth O. Stanley, Joel Lehman, and Lisa Soros. Open-endedness: The Last Grand Challenge You've Never Heard Of. <https://www.oreilly.com/radar/open-endedness-the-last-grand-challenge-youve-never-heard-of/>, May 2023. Accessed: 2023-05-15.
- Potter Stewart. *Jacobellis v. Ohio*, 378 U.S. 184. United States Supreme Court, 1964.
- Marilyn Strathern. 'Improving ratings': audit in the British University system. *European review*, 5(3):305–321, 1997.
- Daniel Chee Hian Tan, Jenny Zhang, Zhibin Li, et al. Perceptive Locomotion with Controllable Pace and Natural Gait Transitions Over Uneven Terrains. *arXiv preprint arXiv:2301.10894*, 2023.
- Adaptive Agent Team, Jakob Bauer, Kate Baumli, Satinder Baveja, Feryal Behbahani, Avishkar Bhoopchand, Nathalie Bradley-Schmieg, Michael Chang, Natalie Clay, Adrian Collister, et al. Human-Timescale Adaptation in an Open-Ended Task Space. *arXiv preprint arXiv:2301.07608*, 2023.

Alexey Turchin and David Denkenberger. Classification of global catastrophic risks connected with artificial intelligence. *AI & Society*, 35(1):147–163, 2020.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.

Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. Paired open-ended trailblazer (poet): Endlessly generating increasingly complex and diverse learning environments and their solutions. *arXiv preprint arXiv:1901.01753*, 2019.

Rui Wang, Joel Lehman, Aditya Rawal, Jiale Zhi, Yulun Li, Jeffrey Clune, and Kenneth Stanley. Enhanced poet: Open-ended reinforcement learning through unbounded invention of learning challenges and their solutions. In *International Conference on Machine Learning*, pp. 9940–9951. PMLR, 2020.

Sherry Yang, Ofir Nachum, Yilun Du, Jason Wei, Pieter Abbeel, and Dale Schuurmans. Foundation models for decision making: Problems, methods, and opportunities. *arXiv preprint arXiv:2303.04129*, 2023.

Yunzhi Zhang, Pieter Abbeel, and Lerrel Pinto. Automatic curriculum learning through value disagreement. *Advances in Neural Information Processing Systems*, 33:7648–7659, 2020.

A LEARNING PROGRESS CURRICULUM DETAILS

For each task, we calculate the task success rate t_{rdn} achieved by a random action policy. At fixed evaluation intervals during training, the evaluated task success rate t_{eval} of the RL policy is normalized as such:

$$t_{norm} = \frac{t_{eval} - t_{rdn}}{1 - t_{rdn}}$$

The above is our contributed extension to the learning progress method in Kanitscheider et al. (2021). t_{norm} is smoothed with an EMA function to obtain p_{recent} (Figure 5, green). p_{recent} is smoothed with a second identical EMA function to obtain $p_{gradual}$ (Figure 5, brown). The exponential smoothing constant applied in all experiments is 0.1. The bidirectional learning progress measure is given by $LP = |f(p_{recent}) - f(p_{gradual})|$ (Figure 5, blue), where f is the reweighting function:

$$f(p) = \frac{(1 - p_{\theta})p}{p + p_{\theta}(1 - 2p)}$$

with parameter $p_{\theta} = 0.1$.

We employ a sampling function to transform the measure of learning progress into task sampling weights, focusing mostly on tasks with the largest learning progress. The steps are as follows:

- Z-score the reweighted learning progress (subtract mean and divide by standard deviation).
- Apply a sigmoid to the result.
- Normalize resulting weights to sampling probabilities.




Figure 5: The process of determining an agent’s learning progress on a task from its measured success probability on that task in an example fictional problem.

A.1 LEARNING PROGRESS CURRICULUM ABLATION

The learning progress curriculum without the proposed extension (Section 3.2) of normalizing the task success rates to a random action baseline is labelled as **LP-no-norm**. In light of limited compute, we limit ourselves to the Crafter domain and experiments are run for 30 million time steps (vs. 100 million in the main experiments) and are repeated 10 times (same as the main experiments) with different random seeds. The LP curriculum achieves higher average task success rates and learns more tasks with the proposed normalization extension (Figure 6). Across all metrics and in both domains, the differences in performance between LP and LP-no-norm at 50%, 75%, and 100% of the way through training are statistically significant (all $p < 0.05$, Mann Whitney U test).




Figure 6: **Performance of learning progress curriculum in Crafter with and without the proposed normalization extension.** Average task success rates and number of tasks with success rates more than 0.2 for learning progress curriculum with and without the proposed normalization extension across training steps. Metrics are calculated on all tasks. LP performs better with the proposed normalization extension.

B CRAFTER PROMPT

We give GPT-3 three examples as part of the prompt:

You are a player in a game. You want to learn as many different skills as possible. You can do this task well: collect wood.
Suggest whether the given tasks are interesting: collect drink, collect wood, make stone sword, make wood pickaxe, place furnace.
collect drink: True
collect 2 drink: True
collect 3 drink: True
collect wood: False
collect 2 wood: False
collect 3 wood: False
make stone sword: True
make 2 stone sword: True
make 3 stone sword: True
make wood pickaxe: True
make 2 wood pickaxe: True
make 3 wood pickaxe: True
place furnace: True
place 2 furnace: True
place 3 furnace: True

You are a player in a game. You want to learn as many different skills as possible. You can do this task well: make 2 iron pickaxe.
Suggest whether the given tasks are interesting: collect coal, collect iron, make iron pickaxe, make iron sword, place table.
collect coal: True
collect 2 coal: True
collect 3 coal: True
collect iron: True
collect 2 iron: True
collect 3 iron: True
make iron pickaxe: False
make 2 iron pickaxe: False
make 3 iron pickaxe: False
make iron sword: True

```
make 2 iron sword: True
make 3 iron sword: True
place table: True
place 2 table: True
place 3 table: True
```

```
You are a player in a game. You want to learn as many different
skills as possible. You can do this task well: place 3 stone.
Suggest whether the given tasks are interesting: collect diamond,
collect stone, make stone pickaxe, make wood sword, place stone.
collect diamond: True
collect 2 diamond: True
collect 3 diamond: True
collect stone: True
collect 2 stone: True
collect 3 stone: True
make stone pickaxe: True
make 2 stone pickaxe: True
make 3 stone pickaxe: True
make wood sword: True
make 2 wood sword: True
make 3 wood sword: True
place stone: False
place 2 stone: False
place 3 stone: False
```

Given the set of tasks that the agent can do relatively well and the set of tasks that needs to be determined as interesting or not, the additional prompt is:

```
You are a player in a game. You want to learn as many different
skills as possible. You can do these tasks well: <tasks done
well>.
Suggest whether the given tasks are interesting: <tasks to be
determined>.
```

In the few-shot examples, we do not include all possible tasks in <tasks to be determined> to reduce token usage. While each example sets different tasks for <tasks to be determined>, during inference, all tasks needing classification as interesting or not (Section 3.3) are inputted as <tasks to be determined> without any additional filtering.

GPT-3 predicts whether each task in <tasks to be determined> is interesting or not with True/ False, following the format in the few-shot examples. In our experience, GPT-3 nearly always conforms to the requested output format. However, in the rare cases that GPT-3 deviates from the expected output format, the responses are regenerated at a higher temperature, making the output less deterministic. For tasks where GPT-3 did not provide an answer, we modify the <tasks to be determined> input to include only the tasks lacking responses, and then regenerate these responses.

We access GPT-3 through OpenAI’s APIs, opting for the Davinci model in our experiments, which costs \$0.02 per 1000 tokens when these experiments were run. Caching significantly reduces the number of API queries. We extensively cache GPT-3 prompts and responses and consistently reuse this cache across multiple runs.

C BABYAI PROMPT

Full system prompt to GPT-4:

```
You are a helpful assistant that tells an AI agent the next
tasks to do in this 2D grid environment. The ultimate goal that
it would like your help with is to learn as many interestingly
different skills as possible, meaning a wide diversity of
```

different skills that would help it be ready to solve new skills someone might ask it to perform, or to transfer what it has learned in this environment to other environments.

The grid world has objects in six distinct colors - "red", "green", "blue", "purple", "yellow", and "grey" - and of four types - "key", "ball", "box", and "door". Each task is a sequence of instructions, connected using "then", specifying the order of instructions to follow. Instructions include interactions like going to objects, picking up items, opening doors (which requires the appropriately colored key if the door is locked), and putting objects next to another. Object placements are randomized.

I will give you the following information:
Tasks the agent currently do well: ...
Predict which of these tasks are interesting: ...

You must follow the following criteria:

- 1) You should act as a mentor and guide the AI agent to the next most interesting tasks.
- 2) Interesting tasks are roughly those that are sufficiently different from the ones that it can already do, and should be novel, diverse and at least worth learning.

You should only respond in the format as described below:

RESPONSE FORMAT:

Reasoning: Based on the information given, do reasoning about why each task is interesting or not.

Predictions: The list of predictions. For each line, put the task, then a colon, then True for interesting or False for boring.

Here are some example responses:

Tasks the agent currently do well: "open some door", "go to some object"

Predict which of these tasks are interesting: "pick up some object", "put some object next to some other object", "go to some object, then put some object next to some other object", "go to some object, then go to some object", "go to some object, then open some door", "open some door, then go to some object"

Reasoning: Tasks that are recombinations of "go to some object" or "open some door" are not interesting. The tasks that introduce something different from the tasks done well are "pick up some object", "put some object next to some other object", and "go to some object, then put some object next to some other object", having a new actions of picking up an object, or putting an object next to another.

Predictions:

"pick up some object": True

"put some object next to some other object": True

"go to some object, then put some object next to some other object": True

"go to some object, then go to some object": False

"go to some object, then open some door": False

"open some door, then go to some object": False

Tasks the agent currently do well: "go to some object, then open some door", "go to some object", "pick up some object"

Predict which of these tasks are interesting: "pick up some object", "go to some object, then open some door, then open some

door", "pick up some object, then put some object next to some other object, then put some object next to some other object", "put some object next to some other object, then go to some object, then put some object next to some other object, then open some door, then open some door"

Reasoning: Tasks that are recombinations of "go to some object, then open some door", "go to some object", or "pick up some object" are not interesting. The tasks that introduce something different from the tasks done well are "pick up some object, then put some object next to some other object, then put some object next to some other object" and "put some object next to some other object, then go to some object, then put some object next to some other object, then open some door, then open some door", having a new action of putting some object next to another.

Predictions:

"pick up some object": False

"go to some object, then open some door, then open some door":

False

"pick up some object, then put some object next to some other object, then put some object next to some other object": True

"put some object next to some other object, then go to some object, then put some object next to some other object, then open some door, then open some door": True

Tasks the agent currently do well: "go to some object", "put some object next to some other object", "pick up some object"

Predict which of these tasks are interesting: "pick up some object, then go to some object", "go to some object, then open some door, then open some door", "pick up some object, then put some object next to some other object, then put some object next to some other object", "open some door, then put some object next to some other object, then go to some object, then put some object next to some other object, then open some door", "go to some object, then pick up some object, then pick up some object, then put some object next to some other object, then go to some object"

Reasoning: Tasks that are recombinations of "go to some object", "put some object next to some other object", or "pick up some object" are not interesting. The only tasks that introduce something different from the tasks done well are "go to some object, then open some door, then open some door" and "open some door, then put some object next to some other object, then go to some object, then put some object next to some other object, then open some door", having a new action of opening a door.

Predictions:

"pick up some object, then go to some object": False

"go to some object, then open some door, then open some door":

True

"pick up some object, then put some object next to some other object, then put some object next to some other object": False

"open some door, then put some object next to some other object, then go to some object, then put some object next to some other object, then open some door": True

"go to some object, then pick up some object, then pick up some object, then put some object next to some other object, then go to some object": False

We access GPT-4 through OpenAI's APIs, which costs \$0.03 per 1000 tokens when these experiments were run. Caching significantly reduces the number of API queries. We extensively cache GPT-4 prompts and responses and consistently reuse this cache across multiple runs.

D PARTITIONING ALL TASKS INTO INTERESTING OR BORING

The MoI partitions all tasks as either “interesting” or “boring” based on their success rates and the LM’s assessment of their relation to tasks already classified as “interesting” (Algorithm 1). Since the raw evaluation of task success rates can be noisy, the task success rates referenced in this section are smoothed with an exponential moving average function. Algorithm 1 iteratively selects the task with the highest success rate not yet categorized, adds it to the “interesting” set (Step 3), prompts the LM to identify boring tasks from the remaining tasks in relation to the “interesting” set (Step 4), and updates the “boring” set (Step 5), repeating until all tasks are categorized. In Step 3, the rationale for adding the task with the highest success rate, that is not yet categorized, to the “interesting” set is twofold: first, the LM considers this task to be sufficiently distinct from those already in the “interesting” set, otherwise, it would have already been considered “boring” in Step 4; second, its relatively higher success rate indicates that it aligns closer to the agent’s current skill level.

To illustrate in Crafter, assume that the “collect wood” task has the highest success rate. It is added to the “interesting” set (Step 3). “collect wood” is interestingly different from all tasks currently declared as “interesting” because there are no tasks in that set yet. This is an initialization step for the algorithm. Then the LM deems “collect 2..10 wood” tasks as boring (Step 4). Now, when repeating Step 3, the algorithm will add the next task with the highest success rate and not yet categorized, e.g., “place table”, to the “interesting” set, whereas without the filter in Step 4, “collect 2 wood” might have been selected as interesting instead.

Algorithm 1 Mechanism to partition the task set into interesting and boring sets.

- 1: Sort the tasks based on the evaluated task success rates.
 - 2: Create two empty sets, one to track the interesting tasks and one to track the boring tasks.
 - 3: Identify the task with highest success rate and not in any of the sets. Add it to the interesting set.
 - 4: Prompt the LM to determine if any of the remaining tasks are boring, contexted on the current set of interesting tasks. Tasks in the interesting set are input as tasks that the agent can do well, and tasks yet to be categorized are asked to be predicted as interesting or not in the LM prompt.
 - 5: Update the boring set with tasks that the LM has determined as boring.
 - 6: Repeat steps 3 - 5 until all tasks are in either set.
-

E CRAFTER BORING TASKS

The 90 boring tasks in main experiments of the Crafter environment are:

- collect N drink, where $N \in [2, 10]$
- collect N wood, where $N \in [2, 10]$
- collect N coal, where $N \in [2, 10]$
- collect N stone, where $N \in [2, 10]$
- collect N iron, where $N \in [2, 10]$
- collect N diamond, where $N \in [2, 10]$
- place N table, where $N \in [2, 5]$
- place N furnace, where $N \in [2, 5]$
- place N stone, where $N \in [2, 5]$
- make N wood pickaxe, where $N \in [2, 5]$
- make N wood sword, where $N \in [2, 5]$
- make N stone pickaxe, where $N \in [2, 5]$
- make N stone sword, where $N \in [2, 5]$
- make N iron pickaxe, where $N \in [2, 5]$
- make N iron sword, where $N \in [2, 5]$

F CRAFTER POLICY AND OPTIMIZATION DETAILS

The model architecture is similar to those in previous learning progress works (Kanitscheider et al., 2021). The RGB inputs are passed through a 2-layer convolution network with ReLU activations. The RGB convnet is followed by a fully connected layer of size 256. The visual embeddings are concatenated with the task encoding before being passed into an LSTM cell of size 256. The network output is given by a 2-layer linear action head for the policy and a 2-layer linear layer for the value function. Both have Tanh activation functions. Optimization is performed with Proximal Policy Optimization (Schulman et al., 2017) and General Advantage Estimation (Schulman et al., 2015).

Parameter	Value
Discount factor	0.99
Learning rate	1e-4
PPO clip threshold	0.2
GAE lambda	0.95
Entropy coefficient	0.01
Batch size	2048
Epochs	4
Max episode length	1500

G BABYAI POLICY AND OPTIMIZATION DETAILS

The model architecture is similar to those in previous works done on BabyAI (Chevalier-Boisvert et al., 2018). The symbolic grid cell observations are passed through a 3-layer convolution network with ReLU activations. The natural language task is encoded with a GRU (Dey & Salem, 2017). The grid cell embeddings and the task embeddings are jointly processed through a convolutional network with two batch-normalized FiLM (Perez et al., 2018) layers, then through an LSTM cell of size 128. The network output is given by a 2-layer linear action head for the policy and a 2-layer linear layer for the value function. Both have Tanh activation functions. Optimization is performed with Proximal Policy Optimization (Schulman et al., 2017) and General Advantage Estimation (Schulman et al., 2015).

Parameter	Value
Discount factor	0.99
Learning rate	1e-4
PPO clip threshold	0.2
GAE lambda	0.95
Entropy coefficient	0.01
Batch size	2048
Epochs	4
Max episode length	depends on task difficulty 100 - 1000

H SUPPLEMENTARY METRICS AND PLOTS

The metrics and plots in this section complement the results shown in Section 4.3. Figure 10 shows the task sampling rates of all tasks in Crafter across different methods. To better see how different methods perform on interesting tasks in Crafter, Figure 7 offers the same metrics as Figures 3 and 4 but calculated on the subset of interesting tasks only. Figures 8 and 9 are subsets of Figures 3 and 10 respectively, zoomed into the set of interesting tasks only. Figure 4 and Figure 12 show the task success rates and task sampling rates of all tasks in BabyAI across different methods.




Figure 7: **Performance in Crafter on interesting tasks only.** Average task success rates and the number of tasks with success rates more than 0.2 for each method across training steps. OMNI achieves much higher average task success rates and learns more interesting tasks than uniform sampling or choosing tasks based on learning progress alone.

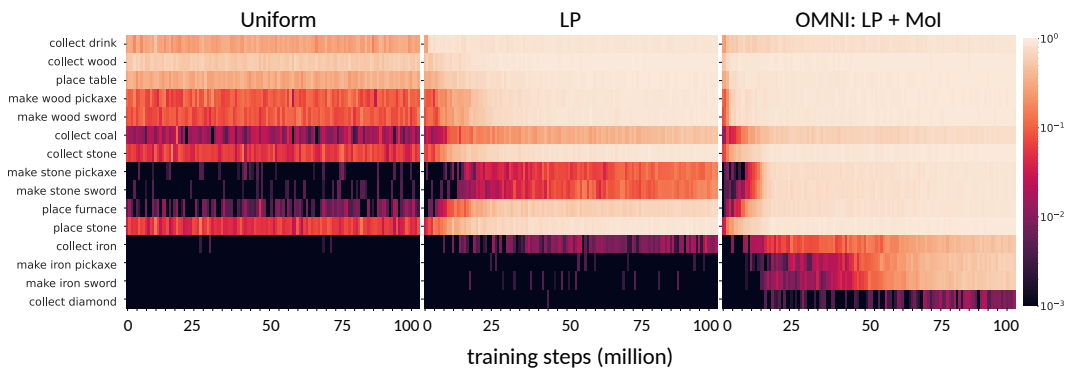


Figure 8: **Conditional success probabilities of interesting tasks only in Crafter.** This figure is a zoomed in subset of Figure 3. OMNI achieves higher task success rates on interesting tasks than uniform sampling or choosing tasks based on learning progress alone.




Figure 9: **Sampling probabilities for interesting tasks only in Crafter.** This figure is a zoomed in subset of Figure 10. OMNI exhibits more intense sampling of interesting tasks than uniform sampling or choosing tasks based on learning progress alone.




Figure 10: **Sampling probabilities for all tasks in Crafter.** Tasks are ordered as in Figure 3. Uniform sampling per task is barely visible in the heatmap. LP accurately tracks and samples tasks whose success probabilities change the most. OMNI samples tasks with high learning progress, but narrows its focus to only the interesting ones within that set.




Figure 11: **Conditional success probabilities of all tasks in BabyAI.** Tasks are organized from simple to complex based on the instruction length. Most tasks are very difficult for the agent to learn on, hence having low success rates. This mirrors the challenges in the real world where there can be countless difficult or infeasible tasks. Nevertheless, OMNI achieves higher task success rates than uniform sampling or choosing tasks based on learning progress alone.




Figure 12: **Sampling probabilities for all tasks in BabyAI.** Tasks are ordered as in Figure 4. Uniform sampling per task is barely visible in the heatmap. LP accurately tracks and samples tasks whose success probabilities change the most. OMNI samples tasks with high learning progress, but narrows its focus to only the interesting ones within that set.

I ORACLE MODEL OF INTERESTINGNESS

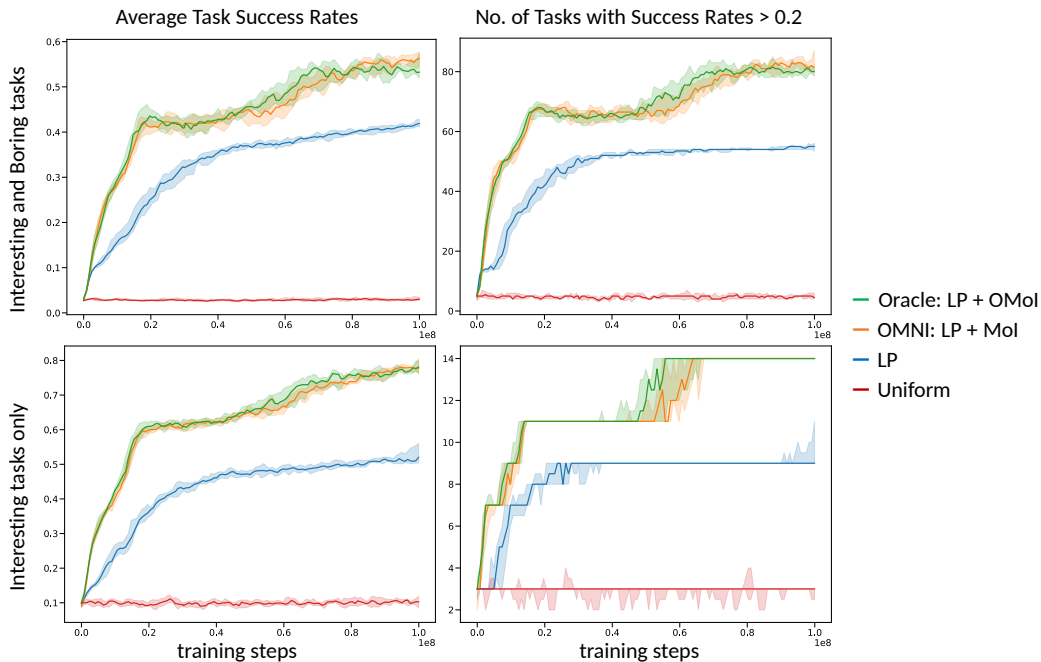


Figure 13: **Performance in Crafter on all tasks.** Average task success rates and the number of tasks with success rates more than 0.2 for each method across training steps. This figure is the same as Figures 3 and 4, with an additional oracle method added for comparison. OMNI achieves comparable performance to the oracle, with no statistically significant difference.

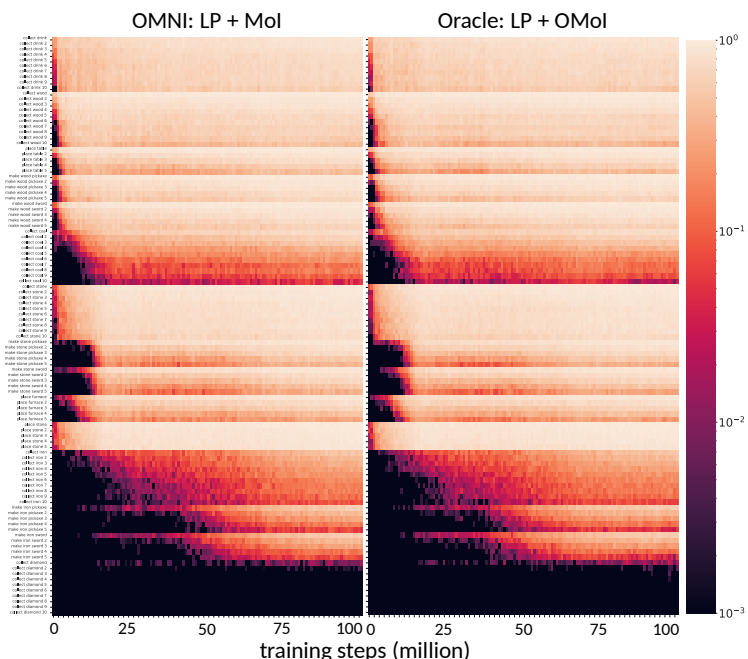


Figure 14: **Conditional success probabilities of all tasks for OMNI and the oracle in Crafter.** Tasks are ordered as in Figure 3. OMNI achieves comparable task success rates to the oracle.

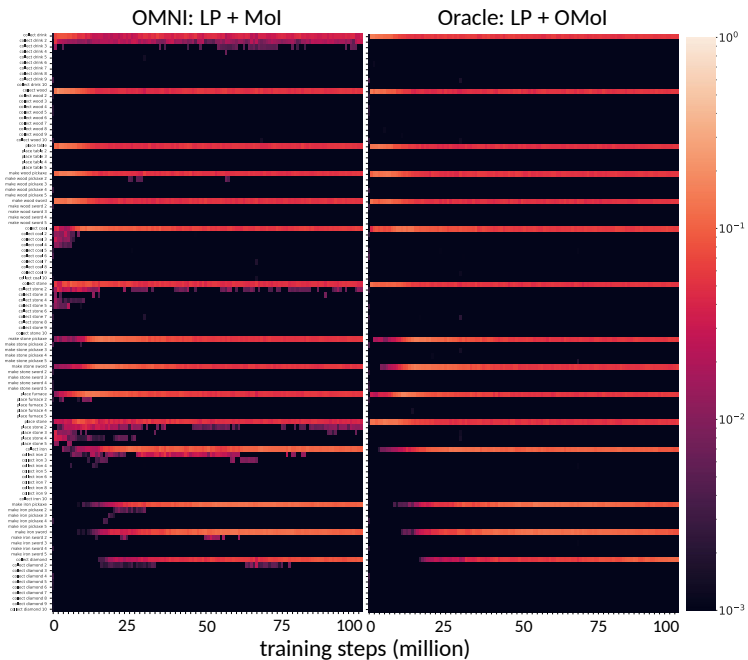


Figure 15: **Sampling probabilities for all tasks for OMNI and the oracle in Crafter.** Tasks are ordered as in Figure 3. OMNI focuses on interesting tasks with high learning progress, similar to that of the oracle.

The Oracle Model of Interestingness (**OMoI**) is a hand-designed model that assigns sampling weights of 1.0 to interesting tasks and 0.001 to boring tasks. In Crafter, interesting tasks are considered to be the set of 15 tasks shown in Figure 2, and all other tasks are considered boring. In BabyAI, interesting tasks are those with single instructions. Single-instruction tasks are the foundational building blocks that should be learned first, so we consider those interesting in the OMoI. In theory, a more skilled agent would eventually move beyond them to two or more instruction tasks, but in our experiments,

our agents do not even learn this first set (despite substantial compute budgets; BabyAI is a very hard RL task environment if one does not take the shortcut of providing solution demonstrations and doing imitation learning (Chevalier-Boisvert et al., 2018)), meaning that these tasks are uninteresting for them to be focusing on. The OMol functions as an oracle that accurately discerns which tasks humans would typically find interesting for the agent to learn the most skills in this environment. Although we were able to design an oracle MoI in these domains, it is important to note that constructing such models in more complex domains will not always be feasible, nor would it scale well given the human labor required.

Across all metrics and in both domains, most of the differences in performance between OMNI and the oracle at 25%, 50%, 75%, and 100% of the way through training are not statistically significant (all but one $p > 0.05$, Mann Whitney U test). OMNI’s performance is comparable to that of the oracle (Figures 13 and 16). OMNI and the oracle achieve similar task success rates for each task (Figures 14 and 17), and induce similar patterns in task sample rates (Figures 15 and 18). This suggests that LMs can capture key aspects of what humans typically find interesting.



Figure 16: **Performance in BabyAI on all tasks.** Average task success rates and the number of tasks with success rates more than 0.2 for each method across training steps. This figure is the same as Figures 3 and 4, with an additional oracle method added for comparison. OMNI achieves comparable performance to the oracle, with almost no statistically significant difference.




Figure 17: **Conditional success probabilities of all tasks for OMNI and the oracle in BabyAI.** Tasks are ordered as in Figure 4. OMNI achieves comparable task success rates to the oracle.

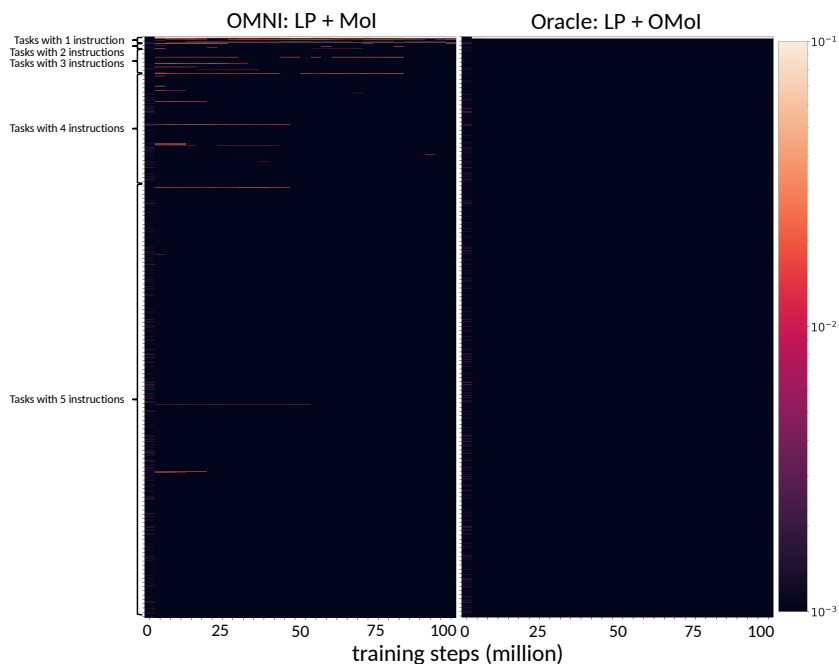


Figure 18: **Sampling probabilities for all tasks for OMNI and the oracle in BabyAI.** Tasks are ordered as in Figure 4. OMNI focuses on interesting tasks with high learning progress, overlapping with that of the oracle, which only samples tasks with one instruction.

J USING COMPOUNDS AS BORING TASKS IN CRAFTER

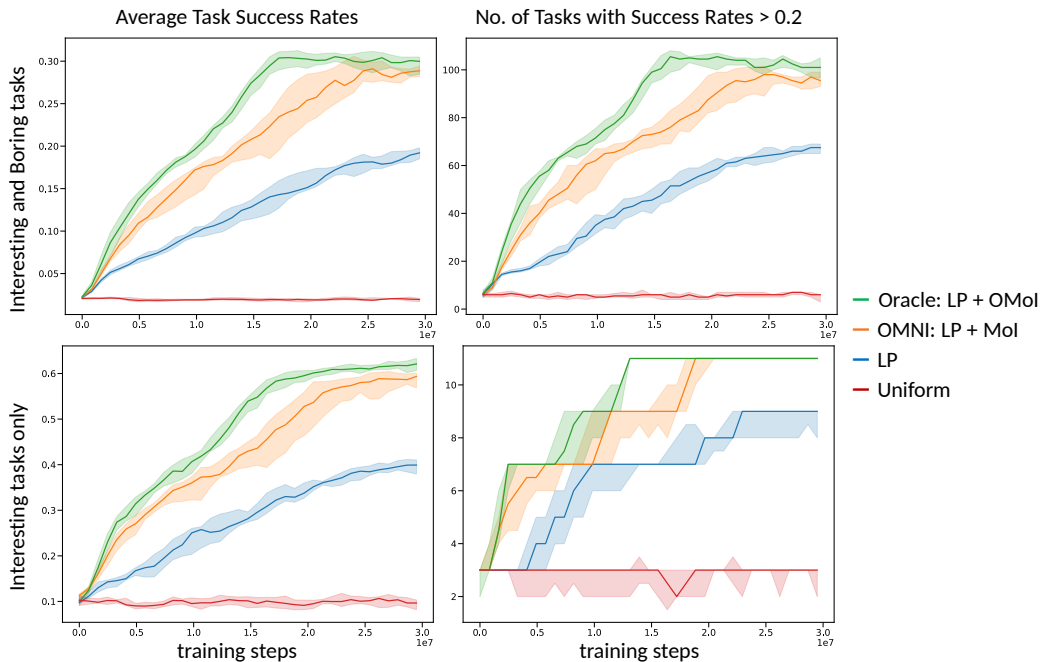


Figure 19: **Performance in Crafter on all tasks, including compound tasks.** Average task success rates and the number of tasks with success rates more than 0.2 for each method across training steps. Compound and repetitive tasks are those considered boring. OMNI achieves much higher average task success rates and learns more tasks than LP or Uniform, close to that of the oracle.




Figure 20: **Conditional success probabilities of all tasks, including compound tasks, in Crafter.** Agents are trained with compound and repetitive tasks as boring tasks. Tasks are ordered as in Figure 3, with additional compound tasks at the bottom. OMNI achieves higher task success rates than LP and Uniform in a wide range of tasks, comparable to that of the oracle.

In this Crafter setup, compound and repetitive tasks are those considered boring. In addition to the repetitive boring tasks (Appendix E), compound tasks are generated by combining any two of the 15 interesting tasks (Figure 2). Hence, there is a total of 15 interesting tasks, 195 boring tasks, and 1023 extremely challenging tasks. In light of limited compute, experiments in this setting are run for 30 million time steps (vs. 100 million in the main experiments) and are repeated 10 times (same as the main experiments) with different random seeds. We compare the performance of agents trained with: (1) **Uniform** sampling, (2) Learning Progress (**LP**) only, (3) OMNI: Learning Progress with additional filtering by a Model of Interestingness (**OMNI: LP + MoI**), and (4) the oracle: Learning Progress with additional filtering by an Oracle Model of Interestingness (**Oracle: LP + OMoI**). The high-level summary of the results from these experiments is that they are qualitatively similar to when boring tasks are repetitive tasks only (Section 4.3).

Uniform sampling. The results with uniform sampling are poor (Figure 19). The agent learns 6 (CI: 3 – 6) tasks (interesting or boring) and 3 (CI: 2 – 3) interesting tasks, achieving an average task success rate of 0.019 (CI: 0.017 – 0.021) on interesting and boring tasks, and 0.097 (CI: 0.082 – 0.105) on interesting tasks only.

Learning Progress Curriculum. The agent learns 67 (CI: 65 – 69) tasks (interesting or boring) and 9 (CI: 8 – 9) interesting tasks, achieving an average task success rate of 0.19 (CI: 0.19 – 0.20) on interesting and boring tasks, and 0.40 (CI: 0.38 – 0.42) on interesting tasks only. Across all metrics, the differences between LP and Uniform at 25%, 50%, 75%, and 100% of the way through training are statistically significant (all $p < 1e-3$, Mann Whitney U test), showing that LP significantly outperforms uniform sampling (Figure 19). LP accurately tracks and samples tasks with high learning progress but is distracted by the many boring tasks (Figure 21).

OMNI: Learning Progress + a Model of Interestingness. The agent learns 96 (CI: 93 – 99) tasks (interesting or boring) and 11 (CI: 11 – 11) interesting tasks, achieving an average task success rate

of 0.29 (CI: 0.28 – 0.30) on interesting and boring tasks, and 0.59 (CI: 0.56 – 0.60) on interesting tasks only. Across all metrics, the differences between OMNI and LP at 25%, 50%, 75%, and 100% of the way through training are statistically significant (all $p < 1e-3$, Mann Whitney U test), showing that OMNI significantly outperforms LP (Figure 19). Since OMNI focuses on interesting tasks with high learning progress (Figure 21), the agent achieves higher average task success rates and learns more challenging tasks faster (Figure 20).

Oracle: Learning Progress + an Oracle Model of Interestingness. The agent learns 101 (CI: 97 – 105) tasks (interesting or boring) and 11 (CI: 11 – 11) interesting tasks, achieving an average task success rate of 0.30 (CI: 0.29 – 0.31) on interesting and boring tasks, and 0.62 (CI: 0.61 – 0.63) on interesting tasks only. The differences between the oracle and OMNI at 25%, 50%, 75%, and 100% of the way through training are sometimes statistically significant (some $p < 0.05$, Mann Whitney U test), showing that OMNI does not always achieve comparable performance to the oracle (Figure 19). However, OMNI emulates similar task success rates (Figure 20) and task sample rates to the oracle (Figure 21). Having OMNI approach the performance benchmark set by the oracle is already an indicator of success.




Figure 21: **Sampling probabilities for all tasks, including compound tasks, in Crafter.** Agents are trained with compound and repetitive tasks as boring tasks. Tasks are ordered as in Figure 20. LP accurately tracks, and thus samples, tasks whose success probabilities change the most. OMNI narrows its focus to only the interesting ones within the set of high learning progress tasks.

K USING SYNONYMS AS BORING TASKS IN CRAFTER

In this Crafter setup, synonymous and repetitive tasks are those considered boring. In addition to the repetitive boring tasks (Appendix E), synonymous tasks are those with different task representations (i.e., synonymous descriptions of tasks like “collect wood” and “gather wood”) but the same success conditions. In this setting, the OMoI regards the 15 tasks shown in Figure 2 and their synonyms as interesting. The repetitive tasks and their synonyms are also regarded as boring. Therefore, the OMoI identifies 90 interesting tasks, 540 boring tasks, and 1023 extremely challenging tasks. The same classification is utilized for subsequent analysis. In light of limited compute, experiments in this

setting are run for 30 million time steps (vs. 100 million in the main experiments) and are repeated 10 times (same as the main experiments) with different random seeds. We compare the performance of agents trained with: (1) **Uniform** sampling, (2) Learning Progress (**LP**) only, (3) OMNI: Learning Progress with additional filtering by a Model of Interestingness (**OMNI: LP + MoI**), and (4) the oracle: Learning Progress with additional filtering by an Oracle Model of Interestingness (**Oracle: LP + OMoI**). The high-level summary from these experiments is that, when furnished with adequate information about the agent, OMNI delivers qualitatively similar results to when boring tasks are repetitive tasks only (Section 4.3), or compound and repetitive tasks (Section J).

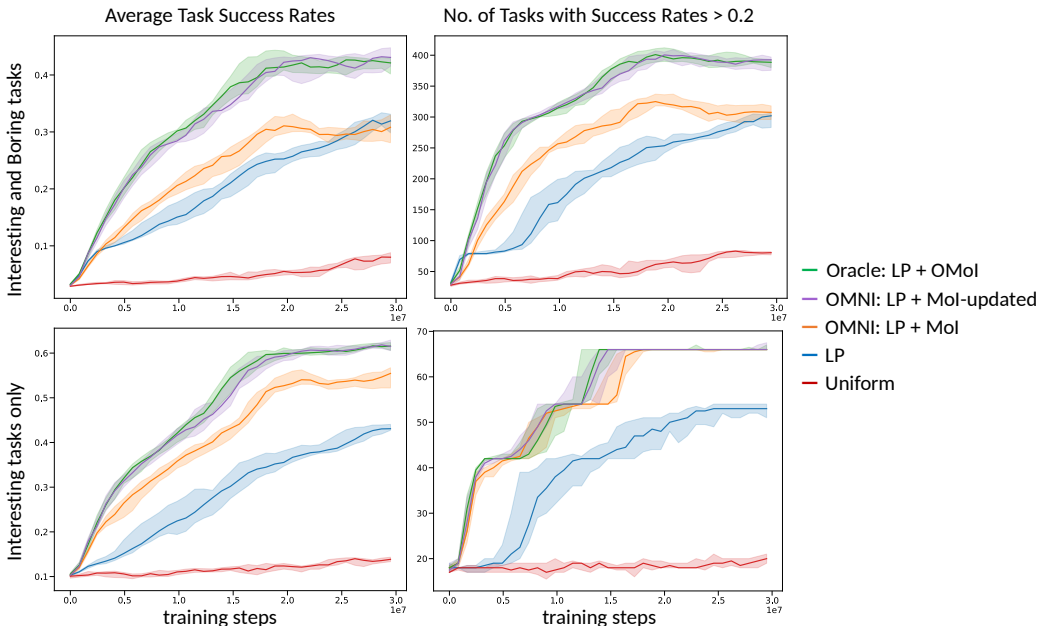


Figure 22: **Performance in Crafter on all tasks, including synonymous tasks.** Average task success rates and the number of tasks with success rates more than 0.2 for each method across training steps. Agents are trained with synonymous and repetitive tasks as boring tasks. OMNI: LP + MoI-updated achieves much higher average task success rates and learns more tasks than LP, Uniform, or LP + MoI, comparable to the oracle.

Original Verb	Synonyms
collect	gather harvest procure acquire amass
make	craft acquire build construct create
place	put deploy install putdown position

Table 1: Synonyms used to generate boring tasks in the Crafter environment

Uniform sampling. The results with uniform sampling are poor (Figure 22). The agent learns 80 (CI: 77 – 82) tasks (interesting or boring) and 20 (CI: 19 – 21) interesting tasks, achieving an average task success rate of 0.080 (CI: 0.070 – 0.088) on interesting and boring tasks, and 0.14 (CI: 0.13 – 0.14) on interesting tasks only.

Learning Progress Curriculum. The agent learns 302 (CI: 287 – 306) tasks (interesting or boring) and 53 (CI: 52 – 54) interesting tasks, achieving an average task success rate of 0.32 (CI: 0.31 – 0.33) on interesting and boring tasks, and 0.43 (CI: 0.42 – 0.44) on interesting tasks only. Across all metrics, the differences between LP and Uniform at 25%, 50%, 75%, and 100% of the way through training are statistically significant (all $p < 1e-2$, Mann Whitney U test), showing that LP significantly outperforms uniform sampling (Figure 22).

OMNI: Learning Progress + a Model of Interestingness. The agent learns 307 (CI: 296 – 318) tasks (interesting or boring) and 66 (CI: 66 – 66) interesting tasks, achieving an average task success rate of 0.31 (CI: 0.28 – 0.33) on interesting and boring tasks, and 0.56 (CI: 0.52 – 0.57) on interesting tasks only. On the interesting tasks only, the differences between OMNI and LP at 25%, 50%, 75%, and 100% of the way through training are statistically significant (all $p < 1e-2$, Mann Whitney U test), showing that OMNI significantly outperforms LP (Figure 22). However, on all tasks (interesting and boring), the differences between OMNI and LP at 25%, 50%, 75%, and 100% of the way through training are not always statistically significant (not all $p < 0.05$, Mann Whitney U test), showing that OMNI does not outperform LP on all metrics (Figure 22). We further investigate this by comparing OMNI with the oracle. In the subsequent paragraphs, we introduce a minor modification to OMNI, which then significantly outperforms LP (as occurred in the previous experimental settings).

Oracle: Learning Progress + an Oracle Model of Interestingness. The agent learns 389 (CI: 380 – 393) tasks (interesting or boring) and 66 (CI: 66 – 67) interesting tasks, achieving an average task success rate of 0.42 (CI: 0.41 – 0.43) on interesting and boring tasks, and 0.62 (CI: 0.61 – 0.62) on interesting tasks only. On all tasks (interesting and boring), the differences between the oracle and OMNI: LP + MoI at 25%, 50%, 75%, and 100% of the way through training are statistically significant (all $p < 1e-3$, Mann Whitney U test), showing that the oracle significantly outperforms OMNI: LP + MoI on the set of all tasks (Figure 22). This difference in performance can be attributed to the different interestingness notions held by the OMoI vs. the MoI. The current LM-based MoI in OMNI regards all synonymous tasks as boring, even if they are not repetitive tasks. This is because the LM-based MoI assumes that the RL agent already has a language prior and that sampling synonymous tasks would not enable the agent to learn more skills. In other words, the LM-based MoI incorrectly assumes that the RL agent knows how to do the task “gather wood” if it knows how to do the task “collect wood”, not understanding that the RL agent has not yet learned that these two tasks are actually the same thing. Recall that the RL agent in our setup does not have a language prior yet, because the tasks are encoded as a bag-of-words instead of using a pre-trained natural language encoder (which is likely to encode “gather wood” and “collect wood” closer in the embedding space than other non-synonymous tasks).

OMNI: Learning Progress + an updated Model of Interestingness We update the prompt input to GPT-3, adding information that the agent currently lacks a language prior and perceives synonymous tasks as completely different (Figure 23). The MoI with updated prompt is labelled as **MoI-updated**. The MoI-updated now predicts that synonymous tasks are still interesting. The agent trained with OMNI: LP + MoI-updated learns 393 (CI: 376 – 397) tasks (interesting or boring) and 66 (CI: 66 – 67) interesting tasks, achieving an average task success rate of 0.43 (CI: 0.41 – 0.45) on interesting and boring tasks, and 0.62 (CI: 0.61 – 0.63) on interesting tasks only. Across all metrics, the differences in performance between OMNI: LP + MoI-updated and the oracle at 25%, 50%, 75%, and 100% of the way through training are not statistically significant (all $p > 0.05$, Mann Whitney U test). This shows that OMNI: LP + MoI-updated’s performance is comparable to that of the oracle (Figure 22). OMNI: LP + MoI-updated and the oracle achieve similar task success rates for each task (Figure 25), and induce similar patterns in task sample rates (Figure 26).

These experiments demonstrate that OMNI can be effectively utilized across a range of task settings when provided with sufficient information. Initially, the MoI was not aware that the RL agent lacked a language prior, which led it to categorize synonymous tasks as boring. However, by supplementing the prompt with additional information about the agent (i.e., its lack of a language prior), OMNI was able to recognize synonymous tasks as interesting, thereby accelerating the agent’s learning.

An alternative approach could involve providing the MoI with more information about the agent’s performance, so that it can analyze the data, identify and adapt to the agent’s inherent limitations (e.g., the absence of a language prior). LMs can potentially automatically analyze the task success rates and update their predictions of what is interesting without human intervention. Preliminary investigations with both GPT-3 and GPT-4 have shown promising results. These models generated plausible analyses

of the differences in task success rates between synonymous tasks, and adjusted their predictions of what is interesting based on these analyses (Figures 24 and 27). These observations imply that LMs may already have the capacity to analyze data autonomously and adapt their outputs accordingly.



Figure 23: GPT-3 output (**Left**) from original prompt (synonymous tasks are determined as not interesting) and (**Right**) after updating prompt (synonymous tasks are determined as interesting). GPT-3 can adapt its prediction of what tasks are interesting to the prompt update.



Figure 24: GPT-3 output (**Left**) before analysis (synonymous tasks are deemed not interesting), (**Middle**) analysis of task success rates (success rates of synonymous tasks are different), and (**Right**) after analysis (synonymous tasks are deemed interesting). GPT-3 generated a plausible analysis for the agent’s achieved task success rates, and automatically updated its interestingness predictions.




Figure 25: **Conditional success probabilities of all tasks, including synonymous tasks, in Crafter.** Agents are trained with synonymous and repetitive tasks as boring tasks. Tasks are organized based on the prerequisite tasks that must be accomplished in order to complete the target task. Task names (left of each row) are readable in a digital format with zoom. OMNI: LP + Mol-updated achieves higher task success rates than LP and Uniform in a wide range of tasks, comparable to that of the oracle.




Figure 26: **Sampling probabilities for all tasks, including synonymous tasks, in Crafter.** Agents are trained with synonymous and repetitive tasks as boring tasks. Tasks are ordered as in Figure 25. LP accurately tracks, and thus samples, tasks whose success probabilities change the most. OMNI: LP + MoI-updated focuses on interesting tasks with high learning progress, similar to that of the oracle.

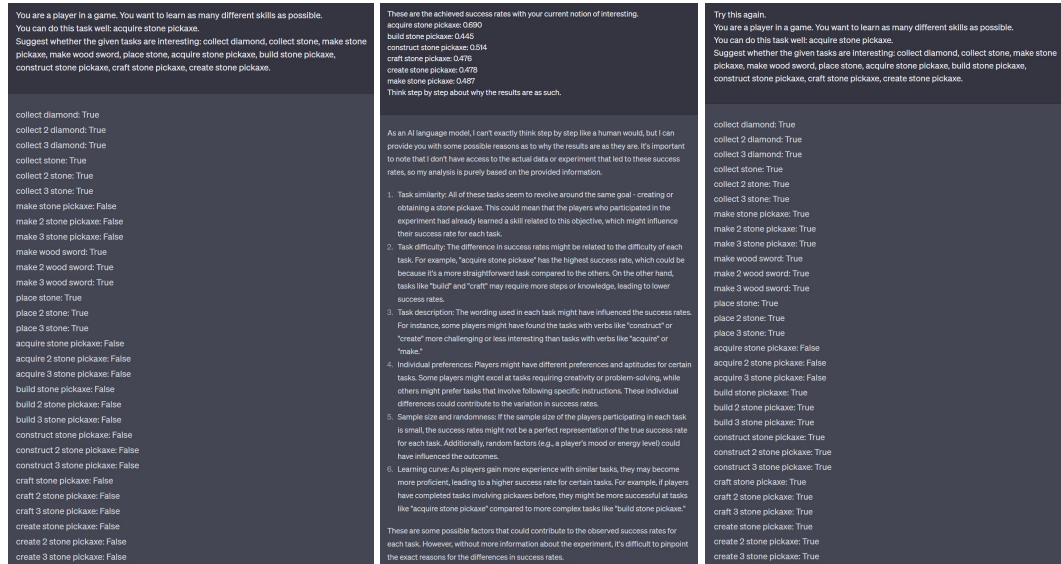


Figure 27: GPT-4 output (**Left**) before analysis (synonymous tasks are deemed not interesting), (**Middle**) analysis of task success rates (success rates of synonymous tasks are different), and (**Right**) after analysis (synonymous tasks are deemed interesting). The darker shaded area is the input prompt, the lighter shaded area is the generated output. GPT-4 generated a plausible analysis for the agent’s achieved task success rates, and automatically updated its interestingness predictions.

L INTEGRATING MORE ASPECTS OF INTERESTINGNESS INTO LMS

For future research, we plan to explore the potential of leveraging LMs to address multiple aspects of interestingness, including learning progress. By examining these aspects individually or in a combined manner, we aim to develop more robust and adaptive models that can efficiently cater to a wide range of tasks and applications. Preliminary evidence suggests that LMs, such as GPT-4, possess the capacity to integrate multiple aspects of interestingness (Figure 28). Further investigation and experimentation in this direction could lead to the development of AI models that exhibit a more sophisticated approach to learning by effectively capturing and integrating various aspects of interestingness.

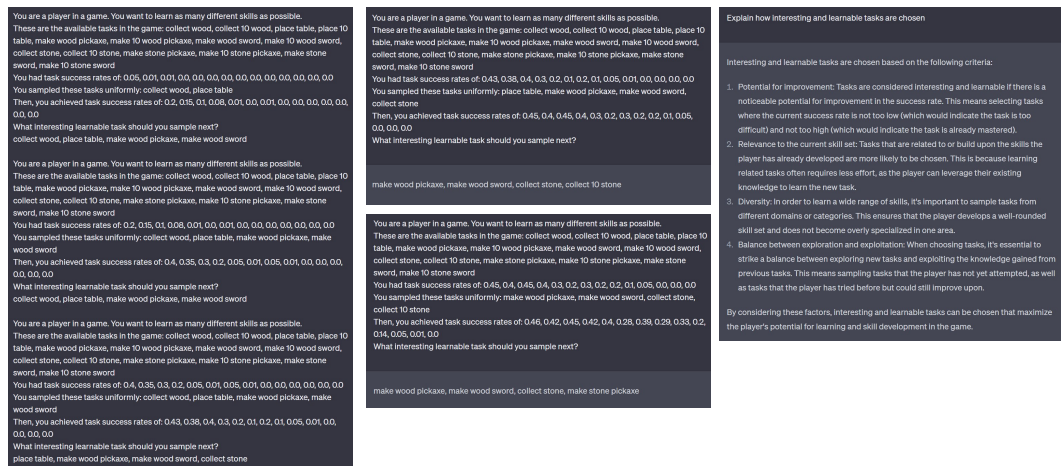


Figure 28: (**Left**) Few-shot examples part of first prompt input to GPT-4. (**Middle**) GPT-4’s prediction of learnable and interesting tasks. (**Right**) GPT-4 explanation of how interesting and learnable tasks are chosen. The darker shaded area is the input prompt, the lighter shaded area is the generated output. By providing GPT-4 with a few examples of what learnable and interesting tasks are, the model can synthesize these concepts and consistently generate a set of tasks that meet the desired criteria.