

Playing NetHack with LLMs: Potential & Limitations as Zero-Shot Agents

Dominik Jeurissen, Diego Perez-Liebana, Jeremy Gow
Queen Mary University of London
{d.jeurissen, diego.perez, jeremy.gow}@qmul.ac.uk

Duygu Çakmak, James Kwan
Creative Assembly
{duygu.cakmak, james.kwan}@creative-assembly.com

Abstract—Large Language Models (LLMs) have shown great success as high-level planners for zero-shot game-playing agents. However, these agents are primarily evaluated on games where long-term planning is relatively straightforward. In contrast, agents tested in more dynamic environments face limitations due to simplistic environments with only a few objects and interactions. To fill this gap in the literature, we present NetPlay, the first LLM-powered zero-shot agent for the challenging roguelike NetHack. NetHack is a particularly challenging environment due to its diverse set of items and monsters, complex interactions, and many ways to die.

NetPlay uses an architecture designed for dynamic robot environments, modified for NetHack. Like previous approaches, it prompts the LLM to choose from predefined skills and tracks past interactions to enhance decision-making. Given NetHack’s unpredictable nature, NetPlay detects important game events to interrupt running skills, enabling it to react to unforeseen circumstances. While NetPlay demonstrates considerable flexibility and proficiency in interacting with NetHack’s mechanics, it struggles with ambiguous task descriptions and a lack of explicit feedback. Our findings demonstrate that NetPlay performs best with detailed context information, indicating the necessity for dynamic methods in supplying context information for complex games such as NetHack.

Index Terms—NetHack, Large Language Models, Zero-Shot Agent.

I. INTRODUCTION

Recently, agents based on Large Language Models (LLMs) [1] have been successfully applied to robot environments [2], [3] and Minecraft [4]–[6], among others. These agents do not require pre-training and typically involve prompting an LLM to solve tasks by choosing from predefined skills.

LLM agents utilizing predefined skills are particularly promising for game development as developing a set of skills is often more feasible than designing an entire agent. However, existing studies predominantly focus on the capabilities of LLMs for game-playing, neglecting to address their limitations. Evaluations typically focus on predictable tasks, like finding a diamond in Minecraft, which can consistently be achieved through strip mining. Many games require more dynamic decision-making, where long-term planning is challenging, and the correct course of action is more ambiguous. While evaluations have been done on more dynamic environments,

these environments often contain only a handful of objects and lack complex interactions.

We build upon existing literature by evaluating an LLM agent in the context of the complex and unpredictable roguelike NetHack [7]. NetHack is a challenging game with many monsters, items, interactions, partial observability, and an intricate goal condition. The sheer size of NetHack, paired with the many sub-systems the player has to understand, make it an excellent candidate for evaluating the limitations of LLM agents. NetHack’s description files also allow us to define levels, enabling us to evaluate the agent’s abilities in isolation.

In the following, we present NetPlay, a LLM powered agent designed to tackle a wide range of tasks in NetHack. NetPlay is inspired by *autoascend* [8] a handcrafted agent that won the NetHack Challenge 2021 [9]. While *autoascend* relied on a large network of handcrafted rules to handle the complexity of NetHack, NetPlay only requires a set of isolated skills. Our experiments show that NetPlay can interact with most of NetHack’s game mechanics and that it excels in following detailed instructions. Additionally, the agent exhibits creative behavior when focusing its attention on a specific problem. However, when tasked to play autonomously, NetPlay is far outperformed by *autoascend*. Consequently, this paper delves into reasons for this, such as the agent’s struggles to handle ambiguous instructions, confusing observations, and a lack of explicit feedback.

We begin in section II with an overview of NetHack and a review of existing work on LLM-powered agents. Section III discusses the architecture of NetPlay, including many of the design decisions we had to make due to limitations caused by the LLM. In section IV, we first evaluate NetPlay’s ability to autonomously play the game and compare its performance with a simple handcrafted agent and *autoascend*. We follow this up with an in-depth analysis of the agent’s behavior across various isolated scenarios. Subsequently, we analyze the experiment results in section V and conclude this study in section VI. The source code can be found on GitHub¹.

II. BACKGROUND

A. NetHack

NetHack [7], released in 1987, is an extremely challenging turn-based roguelike that continues to receive updates to this

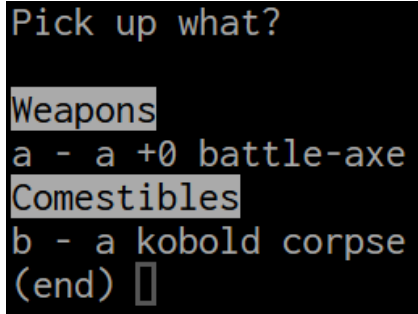
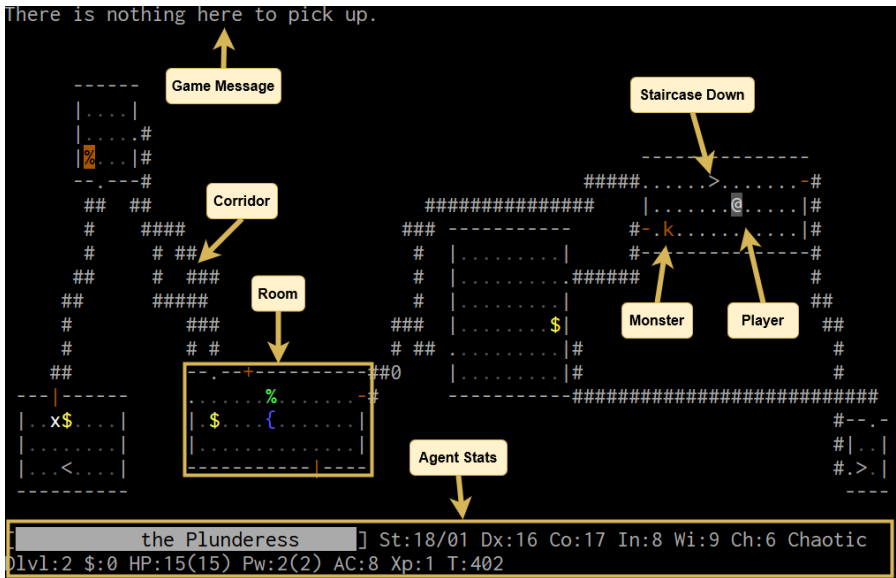


Fig. 1: The terminal view of the game NetHack. The left image presents an annotated view of the in-game screen, featuring the game’s map, an example of a game message, and the agent’s stats. The right image showcases a menu for picking up items from a tile containing multiple objects. Image source: alt.org/nethack

date. The objective is to traverse 50 procedurally generated levels, retrieving the Amulet of Yendor and successfully returning to the surface. Doing so unlocks the final challenge of the game: the four elemental planes, followed by the astral plane, where players must present the Amulet to their deity. See fig. 1 for a screenshot of the game.

Most aspects of the game are generated, such as level layouts, the player’s starting class, and the inventory. The levels follow a somewhat linear structure with many branches and sub-dungeons in between. For instance, the entrance to the gnomish mines always spawns somewhere between depth 2 and 4, giving the player the option to explore them immediately or to postpone exploration until they are stronger.

NetHack encompasses a diverse array of monsters, items, and interactions. Players must skillfully utilize their resources while avoiding many of the game’s lethal threats. Even for seasoned players possessing extensive knowledge of the game, victory is far from guaranteed. The game’s inherent complexity requires players to continuously re-assess their situation to adapt to the unpredictability of the elements at play.

NetHack uses description files (**des-files**) to describe special levels like the oracle level that always contains a room with an oracle monster, centaur statues, and four fountains. Des-files offer extensive control over the level-generation process, allowing entirely handcrafted levels or a slightly constrained level-generation process.

B. NetHack Learning Environment

The NetHack Learning Environment **NLE** [10] serves as a reinforcement learning environment for playing NetHack 3.6.6. NLE offers easy access to most aspects of the game, such as the map, the agent’s inventory, game messages, and the player’s stats. While NLE provides simplified environments

for learning purposes, it also allows users to play the entire game without any restrictions.

MiniHack [11] utilizes NLE alongside des-files to construct small-scale environments that isolate specific challenges that agents will encounter in NetHack. Although MiniHack provides a list of challenges, its primary purpose is to streamline the process of designing new challenges.

C. Autoascend

In the 2021 NeurIPS NetHack Challenge [9], participants tackled the symbolic and neural tracks, where solutions were either handcrafted or designed using machine learning. Notably, the top-performing agents were exclusively symbolic, with the autoascend agent emerging as the frontrunner.

The autoascend agent [12] succeeded by meticulously parsing observations and creating an internal state representation to track essential information. The agent utilized the enriched data to implement a behavior tree by hierarchically combining strategies representing specific behaviors, like fighting, picking up objects, or exploring levels. Overall, autoascend’s strategy consists of staying on the first dungeon level until reaching experience level 8, after which it will rapidly progress deeper into the dungeon. While following this general strategy, autoascend uses many sub-strategies to improve its chance of success, such as a solver for solving the Sokoban levels, using altars for farming or identifying items, or dipping a long sword into a fountain to gain Excalibur.

Despite its victory, autoascend’s performance depended heavily on its starting class, demonstrating optimal results with the Valkyrie class. The agent occasionally descended to depth 10 and reached experience level 10. However, it is crucial to highlight that reaching around depth 50 is only one of the

objectives to beat NetHack, emphasizing how challenging the environment still is.

D. LLM Agents

Recently, a plethora of LLM-based agents have emerged, aiming to leverage the planning capabilities of these models. A prominent testbed for these agents is Minecraft [4]–[6], primarily focusing on the agent’s ability to obtain the various items in the game. While the details vary, most approaches implement a closed-loop planning system in which the LLM generates a plan consisting of a sequence of predefined skills. The plan is then executed and, in case of failure, the agent will re-plan using only feedback from the previous plan. A noteworthy aspect of these agents is the storage and reuse of successful plans, significantly enhancing overall performance due to the hierarchical nature of obtaining items like a diamond pickaxe. The agents primarily utilize an LLM for their knowledge of how to acquire items. However, one agent has demonstrated the ability to construct structures with human feedback [6].

Other popular applications are robot environments, where tasks include rearranging objects on a tabletop, interacting within a kitchen, or engaging in simulated household activities [2], [3], [13], [14]. Because these environments require more dynamic decision-making compared to acquiring items in Minecraft, agents like DEPS [13] and Inner Monologue [2] adopt a distinctive approach. Instead of relying solely on feedback from the last failed plan, they re-plan by considering a substantial portion of their recent interaction history. Similar to our approach, Inner Monologue models the interaction history as a chat containing the LLM’s actions and thoughts, human feedback, and feedback from the environment, such as scene descriptions and if an action was successful. While the robot environments require more dynamic decision-making, the complexity of the observations is limited, usually consisting of a list of visible objects with spatial information being omitted as the low-level skills are handling it.

An alternative use of LLMs involves employing them to design reward functions, which are then used to train reinforcement learning agents [15]–[17]. Most relevant to our work, Motif employs an LLM to learn various playstyles in NetHack. It achieves this by tasking the LLM to decide which of NetHack’s game messages it prefers. Motif can leverage these preferences to learn reward functions for different playstyles by conditioning the LLM to prefer game messages associated with a specific playstyle, such as fighting monsters.

III. NETPLAY

This section discusses our LLM-powered Nethack agent NetPlay. See fig. 2 for an overview of the architecture. Long-term planning in NetHack proves challenging due to its unpredictability, as we cannot know when, where, or what will appear as we explore. Consequently, our agent shares many similarities with Inner Monologue, designed for dynamic environments. It implements a closed-loop system where the LLM selects skills sequentially while accumulating feedback

in the form of game messages, errors, or manually detected events. Although we avoid constructing entire plans, the LLM’s thoughts are included for future prompts, allowing for strategic planning if deemed necessary by the LLM. Notably, while the presented architecture enables an LLM to play NetHack, it is still too limiting to allow for detailed strategic decision-making.

A. Prompting

We prompt the LLM to choose a skill from a predefined list. The prompt comprises three components: (a) a description of the observation, (b) the agent’s short-term memory, and (c) a task description alongside the output format.

(a) The observation description primarily focuses on the current level alongside additional data like context, inventory, and the agent’s stats. Because we do not use a multi-modal LLM, we attempt to convey spatial information by dividing the level into structures like rooms and corridors. Each structure is described using a unique identifier, the number of steps to reach it, the objects it contains with their respective positions, and the number of steps to reach each object. Monsters are described separately from the structures by categorizing them as close or distant, indicating their potential threat level. Each monster is described using its name, position, and number of steps to reach it. For close monsters, we also include compass coordinates. The LLM is also informed about which structures can be further explored alongside the positions of boulders and doors that block exploration progress. Note that despite our emphasis on providing spatial information, navigating the environment proved challenging for the LLM. Consequently, we automated a large portion of the exploration process using a single skill, potentially rendering certain aspects of this observation description obsolete.

(b) The short-term memory is implemented using a list of messages representing the timeline of events. Each message is either categorized as system, AI, or human. System messages convey feedback from the environment like game messages or errors, AI messages capture the LLM’s responses, and new tasks are indicated by human messages. Note that while it is possible for a human to provide continual feedback, we only study the case where the agent is given a task at the start of the game. The memory size is capped at 500 tokens, with older messages being deleted first. Observation descriptions are not stored in the memory due to their size.

(c) The task description includes details about the current task, available skills, and a JSON output format. We employ chain-of-thought prompting [18] to guide the LLM to a skill choice.

B. Skills

Skills, similar to strategies in `autoascend`, implement specific behaviors by returning a sequence of actions. They accept both mandatory and optional parameters as input. Skills can generate messages as feedback, which are stored in the agent’s memory. Messages are often used, for example, to report why a skill failed. An excerpt of skills can be found in table I.

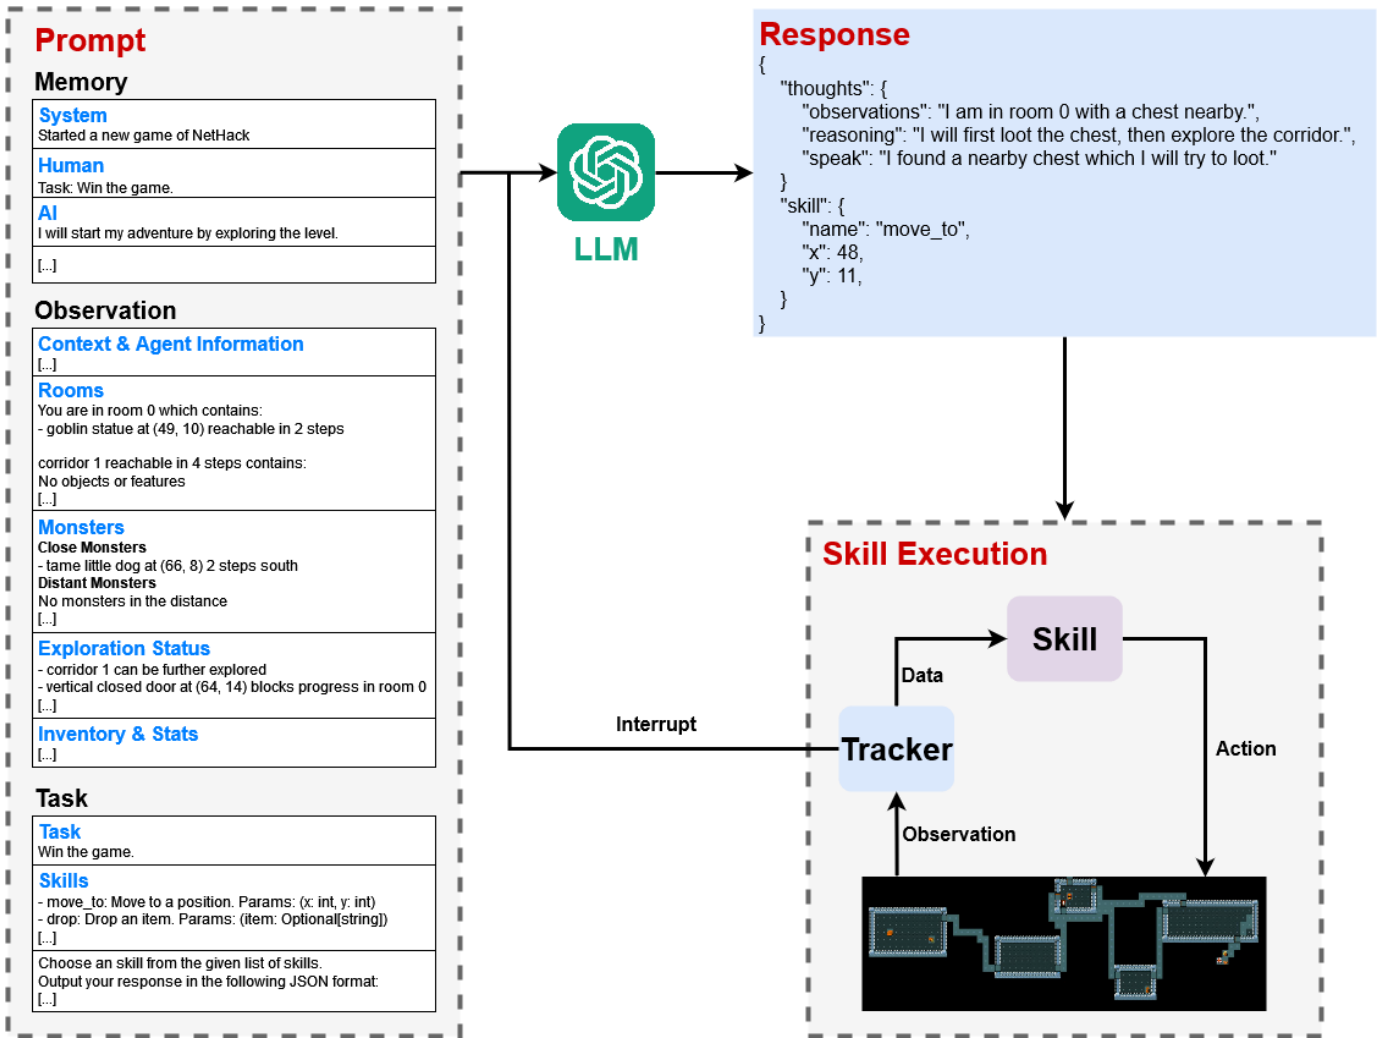


Fig. 2: Illustration of NetPlay playing NetHack. A prompt is constructed using messages containing the agent’s memory, a description of the current observation, and a task description with available skills and the desired output format. The LLM response is then parsed to retrieve and execute the chosen skill. While executing the selected skill, a tracker enriches the given observations and detects important events, such as when a new monster appears. When the skill is done or events interrupt the skill execution, the agent will restart the prompting process. An example prompt can be found on Github: https://github.com/CommanderCero/NetPlay/blob/main/example_prompt_plus_messages.txt

Navigation is automated through skills like “move_to x y” or “go_to room_id”. However, exploring levels with only these skills proved challenging for the LLM. To address this, we introduced the “explore_level” skill, which uses the exploration strategy from `autoascend`. This skill explores the current level by uncovering tiles, opening doors, and searching for hidden corridors. We removed the ability to kick open doors to avoid potential issues such as aggravating shopkeepers. Note that the agent can still decide to kick open doors using a separate “kick” skill. All movement-related skills will attack monsters that are in the way. The LLM can turn off this behavior using the “set_avoid_monster_flag” skill.

To indicate when the agent is done with a given task, it has access to the “finish_task” skill. Additionally, the LLM is equipped with the “press_key” and “type_text” skills for

navigating NetHack’s various game menus. While a menu is open, only the “finish_task” and text input skills remain available.

The remaining skills are thin wrappers around NetHack commands, such as `drink` or `pickup`. However, these commands often involve multiple steps, such as confirming which item to drink or first positioning the agent correctly to then pick up an item. Consequently, the LLM often assumed that the “drink” command accepts an item parameter or that “pickup” works seamlessly regardless of the agent’s current position. To mitigate these issues, we implemented four types of command skills. Base commands only invoke the command. Position commands offer the option to first move to the desired location. Inventory commands accept an item parameter to resolve the following popup menu. Finally, direction commands like “kick”

TABLE I: **Skill Examples:** Skills represent parametrizable behaviors that the LLM uses to play the game. The name, parameters, and descriptions help to understand what each skill does. For some skills, the LLM can omit optional parameters marked in [square brackets]. Note that the skill type is only used internally and does not matter for the final agent.

Type	Name	Parameters	Description
Special	explore_level		Explores the level to find new rooms, as well as hidden doors and corridors.
Special	set_avoid_monster_flag	value: bool	If set to true skills will try to avoid monsters.
Special	press_key	key: string	Presses the given letter. For special keys only ESC, SPACE, and ENTER are supported.
Position	pickup	[x: int, y: int]	Pickup things at your location or specify where you want to pickup an item.
Position	up	[x: int, y: int]	Go up a staircase at your location or specify the position of the staircase you want to use.
Inventory	drop	item_letter: string	Drop an item.
Inventory	wield	item_letter: string	Wield a weapon.
Direction	kick	x: int, y: int	Kick something.
Basic	cast		Opens your spellbook to cast a spell.
Basic	pay		Pay your shopping bill.

move the agent close to a desired position before executing the command in the correct direction.

C. Agent Loop

Upon receiving a new task, the agent prompts the LLM to select the first skill to execute. The LLM’s thoughts and the selected skill are stored in the agent’s memory as feedback. While executing the chosen skill, a data tracker observes and records details such as found structures, features hidden by monsters or items, which tiles the agent has already seen or searched, and events. The information collected by the data tracker is used by skills to make decisions.

The data tracker also looks for specific events in the game to provide additional feedback to the LLM. Events include new in-game messages, newly discovered structures, level changes or teleports, stat changes, low health, and the discovery of new monsters, items, and some map features such as fountains or altars.

A skill continues to run until completion or interruption. Skills are interrupted when specific events occur, such as changing the level, teleporting, discovering new objects, and reaching low health. In addition to events, many skills are interrupted when a menu shows up due to their inability to handle them. Regardless of why a skill stopped, the agent then prompts the LLM to select the next skill. The sole exception is when the “finish_task” skill is selected, or the game has ended, at which point the agent will stop until it receives a new task.

D. Handcrafted Agent

To assess the impact of the LLM in contrast to the predefined skills, we implemented a handcrafted agent that aims to replicate the behavior of `NetPlay` with the task set to “Win the Game”. The following list shows a breakdown of the agent’s decision-making process.

- 1) Abort any open menu, as we did not implement a way to navigate them.
- 2) If there are hostile monsters nearby, fight them.
- 3) If health is below 60%, try healing with potions or by praying.
- 4) Eat food from the inventory when hungry.
- 5) Pick up items, which in this case are potions and food.
- 6) If nothing to explore, move to the next level if possible.

- 7) If nothing else to do, explore the level and try kicking open doors.

All the conditions are evaluated in sequence. Once a condition is met, a corresponding skill is executed. The selected skill will be interrupted in the same way as `NetPlay`. Once a skill is interrupted, the agent will choose the next skill by again checking all conditions in order starting from the first. Note that although we aimed to imitate `NetPlay`’s behavior, the provided rules are too simplistic to capture all the nuances.

IV. EXPERIMENTS

Our goals for the experiments were two-fold. First, to evaluate the ability of `NetPlay` to play `NetHack`. Second, to provide an analysis of the agent’s strengths and weaknesses, focusing on identifying which aspects are influenced by the LLM.

A. Setup

All of our experiments used OpenAI’s GPT-4-Turbo (GPT-4-1106-PREVIEW) API as LLM with the temperature set to 0 and the response format set to JSON. Other models were not considered as initial tests revealed that models like GPT-3.5 and a 70B parameter instruct version of LLAMA 2 [19] could not correctly utilize our skills. The agent’s memory size was set to 500 tokens.

The agent had access to most commands that interact with the game directly, except for some rarely relevant commands, like turning undead or using a monster’s special ability. All control and system commands, like opening the help menu or hiding icons on the map, were excluded. We also implemented a time limit of 10 LLM calls, at which point the experiment would terminate if the in-game time did not advance.

B. Full Runs

We started evaluating `NetPlay` by letting it play `NetHack` without any constraints, tasking it to win the game. We will refer to this agent as the “unguided agent.” Although the task was to play the entire game, the agent occasionally confused its own objectives with the assigned task, resulting in the agent marking the task as done too early. To address this issue, we disabled the “finish_task” skill for this experiment.

TABLE II: **Results summary** of the mean and standard error for the agents achieved score, depth, experience level, and game time.

Metric	NetPlay (Unguided)	NetPlay (Guided)	autoascend	handcrafted
Score	284.85 ± 222.10	405.00 ± 216.38	11341.94 ± 11625.39	250.24 ± 159.17
Depth	2.60 ± 1.39	2.00 ± 1.05	4.01 ± 3.04	2.35 ± 0.93
Level	2.40 ± 1.23	3.30 ± 0.95	3.34 ± 7.69	2.39 ± 1.05
Time	1292.10 ± 942.74	2627.40 ± 1545.12	21169.81 ± 9155.59	1306 ± 924.17

Due to budget limitations, we evaluated all agents using only the Valkyrie role, as most agents performed best with this class during the NetHack 2021 challenge. We conducted 20 runs with the unguided agent. Additionally, we performed 100 runs each with `autoascend` and the handcrafted agent for comparison. After evaluating the unguided agent, we carried out an additional 10 runs employing a “guided agent” who was informed on how to play better. A detailed description of the guided agent will be provided below. For now, a summary of the results can be found in table II.

Table II shows that `autoascend` far outperforms both `NetPlay` and the handcrafted agent. While `NetPlay` managed to beat the handcrafted agent by a small margin, it is likely that with a few tweaks, the handcrafted agent can also outperform `NetPlay`.

The unguided agent primarily failed due to timeouts, followed by deaths caused by eating rotten corpses, fighting with low health, or being overwhelmed by enemies. Many timeouts were caused by the agent attempting to move past friendly monsters, such as a shopkeeper. By default, bumping into monsters attacks them, but for passive monsters, the game prompts the player before initiating an attack. The agent’s refusal to attack these monsters often leads to a loop of canceling the prompt and moving, resulting in eventual timeouts. A similar loop took place when the agent attempted to pick up an item with a generic name on the map but a detailed name in the game’s menu. This confusion led the agent to repeatedly close and reopen the menu, unable to locate the desired item.

Based on the results of the unguided agent, we constructed a guide that included strategies from `autoascend`, such as staying on the first two dungeon levels until reaching experience level 8, consuming only freshly slain corpses to avoid eating rotten ones, and leveraging altars to acquire items. Furthermore, we provided tips for common mistakes by the unguided agent, such as avoiding getting stuck behind passive monsters and informing the agent about the time limit to avoid timeouts.

The guided agent often managed to stay alive longer by consuming freshly killed corpses and praying when hungry or at low health. Its causes of death have been a mixture of timeouts, starvation, and dying in combat. Most of the timeouts stemmed from a bug with our tracker, which fails to detect when an object disappears while being obscured by a monster. For example, the agent repeatedly attempted to pick up a dagger already taken by its pet due to the tracker’s misleading observation. Despite receiving game messages indicating the absence of the item, the agent failed to recognize the situation accurately.

Because we tasked the guided agent to stay on the first two dungeon levels, its average depth is lower than that of the unguided agent. However, because monsters keep spawning over time, staying on the first levels is an excellent way to grind experience. This results in the guided agent gaining more experience than the unguided agent. Nevertheless, the agent’s tendency to stay on the first dungeon levels frequently caused it to die of starvation due to not finding enough monster corpses to eat. Note that `autoascend` had a similar starvation issue.

C. Scenarios

After conducting the full runs, we hypothesized that although `NetPlay` can be creative and interact with most mechanics in the game, it tends to fixate on the most straightforward approach for a given task. To confirm this hypothesis, we constructed various small-scale scenarios using `des-files` and a corresponding task description. Note that we excluded the handcrafted agent and `autoascend` for this experiment as they cannot easily alter their behavior.

The tested scenarios evaluated `NetPlay`’s ability to interact with game mechanics, follow instructions, and its creativity. We conducted five runs for each scenario, with all roles and the “finish_task” skill enabled. We also repeated some scenarios where the agent performed poorly with additional guidelines. We censored the word `NetHack` for the scenarios to evaluate the agent’s ability independently of its knowledge about the game. To avoid the agent never using the “finish_task” skill, we set a time limit of 500 timesteps for creative scenarios and 200 for the others. See table III for a summary of the tested scenarios and their results.

The tested scenarios show that `NetPlay` performs best when provided with concrete instructions. The `FOCUSED BOULDER` task and both `ESCAPE` tasks, in particular, highlight how the agent can act creative if we focus its attention on a specific problem. However, without very detailed instructions, the agent often fails to do what it wants due to incorrect actions and a lack of explicit feedback.

The agent’s struggle with explicit feedback is particularly evident in the `BAG` and `MULTIPICKUP` scenarios, where the agent often failed to navigate the menus correctly. While it understood the menus and often chose the correct course of action, it often failed by forgetting a crucial step, such as closing the menu.

V. POTENTIAL AND LIMITATIONS

`NetPlay` uses a similar architecture to `Inner Monologue` and `DEPS`, which have shown promising results for simple

TABLE III: **Scenarios:** A detailed description of all the tested scenarios, their results, and the agent’s success rate. Note that in some scenarios, the agent did not use the “finish_task” skill, even after completing it. We still count these as success.

Scenario	Success	Description	Results
Game Mechanics			
BAG	1/5	A room with four random objects and a bag of holding with the task of stuffing all objects into the bag.	The bag of holding menu is quite complex. The agent was only successful when using the option that automatically stuffs all items into the bag. In the other cases, the agent forgot to mark an item or to confirm its selection.
GUIDED BAG	3/5	Same as BAG, but we told the agent the quickest way to pick up items and to navigate the bag’s menu.	The agent used the automatic option three times. In the other cases, the agent marked the task done too early, stating that it would pick up the remaining items next.
MULTIPICKUP	3/5	A room with 2-5 objects on the same spot, challenging the agent to navigate the multipickup menu.	The agent often picked up items inefficiently by opening the pickup menu multiple times. It failed twice by forgetting to confirm its item selection.
WAND	1/5	A room with a statue and a wand with the task of hitting the statue with the wand.	The agent often failed by standing atop the statue and casting the wand onto itself. Only once did the wand spawn next to the statue, causing the agent to cast the wand towards the statue.
GUIDED WAND	5/5	Same task as WAND, but we asked the agent to stand next to the statue instead of on top of it and fire in the statue’s direction.	Most of the time, the agent succeeded on the first try, except once when he got it on the second try after repositioning himself.
Instructions			
ORDERED	5/5	A room with the task to pick up two wands, then a scroll of identification, and finally to identify one wand.	The agent executed the tasks accurately in the given order.
UNORDERED	3/5	A room with the task to drink from a fountain, open a locked and a closed door, and kill a monster in any order.	The agent completed the tasks in no particular order. One fail stemmed from high-level mobs spawning from the fountain, and one from incorrectly using the lockpick.
ALTERNATIVE	5/5	Three rooms with a fountain and a potion somewhere. The task was to drink from a fountain or a potion.	The agent always drank from the fountain, which in all cases was found first or was closest to the agent.
CONDITIONAL	4/5	Three rooms, with only a single potion hidden in one of the rooms. The task was to drink from a fountain, or if unavailable a potion.	The agent always drinks the first potion it finds without exploring further. In one case, it deemed the task impossible due to spawning with no fountain or potion in sight.
Creativity			
CARRY	1/5	The agent has to carry two very heavy objects through a monster-filled room. We also provided tools such as a bag of holding, a teleportation wand, and an invisibility cloak.	The agent often refused to play because it could not see the required items or it dropped them in the wrong room.
GUIDED CARRY	4/5	Same task as CARRY, but we told the agent to prioritize killing monsters first, to carry only one of the heavy items at a time, and to use the teleportation wand for easier travel.	Most of the time, the agent carried only one item, and it often used the wand to teleport. It failed once by dropping one item in the incorrect room.
BOULDER	1/5	Two rooms connected by a corridor with a boulder. The agent starts either with pickaxes or wands to remove the boulder.	When given only wands, the agent only used explore_level and ignored the boulder. Only once did it start with a pickaxe that it used to mine the boulder.
FOCUSED BOULDER	3/5	Same task as BOULDER, but the agent was told to remove any boulders blocking its path.	The agent often tried kicking the boulder, which failed, after which it then used a pickaxe or a wand. It failed twice due to not correctly utilizing the available tools.
GUIDED BOULDER	5/5	Same task as FOCUSED BOULDER, but the agent was told explicitly to remove the boulder with the wands or pickaxes. We also provided directions on how to utilize the tools.	In all cases, the agent quickly used the pickaxe or a wand to remove the boulder.
ESCAPE	3/5	The agent must escape from a stone-walled room. Escape methods: Digging with a wand through a wall, teleporting with a wand, or morphing into a wall-phasing monster using a polymorph control ring with a polymorph wand.	The agent escaped twice by teleporting, despite initial teleport failure. It also experimented with the wand of digging, casting it in all directions to find an exit. It failed twice due to incorrectly using the wands.
HINT ESCAPE	5/5	Same as ESCAPE, with a hint engraved on the floor. The hint either reveals which wall is brittle and leads to an escape or hints at the name of the wall-phasing monster.	After finding the hint, the agent often used the suggested escape method, except for one occasion when it teleported instead. In one instance, the initial attempts to dig through the wall failed, so it resorted to exploring other methods.

dynamic environments. Our experiments show that despite the complexity of NetHack, the agent can fulfill a wide range of tasks given enough context information. To our knowledge, this is the first NetHack agent to exhibit such flexible behavior. However, the benefits of the presented approach seem to diminish the more ambiguous a given task is.

A promising use case of the presented architecture is regression testing during game development. Game developers could test specific aspects of their game by providing NetPlay with detailed instructions on what to test. This approach could not only streamline the testing process, but it would also benefit from NetPlay’s flexibility, enabling the tests to adapt dynamically as the game evolves.

Given NetPlay’s proficiency when given detailed context information, an obvious extension to our approach would be granting the agent access to the NetHack Wikipedia. This could be done using a skill that accepts a query and adds the resulting information to the agent’s short-term memory. While we think this can improve the results at the cost of more LLM calls, finding the most relevant information for a given situation would be tricky. Instead, we recommend investing future research into automated methods for finding relevant context information, with a particular focus on finding the most successful past interactions as guidelines on how to play.

A significant limitation of our approach lies in our skills and observation description, which struggle to encompass NetHack’s complexity. Designing the agent to handle all potential edge cases proved challenging, as it is difficult to anticipate every scenario. While the premise of this approach is that the LLM can handle these edge cases, this only works if we have a comprehensive description of the environment and flexible skills. In practice, this would require an ever-growing repertoire of skills and an extremely long observation description.

VI. CONCLUSION

In this work, we introduce NetPlay, the first LLM-powered zero-shot agent for the challenging roguelike NetHack. Building upon an existing approach for simple dynamic environments, we extended its capabilities to address the complexities of NetHack. We evaluated the agent’s performance on the whole game and analyzed its behavior in various isolated scenarios.

NetPlay demonstrates proficiency in executing detailed instructions but struggles with more ambiguous tasks, such as winning the game. Notably, a simple rule-based agent can achieve comparable performance in playing the game. NetPlay’s strength lies in its flexibility and creativity. Our experiments show that, given enough context information, NetPlay can perform a wide range of tasks. Moreover, by focusing its attention on a particular problem, NetPlay is adept at exploring a wide range of potential solutions but often with limited success due to a lack of explicit feedback.

VII. ACKNOWLEDGEMENTS

This work was supported by the EPSRC Centre for Doctoral Training in Intelligent Games & Games Intelligence (IGGI) (EP/S022325/1).

This work was supported by Creative Assembly.

REFERENCES

- [1] H. Naveed and et al., “A comprehensive overview of large language models,” 2023.
- [2] W. Huang and et al., “Inner monologue: Embodied reasoning through planning with language models,” in *Proceedings of The 6th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, vol. 205. PMLR, 14–18 Dec 2023, pp. 1769–1782. [Online]. Available: <https://proceedings.mlr.press/v205/huang23c.html>
- [3] C. H. Song and et al., “Llm-planner: Few-shot grounded planning for embodied agents with large language models,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023.
- [4] Z. Wang and et al., “Jarvis-1: Open-world multi-task agents with memory-augmented multimodal language models,” 2023.
- [5] X. Zhu and et al., “Ghost in the minecraft: Generally capable agents for open-world environments via large language models with text-based knowledge and memory,” 2023.
- [6] G. Wang, Y. Xie, Y. Jiang, A. Mandlekar, C. Xiao, Y. Zhu, L. Fan, and A. Anandkumar, “Voyager: An open-ended embodied agent with large language models,” *Transactions on Machine Learning Research*, 2024. [Online]. Available: <https://openreview.net/forum?id=ehfRiF0R3a>
- [7] K. Lorber, “Nethack home page.” [Online]. Available: <https://nethack.org/>
- [8] “autoascend,” GitHub, 10 2023. [Online]. Available: <https://github.com/maciej-sypetkowski/autoascend>
- [9] Hambro and et al., “Insights from the neurips 2021 nethack challenge,” in *NeurIPS 2021 Competitions and Demonstrations Track*, ser. Proceedings of Machine Learning Research, vol. 176. PMLR, 06–14 Dec 2022, pp. 41–52. [Online]. Available: <https://proceedings.mlr.press/v176/hambro22a.html>
- [10] H. Küttler, N. Nardelli, A. Miller, R. Raileanu, M. Selvatici, E. Grefenstette, and T. Rocktäschel, “The nethack learning environment,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 7671–7684. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/569ff987c643b4bedf504efda8f786c2-Paper.pdf
- [11] M. Samvelyan and et al., “Minihack the planet: A sandbox for open-ended reinforcement learning research,” in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. [Online]. Available: <https://openreview.net/forum?id=skFwlyefkWJ>
- [12] maciej sypetkowski, “Autoascend – 1st place nethack agent for the nethack challenge at neurips 2021,” GitHub, 01 2024. [Online]. Available: <https://github.com/maciej-sypetkowski/autoascend>
- [13] Z. Wang, S. Cai, G. Chen, A. Liu, X. Ma, and Y. Liang, “Describe, explain, plan and select: Interactive planning with LLMs enables open-world multi-task agents,” in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Online]. Available: <https://openreview.net/forum?id=KtvPdGb31Z>
- [14] J. Liang and et al., “Code as policies: Language model programs for embodied control,” in *arXiv preprint arXiv:2209.07753*, 2022.
- [15] Y. J. Ma and et al., “Eureka: Human-level reward design via coding large language models,” in *The Twelfth International Conference on Learning Representations*, 2024. [Online]. Available: <https://openreview.net/forum?id=IEduRUO55F>
- [16] M. Klissarov and et al., “Motif: Intrinsic motivation from artificial intelligence feedback,” in *The Twelfth International Conference on Learning Representations*, 2024. [Online]. Available: <https://openreview.net/forum?id=tmBKIEcDE9>
- [17] M. Kwon, S. M. Xie, K. Bullard, and D. Sadigh, “Reward design with language models,” in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=10uNUgI5KI>
- [18] J. Wei, X. Wang, D. Schuurmans, M. Bosma, brian ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou, “Chain of thought prompting elicits reasoning in large language models,” in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022. [Online]. Available: https://openreview.net/forum?id=_YjQlMeSB_J
- [19] H. Touvron and et al., “Llama 2: Open foundation and fine-tuned chat models,” 2023.