What's the Plan? Evaluating and Developing Planning-Aware Techniques for LLMs

Anonymous ACL submission

Abstract

Planning is a fundamental task in artificial in-002 telligence that involves finding a sequence of actions that achieve a specified goal in a given environment. Large language models (LLMs) are increasingly used for applications that require planning capabilities, such as web or embodied agents. In line with recent studies, 007 we demonstrate through experimentation that LLMs lack necessary skills required for planning. Based on these observations, we advocate for the potential of a hybrid approach that combines LLMs with classical planning methodology. Then, we introduce SimPlan, a novel 013 hybrid-method, and evaluate its performance in 014 a new challenging setup. Our extensive experiments across various planning domains demonstrate that SimPlan significantly outperforms 017 existing LLM-based planners.

1 Introduction

022

Planning is a crucial aspect of artificial intelligence, which involves finding a sequence of actions that can achieve a specific goal in a given environment. Traditionally, planning problems are solved with graph-search heuristics which are based on solving relaxed problems (Blum and Furst, 1997; Hoffmann, 2001) or creating sub-goals by backtracking from the desired end-goals (Hoffmann et al., 2004; Helmert and Domshlak, 2009).

With the success of LLMs in various natural language processing tasks, there has been an increasing interest in utilizing them for applications that involve planning and reasoning, such as web agents (Yang et al., 2023; Yin et al., 2023; Huang et al., 2022), embodied agents (Huang et al., 2022; Song et al., 2023; Ahn et al., 2022) or open-world games (Wu et al., 2023; Wang et al., 2023). A compelling advantage of using Large Language Models (LLMs) for planning tasks lies in their ability to handle problems in natural language. This elimi-



Figure 1: In the Ferry planning domain, an example of a problem instance involves an initial state comprising the location of a ferry and several cars, with specified goals for placing the cars in specific locations (top). The ferry is capable of boarding a car and transporting it between locations. The planning task entails generating a sequence of actions (i.e., a *plan*) such that executing them leads to reaching a goal state (bottom).

nates the need for formalizing them in strict languages like the Planning Domain Definition Language (PDDL; Malik Ghallab et al., 1998). However, LLMs have been found to lack the ability to solve classic planning tasks (Liu et al., 2023; Valmeekam et al., 2023a,b), and the magnitude of their failure in such tasks highlights the need for alternative approaches that can enhance their planning capabilities.

In our work, we start by examining whether LLMs possess capabilities that we deem as nec-

040

041

048

essary for planning. We expand on the analyses of LLMs as planners done by Valmeekam et al. (2023b) and Silver et al. (2022), by focusing on specific capabilities that test their ability to comprehend diverse domains. Our analysis reveals limitations in the ability of LLMs to effectively model past actions and their subsequent impact on the current state. Additionally, the models struggle to identify the full spectrum of applicable actions and lack the capacity to prioritize sub-goals effectively, considering both necessary and optimal ordering. These findings highlight potential shortcomings in the planning and decision-making mechanisms employed by these models.

051

057

061

062

063

064

074

077

084

087

880

094

097

Motivated by our findings, we advocate for a hybrid methodology that integrates traditional planning algorithms with LLMs. We propose a similarity-based planner, SimPlan, a method which combines an action-ranking model with a greedy best-first search algorithm. This approach, as discussed in Section 4, addresses the shortcomings of LLMs described above and enhance the exploration of the state space. We evaluate our approach across various planning domains, comparing its efficacy with that of existing LLM-based planning strategies. The experiments reveal that our hybrid model surpasses traditional methods, underscoring the viability of combining LLMs with conventional planning techniques. We can summarize the contributions of our paper as follows:

- We evaluate LLMs shortcomings in reasoning about planning problems, underlying key missing capabilities.
- We introduce a new hybrid LLM-based planner that integrates classic planning algorithms and tools to overcome these limitations, $SimPlan^1$.
- We propose a novel generalized planning setup and demonstrate significant performance gains over prior work. However, the setup presents ongoing challenges.

The following sections are organized as follows. In Section 2, we provide a necessary overview of classic planning. Section 3 assesses the planning and reasoning capabilities of LLMs, evaluating their performance on various tasks. In Section 4, we introduce our novel approach, which leverages the strengths of traditional planning techniques to improve the efficiency of planning. Finally, in Section 5 we conduct experiments on several diverse planning domains, demonstrating the benefits of our hybrid approach.

2 Classical Planning

The core objective of a planning task is to construct a sequence of actions (i.e., a plan) that transitions from the initial state to a desired goal state. In classical planning, this process relies on a formal representation of the planning domain and the problem instance, encompassing the state space, actions, their preconditions and effects, and the desired goals. A widely used representation is the Planning Domain Definition Language (PDDL; Malik Ghallab et al., 1998). See Appendix A for more details about PDDL.

We now move to describe a problem instance within the Ferry planning domain as illustrated in Figure 1. This instance features a ferry and two cars, car1 and car2, initially positioned at separate docks, loc1 and loc2 respectively. The goal state is that car1 will be at loc2 and car2 will be at loc1. Each state is defined by the locations of the ferry and the cars (e.g., ferry is at loc1, car1 is at loc1, car2 is at loc2). Actions possess preconditions, outlining the state requirements for their execution, and effects, altering the state upon execution. For instance, the "board car1 loc1" action requires the ferry and car1 to be at loc1, and its execution results in car1 being on the ferry. The visualized plan involves six actions. Initially, "board car1 loc1" loads car1 onto the ferry. Then, "sail loc1 loc2" navigates the ferry to loc2. Subsequently, "debark car1" unloads car1 at loc2. Next, "board car2 loc2" loads car2, followed by "sail loc2 loc1" returning to loc1. Finally, "debark car2 loc1" unloads car2, achieving the goal.

Traditionally, classic planning problems have been solved using informed graph-search algorithms, such as greedy-best first search and the seminal A* algorithm (Doran and Michie, 1966; Richter and Westphal, 2010). To guide their search towards the goal state, these algorithms leverage several heuristics. Two prominent examples include the FF heuristic, known for its efficiency and scalability (Blum and Furst, 1997; Hoffmann, 2001), and the landmarks heuristic, which focuses on identifying critical states on the path to the goal (Hoffmann et al., 2004; Helmert and Domshlak,

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145

146

147

148

100

101

¹Dataset and code will be released upon publication.

MODEL	GRIPPERS	DEPOTS	FERRY	MINIGRID	BLOCKS
FALCON-180B	.11	.07	.45	.00	.24
MISTRAL-7B	.00	.07	.31	.39	.16
MIXTRAL-8X7B	.04	.14	.22	.35	.26
LLAMA-2-70B	.60	.44	.50	.26	.52
CODELLAMA-34B	.01	.04	.43	.28	.14
GPT-4 Turbo	.42	.26	.35	.59	.31

Table 1: Success rates for describing the current state given an initial state and a sequence of executed actions.

MODEL	GRIPPERS	DEPOTS	FERRY	MINIGRID	BLOCKS
FALCON-180B	.63	.66	.68	.32	.57
MISTRAL-7B	.30	.24	.41	.26	.18
MIXTRAL-8X7B	.48	.51	.62	.41	.72
LLAMA-2-70B	.60	.44	.50	.26	.52
CODELLAMA-34B	.64	.51	.61	.67	.43
GPT-4 TURBO	.87	.84	.88	.86	.91

 Table 2: Success rates for predicting the applicable actions given a state.

2009).

Building on this, our work leverages generalized planning, a subfield of classical planning, as it offers a more apt framework for evaluating LLM capabilities. Within this paradigm, we learn a heuristic from training examples and subsequently assess its performance on a distinct set of instances, following the methodologies established by Yang et al. (2022); Silver et al. (2023). As detailed in Section 5, our evaluation hinges on a specific variant of generalized planning where the test set comprises problem instances exhibiting greater complexity compared to those in the training set.

3 The Shortcomings of LLMs as Planners

Several recent studies (Valmeekam et al., 2023a,b; Stein and Koller, 2023) have demonstrated limitations in the planning capabilities of LLMs. This section delves into these shortcomings, focusing on a targeted analysis of specific cognitive skills considered essential for successful planning. We start our evaluation by assessing the model's capacity to grasp the operating principles governing the specific domain under consideration. These principles can be reasoned from the preconditions and effects associated with various actions. Next, we test the ability of LLMs to prioritize interim states required for achieving the goal states.

We evaluate the performance of various LLMs
on these skills using an In-Context Learning approach (ICL). In particular, we test the following
models: FALCON-180B (Almazrouei et al., 2023),
LLAMA-2-70B (Touvron et al., 2023), CODELLAMA-34B (Rozière et al., 2024), MISTRAL-7B



Figure 2: The maximum success rate across the different LLMs in inferring the new state, as a function of the number of actions taken.

(Jiang et al., 2023), MIXTRAL-8x7B (Jiang et al., 2024), and GPT-4 TURBO (OpenAI, 2023). The prompts for all the following experiments are provided in Appendix B. The examples for these experiments were randomly generated from diverse, well-known planning domains. Further details regarding the specific domains employed and implementation considerations are provided in Appendix C.

182

183

184

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

202

203

204

205

206

207

208

209

210

211

212

213

214

3.1 Observation

The actions' effects. The model is given an initial state and a sequence of actions, and is asked to infer the new state after executing the sequence. This experiment was also conducted recently by Valmeekam et al. (2023b), however it was limited to a single domain and to GPT models.

The results of this experiment are demonstrated in Table 1. Overall we can observe that LLMs perform poorly in understanding how a sequence of actions affects the environment. Moreover, the results show that the models' performance varies significantly depending on the complexity of the environment. For example, in the Minigrid environment, the Mistral-7B model performs poorly, with an accuracy of 0.14. In contrast, the Falcon-180B model performs much better, with an accuracy of 0.21. Moreover, we can observe that GPT-4-Turbo achieves the highest overall performance across all scenarios except Ferry, with percentages ranging from 30% to 58%. Lastly, Figure 2 indicates that LLMs face increasing difficulty in accurately inferring the task state as the number of actions taken increases.

Applicable actions.In this experiment, we eval-215uate the model's understanding of the actions' pre-216

169

170

172

173

174

175

149

150

conditions. This is done by providing the model with a state and a set of four possible actions and asking the model which action is applicable.

Table 2 presents the results of the applicable actions experiment. Although this task does not require complex reasoning abilities, the only model which exhibited consistently high accuracy across all evaluated environments is GPT-4 Turbo.

3.2 Goal Ordering

217

218

219

221

232

236

241

242

244

245

246

248

249

250

257

259

260

261

262

265

Reasonable and necessary ordering. Within specific problem-solving domains, like Blocksworld and Minigrid, the sequential achievement of goals becomes critical due to the inherent dependencies between them. This necessitates the completion of certain goals prior to attempting others. For example, consider the Blocksworld domain where blocks are stacked on top of one another and at each step we can only remove the top block from each stack. Assume we are given the following goals: should be stacked on \bigcirc and \bigcirc should be stacked on \bigcirc . Within this scenario, prioritizing the completion of the latter goal would constitute a reasonable order. Otherwise, we would have to remove the blocks on top of **b** before we can pick it up, and this would necessitate the destruction of the previously achieved goal.

Consider now the Minigrid domain, a domain where the agent has to pick-up keys and unlock doors in order to reach a certain location in a 2D map. We can observe that it is *necessary* to unlock some doors before others, since there are inner doors we can't reach until we unlock the outer doors that are blocking the way. Accordingly, the ability of the model to prioritize goals in a reasonable and necessary fashion is crucial for planning, and has long been the subject of research in classical planning (Hoffmann and Koehler, 2000; Hoffmann et al., 2004).

Therefore, in this set of experiments, we present the model with two goals and ask it to choose which of the goals should be completed first. We choose to test this for the Blocksworld and Minigrid domains, since these domains require a reasonable or necessary order between the goals.

Optimal ordering. In optimal planning, the current state imposes an ordering between the goals in order to avoid redundant actions. For example, in Figure 1, since the ferry is already at location

	REASON.	NECES.	Opt	IMAL
MODEL	BLOCKS	MINIGRID	Ferry	GRIPPERS
FALCON-180B	.71	.69	.65	.65
MISTRAL-7B	.58	.85	.63	.59
MIXTRAL-8X7B	.04	.62	.64	.59
Meta	.21	.69	.63	.59
CODELLAMA-34B	.13	.85	.67	.70
GPT-4 Turbo	.75	.78	.63	.59

Table 3: Success rate results for the goal planning experiments, testing for the model's ability to prioritize goals when imposed with reasonable, necessary or optimal orderings.

1, it should board car 1 so it can debark it at location 2, before moving to location 2. For this experiment we introduce a scenario where the model is provided with an initial state and a set of goals. The model then has to reason about which goal can be achieved with a lower number of actions. We use the Ferry and Grippers domains, as these domains do not have necessary or reasonable ordering between their goals. As such, the main consideration is the optimality of the plan. 266

267

269

270

271

272

273

274

275

276

277

278

279

281

282

284

285

286

287

289

290

291

292

293

294

295

296

297

300

301

The results for the goal planning experiments are presented in Table 3. Overall, Codellama-34B exhibits superior performance over the other models in 3 out-of 4 domains. However, in terms of adhering to optimal goal orderings, the models generally show similar performance, with scores ranging from 0.63 to 0.67. These findings suggest that models struggle with goal reasoning.

4 SimPlan

In this section we present *SimPlan*, an actionranking model combined with a greedy best-first search (GBFS) algorithm designed to address the limitations observed in Section 3. First, We describe the data generation and augmentation process (Section 4.1). Next, we describe the architecture and the training process (Section 4.2). Finally, we discuss inference-time decoding (Section 4.3).

4.1 Data Generation and Augmentation

Planning problems formulation has a symbolic nature, which significantly simplifies training data generation². However, since this formulation uses unique identifiers for objects (e.g., blocks in Blockworld denoted as b1-b5), there is a potential bias risk. While seemingly arbitrary, these identifiers can inadvertently lead LLMs to develop preferences based solely on their frequency in the train-

²Can be generated using libraries such as PDDL generators Seipp et al., 2022.



(a) The actions ranking model.

(b) The planning process.

334

335

336

Figure 3: Our proposed *SimPlan* architecture. Colors green, yellow and orange are used to denote states, goals and actions, respectively. (a) At each iteration, a bi-encoder is used to generate contextualized token-level representations for the concatenated current state and goals, as well as for each applicable action. The actions' representations can be extracted once in an offline process. Then, the set of applicable actions are ranked based on their similarity with the state and goals representation, using the late-interaction architecture of ColBERT (Khattab and Zaharia, 2020) (b) At inference time, the scores are inserted into a prioritized queue for choosing the next state.

ing data, To mitigate this issue, we augment the training dataset with 100 permutations per instance. These permutations involve replacing all objects indices with new sampled indices, essentially "deidentifying" the objects. This augmentation promotes LLM impartiality towards object indexing and enhances their ability to generalize to unseen indices during testing, resulting in a more robust and versatile model.

4.2 Learning to Rank Actions

302

310

311

312

313

315

317

321

This section introduces our action-ranking model, a bi-encoder model tasked with ranking applicable actions based on the current state and desired goals, as depicted in Figure 3a. Recognizing the inherent challenges faced by LLMs in generating applicable actions or even inferring what the applicable actions are, as observed in Section 3.1, we formulate the task of selecting the next action given a state and a goal as a retrieval problem. This retrieval-based approach ensures that the model only considers applicable actions.

Leveraging the late interaction architecture in-323 troduced by Khattab and Zaharia (2020), we propose an analogous schema wherein the goal and 325 state serve as the query, while the action acts as the context. As illustrated in Figure 3a (top), this 327 schema demonstrably enhanced predictive power 329 by computing cosine similarity between individual tokens in the query and context, rather than relying on pooled representations. For each query token the maximum similarity score is selected, and the overall similarity score is calculated as the sum of 333

these individual maxima. This similarity score can be regarded as the confidence of the model in the current path.

To optimize the model, we employ cross-entropy 337 loss, comparing the top ranked action with the gold 338 next action. To prevent the collapse of action repre-339 sentations where all actions become indistinguish-340 able from each other, ranking actions based on their 341 similarity requires the inclusion of negative exam-342 ples during training. We employ two methods for 343 incorporating negative examples to enhance model performance. Firstly, we leverage the in-batch neg-345 ative sampling approach described by Henderson 346 et al. (2017). Within each training batch containing 347 n action-state pairs, every other pair serves as a negative example for every original pair. Secondly, 349 we create "hard negative" examples using three 350 distinct techniques: (1) action name replacement, 351 where actions are swapped (e.g., changing the ac-352 tion "debark car1 loc1" to the opposite action 353 "board car1 loc1"); (2) subterm swapping, where 354 the order of subterms is changed (e.g., changing 355 "sail loc1 loc2" to "sail loc2 loc1"); and (3) random subterm sampling, where individual 357 subterms are replaced with random values (e.g., 358 changing "board car2 loc2" to "board car1 359 loc1"). These techniques ensure the generated 360 negatives are syntactically similar to the positive 361 examples, forcing the model to learn finer distinctions between actions. 363

4.3 Planning

364

367

370

371

373

374

375

410

411

412

413

Our proposed planning approach addresses two key challenges described earlier: (1) the negative correlation between the number of executed actions and the accuracy of describing the current state; and (2) the limitations of beam search used by LLMs for planning tasks.

State Updates. In Section 3.1 we demonstrated a negative correlation between the number of actions taken and the accuracy of a large language model (LLM) in describing the current state. To circumvent this issue and align with the methodology employed in Hao et al. (2023), we adopt a state-update strategy that calculates the new state directly from the current state and the chosen action, bypassing the LLM's potentially inaccurate inference.

Greedy Best-First Search Algorithm. The com-381 mon decoding algorithm used for natural language generation is the beam search algorithm, which is a local search algorithm that keeps only a fixed number k of promising paths as candidates. While effective for optimization problems, the beam search 386 algorithm has several limitations related to planning problems. First, a known issue associated with beam search is that it "can suffer from a lack of diversity among the k paths-they can quickly be-390 come concentrated in a small region of the search space" (Russell and Norvig, 2009). Second, unless they are implemented with backtracking, local search algorithms lack the ability to continue from an earlier generated path and explore taking alternative actions. Drawing inspiration from well-known classic planners (e.g., the LAMA planner (Richter and Westphal, 2010)), we employ a graph-based algorithm to address this limitation. GBFS chooses the next node to expand based on a cost function. 400 Similar to the implementation of beam search for 401 language generation, our cost function is imple-402 mented as an aggregated score of all the probabili-403 ties extracted for actions participating in the path. 404 However, to avoid penalizing long sequences, we 405 replace the sum log probability with an average log 406 probability. Overall, we suggest that GBFS will fa-407 cilitate greater exploration of high-potential paths, 408 making it better suited for goal-directed planning. 409

We now describe the planning process of our approach, as depicted in Figure 3(b). We implement the GBFS algorithm using a priority queue, which is intended to manage all explored states and priori-



Figure 4: An example of our experimental setting for the Blocksworld domain. The train and development sets include problem instances from a simple configuration (top), and the test set include problem instances from a complex configuration (bottom), where the difference between the two configurations is the number of blocks.

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

430

431

432

433

434

435

436

437

438

439

440

441

442

tize their subsequent expansion. We limit the queue size to 1000 to avoid incurring memory issues. We start by generating and encoding all applicable actions using the trained bi-encoder and store their representations in memory. Afterwards, the initial state is added to the priority queue. At each step, the state with the highest priority, alongside the goals predicates, is encoded using the trained biencoder to extract its latent tokens representation. Similarity scores are extracted by comparing the state's latent tokens representation with all applicable actions, as described in Section 4.2. Using the acquired scores, a heuristic value is calculated for each explored path. This iterative process continues until the desired goal state is reached.

5 Experiments

To assess the effectiveness of our proposed method, we conduct experiments testing the capability of the SimPlan algorithm to generalize from simplified tasks to more complex scenarios, as demonstrated in Figure 4. The domains used in the experiments were the ones covered in Section 3.

5.1 Datasets

To systematically evaluate planning model performance across varying difficulties, we generated two problem configurations for each domain: simple and complex. The configurations vary in their number of objects. For example, in Blocksworld, simple configurations consists of 2-

	Blo	cks	Fe	rry	Grip	pers	Dep	oots	Min	igrid
Method	S	С	S	Ċ	S	C	S	С	S	C
LLM4PDDL GPT-4 Turbo (no validation)	.26	.0	.20	.0	.60	.0	.30	.0	.13	.08
LLM4PDDL CodeLlama-34b-instruct	.07	.0	.25	.0	.33	.0	.03	.0	.07	.0
Plansformer CodeLlama-7b-instruct	.96	.0	.80	.0	.90	.0	.53	.0	.93	.88
Random	.33	.0	.03	.0	.20	.0	.13	.0	.16	.22
<pre># Goals Completed Heuristic</pre>	.25	.0	.63	1.0	.76	.96	.43	.0	.16	.22
SimPlan (ours)	1.0	.56	1.0	1.0	1.0	1.0	1.0	.12	1.0	.86

Table 4: The fraction of test problems solved. We report the results both on a test set of simple problem instances (S), which have the same configuration as the problems used in the training set, and complex problem instances (C) with an increased number of objects. The bolded results are significant (paired student's t-test, p < 0.05).

5 blocks, while complex configurations featured 11-25 blocks. Each instance includes a plan generated by the LAMA planner (Richter and Westphal, 2010) for training purposes.

This manipulation impacts plan length in two key ways: (1) More objects introduce more goals, and (2) new objects can act as obstacles, requiring further actions for removal. For example, in Blocksworld (Figure 4), stacking a new atop 2 require two additional actions to remove it before unstacking 2 from 3. In addition, a new goal can be added, such as placing a on top of 3, which would require two additional steps to pick it up and stack it, totaling in increasing the plan length by four actions. Consequently, simple problems had an average plan length of 13.5 actions, while complex problems had a significantly higher average of 357.2 actions. More details about the datasets can be found in Appendix C.

5.2 Baselines

The baselines used in our experiments can be categorized as follows:

- In-context Learning: The first baseline utilizes vanilla GPT-4 Turbo and directly instructs it to generate a plan. The second baseline incorporates CodeLlama-34b-instruct with a soft-validation strategy³ to address poorly formatted LLM outputs. Both adhere to the twoexample few-shot prompt design established by LLM4PDDL (Silver et al., 2022).
- 2. Fine-tuning: Code-llama-7b-instruct model, trained for code instruction tasks, was

fine-tuned using the approach proposed by Plansformer (Pallagani et al., 2022).

475

476

477

478

479

480

481

482

483

484

485

486

487

488

489

490

491

492

493

494

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

Naive Baselines: We added two naive approaches: (1) Random, samples a subsequent action from a list of applicable actions; and (2) # Goals Completed Heuristic, a GBFS with a heuristic considering the number of fulfilled goals.

Experimental Setup. All models are trained using instances from the simple configuration and are then tested separately on unseen instances from the simple and complex configurations. In order to produce comparable results we set a fixed number of next action predictions for each problem instance based on the number of states that a classic planner expanded, further elaborated in Appendix D.2. Following Valmeekam et al. (2023b), we measured the model's ratio of solved problems instances.

5.3 Results

The performance of the different baselines and SimPlan is presented in Table 4. First we can observe that SimPlan surpasses all other models across all problems and configurations except for comparable results on the complex configuration in Minigrid. Moreover, we can observe that all models struggled in the challenging Depots domain. This domain necessitates object stacking (akin to Blocksworld) while introducing the requirement of object relocation (similar to Ferry). This combination leads to actions with many subterms (e.g., unload hoist3 crate14 truck0 distributor1) and difficult reasoning challenges.

As anticipated and aligned with the expectations established in Section 3, LLM-based baseline planners struggle in tackling complex problem instances. This finding diverges from the re-

443

444

445

- 459
- 461

462 463

464

465 466

467

468 469

470

471 472

473

³This strategy leverages a planning tool to identify all applicable actions at each step and replaces invalid actions with the most similar valid alternative based on cosine similarity.

Method	Blocks	Ferry	Minigrid
SimPlan	.56	1.0	.86
w/o Hard Negatives	.36	1.0	.78
w/o Data Augmentation	.0	.0	.56
w/o Updating State	.0	.0	.64

Table 5: Ablation results, reporting the fraction of test problems solved from the complex configuration.

sults reported by LLM4PDDL (Silver et al., 2022), which demonstrated strong generalization capabilities within the Grippers domain. We believe that this discrepancy is attributed to the increased complexity of our Grippers setup, featuring five rooms compared to the two employed in LLM4PDDL (see Appendix C).

The Random model's similar performance across simple and complex Minigrid tasks suggests minimal complexity increase, suggesting a possible explanation to Plansformer's performance in this domain. The # Goals Completed baseline exhibits strong performance in Ferry and Grippers domains, highlighting the potential benefits of a simple goal counting heuristic in specific domains, particularly in the domains where the goals are independent. Conversely, the Blocksworld domain favors the Random baseline over the goal-oriented approach, underscoring the critical role of goal ordering in effective planning.

5.4 Ablations

511

512

513

514

515

516

517

519

520

521

522

524

525

526

529

532

533

534

535

537

538

540

541

545

546

We conducted an ablation study analyzing the contributions of various components in our *SimPlan*.
We evaluated the contributions of hard negatives, data augmentation and the state updates component.
See Appendix E for implementation details.

The results are presented in Table 5. We can observe that hard negatives provided a significant improvement in the results. Data augmentation, however, was crucial for generalization, as its removal hindered handling previously unseen objects indices. Finally, eliminating state updates led to poor performance, highlighting the difficulty for LLMs to manage state solely through action sequences.

6 Related work

547 Valmeekam et al. (2023b) presented a benchmark
548 for evaluating LLM-based planners and concluded
549 that such models display subpar performance.
550 Follow-up works suggested enriching the planning
551 process with multiple reasoning steps inspired by

Chain-of-Thought and Tree of Thoughts (Wei et al., 2022; Yao et al., 2023). Hu et al. (2023); Stein and Koller (2023) prompt the LLM to reason how the goals can be achieved before generating a plan. Hao et al. (2023) prompt the LLM to describe the updated state between each executed action. Valmeekam et al. (2023c) proposed that models will self-critique their own generated plan and fix them, similar to (Madaan et al., 2023). However, many of these works concluded that models tend to struggle with such reasoning tasks. Our analysis in Section 3 is motivated by this line of research, underlying missing reasoning capabilities in LLM-based planners.

Recognizing LLMs' limitations, follow up work explored hybrid approaches of using LLMs along with classic planning tools. Valmeekam et al. (2023c) tested the use of a planning verification tool instead of LLM to provide feedback for generated plans. Silver et al. (2022) also used a planning verification tool to find and fix inapplicable actions, as described in Section 5. Pallagani et al. (2022) used a classic planner to generate a large training dataset of solved problems in order to fine-tune code models for planning. Our work follows this hybrid approach, assuming access to planning tools (Section 4). However, we diverge in the experimental setup (Section 5) by focusing on generalized planning. In other words, we examine how models cope with significantly larger problems than seen during fine-tuning, a novel aspect compared to prior works.

7 Conclusions

Despite the encouraging outcomes achieved by our approach, a notable gap exists between the efficency of traditional planners and LLM-based planners. We aim to narrow this gap by exploring new hybrid techniques, particularly focusing on areas such as goal ordering where the LLM exhibited poor performance in our evaluations. In future work, we intend to broaden the scope of our methodology to address real-world planning challenges that deviate from the structured PDDL format, including scenarios relevant to web and embodied agents. Through these efforts, we aim not only to refine the efficacy of our methodology but also to extend its utility across a more diverse array of planning tasks. To facilitate these advancements, we advocate for the employment of our dataset as a robust platform for testing planning algorithms.

582

583

584

585

586

587

588

589

590

591

592

593

594

595

596

597

598

599

600

601

552

553

606

611

612

613

614

617

618

619

622

624

625

630

631

632

633

634

641

642

648

652

8 Limitations

Classic planning problems have been developed for over two decades and include many domains and variants of planning, such as numeric planning and optimal planning. Our work focuses on 5 domains which we found interesting, such as a navigation task, and a domain which combines the challenges from two other domains in our set (Depots). However, this work is not an exhaustive evaluation of LLM-based planners. Regarding SimPlan, as described in the paper, we found that it degrades in performance as we add more objects in Blocksworld, which indicates a weakness when adding more blocks with dependencies amongst them. In addition, our approach is limited by the max sequence length of the LLM, which would make it difficult to handle significantly larger problems.

9 Ethics

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil J. Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Mengyuan Yan, and Andy Zeng. 2022. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Étienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. The falcon series of open language models.
 - Avrim L Blum and Merrick L Furst. 1997. Fast Planning Through Planning Graph Analysis.
 - James E Doran and Donald Michie. 1966. Experiments with the graph traverser program. *Proceedings of the*

Royal Society of London. Series A. Mathematical and Physical Sciences, 294(1437):235–259. 653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

- Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. 2023. Reasoning with Language Model is Planning with World Model.
- Malte Helmert and Carmel Domshlak. 2009. Landmarks, Critical Paths and Abstractions: What's the Difference Anyway? Proceedings of the International Conference on Automated Planning and Scheduling, 19:162–169.
- Matthew Henderson, Rami Al-Rfou, Brian Strope, Yunhsuan Sung, Laszlo Lukacs, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. 2017. Efficient Natural Language Response Suggestion for Smart Reply.
- J. Hoffmann and J. Koehler. 2000. On Reasonable and Forced Goal Orderings and their Use in an Agenda-Driven Planning Algorithm. *Journal of Artificial Intelligence Research*, 12:338–386.
- J. Hoffmann, J. Porteous, and L. Sebastia. 2004. Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research*, 22:215–278.
- Joerg Hoffmann. 2001. FF: The Fast-Forward Planning System. *AI Magazine*, 22(3):57–57.
- Hanxu Hu, Hongyuan Lu, Huajian Zhang, Yun-Ze Song, Wai Lam, and Yue Zhang. 2023. Chain-of-symbol prompting elicits planning in large langauge models.
- Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. 2022. Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, Lélio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2024. Mixtral of experts.
- Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert.

813

814

815

760

Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone. 2023. LLM+P: Empowering Large Language Models with Optimal Planning Proficiency.

706

707

710

711

712

713

714

716

719

723

729

732

733

734

736

737

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754 755

756

- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. Self-refine: Iterative refinement with self-feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Malik Ghallab, Adele Howe, Craig Knoblock, Drew McDermott, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. 1998. PDDL - The Planning Domain Definition Language.pdf.
- OpenAI. 2023. GPT-4 Technical Report.
- Vishal Pallagani, Bharath Muppasani, Keerthiram Murugesan, Francesca Rossi, Lior Horesh, Biplav Srivastava, Francesco Fabiano, and Andrea Loreggia. 2022. Plansformer: Generating Symbolic Plans using Transformers.
- Silvia Richter and Matthias Westphal. 2010. The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks. *Journal of Artificial Intelligence Research*, 39:127–177.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2024. Code llama: Open foundation models for code.
- Stuart Russell and Peter Norvig. 2009. Artificial Intelligence: A Modern Approach, third edition. Pearson.
- Jendrik Seipp, Álvaro Torralba, and Jörg Hoffmann. 2022. PDDL generators. https://doi.org/10. 5281/zenodo.6382173.
- Tom Silver, Soham Dan, Kavitha Srinivas, Joshua B. Tenenbaum, Leslie Pack Kaelbling, and Michael Katz. 2023. Generalized Planning in PDDL Domains with Pretrained Large Language Models.
- Tom Silver, Varun Hariprasad, Reece S. Shuttleworth, Nishanth Kumar, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. 2022. PDDL Planning with Pretrained Large Language Models.
- Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M. Sadler, Wei-Lun Chao, and Yu Su. 2023.
 LLM-Planner: Few-Shot Grounded Planning for Embodied Agents with Large Language Models.

- Katharina Stein and Alexander Koller. 2023. Autoplanbench: : Automatically generating benchmarks for llm planners from pddl.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and finetuned chat models.
- Karthik Valmeekam, Matthew Marquez, and Subbarao Kambhampati. 2023a. Can large language models really improve by self-critiquing their own plans?
- Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2023b. PlanBench: An Extensible Benchmark for Evaluating Large Language Models on Planning and Reasoning about Change.
- Karthik Valmeekam, Sarath Sreedharan, Matthew Marquez, Alberto Olmo, and Subbarao Kambhampati. 2023c. On the Planning Abilities of Large Language Models (A Critical Investigation with a Proposed Benchmark).
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An Open-Ended Embodied Agent with Large Language Models.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. In Advances in Neural Information Processing Systems.
- Yue Wu, Shrimai Prabhumoye, So Yeon Min, Yonatan Bisk, Ruslan Salakhutdinov, Amos Azaria, Tom Mitchell, and Yuanzhi Li. 2023. Spring: Studying the paper and reasoning to play games.
- Rui Yang, Lin Song, Yanwei Li, Sijie Zhao, Yixiao Ge, Xiu Li, and Ying Shan. 2023. GPT4Tools: Teaching Large Language Model to Use Tools via Selfinstruction.

Ryan Yang, Tom Silver, Aidan Curtis, Tomas Lozano-Perez, and Leslie Pack Kaelbling. 2022. PG3: Policy-Guided Planning for Generalized Policy Generation.

816

817

- 819 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran,
 820 Thomas L. Griffiths, Yuan Cao, and Karthik
 821 Narasimhan. 2023. Tree of thoughts: Deliberate
 822 problem solving with large language models.
- Bayin, Faeze Brahman, Abhilasha Ravichander, Khyathi Chandu, Kai-Wei Chang, Yejin Choi, and
 Bill Yuchen Lin. 2023. Lumos: Learning agents
 with unified data, modular design, and open-source
 llms.

829

834

836

838

840

841

842

845

851

857

861

A PDDL

```
(define (problem ferry-l3-c2)
  (:domain ferry)
  (:objects l0 l1 l2 c0 c1 )
  (:init
      (empty-ferry)
      (at c0 l1)
      (at c1 l2)
      (at-ferry l2)
  )
  (:goal
      (and
           (at c0 l0)
           (at c1 l0)
      )
  )
```

Figure 5: An example PDDL problem instance definition for the Ferry domain.

Classical planning is based on the notion that the entire planning domain and problem instance are described in a formalised, machine-readable format. One common format is the Planning Domain Definition Language (PDDL; Malik Ghallab et al., 1998). The domain definition (Figure 6) encodes the physics of the world, while the problem instance definition (Figure 5) specifies the initial state and desired goals, thereby customizing the domain to a specific scenario. The planning task entails generating a sequence of actions that facilitate a transition from the initial state to a goal state, thereby constituting a plan.

A domain consists of a pair <P, A>, where P represents a set of predicates, and A is a set of actions. Each predicate $p \in P$ includes a name and a set of variables denoted by a question mark. For instance, at-ferry ?1 is a predicate, while at-ferry(loc1) is a specific truth-assignment to the predicate, indicating the ferry's presence at location 1. Predicates can also be negated (e.g., not(at-ferry(loc1))). An action $a \in A$ consists of a name, a set of variables, a set of effect predicates, and a set of precondition predicates. For instance, board(?car ?loc) is an action denoting the boarding of car ?car on the ferry at location ?loc. The action's preconditions include predicates such as at_ferry(?loc), at(?car,?loc), and empty_ferry, signifying that both the ferry and the car ?car are at location ?loc, and that the ferry is empty. The action's effects include predicates such as not(empty_ferry), not(at(?car,?loc)), and on(?car), indicating

```
(define (domain ferry)
   (:predicates (at-ferry ?1) (at ?c ?1)
        (empty-ferry) (on ?c))
   (:action sail
       :parameters (?from ?to)
       :precondition (at-ferry ?from)
       :effect (and (at-ferry ?to) (not
           (at-ferry ?from))))
   (:action board
       :parameters (?car ?loc)
       :precondition (and (at ?car ?loc)
           (at-ferry ?loc) (empty-ferry))
       :effect (and (on ?car) (not (at
           ?car ?loc)) (not
           (empty-ferry))))
   (:action debark
       :parameters (?car ?loc)
       :precondition (and (on ?car)
           (at-ferry ?loc))
       :effect (and (at ?car ?loc)
           (empty-ferry) (not (on ?car))))
```

Figure 6: The PDDL Ferry domain definition. The domain definition specifies the predicates and actions, encapsulating the physics of the domain.

that the ferry is no longer empty, that the car ?car is no longer at location ?loc, and that the car ?car is now on the ferry.

A problem is a triple <0, I, G>, where 0 is a set of *objects*, I is a set of truth-assigned predicates that are currently true in the world model, and G is a set of truth-assigned predicates designated to be achieved.

B Shortcomings of LLMs Experiment Details

In Section 3 we described a set of controlled experiments which test for specific reasoning abilities of LLM-based planners. Our prompt is constructed as following: for each domain we start by describing the domain, such as exemplified by Figure 7. Then, using the relevant prompt for each experiment (Figures 8 to 11), we add 2 example questions which contain an answer, followed by an unsolved example.

We are dealing with the Blocksworld problem. In this domain we have 4 possible actions: 1. pickup - (?ob - object) - Preconditions: The object (?ob) must be clear, on the table, and the arm must be empty. - Effects: After executing the action, the object is now held, not clear, and not on the table. - Example: If you execute (pickup blockA), it means you pick up "blockA" from the table. 2. putdown - (?ob - object) - Preconditions: The object (?ob) must be currently held. - Effects: After executing the action, the object is now clear, the arm is empty, and the object is on the table. - Example: If you execute (putdown blockB), it means you put down
"blockB" on the table. 3. stack - (?ob - object, ?underob - object) - Preconditions: The object that you want to stack (?ob) must be held, and the object underneath (?underob) must be clear. - Effects: After executing the action, the arm is empty, the stacked object (?ob) is clear, and it is now on top of the underneath object (?underob). The underneath object is no longer clear. - Example: If you execute (stack blockC blockD), it means you stack "blockC" on top of "blockD". 4. unstack - (?ob - object, ?underob object) - Preconditions: The object that you want to unstack (?ob) must be on top of another object (?underob), and both objects must be clear. Additionally, the arm must be empty. - Effects: After executing the action, the object (?ob) is now held, the underneath object (?underob) is clear, and the relationship "on" between (?ob) and (?underob) is broken. Also, (?ob) and (?underob) are no longer clear, and the arm is not empty. - Example: If you execute (unstack blockC blockD), it means you unstack "blockC" from on top of "blockD".

Figure 7: Explanation provided about the Blocksworld domain before each Blocksworld experiment described in Section 3.

Given this initial state: <STATE> and the following actions: <ACTIONS>

What is the new state? Answer:

Figure 8: Prompt template used for the actions' effects experiment.

Given this state: <STATE>

Which one of the following actions can be performed from the above state? The possible actions are: <ACTIONS>

Answer:

Figure 9: Prompt template used for the applicable actions experiment.

Given this state: <STATE>

Which one of the following goals can be achieved with fewer actions? The goals are: <GOALS>

Answer:

Figure 10: Prompt template used for the optimal goals ordering experiment.

Given this state: <STATE>

In order to complete the task both goals
 must be accomplished simultaneously.
Which one of the following goals needs to
 be achieved first?
The goals are: <GOALS>

Answer:

Figure 11: Prompt template used for the reasonable or necessary goals ordering experiment.

975

927

928

929

930

881 882

887

897

900

901

903

904

905

907 908

909

910

911

912

913

914

915

917

918

919

921

922

923

925

926

С **Generating Problem Instances and** Plans

In this section we provide more details about the generated problem instances and the plans to solve them. We start by generating problem instances with the PDDL generators library (Seipp et al., 2022), which we use to generate our experiments' data, as described in Section 3, as well as the training, validation and test datasets, as described in Section 5. When generating problems, the amount of objects for each generated problem instance is controllable. For example, the Grippers domain problem generator configuration controls the amount of robots, balls and rooms. For each domain, we choose one object to tweak its amount, and fix the amount of all other objects. For example, in the 896 Grippers domain we only change the number of balls, and we fix the number of robots to 1 and the number of rooms to 5. We then increase the number of instances of that object between the simple and complex configurations to create longer plans. 902 The configurations used are described in Table 6and the size of the datasets is described in Table 7. An example of the complex configuration plan is provided in Figure 18 to illustrate the difficulty of the task.

We then solve the generated problem instances with the LAMA planner⁴ and save the found plans. Lastly, we remove problem instances based on the following filters:

> • Timeout or search unsolved: Problem instances that the LAMA planner could not solve after a ten-minute timeout or after exploring the entire search space were removed. If there were many such errors, we tweaked the configurations such that the LAMA planner can solve the problems, until we reached the configurations described in Table 6.

- Empty plans: Instances where their plans have no steps were removed. This can happen in cases where the initial state is a goal state.
- Duplicates: Instances with the same initial state and goals are removed.

C.1 Domains

Blocksworld. In this domain, the objects are blocks and the agent is a robotic arm that can pick them up and put them down on a table. Blocks can be stacked, and only a block that is clear (i.e., doesn't have a block on top of it) can be picked up. The goal predicates are to stack blocks on top of other blocks.

Ferry. In this domain the objects are cars and locations, and the agent is a ferry that can move between locations, board a car in one location or debark it in another location. The ferry can only board one car at a time. The problem requires the ferry to transport the cars to their designated locations.

Grippers. In this domain the objects are balls and rooms, and the agent is a robot that can move between rooms, pick-up balls or drop a picked-up ball. The robot can hold two balls simultaneously with its left and right grippers. The problem requires the robot to transport the balls to their designated rooms.

Depots. Similar to Blocksworld, the objects are crates and the agent is a hoist that can pick them up or put them down. The difference from Blocksworld is that there are multiple "tables", which are called pallets, and each pallet has its own hoist. The pallets are located in different locations, called depots and distributors. The truck is another agent that can move the crates between locations. This is similar to Ferry and Grippers. The goal predicates are to stack the crates at specific locations, or on top of other crates.

Minigrid. In this grid domain, the objects are walls, keys and doors, and the agent is a robot that can traverse the grid, pick up keys and unlock doors. Doors can be locked, and their locks have certain shapes with keys having matching shapes. Only one key can be picked up at a time, thus the robot must first drop a key in order to pick up a different key. The more number of shapes there are, the more often the robot has to switch the key it is holding. The goal predicate is to reach a certain location in the board. The floor plan that we used for all Minigrid problems is depicted in Figure 12.

C.2 Domains Modifications

Some domains have a lot of predicates to describe the current state, thus creating a very long prompt for the LLM. This is a problem mostly in the Minigrid domain, which as seen in Figure 12, has a large 2D map which contains 64 places. This map is meticulously described in the PDDL format. To

⁴To run the LAMA planner, we use the open-source Fast Downward planning system https://www.fast-downward. org

	Bl	ocks	F	ERRY	Gri	PPERS	De	Depots Minig		NIGRID
	SIMPLE	COMPLEX	SIMPLE	COMPLEX	SIMPLE	COMPLEX	SIMPLE	COMPLEX	SIMPLE	COMPLEX
# OBJECTS	2-5	11-25	2-5	11-25	2-5	11-25	2-5	11-15	1-2	3-8
AVG. PLAN	8.3	88.5	11.3	50.2	10.1	52.5	10.4	132.1	19.5	27.3
LENGTH										
CHANGED	Bı	OCKS			VII S	CDATES		Su	ADEC	
Object	DL	OCKS	C	AKS	BALI		CRATES		SHAPES	

Table 6: Changes in the number of objects and average plan lengths between the simple and complex configurations.

S	IMPLE	COMPLEX	
TRAIN	DEV	Test	Test
100	30	30	50

Table 7: The size of the sets generated for each domain.

Lw
wL
LwwwwwL
Lw
wL
wLwwwww
wL
Lw

Figure 12: Minigrid floorplan that we used for generating problem instances. L represents locked doors, W represents blocked places which can't be crossed, and all other dots represent places where the agent can walk. In each problem instance, the robot start and end location are chosen randomly. The keys are spread around the map, not depicted in this figure.

976 sidestep this problem of having extremely long prompts, we make several changes to their PDDL 977 definition which allow us to create a shorter state 978 representation. Firstly, we remove predicates that contain typing information of objects. For example, we remove predicates that mention that the ob-981 ject p1 is a place (place p1), that the object key1 982 is a key (key key1), and that the object shape1 983 is a shape (shape shape1). Given this change, it should still be possible to calculate applicable 985 actions, as LLMs have different representations for each token, and they can infer the type of the object from the prefix, context and provided ex-989 amples. In addition, we remove predicates which contain information about connected paths, such 990 as p1 is connected to p2 (conn p1 p2). Originally, 991 this is used in Minigird to represent the locations 992 of walls. We remove this only for the fine-tuned 993

models which can learn the map during training. Finally, in Minigrid, we remove predicates which indicate that a door is open (e.g., open p3). The predicates already describe which doors are locked (e.g., locked p16), and it should be possible for the LLM to see which places are locked and from it to infer the open doors. 994

995

996

997

998

999

1000

1001

1023

1024

1025

1026

C.3 Goal Planning Experiment

As described in Section 3.2, we devise experiments 1002 in which there is one goal that should reasonably or necessarily be prioritized over a different goal. To 1004 identify such instances, we employ the landmarks 1005 graph generated by the LAMA planner (Richter 1006 and Westphal, 2010). Landmarks are sub-goals 1007 that the LAMA planner algorithm identified as 1008 necessary to complete in any possible plan that 1009 solves the given problem instance. This graph's 1010 nodes represent goals or intermediate sub-goals, 1011 with directed edges denoting the order in which 1012 these goals should be approached. The edges are 1013 labeled to indicate the type of ordering, such as 1014 "n" for necessary and "r" for reasonable. Utilizing 1015 these labels, we construct prompts for the language 1016 model that requires prioritizing between 2 distinct 1017 goals, as described in Appendix B. In the context of 1018 the Blocksworld domain, our selection criteria for 1019 goals involve choosing those linked by a minimum 1020 of two edges, ensuring that the same block does 1021 not appear in both goals. 1022

D Training Details

This section provides further details about the training of our SimPlan model and the baselines.

D.1 Approaches

SimPlan.We start by reporting results for our1027similarity-based retrieval model, SimPlan.For1028evaluating the retrieval model, we use the met-1029ric MAP@100.The performance of the retrieval1030

model on the intrinsic development set (comprised 1031 of simple difficulty problems) ranges from 60% to 1032 90% between domains. For all domains we used 1033 a batch size of 32, except for the Depots and Min-1034 igrid domains which we used a batch size of 16. Then, an extra 2 hard negatives are added per ex-1036 ample in the batch, sampled from the pool of hard 1037 negatives which was extracted using the techniques 1038 described in Section 4.2. We used one A100 80-1039 GB GPU to train each model for 10 epochs with 1040 a training time of 12 hours, and selected the best checkpoint. For hyperparameter tuning we used a 1042 grid search with the following parameters: learn-1043 ing rate varied between [4e-5, 4e-4, 4e-3], warmup 1044 steps varied between [0, 100, 500] and weight de-1045 cay varied between [0.01, 0.001, 0.0001]. After hyperparameter tuning, we fixed the parameters to a learning rate of 4e-4, warmup steps of 100 and 1048 a weight decay of 0.001. Our implementation is 1049 based on the sentence transformers library⁵ adapted with code from the ColBERT library. We imple-1051 mented our own simple planning translator based on the PDDL domains definitions. See Figure 14 for an example input / output.⁶ 1054

LLM4PDDL. Our first two baselines are based on the in-context learning approach, for which we follow the few-shot prompt design from LLM4PDDL (Silver et al., 2022) with 2 training examples. See Figure 15 for an example prompt. For GPT-4 Turbo, we use the vanilla strategy where we simply prompt the model to generate a plan with temperature 1.0. For CodeLlama-34b-instruct, we use the soft-validation strategy for dealing with malformed LLM outputs. In the soft-validation strategy, a planning translator is used to find all applicable actions at each step. After each action is generated, it is validated that it is applicable in the current state. If it is not, it is replaced with the closest valid action, based on a cosine similarity distance in a pre-trained embedding space of the paraphrase-MiniLM-L6-v2 model. For CodeLlama-34b-instruct, we use a beam size of 1, similar to LLM4PDDL. To allow more exploration of the search space, we sample 16 different plans from the LLM with a temperature of 0.5.

1055

1057

1058

1059

1061

1062

1063

1064

1065

1066

1067

1069

1070

1071

1074

1075

1078

Plansformer. We fine-tune an LLM on the training set of problem instances from the simple configuration. The following fine-tuning settings de-

```
<INITIAL> the hand is holding b1 <INITIAL>
b24 is clear <INITIAL> b24 is on the
table
<GOAL> b24 is on top of b1
```

Figure 13: Example *SimPlan* state parsed and linearized into tokens. During training, for each problem instance the blocks names are scrambled in order for the model to train the embeddings of all relevant tokens. In this example, b2 was converted into b24.

<goal> on b2 b1 <init> arm-empty , clear b1, on b1 b2, on-table b2 <action> pickup</action></init></goal>
<pre> clear x, on-table x, arm-empty</pre>
<pre><effect> holding x, not clear x, not</effect></pre>
on-table x, not arm-empty
<action> putdown</action>
<pre> holding x</pre>
<effect> clear x, on-table x,</effect>
arm-empty, not holding x
<action> stack</action>
<pre> clear y, holding x</pre>
<effect> arm-empty, clear x, on x y,</effect>
not clear y, not holding x
<action> unstack</action>
<pre> on x y, clear x, arm-empty</pre>
<pre><effect> clear y, holding x, not</effect></pre>
arm-empty, not clear x, not on x y

Figure 14: Example Plansformer instance.

scribed is based on Pallagani et al. (2022). We use CodaLlama-7b as the base model, which follows the intuition by Plansformer to use code models. The input to the model contains the initial state and the goals, along with the description of the domain's actions. We add special tokens between the different parts, such as <GOAL>, <INIT>, <ACTION>, <PRE>, <EFFECT> to describe the goals, the initial state, and the actions, along with their preconditions and effects. The gold output is a sequence of actions which make up a plan. See Figure 13 for an example of an input / output. For early stopping, we use ROUGE-L to compare the gold output with the predicted output on a development set of easy problems. During generation, we use a beam search of 16 to allow for a large exploration of the state space.

1079

1080

1081

1082

1083

1084

1085

1086

1087

1088

1089

1092

1094

D.2 Inference-time constraints

To produce comparable results between the differ-
ent models, for each model we limit the number1097of next actions predictions. For SimPlan and the
random baseline, the number of predictions is sim-1098

⁵https://www.sbert.net/

⁶https://github.com/stanford-futuredata/ColBERT

```
0:
(:objects b1 b2 - object)
(:init (arm-empty) (on b1 b2) (on-table b2) (clear b1))
(:goal (on b2 b1))
Α:
(unstack b1 b2) (putdown b1) (pickup b2) (stack b2 b1)
(:objects b1 b2 b3 b4 b5 - object)
(:init (arm-empty) (on b1 b4) (on-table b2) (on b3 b1) (on-table b4) (on b5 b3) (clear b2)
    (clear b5))
(:goal (on b1 b2) (on b2 b3) (on b3 b5) (on b4 b1))
(unstack b5 b3) (putdown b5) (unstack b3 b1) (stack b3 b5) (pickup b2) (stack b2 b3) (unstack
    b1 b4) (stack b1 b2) (pickup b4) (stack b4 b1)
(:objects: b1 b10 b11 b12 b13 b14 b15 b2 b3 b4 b5 b6 b7 b8 b9 - object)
(:init (arm-empty) (on b1 b5) (on b2 b6) (on-table b3) (on b4 b13) (on b5 b15) (on b6 b1) (on
    b7 b3) (on-table b8) (on b9 b2) (on b10 b8) (on-table b11) (on-table b12) (on b13 b9) (on
    b14 b12) (on-table b15) (clear b4) (clear b7) (clear b10) (clear b11) (clear b14))
(:goal (on b1 b13) (on b2 b10) (on b3 b5) (on b4 b3) (on b8 b12) (on b9 b4) (on b10 b14) (on
    b11 b6) (on b12 b15) (on b13 b8) (on b15 b7))
```

Figure 15: Example LLM4PDDL baseline prompt for a Blocksworld problem instance prepended with two training examples. This format is based on Silver et al. (2022), where newlines separating the predicates were replaced with spaces in this figure for the sake of brevity.

1101 ply the count of calls to the model, since we predict one next action in each model call. For LLM4PDDL 1102 and Plansformer, after each token is generated, 1103 we implemented a check if this token is one of the 1104 following tokens: ")" or "), (". If yes, we identify 1105 this as a generated action and count this as one 1106 prediction. An important aspect of beam search 1107 is that each beam generating an action increases 1108 this count by one. A stop criteria is applied once 1109 the number of predictions exceeds the limitation. 1110 We note that this limitation is not implemented for 1111 GPT-4, as the amount of beams that GPT-4 uses is 1112 not information that we found available. 1113

The value of the limit for next action predictions 1114 is configured differently per problem instance. We 1115 first calculate how many states were expanded by 1116 the LAMA planner when solving each problem in-1117 stance. However, since the LAMA planner uses 1118 greedy best-first search, it is not limited to sequen-1119 tial exploration like beam search. It is thus possible 1120 that this limitation is too harsh for beam-search de-1121 coding algorithms which sequentially extract multi-1122 1123 ple paths, reducing the overall maximum explored plan length. We thus multiply this limitation by 16, 1124 which is the beam size we used for beam-search 1125 decoding algorithms. Finally, we also set a timeout 1126 of 5 minutes for all models. 1127

E Ablation Experiment Details

In this section we provide more information about 1129 the ablation experiments in Section 5.4, specifically 1130 about the state updates experiment. As mentioned 1131 in Section 4.3, our SimPlan method involves up-1132 dating the state at each step, and the history of the 1133 actions taken is not included in the input to the en-1134 coder. For the state updates experiment, we train a 1135 variant of our model where the input to the model is 1136 fixed to the initial state as defined in the problem in-1137 stance, the set of goals, and the sequence of actions 1138 taken so far. Special tokens are used to separate 1139 between the state predicates, goal predicates and 1140 the actions. The model is then expected to infer 1141 the current state based on the initial state and the 1142 sequence of actions, similar to the LLM4PDDL 1143 and Plansformer baselines. 1144

1128

1145

F Learned Representation Analysis

We analyze the learned representation of the actions1146and the states, described in Section 4, by using the1147T-SNE algorithm (Figures 16 and 17).7 Since we1148are using a late interaction schema, each individual1149state or action representation has multiple token1150embeddings. For our analysis, we chose to take the1151

⁷Implemented with https://scikit-learn.org



Figure 16: T-SNE plot of the learned representation of actions in the Ferry domain. The shapes depict the different action types, and the colors depict the different locations where the ferry is at when taking the action. The visible clusters of colors indicate that the model learned the concept of applicable actions.



Figure 17: T-SNE plot of the learned representation of sampled states in the Ferry domain. Colors and shapes are determined based on the next action to take in this state according to the gold plan. The visible clusters of color-shape pairs indicate that the model is taking goal-informed decisions when representing the current state.

maximum token embedding. To create states to use for our analysis, we sampled 5 problem instances from our training dataset and extracted a state after the execution of each action in the accompanied plan. In the visualization, each state is assigned an action based on the next action in the plan.

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

The figures in our analysis show that the model learns the concept of applicable actions and next action planning. In our settings of the Ferry domain, there are overall 5 locations. In the actions representations (Figure 16), the colors represent the location of the ferry. Since the colors are wellclustered, and the location of the ferry determines the applicable actions, this demonstrates that the 1165 model learned the concept of applicable actions. In 1166 addition, in the states representations visualization 1167 (Figure 17), we can see that different states which 1168 have similar next actions are clustered together. 1169 This demonstrates that the action-ranking model 1170 is taking goal-informed decisions. Interestingly, 1171 in the states representation, the model clusters to-1172 gether the debark actions (star shape). 1173 (lift hoist2 crate11 crate7 distributor0),(lift hoist1 crate14 crate13 depot1),(load hoist1 crate14 truck0 depot1),(lift hoist0 crate12 crate3 depot0),(drive truck0 depot1 depot0),(load hoist0 crate12 truck0 depot0),(drive truck0 depot0 depot1),(unload hoist1 crate12 truck0 depot1),(drive truck0 depot1 distributor0),(load hoist2 crate11 truck0 distributor0),(drive truck0 distributor0 depot0),(unload hoist0 crate14 truck0 depot0),(drive truck0 depot0),(drive truck0 depot0),(drive truck0 depot1),(drive truck0 depo hoist2 crate7 truck0 distributor0),(lift hoist2 crate4 crate2 distributor0),(load hoist2 crate4 truck0 distributor0),(drive tru distributor0 distributor1),(lift hoist2 crate2 crate1 distributor0),(unload hoist3 crate4 truck0 distributor1),(drive truck0 distributor0),(load hoist2 crate2 truck0 distributor0),(drive truck0 distributor0),(load hoist3 crate4 truck0 distributor1), (unload hoist3 crate2 truck0 distributor1), (drive truck0 distributor1 depot1), (drive truck0 depot1 distributor0), (unload hoist2 crate7 truck0 distributor0),(drop hoist2 crate7 crate1 distributor0),(drive truck0 distributor0 distributor1),(drop hoist1 crate12 crate13 depot1),(drop hoist3 crate2 crate9 distributor1),(unload hoist3 crate11 truck0 distributor1),(drive truck0 distributor1 distributor0),(lift hoist2 crate7 crate1 distributor0),(load hoist2 crate7 truck0 distributor0),(lift hoist2 crate1 pallet2 distributor0),(drive truck0 distributor0 depot1),(lift hoist1 crate12 crate13 depot1),(load hoist1 crate12 truck0 depot1),(lift hoist1 cratel3 cratel0 depot1),(load hoist1 cratel3 truck0 depot1),(lift hoist1 crate10 crate0 depot1),(load hoist1 crate10 truck0 depot1),(drive truck0 depot1 distributor0),(lift hoist1 crate0 pallet1 depot1),(drop hoist2 crate1 pallet2 distributor0),(unload hoist2 crate10 truck0 distributor0),(drop hoist2 crate10 crate1 distributor0),(unload hoist2 crate13 truck0 distributor0),(drive truck0 distributor0 depot1),(load hoist1 crate0 truck0 depot1),(unload hoist1 crate12 truck0 depot1),(drive truck0 depot1 distributor1),(drop hoist3 crate11 crate2 distributor1), (unload hoist3 crate4 truck0 distributor1), (drop hoist1 crate12 pallet1 depot1), (drop hoist2 crate13 crate10 distributor0), (drop hoist3 crate4 crate11 distributor1), (unload hoist3 crate7 truck0 distributor1), (drive truck0 distributor1) distributor0), (unload hoist2 crate0 truck0 distributor0),(drop hoist2 crate0 crate13 distributor0),(drive truck0 distributor0 distributor1),(load hoist3 crate7 truck0 distributor1),(lift hoist3 crate4 crate11 distributor1),(load hoist3 crate4 truck0 distributor1),(lift hoist3 crate11 crate2 distributor1),(drive truck0 distributor1 distributor0),(unload hoist2 crate4 truck0 distributor0),(drop hoist2 crate4 crate0 distributor0),(drive truck0 distributor0 depot1),(unload hoist1 crate7 truck0 depot1),(drive truck0 depot1 distributor0),(drop hoist1 crate7 crate12 depot1),(lift hoist2 crate4 crate0 distributor0),(load hoist2 crate4 truck0 distributor0),(drive truck0 distributor0 depot1),(unload hoist1 crate4 truck0 depot1),(drive truck0 depot1 distributor0),(lift hoist2 crate0 crate13 distributor0),(drop hoist1 crate4 crate7 depot1),(load hoist2 crate0 truck0 distributor0),(lift hoist2 crate13 crate10 distributor0),(drive truck0 distributor0 depot1),(lift hoist1 crate4 crate7 depot1),(load hoist1 crate4 truck0 depot1),(drive truck0 depot1 distributor0), (lift hoist1 crate7 crate12 depot1), (load hoist2 crate13 truck0 distributor0), (drive truck0 distributor0 depot1), (load hoist1 crate7 truck0 depot1), (unload hoist1 crate13 truck0 depot1), (drop hoist1 crate12 depot1), (unload hoist1 crate0 truck0 depot1), (drop hoist1 crate0 crate13 depot1), (unload hoist1 crate4 truck0 depot1), (drop hoist1 crate4 crate0 depot1), (drive truck0 depot1 distributor1),(lift hoist2 crate10 crate1 distributor0),(drop hoist3 crate11 crate2 distributor1),(unload hoist3 crate7 truck0 distributor1),(drive truck0 distributor1 distributor0),(load hoist2 crate10 truck0 distributor0),(lift hoist2 crate1 pallet2 distributor0),(drop hoist3 crate7 crate11 distributor1),(load hoist2 crate1 truck0 distributor0),(unload hoist2 crate10 truck0 distributor0), (drive truck0 distributor0 distributor1), (drop hoist2 crate10 pallet2 distributor0), (drive truck0 distributor1 distributor0), (unload hoist2 crate1 truck0 distributor0), (drop hoist2 crate1 crate10 distributor0), (drive truck0 distributor0 distributor1),(lift hoist3 crate7 crate11 distributor1),(load hoist3 crate7 truck0 distributor1),(lift hoist3 crate11 crate2 distributor]),(drive truck0 distributor1 distributor0),(unload hoist2 crate7 truck0 distributor0),(drop hoist2 crate7 crate1 distributor0),(drive truck0 distributor0),(drive truck0 distributor0),(load hoist0 crate14 truck0 depot0),(lift hoist0 crate3 pallet0 depot0),(load hoist0 crate3 truck0 depot0),(unload hoist0 crate14 truck0 depot0),(drive truck0 depot0 depot1),(drop hoist0 crate14 pallet0 depot0),(drive truck0 depot1 distributor1),(load hoist3 crate11 truck0 distributor1),(drive truck0 distributor1 depot1),(unload hoist1 crate11 truck0 depot1),(drive truck0 depot1 distributor1),(lift hoist3 crate2 crate9 distributor1),(load hoist3 crate2 truck0 distributor1),(drive truck0 distributor1 depot0),(unload hoist0 crate2 truck0 depot0),(drop hoist0 crate2 crate14 depot0),(drive truck0 depot0 distributor1),(lift hoist3 crate9 crate8 distributor1),(load hoist3 crate9 truck0 distributor1),(lift hoist3 crate6 crate6 distributor1),(load hoist3 cr truck0 distributor1),(lift hoist3 crate6 crate5 distributor1),(drive truck0 distributor1 distributor0),(unload hoist2 crate8 truck0 distributor0),(drive truck0 distributor0 distributor1),(load hoist3 crate6 truck0 distributor1),(lift hoist3 crate5 pallet3 distributor1),(drive truck0 distributor1 distributor0),(load hoist2 crate8 truck0 distributor0),(unload hoist2 crate6 truck0 distributor0),(drop hoist2 crate6 crate7 distributor0),(unload hoist2 crate8 truck0 distributor0),(drive truck0 distributor0),(lift hoist0 crate2 crate14 depot0),(load hoist0 crate2 truck0 depot0),(drive truck0 depot0 distributor1),(load hoist3 crate5 truck0 distributor1),(unload hoist3 crate2 truck0 distributor1),(drive truck0 distributor1 depot1),(drop hoist3 crate2 pallet3 distributor1),(drive truck0 depot1 distributor1),(unload hoist3 crate5 truck0 distributor1),(drive truck0 distributor1 distributor0),(load hoist2 crate8 truck0 distributor0),(drive truck0 distributor0 distributor1),(load hoist3 crate5 truck0 distributor1),(unload hoist3 crate8 truck0 distributor1),(drive truck0 distributor1 depot1),(drop hoist3 crate8 crate2 distributor1),(drive truck0 depot1 distributor1),(unload hoist3 crate9 truck0 distributor1),(drop hoist3 crate9 crate8 distributor1),(unload hoist3 crate5 truck0 distributor1),(drive truck0 distributor1 depot1),(load hoist1 crate11 truck0 depot1),(drive truck0 depot1 distributor1),(load hoist3 crate5 truck0 distributor1),(drive truck0 distributor1 depot1),(unload hoist1 crate5 truck0 depot1),(drop hoist1 crate5 crate4 depot1),(unload hoist1 crate11 truck0 depot1),(drop hoist1 crate11 crate5 depot1)

Figure 18: Example plan with 184 actions for a problem instance from the Depots' complex configuration.