

Learning to Draw ASCII Improves Spatial Reasoning in Language Models

Anonymous ACL submission

Abstract

Human spatial cognition involves both perceiving spatial configurations and reconstructing them through sketching or mental imagery. Although perceiving is often intuitive, reconstruction is substantially more demanding. To facilitate this process, humans frequently rely on sketching as a “read-write loop” to visualize ideas and anchor their reasoning when solving complex spatial problems. In this work, we investigate whether Large Language Models (LLMs) exhibit a similar dependency when processing spatial topology. We introduce TEXT2SPACE, a dataset that pairs natural language descriptions with ground-truth ASCII representations and QA pairs, using ASCII as a diagnostic probe to decouple representation failures from reasoning failures. Our evaluation reveals a pronounced “Read-Write Asymmetry”: LLMs interpret ASCII representations effectively but struggle to reconstruct them, and this generation bottleneck propagates to the subsequent reasoning phase, producing flawed intermediate representations that distort subsequent reasoning. Training models on Text $\leftrightarrow$ ASCII alignment significantly improves performance on text-only spatial reasoning tasks, even without generating ASCII at inference time. Our findings show that enabling LLMs to “learn to draw” not only allows them to better leverage sketching as a reasoning strategy but also fosters deeper internalization of spatial understanding.

1 Introduction

Spatial reasoning (Byrne and Johnson-Laird, 1989; Newcombe and Huttenlocher, 2000) is fundamental to human cognition (Malanchini et al., 2020). As language models (LLMs) are increasingly applied to spatial reasoning tasks, evaluating their capabilities has become a critical priority (Li et al., 2024; Cohn and Hernandez-Orallo, 2023; Wang et al., 2024; Liu et al., 2023; Cheng et al., 2024).

However, existing benchmarks are limited to evaluating final answers (Yamada et al., 2023; Liu et al., 2023). When a model fails, the absence of interpretability makes it unclear whether the error stems from an inaccurate internal representation or from insufficient reasoning over an otherwise correct representation. This lack of observability into intermediate spatial representations during the reasoning process creates a diagnostic gap, hindering the identification of effective ways to achieve faithful spatial reasoning in LLMs (Singh et al., 2024; Parcalabescu and Frank, 2024; Jacovi and Goldberg, 2020; Creswell and Shanahan, 2022).

We address this gap by introducing ASCII representations that encode two dimensional structure as explicit text based grids, reducing the ambiguity and variability inherent in sequential natural language descriptions (Piantadosi et al., 2012). ASCII diagrams offer an explicit and verifiable way to represent spatial layouts, allowing us to examine the intermediate representation a model produces and directly assess its correctness. To support this analysis, we introduce TEXT2SPACE, a benchmark of diverse spatial configurations paired with natural language descriptions, images, and three styles of ASCII grids. Each instance includes a ground truth ASCII representation, enabling evaluation of intermediate spatial representations during the reasoning process and improving interpretability in spatial reasoning tasks. The dataset further provides fine-grained annotations, including spatial complexity, terminology, and transitive inference requirements, to support systematic analysis.

Using TEXT2SPACE, we contribute three analyses. First, we **benchmark** language models on ASCII comprehension, generation, and spatial reasoning. Second, we examine the **reasoning process** in spatial tasks involving ASCII representations, assessing the extent to which these representations support model reasoning. Finally, we investigate whether enhancing models’ understanding




Figure 1: Overview of the study design. Inspired by empirical cognitive human strategies, we utilize Natural Language (NL) and ASCII representations as independent reasoning sources. This bidirectional cycle connects the two modalities via Parsing, Drawing, Interpreting, and Describing steps. We find that training LLMs to mimic this human-like, sketch-based reasoning process significantly improves accuracy on complex spatial reasoning tasks.

of spatial relations through ASCII-based learning generalizes to improved performance on purely linguistic spatial reasoning tasks.

Our evaluation yields three principal findings: **Finding 1: Comprehension exceeds generation.** Models can interpret spatial information from ASCII diagrams more effectively than they can produce such diagrams from text. This suggests that constructing a structured spatial representation from abstract descriptions is a more demanding cognitive process than reading and understanding one. **Finding 2: Flawed representations impair reasoning.** Most models achieve higher accuracy when answering directly from text without generating ASCII diagrams. Errors in constructing the intermediate spatial representation can propagate to the final answer. This highlights that reasoning performance is fundamentally constrained by the quality of the underlying representation. **Finding 3: Spatial knowledge generalizes from learning with ASCII.** Fine-tuning on ASCII-based tasks significantly enhances spatial reasoning from natural language descriptions. This indicates that models can acquire abstract spatial representations rather than merely memorizing ASCII syntax. The knowledge learned in the ASCII domain generalizes to broader spatial reasoning tasks.

We present three contributions: (1) a benchmark for evaluating spatial reasoning using ASCII as a reference; (2) a framework for assessing through the construction and interpretation of ASCII spatial representation; and (3) empirical evidence showing that spatial knowledge acquired from ASCII representations generalizes effectively to a wide range of spatial reasoning tasks.

2 Related Works

Linguistic Spatial Reasoning Benchmarks. Early work on spatial language processing focused on extracting formal representations from natural language text (Kordjamshidi et al., 2011). Subsequent benchmarks like bAbI (Weston et al., 2015) introduced positional reasoning, while STEPGAME (Shi et al., 2022) and SPARTQA (Mirzaee et al., 2021) advanced the field with more complicated reasoning and broader linguistic coverage. However, a critical limitation remains across these benchmarks: they rely exclusively on final-answer accuracy. This "black-box" evaluation obscures the underlying inference process, making it impossible to distinguish between failures in representation and failures in reasoning. Our work addresses this methodological gap by introducing ground-truth intermediate representations, enabling precise diagnosis of spatial reasoning errors.

While integrating visual modalities (VLMs) offers a potential remedy for spatial understanding, recent studies indicate that visual perception introduces its own challenges (Liu et al., 2023, 2024; Song et al., 2024; Kamath et al., 2023) and spatial reasoning (Chen et al., 2024; Zhang et al., 2025a; Ogezi and Shi, 2025). In contrast, our work decouples spatial reasoning from visual perception by utilizing ASCII grids. These grids serve a dual purpose: they act not only as a reasoning source within the context but also as a verifiable representation of all spatial relations.

ASCII and Spatial Representation. The utility of intermediate reasoning steps is well-established through Chain-of-Thought (CoT) prompting (Wei




Figure 2: An example in our TEXT2SPACE dataset: we have natural language description, Q-A pairs, three types of ASCII, and a rendered image.

et al., 2022) and its extensions, such as Tree-of-Thoughts (Yao et al., 2023) and latent reasoning (Hao et al., 2024). Beyond purely linguistic traces, prior work has shown that symbolic, text-native spatial representations can effectively support spatial reasoning in language models, including grid-based worlds and discrete relational layouts (Côté et al., 2018). Visualization-of-Thought (Wu et al., 2024; Li et al., 2025) draws on human mental imagery to generate symbolic or pixel-level visualizations for spatial reasoning. We propose a text-native alternative using ASCII to visualize spatial representations. Unlike natural language reasoning chains, which are hard to evaluate, ASCII creates a fully verifiable, spatially explicit “mental map” directly within the language model.

Generalization of Spatial Reasoning Recent research demonstrates that enhancing specific capabilities in LLMs often yields transfer benefits to other domains. For instance, Shao et al. (2024) indicate that training on mathematics improves general reasoning benchmarks, not just quantitative tasks. Additionally, in the field of spatial understanding, numerous studies have shown that a robust grasp of space benefits visual reasoning tasks (Wei et al., 2025; Han et al., 2025; Wang et al., 2025; Hu et al., 2025). Complementing these findings, our approach emphasizes building spatial cognition via ASCII generation and comprehension. We demonstrate that improvements in these seemingly non-reasoning tasks effectively generalize to text-based reasoning benchmarks.

3 TEXT2SPACE Dataset

We introduce TEXT2SPACE, a spatial reasoning QA dataset with text-based graph (ASCII) representations designed to validate the reasoning process in LLMs. This dataset evaluates how well LLMs interpret, align, and reason about spatial relationships. It presents diverse spatial configurations across four formats: ASCII art, natural language, rendered images, and Query-Answer pairs. An example is shown in Fig 2.

3.1 Data Generation

The dataset is constructed from connected spatial graphs where entities are linked via eight directional relationships: *above*, *below*, *left*, *right*, *upper-left*, *upper-right*, *lower-left*, *lower-right*. All relations adhere to a unified coordinate system, serving as the ground truth for both ASCII diagrams and rendered images.

Instances are generated using a breadth-first search (BFS) procedure to guarantee connectivity. We enforce strict validation rules: self-loops and duplicate edges are prohibited, and any configuration conflicting with existing spatial constraints is rejected. These safeguards ensure every instance represents a valid, internally consistent topology. A detailed data generation process is presented in Appendix F.

We vary complexity from simple two-component structures to intricate graphs with up to eight components and twelve relational statements, capturing configurations that demand transitive reasoning across multiple inference steps.

Each instance includes natural language descriptions synthesized from the spatial graph, with a wide range of synonyms used to ensure linguistic diversity. Corresponding queries are designed so that the target object pairs admit unique inferability, allowing only a single, unambiguous spatial relation to be inferred.

3.2 Multi-Format Representations

We derive both ASCII and image representations from the same underlying structure to facilitate multimodal evaluation and cross-format consistency.

ASCII Visualization: We include three rendering styles to test robustness against formatting variations. The *simple format* uses minimal spacing without borders for clarity. The *grid format* introduces structure using +, -, and | characters. The

panel format employs box-drawing characters for a visually distinct layout.

Image Representation: The spatial arrangement is also rendered in a standardized image format with configurable grid lines and directional arrows, serving as both a multimodal evaluation target and a human-readable reference. Further dataset details and data examples are provided in Appendix A.

4 Experiments

We split TEXT2SPACE into training and test sets, with the test set consisting of 1,000 samples used in the following experiments. We choose the ASCII *grid format* throughout our experiments; details of this format are provided in Appendix B. **Benchmarking:** How effectively do LLMs comprehend, generate, and utilize ASCII-based graph representations for spatial reasoning? **Internal Spatial Processing:** Does an ASCII intermediate step encourage structured, human-like reasoning, and how does this affect downstream performance? **Learning Spatial Relations with ASCII:** Does training LLMs to convert natural language descriptions into ASCII structures enhance spatial reasoning, and does this benefit transfer to pure natural language tasks?

4.1 Models

We evaluate two model categories. **Instruct models:** including LLAMA-3-70B-INSTRUCT (AI@Meta, 2024), QWEN2.5-7B-INSTRUCT (Yang et al., 2024; Team, 2024), QWEN3-30B-A3B-INSTRUCT-2507, QWEN3-235B-A22B-INSTRUCT-2507 (Team, 2025), GPT-4.1 (OpenAI, 2025). **Thinking model:** GPT-OSS-20B (OpenAI, 2025). Unless otherwise specified, we use default settings and apply 5-shot prompting (Brown et al., 2020) to ensure consistency.

4.2 Task Setups

Since ASCII is a unique representation that encodes two-dimensional spatial structure within a one-dimensional token sequence, we first evaluate language models’ ability to understand and generate ASCII representations. Specifically, we assess two complementary capabilities: generating an ASCII representation from a textual description (*Desc*→*ASCII*) and generating a natural-language description from an ASCII representation (*ASCII*→*Desc*).

We then move to spatial reasoning tasks, denoted by $\mathcal{R}$, in which the model is given a reasoning source and asked to answer queries about spatial relationships. $\mathcal{R}_{\text{Desc}}$ refers to spatial reasoning based solely on textual descriptions, while $\mathcal{R}_{\text{ASCII}}$ refers to reasoning based solely on ASCII representations.

Beyond reasoning from a single reference, we investigate the role of ASCII as an intermediate sketching process alongside descriptions, which follows a workflow sequence of *Desc* → *ASCII* → *Answer*. In $\mathcal{R}_{\text{Desc}:(\text{Ans}, \text{ASCII})}$, the model is given a description and first generates the spatial reasoning answer, and then produces an ASCII representation. Conversely, in $\mathcal{R}_{\text{Desc}:(\text{ASCII}, \text{Ans})}$, the model first generates an ASCII representation and then uses it to answer the query. For these settings, we consider both *one-turn generation*, where the model produces the ASCII and the answer within a single response, and *two-turn generation*, where the second output is elicited via a follow-up prompt. In this sense, ASCII functions as a form of self-explanation (Huang et al., 2023; Madsen et al., 2024), and the two generation orders mirror established paradigms: making a prediction and then explaining it, or generating an explanation and using it to make a prediction (Camburu et al., 2018).

Finally, we evaluate $\mathcal{R}_{\text{Desc}+\text{ASCII}}$. In this phase, we substitute the model-generated ASCII with the ground-truth ASCII before the model answers the query. This replacement strategy simulates an ideal reasoning process, allowing us to assess the upper-bound performance of the model when guided by an ideal spatial representation.

4.3 Evaluation Methods

Our evaluation protocol targets three key aspects of the model’s performance: **QA Correctness:** For the spatial reasoning QA tasks, we validate the prediction against a predefined set of relation aliases. **Representation Alignment:** For the *ASCII*↔*Desc* conversion tasks, we measure the correspondence between converted representations. For *ASCII*→*Desc*, the ASCII diagram is treated as the reference. We evaluate, across all generated descriptions, the percentage that can be correctly inferred from the given ASCII representation. Conversely, for *Desc*→*ASCII*, we evaluate the generated ASCII diagrams, $\mathcal{R}_{\text{Desc}:(\text{Ans}, \text{ASCII})}$ and $\mathcal{R}_{\text{Desc}:(\text{ASCII}, \text{Ans})}$, by taking the description as the reference and measuring the proportion of the description that is correctly reflected in the ASCII output. **Reasoning Consistency:** To analyze internal

coherence across the sketching tasks, we introduce two complementary metrics. Consistency: Measures the alignment between the generated ASCII and the generated Answer. This determines if the model’s intermediate output supports its final prediction. Consistency Accuracy: Measures answer accuracy strictly on instances where the ASCII and Answer are consistent. By controlling for internal coherence, this isolates the model’s reasoning capability, helping distinguish whether errors arise from failures in constructing spatial representations or from limitations in reasoning over them.

4.4 Learning Spatial Relations with ASCII

We propose that ASCII representations function as a structured reasoning intermediate: a verifiable “cognitive map” that explicates spatial relationships during model reasoning. We hypothesize that the active *process* of constructing such verifiable maps (Lightman et al., 2023; Zhang et al., 2025b) consolidates spatial logic during training, leading to improved performance on spatial reasoning tasks even when explicit ASCII representations are not produced at inference time.

To investigate the distinct roles of understanding versus generating ASCII representations, we fine-tune the models with three configurations: **Learning to Draw** ($Desc \rightarrow ASCII$), which forces the model to actively construct spatial layouts from text (performing “mental mapping”); **Learning to Read** ($ASCII \rightarrow Desc$), which focuses on parsing and verbalizing ASCII representations; and **Bidirectional Alignment** ($Desc \leftrightarrow ASCII$), a combined approach training on both tasks simultaneously.

We fine-tuned QWEN3-30B-A3B on 4x NVIDIA A40 GPUs using the SFTTrainer from the HuggingFace TRL library (von Werra et al., 2020). We employed Low-Rank Adaptation (LoRA) (Hu et al., 2021; Mangrulkar et al., 2022) for parameter-efficient training, utilizing a split of 4,000 training and 500 validation examples. We adhered to default library configurations for all other settings; full training details are provided in Appendix C.

We evaluate performance on the Spatial Reasoning Task ($\mathcal{R}_{Desc}$) where models predict spatial relations solely from natural language descriptions without explicit ASCII input. We specifically investigate whether improving the ability to *construct* spatial layouts (via ASCII) enhances reasoning transfer on $\mathcal{R}_{Desc}$ without direct supervision, and whether these gains generalize to external benchmarks, such as STEP GAME (Shi et al., 2022), under

Model	$Desc \rightarrow ASCII$	$ASCII \rightarrow Desc$
QWEN2.5-7B	16.1	38.4
LLAMA-3-70B	43.9	57.4
GPT-OSS-20B	44.0	58.1
GPT-4.1	53.6	72.5
QWEN3-235B	46.0	59.7
QWEN3-30B	29.2	55.8
QWEN3-30B FT $_{D \rightarrow A}$	67.9	53.5
QWEN3-30B FT $_{D \leftarrow A}$	24.1	66.5
QWEN3-30B FT $_{D \leftrightarrow A}$	64.6	68.4

Table 1: Performance on bidirectional mapping between ASCII and description. Reported scores are the average percentage of correctly aligned representations across all test samples for mapping textual descriptions to ASCII ($Desc \rightarrow ASCII$) and for identifying descriptions from ASCII ($ASCII \rightarrow Desc$).

a zero-shot prompting setting. Our hypothesis is that the “Learning to Draw” process fosters a generalized spatial understanding.

5 Results and Discussion

5.1 Benchmarking

ASCII Understanding and Generation Table 1 presents the results for the $Desc \leftrightarrow ASCII$ conversion tasks. We identify a consistent asymmetry across all models, where performance on the interpretation task ($ASCII \rightarrow Desc$) significantly exceeds that of the generation task ($Desc \rightarrow ASCII$). This suggests that while models are relatively proficient at “reading” or identifying patterns in ASCII, they struggle significantly more with “drawing” it correctly. Even the reasoning model GPT-OSS-20B fails to construct accurate layouts in over 50% of cases, despite robust interpretative success. We attribute this asymmetry to the inherent difficulty of grounding abstract semantic relations into concrete 2D layouts, a generative process demanding global planning capabilities that standard pre-training does not explicitly cultivate.

Spatial Reasoning Performance Table 2 reports accuracy scores in spatial reasoning tasks ($\mathcal{R}$) under different settings. While GPT-OSS-20B leads in overall performance and QWEN2.5-7B trails, a consistent pattern emerges across all models regarding input format.

When ASCII serves as a provided input signal, it significantly boosts performance. Most models reason more accurately from ASCII alone ($\mathcal{R}_{ASCII}$) than from text descriptions ($\mathcal{R}_{Desc}$), and the combination of both ($\mathcal{R}_{Desc+ASCII}$) yields the highest accuracy. This demonstrates that as an input, ASCII

Model	$\mathcal{R}_{\text{Desc}}$	$\mathcal{R}_{\text{ASCII}}$	$\mathcal{R}_{\text{Desc(Ans, ASCII)}}$		$\mathcal{R}_{\text{Desc(ASCII, Ans)}}$		$\mathcal{R}_{\text{Desc+ASCII}}$
	Ans	Ans	Ans	ASCII	Ans	ASCII	Ans
<i>One-turn generation</i>							
QWEN2.5-7B	42.0	39.9	<u>43.6</u>	12.5	36.7	12.0	55.5
LLAMA-3-70B	71.7	<u>75.3</u>	71.1	42.2	66.9	41.7	78.4
GPT-4.1	85.8	<u>90.1</u>	86.5	50.6	83.8	50.5	92.8
GPT-OSS-20B	93.3	<u>92.7</u>	<u>94.3</u>	38.9	91.9	42.8	97.0
QWEN3-235B-A22B	84.8	<u>89.0</u>	84.2	45.3	82.3	45.4	90.1
QWEN3-30B-A3B	75.0	<u>83.2</u>	71.1	27.9	66.6	27.5	83.7
QWEN3-30B-A3B FT _{Desc→ASCII}	82.2	78.6	<u>82.5</u>	64.2	79.1	64.7	84.9
QWEN3-30B-A3B FT _{Desc←ASCII}	73.7	87.4	68.0	22.0	70.2	20.0	87.2
QWEN3-30B-A3B FT _{Desc↔ASCII}	82.7	81.3	82.9	60.2	<u>83.1</u>	61.0	86.2
<i>Two-turn generation</i>							
QWEN2.5-7B			<u>43.6</u>	13.5	42.1	13.7	52.3
LLAMA-3-70B			<u>70.9</u>	42.8	69.0	42.3	78.9
GPT-4.1			<u>85.9</u>	50.8	81.0	50.6	93.0
GPT-OSS-20B			<u>93.1</u>	45.0	87.7	47.9	94.3
QWEN3-235B-A22B			<u>85.6</u>	45.1	77.8	45.3	89.0
QWEN3-30B-A3B			<u>74.3</u>	27.4	69.3	26.4	82.6
QWEN3-30B-A3B FT _{Desc→ASCII}			80.7	65.1	<u>81.3</u>	64.7	83.7
QWEN3-30B-A3B FT _{Desc←ASCII}			<u>70.1</u>	22.1	<u>65.2</u>	20.2	87.7
QWEN3-30B-A3B FT _{Desc↔ASCII}			<u>83.1</u>	61.4	82.7	61.8	87.5

Table 2: Spatial reasoning results across reasoning sources and settings. The $\mathcal{R}_{\text{Desc(Ans, ASCII)}}$ and $\mathcal{R}_{\text{Desc(ASCII, Ans)}}$ representations denote generation orders in which the final answer precedes or follows the ASCII representation. “One-turn” refers joint generation, while “two-turn” indicates sequential generation in separate inference calls. In the two-turn $\mathcal{R}_{\text{Desc+ASCII}}$ setting, gold-standard ASCII is provided. Answer accuracy is determined by label matching; ASCII accuracy is measured as the proportion of correctly depicted relations in the generated ASCII output. For each language model, we **bold the best-performing reasoning source** and underline the second best.

serves as a critical signal that resolves spatial ambiguities in natural language.

When ASCII is placed as an intermediate generation step (“sketching”), the results are not showing significant improvement. In both *one-turn* and *two-turn* settings, generating the layout prior to prediction ($\mathcal{R}_{\text{Desc(ASCII, Ans)}}$) consistently yields lower accuracy than generating the answer first ($\mathcal{R}_{\text{Desc(Ans, ASCII)}}$). Notably, this approach degrades performance below the text-only baseline ($\mathcal{R}_{\text{Desc}}$). Unlike textual CoT (Wei et al., 2022), where intermediate steps are difficult to verify, ASCII allows us to directly observe that flawed spatial construction propagates errors to subsequence reasoning. The benefit of explicit intermediates depends critically on their accuracy.

The specific drop in answer accuracy implies that **generated ASCII representations are insufficiently accurate to support spatial reasoning**, serving instead as misleading information. This mirrors prior findings suggesting a trade-off between generating supporting artifacts and task accuracy (Camburu et al., 2018; Huang et al., 2023). However, unlike the unstructured text in previous studies, our intermediate artifact is verifiable. This

distinction allows us to pinpoint the source of degradation rather than treating the performance loss as a black-box effect.

Finally, the performance degradation caused by generating ASCII first persists in the *two-turn* generation setting, and does not appear to be amplified relative to the *one-turn* case. This pattern motivates a deeper interpretability analysis. If generating ASCII does not improve performance and can even degrade it, does the resulting representation actually support the final spatial reasoning answer, or is it largely unrelated to the model’s prediction?

5.2 Consistency Analysis in Spatial Reasoning

Table 3 shows the alignment between ASCII representations and textual answers across both output orders and turn settings. We report three key metrics: (1) Answer Accuracy measures overall task performance; (2) Consistency quantifies the logical alignment between the generated ASCII representation and the textual answer; and (3) Consistent Accuracy evaluates correctness specifically on instances where the representation and answer correspond.

In the *one-turn* generation setting, producing the

ASCII representation before the answer does not improve consistency. Similar consistency across output orders suggests that the model does not clearly use the ASCII as a reasoning scaffold, with both outputs likely reflecting the same underlying representation.

In the *two-turn* generation setting, we observed an improved consistency rate, where the model first generates either the answer or the ASCII representation, and is then prompted to produce the other. The consistency gap between $\mathcal{R}_{\text{Desc}:(\text{Ans}, \text{ASCII})}$ and $\mathcal{R}_{\text{Desc}:(\text{ASCII}, \text{Ans})}$ indicates that the response to the follow-up query is strongly influenced by the previously generated output. Despite this notable improvement in consistency, answer accuracy increases only marginally. This limited gain may be partly due to the suboptimal quality of the generated ASCII representations, which correlates with lower answer accuracy (Table 2).

Nevertheless, consistency remains a valuable indicator, particularly given that it operates without requiring ground-truth annotations for either the spatial layout or the final answer. While prior work has leveraged Natural Language Inference (NLI) to assess reasoning via consistency (Wang et al., 2022; Lyu et al., 2023; Kumar and Talukdar, 2020), our approach benefits from the explicit verifiability of the generated artifact. Unlike unstructured text, ASCII representations enforce strict spatial constraints, allowing us to objectively assess the alignment between the representation and the answer. Although faulty representations can occur, we observe that accuracy on consistent instances significantly exceeds the accuracy over all instances. This demonstrates that the internal coherence between the verifiable layout and the text serves as a powerful unsupervised signal for reliability, enabling the identification of high-quality answers even in the absence of external labels.

From a forward-looking perspective, our results shows “ASCII as a Verifier” as a promising direction for post-CoT reasoning research. The structural constraints imposed by ASCII representations suggest that cross-representation consistency may provide a more reliable signal of reasoning validity than text-only self-consistency. Future work may explore integrating such structured verification signals into reasoning pipelines, both as an error-detection mechanism and as a guiding objective for improving intermediate representation construction.

Model	$\mathcal{R}_{\text{Desc}:(\text{Ans}, \text{ASCII})}$			$\mathcal{R}_{\text{Desc}:(\text{ASCII}, \text{Ans})}$		
	Ans	Cons	C_Acc	Ans	Cons	C_Acc
<i>One-turn generation</i>						
QWEN2.5-7B	43.6	23.0	65.2	36.7	20.5	64.4
LLAMA-3-70B	71.1	65.2	87.6	66.9	65.4	83.8
GPT-OSS-20B	94.3	75.9	97.0	91.9	75.5	96.9
QWEN3-235B	84.2	67.8	95.0	82.3	72.5	89.9
QWEN3-30B	71.1	46.1	88.5	66.6	46.3	82.5
QWEN3-30B FT _{D→A}	82.5	73.2	97.5	79.1	71.9	96.5
QWEN3-30B FT _{D←A}	68.0	39.2	84.9	70.2	46.4	81.2
QWEN3-30B FT _{D↔A}	82.9	73.7	97.4	83.1	72.9	96.6
<i>Two-turn generation</i>						
QWEN2.5-7B	43.6	24.4	69.3	42.1	31.1	56.9
LLAMA-3-70B	70.9	63.9	89.9	69.0	71.3	81.9
GPT-OSS-20B	93.1	71.8	96.7	87.7	77.2	91.6
QWEN3-235B	85.6	70.6	93.2	77.8	78.4	84.1
QWEN3-30B	74.3	50.8	88.8	69.3	58.6	75.7
QWEN3-30B FT _{D→A}	80.7	74.1	98.1	81.3	74.9	96.7
QWEN3-30B FT _{D←A}	70.1	42.8	85.7	65.2	55.7	69.7
QWEN3-30B FT _{D↔A}	83.1	73.8	96.9	82.7	76.1	94.2

Table 3: Answer accuracy (Ans), consistency between the generated ASCII and the generated answer (Cons), and accuracy on consistent instances across different settings (C_Acc). For each model, the higher Cons value is in **bold**.

Model	STEPGAME
QWEN3-30B-A3B	41.1
QWEN3-30B-A3B FT _{DESC→ASCII}	44.9
QWEN3-30B-A3B FT _{DESC←ASCII}	39.4
QWEN3-30B-A3B FT _{DESC↔ASCII}	40.4

Table 4: STEPGAME results under zero-shot prompting. The best performing system is shown in **bold**.

5.3 Learning Spatial Relations with ASCII

Table 1 details the performance of conversion tasks after unidirectional and bidirectional fine-tuning. Specifically, *Desc*→*ASCII* accuracy increases from 29.2% to 67.9%, whereas *Desc*←*ASCII* improves from 55.8% to 66.5%. This asymmetry improvement supports our earlier claim that the primary bottleneck lies in ASCII representation *generation* rather than ASCII *understanding*. Moreover, this gap between generation and understanding further widens after fine-tuning on a unidirectional task: as performance on the targeted task improves significantly, the non-targeted task inevitably degrades slightly. In contrast, joint fine-tuning on both ASCII generation and understanding (*Desc*↔*ASCII*) leads to a more balanced capability, yielding an improvement to 64.6% in ASCII generation and an improvement to 68.4% in ASCII understanding.

We examine whether balanced or targeted fine-tuning strategies yield better spatial capability in

QWEN3-30B. Table 2 reports the results on the unseen spatial reasoning task ($\mathcal{R}_{\text{Desc}}$) alongside the supervised objectives to probe whether the model has internalized a generalized spatial understanding beyond simple format conversion. Compared with the base model, both $\text{FT}_{D \rightarrow A}$ and $\text{FT}_{D \leftrightarrow A}$ achieve approximately a 9% improvement in accuracy on the description-only reasoning task ($\mathcal{R}_{\text{Desc}}$) with a similarly narrow margin. This indicates that **learning ASCII representations generalizes well and improves spatial reasoning even when ASCII is not explicitly provided**. In contrast, fine-tuning solely on ASCII understanding exacerbates the asymmetry, resulting in a slight drop in answer accuracy compared with the base model.

For $\mathcal{R}_{\text{Desc}}(\text{Ans}, \text{ASCII})$ and $\mathcal{R}_{\text{Desc}}(\text{ASCII}, \text{Ans})$, models fine-tuned on ASCII generation demonstrate significant improvements in ASCII accuracy. Consequently, we observe a reversal in the relationship between generation and reasoning: whereas low-quality layouts in other models hinder answer accuracy, the high-quality intermediate representations here provide a valid scaffold that meaningfully contributes to the final prediction. This suggests that a reliable intermediate state is a prerequisite for improving downstream reasoning.

In tasks $\mathcal{R}_{\text{ASCII}}$ and $\mathcal{R}_{\text{Desc+ASCII}}$ where ASCII is provided as input, fine-tuning on ASCII understanding generally improves answer accuracy and achieves the best performance among all other training settings. This further suggests that **ASCII representations provide a stronger reasoning substrate than natural language descriptions alone**. This mirrors human reasoning. While humans are capable of inferring spatial relationships from textual descriptions, they make substantially fewer errors when a complete visual layout is externally provided, rather than reconstructed mentally. We observe the same pattern in our human validation study (Appendix I), where participants achieve higher accuracy and greater consistency when solving spatial tasks with access to ASCII layouts compared to text-only descriptions.

Table 3 shows that fine-tuning on ASCII generation improves consistency much more than fine-tuning on ASCII understanding. We hypothesize that training solely on interpretation (Understanding) while retaining flawed output capabilities (Generation) exacerbates internal conflict. Specifically, the model becomes hyper-sensitive to discrepancies between its own noisy “sketch” and the textual description. In contrast, fine-tuning on

generation resolves this conflict at the source. By ensuring the intermediate representation is accurate and aligns with the text, the model creates more reliable reference. Consequently, it results in higher consistency, and consistent instances are more likely to be correct.

Building on the strong performance observed on our TEXT2SPACE dataset, we evaluate models on the STEPGAME benchmark (Shi et al., 2022) to test generalization. We observe that training models to generate ASCII representations yields higher accuracy on this external task (Table 4). This suggests that the core bottleneck lies in representation construction, and explicitly supervising structured intermediate representations helps models internalize abstract spatial constraints, enabling transfer to broader tasks.

6 Conclusion

In this paper, we investigate spatial reasoning in large language models through the lens of explicit representation, ASCII layouts. We introduce TEXT2SPACE, a benchmark that uses ASCII layouts to disentangle representation failures from reasoning failures. Our results reveal a clear read–write asymmetry: while models can reliably interpret spatial information from structured ASCII inputs, they struggle to reconstruct such representations from text, and these errors propagate to downstream reasoning.

Additionally, our consistency analysis demonstrates that the alignment between a model’s generated sketch and its final answer serves as a robust unsupervised indicator of correctness. We find that accuracy is significantly higher on instances where the intermediate representation remains consistent with the textual output, even in the absence of ground truth labels.

We further show that reasoning accuracy is primarily constrained by the quality of intermediate spatial representations. Crucially, fine-tuning on Text $\leftrightarrow$ ASCII alignment improves spatial reasoning from natural language even without explicit sketching at inference time, suggesting that the training process instills spatial understanding that generalizes beyond the ASCII format itself. Together, these findings demonstrate that learning to construct spatial representations offers a promising path toward more robust spatial cognition in language models.

Limitations

This work focuses on spatial reasoning tasks that can be naturally expressed using grid-based ASCII representations. Alternative structured formats, such as coordinate lists or scene graphs, are not explored in this study. Spatial problems involving continuous geometry, pixel-level visual perception, or real-world 3D environments are not explored in this work. While ASCII serves as a text-native visual proxy that leverages the inherent generation capabilities of standard language models without requiring external visual encoders, it remains a discrete abstraction and does not capture the richness of continuous spatial perception, and our synthetic TEXT2SPACE benchmark may not fully represent the noise and ambiguity in real-world scenarios.

Our evaluation relies on answer accuracy and the consistency between generated diagrams and final predictions. Although these metrics provide observable and verifiable signals of spatial reasoning behavior, they do not directly reveal the internal representations used by the model. While the explicit structure of ASCII offers a foundation for iterative verification and self-correction, this work does not explore model self-reflection or corrective mechanisms, which we leave to future research.

References

- AI@Meta. 2024. [Llama 3 model card](#).
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, and 12 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Ruth M.J Byrne and P.N Johnson-Laird. 1989. [Spatial reasoning](#). *Journal of Memory and Language*, 28(5):564–575.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-snli: Natural language inference with natural language explanations. *Advances in Neural Information Processing Systems*, 31.
- Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. 2024. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14455–14465.
- An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang, and Sifei Liu. 2024. Spatialrgpt: Grounded spatial reasoning in vision-language models. *Advances in Neural Information Processing Systems*, 37:135062–135093.
- Anthony G Cohn and Jose Hernandez-Orallo. 2023. Dialectical language model evaluation: An initial appraisal of the commonsense spatial reasoning abilities of llms. *arXiv preprint arXiv:2304.11164*.
- Marc-Alexandre Côté, Ákos Kádár, Xingdi Yuan, Ben Anderson Kybartas, Tavian Barnes, Emery Fine, James Moore, Matthew J. Hausknecht, Layla El Asri, Mahmoud Adada, Wendy Tay, and Adam Trischler. 2018. Textworld: A learning environment for text-based games. In *Workshop on Computer Games*, pages 41–75. Springer.
- Antonia Creswell and Murray Shanahan. 2022. Faithful reasoning using large language models. *arXiv preprint arXiv:2208.14271*.
- Junlin Han, Shengbang Tong, David Fan, Yufan Ren, Koustuv Sinha, Philip Torr, and Filippos Kokkinos. 2025. Learning to see before seeing: Demystifying llm visual priors from language pre-training. *arXiv preprint arXiv:2509.26625*.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. 2024. Training large language models to reason in a continuous latent space. *arXiv preprint arXiv:2412.06769*.

717	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan	Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang,	773
718	Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and	Delip Rao, Eric Wong, Marianna Apidianaki, and	774
719	Weizhu Chen. 2021. Lora: Low-rank adaptation of	Chris Callison-Burch. 2023. Faithful chain-of-	775
720	large language models . <i>Preprint</i> , arXiv:2106.09685.	thought reasoning. In <i>The 13th International Joint</i>	776
		<i>Conference on Natural Language Processing and the</i>	777
721	Zhe Hu, Jing Li, Zhongzhu Pu, Hou Pong Chan, and	<i>3rd Conference of the Asia-Pacific Chapter of the</i>	778
722	Yu Yin. 2025. Praxis-vlm: Vision-grounded decision	<i>Association for Computational Linguistics (IJCNLP-</i>	779
723	making via text-driven reinforcement learning. <i>arXiv</i>	<i>AACL 2023)</i> .	780
724	<i>preprint arXiv:2503.16965</i> .		
725	Shiyuan Huang, Siddarth Mamidanna, Shreedhar	Andreas Madsen, Sarath Chandar, and Siva Reddy. 2024.	781
726	Jangam, Yilun Zhou, and Leilani H Gilpin. 2023.	Are self-explanations from large language models	782
727	Can large language models explain themselves? a	faithful? <i>arXiv preprint arXiv:2401.07927</i> .	783
728	study of llm-generated self-explanations. <i>arXiv</i>		
729	<i>preprint arXiv:2310.11207</i> .	Margherita Malanchini, Kaili Rimfeld, Nicholas G	784
		Shakeshaft, Andrew McMillan, Kerry L Schofield,	785
730	Alon Jacovi and Yoav Goldberg. 2020. Towards faith-	Maja Rodic, Valerio Rossi, Yulia Kavas, Philip S	786
731	fully interpretable nlp systems: How should we	Dale, Elliot M Tucker-Drob, and Robert Plomin.	787
732	define and evaluate faithfulness? <i>arXiv preprint</i>	2020. Evidence for a unitary structure of spatial	788
733	<i>arXiv:2004.03685</i> .	cognition beyond general intelligence. <i>npj Science</i>	789
		<i>of Learning</i> , 5(1):9.	790
734	Amita Kamath, Jack Hessel, and Kai-Wei Chang. 2023.	Sourab Mangrulkar, Sylvain Gugger, Lysandre De-	791
735	What's "up" with vision-language models? investi-	but, Younes Belkada, Sayak Paul, Benjamin Bossan,	792
736	gating their struggle with spatial reasoning. <i>arXiv</i>	and Marian Tietz. 2022. PEFT: State-of-the-art	793
737	<i>preprint arXiv:2310.19785</i> .	parameter-efficient fine-tuning methods. https://github.com/huggingface/peft .	794
738	Parisa Kordjamshidi, Martijn Van Otterlo, and Marie-		795
739	Francine Moens. 2011. Spatial role labeling: To-	Roshanak Mirzaee, Hossein Rajaby Faghihi, Qiang	796
740	wards extraction of spatial relations from natural	Ning, and Parisa Kordjamshidi. 2021. Spartqa: A	797
741	language. <i>ACM Transactions on Speech and Language</i>	textual question answering benchmark for spatial	798
742	<i>Processing (TSLP)</i> , 8(3):1–36.	reasoning. <i>arXiv preprint arXiv:2104.05832</i> .	799
743	Sawan Kumar and Partha Talukdar. 2020. NILE : Natu-	Nora Newcombe and Janellen Huttenlocher. 2000. <i>Mak-</i>	800
744	ral language inference with faithful natural language	<i>ing space: The development of spatial representation</i>	801
745	explanations . In <i>Proceedings of the 58th Annual</i>	<i>and reasoning</i> . MIT press.	802
746	<i>Meeting of the Association for Computational Lin-</i>		
747	<i>guistics</i> , pages 8730–8742, Online. Association for	Michael Ogezi and Freda Shi. 2025. Spare: Enhanc-	803
748	Computational Linguistics.	ing spatial reasoning in vision-language models with	804
749	Chengzu Li, Wenshan Wu, Huanyu Zhang, Yan Xia,	synthetic data. <i>arXiv preprint arXiv:2504.20648</i> .	805
750	Shaoguang Mao, Li Dong, Ivan Vulić, and Furu		
751	Wei. 2025. Imagine while reasoning in space: Mul-	OpenAI. 2025. Gpt-4.1 model documentation. https://platform.openai.com/docs/models/gpt-4.1 .	806
752	timodal visualization-of-thought. <i>arXiv preprint</i>	Accessed: 2025-12-20.	807
753	<i>arXiv:2501.07542</i> .		808
754	Fangjun Li, David C Hogg, and Anthony G Cohn. 2024.	OpenAI. 2025. gpt-oss-120b & gpt-oss-20b model card .	809
755	Advancing spatial reasoning in large language mod-	<i>Preprint</i> , arXiv:2508.10925.	810
756	els: An in-depth evaluation and enhancement us-		
757	ing the stepgame benchmark. In <i>Proceedings of</i>	Letitia Parcalabescu and Anette Frank. 2024. On mea-	811
758	<i>the AAAI Conference on Artificial Intelligence</i> , vol-	suring faithfulness or self-consistency of natural	812
759	ume 38, pages 18500–18507.	language explanations. In <i>Proceedings of the 62nd An-</i>	813
760	Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harri-	<i>annual Meeting of the Association for Computational</i>	814
761	son Edwards, Bowen Baker, Teddy Lee, Jan Leike,	<i>Linguistics (Volume 1: Long Papers)</i> , pages 6048–	815
762	John Schulman, Ilya Sutskever, and Karl Cobbe.	6089.	816
763	2023. Let's verify step by step. In <i>The Twelfth Inter-</i>		
764	<i>national Conference on Learning Representations</i> .	Steven T Piantadosi, Harry Tily, and Edward Gibson.	817
765	Fangyu Liu, Guy Emerson, and Nigel Collier. 2023.	2012. The communicative function of ambiguity in	818
766	Visual spatial reasoning. <i>Transactions of the Associ-</i>	language. <i>Cognition</i> , 122(3):280–291.	819
767	<i>ation for Computational Linguistics</i> , 11:635–651.		
768	Li Liu, Diji Yang, Sijia Zhong, Kalyana Suma Sree	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Jun-	820
769	Tholeti, Lei Ding, Yi Zhang, and Leilani Gilpin. 2024.	Mei Song, Mingchuan Zhang, Y. K. Li, Yu Wu, and	821
770	Right this way: Can vlms guide us to see more to	Daya Guo. 2024. Deepseekmath: Pushing the limits	822
771	answer questions? <i>Advances in Neural Information</i>	of mathematical reasoning in open language models .	823
772	<i>Processing Systems</i> , 37:132946–132976.	<i>ArXiv</i> , abs/2402.03300.	824

- Zhengxiang Shi, Qiang Zhang, and Aldo Lipani. 2022. [Stepgame: A new benchmark for robust multi-hop spatial reasoning in texts](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11321–11329.
- Chandan Singh, Jeevana Priya Inala, Michel Galley, Rich Caruana, and Jianfeng Gao. 2024. Rethinking interpretability in the era of large language models. *arXiv preprint arXiv:2402.01761*.
- Yuda Song, Hanlin Zhang, Carson Eisenach, Sham Kakade, Dean Foster, and Udaya Ghai. 2024. Mind the gap: Examining the self-improvement capabilities of large language models. *arXiv preprint arXiv:2412.02674*.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- Qwen Team. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Galou  dec. 2020. Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl>.
- Jiayu Wang, Yifei Ming, Zhenmei Shi, Vibhav Vineet, Xin Wang, Sharon Li, and Neel Joshi. 2024. Is a picture worth a thousand words? delving into spatial reasoning for vision language models. *Advances in Neural Information Processing Systems*, 37:75392–75421.
- Sophie L Wang, Phillip Isola, and Brian Cheung. 2025. Words that make language models perceive. *arXiv preprint arXiv:2510.02425*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- 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. *Advances in neural information processing systems*, 35:24824–24837.
- Yana Wei, Liang Zhao, Jian-Yuan Sun, Kangheng Lin, Jisheng Yin, Jingcheng Hu, Yinmin Zhang, En Yu, Haoran Lv, Zejia Weng, Jia Wang, Chunrui Han, Yuang Peng, Qi Han, Zheng Ge, Xiangyu Tony Zhang, Daxin Jiang, and Vishal M. Patel. 2025. Open vision reasoner: Transferring linguistic cognitive behavior for visual reasoning. *arXiv preprint arXiv:2507.05255*.
- Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart Van Merri  nboer, Armand Joulin, and Tomas Mikolov. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.
- Wenshan Wu, Shaoguang Mao, Yadong Zhang, Yan Xia, Li Dong, Lei Cui, and Furu Wei. 2024. Mind’s eye of llms: visualization-of-thought elicits spatial reasoning in large language models. *Advances in Neural Information Processing Systems*, 37:90277–90317.
- Yutaro Yamada, Yihan Bao, Andrew K Lampinen, Jungo Kasai, and Ilker Yildirim. 2023. Evaluating spatial understanding of large language models. *arXiv preprint arXiv:2310.14540*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 40 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822.
- Hongjie Zhang, Hourui Deng, Jie Ou, and Chaosheng Feng. 2025a. Mitigating spatial hallucination in large language models for path planning via prompt engineering. *Scientific Reports*, 15(1):8881.
- Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. 2025b. The lessons of developing process reward models in mathematical reasoning. *arXiv preprint arXiv:2501.07301*.

Appendix

A Dataset Details

Table 5 summarizes the possible values for each field in TEXT2SPACE. The dataset is in English, covering various terminologies including cardinal directions (e.g., ‘North’), clock-face references (e.g., ‘3 o’clock’), and standard spatial prepositions (e.g., ‘left of’). Figure 3 provides an overview of the full dataset statistics, and Figure 4 shows the corresponding statistics for the test set.

Field	Possible Values
query_type	full, vertical, horizontal
label (full)	above, below, left, right, upper-left, upper-right, lower-left, lower-right
label (vertical)	above, below, same level
label (horizontal)	left, right, same column
terminology_used	spatial, cardinal, clock, spatial+cardinal, spatial+clock, cardinal+clock, spatial+cardinal+clock
ascii	simple, grid, panel
num_components	2–8
num_relations	1–12
is_directly_stated	true, false
has_unique_layout	true, false
ambiguous_stages	0–11

Table 5: TEXT2SPACE Dataset Schema, including spatial query types, terminology variations, and structural constraints (components and relations).

B ASCII Format Choice

Using three different ASCII representations as input, we conduct two experiments to determine the most suitable format. The first experiment evaluates the interpretation of natural language descriptions from ASCII ($ASCII \rightarrow Desc$), while the second assesses spatial reasoning directly from ASCII ($\mathcal{R}_{ASCII}$). As shown in Table 6, performance on $ASCII \rightarrow Desc$ is comparable across all three formats. In contrast, for $\mathcal{R}_{ASCII}$, the grid-based representation yields substantially higher performance. Accordingly, we adopt the grid format in all subsequent experiments.

ASCII Format	QWEN3-30B-A3B
$ASCII \rightarrow Desc$	
simple	57.1
grid	55.8
panel	54.4
$\mathcal{R}_{ASCII}$	
simple	76.8
grid	83.2
panel	79.2

Table 6: ASCII Format Result

C Training configurations

Configuration	Value
Base Model	QWEN3-30B-A3B
Frameworks	HuggingFace Transformers, TRL, PEFT
Fine-tuning Method	LoRA (Rank $r = 16$, $\alpha = 32$)
Target Modules	$q, k, v, o, gate, up, down$ projections
Precision	bfloat16
Optimizer	AdamW
Learning Rate	2×10^{-4} (Cosine Schedule)
Batch Size	4 (per device)
Gradient Accumulation	4 steps
Epochs	1
Dataset Split	4,000 (Train) / 500 (Validation)
Hardware	4× NVIDIA A40

Table 7: Training configuration details.

D Confusion matrices for spatial reasoning predictions

Figure 5 presents the confusion matrices for spatial reasoning predictions before and after fine-tuning on the ASCII–description conversion tasks. Fine-tuning results in stronger concentration along the diagonal and reduces systematic confusions between spatially related classes, indicating improved internal organization of spatial representations rather than uniform scaling of prediction confidence.

E F1 Scores

Table 8 reports the F1 scores for spatial reasoning tasks. In line with Table 2, reasoning from ASCII alone generally outperforms reasoning from descriptions alone. While generating ASCII before answering hurts performance without fine-tuning, fine-tuning with explicit ASCII generation, either unidirectional ($Desc \rightarrow ASCII$) or bidirectional ($Desc \leftrightarrow ASCII$), consistently improves performance when only descriptions are available at inference time.



Figure 3: **Statistical overview of the Text2Space dataset.** The figure illustrates the distribution of components and relations per instance (top left), followed by the balanced distribution of query types and spatial terminology (top right). The bottom row details the ratio of direct to indirect relations, the proportion of unique layouts, and the distribution of ambiguous stages. The summary box (bottom right) highlights key metrics, including a total of 20,000 instances with an average of 5.7 components and 6.4 relations.



Figure 4: **Statistical overview of the Test Set.** The figure illustrates the distributions for the 1,000 held-out test instances. The dataset maintains a balanced distribution of query types and spatial terminology (top right).




Figure 5: Confusion matrices for spatial reasoning predictions by Qwen3-30B-A3B before and after fine-tuning (Qwen3-30B-A3B $\text{FT}_{\text{D} \rightarrow \text{A}}$). The top row ((A)-(B)) shows results on the text-only spatial reasoning task ($\mathcal{R}_{\text{Desc}}$), while the bottom row ((C)-(D)) shows results when ASCII is generated as an intermediate step ($\mathcal{R}_{\text{Desc}(\text{Ans}, \text{ASCII})}$). Left panels ((A),(C)) correspond to the base model, and right panels ((B),(D)) to the fine-tuned model. Rows indicate ground-truth spatial relations and columns indicate predicted relations. Darker diagonal cells indicate correct predictions.

Model	$\mathcal{R}_{\text{Desc}}$	$\mathcal{R}_{\text{ASCII}}$	$\mathcal{R}_{\text{Desc}(\text{Ans}, \text{ASCII})}$	$\mathcal{R}_{\text{Desc}(\text{ASCII}, \text{Ans})}$	$\mathcal{R}_{\text{Desc}+\text{ASCII}}$
<i>One-turn generation</i>					
QWEN2.5-7B	37.7	33.4	<u>38.6</u>	33.9	48.3
LLAMA-3-70B	<u>63.3</u>	62.2	62.3	52.2	68.8
GPT-4.1	81.2	<u>85.8</u>	82.2	77.5	89.6
GPT-oss-20B	91.9	92.0	<u>93.2</u>	90.5	96.8
QWEN3-235B-A22B	81.9	<u>86.4</u>	81.5	79.6	88.0
QWEN3-30B-A3B	72.3	<u>80.3</u>	69.6	66.2	81.3
QWEN3-30B-A3B $\text{FT}_{\text{Desc} \rightarrow \text{ASCII}}$	79.4	72.7	79.5	<u>76.9</u>	81.9
QWEN3-30B-A3B $\text{FT}_{\text{Desc} \leftarrow \text{ASCII}}$	71.4	<u>84.1</u>	67.4	69.4	84.7
QWEN3-30B-A3B $\text{FT}_{\text{Desc} \leftrightarrow \text{ASCII}}$	78.2	73.3	<u>80.0</u>	79.3	82.6
<i>Two-turn generation</i>					
QWEN2.5-7B			<u>38.6</u>	<u>38.6</u>	46.9
LLAMA-3-70B			61.0	56.2	67.2
GPT-4.1			<u>81.0</u>	76.1	90.0
GPT-oss-20B			<u>92.3</u>	86.8	94.0
QWEN3-235B-A22B			<u>83.0</u>	73.1	85.0
QWEN3-30B-A3B			<u>71.8</u>	64.1	78.8
QWEN3-30B-A3B $\text{FT}_{\text{Desc} \rightarrow \text{ASCII}}$			<u>76.2</u>	74.9	78.3
QWEN3-30B-A3B $\text{FT}_{\text{Desc} \leftarrow \text{ASCII}}$			<u>68.2</u>	57.8	82.6
QWEN3-30B-A3B $\text{FT}_{\text{Desc} \leftrightarrow \text{ASCII}}$			<u>77.3</u>	74.0	80.1

Table 8: F1-score results. The best result is shown in **bold**, and the second-best is underlined.

F Data Generation Pipeline

Pseudo-code for the spatial dataset generation pipeline with multi-dimensional balancing.

Algorithm 1 Text2Space Dataset Generation

Require:

- 1: N_{target} : Target count
- 2: R_{comp}, R_{rel} : Complexity ranges (components, relations)

Ensure:

- 3: $\mathcal{D}$: Final spatial dataset
- 4: $\mathcal{D} \leftarrow \emptyset$
- 5: Initialize *Tracker* to monitor distribution
- 6: **while** $|\mathcal{D}| < N_{target}$ **do**
1. **Select Parameters**
- 7: Identify imbalanced categories using *Tracker*
- 8: $P \leftarrow$ select parameters (e.g., node count) to balance
2. **Build Connected Structure**
- 9: $V \leftarrow \{A, B, C, \dots\}$
- 10: $G \leftarrow (V, E = \emptyset)$
- 11: $C \leftarrow \{\text{Random}(V)\}$
- 12: **while** $V \setminus C \neq \emptyset$ **do**
- 13: $u \leftarrow \text{Random}(V \setminus C)$
- 14: $v \leftarrow \text{Random}(C)$
- 15: $dir \leftarrow \text{Random}(8 \text{ Compass Directions})$
- 16: **if** edge (u, v, dir) is valid **then**
- 17: $E \leftarrow E \cup \{(u, v, dir)\}$
- 18: $C \leftarrow C \cup \{u\}$
- 19: **end if**
- 20: **end while**
- 21: Add extra relations to E to reach target complexity
3. **Generate Description**
- 22: $S \leftarrow \text{Random}(\{\text{Spatial}, \text{Cardinal}, \text{Clock}\})$
- 23: $Desc \leftarrow \text{GRAPHToNATURALLANGUAGE}(G, S)$
4. **Create Query**
- 24: $\mathcal{P}_{infer} \leftarrow$ Identify uniquely inferable pairs in G
- 25: $(Q, A) \leftarrow$ Sample pair and Type from $\mathcal{P}_{infer}$
- 26: Determine if relationship is Direct vs. Inferred
5. **Render Visualizations**
- 27: $Vis \leftarrow \text{RENDERASCII}(G, \{\text{Simple}, \text{Grid}, \text{Panel}\})$
6. **Apply Rejection Sampling**
- 28: $I \leftarrow \{G, Desc, Q, A, Vis, S\}$
- 29: $\Delta \leftarrow \text{CALCBALANCEDEVIATION}(I, \text{Tracker})$
- 30: **if** I worsens balance **then**
- 31: **Reject** with probability $P \propto \Delta^2$
- 32: **else**
- 33: $\mathcal{D} \leftarrow \mathcal{D} \cup \{I\}$
- 34: Update *Tracker* with instance I
- 35: **end if**
- 36: **end while**
- 37: **return** $\mathcal{D}$

G Evaluation Algorithm for ASCII and Descriptions

Pseudo-code for evaluating generated descriptions and ASCII layouts.

Algorithm 2 Evaluate Description vs. ASCII Layouts

Require:

- 1: G : ASCII Grid (Model Output)
- 2: D : Natural Language Description
- 3: $Mode$: "verify_ascii" or "verify_desc"

Ensure:

- 4: R : Result tuple (Pass, Accuracy, Breakdown)
1. **Parse Description**
- 5: $Segs \leftarrow \text{Split}(D, \{.,, \})$
- 6: $Claims \leftarrow \emptyset, Bad \leftarrow \emptyset$
- 7: **for** $s \in Segs$ **do**
- 8: Match s with patterns (e.g., "A is left of B")
- 9: **if** match found **then**
- 10: $(u, v, dir) \leftarrow \text{Extract}(s)$
- 11: $Claims.add((u, v, dir))$
- 12: **else**
- 13: $Bad.add(s)$
- 14: **end if**
- 15: **end for**
- 16: $Obj_{ST} \leftarrow \text{UniqueObjects}(Claims)$
2. **Build Grid Graph**
- 17: $Pos \leftarrow \text{ScanCoordinates}(G)$
- 18: Normalize Pos (top-left at 0,0)
- 19: $Truth \leftarrow \emptyset$
- 20: **for** $(u, v) \in Pos$ **do**
- 21: $\vec{d} \leftarrow Pos[u] - Pos[v]$
- 22: $Truth[u, v] \leftarrow \text{VectorToDir}(\vec{d})$
- 23: **end for**
3. **Validate Relations**
- 24: $Corr \leftarrow \emptyset, Err \leftarrow \emptyset$
- 25: **for** $(u, v, dir) \in Claims$ **do**
- 26: **if** $u \notin Pos$ or $v \notin Pos$ **then**
- 27: $Err.add((u, v, "Missing"))$
- 28: **else**
- 29: **if** $dir == Truth[u, v]$ **then**
- 30: $Corr.add((u, v, dir))$
- 31: **else**
- 32: $Err.add((u, v, Truth[u, v]))$
- 33: **end if**
- 34: **end if**
- 35: **end for**
4. **Final Scoring**
- 36: **if** $Mode == \text{"verify_desc"}$ **then** ▷ ASCII is Truth
- 37: $Pass \leftarrow (Err = \emptyset)$
- 38: **else** ▷ Text is Truth
- 39: $Extra \leftarrow Pos.keys \setminus Obj_{ST}$
- 40: $Pass \leftarrow (Err = \emptyset \wedge Bad = \emptyset \wedge Extra = \emptyset)$
- 41: **end if**
- 42: $Acc \leftarrow |Corr|/|Claims|$
- 43: **return** $\{Pass, Acc, Corr, Err, Bad\}$

H Evaluation Algorithm for Consistency

Pseudo-code for evaluating the consistency between generated ASCII layouts and query relations.

Algorithm 3 Evaluate Consistency Between ASCII Grid and Reference relations

Require:

- 1: G : ASCII representation (Model Grid)
- 2: Q : Natural language query
- 3: L : Reference label (relation)

Ensure:

- 4: R : Evaluation result (correctness, objects, relations)
1. **Parse the Query**
- 5: Match Q vs. patterns: "Vertical", "Horizontal", or "Full"
- 6: Extract target objects X and Y from Q
- 7: **if** no valid pattern found **then**
- 8: **return** ERROR("Query parsing failed")
- 9: **end if**
2. **Parse the ASCII Grid**
- 10: Scan G to find coordinates for all characters
- 11: $Pos \leftarrow \{(obj, x, y) \mid obj \text{ found in } G\}$
- 12: **if** $X \notin Pos$ or $Y \notin Pos$ **then**
- 13: **return** ERROR("Queried object missing in grid")
- 14: **end if**
- 15: Normalize Pos so top-left object is at (0, 0)
3. **Infer Relation from Grid**
- 16: $(x_1, y_1) \leftarrow Pos[X]$, $(x_2, y_2) \leftarrow Pos[Y]$
- 17: $Rel_{inf} \leftarrow \emptyset$
- 18: **if** Type is "Vertical" **then**
- 19: **if** $y_1 < y_2$ **then** $Rel_{inf} \leftarrow$ "above"
- 20: **else if** $y_1 > y_2$ **then** $Rel_{inf} \leftarrow$ "below"
- 21: **else** $Rel_{inf} \leftarrow$ "same level"
- 22: **end if**
- 23: **else if** Type is "Horizontal" **then**
- 24: **if** $x_1 < x_2$ **then** $Rel_{inf} \leftarrow$ "left"
- 25: **else if** $x_1 > x_2$ **then** $Rel_{inf} \leftarrow$ "right"
- 26: **else** $Rel_{inf} \leftarrow$ "same column"
- 27: **end if**
- 28: **else** ▷ Full Spatial Query
- 29: Compute vector $\vec{v} = (x_1 - x_2, y_1 - y_2)$
- 30: $Rel_{inf} \leftarrow$ Map $\vec{v}$ to {upper-left, below, right, ...}
- 31: **end if**
4. **Compare with Reference Label**
- 32: $Match \leftarrow (\text{LowerCase}(Rel_{inf}) == \text{LowerCase}(L))$
- 33: $R \leftarrow \{\text{Pass} : Match, \text{Exp} : L, \text{Act} : Rel_{inf}\}$
- 34: **return** R

I Human Validation

Human Validation Setup. To quantify the impact of different input formats on spatial reasoning accuracy, we conducted a controlled human validation study with 11 evaluators divided into two groups: Team A ($n = 7$) and Team B ($n = 4$). We curated three distinct problem sets (Set X, Set Y, and Set Z), each containing 10 spatially complex examples. The design details is shown in Table 9

The study utilized a counterbalanced within-subjects design to disentangle the effect of input format from question difficulty. All participants completed three sessions:

1. **Session 2 (Visual Baseline):** Both teams eval-

uated Set Z using the **ASCII-only** format. Performance was highly consistent across groups (Team A: 0.986, Team B: 0.950), confirming that both teams possessed comparable visual reasoning baselines.

2. **Sessions 1 & 3 (Crossover):** We swapped the assignment of problem sets and formats between teams. Team A solved Set X using **Description-only** (Session 1) and Set Y using **Description+ASCII** (Session 3). Conversely, Team B solved Set Y using **Description-only** (Session 1) and Set X using **Description+ASCII** (Session 3).

This cross-over design ensures that Sets X and Y were evaluated in both text-only and combined formats, allowing for a direct comparison of performance on semantically identical questions.

Task Sessions	Team A ($n = 7$)	Team B ($n = 4$)
Session 1 (Single Format)	Description Only (Set X)	Description Only (Set Y)
Session 2 (Visual Baseline)	ASCII Only (Set Z)	ASCII Only (Set Z)
Session 3 (Combined)	Desc + ASCII (Set Y)	Desc + ASCII (Set X)

Table 9: Human Validation Study Design. The experiment followed a three-session protocol with counterbalanced question sets.

Validation Results Table 10 shows the substantial improvement achieved when reasoning from an ASCII-augmented description compared to reasoning from the description alone.

Question Set	Desc Only (NL Format)	Desc + ASCII (Combined)	Improvement
Set X	0.857 (Team A)	0.975 (Team B)	+11.8%
Set Y	0.875 (Team B)	0.986 (Team A)	+11.1%
Average	0.866	0.981	+11.5%

Table 10: Human performance comparison across different input formats. By counterbalancing Sets X and Y, we isolate the improvement attributed to the format.

Participants and Ethics All participants were paid at an hourly rate consistent with local labor standards. The research involves benign behavioral tasks with no collection of personally identifiable information (PII) and currently under review for an IRB exemption.

Please answer by choosing the most appropriate option from the drop-down menu. Specifically, if A is upper-right of B, then "upper-right" is a more accurate answer than simply "above" or "right."

Important: Each question is independent of the others.

You will be given a text description, an ASCII representation, and a query.

- In the first part, you will see only the text description and the query.
- In the second part, you will see only the ASCII representation and the query.
- In the third part, you will see all three components together.

Task format:

Your question zone

Example 1:
Text Description:
 E is at the 9:00 position from A. B is at the 6 o'clock position from A. C is at 12 o'clock from A. F is at the 1 o'clock position from B. D is at 4:30 from B. A is at the 9:00 position from F. F is at 12 o'clock from D.

ASCII Representation:

```

  C
 | E A F |
 | B     |
 | D     |
  
```

Query:
 Where is E relative to B? (E->B)

Your answer zone

Choose

Choose

Above ☒

Below

Right

Left ☒

Upper-Right

Upper-Left ☒

Lower-Right

Lower-Left

Task format:

Horizontal directions

Choose

Choose

Right

Left

Same Column (Horizontal-only relations)

Vertical directions

Choose

Choose

Above

Below

Same Level (Vertical-only relations)

Full directions

Choose

Choose

Above

Below

Right

Left

Upper-Right

Upper-Left

Lower-Right

Lower-Left

Text Description + Query

Session 1

Example 1:
Text Description:
 E is at the 9:00 position from A. B is at the 6 o'clock position from A. C is at 12 o'clock from A. F is at the 1 o'clock position from B. D is at 4:30 from B. A is at the 9:00 position from F. F is at 12 o'clock from D.

Query:
 Where is B relative to C? (B->C)

Choose

ASCII visualization + Query

Session 2

Example 1:

```

  C
 | E A F |
 | B     |
 | D     |
  
```

Query:
 Where is B relative to C? (B->C)

Choose

Text + ASCII visualization + Query

Session 3

Example 1:
Text Description:
 E is at the 9:00 position from A. B is at the 6 o'clock position from A. C is at 12 o'clock from A. F is at the 1 o'clock position from B. D is at 4:30 from B. A is at the 9:00 position from F. F is at 12 o'clock from D.

```

  C
 | E A F |
 | B     |
 | D     |
  
```

Query:
 Where is B relative to C? (B->C)

Choose

Figure 6: Overview of the human validation interface and instruction protocol. Participants solved spatial reasoning tasks across three counterbalanced sessions: (1) Text Description Only, (2) ASCII Visualization Only, and (3) Combined Text + ASCII. We use Google Form to collect the response from the participants. The interface requires participants to select the most precise directional term from a standardized drop-down menu.

J Prompts and Outputs

J.1 Task Description → ASCII

SYSTEM

Assume you have expertise in spatial reasoning. Your task is to answer the following question to the best of your knowledge.
You will receive a scene description. Your job is to generate the requested ASCII visualization and return the result as a compact JSON object.

Description -> ASCII Diagram.
Read the textual scene description and render the requested ASCII layout.

- Treat the canvas like a map: the first printed row is the top, and the last row is the bottom.
- The first column is the left edge; the last column is the right edge.
- Above = higher row; below = lower row (do not invert the vertical axis).
- Left = smaller column index; right = larger column index (do not mirror horizontally).

Formats:
- grid: <grid using + - | for borders>

Rendering & Positioning Rules

Use given uppercase letters (A, B, C, ...); don't invent new ones.
Grid size is flexible - expand rows/cols as needed to fit all points.
Place all letters; if two overlap, stack them (e.g., "AF").

Relations are first -> second (e.g., "D to F" means D relative to F).

Positioning:
above/below -> same column
left/right -> same level
diagonals (upper/lower-left/right) -> one cell diagonally.
Keep perfect alignment for cardinal directions.
Adjacent items occupy neighboring cells; no diagonal drift.
Maintain uniform cell sizes and consistent borders.

Keep row alignment when depicting left/right relations; do not shift labels vertically.
Keep column alignment when depicting above/below relations; do not shift labels horizontally.
For diagonal relations, ensure the two entities occupy different rows and different columns.
Avoid inconsistent cell widths and do not introduce extra or missing letters.

Additional ASCII Guidance:
- Label every entity exactly as named in the description.

Input Notes:
- Use only the provided description to infer entity positions.

- Do not invent additional entities or relations when information is missing.

Output Notes:

- Produce ASCII views that faithfully match the described scene.

OUTPUT SCHEMA (JSON):

```
{
  "ascii": {
    "grid": "<grid using + - | for borders>"
  }
}
```

Return ONLY a compact JSON object.
Do not include markdown fences or any extra text.

USER

Description:

B is to the lower-right of A. D is to the right of B. C is at the bottom-right of A. C is located below B.

ASSISTANT

```
{"ascii":{"grid":"+++++\\n|A | |\\n+++++\\n| |B |D |\\n+++++\\n|C | |\\n+++++\\n"}}}
```

J.2 Task ASCII → Description

SYSTEM

Assume you have expertise in spatial reasoning. Your task is to answer the following question to the best of your knowledge.

You will receive ASCII diagram(s). Your job is to write a concise spatial summary and return the result as a compact JSON object.

ASCII Diagram -> Description.

Interpret the ASCII evidence and rewrite it as a concise natural-language description.

- Treat the canvas like a map: the first printed row is the top, and the last row is the bottom.
- The first column is the left edge; the last column is the right edge.
- Above = higher row; below = lower row (do not invert the vertical axis).
- Left = smaller column index; right = larger column index (do not mirror horizontally).

Formats:

- grid: <grid using + - | for borders>

Input Notes:

- Treat the ASCII diagram as authoritative evidence.

Output Notes:

- Summarize the spatial arrangement in clear declarative sentences.

Description Guidelines:

- Mention every entity shown in the ASCII input and describe their relative

1159	placement.	horizontal offset: upper-left, upper-	1235
1160	- Keep the description short and factual	right, lower-left, lower-right	1236
1161	; avoid speculative language.	Horizontal Alignment ONLY: same level;	1237
1162	- The spatial summary will consist of	Vertical Alignment ONLY: same column	1238
1163	sentences in the format "A is [relation]		1239
1164	B." Each sentence should describe the	Input Notes:	1240
1165	relationship between exactly two objects	- Use the description to reason about	1241
1166	.	the positions of the queried entities.	1242
1167			1243
1168	OUTPUT SCHEMA (JSON):	Output Notes:	1244
1169	{	- Return the spatial relation as a	1245
1170	"description": "<natural language	single direction string.	1246
1171	spatial summary>"		1247
1172	}	Direction Notes:	1248
1173		- Ensure the answer directly references	1249
1174	Return ONLY a compact JSON object.	the entities named in the question.	1250
1175	Do not include markdown fences or any		1251
1176	extra text.	OUTPUT SCHEMA (JSON):	1252
1177		{	1253
1178	-----	"answer": "<direction>"	1254
1179	USER	}	1255
1180	-----		1256
1181	ASCII Reference:	Return ONLY a compact JSON object.	1257
1182	grid:	Do not include markdown fences or any	1258
1183	+---+---+	extra text.	1259
1184	A		1260
1185	+---+---+	-----	1261
1186	B D	USER	1262
1187	+---+---+	-----	1263
1188	C	Description:	1264
1189	+---+---+	B is to the lower-right of A. D is to	1265
1190		the right of B. C is at the bottom-right	1266
1191	-----	of A. C is located below B.	1267
1192	ASSISTANT		1268
1193	-----	Question: Where is C horizontally	1269
1194	{"description": "B is to the lower-right	relative to D?	1270
1195	of A. D is to the right of B. C is at		1271
1196	the bottom-right of A. C is located	-----	1272
1197	below B."}	ASSISTANT	1273
		-----	1274
		{"answer": "left"}	1275
1198	J.3 Task $\mathcal{R}_{Desc}$		
1199	-----	J.4 Task $\mathcal{R}_{ASCII}$	1276
1200	SYSTEM	-----	1277
1201	-----	SYSTEM	1278
1202	Assume you have expertise in spatial	-----	1279
1203	reasoning. Your task is to answer the	Assume you have expertise in spatial	1280
1204	following question to the best of your	reasoning. Your task is to answer the	1281
1205	knowledge.	following question to the best of your	1282
1206	You will receive a directional question	knowledge.	1283
1207	and a scene description.	You will receive a directional question	1284
1208	Your job is to produce the correct	and ASCII diagram(s).	1285
1209	directional answer and return the result	Your job is to produce the correct	1286
1210	as a compact JSON object.	directional answer and return the result	1287
1211		as a compact JSON object.	1288
1212	Note that the possible directions are		1289
1213	defined as {left, right, above, below,	Note that the possible directions are	1290
1214	upper-left, upper-right, lower-left,	defined as {left, right, above, below,	1291
1215	lower-right, same column, same level},	upper-left, upper-right, lower-left,	1292
1216	and there will be only one correct	lower-right, same column, same level},	1293
1217	answer.	and there will be only one correct	1294
1218		answer.	1295
1219	Description + Query -> Answer.		1296
1220	Answer the spatial relationship question	ASCII + Query -> Answer.	1297
1221	using the textual description.	Answer the spatial relationship question	1298
1222		using the ASCII diagram as the primary	1299
1223	10 Distinct (Mutually Exclusive)	evidence.	1300
1224	Directions		1301
1225	Pure Directions - perfectly aligned:	10 Distinct (Mutually Exclusive)	1302
1226	- left: directly left, no vertical	Directions	1303
1227	offset	Pure Directions - perfectly aligned:	1304
1228	- right: directly right, no vertical	- left: directly left, no vertical	1305
1229	offset	offset	1306
1230	- above: directly above, no horizontal	- right: directly right, no vertical	1307
1231	offset	offset	1308
1232	- below: directly below, no vertical	- above: directly above, no horizontal	1309
1233	offset	offset	1310
1234	Diagonal Directions - both vertical and		

1311 - below: directly below, no vertical
1312 offset
1313 Diagonal Directions - both vertical and
1314 horizontal offset: upper-left, upper-
1315 right, lower-left, lower-right
1316 Horizontal Alignment ONLY: same level;
1317 Vertical Alignment ONLY: same column
1318
1319 - Treat the canvas like a map: the first
1320 printed row is the top, and the last
1321 row is the bottom.
1322 - The first column is the left edge; the
1323 last column is the right edge.
1324 - Above = higher row; below = lower row
1325 (do not invert the vertical axis).
1326 - Left = smaller column index; right =
1327 larger column index (do not mirror
1328 horizontally).
1329
1330 Formats:
1331 - grid: <grid using + - | for borders>
1332
1333 Input Notes:
1334 - Use the ASCII diagram as ground truth;
1335 the query specifies the entities to
1336 compare.
1337
1338 Output Notes:
1339 - Provide a single directional answer
1340 from the defined set.
1341
1342 Direction Notes:
1343 - Base the answer on the ASCII layout;
1344 ignore assumptions that contradict the
1345 diagram.
1346
1347 OUTPUT SCHEMA (JSON):
1348 {
1349 "answer": "<direction>"
1350 }
1351
1352 Return ONLY a compact JSON object.
1353 Do not include markdown fences or any
1354 extra text.
1355
1356 -----
1357 USER
1358 -----
1359 ASCII Reference:
1360 grid:
1361 +---+---+
1362 |A| | |
1363 +---+---+
1364 | |B|D|
1365 +---+---+
1366 | |C| |
1367 +---+---+
1368
1369 Question: Where is C horizontally
1370 relative to D?
1371
1372 -----
1373 ASSISTANT
1374 -----
1375 {"answer": "left"}

1376 J.5 Task $\mathcal{R}_{\text{Desc}}(\text{Ans}, \text{ASCII})$

1377 One-turn Generation

1378 -----
1379 SYSTEM
1380 -----
1381 Assume you have expertise in spatial
1382 reasoning. Your task is to answer the
1383 following question to the best of your
1384 knowledge.

You will receive a directional question
and a scene description.
Your job is to produce the correct
directional answer and generate the
requested ASCII visualization and return
the result as a compact JSON object.

Note that the possible directions are
defined as {left, right, above, below,
upper-left, upper-right, lower-left,
lower-right, same column, same level},
and there will be only one correct
answer.

Description + Query -> ASCII + Answer.
Answer the spatial question and draw
ASCII that matches the same layout.

10 Distinct (Mutually Exclusive)

Directions
Pure Directions - perfectly aligned:
- left: directly left, no vertical
offset
- right: directly right, no vertical
offset
- above: directly above, no horizontal
offset
- below: directly below, no vertical
offset
Diagonal Directions - both vertical and
horizontal offset: upper-left, upper-
right, lower-left, lower-right
Horizontal Alignment ONLY: same level;
Vertical Alignment ONLY: same column

- Treat the canvas like a map: the first
printed row is the top, and the last
row is the bottom.
- The first column is the left edge; the
last column is the right edge.
- Above = higher row; below = lower row
(do not invert the vertical axis).
- Left = smaller column index; right =
larger column index (do not mirror
horizontally).

Formats:
- grid: <grid using + - | for borders>

Rendering & Positioning Rules

Use given uppercase letters (A, B, C,
...); don't invent new ones.
Grid size is flexible - expand rows/cols
as needed to fit all points.
Place all letters; if two overlap, stack
them (e.g., "AF").

Relations are first -> second (e.g., "D
to F" means D relative to F).

Positioning:
above/below -> same column
left/right -> same level
diagonals (upper/lower-left/right) ->
one cell diagonally.

Keep perfect alignment for cardinal
directions.
Adjacent items occupy neighboring cells;
no diagonal drift.
Maintain uniform cell sizes and
consistent borders.

Keep row alignment when depicting left/
right relations; do not shift labels
vertically.
Keep column alignment when depicting
above/below relations; do not shift

1463	labels horizontally.	Then Generate ASCII.	1540
1464	For diagonal relations, ensure the two		1541
1465	entities occupy different rows and	- Turn 1: Answer the spatial question	1542
1466	different columns.	based on reasoning.	1543
1467	Avoid inconsistent cell widths and do	- Turn 2: Generate an ASCII	1544
1468	not introduce extra or missing letters.	visualization that matches your answer.	1545
1469			1546
1470	Additional ASCII Guidance:	## Turn 1: Answer the Question	1547
1471	- Ensure the ASCII visualization aligns		1548
1472	with the answer you provide.	The possible directions are {left, right	1549
1473		, above, below, upper-left, upper-right,	1550
1474	Determine the correct answer first, then	lower-left, lower-right, same column,	1551
1475	draw ASCII that matches it.	same level}. There is only one correct	1552
1476		answer.	1553
1477	Input Notes:		1554
1478	- Use the description to understand the	10 Distinct (Mutually Exclusive)	1555
1479	scene before answering.	Directions	1556
1480	- Ensure the ASCII rendering is	Pure Directions - perfectly aligned:	1557
1481	consistent with the textual evidence.	- left: directly left, no vertical	1558
1482		offset	1559
1483	Output Notes:	- right: directly right, no vertical	1560
1484	- Provide both the answer and the	offset	1561
1485	requested ASCII views in the JSON output	- above: directly above, no horizontal	1562
1486	.	offset	1563
1487		- below: directly below, no vertical	1564
1488	Direction Notes:	offset	1565
1489	- Check the ASCII you generate so it	Diagonal Directions - both vertical and	1566
1490	supports the chosen answer.	horizontal offset: upper-left, upper-	1567
1491		right, lower-left, lower-right	1568
1492	OUTPUT SCHEMA (JSON):	Horizontal Alignment ONLY: same level;	1569
1493	{	Vertical Alignment ONLY: same column	1570
1494	"answer": "<direction>",		1571
1495	"ascii": {	## Turn 2: Generate ASCII Visualization	1572
1496	"grid": "<grid using + - for		1573
1497	borders>"	Formats:	1574
1498	}	- grid: <grid using + - for borders>	1575
1499	}		1576
1500		Rendering & Positioning Rules	1577
1501	Return ONLY a compact JSON object.		1578
1502	Do not include markdown fences or any	Use given uppercase letters (A, B, C,	1579
1503	extra text.	...); don't invent new ones.	1580
1504		Grid size is flexible - expand rows/cols	1581
1505	-----	as needed to fit all points.	1582
1506	USER	Place all letters; if two overlap, stack	1583
1507	-----	them (e.g., "AF").	1584
1508	Description:		1585
1509	B is to the lower-right of A. D is to	Relations are first -> second (e.g., "D	1586
1510	the right of B. C is at the bottom-right	to F" means D relative to F).	1587
1511	of A. C is located below B.		1588
1512		Positioning:	1589
1513	Question: Where is C horizontally	above/below -> same column	1590
1514	relative to D?	left/right -> same level	1591
1515		diagonals (upper/lower-left/right) ->	1592
1516	-----	one cell diagonally.	1593
1517	ASSISTANT	Keep perfect alignment for cardinal	1594
1518	-----	directions.	1595
1519	{"answer": "left", "ascii": {"grid	Adjacent items occupy neighboring cells;	1596
1520	": "++---+++\n A \n+---+++\n	no diagonal drift.	1597
1521	B D \n+---+++\n C \n	Maintain uniform cell sizes and	1598
1522	++---+++"}]}	consistent borders.	1599
			1600
1523	<i>Two-turn Generation</i>	Keep row alignment when depicting left/	1601
1524	-----	right relations; do not shift labels	1602
1525	SYSTEM	vertically.	1603
1526	-----	Keep column alignment when depicting	1604
1527	Assume you have expertise in spatial	above/below relations; do not shift	1605
1528	reasoning. Your task is to answer the	labels horizontally.	1606
1529	following question to the best of your	For diagonal relations, ensure the two	1607
1530	knowledge.	entities occupy different rows and	1608
1531	You will receive a directional question	different columns.	1609
1532	and a scene description.	Avoid inconsistent cell widths and do	1610
1533	Your task is to provide both the correct	not introduce extra or missing letters.	1611
1534	directional answer and the		1612
1535	corresponding ASCII visualization,	- Treat the canvas like a map: the first	1613
1536	including both in a single compact JSON	printed row is the top, and the last	1614
1537	object.	row is the bottom.	1615
1538		- The first column is the left edge; the	1616
1539	This is a two-turn task: Answer First,	last column is the right edge.	1617

1771
1772 Additional ASCII Guidance:
1773 - Ensure the ASCII visualization aligns
1774 with the answer you provide.
1775
1776 Generate the ASCII visualization before
1777 finalizing the answer.
1778
1779 Input Notes:
1780 - Use the description to understand the
1781 scene before answering.
1782 - Ensure the ASCII rendering is
1783 consistent with the textual evidence.
1784
1785 Output Notes:
1786 - Provide both the answer and the
1787 requested ASCII views in the JSON output
1788 .
1789
1790 Direction Notes:
1791 - Check the ASCII you generate so it
1792 supports the chosen answer.
1793
1794 OUTPUT SCHEMA (JSON):
1795 {
1796 "ascii": {
1797 "grid": "<grid using + - | for
1798 borders>"
1799 },
1800 "answer": "<direction>"
1801 }
1802
1803 Return ONLY a compact JSON object.
1804 Do not include markdown fences or any
1805 extra text.
1806
1807 -----
1808 USER
1809 -----
1810 Description:
1811 B is to the lower-right of A. D is to
1812 the right of B. C is at the bottom-right
1813 of A. C is located below B.
1814
1815 Question: Where is C horizontally
1816 relative to D?
1817
1818 -----
1819 ASSISTANT
1820 -----
1821 {"ascii":{"grid":"+++++\\n|A | |
1822 |\\n+---+---+\\n| |B |D |\\n+---+---+\\n|
1823 |C | |\\n+---+---+"},"answer":"left"}
1824
1825 *Two-turn Generation*
1826 -----
1827 SYSTEM
1828 -----
1829 Assume you have expertise in spatial
1830 reasoning. Your task is to answer the
1831 following question to the best of your
1832 knowledge.
1833 You will receive a directional question
1834 and a scene description.
1835 Your task is to provide both the correct
1836 directional answer and the
1837 corresponding ASCII visualization,
1838 including both in a single compact JSON
1839 object.
1840
1841 This is a two-turn task: Generate ASCII
1842 First, Then Answer.
1843
1844 - Turn 1: Generate an ASCII
1845 visualization from the scene description
1846 .
1847 - Turn 2: Answer the spatial question
using the ASCII you generated as a

reference.
Turn 1: Generate ASCII Visualization
Formats:
- grid: <grid using + - | for borders>
Rendering & Positioning Rules
Use given uppercase letters (A, B, C,
...); don't invent new ones.
Grid size is flexible - expand rows/cols
as needed to fit all points.
Place all letters; if two overlap, stack
them (e.g., "AF").
Relations are first -> second (e.g., "D
to F" means D relative to F).
Positioning:
above/below -> same column
left/right -> same level
diagonals (upper/lower-left/right) ->
one cell diagonally.
Keep perfect alignment for cardinal
directions.
Adjacent items occupy neighboring cells;
no diagonal drift.
Maintain uniform cell sizes and
consistent borders.
Keep row alignment when depicting left/
right relations; do not shift labels
vertically.
Keep column alignment when depicting
above/below relations; do not shift
labels horizontally.
For diagonal relations, ensure the two
entities occupy different rows and
different columns.
Avoid inconsistent cell widths and do
not introduce extra or missing letters.
- Treat the canvas like a map: the first
printed row is the top, and the last
row is the bottom.
- The first column is the left edge; the
last column is the right edge.
- Above = higher row; below = lower row
(do not invert the vertical axis).
- Left = smaller column index; right =
larger column index (do not mirror
horizontally).
Turn 2: Answer the Question
The possible directions are {left, right
, above, below, upper-left, upper-right,
lower-left, lower-right, same column,
same level}. There is only one correct
answer.
10 Distinct (Mutually Exclusive)
Directions
Pure Directions - perfectly aligned:
- left: directly left, no vertical
offset
- right: directly right, no vertical
offset
- above: directly above, no horizontal
offset
- below: directly below, no vertical
offset
Diagonal Directions - both vertical and
horizontal offset: upper-left, upper-
right, lower-left, lower-right
Horizontal Alignment ONLY: same level;
Vertical Alignment ONLY: same column

2079
2080 Question: Where is C horizontally
2081 relative to D?
2082
2083 -----
2084 ASSISTANT
2085 -----
2086 {"answer": "left"}

2087 *Two-turn Generation with ground truth ASCII*
2088 -----
2089 SYSTEM
2090 -----
2091 Assume you have expertise in spatial
2092 reasoning. Your task is to answer the
2093 following question to the best of your
2094 knowledge.
2095 You will receive a directional question
2096 and a scene description.
2097 Your task is to provide both the correct
2098 directional answer and the
2099 corresponding ASCII visualization,
2100 including both in a single compact JSON
2101 object.
2102
2103 This is a two-turn task: Generate ASCII
2104 First, Then Answer.
2105
2106 - Turn 1: Generate an ASCII
2107 visualization from the scene description
2108 .
2109 - Turn 2: Answer the spatial question
2110 using the ASCII you generated as a
2111 reference.
2112
2113 ## Turn 1: Generate ASCII Visualization
2114
2115 Formats:
2116 - grid: <grid using + - | for borders>
2117
2118 Rendering & Positioning Rules
2119
2120 Use given uppercase letters (A, B, C,
2121 ...); don't invent new ones.
2122 Grid size is flexible - expand rows/cols
2123 as needed to fit all points.
2124 Place all letters; if two overlap, stack
2125 them (e.g., "AF").
2126
2127 Relations are first -> second (e.g., "D
2128 to F" means D relative to F).
2129
2130 Positioning:
2131 above/below -> same column
2132 left/right -> same level
2133 diagonals (upper/lower-left/right) ->
2134 one cell diagonally.
2135 Keep perfect alignment for cardinal
2136 directions.
2137 Adjacent items occupy neighboring cells;
2138 no diagonal drift.
2139 Maintain uniform cell sizes and
2140 consistent borders.
2141
2142 Keep row alignment when depicting left/
2143 right relations; do not shift labels
2144 vertically.
2145 Keep column alignment when depicting
2146 above/below relations; do not shift
2147 labels horizontally.
2148 For diagonal relations, ensure the two
2149 entities occupy different rows and
2150 different columns.
2151 Avoid inconsistent cell widths and do
2152 not introduce extra or missing letters.
2153
2154 - Treat the canvas like a map: the first
2155 printed row is the top, and the last

row is the bottom.
- The first column is the left edge; the
last column is the right edge.
- Above = higher row; below = lower row
(do not invert the vertical axis).
- Left = smaller column index; right =
larger column index (do not mirror
horizontally).

Turn 2: Answer the Question

The possible directions are {left, right
, above, below, upper-left, upper-right,
lower-left, lower-right, same column,
same level}. There is only one correct
answer.

10 Distinct (Mutually Exclusive)
Directions
Pure Directions - perfectly aligned:
- left: directly left, no vertical
offset
- right: directly right, no vertical
offset
- above: directly above, no horizontal
offset
- below: directly below, no vertical
offset
Diagonal Directions - both vertical and
horizontal offset: upper-left, upper-
right, lower-left, lower-right
Horizontal Alignment ONLY: same level;
Vertical Alignment ONLY: same column

Output Notes:
- Generate the ASCII visualization
before attempting to answer.
- Use the ASCII you generated to
determine the correct answer.
- Turn 1: Provide the ASCII views as a
JSON object with "ascii" key.
- Turn 2: Provide the directional answer
as a JSON object with "answer" key.

OUTPUT SCHEMA (JSON):
{
 "answer": "<direction>",
 "ascii": {
 "grid": "<grid using + - | for
 borders>"
 }
}

Return ONLY a compact JSON object.
Do not include markdown fences or any
extra text.

USER

Description:
B is to the lower-right of A. D is to
the right of B. C is at the bottom-right
of A. C is located below B.

Question: Where is C horizontally
relative to D?

Please generate an ASCII visualization
of this scene.

ASSISTANT

{"ascii":{"grid":"+-+--+--+\\n|A | |
|\\n+-+--+--+\\n| |B |D |\\n+-+--+--+\\n|
|C | |\\n+-+--+--+"}}}

2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244

2245

2246
2247
2248
2249
2250

2251
2252

2253
2254
2255
2256
2257
2258
2259
2260
2261
2262

2263
2264
2265
2266
2267
2268

2269
2270
2271
2272
2273
2274
2275
2276
2277
2278

2279

2280
2281
2282

```
-----  
USER  
-----  
Now, based on the ASCII visualization  
you generated, answer the question:  
Where is C horizontally relative to D?  
  
-----  
ASSISTANT  
-----  
{ "answer": "left" }
```

K Artifacts, Licensing, and Intended Use

In this work, we introduce a new dataset and utilize several existing open-source artifacts. Below, we detail the licensing, terms of use, and intended usage for each component to ensure transparency and compliance with responsible AI research practices.

K.1 Our Contributions: The TEXT2SPACE Dataset

License and Terms The TEXT2SPACE dataset will be released under the **CC-BY 4.0 International License**, allowing for free use and redistribution with appropriate attribution. The accompanying code base, including the data generation pipeline and evaluation scripts, is released under the **MIT License**. We grant permission for researchers to modify and use these resources for both academic and commercial purposes, provided that the original paper is cited.

Intended Use The dataset is designed primarily to evaluate and enhance the spatial reasoning capabilities of Large Language Models (LLMs). It serves as a diagnostic benchmark for assessing specific spatial skills such as orientation, relative positioning, and transitive inference.

Limitations and Misuse Potential As the dataset consists entirely of synthetically generated spatial descriptions and ASCII diagrams using geometric primitives, it contains no personally identifiable information (PII) or offensive content. We foresee no substantial risk of misuse. However, users should note that the discrete grid representation is an abstraction and may not fully transfer to continuous real-world spatial tasks, such as autonomous driving or physical navigation validation.

K.2 Use of Existing Artifacts

We utilize publicly available models, datasets, and software libraries in accordance with their respective licenses and terms of use.

Models We employ the following models for training and evaluation:

- **Open-sourced Models:** We utilize LLAMA-3-70B-INSTRUCT (AI@Meta, 2024), GPT-OSS-20B (OpenAI, 2025), and Qwen Models (QWEN2.5-7B-INSTRUCT (Yang et al., 2024; Team, 2024), QWEN3-30B-A3B-INSTRUCT-2507, QWEN3-235B-A22B-INSTRUCT-2507 (Team, 2025)). These are open-weights models, and we adhere to their community license agreements, utilizing them strictly for academic research purposes and complying with its open-source license for research evaluation.
- **Proprietary Models:** We utilize GPT-4.1 (OpenAI, 2025) via its official API as a baseline and evaluator. We comply with the provider’s Terms of Service regarding data usage and content generation.

Computation Costs

- **Model Fine-tuning:** We conducted our fine-tuning experiments on a node equipped with 4 NVIDIA A40 GPUs. The training process for a single model fine-tuning required approximately 4.5 hours, resulting in a total computational cost of 18 GPU hours per model fine-tuning.
- **Inference:** Each inference task takes approximately 100 minutes for a total of 1000 samples.

Datasets To evaluate generalization capabilities, we utilize the STEPGAME benchmark (Shi et al., 2022). We use this dataset consistent with its original release terms, which are intended for academic research on spatial reasoning.

Software Libraries Our training and evaluation pipeline leverages the Hugging Face ecosystem, including transformers, trl, and peft (von Werra et al., 2020; Hu et al., 2021; Mangrulkar et al., 2022). These libraries are distributed under the Apache 2.0 license, and we use them in accordance with their documentation.