From Language to Pixels: Task Recognition and Task Learning in LLMs

Anonymous ACL submission

Abstract

Large language models (LLMs) can perform unseen tasks by learning from a few in-context examples, appearing to comprehend instructions like humans. We investigate how in-context learning works on a challenging non-language task. The task requires the LLM to generate 8x8 pixel matrices representing images of simple objects. We analyze if this task is solved by recognizing similar formats from the training data (task recognition) or by understanding the instructions and learning the skill de novo during inference (task learning). Our experiments demonstrate that LLMs use task recognition to generate meaningful pixel matrices and fail to learn such tasks when encountering unfamiliar formats. In that context, we introduce a simple framework to explain LLMs' capabilities and problem-solving processes that avoids anthropomorphism and complex descriptions of LLMs' inner workings. Our findings offer insights into LLMs' learning mechanisms to guide future research on their seemingly human-like behavior.

1 Introduction

The rapid development of LLMs increases the chance of misinterpreting their capabilities. The outputs generated by LLMs like ChatGPT can be astounding and are already used by millions. Exploring the potential of LLMs and researching their underlying mechanisms is crucial to understand how they accomplish complex tasks using next-token prediction alone. This not only helps in interpreting LLMs outputs correctly but also guides future research to improve language models.

Since Brown et al. (2020) demonstrated the abilities of GPT-3, there has been a growing interest in researching LLMs potential to solve downstream tasks. Many studies demonstrated LLMs' strong abilities in solving difficult reasoning tasks (Li et al., 2021; Srivastava et al., 2023; Huang and Chang, 2023). Letting the model explain each reasoning step has been shown to drastically improve results (Lampinen et al., 2022; Kojima et al., 2022; Wei et al., 2022). These findings suggest the potential of LLMs to solve complex unseen problems and break down a task into logical subtasks to achieve accurate results. Contrary, it has been shown that LLMs still struggle to reason about seemingly simple tasks where humans perform superior (Valmeekam et al., 2022; Binz and Schulz, 2023; Huang and Chang, 2023). These findings demonstrate ambiguity in LLMs' general abilities and the need to further probe challenging tasks.

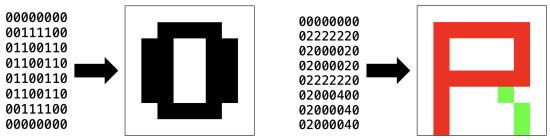

Figure 1: Two pixel matrices generated by an LLM representing simple 8x8 images.

Besides the ambiguity about general capabilities, it is also uncertain how and why LLMs benefit from in-context learning, i.e., how they can learn from instructions and examples in the prompt. The GPT-3 paper (Brown et al., 2020) suggests that in-context learning could stem from genuinely learning new tasks, identifying tasks learned during training, or a mix of both approaches. The authors emphasize the need to explore the precise mechanics of in-context learning as a crucial and unexplored avenue for future research. We adopt the terminology proposed by Pan et al. (2023) and distinguish between **task learning** (TL) and **task recognition** (TR) to thoroughly investigate LLMs' capabilities.

When investigating LLMs' capabilities, philosophically loaded terms like "learning" or "understanding" should be used with caution (Bender et al., 2021; Shanahan, 2022). LLMs lack human-like comprehension and reasoning because their training misses context, such as communicative intent or other modalities (Bender and Koller, 2020; Asher et al., 2023). Nevertheless, it remains in-

triguing to what degree current LLMs can mimic human-like learning, reasoning, and understanding to solve tasks. Frameworks and metaphors that avoid intricate and resource-intensive empirical analysis of LLMs' inner workings can help to explain LLM behaviors and to analyze their potential (Shanahan et al., 2023).

To investigate the learning capabilities of LLMs, we instruct a model to generate pixel art as 8x8 pixel matrices (see Figure 1) solely using the model's inherent tokens without additional fine-tuning or incorporating image models. While investigating this task, we demonstrate unique capabilities, pinpoint where the pixel matrix task falls along the spectrum between TL and TR, and establish a straightforward framework that breaks complex tasks into subtasks to clarify how LLMs solve such tasks.

Our main contributions are summarized as follows:

- We show that LLMs can create pixel matrices of digits, letters, and simple everyday objects, and compare different models on this task.
- We demonstrate that LLMs rely on task recognition to generate meaningful pixel matrices, which indicates strong task recognition abilities but limited task learning for uncommon non-language tasks in LLMs.
- We propose a framework for breaking tasks into subtasks to explain the capabilities of LLMs, enabling a better separation between TL and TR.

The next Section provides an overview of related works. Following that, we outline our experimental setup. Section 4 presents the results, showcasing LLM capabilities in the pixel matrix task, testing TR and TL through various experiments, and explaining our framework. Finally, the results are discussed, and we conclude our findings.

2 Related Work

Text-to-image models The success of LLMs has influenced the research on multi-modal models for vision and language. Many text-to-image models use LLM text encoders and combine them with an image encoder (e.g., Radford et al., 2021; Alayrac et al., 2022; Saharia et al., 2022). Others adhere closer to the LLM architecture. Koh et al. (2023) have used visual encoders to ground an LLM in the visual domain, enabling image captioning and text-to-image tasks. Similarly, Pourreza et al. (2023) generate images represented as brush strokes. They add an image feature extractor and cross-attention blocks to provide visual feedback during stroke generation. These studies demonstrate that LLMs possess implicit knowledge about images and the visual domain. Our task is distinct in that it is solved without extra image encoders or training.

Pure LLMs on image generation Probing the ability of LLMs to create images without fine-tuning or adding layers has been done by benchmarks that assemble tasks to assess LLM performances. While BIG-bench (Srivastava et al., 2023) includes some tasks that involve ASCII art, Bang et al. (2023) let ChatGPT draw country flags with Scalable Vector Graphics (SVG) code. In addition, there have been several blog posts exploring the ability of ChatGPT to draw images with SVG (Pu, 2022; Shahir, 2023 Shiryaev, 2022). To our knowledge, no comprehensive study has explored LLMs' capacity for generating visuals, and we are the first to assess the pixel matrix image format.

In-context learning How in-context learning works is still disputed. Some recent studies compared in-context learning to implicit fine-tuning or Bayesian inference (Dai et al., 2023; Von Oswald et al., 2023; Xie et al., 2022). Others have investigated whether this capability arises from comprehending the instruction or recognizing the task from the training data. Pan et al. (2023) introduced the terminology employed in our paper and distinguished cases where TL and TR were applied. Some studies have experimented with modifying the prompt or comparing tasks, aiming to ascertain whether LLMs employ TR or TL (Reynolds and McDonell, 2021; Min et al., 2022). We diverge in our approach by closely examining one particular task concerning TL or TR to contribute insights to the broader understanding of in-context learning and to provide a model for future task evaluations.

3 Experimental Setup

Our experimental setup aims to assess LLMs' general capabilities on the complex pixel-matrix task and determine whether it is addressed through TR and TL. Simultaneously, we demonstrate the utility of our framework of breaking down tasks.

3.1 Experiments

For the experiments, we adjusted the pixel symbols, the examples, and the task description of the prompt.

Testing general capabilities We conducted a *Baseline* experiment to show the general capabilities and act as a baseline for the other experiments. The pixel symbols are 0 and 1 and represent white and black. Furthermore, we conducted experiments with color representations to evaluate GPT-3.5's pixel matrix capabilities and determine the limits of the pixel matrix task. The *Color Digits* experiment added four additional digits for red, blue, yellow, and green, while the *Color RGB* experiment used RGB code values as pixel symbols.

Task learning vs. task recognition We follow three hypotheses as indicators for TL. First, performance should be independent of the frequency in the training data and equal across tasks of the same difficulty. Second, performance should deteriorate when providing misleading instructions. Lastly, performance should improve when making the task easier, considering tokenization and comparing how humans learn. Following these hypotheses, we designed three experiments.

The first experiment, (*GK Pixels*), substitutes the pixel values 0 and 1 with the random letters G and K. It is based on the assumption that black-and-white pixel matrices with 0s and 1s are prevalent in the training data while matrices with values G and K are not. Simple objects like digits and letters represented as matrices are likely well represented and learned during training, enabling TR. In contrast, the GK pixel matrix format must be learned from the context.

The second experiment, *Wrong Labels*, uses mislabeled examples in the prompt. We hypothesize that a deteriorating performance would indicate TL because misleading examples make learning tasks more difficult. However, a constant performance in this experiment means that in-context learning helps to identify a given task rather than actually learning it (Min et al., 2022; Pan et al., 2023). We conducted the *Wrong Labels* experiment by mislabeling each of the four prompt examples. We also experimented without any examples (zero-shot) to test if the model could learn from the instruction of the prompt alone.

GPT-3.5's tokenizer[1] combines multiple pixel symbols of our baseline matrix format into one token (see Figure 5). Consequently, a single pixel matrix can be represented by various tokens, expanding the token vocabulary from two to more than ten. Hence, learning the task becomes more challenging because it is necessary to understand the meaning of each potential token. For the third experiment (*One Token*), we inserted spaces between pixel symbols to ensure one token corresponds to a single pixel symbol and make the task easier to learn. Such adjustment enhanced results for straightforward pattern completion tasks (Mirchandani et al., 2023). In the case of TL, we anticipate improved results for our task.

In addition to these three experiments, we compare the results of objects we believe to be similarly challenging to sketch but unevenly represented as pixel matrices in the training data or of different difficulty but evenly represented in the training data. If the results align with the assumed difficulty, it suggests TL without relying on pixel matrices in the training data. For TR, only the training data is relevant. Consequently, when two similarly challenging objects have uneven representation, the one more prevalent during training should yield more accurate results.

Breaking tasks into subtasks With this set of experiments, we demonstrate our framework of breaking tasks into subtasks to explain and evaluate the capabilities of LLMs. Details of the framework are discussed in Section 5.4.

First, we generated textual descriptions of objects' shapes and visual appearance using a simple prompt (compare Appendix A) as this is a possible subtask serving as a first step for solving the pixel matrix task. Based on these self-generated descriptions, we let the model generate a 01 pixel matrix for the respective task, as this is the most common representation in the training set. Next, we tested the explicit translation from 01 to GK pixel matrices for the GK task with a simple prompt including four translation examples (compare Appendix A). If LLMs can generate GK pixel matrices correctly, one could assume they first recognize the 01 pixel matrices from training and then use the prompt to translate 01 pixel matrices to GK pixel matrices, solving two subtasks and combining the results.

We also tested our pixel matrix task with a different image format. Instead of pixel matrices, we used SVG code to generate the images. This task allows for a more straightforward combination of

[1] https://platform.openai.com/tokenizer

different shapes than with a pixel matrix because overlapping shapes do not influence each other. As the SVG code enables the inclusion of text, we did not test for digits, letters, and punctuation symbols. Examining the generation of real-world objects can reveal whether the model predominantly replicates SVG code for the specific object or decomposes the object into subparts and combines them.

Model comparison By default, we employed the GPT-3.5-turbo chat completion model, fine-tuned for instruction-based tasks accessed through the OpenAI API. In this paper, we refer to this model as GPT-3.5. We further compared various models (Bloom (Scao et al., 2022), GPT-Neox20B (Black et al., 2022), Starcoder (Li et al., 2023)) on the pixel matrix task. Our primary aim is to contrast models fine-tuned on code with traditional LLMs based on our hypothesis that with TL, the performance should be independent of the frequency in the training data. Anticipating a higher occurrence of pixel matrices in code-based data and better instruction comprehension by classic LLMs, we hypothesize that the traditional models will excel if TL is applied. Conversely, with TR, the code models should demonstrate superior performance.

We also compared GPT-4 (OpenAI, 2023) to its predecessor GPT-3.5.

3.2 Prompt Structures

The prompts in the pixel matrix experiments consist of a task description and demonstrations. The description introduces the concept of pixel symbols and pixel matrices. The four demonstration show examples with labels. These 8x8 pixel matrices are depicted with 8 rows, each containing 8 pixel symbols (e.g., 0 and 1 or G and K). Compare Figure 1 for one pixel matrix. After the four examples, another label specifying the object to be sketched is added. The model is expected to generate a corresponding pixel matrix.

For each experiment, we evaluated the performance of creating pixel matrices across four categories: digits, letters, punctuation symbols, and real-world objects. The real-world objects are items displayable on an 8x8 pixel canvas (e.g., chess board, padlock, or sun). In total, we have 930 different objects. We generated ten instances for each object. Appendix 2 includes a complete enumeration of the objects. We tested each of the four object categories separately. The category of the examples always aligned with the requested object's type. Refer to Appendix A for details on the full-length prompts.

Experiments	Digits	Letters	Punct.	Objs.
Baseline	73%	76%	45%	14%
Color Digits	56%	65%	32%	3%
Color RGB	15%	6%	10%	6%
GK Symbols	36%	37%	21%	2%
One Token	72%	72%	36%	12%
Wrong Label [2]	50%	67%	35%	12%
GPT-4 Baseline	79%	91%	66%	23%
GPT-4 RGB	62%	43%	40%	21%

Table 1: Comparing experiment results: Percentage of recognizable images across generation tasks and image categories. Results obtained by human evaluation.

3.3 Evaluation

We converted each pixel matrix to an image with a simple Python script. Then, we conducted a classification study with three annotators. Each annotator classified all generated images without knowing the set of possible generated digits, letters, and punctuation symbols. For real-world objects, we provided the correct answer and let the annotator evaluate whether the respective object is recognizable because even a good image on an 8x8 pixel canvas is challenging to recognize without context. We counted the percentage of generations correctly classified or marked as recognizable for each experiment.

4 Results

This section presents the results of our experiments. The quantitative results are summarized in Table 1.

4.1 General Capabilities

At first, we establish a reference point with the basic experiment and demonstrate the general capabilities of GPT-3.5 to create pixel matrices, including colorful images.

Basic experiment For digits, nearly all generated outputs are easily identifiable, with the only exception being the number 4, which is only recognizable in three out of ten instances. For two-digit numbers (10 and 32), none of the generated examples accurately display the number. Instead, some pixel matrices represent other meaningful objects, like the letters A or H, or the number 9.

Almost all letters are consistently generated correctly and recognizable. Exceptions occurred with

[2] Since the objects of the mislabeled demonstrations are also evaluated in the digits and letters categories, this potentially lowers the score by 40.

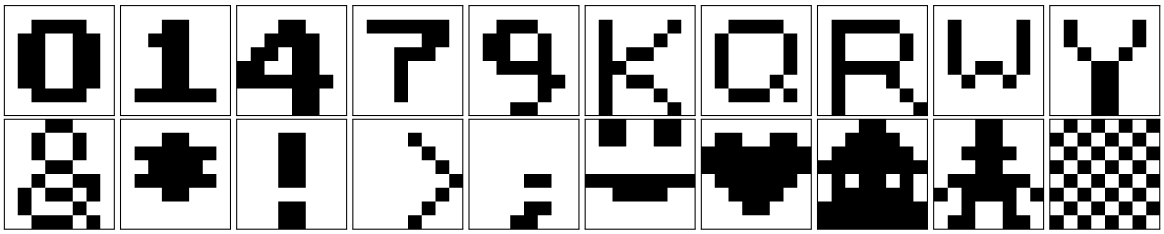

Figure 2: Selected positive examples of the baseline experiment: digits 0, 1, 4, 7, and 9; letters K, Q, R, W, and Y; the symbols ampersand, asterisk, exclamation point, and semi-colon; the objects sad face, heart, house, stick figure, and chess board.

more complex letters. The letter E often features more than three horizontal lines, and W and M are occasionally wrong. Surprisingly, the letter V consistently appears as an X in the pixel matrix. The German umlauts Ä, Ö, and Ü fail to show corresponding vowels, often resulting in seemingly random pixel matrices.

Except for the percent and dollar signs, most punctuation symbols are generated correctly at least once in ten instances. Even complex symbols, such as the ampersand, are successfully abstracted and generated on an 8x8 pixel matrix. More common symbols like the comma and exclamation mark are consistently generated. Left closing symbols yield good results, while not a single right closing symbol is generated correctly (see Figure 6). When requesting a right closing symbol, the output typically corresponds to the left symbol.

Simple and common everyday items like the heart, sad face, and house yield mostly recognizable images. For the stick figure and cat, some instances are of high quality, while others seem completely random. Results for the remaining objects appear mostly random. Compare Figure 2 for selected positive examples.

Adding color Results for digits, letters, and punctuation show slightly lower accuracy than simple black-and-white images. Complex letters such as W and M become even less recognizable than in the basic experiment. Certain letters comprised multiple colors. The real-world objects did not benefit from the added colors. Compare Figure 13 for selected images from these experiments.

The quality of the generated pixel matrices in the real-world object category declines. The images display vibrant colors, but the colors lack an association with the specified objects. The images for the star and sun does not contain more yellow, and those for a window or glass does not contain more blue. Only the cactus images consistently appears green in all ten images.

```
- The fundamental shape is a vertical line
positioned in the center of the 8x8 grid.
- At the top of the line, there is a small
horizontal line extending towards the right
side of the grid, connected to the vertical
line's midsection.
- The bottom part of the vertical line extends
slightly below the grid's baseline, forming a
slight curve.
```

Figure 3: Excerpt from a generated description for the digit 7, showing incompetence in that task.

```
The simplified 8x8 pixel representation of a
house consists of a square shape measuring 6
pixels in height and 6 pixels in width,
representing the main body of the house. On top
of the square, centered horizontally, there is
a triangle shape measuring 4 pixels in height
and 6 pixels in width, representing the roof.
The top row of the triangle is aligned with the
top row of the square.
```

Figure 4: Excerpt from a generated description for a real-world object, illustrating the ability to describe individual components accurately, yet assembling them incohesive.

For the experiment with RGB color codes as pixel symbols, the model often generates a textual response indicating the inability to produce images instead of a pixel matrix. Most of the resulting shapes appear random or replicate an example from the prompt.

4.2 Task Recognition vs. Task Learning

The results demonstrated in this section reveal insights into where the pixel matrix task lies on the spectrum between TR and TL, which will be discussed in Section 5.2.

Different pixel symbols When we substitute the values 0 and 1 with the random letters G and K, the output is always a pixel matrix with G and K. Some instances are generated correctly, but overall, the results are significantly worse across all four categories (see Table 1).

Wrong labels The quality of the results across all four categories remains consistent with the baseline. Even when the model regenerates examples from the prompt, it ignores the incorrect labels and produces correct pixel matrices. Additionally, the zero-shot experiments reveal that the model needs examples to adhere to the format.

One token per pixel symbol We observe small setbacks compared to the basic experiment, particularly for punctuation symbols. The output more often does not conform to the correct format, generating a message stating it is a language model and cannot generate images.

4.3 Decomposing the pixel matrix task

Translating 01 pixel matrices to GK pixel matrices is successful in 51 out of 60 instances. According to our evaluation, 35% of SVG images of real-world objects are correct, a much higher score than for all other experiments. With only a few exceptions, the resulting images display the colors connected to the desired object. Inaccurate images frequently present correct parts of objects, but the model fails to assemble the details in a hierarchical, cohesive way (compare Figure 7).

Generated textual descriptions often lack coherence for digits, letters, and most punctuation symbols. The described shapes appear random for all numbers except for 0, 1, and 8 (see Figure 3). Some level of abstraction is observed for real-world objects, but coherency keeps lacking. The generated texts mention useful shapes for an object but unrealistically combine them (see Figure 4). When using the self-generated descriptions to create pixel matrices, there is a noticeable decline in the overall results for digits, letters, and punctuation. For objects, no clear overall difference emerges between the basic experiment.

4.4 Model Comparison

GPT-3.5, with 175 billion parameters, demonstrates superior performance. In contrast, Bloom, also equipped with 175 billion parameters, adheres to the format but fails to solve any pixel matrix correctly. Gpt-Neox with 20 billion parameters does not generate any meaningful pixel matrix. The smallest model (15.5 billion parameters), Stracoder, fine-tuned specifically for code-related tasks, displays the most robust performance behind GPT-3.5. Its generations were similar to the ones from GPT-3.5 for digits and letters, but it consistently fails to draw punctuation symbols and real-world objects (See Figure 7).

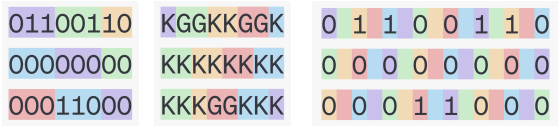

Figure 5: Comparing tokenization of different image formats: each color-coded sequence represents symbols combined into a single token. Images are screenshots from the OpenAI tokenizer webpage.

GPT-4 shows significant improvements compared to GPT-3.5. It creates much more meaningful real-world objects on an 8x8 canvas (e.g., padlock, flash, key, or cactus), and it consistently creates colorful digits with RGB codes as pixel values, which GPT-3.5 cannot generate.

5 Discussion

In this section, we examine our experiment results. We discuss the overall performance in the pixel matrix task, assess if they tackle the task through TL or TR, and evaluate if breaking down the task into subtasks is a valid framework to explain LLM abilities.

5.1 LLMs Image Generation Capabilities

GPT-3.5 showcases a remarkable ability to generate simple images in the form of pixel matrices without fine-tuning and including an image layer. This feat alone would have been shocking just a few years ago and displays the speed at which LLMs evolve. Furthermore, there might be undiscovered potential with better prompt engineering (e.g., more examples in the prompt or explicit chain of thought prompting).

The results indicate that language models already possess information across different modalities. For example, LLMs might outperform traditional image models in describing the abstract appearance of objects since describing visuals through language inherently involves abstraction and omission of details. In contrast, traditional image models learn with explicit images. This leads us to conclude that multi-modal models might reveal strong synergies.

The capability of an LLM to effectively handle non-language tasks, such as generating pixel matrices, also highlights its potential beyond language generation. LLMs could address tasks like pattern completion or adapt to unknown formats through fine-tuning. To improve LLMs for such tasks, it is

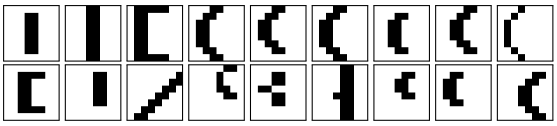

Figure 6: Left (top row) and right (bottom row) closing symbols, emphasizing the discrepancy in image quality for closely related objects.

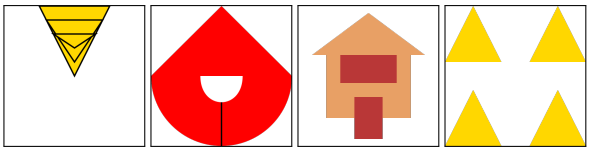

Figure 7: SVG-generated images (crown, wine glass, house, and lightning flash) exhibit a lack of cohesiveness in the arrangement of their components.

crucial to explore whether LLMs have the ability to learn new formats from scratch during inference (TL) or if they rely on TR, manifesting the importance of a large and diverse training corpus.

5.2 Pixel Matrix Task Is Solved by Task Recognition

We found evidence that the pixel matrix task cannot be solved without TR. This evidence includes the deteriorated outcomes with uncommon (GK) pixel values and unaltered results despite incorrect examples in the prompt (compare Table 1). Additionally, our zero-shot experiments confirm that relying solely on the instruction part of the prompt is insufficient for the model to learn the task. Altering the format to represent each pixel by one token does not improve the results despite simplifying the task and making it easier to learn. Also, besides GPT-3.5-turbo, only Starcoder, fine-tuned on code, generates some meaningful results, providing evidence for TR.

Further, several pairs of objects with the same difficulty level are solved in only one of two instances. These object pairs include 1) digits vs. number 10, 2) letter A vs. letter Ä, and 3) left closing symbols vs. right closing symbols (see Figure 6). If the model *truly* understands the tasks (rather than only replicating training instances), it would be able to solve either both or none of the pair's instances. We assume that for all three instances, the former is significantly more present during training than the latter. Similarly, drawing the complex shapes of some digits, punctuation symbols, and letters appears more challenging than drawing simple real-world objects. For instance, a child might find it challenging to copy an ampersand but not a simple smiley face. However, the results are consistently better for such complex shapes than for simple real-world objects.

In summary, our findings indicate that the pixel matrix task is solved through TR, not TL. Thus, it shows that LLMs lack human-like understanding and generalization, despite occasional impressions. Rather, the large data corpus is extremely powerful and includes samples of uncommon tasks. Given the current architecture, we assert that LLMs will not attain human-like understanding and reasoning. Enhancing current LLMs requires incorporating even more diverse data and improved training.

5.3 What Task Is Recognized?

Based on the findings above, we conclude that the model primarily relies on TR to solve the pixel matrix task. While it recognizes the task and uses training data to generate pixel matrices, we believe the pixel matrix task has been learned only to a limiting degree during training. Language models are not extensively trained on text-to-pixel matrix data, making it challenging to generalize this task to new images by combining seen objects and concepts. We conclude that it is more appropriate to say that it recognizes the task of drawing a specific object instead of recognizing a general text-to-image task. Therefore, the model cannot generalize to pixel matrices for objects it has not encountered during training, as this ability is neither acquired during training nor during inference. In conclusion, when evaluating LLMs on a task, it's crucial to pinpoint the exact task that is recognized and solved.

5.4 Breaking Tasks into Subtasks

Although the overall results are worse when using uncommon (GK) pixel values, meaningful pixel matrices are still generated in some instances. We investigate the capability of breaking the task into individual steps as a straightforward explanation of how LLMs solve complex and unseen tasks. Some steps may be addressed through TR, while others may involve TL.

Pure TR might have been applied during experiments with the standard format and the object categories letters and digits, thus skipping all intermediate steps. However, pure TR does not explain the results of the pixel matrix task with pixel values G and K, as it is unlikely that it has encountered any such pixel matrices during training. The model seems to be able to break the task down of creating a GK pixel matrix into the subtasks of creating a

01 pixel matrix and exchanging all 0s and 1s with Gs and Ks. We showed that both subtasks can be solved by themselves. We believe that LLMs can recognize the pixel matrices from training and, instead of adhering to a familiar set of symbols, translate them to a different one.

This observation makes us consider whether a model could further break down the pixel matrix task. Could an LLM theoretically solve the pixel matrix task by breaking it down into subtasks and addressing them individually without relying on TR? The pixel matrix task could be broken down into simpler subtasks such as accessing visual information about the world, abstracting object visuals to fit on a small canvas, breaking down object visuals into basic shapes, generating different shapes on a pixel matrix, and combining this information to form a pixel matrix representing a certain object. To solve the pixel matrix task without TR, an LLM would then need to "understand" the meaning of instructions, be able to solve all subtasks and apply "reasoning" to combine them in a meaningful way to create accurate pixel matrices. To a lesser degree, this ability was shown during the G and K pixel matrix tasks.

Our results suggest that current models fail to break down the pixel matrix task accordingly and thus rely on TR to solve the task. One reason might be that the model fails at one of the needed subtasks, as we showed that the generated textual descriptions for digits, letters, and punctuation symbols were insufficient. Also, it is very challenging to combine different shapes on a pixel matrix, as overlapping parts might result in completely new tokens.

However, for real-world objects, the visual information about the shapes might be enough. Further, in SVG code, combining various elements and shapes is accomplished by concatenating corresponding lines of code. Thus, a real-world world object with SVG code might be generated by breaking down the object into its visual parts and combining them into one image. The examples generated by Pu (2022) and our SVG experiment showed that the model draws different parts of objects but fails to put them together with correct dimensions and spatial correlations (see Figure 7). This suggests that an LLM solves image generation with SVG by breaking it down into subtasks and combining the results accordingly.

The SVG and GK experiments showcase how our framework of analyzing the subtask helps to evaluate and explain the capabilities of LLMs. This practice prevents misleading conclusions, sheds light on nuanced abilities, and helps us to predict what is possible. When analyzing tasks like the GK pixel matrix, one might mistakenly conclude that the model is learning from scratch. However, breaking the task down reveals that it might not be that difficult to solve the rather simpler subtask learned during training.

Breaking down tasks into subtasks to describe the capabilities of LLMs makes it more difficult to distinguish between TL and TR. Does true TL at inference exist if it breaks down the task and recognizes these subtasks learned during training? Do we call the concept of breaking down a task and combining the solutions of the subtasks into a new unseen output true TL? If yes, is TL that surprising? We do not have an answer to these questions yet. Future work might compare these thoughts to how humans learn.

6 Conclusion

Our study demonstrates that LLMs can generate pixel matrices representing objects such as digits, letters, and simple real-world items. Our experiments show further evidence of strong task recognition and limited task learning ability. We argue that breaking down complex tasks into smaller subtasks is a useful framework for evaluating and explaining LLM capabilities, preventing misleading conclusions when assessing the task's overall performance and helping to distinguish between TL and TR.

While we believe a more advanced pixel matrix task is challenging for LLMs and probes some interesting capabilities, we expect further and probably fast improvements for it with new models. These results may be surprising, potentially convincing one to believe that the model comprehends the task in a human-like manner. However, one should be careful not to be fooled that the results stem from truly understanding the task. Instead, LLMs will have seen more data, improve the ability to recognize tasks, locate appropriate outputs, break down complex tasks, and translate learned formats to new ones. In summary, LLMs will excel at predicting the next token in a sequence of words given a learned distribution, which may seem human but very much isn't.

7 Limitations

The lack of a more objective and robust evaluation is one of the limitations of this study. Nevertheless, our evaluation captures broader trends, and our discussion is based on conclusions drawn from clear tendencies in the results rather than small differences. As the outputs depend on prompts, more prompt engineering and other prompts could have yielded different results.

We also rely on the generated outputs and general knowledge about LLMs for interpreting their results because empirical analysis of the model's internal workings (e.g., evaluating attention patterns) is extremely resource-intensive. Thus, we interpret LLMs, essentially a highly complex statistical distribution, with simple human-like concepts (e.g., breaking down tasks). The discrepancy in complexity between human concepts and the inner workings of LLMs may lead to wrongful conclusions.

References

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, et al. 2022. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*.

Nicholas Asher, Swarnadeep Bhar, Akshay Chaturvedi, Julie Hunter, and Soumya Paul. 2023. Limits for learning with language models. In *Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (*SEM 2023)*, pages 236–248, Toronto, Canada. Association for Computational Linguistics.

Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.

Emily M. Bender and Alexander Koller. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.

Marcel Binz and Eric Schulz. 2023. Using cognitive psychology to understand gpt-3. *Proceedings of the National Academy of Sciences*, 120(6):e2218523120.

Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, Usvsn Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. GPT-NeoX-20B: An open-source autoregressive language model. In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136, virtual+Dublin. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. 2023. Why can GPT learn in-context? language models secretly perform gradient descent as meta-optimizers. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4005–4019, Toronto, Canada. Association for Computational Linguistics.

Jie Huang and Kevin Chen-Chuan Chang. 2023. Towards reasoning in large language models: A survey. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1049–1065, Toronto, Canada. Association for Computational Linguistics.

Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. 2023. Grounding language models to images for multimodal generation. *arXiv preprint arXiv:2301.13823*.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.

Andrew Lampinen, Ishita Dasgupta, Stephanie Chan, Kory Mathewson, Mh Tessler, Antonia Creswell, James McClelland, Jane Wang, and Felix Hill. 2022. Can language models learn from explanations in context? In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 537–563, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Belinda Z. Li, Maxwell Nye, and Jacob Andreas. 2021. Implicit representations of meaning in neural language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1813–1827, Online. Association for Computational Linguistics.

Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, et al. 2023. Starcoder: may the source be with you! *arXiv preprint arXiv:2305.06161*.

Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Suvir Mirchandani, Fei Xia, Pete Florence, brian ichter, Danny Driess, Montserrat Gonzalez Arenas, Kanishka Rao, Dorsa Sadigh, and Andy Zeng. 2023. Large language models as general pattern machines. In *7th Annual Conference on Robot Learning*.

OpenAI. 2023. Gpt-4 technical report. *ArXiv*, abs/2303.08774.

Jane Pan, Tianyu Gao, Howard Chen, and Danqi Chen. 2023. What in-context learning "learns" in-context: Disentangling task recognition and task learning. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8298–8319, Toronto, Canada. Association for Computational Linguistics.

Reza Pourreza, Apratim Bhattacharyya, Sunny Panchal, Mingu Lee, Pulkit Madan, and Roland Memisevic. 2023. Painter: Teaching auto-regressive language models to draw sketches. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 305–314.

Evan Pu. 2022. Probing compositional understanding of chatgpt with svg. https://evantheBouncy.medium.com/probing-compositional-understanding-of-chatgpt-with-svg-74ec9ca106b4 accessed 12/10/2023.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.

Laria Reynolds and Kyle McDonell. 2021. Prompt programming for large language models: Beyond the few-shot paradigm. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI EA '21, New York, NY, USA. Association for Computing Machinery.

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494.

Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.

Jamshaid Shahir. 2023. Writing code to produce images with chatgpt. https://towardsdatascience.com/image-generation-with-chatgpt-68c98a061bec accessed 12/10/2023.

Murray Shanahan. 2022. Talking about large language models. *arXiv preprint arXiv:2212.03551*.

Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. Role play with large language models. *Nature*, pages 1–6.

Denis Shiryaev. 2022. Drawing mona lisa with chatgpt. https://neural.love/blog/chatgpt-svg accessed 12/10/2023.

Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2023. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*.

Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2022. Large language models still can't plan (a benchmark for LLMs on planning and reasoning about change). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.

Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, Joao Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. Transformers learn in-context by gradient descent. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 35151–35174. PMLR.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. An explanation of in-context learning as implicit bayesian inference. In *International Conference on Learning Representations*.

A Prompt Details

This section supplements Section 3 by offering additional details and complete examples of the prompts used in our experiments.

A.1 Main Experiments

All prompts consist of one natural language instruction that describes the pixel matrix task by introducing pixels and the concept of pixel matrices and pixel symbols. This part is adjusted if the pixel symbols change. The second part of the prompts contains the demonstrations adapted to the object category and pixel symbols.

Independent of the pixel format, the prompt includes examples of numbers 2, 3, 8, and 6 for digits. The prompt examples for letters are A, F, G, and H. In the category of punctuation symbols, the examples are the forward slash, division sign, question mark, and the dot symbol, and for real-world objects, the examples are the smiley face, umbrella, tree, and deer.

See Figure 8 for a complete prompt for the baseline experiment.

A.2 Generating Descriptions

Refer to Figure 10 for the prompt used to generate descriptions of the objects' shapes and appearances.

A.3 Translating 01 to GK pixel matrices

Refer to Figure 11 for the prompt used to translate 01 pixel matrices to GK pixel matrices.

B Model Configurations

Pre-experiments revealed that a temperature of 1 was a good balance between randomness to ensure differences between each instance while maintaining the required format. We implemented a stop sequence for OpenAI models and restricted generated tokens to 100 for huggingface models. Apart from these modifications, we adhered to the default settings provided by the respective model APIs.

C Complementary Results

Some figures showin images resulting from our experiments.

C.1 Color Digits

See Figure 13 for selected results of experiments with added digits as pixel symbols to represent more colors.

```
Images displayed on a computer screen are a
collection of color dots, called pixels. If you
look really closely at the screen, you will be
able to see the individual pixels. The
collection of pixels that make up an image are
stored as a matrix.
We can represent different objects (e.g.,
numbers, letters, or shapes) by creating a
pixel matrix which consists of 0s and 1s. The
matrix should be of the size 8 by 8. Each entry
represents a pixel of a black or a white pixel.
That means the image has a display capable of 8
pixels in width and 8 pixels in height. Since
there are only 64 pixels in total the objects
to be displayed are significantly simplified.

Here is an example of an 8 by 8 pixel matrix
showing  three:
00000000
00111110
00000110
00111100
00001110
00000110
00111100
00000000
###
Here is an example of a grid of pixels that
form an image of  two:
00000000
00111100
01100110
00001100
00011000
00110000
01111110
00000000
###
Here is an example of a grid of pixels that
form an image of  eight:
00000000
00111100
01100110
00111100
01100110
00110100
00011000
00000000
###
Here is an example of a grid of pixels that
form an image of  six:
00000000
00111100
01100000
01100000
00111100
01100110
00111100
00000000
###
This is an example of a grid of pixels that
form an image of [object]:
```

Figure 8: The full prompt used to generate digits for the baseline experiments, with a placeholder for the requested digit.

```
Images displayed on a computer screen are
actually a collection of dots of color, called
pixels. If you look really closely at the
screen, you will be able to see the individual
pixels. The collection of pixels that make up
an image are stored as a matrix. Each pixel can
represent a different color.
We can represent different objects (e.g.,
numbers, letters, or shapes) by creating a
pixel matrix which consists different symbols
and each symbol stands for a different color.
In our 8x8 pixel matrix, we use the following
symbols to encode color: white (represented by
"0"), black ("1"), red ("2"), yellow ("3), green
("4"), and blue ("5"). This means each entry of
the matrix is either a 0, a 1, a 2, a 3, a 4, or
a 5. The matrix should be of the size 8 by 8.

Here is an example of an 8 by 8 pixel matrix
showing a smiley face:
31133113
31133113
33333333
33313333
13333331
31333313
33111133
33333333
###

Here is an example of a grid of pixels that
form an image of a tree:
00444400
04444440
44444444
44444444
44444444
00011000
00011000
00011000
###

Here is an example of a grid of pixels that
form an image of an umbrella:
00555500
05555550
55555555
00001000
00001000
00001000
00101000
00011000
###

Here is an example of a grid of pixels that
form an image of a deer:
10001000
11111000
01010000
11110001
01111111
01111111
01010101
01010101
###
```

Figure 9: Complete prompt used for experiments with new symbols representing different colors for the real-world object category. The first part and the demonstrations are adjusted.

```
Describe a simplified visual representation of
[object] which can be used to create an 8x8
pixel artwork of [object]. Emphasize only the
essential features for recognition, omitting
intricate details due to space constraints.
Deliver a concise description of the
fundamental shape and distinctive traits and if
necessary mention proportions, alignments, and
spatial relationships in a simplified rendition
of [object].
```

Figure 10: Prompt for generating descriptions of objects shapes which were added to pixel matrix prompt with the idea to enhance the outputs.

```
I want you to translate a 01 pixel matrix to a
GK pixel matrix. Replace every 0 with a G and
every 1 with a K.
Here is an example of a 01 pixel matrix:
00000000
00111100
01100000
01100000
00111100
01100110
00111100
00000000
###
Translation:
GGGGGGGG
GGKKKKGG
GKKGGGGG
GKKGGGGG
GGKKKKGG
GKKGGKKG
GGKKKKGG
GGGGGGGG
###

[...three more examples...]

Here is an example of a 01 pixel matrix:
[Exmaple 01 pixel matrix]
Translation:
```

Figure 11: Prompt for generating descriptions of objects shapes which were added to pixel matrix prompt with the idea to enhance the outputs.

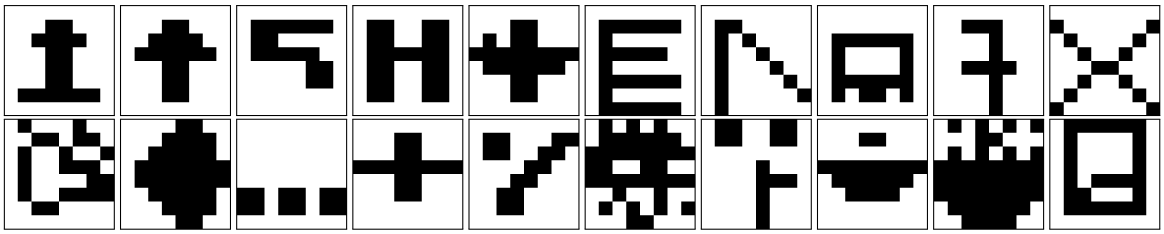

Figure 12: Negative examples from the baseline experiment: digits 1, 4, 5, 10, and 32; letters E, N, Ö, T, and V; the symbols at-sign, dollar sign, double quote, equal sign, and percent sign; the objects boat, cup, fly, star and wine glass.

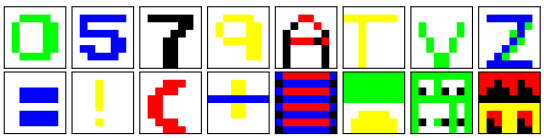

Figure 13: Selected instances of images generated with new pixel values for color, adhering to the format.

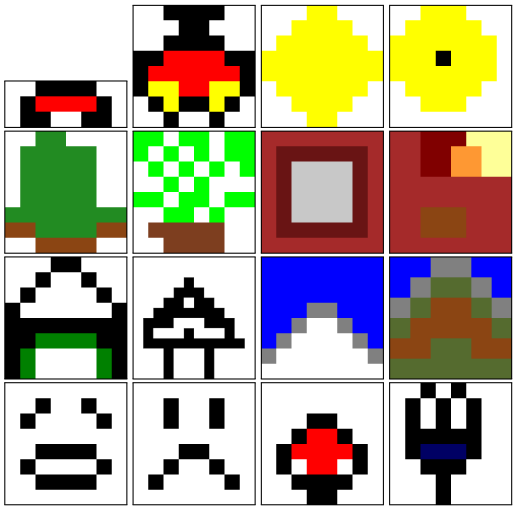

Figure 14: Selected instances of images generated by GPT-4 on an 8 by 8 pixel matrix with RGB values as pixel symbols showing two instances of each object: car, sun, cactus, coffe, house, mountain, sad face, wine glass. The magenta stripes resulting from translating slightly off output textual format to images.

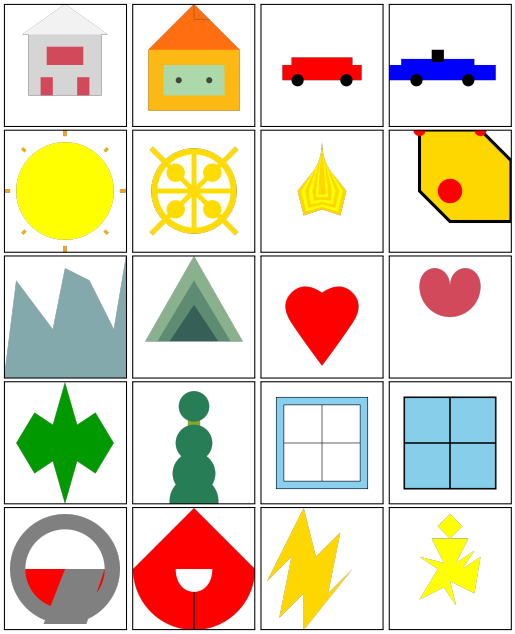

Figure 15: Example results from experiments with SVG code as image format showing two instances of each object: house, car, sun, crown, mountain, heart, cactus, window, wine glass, and lightning flash.

C.2 GPT-4 Color RGB

See Figure 14 for selected images generated by GPT-4 with RGB color codes as pixel symbols.

C.3 Images with Scalable Vector Graphics

Compare Figure 15 for images generated with SVG code.

D Supplementary Findings

One thought was that an 8x8 pixel canvas might be too limiting, hindering the LLM from generating meaningful images of real-world objects like a cat

or a boat. We conducted experiments with a pixel matrix of size 16 by 16. However, the overall results showed a slight degradation and we did not see any improvements for specific objects where a larger canvas may be beneficial.

We have conducted most of our experiments with a format that used more than ten different tokens to represent a pixel matrix as the tokenizer combines sequences of 0s and 1s. During our experiments with G and K as symbols, we assume that it translated seen in 01-matrices to the new format even though the tokens themselves and of tokens changes (see Figure 5). Thus, the token embeddings represent even uncommon non-semantic similarities.

We also attributed some observed behavior to the autoregressive architecture of GPT-like models. For example, the letter V was never correctly generated and always resulted in the letter X. We assume that after correctly generating the first line of the matrix according to V, the subsequent generations favored a pixel matrix showing the letter X due to its higher frequency in training compared to that of a V. If the prompt would be truly "understood", the attention to the previous pixels should not overshadow the requested object.

In the *One Token* Experiment, overall results showed a slight decline, but notable improvements were observed with the new matrix format, especially for the object 'chess board'. An additional experiment with the chess board showed that while 42 out of 60 generations were correct with the basic matrix format, all 60 chess board generations were correct with the new format. We assume this might be because this object resembles a pattern completion task, which is less error-prone with fewer tokens Mirchandani et al. (2023).

E Enumeration of All Objects

We evaluated the performance of creating pixel matrices across four categories: digits, letters, punctuation symbols, and real-world objects. Digits range from 0 to 9, with additional two-digit numbers, 10 and 32. The letter category comprises all Latin alphabet letters, including German umlauts (Ä, Ü, Ö) and eszett (ß), totaling 30 letters. The punctuation symbol category encompasses 25 different symbols, while real-world objects consist of 26 diverse items. All objects are enumerated in Table 2.

digits	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 32
letters	A, B, C, D, E, F, G, H, I, J, K, L, M, N, O, P, Q, R, S, T, U, V, W, X, Y, Z, Ä, Ö, Ü, ß
punct.	comma, semicolon, exclamation point, equal sign, plus sign, hashtag, dollar sign, percent sign, ampersand, asterisk, left parenthesis, right parenthesis, left bracket, right bracket, left curly brace, right curly brace, less than sign, greater than sign, backslash, underscore, colon, single quote, double quote, at sign, caret
objs.	sad face, cup, heart, wine glass half full, cactus, key, skull, mouse, crown, lightning flash, padlock, cat, crab, a chess board, a house, coffee, car, window, chair, star, mountain, sun, boat, stick figure, fly

Table 2: Enumeration of all generated objects in our experiments.

F Evaluation

We recruited 3 students from our university to annotate the images generated in our experiments. Figure 16 shows our whole set of instructions.

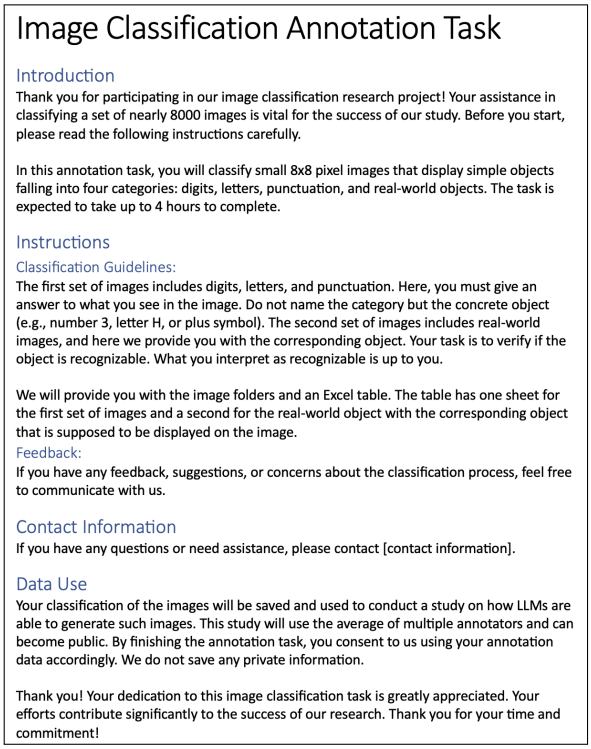

Image Classification Annotation Task

Introduction

Thank you for participating in our image classification research project! Your assistance in classifying a set of nearly 8000 images is vital for the success of our study. Before you start, please read the following instructions carefully.

In this annotation task, you will classify small 8x8 pixel images that display simple objects falling into four categories: digits, letters, punctuation, and real-world objects. The task is expected to take up to 4 hours to complete.

Instructions

Classification Guidelines:

The first set of images includes digits, letters, and punctuation. Here, you must give an answer to what you see in the image. Do not name the category but the concrete object (e.g., number 3, letter H, or plus symbol). The second set of images includes real-world images, and here we provide you with the corresponding object. Your task is to verify if the object is recognizable. What you interpret as recognizable is up to you.

We will provide you with the image folders and an Excel table. The table has one sheet for the first set of images and a second for the real-world object with the corresponding object that is supposed to be displayed on the image.

Feedback:

If you have any feedback, suggestions, or concerns about the classification process, feel free to communicate with us.

Contact Information

If you have any questions or need assistance, please contact [contact information].

Data Use

Your classification of the images will be saved and used to conduct a study on how LLMs are able to generate such images. This study will use the average of multiple annotators and can become public. By finishing the annotation task, you consent to us using your annotation data accordingly. We do not save any private information.

Thank you! Your dedication to this image classification task is greatly appreciated. Your efforts contribute significantly to the success of our research. Thank you for your time and commitment!

Figure 16: The complete instructions given to the human annotators.