

Poetry in Pixels: Prompt Tuning for Poem Image Generation via Diffusion Models

Sofia Jamil¹ Bollampalli Areen Reddy¹ Raghvendra Kumar¹
Sriparna Saha¹ K J Joseph² Koustava Goswami²

¹Department of Computer Science & Engineering, Indian Institute of Technology Patna, India

²Adobe Research, India

¹{sofia_2321cs16, bollampalli_2101cs19, raghvendra_2221cs27, sriparna}@iitp.ac.in

²{koustavag, josephkj}@adobe.com

Abstract

The task of text-to-image generation has encountered significant challenges when applied to literary works, especially poetry. Poems are a distinct form of literature, with meanings that frequently transcend beyond the literal words. To address this shortcoming, we propose a *PoemToPixel* framework designed to generate images that visually represent the inherent meanings of poems. Our approach incorporates the concept of prompt tuning in our image generation framework to ensure that the resulting images closely align with the poetic content. In addition, we propose the *PoeKey* algorithm, which extracts three key elements in the form of emotions, visual elements, and themes from poems to form instructions which are subsequently provided to a diffusion model for generating corresponding images. Furthermore, to expand the diversity of the poetry dataset across different genres and ages, we introduce *MiniPo*, a novel multimodal dataset comprising 1001 children’s poems and images. Leveraging this dataset alongside *PoemSum*, we conducted both quantitative and qualitative evaluations of image generation using our *PoemToPixel* framework. This paper demonstrates the effectiveness of our approach and offers a fresh perspective on generating images from literary sources. *The code and dataset used in this work are publicly available.*¹.

1 Introduction

Poetry, with its layered meanings and expressive language, often paints descriptive images in the mind of the reader. But what if these imagined scenes could be brought to real vision? Visualizing poetry adds a new dimension to the experience, allowing us to see and feel the poem in a way that extends beyond the words. In recent years, the automatic generation of realistic images from arbitrary

text descriptions has attracted a lot of attention. Several initiatives have been made to address the zero-shot text-to-image generation challenge by pre-training large-scale generative models on massive datasets of image-text pairs, such as DALL-E (Ramesh et al., 2021) and CogView (Ding et al., 2021). Unlike typical text, creating images for poems presents a significant challenge, as the deeper meanings embedded in poetry can easily be overlooked if only the literal interpretation is considered. However, with the advancements in these models, an important question arises: *Can expressive forms of language, such as poetry, be visually represented in a way that accurately captures their intended meaning?* To address this, we explore the potential of Diffusion Models and Large Language Models (LLMs) in the domain of poetry. We propose the *PoemToPixel* pipeline, which is built on the summarization capabilities of LLMs along with the image synthesis capabilities of Diffusion Models.

Poems frequently challenge the conventional narrative structures found in stories and paragraphs, prompting a unique approach to analysis and interpretation (Stibbs, 1995). Taking into consideration this attribute of poetry, our *PoemToPixel* framework adopts a two-step strategy for poem visualization, combining poem summarization and poem key element extraction. For summarization, we use the GPT-4o mini (Achiam et al., 2023) Model, which has shown impressive results across a variety of summarization tasks (Zhang et al., 2024; Goyal et al., 2022a). *Despite these favourable outcomes, our analysis revealed limitations in the summaries generated by GPT for creative content, highlighting the need for further refinement. As a result, we employed Prompt Tuning based on human feedback on the GPT-4o mini model to improve the quality of poem summaries.* These generated summaries enhance comprehension by transforming abstract concepts into more concrete insights, mak-

¹<https://github.com/SofeeyaJ/Poetry-In-Pixels-Coling2025>

ing the poet’s message clearer and more engaging. With these textual summaries available, they can be transformed into images that vividly convey the ideas of the poem. However, the complexity of crafting effective prompts poses a significant challenge for the average user. This difficulty arises from the major difference between user-provided natural language prompts and the keyword-enriched prompts required for the system’s high-quality output (Brade et al., 2023a; Witteveen and Andrews, 2022; Chen et al., 2023). *In order to bridge this gap, we propose the **PoeKey** algorithm, which extracts **KEY** elements from **POEM** summaries. These key elements extracted from the poems are then formulated into precise instructions, which are fed into diffusion models for image generation. Furthermore, refining user language into system language is crucial for enhancing the user experience (Hao et al., 2024; Betker et al., 2023; Brade et al., 2023a). With this in mind, we repeatedly modified prompts by utilizing a human feedback loop that incorporated both quantitative and qualitative measures. Based on this feedback, the prompts were further changed to achieve better image generation.* To summarize, we make the following key contributions:

- We present a novel approach of poem-to-image generation through our innovative **PoemToPixel** pipeline, a framework designed to produce multimodal outputs in the form of visual images.
- We introduce **PoeKey**, a novel algorithm designed to extract three key elements from poems, facilitating the creation of high-quality prompts for image generation.
- We implement a two-step prompt tuning process to analyze the summarization capability of GPT and the image generation capability of diffusion models.
- We introduce a novel multimodal dataset **MiniPo** consisting of 1001 children’s rhymes that integrate textual and visual elements, along with metadata, thereby enhancing the available resources for poem analysis across diverse age groups.

2 Related Works

Text-to-Image Prompting: With the emergence of Large Language Models (LLMs), the need for

efficient prompting has become more apparent. However, in tasks like text-to-image generation, the pre-training of LLMs does not fully address how prompts influence the image creation process. Automatically refining user input into system-optimized prompts is a crucial step toward developing more user-friendly text-to-image systems (Xie et al., 2023). Some approaches include interactive systems that assist users by suggesting immediate improvements to help refine their prompts, though manual adjustments are still often necessary (Feng et al., 2023; Brade et al., 2023b; Liu and Chilton, 2022). *To address this, our PoemToPixel approach combines human-guided prompt tuning with the interactive capabilities of LLMs to generate prompts that effectively drive text-to-image models, producing relevant images for poetry.*

Image and Poetry: There has been limited research focused on bridging the gap between poetry and images. A few studies have focused on generating poems from images using a deep coupled visual-poetic embedding model combined with RNN-based adversarial training (Liu et al., 2018). In the study by (Cheng et al., 2018), poetry was generated from images using recurrent neural networks trained on existing poems. By contrast, (Li et al., 2021) worked on generating images from poems, specifically focusing on creating paintings that replicated Feng Zikai’s unique artistic style, representing a distinct form of art altogether. However, current literature lacks approaches for generating images from poems, irrespective of genre, type, or style. *Our work addresses this gap by introducing the PoemToPixel approach, which captures the complex meanings within poems, resulting in relevant image generation.*

Diffusion Models: The field of text-to-image generation has seen significant advancements, particularly with the introduction of various diffusion model techniques (Ho et al., 2020; Ruiz et al., 2023; Saharia et al., 2022; Gu et al., 2022; Nichol et al., 2022). These models have proven to be highly effective at generating realistic and diverse images from textual prompts. However, as highlighted by (Ramesh et al., 2022), diffusion models often struggle with complex prompts that require advanced skills, such as binding attributes to objects and spatial reasoning. Recent works (Epstein et al., 2023; Kang et al., 2023) have aimed to enhance the capabilities of pre-trained diffusion models to better capture intricate details in textual descriptions. Despite these improvements, the models still fall

short in accurately representing the finer details in longer, more complex prompts. *Our PoemToPixel approach addresses this challenge by formulating prompts that convey intricate details in a concise and precise manner, improving the generation of relevant images from poetry.*

3 Corpus

In the existing literature, the datasets that are available for poetry generation are predominantly unimodal (text-based). To accomplish our research objectives, we employed two datasets.

Type	MiniPo	PoemSum
Number of Poems	1001	3011
Max Poem Length (in words)	346	6830
Avg. Poem Length (in words)	50.47	209
Max Summary Length (in words)	491	1104
Avg. Summary length (in words)	105.79	141

Table 1: Statistics of the two Datasets that we have used in our research.

PoemSum: The first dataset, *PoemSum* (Mahbub et al., 2023), contains 3011 poetry samples and their English-language summaries. In the *PoemSum* dataset, the poems are collected from multiple online sources, whereas the corresponding standard reference summaries were obtained from ‘Poem Analysis’², a website known for its comprehensive archive of high-quality poem summaries.

MiniPo: To facilitate and promote poetry research, we introduce a novel dataset, ‘*MiniPo*’, consisting of 1001 samples that address the lack of representation of children’s poetry in the existing dataset. Existing datasets covered various genres but did not include any children’s poems. ‘*MiniPo*’ aims to enhance diversity in poem analysis, allowing for a more comprehensive and inclusive study of poetry across different age groups and genres. Each sample in the dataset includes the poem’s title, the poem itself, and images depicting the poem’s plot. We will now outline the steps taken to create our dataset.

1) Data Collection: The *MiniPo* dataset was curated from various online open sources, with a specific focus on nursery rhymes.^{3 4} The choice to focus on nursery poems was inspired by their concise structure and simple meanings, making them

²<https://poemanalysis.com/>

³<https://www.nurseryrhymes.org/nursery-rhymes.html>

⁴<https://smart-central.com/>

ideal for both present and future poetry research.

To ensure the reliability and accuracy of the data collection process, a group of three English-proficient final-year undergraduate students were selected. The selection criteria ensured that these individuals possessed the necessary technical skills to verify the collected data, specifically poems sourced from various online websites.

2) Data Cleaning: After collecting data from various sources, additional processing was required to ensure dataset consistency. We extracted textual content from nursery rhymes and presented it as metadata to our dataset for future manipulation. The refined records were then verified to ensure that any inconsistencies were addressed. We have provided different statistics of the *PoemSum* and *MiniPo* dataset in Table 1.

4 Proposed Methodology

4.1 Problem Statement

Given a dataset $\mathcal{D} = \{P_1, P_2, \dots, P_n\}$, where each data point P_i represents a poem, the objective is to generate a corresponding set of images $\{I_{P_1}, I_{P_2}, \dots, I_{P_n}\}$ such that each image I_{P_i} visually represents the interpretation of the poem P_i . The textual poems P_i are summarized to S_i using a summarization module, $S_i = f_{summ}(P_i)$, where f_{summ} represents the summarization function of the LLM. Leveraging the summary S_i , key components of poems such as themes, emotions, and visual elements are extracted, $E_i = f_{KeyExtraction}(S_i)$, utilizing the poem embeddings and similarity scores. The extracted key components E_i are then fed into an image generation framework \mathcal{F} to produce the images I_{P_i} , $I_{P_i} = \mathcal{F}(E_i)$, where \mathcal{F} is the image generation function applied to the key components. The aim is to ensure that each generated image I_{P_i} accurately reflects the visual tone of its corresponding poem P_i .

4.2 PoemToPixel Framework

To make poetry more understandable and visually appealing to readers, we introduce the *PoemToPixel* framework, which converts textual poems into visual illustrations. In this section, we present our *PoemToPixel* pipeline, which allows for training-free image generation that accurately captures the meanings embedded in the verses. The overall procedure of our method is divided into the following phases:

4.2.1 Phase 1: Summarization Module

The complex nature of poetic language presents a unique set of challenges during summarization. Poems often contain rich, multifaceted language, such as metaphors, similes, and symbolic imagery, making it difficult to capture their essence in a concise summary without losing the intended meaning or emotional impact. The emergence of Large Language Models (LLMs) such as GPT creates novel possibilities for innovation in this field. While GPT has demonstrated strong performance in a variety of text summarization tasks, like long document summarization (Bajaj et al., 2021), dialogue summarization (Zhu et al., 2024), and news summarization (Goyal et al., 2022b), its capabilities in the genre of poetry are largely unexplored. To address this, we prompted GPT-4o mini model to summarize poems. However, the quality of GPT’s output is highly dependent on the input prompts (Antar, 2023), making the design of effective prompts important. As shown in Figure 1, we implemented prompt tuning based on feedback, aligning our prompts through both qualitative (human feedback) and quantitative measures. The refinement process for these prompts is detailed below, and we repeated the process iteratively based on the results of these evaluations.

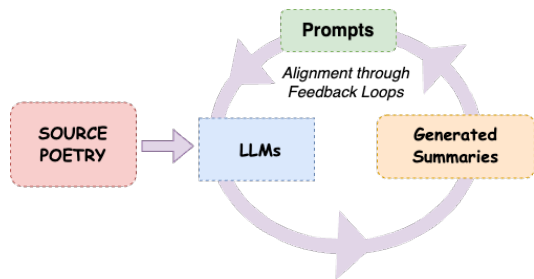


Figure 1: A framework of an iterative process with prompts refined based on feedback for improved summarization.

Prompt Tuning through Human Feedback: Automatic evaluation metrics can often be insufficient for summarization tasks, particularly in the context of poetry, where meaning is deeply embedded and comprehension is varied. To address this issue, we collaborated with a team of four professional experts from the Poetry Society of the Nation, which included one renowned poet, two university professors, and one PhD in visual and literary arts. We chose 10% of the samples from the *PoemSum* dataset for prompt tuning. We asked the experts to rate the summaries generated by GPT-4o mini,

	Expert Scores	Rouge - L	Meteor
Prompt 1	0	0.1720	0.1488
Prompt 2	2	0.1717	0.1564
Prompt 3	2	0.1770	0.1827
Prompt 4	4	0.1811	0.1825
Prompt 5	4	0.1821	0.1855
Prompt 6	4	0.1801	0.2140
Prompt 7	2	0.1779	0.2041

Table 2: Sum of Expert Scores provided for summarization based on quality of summaries generated in each round in comparison to gold reference summaries.

assigning either a +1 reward or a -1 penalty based on the quality of the summaries compared to gold reference summaries. The process was repeated until the overall score decreased after any round. Table 2 shows the total scores after each round of summary generation, given the particular prompt. Based on expert feedback, we refined the prompts for positive scores by modifying keywords, focusing on specific details, or highlighting key aspects. To maintain fairness and avoid bias against AI-generated summaries, we did not inform the experts that the summaries were generated by GPT. After six rounds of iterative feedback, the score decreased in the seventh round, after which we discontinued the process. We concluded that the sixth prompt \mathcal{R}_6 was the most effective for poem summarization using GPT-4o mini.

\mathcal{R}_6 : Summarize below, covering the main theme, mood, and any notable literary devices used by the poet.

All the subsequent prompts used in our Prompt tuning through feedback are listed in the Appendix A.3, with the total expert scores provided in Table 2.

Prompt Tuning through Automated Evaluation: Apart from human evaluation, automated evaluation is often used for its efficiency and reliability. Most summarization tasks are based on metrics like ROUGE (Lin, 2004), BLEU (Papineni et al., 2002), and METEOR (Banerjee and Lavie, 2005) scores. For our automatic evaluation, we used ROUGE and METEOR scores to assess the performance of summaries generated by various prompts, as shown in Table 2. The results indicate that Round 6 achieved impressive scores in both metrics. However, the differences in scores among the seven prompts within the same model were relatively minor. Based on these findings, we conducted further experiments on the entire dataset using the finalized prompt \mathcal{R}_6 .

4.2.2 Key Element Extraction Unit

We implement our proposed *PoeKey* algorithm to extract emotions, visual elements, and themes from poems, which are necessary for shaping and guiding the visual representation of the poem. Emotions within the poems are identified using an emotion classifier (Hartmann, 2022) that assigns an emotion tag based on the most dominant sentiment depicted in the poem summary. Additionally, visual elements are identified by utilizing nouns, pronouns, and named entities described within the poem summaries. For theme extraction, we compute the sentence embedding of the poem summary using Sentence-BERT (Reimers, 2019). Cosine similarity is then calculated between the poem’s embedding and each predefined theme embedding, with the theme corresponding to the highest similarity score assigned to the poem. These themes, which commonly recur across various literary works and genres, were identified using relevant online resources^{5 6} that highlight major themes in poetry. Additionally, we collaborated with literary scholars specializing in literature to create concise descriptions for each theme.

4.2.3 Instruction Generator

The key components extracted from the poems are then provided to an instruction generation module. This module incorporates GPT-4o mini and the SDXL Turbo diffusion model (Sauer et al., 2023). We selected the SDXL Turbo diffusion model for its cost-effective and time-efficient image generation process. As shown in Figure 2, the process begins by providing GPT-4o mini with a prompt and the extracted elements (Visual Elements, Themes, and Emotions) as the context. The output is an instruction for creating an image corresponding to a particular poem, which will be fed to a diffusion model for image generation.

Prompt to GPT: *Create a prompt which can generate an image which will be able to represent the poem using the given emotion and visual elements and theme of the poem, and the prompt should be under 50 words.*

Instruction by GPT: *To convey the emotion of sadness in the context of mortality and the afterlife, I envision an image that captures a dense jungle setting, symbolizing the overwhelming presence of*

⁵<https://www.poetryfoundation.org/topics-themes>

⁶<https://poemanalysis.com/poetry-explained/poetry-themes/>

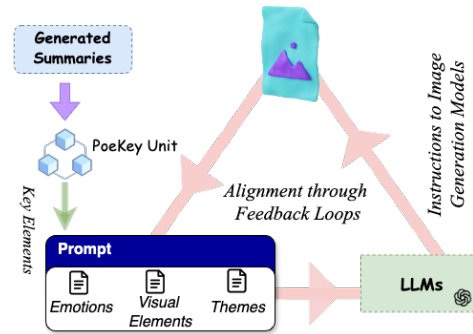


Figure 2: A framework of an iterative process of image prompts refined based on feedback.

life and death intertwined. The focal point would be a solitary parrot perched on a gnarled branch.

The diffusion model then generates an image based on the instructions. We evaluated the images both automatically and manually described below, and we improved our prompts based on feedback.

Human Evaluation Assessing the quality of generated images for poems requires human evaluations, as the images must capture the embedded meanings and emotions of the poetry. To conduct this evaluation, we collaborated with the aforementioned four area experts. We selected 5% of the samples from the *PoemSum* dataset. The experts reviewed the images produced by the diffusion model to identify whether they captured the emotion, themes, and meaning of the poems. Each expert rated the images on a scale of 1 to 5, and the average score was calculated after each round. A higher score indicated better alignment with the poems, whereas a lower score indicated misalignment. Based on expert feedback, we refined the GPT prompts, which in turn adjusted the instruction prompts and generated images based on those instructions as demonstrated in Figure 3.

Updated Prompt to GPT: *Write a prompt for an image that visually represents the poem’s mood and themes using the given emotional tone and visual details. Keep the prompt concise, under 50 words.*

Updated Instruction by GPT: *Create an image of a vibrant parrot perched on a gnarled branch in a dense, shadowy jungle. Surround it with wilted plants and hints of decay, embodying a sense of waiting and longing, reflecting mortality and the somber afterlife.*

This iterative process resulted in a gradual improvement in the coherence of the images and poems as outlined in Table 3. After five rounds of



Figure 3: An instance of image generation phases during instruction tuning.

feedback, the average score began to decline in the sixth round, at which point we stopped the process and finalized the fifth prompt.

\mathcal{I}_5 : *Formulate a prompt to generate an image that reflects the poem’s emotional depth and themes using the specified visual elements, keeping the prompt within 50 words.*

All the prompts used for prompt tuning in image generation are shown in the Appendix A.4, along with the average expert scores outlined in Table 3.

Automated Evaluation: Along with human evaluation, we also performed automated evaluation to reduce the time and cost involved in manual expert review. For this, we used two well-established quantitative metrics: Image-Text Matching Loss (ITM) (Li et al., 2022) and Image-Text Contrastive Loss (ITC) (Li et al., 2022). The results shown in Table 3 demonstrate that the \mathcal{I}_5 prompt achieved the highest scores in both metrics. Based on these results, we proceed with the image generation task on the entire dataset using this finalized prompt.

	Expert Scores	ITM	ITC
Instruction 1	1.75	0.111	0.221
Instruction 2	3	0.330	0.246
Instruction 3	3	0.333	0.272
Instruction 4	3.75	0.310	0.269
Instruction 5	4.25	0.356	0.276
Instruction 6	4	0.334	0.263

Table 3: Expert scores provided for image generation based on alignment of images with respect to poems.

5 Experiment and Results

5.1 Phase 1: Summarization

Comparison with the Baselines: The experiments were conducted using GPT-4o mini and included comparisons with popular summarization models such as T5 (Raffel et al., 2020), BART (Lewis,

2019), Pegasus (Zhang et al., 2020), and Falcon (Almazrouei et al., 2023) on both datasets. In recent years, these models have demonstrated impressive performance across various summarization tasks (Laskar et al., 2022; Ravaut et al., 2022). For the *PoemSum* dataset, we evaluated the model-generated summaries against the gold reference summaries using a range of established metrics, including ROUGE-1 (R1), ROUGE-2 (R2), ROUGE-L (RL), BLEU-1, BLEU-2, BLEU-3, BLEU-4, and METEOR, as shown in Table 4. These metrics collectively provide a comprehensive assessment of the accuracy, fluency, and overall quality of the generated summaries across all models.

Scores	Baselines				GPT-4o mini	
	Pegasus	Bart	T5	Falcon	Without PT	With PT
R1	0.0167	0.2006	0.1915	0.2065	0.2921	0.3188
R2	0.0015	0.0421	0.0356	0.0374	0.0428	0.0586
RL	0.0141	0.1343	0.1279	0.1321	0.1720	0.1801
BLEU-1	0.0495	0.1837	0.1812	0.2313	0.5327	0.6541
BLEU-2	0.0153	0.1323	0.1324	0.1706	0.3986	0.5107
BLEU-3	0.0067	0.0729	0.0703	0.0916	0.2318	0.3102
BLEU-4	0.0036	0.0419	0.0389	0.0502	0.1376	0.1867
METEOR	0.0111	0.1041	0.0998	0.1089	0.1488	0.2140

Table 4: Quantitative evaluation of generated summaries from different models on the *PoemSum* Dataset using Different approaches. Here, PT refers to Prompt Tuning with Feedback Loops

Key Observations: The results presented in Table 4 demonstrate that summaries produced using GPT-4o mini with Prompt Tuning outperformed those generated by all base models, including GPT-4o mini models without Prompt Tuning. Furthermore, as shown in Table 5, applying Prompt Tuning improved the generated summaries, making them more reflective of the poem’s underlying meaning and more similar to the gold reference summaries. In contrast, summaries generated without Prompt Tuning failed to capture the true interpretation of the poem. Additionally, the base models were unable to provide meaningful summaries, often producing identical verses with minimal variation.

5.2 Phase 2: Image Generation

Qualitative Evaluation: Figure 4 presents a detailed comparison of images produced by different approaches for various poems using SDXL Base 1.0 (Podell et al., 2023). *The complete poems mentioned here are provided in the Appendix A.6.* Our qualitative analysis focused on three key aspects:

1) Theme: The *PoemToPixel* Pipeline demonstrated a remarkable ability to capture the emotional theme and setting of poetry verses. As shown

Poem	Reference Summary	Baseline (Falcon) Summary	GPT
<p><i>Lovely dainty Spanish needle With your yellow flower and white, Dew bedecked and softly sleeping, Do you think of me to-night? Shadowed by the spreading mango, Nodding o'er the rippling stream, Tell me, dear plant of my childhood, Do you of the exile dream? Do you see me by the brook's side Catching crayfish 'neath the stone, As you did the day you whispered: Leave the harmless dears alone? Do you see me in the meadow Coming from the woodland spring With a bamboo on my shoulder And a pail slung from a string? Do you see me all expectant Lying in an orange grove, While the swee-swees sing above me, Waiting for my elf-eyed love? Lovely dainty Spanish needle, Source to me of sweet delight, In your far-off sunny southland Do you dream of me to-night?</i></p>	<p>The Spanish Needle' by Claude McKay is a beautiful, nostalgic poem that looks back on a speaker's childhood. In the stanzas, the speaker asks the Spanish needle, a type of plant, if it remembers him during his youth. He has fond memories of spending time outside, near a river, carrying water, and spending time with his love. It's clear he's feeling nostalgic about his past and hoping that through the plant, a tiny bit of himself still exists in the past.</p>	<p>Lovely dainty Spanish needle With your yellow flower and white, Dew bedecked and softly sleeping, Do you think of me to-night? Shadowed by the spreading mango, Nodding o'er the rippling stream, Tell me, dear plant of my childhood . Do you see me in the meadow Coming from the woodland spring With a bamboo on my shoulder And a pail slung from a string?</p>	<p>1) Prompt Tuning through Feedback: In this nostalgic poem, the speaker addresses a dainty Spanish needle flower, reminiscing about their childhood adventures and longing for a connection with the past. The emotional tone is one of wistfulness and yearning, as the speaker questions whether the flower remembers them and dreams of their joyful moments together. Ultimately, the poem reflects a deep sense of nostalgia and a desire for companionship, even across distance and time.</p> <p>2) Without Prompt Tuning: The poem reflects nostalgia and longing through the perspective of a person reminiscing about their childhood in a sunny, southern land. The speaker addresses a Spanish needle plant, reminiscing about their past experiences and childhood activities such as catching crayfish, fetching water, and waiting for a loved one. The plant, with its yellow and white flowers, is seen as a symbol of the speaker's lost homeland and past joys.</p>

Table 5: Sample of the Generated Summary by the Baseline (Falcon) model and GPT-4o mini Model after Prompt Tuning through Feedbacks with respect to gold reference summaries.

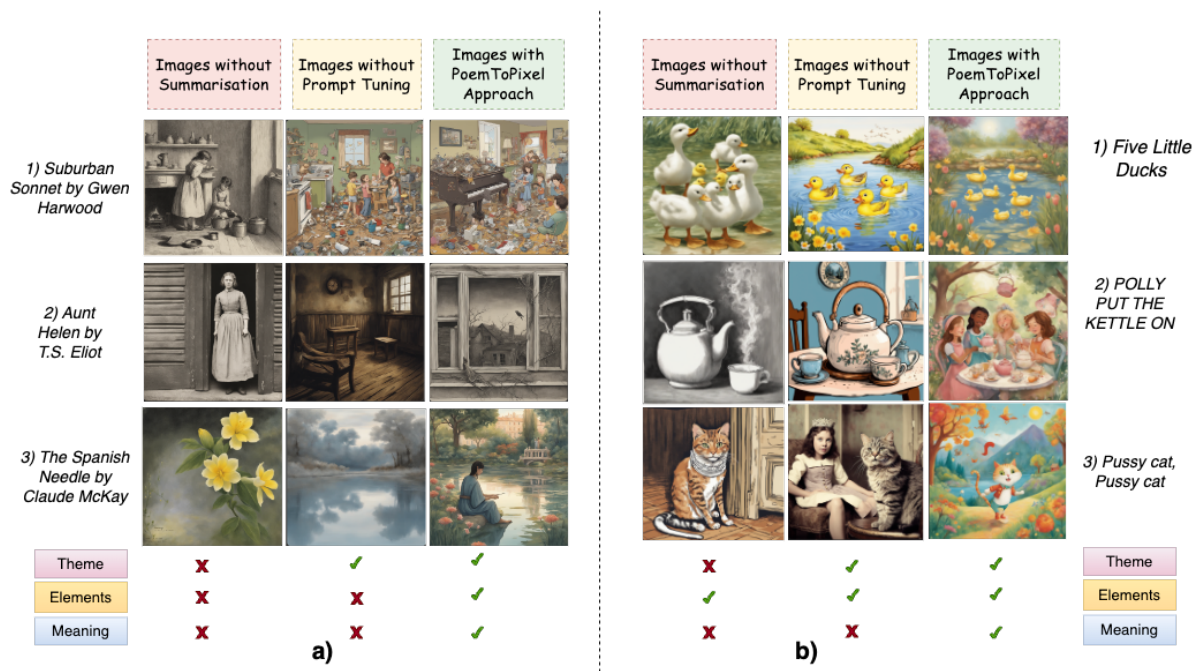


Figure 4: A comparison of Poem to Image generation using SDXL Base Diffusion Model with different methods and *PoemToPixel* Approach on (a) *PoemSum* dataset (b) *MiniPo* ((a collection of Nursery Rhymes) dataset).

in Figure 4 (a), without the summarization approach, the images generated struggled to convey the intended theme or sentiment of the poem. In contrast, when using the summarization approach, the emotional tone of the poem was accurately reflected in the generated images. Similarly, for our proposed *MiniPo* dataset, as shown in Figure 4 (b), images generated without summarization did not reflect the poem's setting as depicted in the verses, whereas images generated with summarization successfully captured the joyous theme (for instance,

Poem 2 and Poem 3 in Figure 4 (b)).

2) Visual Elements: The *PoemToPixel* Framework demonstrated an exceptional ability to capture the visual aspects of the poems due to the incorporation of our proposed *PoeKey* algorithm, which extracted the visual elements described in the poem. For instance, in Figure 4 (b), Poem 3 illustrates its capacity to depict aspects such as flowers, rivers, and the act of sitting by the river-side, precisely representing the imagery described in the poem. In contrast, the images produced

without *PoemToPixel* framework failed to capture these essential visual elements, as shown in Poem 3 in Figure 4(a). **3) Meaning:** Figures 4 (a) and 4 (b) demonstrate the *PoemToPixel* Framework’s potential to efficiently comprehend and convey the deeper meaning of the poems. This effectiveness is primarily due to the summarization process, which captures the complex meanings present in the verses, leading to the formation of precise instructions for image generation models. In contrast, images generated without prompt tuning or summarization failed to visually represent the poem’s meaning. For example, in Figure 4 (b), for the famous nursery rhyme “Polly, Put the Kettle On,” the *PoemToPixel* framework accurately captures the essence of “everyone drinking tea,” as well as the key element—the kettle. On the other hand, other approaches only managed to depict the kettle, missing the broader context of the poem.

Quantitative Evaluation: The results presented in Table 6 provide valuable insights into the performance of various image generation approaches, highlighting the advantages of each method in the context of our task. The key findings include:

1) *PoemToPixel*: Our proposed approach outperformed all other methods, achieving the highest ITM (Image-Text Matching) and ITC (Image-Text Coherence) scores for both datasets. This superior performance can be attributed to our two-phase image generation process, in which summarization effectively captured the poem’s meaning, and our proposed *PoeKey* algorithm extracted the core elements of the poem to incorporate into the generated images. This combination greatly improved the alignment of the poem and its visual representation.

2) Performance of Different Approaches: Besides the *PoemToPixel* method, we compared two other approaches: (i) generating images directly from the poem without summarization, and (ii) generating images without prompt tuning, where the image models received an initial instruction without any feedback mechanism. The results indicate that image generation models, when directly provided with the poem text, faced challenges grasping the intricate and hidden meanings within the verses, leading to irrelevant images. Moreover, without prompt tuning, the models underperformed compared to those with feedback, showing a significant improvement of 34.97% in ITM and 27.4% in ITC scores when using prompt tuning. This highlights the importance of feedback mechanisms in producing accurate and contextually relevant images.

Dataset	Scores	Without summary	Without PT	With <i>PoemToPixel</i>
PoemSum	ITM	0.123	0.251	0.386
	ITC	0.205	0.274	0.312
MiniPo	ITM	0.333	0.508	0.618
	ITC	0.306	0.329	0.364

Table 6: Quantitative Comparisons of Image Generation Approach for Poems on *PoemSum* and *MiniPo* Dataset using SDXL Base Diffusion Model. Here, PT refers to Prompt Tuning.

3) Dataset-Based Performance: We conducted image generation experiments on two datasets, *PoemSum* and *MiniPo*, which varied in size, genre, form, and style. As shown in Table 6, the *PoemToPixel* pipeline achieved 37.54% higher ITM and 14.28% higher ITC scores in the *MiniPo* dataset than in the *PoemSum* dataset. This is possibly due to the simpler themes of nursery rhymes in the *MiniPo* dataset, in contrast to the more complex meanings found in the poems present in the *PoemSum* dataset. Additionally, the shorter length of *MiniPo* poems (an average of 58 words) compared to *PoemSum* (an average of 209 words) made it easier for the model to generate relevant imagery. Despite these advantages in the *MiniPo* dataset, other methods still struggled to produce images that aligned with the rhymes, whereas our *PoemToPixel* approach delivered strong results across both datasets.

6 Conclusion

In conclusion, the *PoemToPixel* framework makes a valuable contribution to the field of poetic image generation. It demonstrates the ability to interpret the intricate meanings embedded in poems and generate corresponding visuals. Our *PoeKey* algorithm further enhances this by accurately capturing the emotions, themes, and visual elements necessary for an authentic representation of poetic imagery. Additionally, we introduce the *MiniPo* Dataset, consisting of 1001 nursery rhymes, each paired with images generated by the *PoemToPixel* framework, establishing a new multimodal dataset for children’s rhymes. Looking ahead, we aim to expand our research by generating multiple images that depict different segments of a single poem.

7 Limitations

While the *PoemToPixel* framework effectively captures the meanings, emotions, and visual elements in poems, it faces challenges when multiple emotions or meanings are expressed within a single

poem, as one image cannot encompass all these aspects. Additionally, the framework lacks language-agnostic capabilities, as it currently only supports English and does not extend to other languages.

8 Ethical Consideration

A key ethical consideration involves the inherent biases present in diffusion models, which may reflect societal, cultural, or data-driven biases from the pre-trained models. These biases can potentially influence the generation of images related to poems on specific topics or forms, resulting in unfair or inappropriate outputs.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ebtessam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, M erouane Debbah,  tienne Goffinet, Daniel Hessel, Julien Launay, Quentin Malartic, et al. 2023. The falcon series of open language models. *arXiv preprint arXiv:2311.16867*.
- D Antar. 2023. The effectiveness of using chatgpt4 in creative writing in arabic: Poetry and short story as a model. *Inf. Sci. Lett.*, 12:2445–2459.
- Ahsaas Bajaj, Pavitra Dangati, Kalpesh Krishna, Pradhiksha Ashok Kumar, Rheeya Uppaal, Bradford Windsor, Eliot Brenner, Dominic Dotterrer, Rajarshi Das, and Andrew McCallum. 2021. Long document summarization in a low resource setting using pretrained language models. *arXiv preprint arXiv:2103.00751*.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, et al. 2023. Improving image generation with better captions. *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2(3):8.
- Stephen Brade, Bryan Wang, Mauricio Sousa, Sageev Oore, and Tovi Grossman. 2023a. Promptify: Text-to-image generation through interactive prompt exploration with large language models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–14.
- Stephen Brade, Bryan Wang, Mauricio Sousa, Sageev Oore, and Tovi Grossman. 2023b. Promptify: Text-to-image generation through interactive prompt exploration with large language models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–14.
- Yang Chen, Yingwei Pan, Yehao Li, Ting Yao, and Tao Mei. 2023. Control3d: Towards controllable text-to-3d generation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 1148–1156.
- Wen-Feng Cheng, Chao-Chung Wu, Ruihua Song, Jianlong Fu, Xing Xie, and Jian-Yun Nie. 2018. *Image inspired poetry generation in xiaoice*. *Preprint*, arXiv:1808.03090.
- Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. 2021. Cogview: Mastering text-to-image generation via transformers. *Advances in neural information processing systems*, 34:19822–19835.
- Dave Epstein, Allan Jabri, Ben Poole, Alexei A. Efros, and Aleksander Holynski. 2023. *Diffusion self-guidance for controllable image generation*. *Preprint*, arXiv:2306.00986.
- Yingchaojie Feng, Xingbo Wang, Kam Kwai Wong, Sijia Wang, Yuhong Lu, Minfeng Zhu, Baicheng Wang, and Wei Chen. 2023. Promptmagician: Interactive prompt engineering for text-to-image creation. *IEEE Transactions on Visualization and Computer Graphics*.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022a. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022b. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.
- Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. 2022. *Vector quantized diffusion model for text-to-image synthesis*. *Preprint*, arXiv:2111.14822.
- Yaru Hao, Zewen Chi, Li Dong, and Furu Wei. 2024. Optimizing prompts for text-to-image generation. *Advances in Neural Information Processing Systems*, 36.
- Jochen Hartmann. 2022. Emotion english distilroberta-base. <https://huggingface.co/j-hartmann/emotion-english-distilroberta-base/>.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. *Denosing diffusion probabilistic models*. *CoRR*, abs/2006.11239.
- Wonjun Kang, Kevin Galim, and Hyung Il Koo. 2023. *Counting guidance for high fidelity text-to-image synthesis*. *Preprint*, arXiv:2306.17567.

- Md Tahmid Rahman Laskar, Enamul Hoque, and Jimmy Xiangji Huang. 2022. Domain adaptation with pre-trained transformers for query-focused abstractive text summarization. *Computational Linguistics*, 48(2):279–320.
- M Lewis. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Dan Li, Shuai Wang, Jie Zou, Chang Tian, Elisha Nieuwburg, Fengyuan Sun, and Evangelos Kanoulas. 2021. Paint4poem: A dataset for artistic visualization of classical chinese poems. *arXiv preprint arXiv:2109.11682*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Bei Liu, Jianlong Fu, Makoto P Kato, and Masatoshi Yoshikawa. 2018. Beyond narrative description: Generating poetry from images by multi-adversarial training. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 783–791.
- Vivian Liu and Lydia B Chilton. 2022. Design guidelines for prompt engineering text-to-image generative models. In *Proceedings of the 2022 CHI conference on human factors in computing systems*, pages 1–23.
- Ridwan Mahbub, Ifrad Khan, Samiha Anuva, Md Shihab Shahrar, Md Tahmid Rahman Laskar, and Sabbir Ahmed. 2023. **Unveiling the essence of poetry: Introducing a comprehensive dataset and benchmark for poem summarization**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14878–14886, Singapore. Association for Computational Linguistics.
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2022. **Glide: Towards photorealistic image generation and editing with text-guided diffusion models**. *Preprint*, arXiv:2112.10741.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. **Hierarchical text-conditional image generation with clip latents**. *Preprint*, arXiv:2204.06125.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr.
- Mathieu Ravaut, Shafiq Joty, and Nancy F Chen. 2022. Summareranker: A multi-task mixture-of-experts re-ranking framework for abstractive summarization. *arXiv preprint arXiv:2203.06569*.
- N Reimers. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023. **Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation**. *Preprint*, arXiv:2208.12242.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. 2022. **Photo-realistic text-to-image diffusion models with deep language understanding**. *Preprint*, arXiv:2205.11487.
- Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. 2023. Adversarial diffusion distillation. *arXiv preprint arXiv:2311.17042*.
- Andrew Stibbs. 1995. The specialness of poetry. *English in Education*, 29(3):14–20.
- Sam Witteveen and Martin Andrews. 2022. Investigating prompt engineering in diffusion models. *arXiv preprint arXiv:2211.15462*.
- Yutong Xie, Zhaoying Pan, Jinge Ma, Luo Jie, and Qiaozhu Mei. 2023. A prompt log analysis of text-to-image generation systems. In *Proceedings of the ACM Web Conference 2023*, pages 3892–3902.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International conference on machine learning*, pages 11328–11339. PMLR.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B Hashimoto. 2024. Benchmarking large language models for news

summarization. *Transactions of the Association for Computational Linguistics*, 12:39–57.

Rongxin Zhu, Jey Han Lau, and Jianzhong Qi. 2024. Factual dialogue summarization via learning from large language models. *arXiv preprint arXiv:2406.14709*.

A Appendix

A.1 MiniPo Dataset

This section presents instances of poems in Table 7 from our proposed multimodal *MiniPo* dataset, along with the corresponding images in Figure 5 and Figure 6 generated by our *PoemToPixel* framework.

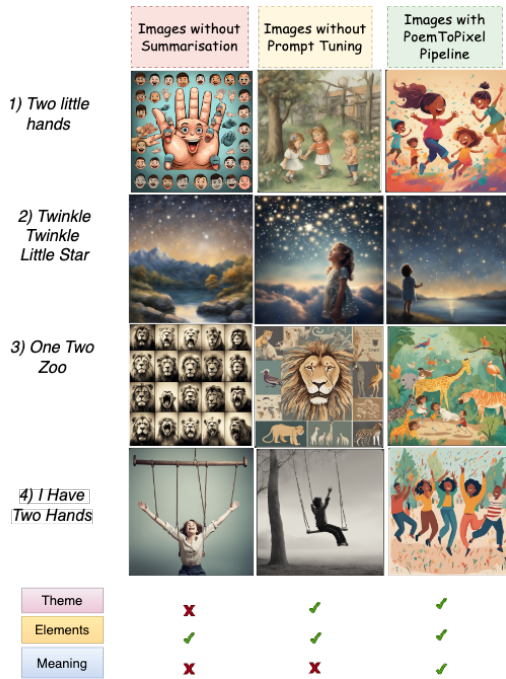


Figure 5: A comparison of poem to image generation using SDXL Base Diffusion Model for poem mentioned in Table 7 with different methods and *PoemToPixel* Approach

A.2 PoemSum Dataset

This section presents instances of poems in Table 8 from *PoemSum* dataset, along with the corresponding images in Figure 7 generated by our *PoemToPixel* framework.

A.3 Prompt Tuning for Summarisation

Summaries were generated using different prompts and compared to the gold reference summaries from the *PoemSum* dataset. Experts evaluated the quality of the generated summaries against the gold references, assigning a +1 for high quality and -1



Figure 6: A comparison of poem to image generation using SDXL Base Diffusion Model for poem mentioned in Table 7 with different methods and *PoemToPixel* Approach

MiniPo Dataset	
Title	Poems
1) Two little hands	Two little hands go Clap clap clap Two little feet goTap tap tap Two little eyes That open wide One little head Nods side to side
2) Twinkle Twinkle Little Star	Twinkle, twinkle little star, how I wonder what you are. Up above the world so high, like a diamond in the sky. Twinkle, twinkle little star, how I wonder what you are. One, two, one, two lets visit the zoo Three, four, three, four hear the lion roar Five, six, five, six watch the monkey tricks Seven, eight, seven, eight peacock looks so great Nine, ten, nine, ten we shall come again.
3) One Two Zoo	I have two arms that swing like this, (thrice) I have two arms that swing like this, Swing, Swing, Swing. I have two hands that clap like this,(thrice) I have two hands that clap like this, Clap, Clap, Clap. I have little fingers that wriggle like this, (thrice) I have little fingers that wriggle like this, Wriggle, Wriggle, Wriggle. I have two legs that jump like this, (thrice) I have two legs that jump like this, Jump, Jump, Jump. I have two feet that march like this, (thrice) I have two feet that march like this, March, March, March. I have little toes that hop like this, (thrice) I have little toes that hop like this,Hop, Hop, Hop.
5) Bits of Paper	Bits of paper, Bits of paper Lying on the floor, Lying on the floor Makes the place untidy, Makes the place untidy Pick them up, Pick them up Collect all the papers, Collect all the papers Where shall we throw?Where shall we throw? Throw them in the dustbin, Throw them in the dustbin Now the place is clean, Now the place is clean
6) See me	See me skip, see me run, I am going to school Like everyone. See me shout, See me wave, I am off to school Bye, everyone. See me hug, See me dance, I see my friends Hello, everyone! See me smile, See me grin, When the bell rings I walk in. See me jump, See me play. I am in S.P.M. Public SchoolHip Hip Hooraay!!!
7) Once a Lion caught a mouse	Once a lion caught a mouse. The mouse begged the lion not to eat him. Someday, he said, he would help the lion. The lion set the mouse free.
8) I have a tooth brush neat and gay	One day, the lion got caught in a hunter's net. The mouse came to help lion. He cut the net with his sharp teeth. The lion was set free. The lion and the mouse became good friends. I have a toothbrush, neat and gay, To brush my teeth with, everyday. I brush them each morning, Then I brush them each night; Till all are shining, Clean and bright.

Table 7: List of Poems of MiniPo Dataset

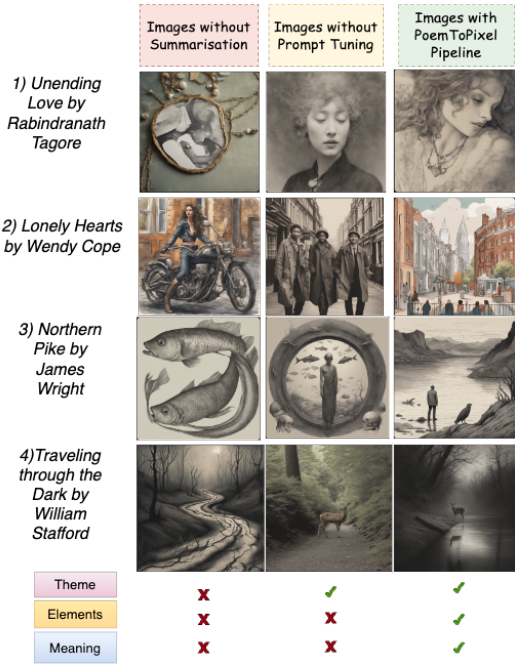


Figure 7: A comparison of poem to image generation using SDXL Base Diffusion Model for poem mentioned in Table 8 with different methods and *PoemToPixel* Approach

for low quality. The expert scores for each round are presented in the Table 9 below.

A.4 Prompt Tuning for Image Generation

Images were generated for 5% of the *PoemSum* dataset using various instructions produced by GPT in response to input prompts, as shown in the Table 10. Experts evaluated the images based on how effectively they conveyed the meaning of the poems, assigning scores on a scale of 1 to 5. Based on the total scores, the input prompts were refined, leading to updated instructions from GPT, which were then used to regenerate the images.

A.5 Description of evaluation metrics used for summarization

ROUGE-1 measures the overlap of unigrams, while ROUGE-2 evaluates bigram overlap, and ROUGE-L assesses the longest common subsequence between the generated and reference texts, focusing on both sequence similarity and fluency. The BLEU scores assess the precision of n-gram overlaps, with BLEU-1 for unigrams, BLEU-2 for bigrams, BLEU-3 for trigrams, and BLEU-4 for fourgrams. Additionally, METEOR computes the alignment of unigrams and synonym matches between the generated and reference texts, incorporating precision, recall, and stemming to assess overall

PoemSum Dataset	
Title	Poems
1) <i>Unending Love</i> by Rabindranath Tagore	<p><i>I seem to have loved you in numberless forms, numberless times... In life after life, in age after age, forever. My spellbound heart has made and remade the necklace of songs. That you take as a gift, wear round your neck in your many forms, In life after life, in age after age, forever.</i></p> <p><i>Whenever I hear old chronicles of love, its age-old pain, Its ancient tale of being apart or together. As I stare on and on into the past, in the end you emerge, Clad in the light of a pole-star piercing the darkness of time: You become an image of what is remembered, forever.</i></p>
2) <i>Lonely Hearts</i> by Wendy Cope	<p><i>You and I have floated here on the stream that brings from the fount. At the heart of time, love of one for another. We have played alongside millions of lovers, shared in the same Sky sweetness of meeting, the same distressful tears of farewell- Old love but in shapes that renew and renew forever.</i></p> <p><i>Today it is heaped at your feet, it has found its end in you The love of all man's days both past and forever: Universal joy, universal sorrow, universal life. The memories of all loves merging with this one love of ours - And the songs of every poet past and forever.</i></p> <p><i>Can someone make my simple wish come true? Make biker seeks female for touring fun, Do you live in North London? Is it you? Gay vegetarian whose friends are few, I'm into music, Shakespeare and the sun. Can someone make my simple wish come true? Executives in search of something new- Perhaps bisexual woman, artsy, young. Do you live in North London? Is it you? Successful, straight and solvent? I am too- Attractive Jewish lady with a son. Can someone make my simple wish come true? I'm Libran, inexperienced and blue- Need slim, non-smoker, under twenty-one. Do you live in North London? Is it you? Please write (with photo) to Box 152. Who knows where it may lead once we've begun? Can someone make my simple wish come true? Do you live in North London? Is it you?</i></p>
3) <i>Northern Pike</i> by James Wright	<p><i>All right. Try this, Then. Every body I know and care for, And every body Else is going To die in a loneliness I can't imagine and a pain I don't know. We had To go on living. We Untangled the net, we slit The body of this fish Open from the hinge of the tail To a place beneath the chin I wish I could sing of, I would just as soon we let The living go on living. An old poet whom we believe in Said the same thing, and so We paused among the dark cattails and prayed For the muskrats, For the ripples below their tails, For the little movements that we knew the crawdads were making under water.</i></p> <p><i>For the right-hand wrist of my cousin who is a policeman. We prayed for the game warden's blindness. We prayed for the road home. We ate the fish. There must be something very beautiful in my body. I am so happy.</i></p>
4) <i>Traveling through the Dark</i> by William Stafford	<p><i>Traveling through the dark I found a deer dead on the edge of the Wilson River road. It is usually best to roll them into the canyon: that road is narrow; to swerve might make more dead. By glow of the tail-light I stumbled back of the car and stood by the heap, a doe, a recent killing; she had stiffened already, almost cold. I dragged her off; she was large in the belt. My fingers touching her side brought me the reason- her side was warm; her fawn lay there waiting, alive, still, never to be born. Beside that mountain road I hesitated. The car aimed ahead its lowered parking lights: under the hood purred the steady engine. I stood in the glare of the warm exhaust turning red; around our group I could hear the wilderness listen. I thought hard for us all- my only swerving- then pushed her over the edge into the river.</i></p>

Table 8: List of Poems of PoemSum Dataset

translation quality.

A.6 Poems used in Qualitative Analysis of Poems

The complete poems, referenced by their titles in Figures 4 (a) and (b), are listed in the Table 11 below.

	Prompts	Score 1	Score 2	Score 3	Score 4	Total Score
\mathcal{R}_1	Summarize the following poem.	1	-1	1	-1	0
\mathcal{R}_2	Summarize the following poem in 2-3 sentences, focusing on the central theme and mood.	1	1	-1	1	2
\mathcal{R}_3	Summarize the poem in 2-3 sentences, ensuring to convey the poet's emotional tone and the underlying message.	1	1	-1	1	2
\mathcal{R}_4	Summarize the following poem, focusing on the central theme and mood.	1	1	1	1	4
\mathcal{R}_5	Summarize below, covering the main emotion, and mood of the poem.	1	1	1	1	4
\mathcal{R}_6	Summarize below, covering the main theme, mood, and any notable literary devices used by the poet.	1	1	1	1	4
\mathcal{R}_7	Summarize the poem below, capturing the poet's emotional tone, theme, mood and the underlying message.	1	1	1	-1	2

Table 9: Expert scores at the end of each round to mark the summaries generated by GPT, comparing them to the gold standard summaries using the provided prompts \mathcal{R}_i .

	Prompts	Score 1	Score 2	Score 3	Score 4	Total Score
\mathcal{I}_1	Create a prompt which can generate an image which will be able to represent the poem using the given emotion and visual elements and theme of the poem and the prompt should be under 50 words	2	2	1	2	1.75
\mathcal{I}_2	Create a prompt to generate an image that captures the poem's essence by combining the specified emotion, visual elements, and theme, ensuring the prompt is concise and under 50 words.	3	3	2	4	3
\mathcal{I}_3	Write a prompt for an image that visually represents the poem's mood and themes using the given emotional tone and visual details. Keep the prompt concise, under 50 words.	3	3	3	3	3
\mathcal{I}_4	Design a prompt that generates an image capturing the essence of the poem through a combination of the specified emotions, visual elements, and thematic undertones.	3	4	4	4	3.75
\mathcal{I}_5	Formulate a prompt to generate an image that reflects the poem's emotional depth and themes using the specified visual elements, keeping the prompt within 50 words.	3	5	4	5	4.25
\mathcal{I}_6	Compose a prompt capable of generating an image that visually interprets the poem, emphasizing the given emotions, key visual elements, and thematic elements.	3	4	4	5	4

Table 10: Expert scores at the end of each round for the images generated based on the instructions provided by GPT through these prompts \mathcal{I}_i .

PoemSum Dataset (Figure a)		MiniPo Dataset (Figure b)	
Title	Poems	Title	Poems
1) Suburban Sonnet by Gwen Harwood	<p><i>She practises a fugue, though it can matter to no one now if she plays well or not. Beside her on the floor two children chatter, then scream and fight. She hushes them. A pot boils over. As she rushes to the stove too late, a wave of nausea overpowers subject and counter-subject. Zest and love drain out with soapy water as she scours the crusted milk. Her veins ache. Once she played for Rubinstein, who yawned. The children caper round a sprung mousetrap where a mouse lies dead. When the soft corpse won't move they seem afraid. She comforts them; and wraps it in a paper featuring: Tasty dishes from stale bread.</i></p>	1) Five Little Ducks	<p><i>Five little ducks went swimming one day Over the hill and far away Mother duck said, "Quack quack quack quack" And only four little ducks came back!</i></p> <p><i>Four little ducks went swimming one day Over the hill and far away Mother duck said, "Quack quack quack quack" And only three little ducks came back!</i></p> <p><i>Three little ducks went swimming one day Over the hill and far away Mother duck said, "Quack quack quack quack" And only two little ducks came back!</i></p> <p><i>Two little ducks went swimming one day Over the hill and far away. Mother duck said, "Quack quack quack quack" And only one little duck came back!</i></p> <p><i>One little duck went swimming one day Over the hill and far away Mother duck said, "Quack quack quack quack" And all her five little ducks came back!</i></p>
2) Aunt Helen by T.S. Eliot	<p><i>Miss Helen Slingsby was my maiden aunt, And lived in a small house near a fashionable square Cared for by servants to the number of four. Now when she died there was silence in heaven And silence at her end of the street. The shutters were drawn and the undertaker wiped his feet— He was aware that this sort of thing had occurred before. The dogs were handsomely provided for, But shortly afterwards the parrot died too. The Dresden clock continued ticking on the mantelpiece, And the footman sat upon the dining-table Holding the second housemaid on his knees— Who had always been so careful while her mistress lived.</i></p>	2) POLLY PUT THE KETTLE ON	<p><i>Polly, put the kettle on, Polly, put the kettle on, Polly, put the kettle on, We'll all have tea. Sukey, take it off again, Sukey, take it off again, Sukey, take it off again, They've all gone away. Polly, put the kettle on, Polly, put the kettle on, Polly, put the kettle on, We'll all have tea. Blow the fire and make the toast, Put the muffins on to roast, Blow the fire and make the toast, We'll all have tea. Polly, put the kettle on, Polly, put the kettle on, Polly, put the kettle on, We'll all have tea. Sukey, take it off again, Sukey, take it off again, Sukey, take it off again, Sukey, take it off again, They've all gone away.</i></p>
3) The Spanish Needle by Claude McKay	<p><i>Lovely dainty Spanish needle With your yellow flower and white, Dew bedecked and softly sleeping, Do you think of me to-night? Shadowed by the spreading mango, Nodding o'er the rippling stream, Tell me, dear plant of my childhood, Do you of the exile dream? Do you see me by the brook's side Catching crayfish 'neath the stone, As you did the day you whispered: Leave the harmless dears alone? Do you see me in the meadow Coming from the woodland spring With a bamboo on my shoulder And a pail slung from a string? Do you see me all expectant Lying in an orange grove, While the swee-swees sing above me, Waiting for my elf-eyed love? Lovely dainty Spanish needle, Source to me of sweet delight, In your far-off sunny southland Do you dream of me to-night?</i></p>	3) Pussy Cat Pussy Cat	<p><i>Pussy cat, pussy cat, Where have you been? "I've been to London to Look at the Queen." Pussy cat, pussy cat, What did you there? "I frightened a little mouse Under the chair."</i></p>

Table 11: Complete list of Poems mentioned in Figure 4.