

# SS-GEN: A Social Story Generation Framework with Large Language Models

Yi Feng<sup>1</sup>, Mingyang Song<sup>2</sup>, Jiaqi Wang<sup>3\*</sup>, Zhuang Chen<sup>4</sup>, Guanqun Bi<sup>5</sup>  
Minlie Huang<sup>5</sup>, Liping Jing<sup>1\*</sup>, Jian Yu<sup>1</sup>

<sup>1</sup>Beijing Key Laboratory of Traffic Data Mining and Embodied Intelligence, Beijing Jiaotong University

<sup>2</sup>Tencent Hunyuan, Beijing

<sup>3</sup>Jarvis Research Center, Tencent YouTu Lab

<sup>4</sup>School of Computer Science and Engineering, Central South University

<sup>5</sup>CoAI Group, DCST, IAI, BNRIST, Tsinghua University

yifeng@bjtu.edu.cn

## Abstract

Children with Autism Spectrum Disorder (ASD) often misunderstand social situations and struggle to participate in daily routines. Social Stories™ are traditionally crafted by psychology experts under strict constraints to address these challenges but are costly and limited in diversity. As Large Language Models (LLMs) advance, there’s an opportunity to develop more automated, affordable, and accessible methods to generate Social Stories in real-time with broad coverage. However, adapting LLMs to meet the unique and strict constraints of Social Stories is a challenging issue. To this end, we propose **SS-GEN**, a Social Story **GEN**eration framework with LLMs. Firstly, we develop a constraint-driven sophisticated strategy named **STAR**SOW to hierarchically prompt LLMs to generate Social Stories at scale, followed by rigorous human filtering to build a high-quality dataset. Additionally, we introduce **quality assessment criteria** to evaluate the effectiveness of these generated stories. Considering that powerful closed-source large models require very complex instructions and expensive API fees, we finally fine-tune smaller language models with our curated high-quality dataset, achieving comparable results at lower costs and with simpler instruction and deployment. This work marks a significant step in leveraging AI to personalize Social Stories cost-effectively for autistic children at scale, which we hope can encourage future research on special groups.

## Introduction

Social Story™<sup>1</sup>, pioneered by Gray and Garand (1993), helps children with autism<sup>2</sup> define the social context of anticipated or experienced situations, putting them into perspective and then developing statements on how they feel

\*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>As the social-skill intervention tool, Social Stories™ are widely used to specifically help individuals with ASD understand and navigate specific social situations. When capitalized, a Social Story refers to a Social Story™ following the criteria in 10.2 Tutorial (Gray 2010).

<sup>2</sup>Autism Spectrum Disorder (ASD) is a lifelong neurodevelopmental disorder characterized by impairments in social interaction, communication, and restricted or repetitive behaviors, affecting millions of children (approximately 1/100) worldwide with the number continuously growing (Zeidan et al. 2022).

**Title:** What Is A Mistake?

**Introduction:**  
A mistake is an answer, idea, or act that is an error. When someone says or does something that is not right, it's a mistake.

**Main Body:**  
There are many examples of mistakes. It's a mistake to misspell a word. It's a mistake to leave a jacket at home on a very cold day. It's a mistake to forget to turn in finished schoolwork. People make many other kinds of mistakes, too.  
As people grow, they learn from their mistakes. They may not make the same mistake again. However, people are always growing, and having new experiences. For this reason, people are always making new mistakes.  
Sometimes, people know they have made a mistake. Other times, they learn that they made a mistake from others. Once in a while, a mistake is made and no one notices it.  
Most people try to answer questions correctly. They try to have good ideas. They try to do the right thing. As hard as people may try, though, they still make mistakes.

**Conclusion:**  
Mistakes are a part of Life on Planet Earth. That is okay.

**\*\* Strict Constraints \*\***

**Structural Clarity:** title; introduction; main body; conclusion  
**Descriptive Orientation:** describe more than direct  
**Situational Safety:** personal narrative, comfortable vocabulary, positive tone, unambiguous expression

Figure 1: A Social Story excerpted from *The New Social Story Book* written by Carol Gray (Gray 2010), which adhere to strict constraints that we summarized based on professional guidelines: Structural Clarity, Descriptive Orientation, Situational Safety.

and what actions to take in response, thereby positively improving their social understanding and behavioral outcomes. For example, the story in Figure 1 helps autistic children understand and manage the stress and anxiety triggered by mistakes.

These stories are traditionally created by psychologists following the 10.2 tutorial (Gray 2010), which we summarize as three strict constraints: **Structural Clarity** to organize the title, introduction, main body and conclusion for clear social cues; **Descriptive Orientation** to describe more than direct, being patient and humility; **Situational Safety** to emphasize an appropriate perspective, positive tone, unambiguous language and precise vocabulary, ensuring the re-

assuring quality and safeguarding the self-esteem of autistic youth. These handcrafted stories serve as a library for caregivers to customize according to the specific needs of ASD children, but are costly and limited in diversity and scope.

With rapid advancements of Large Language Models (LLMs), it’s promising to automatically generate Social Stories with LLMs. Existing research has explored generating suspenseful (Xie and Riedl 2024), playable (Nasir, James, and Togelius 2024), age-appropriate (Valentini et al. 2023; Bhandari and Brennan 2023) or legal-concept informative (Jiang et al. 2024) stories, primarily for adults or general children. However, adapting LLMs to meet the unique and strict constraints of Social Stories is extremely challenging and rarely a topic of interest, despite the importance of this early intervention tool for improving social skills and enhancing the life quality of such special group (Lynn Kern Koegel and Bradshaw 2014).

In this work, we introduce **SS-GEN**, a Social Story **GEN**eration framework based on their unique constraints. (1) Initially, we develop a constraint-driven strategy named **STARSOW** (Figure 2) to construct a dataset of 5K+ Social Stories by prompting advanced LLMs (e.g., GPT-4o) through a breadth-first, hierarchical synthesis from a seed set of manually-written stories, followed by rigorous human filtering. The STARSOW is similar to writing a book or sowing a star-tree. Concretely, we first generate chapters to cover various story themes (roots), then create story titles (branches) for each theme as much as possible, and last generate story content (star fruits) cautiously that meets specific constraints given the title and the theme. The analysis shows that our dataset is more diverse, flexible, and effective, making it acceptable to human evaluators. (2) Additionally, to verify the effectiveness of generated Social Stories and ensure their adherence to the strict constraints above, we propose the **quality assessment criteria** which is employed in both human and GPT evaluations. (3) Finally, we train smaller language models on our crafted dataset, considering that powerful closed-source LLMs require very complex instructions and expensive API fees. Empirical results show that our dataset significantly improves LLMs performance on SS-GEN, at lower costs and with simpler instructions.

In summary, to the best of our knowledge, we are the first to propose SS-GEN task and framework, which aims to make Social Stories automated, affordable, more accessible and of high quality for their use as a social-skill intervention tool within the autism community. Our framework includes:

- Dataset construction by STARSOW: A breadth-first, hierarchical, rapid-scaling, and constraint-driven synthetic strategy focusing on qualified Social Stories.
- Quality assessment criteria: Personalized criteria for evaluating the adherence of generated Social Stories to the strict constraints.
- Empirical results: Experiments by training and testing smaller and cost-effective language models, which show that our dataset significantly improve the LLMs performance on SS-GEN.

We hope this work supports autistic youth’s social learning needs and inspires future research on special groups.




Figure 2: The overview of STARSOW, which expands Social Stories in a diverse and hierarchical manner within constraints, similar to planning a star tree.

## Social Story Generation

In this section, we introduce the development of Social Story Generation (SS-GEN) framework. We begin by formulating the general notation. Next, we explain the dataset synthesis by our refined constraint-driven strategy named STARSOW (Figure 2), which hierarchically generate varied Social Stories using LLMs (eg., GPT-4o) from a seed set of manually-written stories. Additionally, we propose the quality assessment criteria to verify the adherence of generated Social Stories to the strict constraints.

### Overview

A Social Story includes a clear structure: a title that establishes the intervention goal of the story, an introduction that explains the story theme and setting, a main body that depicts social situations and appropriate responses, and a conclusion that summarizes the core idea presented. The last three parts constitute the Social Story content  $y$ , which is written based on the Social Story title  $x$ . For SS-GEN, we define the Social Story pair as  $(x, y)$  and specify that the language model  $M$  is expected to produce the story content  $y$  given the story title  $x$ .

### Dataset Construction

To expand qualified Social Stories with broad coverage, we develop a constraint-driven strategy STARSOW, which hierarchically expands and synthesizes data as if planting a star tree. The dataset we want to generate contains a set of root nodes  $\{(C_t, E_t)\}$  respectively representing the chapters and their explanations that define Social Story themes. Each root node  $(c_t, e_t)$  further branches out into multiple Social Story pairs  $\{x_{t,j}, y_{t,j}\}_{j=1}^{n_t}$ ,  $n_t > 1$ , where  $x_{t,j}$  and  $y_{t,j}$  mean the title and the content of the  $j^{th}$  Social Story in the chapter node  $(c_t, e_t)$ .

The process for STARSOW strategy consists of four steps: 1) Taking root: explain and expand the chapters, 2) Branching Out: generate diverse Social Story titles from the root layer, 3) Bearing Star Fruits: complete the Social Story content given the title, 4) Gardening Work: filter out chapters, titles, and stories which are invalid and redundant.

Category	Specific Questions	Explanation
<b>Structural Clarity</b> (SC.) 1 ~ 5 scale	<b>Q1:</b> Does the Social Story have a clear structure?	A clear structure refers to that there is a straightforward title establishing the intervention goal, an introduction explaining the story theme and setting, a main body depicting the specific situation and appropriate responses, and a conclusion summarizing the core information presented.
	<b>Q2:</b> Do the introduction and the main body show correlation with each other?	The introduction presents the main point, attracts the audience’s attention, and guides them to the main body. The main body develops the core idea mentioned in the introduction in detail.
	<b>Q3:</b> Do the main body and the conclusion show a correlation with each other?	The conclusion should highlight the core information from the main body, reinforcing the primary messages and insights discussed.
	<b>Q4:</b> Do the conclusion and the introduction show a correlation with each other?	The conclusion and the introduction share the same main point, with the conclusion serving as a callback to the introduction.
<b>Descriptive Orientation</b> (DO.) Yes / No	<b>Q1:</b> Does the Social Story describe more than direct? $\frac{\#DescriptiveSentences}{\#CoachingSentences} \geq 2$	The number of descriptive sentences should exceed twice the number of coaching sentences in a Social Story according to GR-Eight Formula(Gray 2010). A descriptive sentence accurately describes the interaction, an event, or an explanation of the rationale that underlies what and why people think, say, or do, including perspective and affirmative sentences. Coaching sentences gently and safely guide the behavior.
<b>Situational Safety</b> (SS.) Yes / No	<b>Q1:</b> Does the Social Story use the appropriate person perspective?	<b>A):</b> The Social Story should never use the second-person perspective to prevent being overly direct. <b>B):</b> When describing negative behaviors, the Social Story should never employ the first-person perspective to safeguard the dignity and esteem of the audience.
	<b>Q2:</b> Does the Social Story consistently convey a positive and patient tone?	The Social Story should always describe situations or guide behaviors in a positive manner, especially when depicting typical or expected behaviors in specific situations.
	<b>Q3:</b> Does the Social Story express accurately?	The Social Story should use language that is as clear and unambiguous as possible because ASD children typically interpret things literally rather than inferring implicit meanings that require social insight.
	<b>Q4:</b> Does the Social Story use exact vocabulary?	The Social Story should choose the most comfortable and accurate vocabulary for the audience. Firstly, use positive verbs while also being mindful of the varying implications of verbs. Avoid using terms that are likely to evoke strong emotional feelings such as “shouldn’t”, “must”, “supposed to” and so on.

Table 1: The quality assessment criteria table outlines a checklist of questions for Structural Clarity (SC.), Descriptive Orientation (DO.), and Situational Safety (SS.), each accompanied by a detailed explanation.

**Taking Root.** STARSOW begins by generating new chapter root nodes using an expert-written seed set as a reference (Gray 2010). This seed set contains 14 chapters and 179 meticulously selected Social Story pairs in total. To ensure the soundness of new chapter themes, we apply an explain-then-generate approach to better leverage the emergent abilities of the LLM. Firstly, we prompt the LLM to obtain broad explanations of the seed chapters, which form the foundational seed set of root nodes. Secondly, we sample 8 “chapter-explanation” pairs as in-context examples to expand each new root node. Of the 8 examples, 4 are from the seed set, and 4 are from the model-generated brand-new nodes in previous steps to promote diversity.

**Branching Out.** As a certain number of new root nodes is reached, we branch out the corresponding Social Story titles for each node to cover broader intervention goals. This requires the LLM to understand the diverse range of issues that autistic youth may encounter within each specific chapter’s theme. We found that the LLM can come up with diverse story titles to a large extent when prompted with 8 “chapter-explanation-titles” in-context examples which randomly sampled from the seed set.

**Bearing Star Fruits.** As the chapter root nodes grow, and the story titles flourish, we ultimately obtain qualified story fruits by carefully controlling the cultivation conditions. To positively improve social understanding and behavioral outcomes for autistic youth, each Social Story must strictly adhere to our summarized specific constraints for SS-GEN, according to the 10.2 Tutorial (Gray 2010) which provides

detailed guidance on writing and applying Social Stories. These constraints include structural clarity, descriptive orientation, and situational safety. Additionally, we sampled 4 “chapter-explanation-title-story” from the seed set as qualified in-context examples. Finally, we designed a constraint-driven prompt to guide the LLM in completing the Social Story given the root nodes with “chapter-explanation” pair and the corresponding Social Story “title”.

**Gardening Work.** To encourage diversity, a new chapter is added to the Social Story dataset only when its ROUGE-L similarity with any existing chapter is less than 0.7. The same applies to a new Social Story title. When generating the Social Story content based on the title goal, we filter out invalid outputs that do not meet the strict constraints of the Social Stories (e.g., using the narrative “you”, or negative words that usually cause anxious or resistant feelings). After manual screening and refinement, only 74.3% of the data was deemed valid to construct the final dataset.

### Quality Assessment Criteria

To evaluate the validity of Social Stories and their adherence to strict constraints, we design the quality assessment criteria in Table 1 based on the constraints of structural clarity, descriptive orientation, and situational safety. Each constraint includes corresponding checklist questions with detailed explanations provided afterward. We define structural clarity as a soft constraint, scored on a 1 to 5 point scale, with 5 being the best. The remaining constraints are considered hard constraints and assessed with “Yes / No”.

Statistic	
# of chapters	57
# of titles in each chapter	$\geq 70$
# of Social Story Pairs in total	5085
ave.chapter length (in words)	2.46
ave.title length (in words)	5.28
ave.story content length (in words)	281.65

Table 2: Statistics of the constructed Social Story dataset by applying STARSOW to GPT-4.



Figure 3: The top 20 most common root verbs (inner circle) and their top 3 direct noun objects (outer circle) in the generated Social Story titles.

Additionally, we have developed a website for manual assessment of Social Stories and preference ranking based on the above rating scale, which will be utilized in assessing the quality of generated data in subsequent sections.

### Assessment of Social Story

This section provides an overview of the Social Story dataset created using STARSOW. We select the engine GPT-4<sup>3</sup> as the data synthesizer based on the quality analysis below.



#### Statistics

Table 2 describes the basic statistics of the generated dataset for SS-GEN. We generated a total of 57 chapter themes and more than 5K Social Stories, with at least 70 in each chapter after filtering. We then divide the data into training, validation, and testing sets with a ratio of 8:1:1.



#### Diversity

We explore the distinctions between the synthesized dataset and the manually written seed set that prompted the generation. For each generated Social Story pair (title, story con-

<sup>3</sup>OpenAI API: <https://openai.com/api/>. Query parameters are detailed in the supplementary.



(a) Length distribution of the generated Social Story pairs.



(b) Distribution of the ROUGE-L scores between the generated Social Stories and their most similar seed stories.

Figure 4: Overview of generated data statistics.

tent), we calculate its highest ROUGE-L overlap with the corresponding data in the seed set. Figure 4b shows the distribution of these ROUGE-L scores, indicating that a significant number of new titles and stories were generated with minimal overlap with the seeds. Additionally, we also examine diversity from the length perspective, with the distribution shown in Figure 4a.

To further analyze what types of intervention goals are generated and how diverse they are, we identify the verb-noun structure in the generated Social Story titles. Through the Berkeley Neural Parser<sup>4</sup> (Kitaev and Klein 2018; Kitaev, Cao, and Klein 2019), we parse the titles and then extract the verb closest to the root along with its first direct noun object. Figure 3 shows the top 20 most common root verbs and their top 3 direct noun objects, which can represent common topics within the dataset. The figure shows these intervention goals exhibit a wide range of intents and textual formats.

#### Quality

We have shown the quantity and diversity of the generated data before, but the quality remains uncertain, which is crucial for its effectiveness in social interventions for

<sup>4</sup><https://parser.kitaev.io/>

Quality Assessment Criteria & Human Preference		GPT3	GPT4
SC.	Does the Social Story have a clear structure?	82.7%	83.0%
	Does the introduction and the main body show a correlation with each other?	86.2%	86.8%
	Does the main body and the conclusion show a correlation with each other?	82.5%	83.7%
	Does the conclusion and the introduction show a correlation with each other?	84.1%	82.9%
DO.	Does the Social Story describe $\zeta$ direct?	56.0%	59.0%
	Does the Social Story use the appropriate person perspective? (A)	90.0%	94.0%
SS.	Does the Social Story use the appropriate person perspective? (B)	91.0%	93.0%
	Does the Social Story consistently convey a positive and patient tone?	97.0%	98.0%
	Does the Social Story express accurately?	91.0%	90.0%
	Does the Social Story use exact vocabulary?	95.0%	92.0%
HIP.	How do you prefer the Social Story?	38.0%	62.0%

Table 3: Data quality review of Social Stories generated from GPT-3 and GPT-4 given titles in the seed set. SC.: Structural Clarity, DO.: Descriptive Orientation, SS.: Situational Safety, HIP.: Human Preference.

autistic youth. To assess the quality of Social Stories generated by the GPT series, we randomly sample 100 story pairs from the original seed set, using their titles to prompt `gpt-3.5-turbo`, `gpt-4o` respectively and regenerate stories as like **Bearing Star Fruits**. The two sets of regenerated stories are then compared with the story references from the sampled pairs. Specifically, we carefully check the two sets based on checklist questions in Table 1, and then rank the two sets based on the preference.

Table 3 presents the assessment results focusing on structural clarity (SC.), descriptive orientation (DO.), situational safety (SS.) and the human preference (HIP.). The SC. depicts the percentage of the average score relative to the full score. The DO. and SS. show the percentage of meeting the corresponding criteria. The HIP. represents the percentage of human preference between the two sets, summing to 1.

The findings reveal that both engines score over 80% on SC. and SS. and can perform well on SS-GEN, indicating GPT series can construct logically structured Social Stories and ensure the narrative’s safety and appropriateness, despite subtle differences in narrative flow and language nuances. Although the two sets perform relatively lower on DO. with scores between 50% to 60%, we found that those stories are indeed not prescriptive or disrespectful, and they can be easily improved with simple modifications. Notably, `gpt-4o` consistently outperforms `gpt-3.5-turbo` in most categories and is preferred more (62%). Therefore we finally choose `gpt-4o` to construct the whole dataset.

## Experiments

In this section, we conduct extensive experiments<sup>5</sup> to evaluate the performance of smaller, cost-effective language models which are trained and tested using our curated dataset for SS-GEN, considering that powerful closed-source LLMs re-

<sup>5</sup>The code, prompt, data and technical appendix are available at <https://github.com/MIMIFY/SS-GEN>

Title-to-Story for SS-GEN
Develop a concise, clear, straightforward, positive and supportive Social Story titled “{title}” for children and teens with autism, 200-300 words, that promotes their social understanding and boosts their participation in daily activities, fostering independence and confidence. Title: {title}

Figure 5: The prompt is designed to enhance the language model’s ability to complete the Social Story content given the Social Story title.

quire very complex instructions and expensive API fees.

## Experimental Settings

**Model selection:** To conduct thorough experiments within the controlled resource, we selected a series of mainstream language models with 2B, 7B and 8B parameters. Specifically, we choose **Mistral** (2B, 7B) (Jiang et al. 2023), **Gemma** (2B, 7B) (Team et al. 2024), **LLama3-8B**<sup>6</sup> along with each variant of instruction tuning (Instruct).

The experiments include two types. Firstly, we test the performance of popular language models on SS-GEN by directly prompting the original models to generate the Social Story content based on the title (Zero-Shot). Next, we fine tune these models with the proposed dataset and then compare the generated Social Stories (SFT).

We utilize the Parameter-Efficient Fine-Tuning (PEFT) strategy, integrated with Low-Rank Adaption (LoRA) using the LLaMA Factory (Zheng et al. 2024), to test and fine-tune these models on four NVIDIA GeForce RTX 4090 GPUs.

**Dataset and prompt:** We utilize the Social Story dataset constructed through STARSOW for SS-GEN. The dataset is split with 80%, 10%, 10% respectively for training, validation and testing.

Besides, we utilize the same precise “title-to-story” prompt as illustrated in Figure 5 for both training and testing. This simple prompt is designed to enhance the model’s capacity to construct a Social Story from the provided title (the intervention goal of the story).

## Evaluation Methods

To comprehensively and reliably evaluate the generated Social Stories, we leverage both objective measures from traditional metrics and qualitative evaluations based on the proposed quality assessment criteria.

**Traditional Metrics:** We use a variety of evaluation metrics, including word-based BLEU-4, ROUGE-1, ROUGE-2, ROUGE-L scores, and embedding-based BertScore, to ensure objective and quantifiable assessments. Higher values in these metrics indicate better model performance and greater N-gram similarity between the generated text and the reference text, showing that the generated data more accurately reflects the patterns found in the reference.

**Quality Assessments:** We incorporate the proposed quality assessment criteria for Social Stories into both GPT-4 and human evaluations.

<sup>6</sup><https://github.com/meta-llama/llama3>




Figure 6: GPT-4 evaluation results of various models on five dimensions on Zero-Shot (blue) and SFT (orange) scenarios for SS-GEN: Coherence ( $\mathcal{CH}$ ), Descriptiveness ( $\mathcal{DC}$ ), Empathy ( $\mathcal{EM}$ ), Grammaticality ( $\mathcal{GA}$ ), and Relevance ( $\mathcal{RE}$ ).

Model	Zero-Shot (%)					SFT (%)				
	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	BertScore	BLEU-4	ROUGE-1	ROUGE-2	ROUGE-L	BertScore
Gemma (2B)	8.86	21.81	4.63	7.22	79.0	<b>51.18</b> <sup>†</sup> <sub>+42.32%</sub>	<b>47.58</b> <sup>†</sup> <sub>+25.77%</sub>	<b>20.13</b> <sup>†</sup> <sub>+15.50%</sub>	<b>27.29</b> <sup>†</sup> <sub>+20.07%</sub>	88.8 <sup>†</sup> <sub>+9.8%</sub>
Gemma-Instruct (2B)	31.68	30.44	6.49	14.17	82.7	50.35 <sup>†</sup> <sub>+18.67%</sub>	46.53 <sup>†</sup> <sub>+16.09%</sub>	19.36 <sup>†</sup> <sub>+12.87%</sub>	26.93 <sup>†</sup> <sub>+12.76%</sub>	88.8 <sup>†</sup> <sub>+6.1%</sub>
Gemma-1.1-Instruct (2B)	<b>42.43</b>	<b>34.98</b>	<b>8.82</b>	<b>17.22</b>	<b>84.9</b>	51.02 <sup>†</sup> <sub>+8.590%</sub>	46.70 <sup>†</sup> <sub>+11.72%</sub>	19.53 <sup>†</sup> <sub>+10.71%</sub>	27.15 <sup>†</sup> <sub>+9.930%</sub>	<b>88.9</b> <sup>†</sup> <sub>+4.0%</sub>
Gemma (7B)	19.24	32.12	9.57	12.48	81.7	54.74 <sup>†</sup> <sub>+35.50%</sub>	<b>52.12</b> <sup>†</sup> <sub>+20.00%</sub>	24.74 <sup>†</sup> <sub>+15.17%</sub>	31.04 <sup>†</sup> <sub>+18.56%</sub>	<b>89.9</b> <sup>†</sup> <sub>+8.2%</sub>
Gemma-Instruct (7B)	39.80	32.56	8.09	16.33	84.1	53.82 <sup>†</sup> <sub>+14.02%</sub>	50.80 <sup>†</sup> <sub>+18.24%</sub>	23.34 <sup>†</sup> <sub>+15.25%</sub>	30.19 <sup>†</sup> <sub>+13.86%</sub>	89.7 <sup>†</sup> <sub>+5.6%</sub>
Gemma-1.1-Instruct (7B)	<b>45.64</b>	37.65	10.00	17.87	84.8	54.04 <sup>†</sup> <sub>+8.400%</sub>	51.04 <sup>†</sup> <sub>+13.39%</sub>	23.71 <sup>†</sup> <sub>+13.71%</sub>	30.40 <sup>†</sup> <sub>+12.53%</sub>	89.8 <sup>†</sup> <sub>+5.0%</sub>
Mistral-v0.1 (7B)	19.76	25.78	5.95	12.47	81.6	<b>54.84</b> <sup>†</sup> <sub>+35.08%</sub>	51.70 <sup>†</sup> <sub>+25.92%</sub>	24.07 <sup>†</sup> <sub>+18.12%</sub>	30.58 <sup>†</sup> <sub>+18.11%</sub>	89.8 <sup>†</sup> <sub>+8.6%</sub>
Mistral-v0.1-Instruct (7B)	40.14	38.23	11.85	<b>20.79</b>	<b>86.5</b>	54.17 <sup>†</sup> <sub>+14.03%</sub>	51.17 <sup>†</sup> <sub>+12.94%</sub>	23.79 <sup>†</sup> <sub>+11.94%</sub>	30.54 <sup>†</sup> <sub>+9.750%</sub>	89.7 <sup>†</sup> <sub>+3.2%</sub>
Mistral-v0.2-Instruct (7B)	37.44	37.70	11.51	19.74	86.3	54.30 <sup>†</sup> <sub>+16.86%</sub>	52.05 <sup>†</sup> <sub>+14.35%</sub>	<b>24.75</b> <sup>†</sup> <sub>+13.24%</sub>	<b>31.18</b> <sup>†</sup> <sub>+11.44%</sub>	<b>89.9</b> <sup>†</sup> <sub>+3.6%</sub>
LLama3 (8B)	9.49	23.46	5.82	8.15	79.3	16.41 <sup>†</sup> <sub>+6.920%</sub>	47.66 <sup>†</sup> <sub>+24.20%</sub>	22.20 <sup>†</sup> <sub>+16.38%</sub>	9.69 <sup>†</sup> <sub>+1.540%</sub>	81.2 <sup>†</sup> <sub>+1.9%</sub>
LLama3-Instruct (8B)	39.04	<b>41.19</b>	<b>12.68</b>	19.37	85.5	54.56 <sup>†</sup> <sub>+15.52%</sub>	51.26 <sup>†</sup> <sub>+10.07%</sub>	23.64 <sup>†</sup> <sub>+10.96%</sub>	30.30 <sup>†</sup> <sub>+10.93%</sub>	89.7 <sup>†</sup> <sub>+4.2%</sub>

Table 4: Average Performance of three inferences on traditional metrics in Zero-Shot and SFT scenarios for SS-GEN. The subscripts on the bottom right indicate the performance change of the same model after fine-tuning with the dataset. The symbol (†) indicates a significant difference between pre- and post-finetuning results ( $p < 0.01$ ) in a paired t-test at  $\alpha = 0.05$ .

For GPT-4 evaluation, we apply the Coherence ( $\mathcal{CH}$ ), Descriptiveness ( $\mathcal{DC}$ ), Empathy ( $\mathcal{EM}$ ) scores to separately evaluate whether the Social Story meets the checklist requirements for structural clarity, descriptive orientation, situational safety as outlined in the established criteria for quality assessment. Additionally, for the text itself, we use the Grammaticality ( $\mathcal{GA}$ ) score to determine grammatical correctness, fluency, and adherence to language standards. We also use the Relevance ( $\mathcal{RE}$ ) score to assess the consistency and alignment between the generated story content and the given title prompt. All scores range from 1 to 5, with 5 representing the highest level of quality, clarity, and adherence to the evaluation standards.

For human evaluation, we invite three psychology students who received comprehensive training on the quality assessment criteria guidelines before their participation in the evaluation. Additionally, we develop a user-friendly website shown in the supplementary to standardize expert evaluations, improve efficiency, and enhance the horizontal comparison of text generated by different models. Based on the criteria, experts use this website to score Q1, Q2, Q3, Q4 of the structural clarity ( $\mathcal{SC}$ ) on a scale of 1 to 5, and to provide “Yes/No” for Q1 of the descriptive orientation ( $\mathcal{DO}$ ) and Q1(A), Q1(B), Q2, Q3, Q4 of the situational safety ( $\mathcal{SS}$ ). Finally, we report the average score of all  $\mathcal{SC}$  sub-items, as well as the proportion of qualified  $\mathcal{SS}$  and  $\mathcal{DO}$ , where the “qualified” means all sub-items are rated as “Yes”.

## Results and Analysis

We evaluate mainstream models of various scales for generating Social Stories in Zero-Shot and SFT scenarios. Table 4 presents the average results of three inferences on the test set using traditional metrics. Figure 6 shows GPT-4 evaluations of 200 randomly selected stories from the test set, focusing on Coherence ( $\mathcal{CH}$ ), Descriptiveness ( $\mathcal{DC}$ ), Empathy ( $\mathcal{EM}$ ), Grammaticality ( $\mathcal{GA}$ ) and Relevance ( $\mathcal{RE}$ ). In Table 5, we select the top 2B and 7B models on both traditional and GPT-4 evaluations for further expert assessment, along with 50 randomly selected stories from the test set for the expert assessment. We present the analysis of the case study in the technical appendix.

**Performance improvement through SFT:** We significantly improve LLMs on SS-GEN through fine-tuning with our proposed Social Story dataset, which can be shown by the positive changes in the subscripts on the bottom right of SFT scenarios in Table 4. For example, SFT scores in Table 4 are markedly higher than Zero-Shot scores every row. The fine-tuned Gemma (2B) improves its performance by 42.32%<sup>†</sup>(BLEU-4), 25.77%<sup>†</sup>(ROUGE-1), 15.50%<sup>†</sup>(ROUGE-2), 20.07%<sup>†</sup>(ROUGE-L) and 9.8%<sup>†</sup>(BertScore). Multiple inferences (3 times) were conducted for each model before and after fine-tuning, followed by a t-test ( $\alpha = 0.05$ ), with † in Table 4 indicating a significant difference between pre- and post-fine-tuning results ( $p < 0.01$ ).

Scenarios	Model	Quality Assessment		
		SC.	DO.	SS.
SFT	Gemma (7B)	4.90	86.7%	83.3%
	Gemma (2B)	4.83	83.3%	80.0%
Zero-Shot	Gemma-1.1-Instruct (7B)	3.89	93.3%	10.0%
	Gemma-1.1-Instruct (2B)	3.28	90.0%	6.00%

Table 5: Results of the human evaluation on quality assessment and preference ranking for the best Gemma models (2B, 7B) on both Zero-Shot and SFT scenarios.

Scenario	Model	SC. / CH.	DO. / DC.	SS. / EM.
SFT	Gemma (7B)	0.52	0.43	0.47
	Gemma (2B)	0.49	0.46	0.34
Zero-Shot	Gemma-1.1-Instruct (7B)	0.54	0.44	0.41
	Gemma-1.1-Instruct (2B)	0.51	0.48	0.36

Table 6: Pearson correlation coefficients between GPT and human Ratings, showing a positive correlation trend.

As for GPT-4 quality evaluation, Figure 6 shows that the yellow line representing SFT models is consistently higher than the blue line representing Zero-Shot models. Moreover, we select the overall best-performing 2B/7B models on both traditional evaluations and GPT-4 evaluations for further expert assessment, which are fine-tuned Gemma (2B/7B), zero-shot Gemma-1.1-Instruct (2B/7B) for human evaluation, shown in Table 5. Results from human evaluation indicates that the stories generated by Zero-Shot models have a much lower situational safety (SS.) compliance rate compared to those generated by SFT models. Besides, Table 6 shows the correlation between ratings provided by GPT evaluators (CH., DC., EM.) and human evaluators (SC., DO., SS.) on the same set of data as Table 5, indicating a positive correlation trend between the GPT evaluation and human evaluation.

**Instruct Models on Zero-Shot:** In Zero-shot scenarios, Instruct models consistently demonstrate superior performance compared to their original pre-trained base counterparts. This is evident in both traditional metrics and GPT-4 evaluation. For instance, Figure 6 shows that all peaks of the blue lines correspond to Instruct models, while all troughs indicate base models. Besides, Table 4 demonstrates that under Zero-Shot settings, Gemma-1.1-Instruct (2B) improves performance by 33.57% $\uparrow$ (BLEU-4), 13.17% $\uparrow$ (ROUGE-1), 4.19% $\uparrow$ (ROUGE-2), and 10% $\uparrow$ (ROUGE-L) compared to Gemma (2B).

**Base Models on SFT:** In SFT scenarios, base models demonstrate greater potential for adapting to domain-specific data. For example, fine-tuned base models outperform fine-tuned Instruct models, as seen in both GPT evaluations and traditional metrics. This trend holds true for the models of Gemma and Mistral series, but not for the LLama3 series. It is speculated that the LLama3 base model, pre-trained with 15T data, may need a larger and more targeted fine-tuning dataset to effectively leverage its pre-learned complex patterns and features, thus reaching optimal performance on specific tasks like SS-GEN.

## Related Work

**LLMs for Story Generation.** The advent of LLMs has significantly advanced automatic story generation, enabling tailored narratives with minimal manual intervention to mimic human-written stories. Researchers have utilized storytelling to enhance various fields, including creating playable 2D story-games (Nasir, James, and Togelius 2024), crafting suspenseful stories (Xie and Riedl 2024), generating longer and more creative narratives (Yang et al. 2023), and conveying legal concepts (Jiang et al. 2024). However, these stories focus on adult audiences, emphasizing narrative flow, entertainment, and emotional engagement. This free-form approach contrasts sharply with the requirements of Social Stories, which need clear structure and accurate description. (Valentini et al. 2023; Bhandari and Brennan 2023) explored generating age-appropriate stories for typical children, assessing the trustworthiness and suitability of LLM-generated content. Nonetheless, these stories do not meet the stringent demands of situational safety in Social Stories, including constraints on narrative perspective, positive tone, vocabulary selection, and precise expression. In this work, we explore the constraint-driven strategy of generating and evaluating Social Stories, focusing on the audience of children with autism.

**LLMs for Data Augmentation.** Built upon massive corpora and qualified alignment, the increasingly powerful LLMs (Achiam et al. 2023) are promising at producing texts of eligible quality, resulting in a surge of research interests in exploring using LLMs to create synthetic data for tasks or to augment existing datasets to boost task performance (Patel, Raffel, and Callison-Burch 2024). Such applications include generating tabular data (Borisov et al. 2022), relation triplets (Chia et al. 2022), sentence pairs (Schick and Schütze 2021), open-domain dialogue (Zheng et al. 2023), instruction data (Peng et al. 2023; Shao et al. 2023; Wang et al. 2023; Sun et al. 2024), etc.. Enhanced task performance validates the effectiveness of LLM-generated data. In this work, we design a constraint-driven strategy named STARSOW to use LLMs for constructing qualified Social Story dataset, aiming to make Social Stories more accessible, affordable, and of high quality for their use as a social-skill intervention tool within the autism community.

## Conclusion

We propose **SS-GEN**, an innovative Social Story Generation framework that includes: (1) Qualified dataset construction for SS-GEN through a breadth-frist, hierarchical and constraint-driven strategy **STARSOW**. (2) The personalized **quality assessment criteria** for evaluating the adherence of generated Social Stories to their unique and strict constraints. (3) Extensive experiments by training and testing smaller, cost-effective language models, which show that our dataset significantly improve the LLMs performance on SS-GEN. We hope SS-GEN makes Social Stories more accessible, affordable and of high quality for their use as a social-skill intervention tool within the autism community.

## Ethical Statement

The generation of Social Stories entails a range of ethical considerations and limitations that require careful attention to ensure their responsible application, particularly in sensitive contexts involving neurodiverse children. The primary ethical issues are outlined as follows. (1) Bias and Stereotypes: training data may carry social biases although we have rigorous human review to ensure quality. (2) Caregiver supervision: AI-drafted stories need caregiver’s refinement to suit each child’s unique experiences, ensuring ethical and emotional suitability. (3) Ethical risks: limited transparency in generation may affect long-term predictability. Expert feedback and heuristic controls are advised to mitigate emotional insensitivity risks.

## Acknowledgments

This work was partly supported by the National Key Research and Development Program of China under Grant 2024YFE0202900; the National Natural Science Foundation of China under Grant (62436001, 62176020); the Joint Foundation of the Ministry of Education for Innovation team (8091B042235); the Fundamental Research Funds for the Central Universities (2019JBZ110); and the State Key Laboratory of Rail Traffic Control and Safety (Contract No. RCS2023K006), Beijing Jiaotong University. We sincerely thank our collaborators for their invaluable contributions, including insightful feedback and support in refining the paper and addressing key challenges. Their dedication and expertise greatly enhanced the quality and impact of this work. We thank the anonymous reviewers for carefully reading our paper and their insightful comments and suggestions.

## References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bhandari, P.; and Brennan, H. 2023. Trustworthiness of Children Stories Generated by Large Language Models. In *Proceedings of the 16th International Natural Language Generation Conference*, 352–361.
- Borisov, V.; Sessler, K.; Leemann, T.; Pawelczyk, M.; and Kasneci, G. 2022. Language Models are Realistic Tabular Data Generators. In *The Eleventh International Conference on Learning Representations*.
- Chia, Y. K.; Bing, L.; Poria, S.; and Si, L. 2022. Relation-Prompt: Leveraging Prompts to Generate Synthetic Data for Zero-Shot Relation Triplet Extraction. In *Findings of the Association for Computational Linguistics: ACL 2022*, 45–57.
- Gray, C. 2010. *The new social story book*. Future Horizons.
- Gray, C. A.; and Garand, J. D. 1993. Social stories: Improving responses of students with autism with accurate social information. *Focus on autistic behavior*, 8(1): 1–10.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; Casas, D. d. l.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825*.
- Jiang, H.; Zhang, X.; Mahari, R.; Kessler, D.; Ma, E.; August, T.; Li, I.; Pentland, A.; Kim, Y.; Kabbara, J.; et al. 2024. Leveraging Large Language Models for Learning Complex Legal Concepts through Storytelling. *arXiv preprint arXiv:2402.17019*.
- Kitaev, N.; Cao, S.; and Klein, D. 2019. Multilingual Constituency Parsing with Self-Attention and Pre-Training. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3499–3505.
- Kitaev, N.; and Klein, D. 2018. Constituency Parsing with a Self-Attentive Encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2676–2686.
- Lynn Kern Koegel, K. A., Robert L. Koegel; and Bradshaw, J. 2014. The importance of early identification and intervention for children with or at risk for autism spectrum disorders. *International journal of speech-language pathology*, 16(1): 50–56. PMID: 24328352.
- Nasir, M. U.; James, S.; and Togelius, J. 2024. Word2World: Generating Stories and Worlds through Large Language Models. *arXiv preprint arXiv:2405.06686*.
- Patel, A.; Raffel, C.; and Callison-Burch, C. 2024. DataDreamer: A Tool for Synthetic Data Generation and Reproducible LLM Workflows. *arXiv preprint arXiv:2402.10379*.
- Peng, B.; Li, C.; He, P.; Galley, M.; and Gao, J. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*.
- Schick, T.; and Schütze, H. 2021. Generating Datasets with Pretrained Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 6943–6951.
- Shao, Z.; Gong, Y.; Shen, Y.; Huang, M.; Duan, N.; and Chen, W. 2023. Synthetic prompting: Generating chain-of-thought demonstrations for large language models. In *International Conference on Machine Learning*, 30706–30775. PMLR.
- Sun, Z.; Shen, Y.; Zhou, Q.; Zhang, H.; Chen, Z.; Cox, D.; Yang, Y.; and Gan, C. 2024. Principle-driven self-alignment of language models from scratch with minimal human supervision. *Advances in Neural Information Processing Systems*, 36.
- Team, G.; Mesnard, T.; Hardin, C.; Dadashi, R.; Bhupatiraju, S.; Pathak, S.; Sifre, L.; Rivière, M.; Kale, M. S.; Love, J.; et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Valentini, M.; Weber, J.; Salcido, J.; Wright, T.; Colunga, E.; and von der Wense, K. 2023. On the Automatic Generation and Simplification of Children’s Stories. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 3588–3598.
- Wang, Y.; Kordi, Y.; Mishra, S.; Liu, A.; Smith, N. A.; Khoshabi, D.; and Hajishirzi, H. 2023. Self-Instruct: Aligning Language Models with Self-Generated Instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 13484–13508.



Xie, K.; and Riedl, M. 2024. Creating Suspenseful Stories: Iterative Planning with Large Language Models. *arXiv preprint arXiv:2402.17119*.

Yang, K.; Klein, D.; Peng, N.; and Tian, Y. 2023. DOC: Improving Long Story Coherence With Detailed Outline Control. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 3378–3465.

Zeidan, J.; Fombonne, E.; Scolah, J.; Ibrahim, A.; Durkin, M. S.; Saxena, S.; Yusuf, A.; Shih, A.; and Elsabbagh, M. 2022. Global prevalence of autism: A systematic review update. *Autism research*, 15(5): 778–790.

Zheng, C.; Sabour, S.; Wen, J.; Zhang, Z.; and Huang, M. 2023. AugESC: Dialogue Augmentation with Large Language Models for Emotional Support Conversation. In *Findings of the Association for Computational Linguistics: ACL 2023*, 1552–1568.

Zheng, Y.; Zhang, R.; Zhang, J.; Ye, Y.; and Luo, Z. 2024. Llamafactory: Unified efficient fine-tuning of 100+ language models. *arXiv preprint arXiv:2403.13372*.