

From We to Me: Theory Informed Narrative Shift with Abductive Reasoning

Jaikrishna Manojkumar Patil^{1*}, Divyagna Bavikadi^{1*}, Kaustuv Mukherji¹,
Ashby Steward-Nolan¹, Peggy-Jean Allin², Tumininu Awonuga², Joshua Garland²
and Paulo Shakarian¹

¹Syracuse University

²Arizona State University

{jpatil01, dbavikad, kmukherj, astewa29}@syr.edu, {pallin, osawonug, Joshua.Garland}@asu.edu,
pashakar@syr.edu

Abstract

Effective communication often relies on aligning a message with an audience’s narrative and worldview. Narrative shift involves transforming text to reflect a different narrative framework while preserving its original core message—a task we demonstrate is significantly challenging for current Large Language Models (LLMs). To address this, we propose a neurosymbolic approach grounded in social science theory and abductive reasoning. Our method automatically extracts rules to abduce the specific story elements needed to guide an LLM through a consistent and targeted narrative transformation. Across multiple LLMs, abduction-guided transformed stories shifted the narrative while maintaining the fidelity with the original story. For example, with GPT-4o we outperform the zero-shot LLM baseline by 55.88% for collectivistic to individualistic narrative shift while maintaining superior semantic similarity with the original stories (40.4% improvement in KL divergence). For individualistic to collectivistic transformation, we achieve comparable improvements. We show similar performance across both directions for Llama-4, and Grok-4 and competitive performance for Deepseek-R1.

1 Introduction

Narratives function as cognitive constructs that shape social reality and influence how individuals interpret events [Javadi *et al.*, 2025; Li *et al.*, 2024; Bruner, 1991]. Cultural orientations—particularly the distinction between individualism and collectivism systematically affect how people attribute causality, assign moral responsibility, and evaluate success or failure [Choi, 2025; Markus and Kitayama, 2014; Triandis, 1995]. While LLMs have been applied to narrative-adjacent task such as story generation [Roeein *et al.*, 2025; Tian *et al.*, 2024; Pei *et al.*, 2024], narrative discourse analysis [Jenner *et al.*, 2025; Piper and Bagga, 2024], text style transfer [Tao *et al.*, 2025; Mukherjee *et al.*, 2024], they

*equal contribution

Original (Collectivistic):

“But it will not rise for the feet of a few. *Only when all hands dig—young, old, weak, strong—will the spring awaken.*”

Baseline (Zero-shot LLM): × Failed to transform

“But it will not rise for the feet of a few. Only when all hands dig—young, old, weak, strong—will the spring awaken.”

Our Method: ✓ Individualistic Shift

“But it will not rise for the feet of a few. *Only when one determined soul digs will the spring awaken.*”

Figure 1: Abduction-guided LLM transformation correctly identifies narrative element makes the target narrative shift.

have not been shown to transform a narrative’s cultural orientation while preserving its core semantics. This problem is non-trivial—as illustrated in Figure 1, a zero-shot LLM prompt fails to identify and modify the key narrative marker (*‘all hands dig—young, old, weak, strong’*) that signals collectivism. We propose a neurosymbolic abduction-guided approach grounded in social-science theory that identifies the specific story elements that require transformation. As seen in Figure 1, our abduction-guided transformation correctly identifies the key collectivistic phrase and modifies it into *‘only when one determined soul digs.’* Across multiple LLMs, our approach outperforms the zero-shot baseline by up to 55.88% in shifting narrative while maintaining up to 40.4% better semantic fidelity with the original story. Note that we learn rules from social-science theory to ensure narrative transformations are aligned with the target corpus, rather than simply maximizing narrative orientation ratings.

The rest of the paper is organized as follows: Section 2 provides background on individualism-collectivism as cultural dimensions and introduces our social-science grounded diagnostic survey. Section 3 reviews related work in text style transfer, theory-grounded LLM systems, and abductive reasoning. Section 4 provides a logical reasoning framework for narrative shift and formalizes it as an abduction problem. Section 5 presents our two-phase methodology—rule learning from training corpus and iterative abductive transformation. Section 6 demonstrates the effectiveness of our approach against a zero-shot baseline across different LLMs and narrative shift directions. Finally, we conclude with our findings and future work in Section 7.

2 Background

Research in social sciences has shown that effective communication relies on aligning a message with an audience’s narrative and worldview [Green and Appel, 2024; Green and Brock, 2000]. This principle extends across diverse contexts: diplomacy, journalism, and efforts to mitigate political polarization all depend on understanding how different communities interpret the same events through distinct cultural lenses. In military and intelligence contexts, information operations similarly tailor messages to specific populations. For example, state and non-state actors involved in contemporary conflicts frame the same events differently for domestic, allied, and adversary audiences based on historical memory and collective identity [Paul and Matthews, 2016]. Individualism and collectivism are foundational cultural dimensions that shape how the self is defined, how goals are prioritized, and how appropriate action is determined [Choi, 2025; Triandis, 1995; Hofstede, 1984]. Individualism is associated with an independent self-construal emphasizing autonomy, personal traits, and individual achievement, whereas collectivism is associated with an interdependent self-construal emphasizing social roles, relationships, and group belonging [Markus and Kitayama, 2014; Singelis, 1994].

We define narrative shift to be the transformation of a text’s narrative orientation while preserving the original narrative’s core message and events. By refining LLMs to systematically alter narrative orientation along this cultural dimension while preserving fidelity—a capability currently absent in existing literature—we can facilitate more impactful communication across cultural divides. This allows for the intentional creation of messages that either align with or strategically contrast the narratives of specific communities. To operationalize these cultural dimensions for computational analysis, our team that includes social science experts developed a structured 20-item diagnostic survey grounded in cross-cultural psychology and narrative theory. The survey evaluates narratives across 20 distinct feature dimensions, with each feature paired bidirectionally—one question assesses individualistic framing and its complement assesses collectivistic framing. This results in 40 diagnostic questions (20 individualistic, 20 collectivistic). Individualistic questions emphasize personal goal prioritization, action guided by individual preferences, individual responsibility over outcomes, and recognition based on personal achievement [McAdams, 1993; Bercovitch, 1975; Moon, 2020]. Collectivistic questions emphasize group goal prioritization, action guided by social norms and group expectations, collective responsibility, and recognition based on shared contribution [Singelis and Brown, 1995]. Each question targets a specific narrative element—such as conflict, actors, actions, goals, relationships, resolution, or moral evaluation—such that cumulative responses determine the degree to which a narrative aligns with individualistic or collectivistic orientations. We describe the diagnostic in more detail in Section 6, and include the full diagnostics in Appendix B.

3 Related Work

Recent work on zero-shot text-style transfer with LLMs [Reif *et al.*, 2022] show that diverse rewriting examples in prompts can eliminate the need for complex task-specific training data. However they lack systematic guidance on what specific style elements to modify and thus suffer from hallucinations and incoherent generation. Another recent approach [Suzgun *et al.*, 2022] attempt to improve style-transfer through ranking and selection of candidate transformations using contrastive prompts. However it seems to show directional biases—positive to negative transformation was easier than negative to positive. Our method replaces post-hoc ranking by abductive reasoning to identify which chunks require modification and allows for bidirectional coherent narrative shift without expensive candidate generation. An alternative line of research focuses on fine-tuning the LLMs on style-transfer data. Recent study [Mukherjee *et al.*, 2024] comparing zero-shot, few-shot, and finetuning approaches shows that task-specific finetuning substantially improves multilingual style transfer. However, finetuning requires large amount of labeled data. Our approach does not require this task-specific finetuning and works well even with small language models.

Recent work [Singh *et al.*, 2025; Mukherjee *et al.*, 2025] show that combining multiple evaluation metrics with LLM-based judgment provide more reliable assessments of narrative quality instead of metrics or LLM judge alone. Our work extends this principle to the domain of narrative transformation. We use KL divergence to assess semantic fidelity and LLM-diagnosis to measure narrative shift. Theory-grounded LLM systems or generative agents [Yue *et al.*, 2025; Chun *et al.*, 2025; Park *et al.*, 2023] across domains have showed that constraining LLM outputs with domain specific social science principles improves coherency as well as real-world alignment. We extend this principle to narrative transformation by grounding our approach in social science theory—specifically, individualism and collectivism. Early work in abductive inference has been applied for movement data. It was used to generate faux trajectories that meet spatio-temporal constraints [Bavikadi *et al.*, 2025a] and identifying future regions of an agent [Bavikadi *et al.*, 2025b]. On the contrary, we look at narrative shift on natural language text and use optimization to guide an LLM-based transformation.

4 Technical Preliminaries

We consider a logical language with a sets of constants \mathcal{C} , variables \mathcal{V} , and predicates \mathcal{P} . Let $\mathcal{C}_s \cup \mathcal{C}_c \cup \mathcal{C}_n \cup \mathcal{C}_f \subseteq \mathcal{C}$, where the domain \mathcal{C}_s is a set of constants denoting narrative stories, domain \mathcal{C}_c a set of constants representing unique chunks (subsets of a story) from all stories (additionally, $\mathcal{C}_c^s \subseteq \mathcal{C}_c$ denotes the set of unique chunks for the story $s \in \mathcal{C}_s$), domain $\mathcal{C}_n = \{ind, col\}$ is a set of constants indicating the individualistic and collectivistic narrative and domain $\mathcal{C}_f = \{f_1, \dots, f_{20}, \dots, f_{40}\}$ is a set of constants indicating various narrative features. Correspondingly, the sets of variables are, $\mathcal{V}_s^{var} \cup \mathcal{V}_c^{var} \cup \mathcal{V}_n^{var} \cup \mathcal{V}_f^{var} \subseteq \mathcal{V}$.

In addition to first-order logic syntax and semantics, we allow for annotations on atoms. Annotations are elements of a lower semi-lattice structure \mathcal{L} , where each element is a

subset interval of a unit interval $\mu \subseteq [0, 1]$, this generalizes fuzzy logic [Kifer and Subrahmanian, 1992]. An atom can be represented as $a = p(t_1, \dots, t_n)$, where $p \in \mathcal{P}$ is an n -ary predicate and $t_1, \dots, t_n \in \mathcal{C} \cup \mathcal{V}$. A fact or a ground atom is an atom without any variables. An annotated ground atom a_μ means a is associated with an annotation μ . Consider \mathcal{G} to a set of all ground atoms. We use an annotated logic framework where the desirable annotation is positioned at top of the lattice and $[0, 1]$ is the bottom element (\perp) of the lattice. If $\alpha, \alpha_1, \dots, \alpha_m$ are atoms, then $\alpha \leftarrow \alpha_1 \wedge \dots \wedge \alpha_m$ is called a rule. A program Π is defined as set of facts and rules. We use the notation Γ_Π^* to refer to the deductive closure with the logic program Π and initial set of facts with assigned annotations to \perp (meaning that we start deduction under the assumption of uncertainty). We perform deductive inference with logic program Π [Shakarian and Simari, 2022] and is efficiently implemented using PyReason [Aditya *et al.*, 2023].

Example 4.1 (Language). A narrative story $s \in \mathcal{C}_s$ can be expressed as a series of chunks $\langle c_1, c_2, \dots \rangle$ that have certain narrative characteristics, where $\{c_1, c_2, \dots\} \in \mathcal{C}_c^s$. We represent the chunks as $\{\text{contains}(s, c_1)_1, \text{contains}(s, c_2)_1 \dots\}$ and Ch_s is a set of chunks from \mathcal{C}_c in story s . Suppose s is a collectivistic narrative story. We assume binary predicates $\{s_feat, c_feat, corpus_sim, contains, associated\} \in \mathcal{P}$. For our use-case, we assign the annotation with a scalar- the lower bound of the annotation intervals, as they always have the upper bound as 1 or is the same as the lower bound. Each natural language text or chunk can be related with multiple narrative features denoted by constants \mathcal{C}_f . We use a trainable generative function $g_\theta(x) = \mu; x \in \mathcal{C}_s \cup \mathcal{C}_c$ to obtain a rating $\mu \in \{0.2, 0.4, 0.6, 0.8, 1\}$ on features for a given text. We notate this in our language using ground atoms $c_feat(s, ind)_{0.6}$ for $s \in \mathcal{C}_s, ind \in \mathcal{C}_n$ showing that the story is of a individualistic narrative with a confidence of 0.6, and $c_feat(c, f_{20})_{0.2}$ for $c \in \mathcal{C}_c, f_{20} \in \mathcal{C}_f$ showing that the story has narrative characteristic f , with a confidence of 0.2. Also, we use $associated(ind, f_{20})_1$ to show that the feature f_{20} is associated with the individualistic narrative. We have $\{f_1, \dots, f_{20}\}$ associated with the individualistic narrative and $\{f_{20}, \dots, f_{40}\}$ associated with the collectivistic narrative. Finally, to represent the confidence of the story having an individualistic narrative, we use ground atom, $corpus_sim(x, f_{20})_{conf}$ for $x \in \mathcal{C}_c \cup \mathcal{C}_s$, and $conf \in [0, 1]$.

Abduction. Given a story, our goal is to transform it's narrative by identifying a set of chunks to modify along with the corresponding characteristics. We formalize our problem as an abduction problem: Given a set of observations \mathcal{O} that are annotated facts depicting the various chunks and corresponding narrative characteristics of the story, hypothesis \mathcal{H} consisting of annotated facts that have the same chunks in \mathcal{O} with different narrative characteristics, a logic program Π governing the narrative structure of the story. This gives us an instance of an abduction problem $\langle \mathcal{O}, \mathcal{H}, \Pi \rangle$. We abduce an explanation $\mathcal{E} \subseteq \mathcal{H}$ to the abduction problem $\langle \mathcal{O}, \mathcal{H}, \Pi \rangle$, such that $\Pi \cup \mathcal{O} \cup \mathcal{E}$ is consistent and \mathcal{E} maximizes the parsimony function σ . We now formally define Observations \mathcal{O} , Hypothesis \mathcal{H} , and Logic Program Π .

Definition 4.1 (Observations (\mathcal{O})). Given $s \in \mathcal{C}_s$, a narrative diagnosis function, $diagnosis : \mathcal{C}_s \rightarrow \mathcal{G}$, generates set of ground atoms, then \mathcal{O} , is defined as:

$$\mathcal{O} = \forall_{c \in \mathcal{C}_c^s} \{\text{contains}(s, c)\} \cup \text{diagnosis}(s)$$

$$\text{diagnosis}(s) = \forall_{c \in \mathcal{C}_c^s} \{c_feat(c, f)_\mu$$

$$| f \in \mathcal{C}_f, \mu \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}\}$$

We further elaborate on the narrative *diagnosis* function in Section 6.

Definition 4.2 (Hypothesis (\mathcal{H})). For a story $s \in \mathcal{C}_s$, \mathcal{H} is the set of all possible atoms (that potentially contribute to the narrative transformation) indicating different chunks with various narrative characteristics. We define \mathcal{H} as follows:

$$\mathcal{H} = \forall_{f \in \mathcal{C}_f} \{c_feat(c, f)_{\mu'} |$$

$$\mu' \geq \max(r, 0), c_feat(c, f)_\mu \in \mathcal{O}\}$$

Definition 4.3 (Logic Program (Π)). For $S \in \mathcal{V}_c^{var}, C \in \mathcal{V}_c^{var}, F \in \mathcal{V}_n^{var}, F' \in \mathcal{V}_f^{var}, \mu \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$, and $c \in [0, 1]$, we define 2 types of non-ground rule templates:

$$R1 : \text{corpus_sim}(S, F)_c \leftarrow s_feat(S, F)_\mu$$

$$R2 : s_feat(S, F)_{\mu'} \leftarrow c_feat(C, F')_\mu$$

$$\wedge \text{contains}(S, C)_1 \wedge \text{associated}(F, F')_1$$

Consider an explanation $\mathcal{E} = p(c, f) : \mu \in \Pi$ formed with a binary predicate $p = c_feat$, and a corresponding set of constants $\mathcal{S} = \mathcal{C}_f \setminus \{f\}$ (which we call suggestive features). We define a parsimony function, σ , that outputs a scalar value (which we call the corpus similarity score) indicating the cultural relevance for a given story $s \in \mathcal{C}_s$.

$$\sigma_{(\Pi, s, f)}(\mathcal{E}) = \Gamma_{\Pi \cup \mathcal{E} \cup \mathcal{O}}^*(\text{corpus_sim}(s, f))$$

$$- \Gamma_{\Pi \cup \mathcal{O}}^*(\text{corpus_sim}(s, f))$$

We define $\mathcal{C}_{c, \mathcal{O}}$ to indicate the set of constants from the ground atoms of \mathcal{O} that belong to \mathcal{C}_c . Similarly, we also define $\mathcal{C}_{f, \mathcal{O}}, \mathcal{C}_{f, \mathcal{E}}$. We define a function ϕ_k as the optimization function to solve our abduction problem that forms a set of new facts of the form $p(c_k, f_i) : \mu_i$ where $f_i \in \mathcal{S}, c_k \in \mathcal{C}_{c, \mathcal{O}}$, that maximizes the corpus similarity score.

$$\phi_k(c, p) = \arg \max_{\substack{\mathcal{E} \subseteq \mathcal{H} \setminus \mathcal{O} \\ |\mathcal{E}|=k, f \in \mathcal{C}_{f, \mathcal{E}}}} \sigma_{(\Pi \cup \mathcal{E}, c, f)}(\mathcal{E})$$

$$s.t. \quad \Pi \cup \mathcal{E} \cup \mathcal{O} \text{ is consistent}$$

We note that we formalize the rule templates and \mathcal{H} to guarantee the consistency of $\Pi \cup \mathcal{O} \cup \mathcal{E}$. As $\forall \mu (a_\mu \in \mathcal{O} \implies \exists \mu' > \mu | a_{\mu'} \in \mathcal{H})$, the explanation will always have an annotation that is above the annotation of the corresponding ground atom in the observation. $\Pi \cup \mathcal{O} \cup \mathcal{E}$ is not inconsistent as $\Gamma_{\Pi \cup \mathcal{O} \cup \mathcal{E}}$ does not result in deducing for an atom assigned with a new annotation that is lower than the initial annotation. Using the lower lattice structure \mathcal{L} , for the annotations and the rule templates, result in annotations of Γ_Π to be higher in the lattice that the initial set of facts.

Example 4.2 (Abduction). *Building upon the Example 4.1, \mathcal{O} is a set of annotated facts with each fact representing a chunk in the given story s , $\{c_feat(c_1, individual_acolades)_0, c_feat(c_1, relationship_framing)_0.4, contains(s, c_1)_1, c_feat(c_2, uniqueness)_0.2, contains(s, c_2)_1 \dots\}$ Hypothesis is also a set of annotated facts, $\mathcal{H} = \{c_feat(c_1, uniqueness)_1\}$, representing hypothetical features of the chunks to change the overall stories narrative to be individualistic. Π is the set of annotated logical rules that specify the individualistic narrative of a given story like:*

$$s_feat(s, ind)_{0.8} \leftarrow c_feat(c_1, uniqueness)_{0.8} \wedge contains(c_1, s)_1 \wedge associated(ind, uniqueness)_1$$

Here, c_1 of the s frames the individual as self-reliant- a quality associated with an individualistic narrative, then the story is of individualistic narrative with a confidence of 0.8. That leads to s being similar to the training corpus with the confidence of 0.56 indicated by the rule:

$$corpus_sim(s, ind)_{0.56} \leftarrow s_feat(s, ind)_{0.8}$$

Explanation $\mathcal{E} = \{c_feat(c_1, uniqueness)_{0.8}\}$ is a suggested set of chunks that to be modified to change the narrative characteristics of the story.

5 Methodology

We consider two narrative transformation tasks both formalized as abduction problem $\langle \mathcal{O}, \mathcal{H}, \Pi \rangle$. Given a story $s \in \mathcal{C}_s$, we search for an explanation $\mathcal{E} \subseteq \mathcal{H}$ such that when chunks from \mathcal{E} are transformed, the story shifts toward a target narrative orientation. Our system combines symbolic reasoning with large language models through a two-phase architecture. Phase 1 uses training corpus to learn rules which are then used in Phase 2 to perform abduction-guided transformation of a test story as seen in Figure 2.

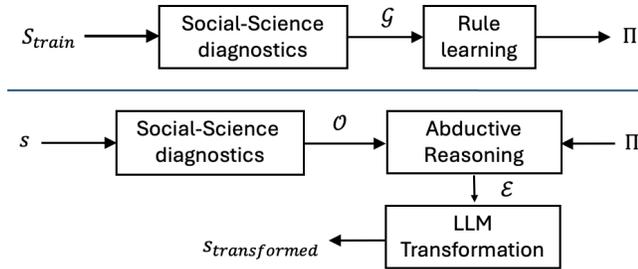


Figure 2: Abduction-guided Transformation (Our Approach).

Phase 1: Rule learning. We have a training corpus $S_{train} \subseteq \mathcal{C}_s$ of stories from a target narrative orientation. We get the observations \mathcal{O} from all the training stories as follows.

$$\mathcal{O}_{train} = \bigcup_{s \in S_{train}} diagnosis(s) \cup \forall_{c \in \mathcal{C}_c^s} \{contains(s, c)\}$$

We then infer the set of ground atoms with a predicate $q = s_feat$ by one application of fixpoint operator, i.e.

$$\Pi_q = \Gamma_{\mathcal{O}_{train} \cup R1}^*(q(s, f))$$

Now, for each feature $f \in \mathcal{C}_f$ and each rating value $\mu \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$, we compute the confidence from observations:

$$conf(f, r) = \frac{|\{s \in S_{train} : \exists q(s, f)_\mu \in \Pi_q\}|}{|S_{train}|}$$

For each (f, μ) pair, we extract a rule of the form $R2$:

$$corpus_sim(X, f)_{conf(f, \mu)} \leftarrow q(X, f)_\mu$$

Phase 2: Iterative Abductive Transformation. For a test story $s_0 \in \mathcal{C}_s$ requiring narrative transformation, we perform iterative abductive reasoning as described in Algorithm 1. At each iteration t , we solve the abduction problem $\langle \mathcal{O}_t, \mathcal{H}, \Pi \rangle$ as formalized in Section 4 to identify the chunks to transform. We define a narrative transformation function that modifies narrative content:

$$llm_transform : \mathcal{C}_c \times \mathcal{C}_f \times \mathcal{C}_s \times [0, 1] \rightarrow \mathcal{C}_s$$

The transformation function takes as input a chunk $c \in \mathcal{C}_c$, a feature $f \in \mathcal{C}_f$, a story $s \in \mathcal{C}_s$, and a target annotation value $\tau \in [0, 1]$. It outputs a modified story $s^{transformed} \in \mathcal{C}_s$ where the selected chunk has been modified by an LLM to shift the narrative toward the target annotation. The LLM prompt is fixed for a given transformation direction but may differ for different directions. The LLM prompt to transform particular identified story chunk given the current narrative and target narrative is provided in the Appendix C.

Algorithm 1 Iterative Abductive Transformation

Require: s_0, Π, T_{max}, k

Ensure: Optimal transformed story s_{t^*} and optimal explanation \mathcal{E}_{t^*}

- 1: $t \leftarrow 0$
 - 2: **while** $t < T_{max}$ **do**
 - 3: $\mathcal{O}_t \leftarrow diagnosis(s_t)$
 - 4: $\mathcal{E}_t \leftarrow \phi_k$ \triangleright Solve $\langle \mathcal{O}_t, \mathcal{H}, \Pi_{train} \rangle$ as defined in Section 4
 - 5: $C_t \leftarrow \{c \in \mathcal{C}_c \mid \exists (f, \mu) c_feature(c, f)_\mu \in \mathcal{E}_t\}$
 - 6: **for** $c \in C_t$ **do**
 - 7: $f, \tau \leftarrow extractFeature(\mathcal{E}_t)$
 - 8: $s_t \leftarrow llm_transform(c, f, s_t, \tau)$
 - 9: **end for**
 - 10: $t \leftarrow t + 1$
 - 11: **end while**
 - 12: $t^* \leftarrow \arg \max_t \{\text{median}(\mathcal{O}_t)\}$
 - 13: **return** $s_{t^*}, \mathcal{E}_{t^*}$
-

Baseline. We implement a baseline transformation function using a zero-shot LLM approach. The baseline prompt is: “*Make the following story more [individualistic/collectivistic]*” where the direction is determined by the target transformation task.

Cost Analysis We show theoretically the bound on the number of LLM calls with Proposition 5.1 (proof in the Appendix A) and that it increases linearly with the number of abduced story chunks.

Proposition 5.1. *Let s be a story represented as a sequence of tokens and C_c^s denote the set of chunks from s . For each chunk $c \in C_c^s$, let $size(c)$ denote the number of tokens in c . Then total number of LLM transformations is bounded by following quantity.*

$$N_c = \sum_{c \in C_c^s} size(c)$$

6 Experimental Results

In this section, we present the experiments conducted to validate our approach. We establish our main evaluation metrics and discuss our findings. We perform a hyperparameter sensitivity study before concluding with our experimental results.

Setup. We create a dataset of 118 natural language narratives (involving 90 individualistic and 28 collectivistic narratives). The individualistic narrative stories in our study were selected based particularly on narratives following a “rags to riches” trajectory common in Western literary traditions [McAdams, 1993] and sourced from classic American literature as well as contemporary biographical narratives of successful figures. Collectivistic narratives, which are less prevalent in Western written fiction, were sourced from diverse geographical and cultural contexts including African, Asian, and Latin American traditions, as well as Western narratives that foreground social relationships and critique individualism [Achebe, 2012; Bâ, 1981; Klinkowitz, 2004].

Narrative Diagnosis. We construct a social science theory based narrative diagnosis survey elaborated in Section 2. The answers to the survey are in the scale of 1-5 with higher scores indicating the higher individualistic or collectivistic perspective of the narrative. Given a candidate story, we use an LLM to answer the survey to obtain ratings of the same story 10 times. We repeat this for some stories and found stable results on the ratings as seen in Figure 3. Note that these ratings are normalized to form the annotations in the output of $diagnosis(s)$ introduced in Section 4.

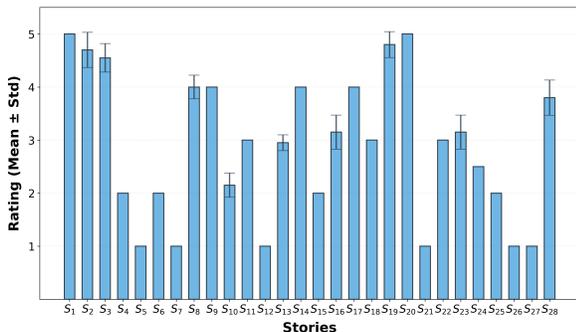


Figure 3: Stability of Social-Science LLM diagnostics over 10 runs for each story.

Evaluation Metrics. We use two metrics. First, to capture the narrative of a story we define diagnosis score— given a narrative, we use a diagnosis to evaluate how individualistic or collectivistic the narrative is based on a survey. We

Table 1: Average improvement over original stories for C→I transformation. Best performance per LLM is underlined; overall best is bolded.

Approach	Base Model	Improvement (%)
Zero-Shot	GPT-4o	26.73
Abduction (Ours)	GPT-4o	<u>97.12</u>
Zero-Shot	Grok-4	61.50
Abduction (Ours)	Grok-4	<u>94.33</u>
Zero-Shot	Llama-4	45.76
Abduction (Ours)	Llama-4	<u>46.54</u>
Zero-Shot	DeepSeek-R1	56.79
Abduction (Ours)	DeepSeek-R1	<u>71.91</u>

compute the median of the 20 diagnosis scores from the narrative diagnosis survey (explained in the Section 2) on transformed stories and report the average improvement over the original in the Table 1 and 2 for different LLMs. We also report the mean, median and mode ratings to compare zero-shot approach and our approach for all the stories (Figures in Appendix D).

The second metric we use is the Kullback-Leibler (KL) divergence to analyze the content similarity between the transformed story and the original story. Let S^{trans} and S be probability distributions of transformed stories and the original story over a shared vocabulary \mathcal{X} , where we employ additive smoothing with $\alpha = 10^{-5}$. Here, $S^{trans}(x) > 0 \implies S(x) > 0$ for all $x \in \mathcal{X}$. While We define KL divergence as:

$$D_{KL}(S^{trans}(x) \parallel S(x)) = \sum_{x \in \mathcal{X}} S_{trans}(x) \log \left(\frac{S^{trans}(x)}{S(x)} \right)$$

We use KL divergence to depict the deviation of the transformed stories through baseline, as well as our approach from the original stories across different language models.

Implementation Details. For all our experiments, we used Apple M4 Pro with 20-core architecture and 24GB RAM. We use reasoning models GPT-4o, Deepseek-R1, Grok-4 fast reasoning model and a non-reasoning model Llama-4 Maverick 17B 128E Instruct v1. The task is to transform a story with a collectivistic to individualistic narrative (C→I) unless specified the other way around (I→C). We use following methods:

Zero-shot Transformation (Baseline). We use a single task instance to transform the story to a different narrative as a prompt to the LLM.

Abduction-guided Transformation. Our algorithm provides suggestion to change characteristics of the story for the transformation. We do this for 3 iterations and report the results of the best of the iterations.

Results of Diagnosis Score and Content Similarity. We analyze the diagnosis scores of the transformed stories from all approaches in comparison to the original stories. Further, we relate these findings to the content similarity with the original stories of the transformations.

For C→I narrative shift, the transformed stories from the abduction-guided method seems to give a higher individualistic rating with an improvement (over the original) of 75.44%

Original (Individualistic):

“She struck another match against the wall. It burned brightly, and when the light hit the wall, it turned see-through like a thin curtain, and she could see inside a room. A snow-white tablecloth was spread out on the table, and on top sat a fancy dinner set. A roast goose steamed beautifully, stuffed with apples and prunes.”

Baseline (Zero-shot LLM): × Maintains individualism

“She struck another match, and this time, the wall became translucent, revealing a room filled with warmth and comfort. A table was set with a snow-white cloth, and a roast goose steamed invitingly. In her vision, the goose approached her, but when the match burned out, the image faded, leaving only the cold wall.”

Our Method: ✓ Collectivistic Shift

“She struck another match against the wall. It burned brightly, and when the light hit the wall, it turned see-through like a thin curtain, and she could see inside a room *where an entire community was gathered, intertwined in a tapestry of shared stories and warmth, their faces aglow with the deep joy that comes from collective unity and mutual support.*”

Original (Individualistic):

“But in the corner, leaning against the wall, sat the little girl—with red cheeks and a smiling face—frozen to death on the last evening of the old year. The New Year’s sun rose on the small, heartbreaking figure.”

Baseline (Zero-shot LLM): × Preserves individualism

“As the community awoke to the New Year, they found the little girl in the corner, her face serene and smiling, though she had succumbed to the cold. The people gathered around, moved by the sight.”

Our Method: ✓ Collectivistic Shift

“But in the corner, leaning against the wall, *the community had gathered around the little girl—forming a circle of warmth and unity, ensuring she was embraced by their collective love and care on the last evening of the old year.*”

Figure 4: Abduction-guided LLM transformation successfully shifts narrative from I→C across multiple story segments.

across all the models. Further, GPT has the highest improvement over original by 97.12% as well as baseline by 55.88%, while maintaining an improvement in content similarity with the original stories by 40.40% than the baseline (as seen in Figure 5). Comparable results were observed for Grok and Deepseek. Although Llama only had an improvement of 7.63% over the baseline, it preserved the similarity with original story by 49.16%. Moreover, the baseline consistently gave a higher divergence from the original stories while our approach preserves more information from the original story.

On the other hand, for I→C narrative shift, we observe a 89.64% average improvement across the models over original stories as seen in Table 2. Specifically, GPT outperforms the baseline by 15% while preserving the similarity with the original stories by 36% than the baseline (as seen in Figure 6). Comparable results were observed for Grok and Llama. Although Deepseek outperformed our approach by 6%, it lost information from the original stories by 39.10% compared to our approach. A study [Segerer, 2025] finds that DeepSeek produces more collectivist-oriented content than other LLMs, indicating a stronger collectivistic narrative ori-

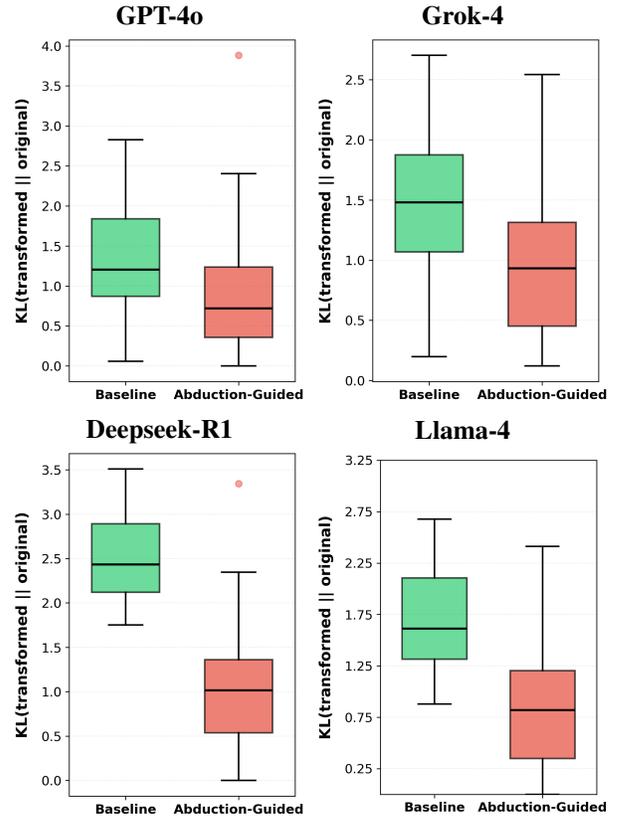


Figure 5: The KL divergence across the stories for C→I narrative shift. Top left: GPT-4o, Top right: Grok-4, Bottom left: Deepseek-R1, Bottom right: Llama-4.

Table 2: Average improvement over original stories for I→C transformation. Best performance per LLM is underlined; overall best is bolded.

Approach	Base Model	Improvement (%)
Zero-Shot	GPT-4o	75.00
Abduction (Ours)	GPT-4o	<u>95.39</u>
Zero-Shot	Grok-4	86.63
Abduction (Ours)	Grok-4	<u>97.32</u>
Zero-Shot	Llama-4	71.43
Abduction (Ours)	Llama-4	<u>85.71</u>
Zero-Shot	DeepSeek-R1	<u>87.67</u>
Abduction (Ours)	DeepSeek-R1	<u>73.68</u>

entation, which is further supported by its baseline performance in collectivistic story transformation in Table 2. A qualitative example of I→C is shown in Figure 4 demonstrating how abduction-guided transformation targets specific narrative chunks while preserving narrative structure. Similar C→I example is provided in Appendix D.

We found that, based on the diagnosis with GPT-4o, for the C→I shift, 9 stories in the collectivistic corpus were already closer to the individualistic end, while 6 were diagnosed as neutral narratives. While, for the I→C, shift, 2 stories were

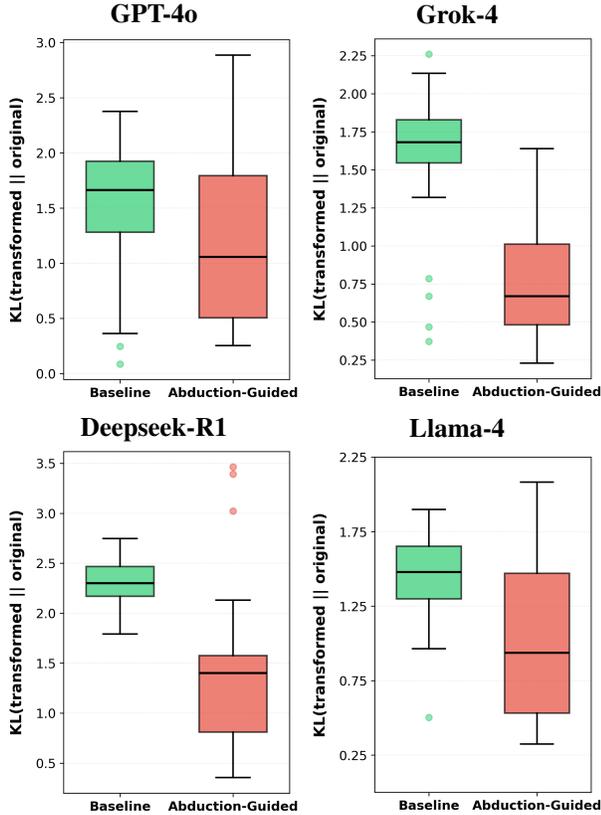


Figure 6: The KL divergence across the stories for I→C narrative shift. Top left: GPT-4o, Top right: Grok-4, Bottom left: Deepseek-R1, Bottom right: Llama-4.

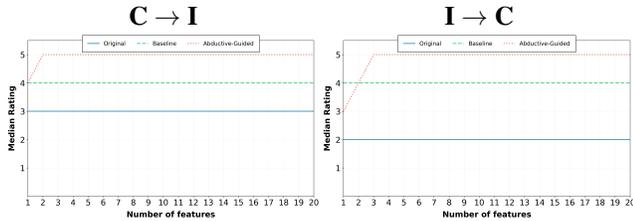


Figure 7: Hyperparameter sensitivity. Left: C→I, Right: I→C.

diagnosed as neutral narratives. All of these stories were successfully transformed into the target narrative. Additionally, our abduction-guided approach demonstrates efficiency by targeting only essential narrative segments.

Overall, for majority models, abduction-guidance improves the baseline for shifting the narrative while maintaining fidelity. Deepseek as a baseline performs competitively with our approach to shift to a collectivistic narrative but loses content similarity from the original story. Moreover, all models perform better at transforming stories into collectivistic than into individualistic narrative.

Hyperparameter Sensitivity. We allow for a hyperparameter, number of features, which is the number of features (ψ) in the solution of $\langle \mathcal{O}, \mathcal{H}, \Pi \rangle$ from Section 4. We report the diagnosis rating across different ψ in Figure 7. For C→I,

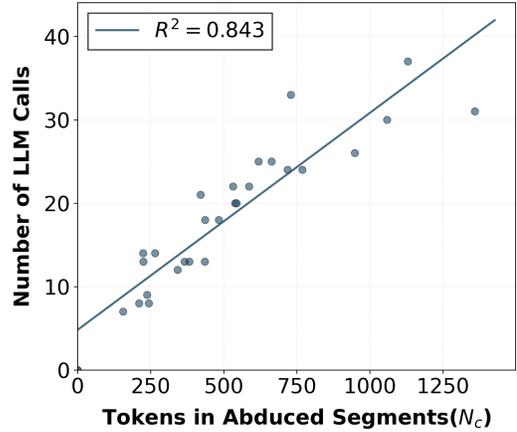


Figure 8: Total tokens of identified segments vs Number of LLM calls for gpt-4o (C→I).

the performance of our approach saturates for $\psi = 2$ at 0.66 times the original and 0.25 times the baseline. For I→C, our approach has a neutral narrative at $\psi = 1$ and shifts toward the target narrative with the increase in ψ while outperforming the baseline before saturating for $\psi = 3$. It achieves the maximum score- 5 while outperforming the original by 150% and the baseline by 25%.

Token usage evaluation. Our approach shows efficiency through targeted chunk modification. Across all stories, only a median of 32.11% of tokens required transformation, with the remaining content preserved unchanged. Figure 8 confirms the theoretical bound of Proposition 5.1, showing a linear relationship between tokens in identified segments (N_c) and number of LLM calls with $R^2 = 0.843$.

7 Conclusion

Aligning text to an audience’s view is significant in journalism, social science and legal communities. While LLMs have been seen to excel in text generation, they seem to lack mechanisms for transforming cultural narrative hidden in the stories while preserving semantic fidelity. Narrative transformation is a high level reasoning task that requires social science domain knowledge, hence we propose a neurosymbolic approach that combines abductive reasoning with LLM based social science diagnostics to shift the narrative of a given story. Our approach shows significant improvement over zero-shot prompting where it either fails to transform the narrative or generate incoherent transformations. Across both reasoning and non-reasoning models, our abduction-guided approach showed significant improvement over zero-shot baseline. For GPT-4o, our approach outperformed zero-shot narrative transformation by 55.88% while maintaining 40.40% more semantic similarity with the original stories for collectivistic-to-individualistic transformation. This work can be extended to transform the story to align it to a specific cultural audience. Another line of direction is to look at a richer temporal rule structure to inform the abductive transformation.

Acknowledgments

This research was funded under the DARPA CONNECT program.

References

- [Achebe, 2012] Chinua Achebe. *Hopes and impediments: Selected essays*. Penguin, 2012.
- [Aditya et al., 2023] Dyuman Aditya, Kaustuv Mukherji, Srikar Balasubramanian, Abhiraj Chaudhary, and Paulo Shakarian. Pyreason: Software for open world temporal logic, 2023.
- [Bâ, 1981] Amadou Hampaté Bâ. The living tradition. *General history of Africa*, 1:166–205, 1981.
- [Bavikadi et al., 2025a] Divyagna Bavikadi, Dyuman Aditya, Devendra Parkar, Paulo Shakarian, Graham Mueller, Chad Parvis, and Gerardo I. Simari. Geospatial trajectory generation via efficient abduction: Deployment for independent testing. *Electronic Proceedings in Theoretical Computer Science*, 416:274–287, February 2025.
- [Bavikadi et al., 2025b] Divyagna Bavikadi, Nathaniel Lee, Paulo Shakarian, and Chad Parvis. Sea-cret agents: Maritime abduction for region generation to expose dark vessel trajectories. *ArXiv*, abs/2502.01503, 2025.
- [Bercovitch, 1975] Sacvan Bercovitch. *The Puritan origins of the American self*. Yale University Press, 1975.
- [Bruner, 1991] Jerome Bruner. The narrative construction of reality. *Critical inquiry*, 18(1):1–21, 1991.
- [Choi, 2025] Hoon-Seok Choi. Individualism–collectivism and intergroup bias. *Asian Journal of Social Psychology*, 28(2):e70010, 2025.
- [Chun et al., 2025] Jiwon Chun, Gefei Zhang, and Meng Xia. ConflictLens: Llm-based conflict resolution training in romantic relationship. In *Adjunct Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*, pages 1–3, 2025.
- [Green and Appel, 2024] Melanie C. Green and Markus Appel. Chapter one - narrative transportation: How stories shape how we see ourselves and the world. volume 70 of *Advances in Experimental Social Psychology*, pages 1–82. Academic Press, 2024.
- [Green and Brock, 2000] Melanie C Green and Timothy C Brock. The role of transportation in the persuasiveness of public narratives. *Journal of personality and social psychology*, 79(5):701, 2000.
- [Hofstede, 1984] Geert Hofstede. *Culture’s consequences: International differences in work-related values*, volume 5. sage, 1984.
- [Javadi et al., 2025] Vahid Sadiri Javadi, Johanne R Trippas, Yash Kumar Lal, and Lucie Flek. Can stories help llms reason? curating information space through narrative. In *The 2nd Workshop on Analogical Abstraction in Cognition, Perception, and Language (Analogy-Angle II)*, page 92, 2025.
- [Jenner et al., 2025] Sarah Jenner, Dimitris Raidos, Emma Anderson, Stella Fleetwood, Ben Ainsworth, Kerry Fox, Jana Kreppner, and Mary Barker. Using large language models for narrative analysis: a novel application of generative ai. *Methods in Psychology*, 12:100183, 2025.
- [Kifer and Subrahmanian, 1992] Michael Kifer and V.S. Subrahmanian. Theory of generalized annotated logic programming and its applications**a preliminary report on this research has appeared in [34]. *The Journal of Logic Programming*, 12(4):335–367, 1992.
- [Klinkowitz, 2004] Jerome Klinkowitz. *The Vonnegut Effect*. Univ of South Carolina Press, 2004.
- [Li et al., 2024] Huihan Li, Yuting Ning, Zeyi Liao, Siyuan Wang, Xiang Lorraine Li, Ximing Lu, Wenting Zhao, Faeze Brahma, Yejin Choi, and Xiang Ren. In search of the long-tail: Systematic generation of long-tail inferential knowledge via logical rule guided search. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2348–2370, Miami, Florida, USA, November 2024. Association for Computational Linguistics.
- [Markus and Kitayama, 2014] Hazel Rose Markus and Shinobu Kitayama. Culture and the self: Implications for cognition, emotion, and motivation. In *College student development and academic life*, pages 264–293. Routledge, 2014.
- [McAdams, 1993] Dan P McAdams. *The stories we live by: Personal myths and the making of the self*. Guilford press, 1993.
- [Moon, 2020] Michael Moon. *A small boy and others: Initiation and initiation in American culture from Henry James to Andy Warhol*. Duke University Press, 2020.
- [Mukherjee et al., 2024] Sourabrata Mukherjee, Atul Kr Ojha, and Ondrej Dusek. Are large language models actually good at text style transfer? *CoRR*, 2024.
- [Mukherjee et al., 2025] Sourabrata Mukherjee, Atul Kr Ojha, John Philip McCrae, and Ondřej Dušek. Evaluating text style transfer evaluation: Are there any reliable metrics? In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 4: Student Research Workshop)*, pages 418–434, 2025.
- [Park et al., 2023] Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023.
- [Paul and Matthews, 2016] Christopher Paul and Miriam Matthews. The russian “firehose of falsehood” propaganda model. *Rand Corporation*, 2(7):1–10, 2016.
- [Pei et al., 2024] Jonathan Pei, Zeeshan Patel, Karim El-Refaï, and Tianle Li. Swag: Storytelling with action guid-

- ance. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14086–14106, 2024.
- [Piper and Bagga, 2024] Andrew Piper and Sunyam Bagga. Using large language models for understanding narrative discourse. In Yash Kumar Lal, Elizabeth Clark, Mohit Iyyer, Snigdha Chaturvedi, Anneliese Brei, Faeze Brahman, and Khyathi Raghavi Chandu, editors, *Proceedings of the 6th Workshop on Narrative Understanding*, pages 37–46, Miami, Florida, USA, November 2024. Association for Computational Linguistics.
- [Reif *et al.*, 2022] Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. A recipe for arbitrary text style transfer with large language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 837–848, 2022.
- [Rooein *et al.*, 2025] Donya Rooein, Vilém Zouhar, Debora Nozza, and Dirk Hovy. Biased tales: Cultural and topic bias in generating children’s stories. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 52–72, 2025.
- [Segerer, 2025] Robin Segerer. Cultural value alignment in large language models: A prompt-based analysis of schwartz values in gemini, chatgpt, and deepseek. 05 2025.
- [Shakarian and Simari, 2022] Paulo Shakarian and Gerardo I. Simari. Extensions to Generalized Annotated Logic and an Equivalent Neural Architecture . In *2022 Fourth International Conference on Transdisciplinary AI (TransAI)*, pages 63–70, Los Alamitos, CA, USA, September 2022. IEEE Computer Society.
- [Singelis and Brown, 1995] Theodore M Singelis and William J Brown. Culture, self, and collectivist communication: Linking culture to individual behavior. *Human communication research*, 21(3):354–389, 1995.
- [Singelis, 1994] Theodore M Singelis. The measurement of independent and interdependent self-construals. *Personality and social psychology bulletin*, 20(5):580–591, 1994.
- [Singh *et al.*, 2025] Jyotika Singh, Weiyi Sun, Amit Agarwal, Viji Krishnamurthy, Yassine Benajiba, Sujith Ravi, and Dan Roth. Can LLMs narrate tabular data? an evaluation framework for natural language representations of text-to-SQL system outputs. In Saloni Potdar, Lina Rojas-Barahona, and Sebastien Montella, editors, *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 883–902, Suzhou (China), November 2025. Association for Computational Linguistics.
- [Suzgun *et al.*, 2022] Mirac Suzgun, Luke Melas-Kyriazi, and Dan Jurafsky. Prompt-and-rerank: A method for zero-shot and few-shot arbitrary textual style transfer with small language models. *arXiv preprint arXiv:2205.11503*, 2022.
- [Tao *et al.*, 2025] Zhen Tao, Dinghao Xi, Zhiyu Li, Liumin Tang, and Wei Xu. Cat-llm: Style-enhanced large language models with text style definition for chinese article-style transfer. *ACM Transactions on Knowledge Discovery from Data*, 19(7):1–33, 2025.
- [Tian *et al.*, 2024] Yufei Tian, Tenghao Huang, Miri Liu, Derek Jiang, Alexander Spangher, Muhao Chen, Jonathan May, and Nanyun Peng. Are large language models capable of generating human-level narratives? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17659–17681, 2024.
- [Triandis, 1995] HC Triandis. Individualism and collectivism west view press. *Boulder, CO*, 1995.
- [Yue *et al.*, 2025] Matthew Yue, Zhikun Xu, Vivek Gupta, Thao Ha, Liesal Sharabi, and Ben Zhou. Relatesim: Leveraging turning point theory and llm agents to predict and understand long-term relationship dynamics through interactive narrative simulations. *arXiv preprint arXiv:2510.00414*, 2025.

A Formal Proof of Runtime Result

We show theoretically the bound on the number of LLM calls with Proposition A.1 and that it increases linearly with the number of abducted story chunks.

Proposition A.1. *Let s be a story represented as a sequence of tokens and \mathcal{C}_c^s denote the set of chunks from s . For each chunk $c \in \mathcal{C}_c^s$, let $\text{size}(c)$ denote the number of tokens in c . Then total number of LLM transformations is bounded by following quantity.*

$$N_c = \sum_{c \in \mathcal{C}_c^s} \text{size}(c)$$

Proof. We define k as the number of LLM transformation calls in a single iteration of abduction-guided approach. Assume BWOC that $k > N_c$.

We define N' as the total tokens in story s :

$$N' = N_c + N_{\bar{c}}$$

where $N_{\bar{c}}$ denotes tokens not in any identified chunk.

$$N_c = N' - N_{\bar{c}}$$

By our assumption, $k > N_c$,

$$k > N' - N_{\bar{c}}$$

Adding $N_{\bar{c}}$ on both sides,

$$k + N_{\bar{c}} > N'$$

However, by definition N' is the maximum possible tokens in the story and thus $k + N_{\bar{c}}$ (LLM calls in addition to non-identified tokens) cannot exceed N' . This is a contradiction. Therefore, $k \leq N_c$. \square

B Narrative Diagnostic Survey

We provide complete set of individualistic narrative based diagnostic questions in Figure 9 and collectivistic narrative based diagnostic questions in Figure 10.

C Details on Systems Prompts

We present the system prompt used for transforming the story with the abducted changes. The LLM prompt to transform particular identified story chunk given the current narrative and target narrative is provided in the Figure 11.

D Additional Experimental Results

In figure 12, 13, 14, 15, we report the median of diagnosis scores and the blue solid bar shows the diagnosis scores, which forms a lower envelope indicating the narrative of the original stories to be on the collectivistic (left column) or individualistic (right column) end. Similarly, 16, 17, 18, 19 shows the mode and 20, 21, 22, 23 median of the diagnosis scores for all of the LLMs.

A qualitative example of C→I is shown in Figure 24 demonstrating how abduction-guided transformation targets specific narrative chunks while preserving narrative structure.

Feature 1: Protagonist-Centered Focus

“Is the narrative driven primarily by one character’s experiences, choices, and inner life?”

Feature 2: Internal Goals

“Does the protagonist pursue personal ambitions (e.g., self-actualization, fame, self-expression) over collective objectives?”

Feature 3: Decision-Driven Plot

“Are key turning points in the story determined by the protagonist’s own decisions rather than group mandates or fate?”

Feature 4: Self-Reliance

“Does the protagonist overcome obstacles through their own resourcefulness rather than relying on communal support?”

Feature 5: Individual Accolades

“Are awards, status, or recognition attributed primarily to the single protagonist rather than to a team or ensemble?”

Feature 6: Meritocracy Emphasis

“Is success portrayed as earned by the protagonist’s talent, hard work, or innate brilliance, rather than by lineage or group standing?”

Feature 7: “Man vs. Self/World” Conflict

“Is the central struggle internal (e.g., self doubt, identity) or between the protagonist and external forces, rather than group conflicts?”

Feature 8: Solo Confrontations

“Do climactic showdowns feature the lone protagonist facing the antagonist or obstacle, rather than a collaborative effort?”

Feature 9: Inner Journey

“Is the character arc centered on the protagonist discovering their own values, strengths, or purpose?”

Feature 10: Uniqueness & Self-Expression

“Are characters celebrated for what makes them unique (quirks, dreams) or for ‘being true to themselves’?”

Feature 11: Self-Construal

“Does the narrative present the self as stable and independent, defined by personal traits rather than social roles?”

Feature 12: Behavioral Guidance

“Are actions in the story guided by the protagonist’s personal attitudes and preferences rather than social norms or group expectations?”

Feature 13: Relationship Orientation

“Does the story depict relationships as optional and based on mutual benefit rather than duty and loyalty?”

Feature 14: Primary Conflict

“Is the central conflict about asserting one’s individual identity or resisting conformity?”

Feature 15: Resolution Style

“Does the story resolve conflicts by standing up for personal rights and achieving justice rather than through compromise and reconciliation?”

Feature 16: Moral Emphasis

“Does the narrative emphasize autonomy, personal integrity, or self-actualization as moral virtues?”

Feature 17: Relationship Framing

“Does the story frame relationships as non-essential, allowing the protagonist to pursue goals independently?”

Feature 18: Vertical Individualism

“Does the story accept social inequality as a natural consequence of individual achievement?”

Feature 19: Personal Ethics over Group Norms

“Does the protagonist’s personal code or conscience take precedence over cultural or familial expectations?”

Feature 20: Self-Actualization Climax

“Does the emotional payoff come from the protagonist’s personal breakthrough instead of restoring group harmony?”

Figure 9: Complete individualistic diagnostic questions: 20 questions assessing individualistic narrative orientation.

Feature 1: Protagonist-Centered Focus

“Is the narrative driven primarily by the group’s or community’s shared experiences, collective choices, and communal identity?”

Feature 2: Internal Goals

“Do the characters pursue group ambitions (e.g., community well-being, family honor, shared success) over individual goals?”

Feature 3: Decision-Driven Plot

“Are key turning points in the story determined by collective decisions, group mandates, or community traditions rather than by a single individual’s decision?”

Feature 4: Self-Reliance

“Do the characters overcome obstacles through communal support, collective action, or shared resources rather than individual effort?”

Feature 5: Individual Accolades

“Are awards, status, or recognition attributed primarily to the group, ensemble, or community effort rather than to an individual?”

Feature 6: Meritocracy Emphasis

“Is success portrayed as resulting from group support, family lineage, or communal contributions rather than solely individual talent?”

Feature 7: “Man vs. Self/World” Conflict

“Is the central struggle between the group and external forces or within group cohesion, rather than an individual’s internal conflict?”

Feature 8: Solo Confrontations

“Do climactic showdowns feature collaborative group efforts or collective confrontation, rather than a lone individual?”

Feature 9: Inner Journey

“Is the character arc centered on the group discovering shared values, collective strengths, or communal purpose?”

Feature 10: Uniqueness & Self-Expression

“Are characters celebrated for conforming to group norms, fulfilling social roles, or contributing to the collective identity rather than uniqueness?”

Feature 11: Self-Construal

“Does the narrative present the self as interdependent, defined by social roles and relationships rather than solely personal traits?”

Feature 12: Behavioral Guidance

“Are actions in the story guided by social norms, group expectations, or communal values rather than personal preferences?”

Feature 13: Relationship Orientation

“Does the story depict relationships as based on duty, loyalty, obligation, and collective well-being rather than solely mutual benefit?”

Feature 14: Primary Conflict

“Is the central conflict about maintaining social harmony, fulfilling collective roles, or adhering to group norms rather than asserting individuality?”

Feature 15: Resolution Style

“Does the story resolve conflicts through compromise, reconciliation, and restoring group harmony rather than individual vindication?”

Feature 16: Moral Emphasis

“Does the narrative emphasize group solidarity, communal responsibility, or collective welfare as moral virtues?”

Feature 17: Relationship Framing

“Does the story frame relationships as essential, requiring individuals to consider the impact of their actions on family and community?”

Feature 18: Vertical Individualism

“Does the story emphasize equality, shared prosperity, and collective welfare over social inequality and individual hierarchy?”

Feature 19: Personal Ethics over Group Norms

“Does adherence to cultural norms, family expectations, or communal codes take precedence over personal preferences?”

Feature 20: Self-Actualization Climax

“Does the emotional payoff come from restoring group harmony, collective well-being, or communal success rather than an individual breakthrough?”

Figure 10: Complete collectivistic diagnostic questions: 20 questions assessing collectivistic narrative orientation.

System: You are given a [source_type] story below.
**** start of story ****
 [story]
**** end of story ****
 Your goal is to make the story more [target_type]. To make it more [target_type], you will update only the selected segment from the story which is provided below.
Selected segment: ** [segment] **
 Now, rewrite the whole story by updating only the selected segment to make the story more [target_type]. Don't change other parts of the story, and just output the rewritten story, nothing else.

Figure 11: Transformation prompt template. Placeholders [source_type], [target_type], [story], and [segment] are instantiated for each identified segment.

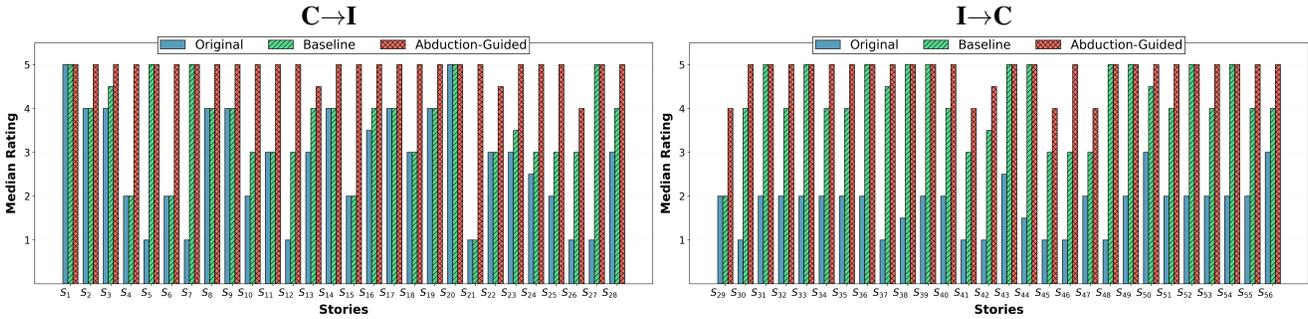


Figure 12: The median of diagnosis ratings for GPT-4o across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

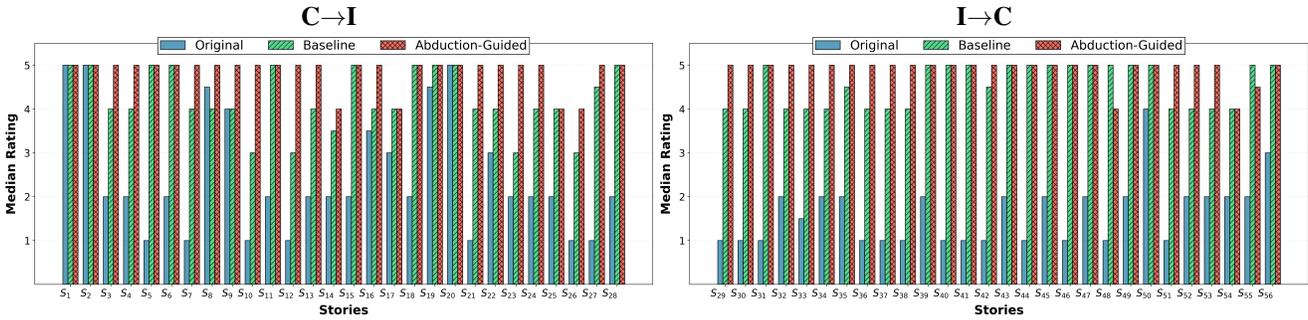


Figure 13: The median of diagnosis ratings for Grok-4 across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

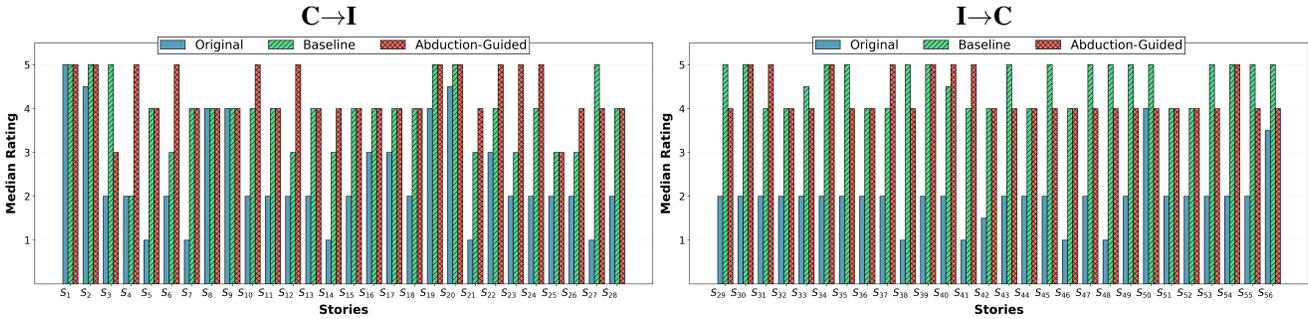


Figure 14: The median of diagnosis ratings for Deepseek-R1 across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

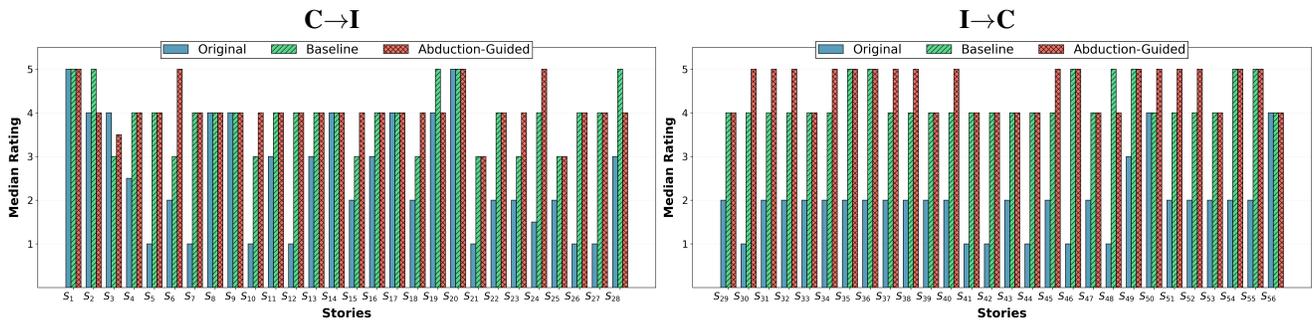


Figure 15: The median of diagnosis ratings for Llama across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

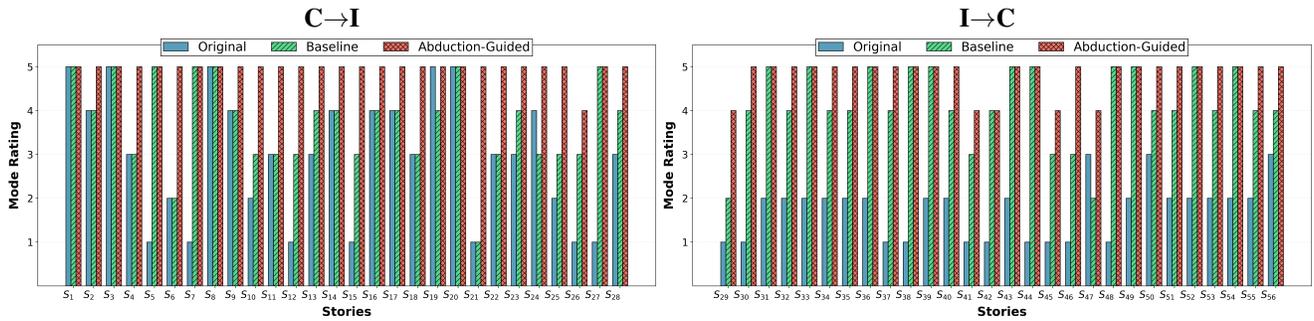


Figure 16: The mode of diagnosis ratings for GPT-4o across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

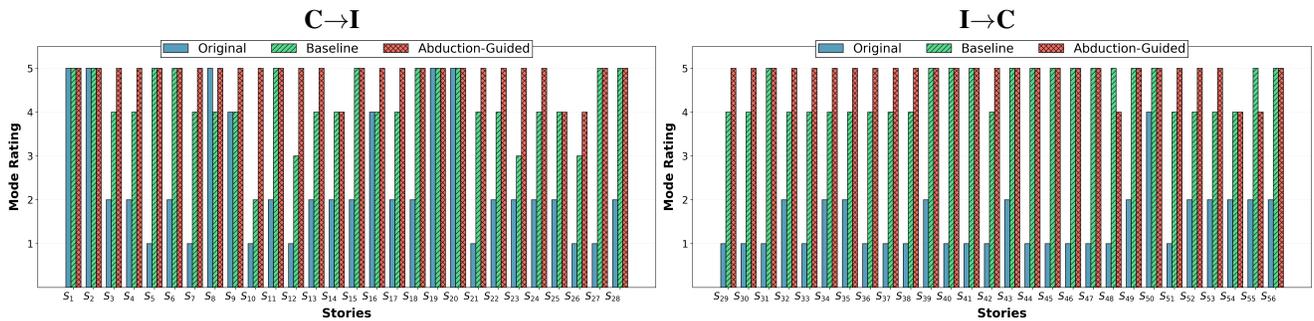


Figure 17: The mode of diagnosis ratings for Grok-4 across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

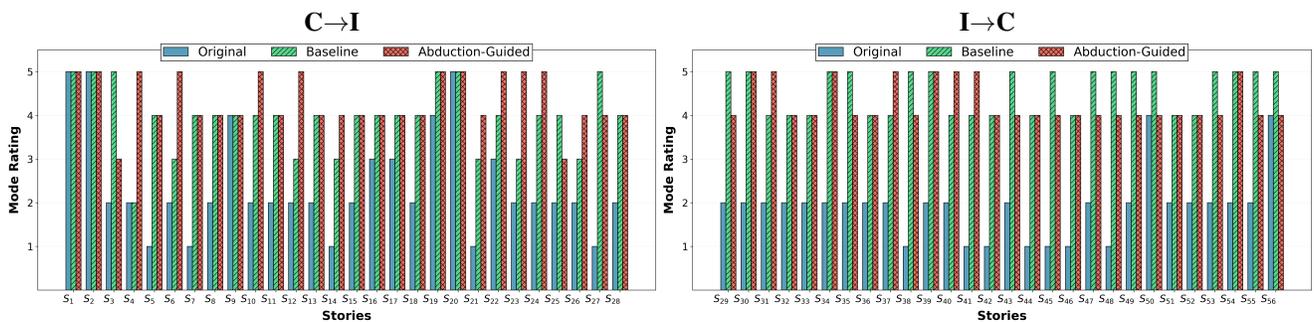


Figure 18: The mode of diagnosis ratings for Deepseek-R1 across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

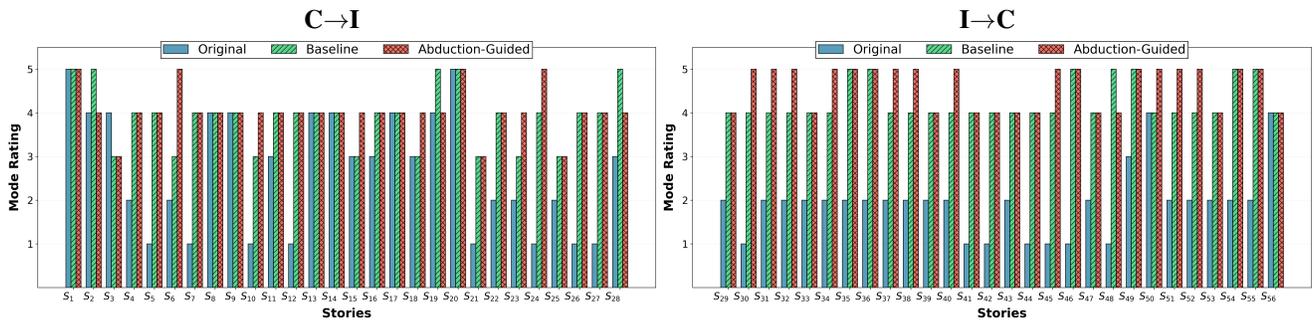


Figure 19: The mode of diagnosis ratings for Llama across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

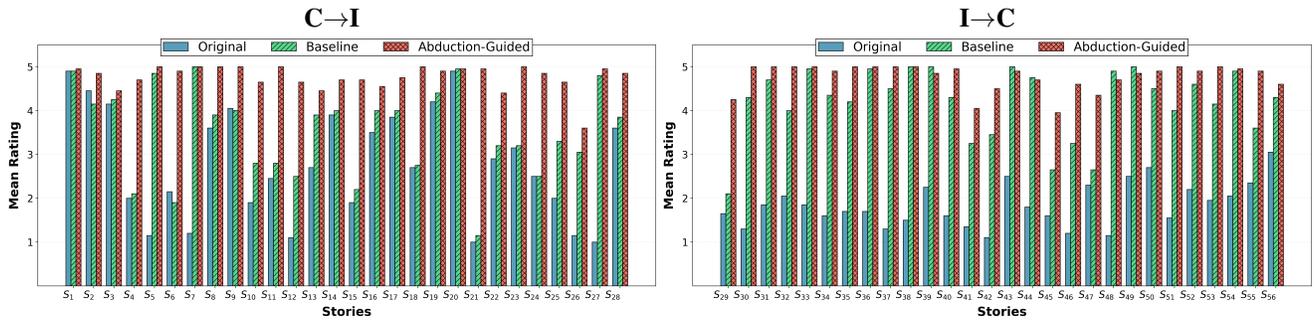


Figure 20: The mean of diagnosis ratings for GPT-4o across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

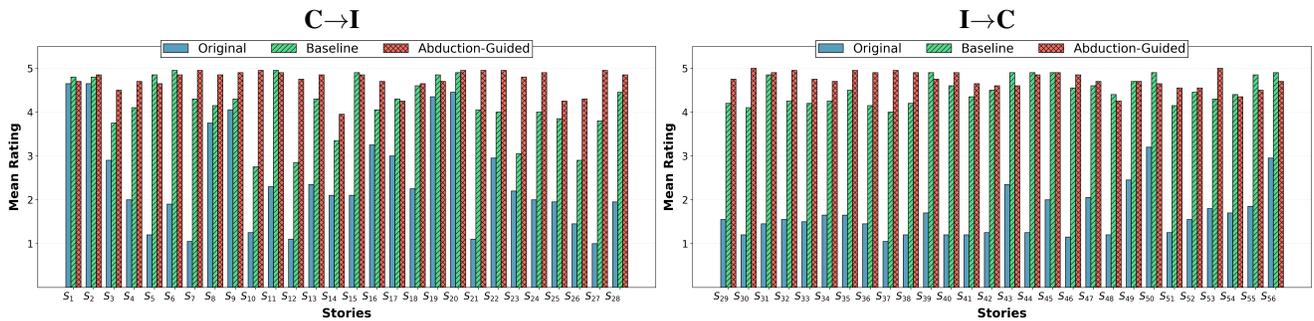


Figure 21: The mean of diagnosis ratings for Grok-4 across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

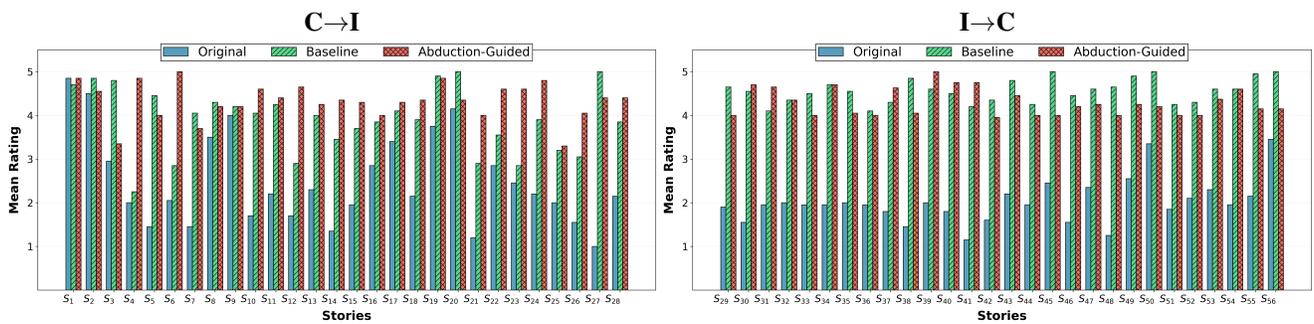


Figure 22: The mean of diagnosis ratings for Deepseek-R1 across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

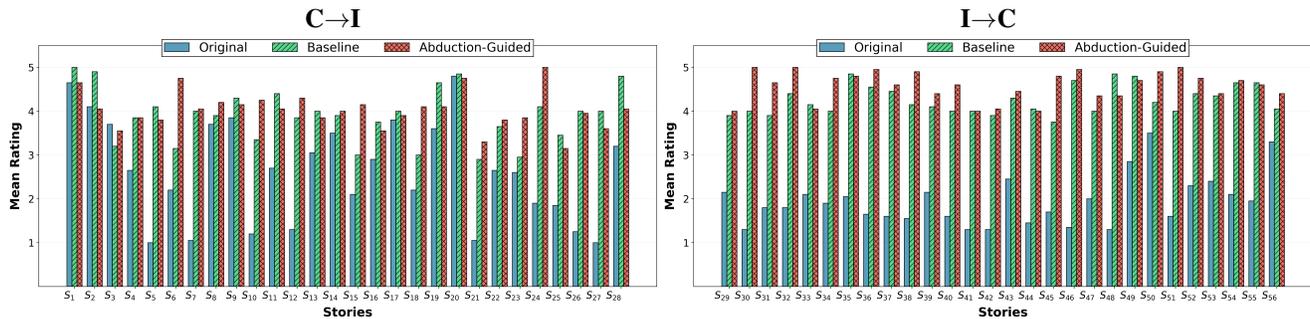


Figure 23: The mean of diagnosis ratings for Llama across the stories in the dataset. Left: C→I, Right: I→C. Note that for both directions higher rating is better.

Original (Collectivistic):

The group of [women] from the community are warriors — women with great courage, enthusiasm, and a spirit of service, said Fernando, a Convoy of Hope agriculture technician in Nicaragua. We have often seen them give up their time to help other people. This work that we have done supporting the Agriculture program has brought us other experiences, the Warrior Women said. Agriculture is not just a job for men. Now we feel empowered to produce our own food. We are grateful to all . . . of Convoy of Hope, especially to the facilitators. We believe that they have given us the opportunity to learn, and that is why we now have the experiences to continue improving our lives, our family, and our community.

Baseline (Zero-shot LLM): ✗ Maintains collectivistic framing

“Andria and her group are true warriors,” praises Fernando, a Convoy of Hope agriculture technician. “Their spirit of service is remarkable, and their contribution invaluable. Through this journey, Andria has gained more than agricultural skills. “Agriculture isn’t just for men,” she asserts. “We’ve learned to produce our own food and feel empowered by it.” *Grateful for the support from Convoy of Hope, Andria is optimistic about the future. “The opportunity to learn has enriched our lives, our families, and our community,” she reflects.*”

Our Method: ✓ Individualistic Shift

“Each woman is a warrior in her own right — an individual with extraordinary courage, enthusiasm, and a personal spirit of service, said Fernando, a Convoy of Hope agriculture technician in Nicaragua. We have often seen them give up their time to help other people. *This journey I have undertaken in supporting the Agriculture program has brought me other valuable experiences,*” one of the Warrior Women said. “Agriculture is not just a job for men. Now I feel empowered to produce my own food. *I am proud of my own determination and efforts. It’s my drive and ambition that have allowed me to learn and gain the experiences needed to improve my life and the lives of those around me.*”

Figure 24: Abduction-guided LLM transformation successfully shifts narrative from C→I.