

# Manzoni AI

## *Bringing 19th-Century Italian Prose to Life through Stylistic Fine-Tuning of a Modern LLM*

**Matteo Bernardi**

*La Sapienza University of Rome*

Dedicated to my professor Pupillo  
who sparked my passion for the works of Alessandro Manzoni

**Abstract**—This work presents **Manzoni AI**, a language model based on Gemma 4B IT, fine-tuned to transform contemporary Italian sentences into a style inspired by Alessandro Manzoni’s *I Promessi Sposi*. The model was trained on a custom parallel dataset featuring modern-to-Manzonian rewritings that were semi-automatically generated and manually refined. Results show that the model effectively reproduces complex syntax, refined vocabulary, and stylistic traits typical of 19th-century Italian. This approach shows promise for applications in education, literary scholarship, and the preservation of linguistic heritage.

### I. MOTIVATION

When I was fifteen, I read *I Promessi Sposi* by Alessandro Manzoni at school. I was lucky enough to have a teacher who didn’t just have us read the novel once or summarize it for oral exams. On the contrary, especially in the early lessons, she guided us through a meticulous, word-by-word analysis of each chapter, teaching us how to read deeply and attentively.

It quickly became clear that the focus wasn’t so much the plot — fascinating, engaging, and certainly instructive for teenagers — but rather the language itself, the way things were said: the irony, the rhetoric, the syntax — everything that made those pages radically different from the everyday speech we were used to. It was, in a word, extraordinary.

Week after week, I found myself absorbing that style. And how could it have been otherwise? At the time, literature and the humanities were my main passion. I realized that when someone is immersed in a certain environment — and especially if they love it — they end up assimilating it, like a sponge. I’ve often wondered whether, beyond the purely educational aim, that painstaking instruction wasn’t also meant to transform how we expressed ourselves — to elevate our language.

Years later, I underwent a surprisingly similar — and deeply human — experience with a “simple” language model with four billion parameters. And just like that fifteen-year-old version of me absorbed Manzoni’s style, this model, once properly fine-tuned, began to imitate it.

To be fair, the results may not always be perfect — but they are often remarkable, and would surely have made that fifteen-year-old boy’s eyes light up. That said, there’s a fundamental difference: the model, in its current form, still can’t generate new Manzonian prose, weaving original stories that faithfully

mimic his stylistic and narrative mechanics. Yet that boy, all those years ago, could — and did.

I still vividly remember a literature test in which the teacher asked us to write an alternative continuation of a chapter. I didn’t just jot down a few lines — I dove in and wrote with enthusiasm, straight into the final draft I intended to hand in. This earned me a spontaneous remark from the teacher, slightly surprised at how quickly it came together. I managed to accomplish what I had set out to do, and in the end, I was truly proud of what I had produced. But upon reflection, that spontaneous writing process closely resembles the operational mechanism of a large language model: generating coherent text incrementally, word by word, in real time from a minimal prompt—except for one key difference: a profound inventiveness.

*Manzoni AI* might not be able to do that — yet — but it’s already capable of producing some truly remarkable “transliterations.”

### II. INTRODUCTION

In recent years, the advent of Large Language Models (LLMs) has profoundly transformed the landscape of computational linguistics and natural language processing. Models such as GPT, PaLM, LLaMA, and Gemma have demonstrated remarkable capabilities in understanding and generating natural language, tackling complex tasks such as machine translation, summarization, semantic classification, question answering, and creative writing. One of the most intriguing developments is the ability to exploit these models for style transfer operations—rewriting texts while preserving semantic content and modifying tone, register, or stylistic identity.

Within this context, *Manzoni AI* is a research project that applies the latest LLM-based technologies to a stylistic and cultural challenge: transforming contemporary Italian sentences into a literary style inspired by Alessandro Manzoni’s *I Promessi Sposi*. The goal is not simply to mimic an archaic or formal register, but to coherently replicate the complex syntactic structures, refined lexical choices, and rhetorical devices that characterize Manzoni’s prose.

The core model used is *Gemma 4B IT*, an open-source language model optimized for Italian, developed by Google. The model was fine-tuned on a custom parallel dataset in

which each example pairs a modern Italian sentence with its Manzonian rewriting. This dataset was initially generated semi-automatically using other LLMs that perform well in style transfer, and then quickly revised manually to ensure stylistic fidelity and linguistic quality.

The scientific interest of this project operates on multiple levels. From a technical perspective, *Manzoni AI* explores the ability of general-purpose language models to learn highly specific stylistic transformations anchored in historical and literary contexts. From a cultural perspective, it aims to establish a dialogue between artificial intelligence and the Italian literary heritage, proposing an innovative approach to the preservation and appreciation of classical authorial language.

Finally, the project opens the door to several potential applications: educational tools for teaching literature and Italian language, writing assistants for stylized expression, creative experiments in augmented literature, and, more broadly, technologies that promote a more active and engaging access to historical texts. With *Manzoni AI*, we seek to contribute meaningfully to the intersection of advanced language modeling, computational creativity, and the enhancement of cultural heritage.

### III. DATASET

To train a language model capable of transforming sentences from contemporary Italian into the Manzonian style, it is necessary to have a parallel dataset containing sentence pairs with equivalent meaning, but expressed in the two different linguistic registers.

The dataset was created through a *reverse dataset creation* process, starting from the original Manzonian text and proceeding towards the modern version of the sentences. Specifically, the full text of *I Promessi Sposi* was used, excluding the introduction. This choice was necessary because the introduction contains excerpts from an original manuscript written in a seventeenth-century register, which is very distant from the nineteenth-century language of the novel. This exclusion was intended to prevent confusion during the model’s training phase.

The text was then split into 38 *chunks*, each corresponding to a chapter of the novel. The next step would ideally have involved a manual and meticulous paraphrasing effort by expert literary scholars, to transform each Manzonian sentence into its modern equivalent while preserving its original meaning. Unfortunately, due to the lack of qualified human resources and time, an automatic approach was adopted.

Since adequate hardware was not available to locally run models large enough to comprehend the complex Manzonian syntax, it was necessary to choose a solution based on cloud resources—and, due to lack of funding, free solutions. The modern rewritings were therefore generated by an external language model, accessed via the API of an online service offering “various” free trial options. Although very large models were available, and initially seemed most appropriate for the task, it became evident that they presented significant issues in terms of response time and their tendency to

not fully comply with the task instructions—often including comments or unnecessary explanations in the output. To balance quality and response time, a medium-sized model, *deephermes-3-mistral-24b*, was chosen, which was available for free during its preview phase. The prompt used was carefully designed to obtain clear, concise rewritings suited to a B1 language level. The original prompt, written in Italian, is shown below, followed by its English translation for clarity:

*Riscrivi il seguente brano tratto dai Promessi Sposi in un italiano moderno e semplice. Mantieni tutte le informazioni e ogni singola frase, ma usa parole facili e frasi brevi, come se parlassi con un amico. Non saltare o riassumere nulla. Italiano livello B1: frasi chiare, niente parole difficili, niente grammatica complicata. Nessun commento, solo il testo riscritto. Semplifica il lessico a un livello B1. Solo il testo riscritto in output. Nessuna maiuscola fuori posto!*

The English translation of the prompt is as follows:

*Rewrite the following excerpt from The Betrothed in modern and simple Italian. Keep all information and every single sentence, but use easy words and short sentences, as if you were talking to a friend. Do not skip or summarize anything. Italian level B1: clear sentences, no difficult words, no complicated grammar. No comments, only the rewritten text. Simplify the vocabulary to B1 level. Only the rewritten text in the output. No misplaced capital letters!*

The resulting dataset, made up of pairs of Manzonian and modern sentences, underwent a revision and cleaning phase to ensure stylistic consistency and linguistic quality before being used for fine-tuning the *Manzoni AI* model.

The dataset was dynamically built in `.csv` format, progressively adding the converted sentences. Each CSV file contains corresponding sentence pairs, with one column for the modern text and one for the Manzonian text.

Below we report some representative examples of phrase pairs taken from various parts of the novel:

#### Example 1 - Chapter I

- **Modern:** Il ramo del lago di Como che si volge a sud, tra due catene di montagne continue, ha una costa sinuosa con seni e golfi a seconda dell’andamento delle montagne. Improvvisamente, la sua larghezza si restringe e assume la forma di un fiume, tra un promontorio a destra e una costa ampia dall’altra parte. Il ponte che unisce le due rive evidenzia questa trasformazione e segna il punto in cui il lago finisce e l’Adda riprende il suo corso, per poi tornare a chiamarsi lago quando le rive si allontanano di nuovo, permettendo all’acqua di espandersi e rallentarsi in nuovi seni e golfi.

- **Manzonian:** Quel ramo del lago di Como , che volge a mezzogiorno, tra due catene non interrotte di monti, tutto a seni e a golfi, a seconda dello sporgere e del rientrare di quelli, vien, quasi a un tratto, a ristringersi, e a prender corso e figura di fiume, tra un promontorio a destra, e un'ampia costiera dall'altra parte; e il ponte, che ivi congiunge le due rive, par che renda ancor più sensibile all'occhio questa trasformazione, e segni il punto in cui il lago cessa, e l'Adda ricomincia, per ripigliar poi nome di lago dove le rive, allontanandosi di nuovo, lascian l'acqua distendersi e rallentarsi in nuovi golfi e in nuovi seni.

### Example 2 - Chapter I

- **Modern:** La costa, formata dal deposito di tre torrenti importanti, scende appoggiandosi a due monti vicini: uno si chiama di San Martino e l'altro, che ha un nome tipico lombardo, si chiama Resegone. Quest'ultimo ha molti picchi allineati, che lo rendono simile a una sega. Chi lo vede per la prima volta, magari dalle mura di Milano che guardano a nord, lo riconosce subito grazie al suo aspetto particolare. In mezzo alla lunga e ampia catena montuosa, lo si distingue facilmente dai monti con un nome meno noto e una forma più comune.
- **Manzonian:** La costiera, formata dal deposito di tre grossi torrenti, scende appoggiata a due monti contigui, l'uno detto di san Martino, l'altro, con voce lombarda, il Resegone, dai molti suoi cocuzzoli in fila, che in vero lo fanno somigliare a una sega: talché non è chi, al primo vederlo, purché sia di fronte, come per esempio di su le mura di Milano che guardano a settentrione, non lo discerna tosto, a un tal contrassegno, in quella lunga e vasta giogaia, dagli altri monti di nome più oscuro e di forma più comune.

### Example 3 – Chapter VIII

- **Modern:** Addio, montagne che sorgono dall'acqua e si innalzano verso il cielo. Cime diverse, familiari a chi è cresciuto tra voi, e impresse nella sua mente, proprio come l'aspetto delle persone che gli stanno vicino. Fiumi e torrenti, i cui suoni sono distinti, come le voci di casa. Case sparse e bianche che sembrano pecore al pascolo lungo il pendio. Addio! È davvero triste lasciare un posto dove si è cresciuti. Anche se qualcuno parte volontariamente, sperando di far fortuna altrove, in quel momento i sogni di ricchezza sembrano meno attraenti. Si meravigliano di aver preso una decisione simile e vorrebbero tornare indietro, se non pensassero che un giorno torneranno come ricchi.
- **Manzonian:** Addio, monti sorgenti dall'acque, ed elevati al cielo; cime inuguali, note a chi è cresciuto tra voi, e impresse nella sua mente, non meno che lo sia l'aspetto de' suoi più familiari; torrenti, de' quali distingue lo scroscio, come il suono delle voci domestiche; ville sparse e biancheggianti sul pendio, come branchi di pecore pascenti; addio! Quanto è tristo il passo di chi, cresciuto tra voi, se ne allontana! Alla fantasia di

quello stesso che se ne parte volontariamente, tratto dalla speranza di fare altrove fortuna, si disabbelliscono, in quel momento, i sogni della ricchezza; egli si meraviglia d'essersi potuto risolvere, e tornerebbe allora indietro, se non pensasse che, un giorno, tornerà dovizioso.

Unfortunately, the chosen model may have made some inaccuracies in the formulation of certain sentences, and we cannot assert with certainty that the entire dataset has been thoroughly refined. The only control performed concerned the identification and removal of so-called *hallucinations*, that is, cases in which the model produced entirely fabricated content or expressed itself in languages other than Italian—sometimes even in a mixture of idioms, including oriental languages. Additionally, we had to clean up some residual occurrences of markers such as `<to_call>` that remained in a few transformations.

As a result, the creation of the dataset was *semi-automatic*, since it required a validation phase—albeit not in-depth—to ensure a minimal quality standard and overall consistency of the linguistic material.

The dataset thus constructed represents a fundamental and valuable resource for training models capable of performing accurate stylistic transformations between contemporary Italian and the Manzonian register. Although the process relied on automatic techniques, it was carefully monitored to ensure the quality and coherence of the data, providing a solid foundation for the fine-tuning of the *Manzoni AI* model and for future research in the field of historical language modeling.

## IV. TRAINING

Training a large language model (LLM) from scratch, i.e., initializing all weights randomly, necessitates computational resources that vastly exceed the scope of the present work. Furthermore, such an approach is markedly inefficient as it neglects the considerable advantages afforded by leveraging pre-trained models, wherein extensive linguistic knowledge is transferred, thereby significantly diminishing both computational demands and convergence time.

All experimental activities were carried out utilizing the free tier of Google Colab, which provides constrained hardware resources, detailed as follows:

- 12.7 GB of system RAM,
- 15 GB of GPU memory (NVIDIA Tesla T4),
- 112.6 GB of available disk storage.

Within these limitations, full fine-tuning—defined as updating all model parameters—was impractical. Consequently, a computationally efficient method known as **Low-Rank Adaptation** (LoRA), introduced by Hu et al. (2021)<sup>1</sup>, was adopted. This technique selectively updates a small subset of parameters by injecting low-rank decomposition matrices into existing weight matrices, while keeping the primary model architecture fixed.

<sup>1</sup>Hu, E., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, L., & Chen, W. (2021). *LoRA: Low-Rank Adaptation of Large Language Models*. <https://arxiv.org/abs/2106.09685>

More formally, given a linear transformation represented by  $W \in \mathbb{R}^{d \times k}$ , the modified transformation is expressed as:

$$Wx \rightarrow Wx + \alpha \cdot (AB)x$$

where  $A \in \mathbb{R}^{d \times r}$ ,  $B \in \mathbb{R}^{r \times k}$ , and  $r \ll d, k$ . Only the matrices  $A$  and  $B$  are updated during fine-tuning, while the original parameter matrix  $W$  remains frozen. This significantly reduces memory consumption and computational overhead, thereby enabling fine-tuning under limited hardware constraints.

Fine-tuning was implemented using the `Unsloth` library, specifically developed to optimize LoRA-based training and facilitate efficient adaptation on consumer-grade GPUs. The library supports a range of open-weight models, including:

- Gemma (1B, 4B, 12B, 27B),
- LLaMA 3 (3B, 8B, 70B),
- Phi,
- Mistral 7B.

Several of these models are also available in 4-bit quantized variants, which substantially reduce GPU memory usage at the expense of generation quality. Preliminary evaluations on quantized models—such as LLaMA 3 8B—revealed notable degradation in semantic coherence and stylistic consistency.

Accordingly, the final model selected was Gemma 3 4B IT in its full precision, non-quantized form. This choice balanced linguistic performance, support for the Italian language, and compatibility with the computational limitations imposed by the Colab environment, rendering it optimal for this study.

The training procedure was executed on the NVIDIA Tesla T4 GPU provisioned by Google Colab, which offered approximately 14.7 GB of VRAM, with an active memory footprint of roughly 9.5 GB during fine-tuning. To accommodate this memory ceiling, hyperparameters were carefully tuned to optimize resource utilization while preserving stable training dynamics.

An **effective batch size of 8** was attained through the combination of:

- `per_device_train_batch_size = 2`, indicating two samples per device per training step;
- `gradient_accumulation_steps = 4`, which accumulates gradients across four steps before executing a parameter update.

This approach emulates a larger batch size without exceeding available memory.

Training was capped at `max_steps = 250`, a quantity deemed sufficient for preliminary assessment of fine-tuning efficacy within constrained computational settings. Although epoch-based termination criteria could have been employed, a step-based cutoff was preferred to grant finer control over training duration.

The learning rate was fixed at  $2e-4$ , a relatively elevated value suitable for short fine-tuning sessions; lower rates (e.g.,  $2e-5$ ) are advisable for extended training to ensure smoother convergence. An initial **warm-up period spanning 5 steps** was introduced to mitigate instability during early optimization phases.

The optimizer employed was `adamw_8bit`, a memory-efficient variant leveraging 8-bit quantization, ideal for environments with limited memory capacity. A **weight decay** coefficient of  $0.01$  was applied to regularize parameter updates and reduce overfitting risk.

A linear learning rate scheduler was utilized, gradually decaying the learning rate throughout training.

For reproducibility, a fixed random seed (3507) was set. Training logs were generated at every step (`logging_steps = 1`), while external logging integrations (e.g., Weights & Biases) were disabled by configuring `report_to = "none"`.

Data loading and preprocessing were parallelized using 2 processes (`dataset_num_proc = 2`) to accelerate tokenization and dataset preparation.

Figure 1 depicts the training loss trajectory over the course of fine-tuning.

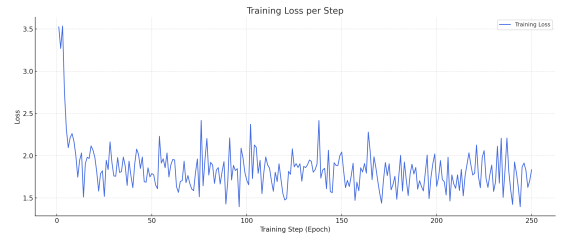


Fig. 1. Training loss progression during the fine-tuning process.

The plot illustrates the **training loss** as a function of training steps (up to 250). A general downward trend is evident, signifying progressive learning and adaptation to the provided dataset. Although fluctuations are observed—consistent with the use of small batch sizes and restricted computational resources—the overall pattern indicates enhanced predictive capability. The relative stabilization of the loss near the conclusion of training suggests an appropriate equilibrium between effective learning and overfitting, given the hardware constraints.

## V. RESULTS

We analyze how the model responds in different contexts and under various configurations.

By *configurations*, we refer to the parameters that influence the generative behavior of the model, specifically **temperature** and **top\_p**, which affect the randomness and diversity of the generated responses.

- **Temperature:** this is a positive real value, typically ranging between 0 and 1 (although it can technically exceed 1). This parameter controls the “creativity” of the model.
  - Low values (close to 0) tend to produce very conservative and repetitive responses, selecting the most probable words.
  - Higher values (up to 1 or beyond) increase the variety and unpredictability, allowing for more creative but potentially less coherent outputs.

- **Top\_p** (nucleus sampling): this parameter ranges between 0 and 1 and limits the selection of candidate words to those whose cumulative probability mass reaches the threshold  $p$ .
  - Low values of top\_p restrict the choice to a small subset of highly probable words, resulting in safer and more linear text.
  - Values close to 1 allow a wider and more diverse selection, increasing the creativity and novelty of the generated text.

By adjusting these parameters, one can control the balance between coherence and creativity in the model’s output, tailoring it to the specific needs of the context or the desired type of output.

#### A. *Manzonian-like context, temperature and top\_p standard*

By “standard” we mean the model’s default configuration chosen for this study, with a `temperature` of 0.85 and a `top_p` of 0.82, which balances creativity and coherence in the generated text.

##### Example 1.1

- **User input**

Anche se tutti speravano che le cose cambiassero, ogni giorno sembrava uguale al precedente. La gente camminava per le strade con il volto stanco, parlava poco e guardava lontano come se aspettasse qualcosa. I più anziani dicevano che bisognava avere pazienza, ma molti iniziavano a pensare che nessuno sarebbe mai venuto ad aiutarli davvero.

- **MANZONI AI**

Ognuno, per dir la verità, sognava che le cose potessero cambiare; ma ogni giorno era, pareva a tutti, uguale al giorno che gli precedeva. Passeggiavano per le strade, con il volto stanco, a poca parola, e con lo sguardo lontano, come se aspettassero qualcosa. I più vecchi, che in generale hanno la pazienza, sognavano che la pazienza era la cosa più bella; ma molti cominciarono a credere che nessuno sarebbe mai venuto a prender loro un po’ di braccio.

##### Example 1.2

- **User input**

Lei rimaneva spesso sola nella stanza, guardando fuori dalla finestra senza dire una parola. Il vento muoveva le tende e sembrava portare via con sé i suoi pensieri. Ogni tanto chiudeva gli occhi, come se volesse ricordare un tempo più felice. Ma quel tempo sembrava ormai lontano, e nessuno sapeva se sarebbe mai tornato.

- **MANZONI AI**

Lei si trovava spesso sola in quella stanza, a contemplare, senza parole, il paesaggio esteriore; e il vento agitava le tende, e sembrava che portasse via con sé i suoi pensieri. Ogni tanto chiudeva gli occhi, come per cercare di ricreare un tempo più felice; ma quel tempo pareva ormai lontano, e nessuno sapeva se sarebbe tornato.

##### Example 1.3

- **User input**

Il paese sembrava tranquillo, ma tutti sapevano che c’era qualcosa che non andava. La gente parlava a bassa voce, e molti evitavano di uscire la sera. Nessuno diceva apertamente cosa temeva, ma lo si capiva dagli sguardi. Solo il prete continuava a parlare di speranza e giustizia, anche se ormai pochi sembravano crederci davvero.

- **MANZONI AI**

Il paese era un po’ quieto, e tutti sapevan che c’era qualcosa che non andava: la gente parlava sottovoce, e molti, in quella sera, si tenevano in casa; nessuno parlava apertamente del male che sentiva, ma lo si sentiva negli occhi. Solo il curato aveva ancora quella parola di speranza e di giustizia, e l’era stata a dirla, per tutto il tempo che gli era rimasto da vivere.

The results obtained in this series of tests demonstrate the model’s capability to effectively emulate the linguistic style and narrative tone characteristic of Manzonian prose. By employing the chosen standard configuration, the model consistently produces outputs that exhibit a high degree of coherence and stylistic fidelity, successfully capturing the nuanced syntactic structures and lexical choices typical of the source material.

Moreover, the generated texts maintain semantic alignment with the original user inputs, preserving the core meaning while enriching the expression with period-appropriate diction and phrasing. This balance between faithfulness to the input content and stylistic adaptation highlights the potential of the model as a valuable tool for literary style transfer and historical language reconstruction.

Overall, these experiments confirm that the calibrated parameters of `temperature` and `top_p` effectively manage the trade-off between creativity and precision, enabling the generation of text that is both imaginative and contextually appropriate within the Manzonian literary framework.

#### B. *Other examples “off the beaten path” but where the harvest is rich.*

In this series, we present examples that are less conventional compared to the Manzonian context, aimed at testing the model’s ability to adapt to shorter sentences and more colloquial expressions.

##### Example 2.1

- **User input**

Il potere logora chi non ce l’ha.

- **MANZONI AI**

Il potere fa la botta a chi non l’ha..

##### Example 2.2

- **User input**

Cosa vorrei sulla mia tomba? Quando sono nato, quando morirò e basta. Non ha senso copiare le bellissime frasi usate dagli altri, perchè quando uno va a vederle, e vede

tante belle cose, si domanda dove sia il cimitero dei cattivi.

- **MANZONI AI**

Cosa vorrei sulla mia lapide? Quando son nato, quando morirò, e basta. Non ha ragione d'esser copiata quella bella prostituzione delle parole che gli altri hanno usate; perché, al veder tante belle cose, si domanda dov'è la lapide de' birboni.

These examples demonstrate the model's versatility in processing inputs that diverge from the canonical Manzonian domain. Despite the informal tone and concise structure of the prompts, the model successfully adapts its language while preserving stylistic coherence and consistency. This flexibility underscores the robustness of the fine-tuning approach in generating outputs that remain faithful to the intended literary style across diverse linguistic registers and contextual variations.

### C. *Testing the temperature and the top\_p*

These parameters are known to influence the randomness and diversity of the output, thus affecting the balance between creativity and coherence. By systematically varying these values, we analyze how the model adapts its lexical and syntactic choices, and the extent to which the generated text maintains fidelity to the original meaning while reflecting stylistic nuances.

- **User input**

Anche se tutti speravano che le cose cambiassero, ogni giorno sembrava uguale al precedente. La gente camminava per le strade con il volto stanco, parlava poco e guardava lontano come se aspettasse qualcosa. I più anziani dicevano che bisognava avere pazienza, ma molti iniziavano a pensare che nessuno sarebbe mai venuto ad aiutarli davvero.

#### **Example 3.1 — Temperature: 0.73, Top\_p: 0.1**

- **MANZONI AI**

Per quanto tutti sperassero che le cose potessero cambiare, ogni giorno pareva uguale al precedente; passeggiavano per le strade, con il viso stanco, con poche parole, con lo sguardo fisso, come se aspettassero qualcosa; i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a sostenerli davvero.

#### **Example 3.2 — Temperature: 0.54, Top\_p: 0.74**

- **MANZONI AI**

Egli, che tutti speravano che la cosa potesse cambiar, ogni giorno pareva agli occhi di tutti, come il giorno che era venuto prima. La gente passeggiava per le strade, con il volto stanco, con poche parole, con lo sguardo lontano, come se aspettasse qualcosa. I più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a far loro davvero il piacere.

#### **Example 3.3 — Temperature: 0.66, Top\_p: 0.57**

- **MANZONI AI**

Per quanto tutti avessero sperato che le cose potessero cambiare, ogni giorno pareva il giorno di prima: passeggiavano per le strade, con la faccia stanca, senza parlare molto, con lo sguardo fisso, come se aspettassero qualcosa; e i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a sostenerli davvero.

#### **Example 3.4 — Temperature: 0.25, Top\_p: 0.67**

- **MANZONI AI**

Ognuno, tuttavia, sperava che le cose mutassero; ma ogni giorno pareva uguale al precedente: passeggiavano per le strade, con il viso stanco, con poche parole, con lo sguardo lontano, come se aspettassero qualcosa: i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a sostenerli davvero.

#### **Example 3.5 — Temperature: 0.73, Top\_p: 0.57**

- **MANZONI AI**

Per quanto tutti avessero sperato che le cose potessero cambiare, ogni giorno pareva uguale al precedente; la gente passeggiava per le strade, con il viso stanco, con poche parole, con lo sguardo lontano, come se aspettasse qualcosa; i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a dar loro davvero una mano.

#### **Example 3.6 — Temperature: 0.1, Top\_p: 0.75**

- **MANZONI AI**

Per quanto tutti sperassero che le cose potessero cambiare, ogni giorno pareva uguale al precedente; passeggiavano per le strade, con il viso stanco, con poche parole, con lo sguardo fisso, come se aspettassero qualcosa; i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a sostenerli davvero.

#### **Example 3.7 — Temperature: 0.36, Top\_p: 0.18**

- **MANZONI AI**

Per quanto tutti sperassero che le cose potessero cambiare, ogni giorno pareva uguale al precedente; passeggiavano per le strade, con il viso stanco, con poche parole, con lo sguardo fisso, come se aspettassero qualcosa; i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a sostenerli davvero.

#### **Example 3.8 — Temperature: 0.73, Top\_p: 0.16**

- **MANZONI AI**

Per quanto tutti sperassero che le cose potessero cambiare, ogni giorno pareva uguale al precedente; passeggiavano per le strade, con il viso stanco, con poche parole,

con lo sguardo fisso, come se aspettassero qualcosa; i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a sostenerli davvero.

### Example 3.9 — Temperature: 0.18, Top\_p: 0.55

- **MANZONI AI**

Per quanto tutti sperassero che le cose potessero cambiare, ogni giorno pareva uguale al precedente; la gente camminava per le strade, con il viso stanco, con poche parole, con lo sguardo lontano, come se aspettasse qualcosa; i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a sostenerli davvero.

### Example 3.10 — Temperature: 0.1, Top\_p: 0.11

- **MANZONI AI**

Per quanto tutti sperassero che le cose potessero cambiare, ogni giorno pareva uguale al precedente; passeggiavano per le strade, con il viso stanco, con poche parole, con lo sguardo fisso, come se aspettassero qualcosa; i più anziani dicevano che bisognava avere pazienza; ma molti cominciarono a credere che nessuno sarebbe mai venuto a sostenerli davvero.

The results from Examples 3.1 through 3.10 clearly illustrate how variations in the `temperature` and `top_p` parameters affect the lexical and syntactic choices of the model. Specifically, lower temperature values, as seen in Examples 3.4 (Temperature 0.25, Top\_p 0.67), 3.6 (Temperature 0.1, Top\_p 0.75), and 3.10 (Temperature 0.1, Top\_p 0.11), tend to produce more conservative and semantically faithful outputs with limited lexical variation, closely mirroring the input sentence’s original structure and meaning.

Conversely, higher temperature values combined with moderate to low `top_p` values, such as in Examples 3.1 (Temperature 0.73, Top\_p 0.1) and 3.5 (Temperature 0.73, Top\_p 0.57), induce a greater degree of lexical creativity and syntactic variation. These examples exhibit more pronounced rephrasings and stylistic nuances while still preserving the core semantic content. For instance, the use of alternative verbs and sentence structures, such as “dar loro davvero una mano,” shows the model’s ability to creatively adapt the text within the Manzonian style.

Moderately low values for `temperature` and `top_p` yield good yet cautious results, as demonstrated in Example 3.2.

Moreover, the consistent presence of Manzonian stylistic elements across all examples — such as the use of archaic verb forms, punctuation mimicking period literature, and the syntactic rhythm typical of Alessandro Manzoni’s prose — attests to the robustness of the model’s style transfer capabilities regardless of sampling parameter configuration.

The systematic variation of `temperature` and `top_p` provides a powerful mechanism for controlling the trade-off between coherence and creativity in literary style emulation.

By selecting appropriate parameter settings, users can tailor the output to prioritize either semantic precision or stylistic inventiveness, as demonstrated by the diverse examples above. This flexibility is crucial for practical applications that require both faithful content reproduction and nuanced literary adaptation.

### D. Inspecting the behaviour of the model in other languages

In order to assess the versatility and robustness of the model’s style transfer capabilities beyond Italian, we examined its output when prompted with sentences in English and French. This cross-linguistic inspection allows us to evaluate whether the Manzonian stylistic imprint is preserved across different source languages, and to what extent the model adapts syntactic and lexical structures accordingly.

#### Example 3.1 - ENGLISH

- **User prompt**

The village was quiet after the storm, with broken branches in the streets and people looking out from their windows in silence.

- **MANZONI AI**

The village was quite still, with the broken boughs in the road, and the inhabitants peering, in silence, from their windows.

#### Example 3.2 - ENGLISH

- **User prompt**

She waited every morning by the same window, hoping to see a familiar face pass by, but the road remained empty.

- **MANZONI AI**

And every morn, she stood by the window, and, with a weary hope, watched for the passing of some well-known face; but the road was empty.

#### Example 3.3 - FRENCH

- **User prompt**

Le village restait silencieux après la tempête, avec des branches cassées partout et des visages inquiets derrière les fenêtres.

- **MANZONI AI**

Le village resterait tranquille sur la vague, et resterait, sur les branches brisées, et sur les visages inquiets derrière les fenêtres.

These examples demonstrate the model’s capacity to transfer the Manzonian style even when the input is provided in languages other than Italian. In the English outputs, the model reproduces characteristic features of Italian-like Alessandro Manzoni’s prose, such as archaic or poetic lexical choices (e.g., “boughs” instead of “branches,” “every morn” instead of “every morning”) and syntactic structures that evoke a literary rhythm reminiscent of 19th-century Italian literature translated into English. Similarly, the French output, while maintaining the source language’s syntactic framework, shows an attempt to mirror Manzoni’s stylistic tone through repetition and a measured cadence.

However, subtle deviations and occasional unnatural constructions (e.g., “resterait tranquille sur la vague”) indicate that the model’s stylistic transfer is less precise outside of Italian, likely due to differences in language structure and the training data’s language distribution, which is exclusively Italian.

These cross-linguistic tests confirm the model’s inherent stylistic robustness, with a clear Manzonian imprint that transcends language boundaries to some extent. This capacity, while strongest in Italian, suggests promising directions for future work on multilingual literary style transfer, emphasizing the importance of language-specific fine-tuning to fully capture the nuances of style across diverse linguistic contexts.

### E. Task focus test

In order to evaluate whether the Manzoni AI model maintains coherence with the task for which it was trained — namely, the stylistic transformation of contemporary Italian into a Manzonian register — we conducted a series of tests based on operational and pragmatic prompts. These inputs, typically unrelated to literary contexts, are designed to assess the model’s ability to stay within its stylistic boundaries and avoid drifting into general-purpose language generation.

Below are some representative examples:

#### Example 4.1

- **User prompt**

Immagina di dover saltare su un masso di pietra alto 3 metri. Cosa faresti?

- **Manzoni AI**

Saltare sur un masso di tre metri: qual sarebbe il mio metodo? A saltar le piante sul tronco, oppure a saltar le piante sul tronco, come facevano i marinai in porto, col piede che stende il gambone, e l’altro che gira il piede, come si fa con la corda di traverso.

#### Example 4.2

- **User prompt**

Devo scrivere una lettera di scuse per la mia professoressa perchè le ho detto che il cane ha mangiato i compiti quando invece non li avevo proprio fatti. Aiutami.

- **Manzoni AI**

Digami una lettera di scuse alla professoressa, per averle detto che il cane aveva mangiato i compiti, quando io non li aveva affatto fatti.

The model’s responses in these scenarios clearly demonstrate that the fine-tuning process successfully constrains its generative behavior to a specific stylistic register. When presented with prompts that are less explicitly literary and instead involve direct or pragmatic requests, the model still attempts to respond to the underlying intent of the input. However, it does so in a limited and stylistically consistent manner.

Specifically, two key behaviors emerge:

- The responses remain partial and non-exhaustive with respect to the task, often avoiding full or practical solutions.

- Even in these atypical contexts, the model consistently adheres to Manzonian stylistic conventions — employing archaic vocabulary, complex syntax, and rhetorical devices.

This behavior is particularly noteworthy, as it suggests that the model does not simply apply a surface-level stylistic filter, but internalizes a broader stylistic identity that guides its outputs even outside canonical literary contexts. Such findings open the possibility for future research and development: with more advanced techniques such as reinforcement learning and the use of larger models, one could envision a version of *Manzoni AI* capable of fully engaging in open-ended dialogue — not only rewriting in Manzonian style, but behaving as a conversational agent that consistently *thinks and speaks* in the literary voice of Alessandro Manzoni. This would parallel the experience of interacting with a general-purpose LLM like ChatGPT, but with a stylistic and rhetorical identity shaped by a specific literary tradition.

### F. Literary Whims: Manzoni AI Rewrites Some Italian Masterpieces

To explore the ability of the Manzoni AI model to reinterpret famous Italian literary texts through the stylistic lens of Manzoni, we submitted to the model some excerpts from well-known works, requesting a rewriting that preserved the original meaning while adapting the form to the distinctive language of Alessandro Manzoni. Below are three representative examples.

#### Example 5.1 - Incipit of Dante Alighieri’s *Commedia*

- **User prompt**

A metà della mia vita, mi trovai smarrito in una foresta oscura e inquietante, perché avevo perso la strada giusta. Era un luogo difficile da attraversare, così pieno di pericoli e confusione che non sapevo dove andare. Volevo tornare indietro, ma la paura e il senso di smarrimento mi bloccavano. Mentre cercavo una via d’uscita, vidi salire verso di me una collina illuminata dal sole, e pensai che quella potesse essere la strada giusta per salvarmi.

- **Manzoni AI**

A metà della sua vita, si trovò smarrito in una foresta oscura e inquietante, per aver intrapreso la strada sbagliata; così densa e percorsa di pericolosità e di confusione, che non sapeva dov’andare; e quando si volesse andar indietro, era pietrificato dalla paura e dalla confusa contemplazione della sua condizione. Sulle pendici, e in un movimento ben disposto verso di lui, vide salire una collina che gli parve quella per cui lui doveva andare; e pensò che era la cosa giusta a salvarlo.

#### Example 5.2 - Incipit of Giovanni Verga’s *Rosso Malpelo*

- **User prompt**

Lo chiamavano Rosso Malpelo perché aveva i capelli rossi, e tutti sapevano che chi ha i capelli rossi è cattivo di natura. Per questo, anche se non aveva ancora fatto nulla di male, lo trattavano già con diffidenza e durezza.

Lo prendevano a calci, gli davano le botte, e nessuno si meravigliava se si difendeva a sua volta con cattiveria. Anzi, dicevano che “il male lo portava dentro, nel sangue”. Perfino sua madre non osava difenderlo, quasi si vergognava di lui. Al massimo diceva che era “un po’ strano”, e lasciava che ognuno lo trattasse come voleva.

- **Manzoni AI**

Lo chiamavano Rosso Malpelo, perché i capelli gli eran rossi; e non c’è gente che non sappia che il rosso è segno di cattiveria. Così già, per averlo visto nascere, e per non averlo visto di spasso, gli si faceva male, e gli si dava prima che non fosse ancora addentato nell’affare. Gli picchiavano, gli stropicciavano, e nessuno si meravigliava se, a sua risposta, si difendesse in malizia. Del male che gli veniva dentro, ch’era nel sangue, non c’era che un non so che che lo rendesse così. Sua madre pure non osava, e già si vergognava a prendersi per lui. O falle un po’ di strano, diceva; e lasciava che ognuno fosse lì a lui.

### Example 5.3 - Giacomo Leopardi’s *L’infinito*

- **User prompt**

C’è una siepe, sopra questo colle solitario, che mi impedisce di vedere una parte dell’orizzonte più lontano. Ma proprio perché non vedo, posso immaginare con la mia mente spazi infiniti oltre la siepe, silenzi profondi e pace assoluta, al punto che quasi mi viene da tremare: mi sembra di perdermi in quell’immensità che non riesco a vedere con gli occhi, ma che sento dentro di me. E, mentre ascolto il fruscio del vento tra le foglie, lo paragono a un altro suono che immagino: il silenzio infinito del tempo passato e del tempo futuro. In quel confronto tra il presente che sento e l’eternità che immagino, la mia mente si perde. Ma proprio questo smarrimento mi è dolce, perché mi fa sentire parte di qualcosa di immenso e senza confini.

- **Manzoni AI**

L’esistenza d’una siepe, su questo colle solitario, che mi priva di una parte dell’orizzonte più lontano, non impedisce che, nel mio pensiero, io possa immaginare desolati spazi infiniti, silenzi infiniti, solenni, come per dirlo: tanto che mi vien quasi la palpazione, di perdermi in quell’immensità che non posso contemplare con gli occhi, ma che sento in me. E, mentre ascolto il fruscio del vento tra le foglie, lo prendo per un altro suono immaginato: il silenzio della grande storia, la grande storia. E, nel confronto tra il presente che sento e l’eternità che sogno, il mio pensiero s’adombra. Ma proprio questo sdoppiarsi è più dolce, perché mi fa sentire parte di quell’immensità inesprensabile.

The analysis of these examples shows that Manzoni AI is capable of re-adapting contemporary or well-known literary texts while preserving their original meaning, yet transforming them into a language that is stylistically Manzonian. This is achieved through the use of archaic expressions, elaborate syn-

tactic constructions, and a rhythm evocative of 19th-century prose.

This creative exercise — a playful yet intellectually stimulating endeavor — allows for a reimagining of foundational literary texts, whose significance extends beyond style to encompass the cultural and philosophical identity of the Italian tradition. By recasting these canonical works in a stylistic register foreign to their original form, Manzoni AI offers a unique perspective: it stages hypothetical encounters between literary voices across time, producing pseudo-historical reinterpretations that, while artificial, invite renewed engagement with the texts.

Beyond its exploratory nature, this stylistic transformation opens a variety of potential applications, enabling new forms of literary analysis, comparative studies, pedagogical tools, and creative experimentation. To the best of our knowledge, such cross-stylistic rewriting — driven by fine-tuned language models — represents a novel direction in computational literature, offering interpretative angles previously inaccessible through traditional methods.

## VI. FINAL CONSIDERATIONS

Manzoni AI demonstrates that, at the current state of technology, even individuals with limited resources can build their own customized language model. More than that, it shows that the fruits of technological progress are not always geared toward the future or a greater rate of innovation, but can also be used to foster a deeper awareness of human reality and its history. The opportunity offered by Manzoni AI to, so to speak, resurrect the refined pen of Manzoni is likely of no “materialistic” interest, but it sheds light on the possibility of a use of technology that is not frantic, but rather pleasurable—almost hedonistically detached from practical concerns. The countless potential applications and developments of such technology are promising not only in the educational sphere but also for communities of passionate readers and connoisseurs of art. The ability to absorb the beautiful art of refined expression—now taught to a machine—is, to the author, a source of great personal satisfaction. Ad maiora.

Rome, June 4, 2025

*Matteo Bernardi*