

S-VoCAL: A Dataset and Evaluation Framework for Inferring Speaking Voice Character Attributes in Literature

Abigail Berthe-Pardo¹, Gaspard Michel^{1,2}, Elena V. Epure^{2,3}, Christophe Cerisara¹

¹LORIA, Vandœuvre-lès-Nancy, France

²Deezer Research, Paris, France

³Idiap Research Institute, Switzerland

abigail.berthe@hotmail.fr, gmichel@deezer.com, elena.epure@idiap.ch, christophe.cerisara@loria.fr

Abstract

With recent advances in Text-to-Speech (TTS) systems, synthetic audiobook narration has seen increased interest, reaching unprecedented levels of naturalness. However, larger gaps remain in synthetic narration systems' ability to impersonate fictional characters, and convey complex emotions or prosody. A promising direction to enhance character identification is the assignment of plausible voices to each fictional character in a book. This step typically requires complex inference of attributes in book-length contexts, such as a character's age, gender, origin or physical health, which in turns requires dedicated benchmark datasets to evaluate extraction systems' performances. We present S-VoCAL (Speaking Voice Character Attributes in Literature), the first dataset and evaluation framework dedicated to evaluate the inference of voice-related fictional character attributes. S-VoCAL entails 8 attributes grounded in sociophonetic studies, and 952 character-book pairs derived from Project Gutenberg. Its evaluation framework addresses the particularities of each attribute, and includes a novel similarity metric based on recent Large Language Models embeddings. We demonstrate the applicability of S-VoCAL by applying a simple Retrieval-Augmented Generation (RAG) pipeline to the task of inferring character attributes. Our results suggest that the RAG pipeline reliably infers attributes such as Age or Gender, but struggles on others such as Origin or Physical Health. The dataset and evaluation code are available at <https://github.com/AbigailBerthe/S-VoCAL>.

Keywords: Digital Humanities, Corpus, Evaluation Methodologies

1. Introduction

Audiobooks have recently attracted growing worldwide interest: the Audiobook industry in the United-States, Europe and China revenue is projected to increase by 20% by the end of 2025 (Statista, 2025a,b,c). Among reasons for success, the ability of Audiobooks to promote literature to individuals that usually do not or cannot engage with the act of reading, such as individuals with reading or visual impairments along with the ability to perform other actions while listening is often mentioned (Roy, 2024; Creamer, 2025). Meanwhile, advances in text-to-speech (TTS) systems' ability to synthesize natural and expressive speech have sparked interest in the synthetic reading of Audiobooks (Walsh et al., 2023; Park et al., 2025), allowing consumers to listen to personalized artificial narrators reading their favorite e-book, and authors to independently promote their work with lower costs than traditional Audiobook production.

A key aspect in performative audiobook narration is the narrator's ability to impersonate fictional characters (Browm, 2021; Goodwin, 2024). For example, narrators typically employ lower pitch when reading utterances from male character or from narrative parts (Pethe et al., 2025). For listeners, this aspect of narration is essential to enhance character identification and reading comprehension (Van Krieken et al., 2017; Wolters et al., 2022) and has been identified as a critical limitation of

current synthetic Audiobooks (Rodero and Lucas, 2023). Besides, modern TTS systems typically suffer from lack of expressivity when synthesizing fictional character utterances (Michel et al., 2025), positioning the enhancement of character impersonation in artificial readers as a central objective for achieving broader advancements in Audiobook narration synthesis.

The alignment between plausible voices and fictional character descriptions yields a promising direction that has been recently explored in the TTS community (Xin et al., 2022; Park et al., 2025). However, finding these character-appropriate voices requires complex Natural Language Processing (NLP) steps to extract voice-specific attributes such as age, gender, region of origin, or physical condition for each character. This information might not always be explicitly stated in books and might be disseminated in distant narrative sections, making the precise extraction of such attributes challenging (Iosif and Mishra, 2014; Papoudakis et al., 2024). To evaluate the ability of NLP systems to accurately extract character information, prior dataset efforts have mostly proposed general descriptions of characters in natural language drawn from literature analysis websites (Yuan et al., 2024; Gurung and Lapata, 2024; Papoudakis et al., 2024). Such descriptions are usually centered around structured profiles or complex literary analysis, involve various side-information that do not correlate with a character's voice or might critically miss important

voice-related attributes.

Hence, we propose S-VoCAL (Speaking Voice Character Attributes in Literature), a dataset and evaluation framework dedicated to the extraction of fictional character attributes that correlate with speaking voice delivery. S-VoCAL entails 8 attributes grounded in phonetic studies, and 952 characters-book pairs distributed among 192 unique novels collected from Project Gutenberg¹.

Attributes are initially extracted by matching characters with their Wikidata² entry, from which the *Age* attribute was complemented with manual annotations. Given the unique and unbalanced nature of each attributes (e.g. *Gender* is categorical ; *Physical Health* is a short text ; the *Type* attribute mostly contains *Humans*), we design a dedicated evaluation framework including weighted metrics, and a new semantic similarity metric based on Qwen-3 8b embeddings (Zhang et al., 2025) that we show correlates strongly with human judgments.

Following Papoudakis et al. (2024), we evaluate a simple Retrieval-Augmented-Generation (RAG) approach on the extraction of S-VoCAL character attributes. The solution involves a Passage Retrieving step based on embeddings derived from instruction-based encoder and an Attribute Inference step performed by Large Language Models (LLMs). We show that the proposed RAG approach typically performs well at extracting closed-class attributes such as *Age*, *Gender* or *Spoken Languages*, but struggles for open-class attributes such as *Physical Health* and *Occupation*.

To summarize, our contributions are as follows:

- We propose S-VoCAL, a dataset and evaluation framework for inferring speaker voice character attributes in novels. The data is publicly available and provides 8 attributes for 952 character-book pairs, over 192 unique books.
- We design an evaluation framework that accounts for the unique and unbalanced nature of each attributes.
- We evaluate a RAG solution for the Attribute Inference task, showing that while it performs well for closed-class attributes, it typically struggle when inferring open-class attributes.

2. Related Works

2.1. Voice Characteristics

Voices can be described through a combination of segmental and suprasegmental features (Nolan, 1987; Di Cristo, 2013). The latter correspond to acoustic features such as pitch, loudness, rhythm,

that can be rendered through various lexical items in literary texts (Parret, 2002). Besides, voices also convey indexical cues that can be used to infer characteristics of the speaker (Laver and Trudgill, 1979; Abercrombie, 1967; Frühholz and Belin, 2018). We find, among features that have been well studied, the speaker gender and age (Nolan, 1987; Hollien and Shipp, 1972a; Sorokowski et al., 2024), physical health (Canter, 1963b) or region of origin and social class (Labov, 1973).

These prior works allow us to identify 8 salient acoustic features that correlate with voice characteristics and that are likely to be identifiable from character descriptions in novels. We describe these attributes in Section 4.

2.2. Character Attributes Datasets

Extracting reliable descriptions of characters from narrative texts is a complex task, as relevant information may be implicit, evolve over time, or distributed over several chapters. Prior works have targeted specific character aspects, such as social roles (Stammach et al., 2022; Gervás and Méndez, 2024), latent personas (Bamman et al., 2014), and more recently, natural language descriptions (Brahman et al., 2021; Yuan et al., 2024; Gurung and Lapata, 2024; Papoudakis et al., 2024). Such descriptions are either general summary of characters, literary analysis, or centered around targeted dimensions (personality, events, ...). In contrast, we define a discrete set of character attributes that correlate with voice delivery, and aim at extracting exactly these attributes rather than high-level character descriptions.

Several works also aim at extracting discrete set of static attributes, such as gender or age (Iosif and Mishra, 2014; Jaipersaud et al., 2024; Baruah and Narayanan, 2024), or dynamic attributes such as emotion arcs (Vishnubhotla et al., 2024). Similar to S-VoCAL, Iosif and Mishra (2014) targets voice-related attributes for storytelling applications, and propose a pipeline to automatically extract these attributes. S-VoCAL extends this approach by covering novels rather than children stories, and by entailing a broader set of attributes derived from sociophonetic studies.

2.3. Evaluating Attribute Inference

Evaluating a system’s ability to infer character attributes is particularly challenging given the unique and unbalanced nature of attributes such as age or physical health. When attributes are described in natural language, such as in Papoudakis et al. (2024), the evaluation becomes even more challenging, as it requires additional reasoning to ensure factual overlaps between the predicted and ground-truth descriptions. Therefore, evaluations

¹<https://www.gutenberg.org/>

²<https://www.wikidata.org/>

Attribute	Voice relevance	Wikidata properties	Type	Coverage	Is New
Name		<i>label</i>	Text	100%	✗
Aliases		<i>also known as, given name, native label, nickname, pseudonym</i>	List	73.5%	✗
Age*	pitch range, formant frequencies, speech rate	<i>relative age</i>	Categorical (4)	4.1% 37.7%	✗
Gender	pitch range, resonance	<i>sex or gender</i>	Categorical (2)	96.5%	✗
Origin	accent, phonetic features, prosody	<i>country of citizenship, place of birth, country of origin</i>	List	58.8%	✓
Residence	accent, speech patterns	<i>residence, work location, country, place of residence</i>	List	26.5%	✓
Occupation	characteristic speech styles	<i>occupation, character type, position held, noble title, military/police rank, crew membership</i>	List	57.7%	✓
Spoken languages	pronunciation, intonation	<i>languages spoken, native language</i>	List	49.2%	✓
Physical health	specific voice qualities (e.g., breathiness)	<i>medical condition</i>	Text	4.1%	✓
Type	human/non-human voice	<i>instance of</i>	Text	100%	✓

Table 1: Attributes overview in S-VoCAL, along with their mapped Wikidata properties. Age* is represented as a relative category inferred from voice features, with coverage reported as (initial | final, after manual annotation). The *Is New* column shows originality compared to (Iosif and Mishra, 2014).

often include fact-based metrics between predicted and reference descriptions, derived from Large Language Models (LLMs) judgments such as PRISMA (Mahon and Lapata, 2024) or Consistency Score (Yuan et al., 2024). For discrete set of attributes, evaluation approaches also include LLM judgments (Jaipersaud et al., 2024), human evaluations (Baruah and Narayanan, 2024), and standard metrics such as F1-score (Iosif and Mishra, 2014).

Given the unique nature of the selected attributes, we also adopt standard strict and soft F1-metrics for closed and semi-closed class attributes (attributes that have a finite set of possible values). For open-class text attributes, we leverage instruction-aware LLM embeddings with Qwen-3 8b (Zhang et al., 2025) to calculate semantic similarities between the predicted and ground-truth attribute, similar to BERTScore (Zhang et al., 2019). Instruction-aware embedding models allows us to derive attribute-specific representations, which can adapt to each attribute particularity. We show that the proposed similarity metric correlates better with human judgments than BERTScore.

3. The Attribute Inference Task

We define a taxonomy of attributes relevant to vocal identity that can be inferred from text, and used to design a general synthetic voice for a fictional

character. We target stable traits, related to the character’s identity, rather than low-level acoustic properties. Voices carry indexical information that enable listeners to infer relatively stable traits, such as age, gender, physical condition, or regional origin (Abercrombie, 1967; Laver and Trudgill, 1991). According to Parret (2002), these cues reflect both *intrinsic* aspects, linked to physiological features, and *extrinsic* aspects, shaped by social and contextual factors. In the context of synthetic book reading, both dimensions are crucial to adapt acoustic parameters of the generated voice, such as pitch range or prosody, to the persona of a fictional character.

These links between a person’s traits and the acoustic realization of their voice have been extensively documented in phonetic and sociophonetic studies. For example, pitch range and fundamental frequency are known to vary with age and gender: male speakers present lower pitch range than females, which in turn exhibit lower pitch range than children (Cruttenden, 1986), and average pitch ranges tends to decrease with age (Hollien and Shipp, 1972b). Health can also impact voice quality and pitch stability (Canter, 1963a), while geographical origin is shown to correlate with segmental realizations (i.e. pronunciation of individual sounds), and accents (Labov, 1973).

We identify eight voice-related attributes that can be inferred from narrative text: *Age, Gender, Ori-*

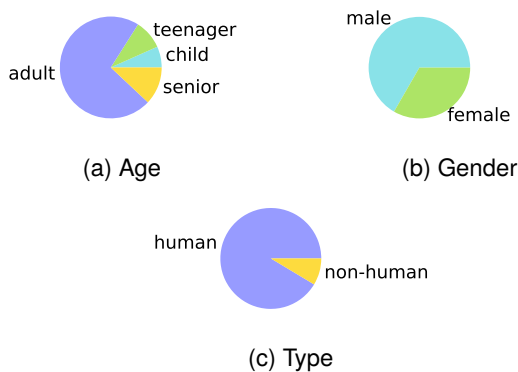


Figure 1: As attributes from S-VoCAL comes from various Gutenberg books, they are unbalanced, especially for *Type* that mostly contains humans and *Age* that mostly contains adults.

gin, *Residence*, *Spoken languages*, *Occupation*, *Physical health*, and *Type*. Their specific influence on vocal characteristics is detailed in Table 1.

The Attribute Inference task aims at extracting these voice-related attributes from long narrative texts: given a character and a novel, we want to infer traits that can influence how that character’s voice would sound. In practice, this corresponds to predicting one or several values for each of the voice-relevant attribute based on the book content only. This task is particularly challenging as literary texts seldom state such information explicitly. Relevant cues are often implicit and distributed across distant passages, requiring reasoning over long context to infer certain attributes indirectly. Accurate reasoning over very long contexts is typically an area where recent NLP models struggle, adding an additional complexity layer to the task of Attribute Inference (Hsieh et al., 2024; Wang et al., 2024a)

To the best of our knowledge, no previous dataset provides explicit labels for speaking voice attributes, nor general evaluation frameworks for evaluating Attribute Inference. Thus, we propose S-VoCAL dataset, a dataset dedicated to the Attribute Inference task, described further.

4. S-VoCAL Dataset

S-VoCAL (Speaking-Voice Character Attributes in Literature) focuses on 8 voice-related attributes defined in Table 1, for 952 characters from 192 novels. These books correspond to works of English or translated world literature published before 1940 and available in full text on Project Gutenberg³. The final selection includes all books that (1) had an English full-text version on Project Gutenberg, (2) had a corresponding entry on Wikidata⁴, and (3)

³<https://www.gutenberg.org/>

⁴<https://www.wikidata.org>

included at least one character on Wikidata (main or secondary), with any of the target properties used to create our attributes (see Table 1). Figure 1 presents the distribution of characters across the main categorical attributes (Age, Gender, and Type), with "adult", "male", and "human" characters largely predominant in the corpus.

4.1. Data Collection

Prior works have mostly relied on character descriptions from literary summary websites or fan-wikis (Papoudakis et al., 2024; Gurung and Lapata, 2024; Yuan et al., 2024). However, these sources vary considerably in the level of detail across books and characters, and exploiting them would have required an additional extraction pipeline with its own validation process. We therefore opted for Wikidata as our main source of character information. Wikidata is a collaborative knowledge graph that provides structured properties about entities, including literary works and fictional characters, filled in by human contributors. Its structured format and frequent human activity make it particularly reliable to align existing properties with our target attributes.

We organized the dataset based on the available characters’ properties in Wikidata. For each character, these properties were heterogeneous and mostly did not match directly our target attributes. We manually selected properties that were semantically very close to the target attribute (either synonyms or narrower subsets) and grouped them under the corresponding predefined attributes. In some cases, multiple complementary values were merged into the same attribute. E.g. *Spoken Languages* aggregates several Wikidata properties, such as *languages spoken*, *native language*, *language from name*, and *writing language*, while the attribute *Gender* is directly mapped to a single Wikidata property, *sex or gender*. Table 1 details the mapping of Wikidata property per attribute.

After collection, the data was normalized by removing duplicates in multilabel attributes, converting similar values into a standardized form, and manually curating properties still present in their Wikidata-codified form.

The character coverage widely varies depending on the attribute, as shown in Table 1. Specifically, the *Age* attribute, which is critical to our task, initially had only 40 instances filled in. For this reason, we proceeded with manually annotating a subset of characters with their *Age* value.

4.2. Age Annotation

Age information within Wikidata is usually provided as a numerical value. However, knowing the exact age of a character is not required to shape its voice,

Age category	F1-score	Instances
Child	0.78	16
Teenager	0.44	10
Adult	0.90	242
Senior	0.71	25

Table 2: Per-class F1-scores of human annotations for the *Age* attribute. *Adult*, the most represented class, shows a high level of agreement, while *Teenager* shows only moderate agreement.

and only approximating an age range typically provides precise enough information. We therefore convert the *Age* attribute into a categorical attribute, consisting in four classes: Child, Teenager, Adult, Senior. The 40 characters whose ages were available as numeric values were automatically mapped into the corresponding class. Then, we manually annotated 319 additional *Age* instances, resulting in a total of 359 annotated cases. Age categories were inferred from websites that provide character analyses (e.g. GradeSaver⁵, LitCharts⁶) We designed guidelines⁷ that prioritize explicit mentions of age, followed by relative age cues, clear references (e.g., "child"), and stage of life indicators. Uncertain cases were left blank.

A total of 293 characters were double-annotated, from which we measured inter-annotator agreement with standard and soft variants of Cohen’s Kappa and F1-score. Unlike other closed-class attributes, *Age* is organized linearly with a natural progression from "Child" to "Senior". Therefore, discrepancies in adjacent classifications (e.g. "teenager" instead of "adult") are less severe than distant ones (e.g. "child" instead of "senior") and could be due to the inherent difficulties in opting between two adjacent categories for certain characters. For this reason, we also report a soft version of both metrics, assigning partial credit when the two annotations differed by only one adjacent age category, with a 0.8 weight.

We report a Cohen’s κ coefficient of 0.56 in its standard form, increasing to $\kappa = 0.74$ when computing the soft variant, and also report a soft F1-score of 0.86. These observations suggest a moderate agreement between annotators, highlighting that most disagreements occur between adjacent classes. Table 2 displays per-class F1-scores for the 293 double-annotated instances, suggesting that "Teenager" is the age category with most disagreement. Annotators finally performed an adjudication step to resolve remaining disagreements and refine the guidelines. After this process, 359

⁵<https://www.gradesaver.com/>

⁶<https://www.litcharts.com/>

⁷Appendix A

characters have a filled *Age* attribute. Age categories distribution in the dataset are reported in Figure 1.

5. S-VoCAL Evaluation Framework

We propose a dedicated evaluation framework tailored to each attribute type and data format (closed, semi-closed, and open). Closed and semi-closed attributes are evaluated using F1-based metrics, while open attributes are evaluated with a similarity metric based on Qwen-3 embeddings, similar to BERTScore, that we validate through correlation with human judgments. These judgments are further used to map similarity values to a human-interpretable scale, allowing a better understanding of models’ performances on open attributes.

5.1. Closed and Semi-Closed Class Evaluation

Closed-class attributes (*Age*, *Gender* and *Type* (*human/non-human*)) take their values into a small set of candidates, thus standard classification metrics can be used normally. Since most characters are of *Type* "human" in S-VoCAL (91%), we conduct a two-stage evaluation for this attribute. The first stage assesses how accurate a system is at distinguishing between human and non-human characters, while the second accounts for semantic similarity between non-human characters. *Type* (*human/non-human*), *Age*, and *Gender* are evaluated using Weighted F1 to account for label imbalance in each categories.

Similarly to the approach taken in the manual annotation of *Age*, we propose a soft F1-score, which penalizes less errors in neighboring categories. The custom soft F1-score uses a weight matrix that encodes linear distances between age groups: 1.0 for exact matches, 0.8 for adjacent categories, and 0.0 otherwise.

The only semi-closed class attribute, *Spoken languages*, takes its values into a broader set of candidates, yet limited to the set of languages names, which reduces the semantic ambiguity compared to open attributes. It can therefore be evaluated with a metric using exact match between values, such as F1-score. We evaluate *Spoken Languages* as a multi-label setup, using binarization per class followed by micro-F1, which is suited to evaluate a system’s ability to infer the set of (potentially multiple) languages spoken by each character.

5.2. Open-Class Evaluation

The open-class attributes (*Origin*, *Residence*, *Occupation*, *Physical Health*, *Type* (*non-human*)) can be expressed in variable ways while bearing a very

Attribute	Prediction	Gold reference	Human score	Interpretation Scale
Origin	America	France	0.0	Prediction is entirely incorrect
Occupation	Writer	Author, nurse	0.5	Partially correct, lacks some occupations
Physical health	Speech disorder	Stuttering	0.66	Not precise but correct
Type	Kitty	Cat	1.0	Correct or synonym

Table 3: Examples of the interpretation scale used by annotators to evaluate each systems’ predictions on open-class attributes.

similar meaning, for instance, “England” vs. “United Kingdom” for *Residence*, or “nobleman” vs. “count” for *Occupation*. BERTScore (Zhang et al., 2019) has been extensively used in the context of the evaluation of generative systems, and could theoretically be applied in our evaluation framework. However, the application of BERTScore within the context of S-VoCAL open-class attributes may entail the following potential limitations. First our gold references and system predictions are typically lists of words or entities, or short phrases, deprived of contextual dependencies, limiting the expressivity of BERT-based representations (Ethayarajh, 2019). A second issue is the reported inability for BERTScore to provide a faithful assessment of text quality in summarization (Fabbri et al., 2021), which led the research community towards LLM-based automatic evaluation metrics (Min et al., 2023).

Qwen3-based Metric To address these limitations, we propose to evaluate semantic similarities between gold and predicted attributes using attribute-aware representations computed from Qwen3-8b Embeddings⁸, a recent instruction-aware embedding model that computes text representations conditioned on an instruction. We leverage this sensitivity to instructions to derive small instructions prompts for each attribute, leading to attribute-conditioned representations of the gold references and predicted outputs. Then, we simply compute the cosine similarity between the attribute-conditioned gold and predicted representations.

Validation We validate the proposed metric quality by computing Spearman ρ correlations between the cosine similarity scores and human judgments. To produce the judgments, three annotators evaluated a subset of the predicted open-class attributes on a discrete scale from 0 to 1, where 0 indicates an entirely incorrect prediction, 1 a correct prediction, and intermediate scores reflect partial correctness. To ensure consistent annotations, we provide to each annotator an interpretable judgment for each score, complemented with examples (some of these examples are displayed in Table 3, and guidelines can be found in Appendix B).

⁸<https://huggingface.co/Qwen/Qwen3-Embedding-8B>

Attribute	α	ρ -Qwen3	ρ -BERT	n
Origin	0.95	0.85	0.44	66
Residence	0.93	0.85	0.31 ($p = 0.013$)	65
P. Health	0.74	0.52	-0.10 ($p = 0.572$)	37
Occupation	0.75	0.66	0.32	121
Type (n-h.)	0.92	0.86	0.63	57

Table 4: Inter-annotator agreement (α) and Spearman correlation (ρ) between human scores and automatic metrics. All correlations are significant at $p < 0.001$, except where specified *Abbreviations*: P. Health = *Physical Health*; Type (n-h.) = *Type (non-human)*.

We report inter-annotator agreement as interval Krippendorff’s α (Krippendorff, 2011) and the Spearman ρ correlation between human judgments and both the Qwen3-based metric and BERTScore in Table 4. Human annotations show a high level of agreement for all attributes, with Krippendorff’s alpha ranging from 0.74 to 0.95. Besides, the proposed metric shows a consistent and significantly higher correlation with human judgments than BERTScore on all attributes, with Spearman ρ ranging from 0.52 to 0.85 while BERTScore achieves zero-correlations for Physical Health and a maximum value of $\rho = 0.63$. Interestingly, correlations between the Qwen-3 based metric and human judgments was found to be lower on the *Physical Health* and *Occupation* attributes that also display lower agreement between annotators, which might indicate that comparing gold and predicted values for these attributes is more challenging.

Human-Aligned Score While the above correlation analysis suggests that our proposed metric is more reliable than BERTScore for evaluating open attributes, the raw cosine similarity scores lack direct human interpretability. Thus, we further leverage human judgments and compute a Human-Aligned Score (HAS) by fitting an isotonic regression line to map the Qwen3 similarity values onto the human evaluation scale for each attribute. This process automatically aligns raw similarity scores to interpretable values, allowing for a clearer understanding of model predictions. Note that this approach only provides a partial interpretation of

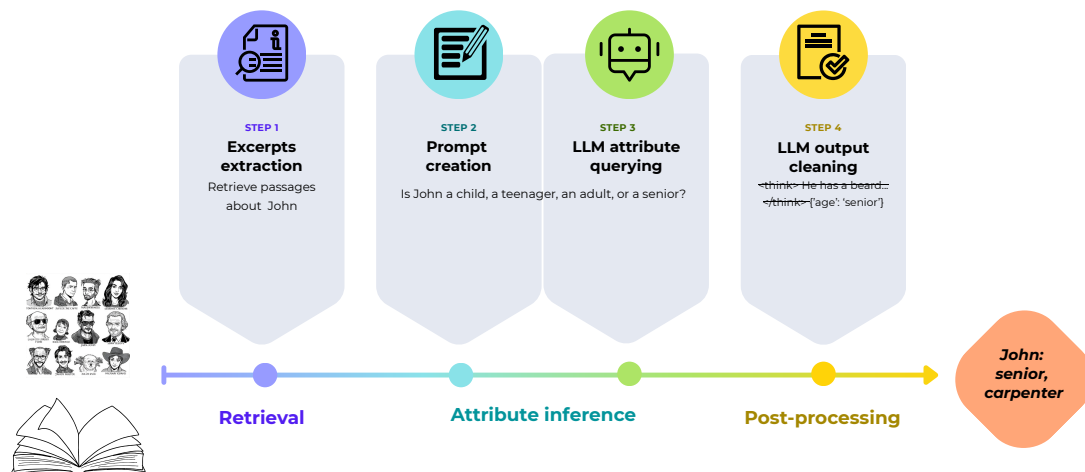


Figure 2: Overview of the RAG pipeline. The first step extracts attribute-specific relevant excerpts for a character. These passages are then prompted to an LLM along with an attribute prompt.

model predictions, subject to confounders in the isotonic regression. However, we found empirically that HAS were typically well aligned with human judgments and easier to interpret than cosine similarity scores, and thus propose to use HAS in S-VoCAL’s evaluation framework.

6. Attribute Inference with RAG

Following prior works (Papoudakis et al., 2024) we perform Attribute Inference on S-VoCAL with Retrieval-Augmented Generation (RAG). Given a novel and the character’s identifiers (name, aliases or nicknames), the RAG pipeline derives character’s attributes from relevant book excerpts. RAG is particularly suited in the Attribute Inference scenario, where relevant information about a character’s attribute might be disseminated in various book sections, or, in the worst-case scenario, might be implicitly mentioned once in the entire narrative. By retrieving and processing attribute-relevant book passages, we address the particular problem of “a needle in a haystack” that occurs when processing an entire book at once. The RAG pipeline is presented in Figure 2, and is composed of three main components described in the following subsections.

6.1. Retrieving Character Excerpts

The retrieval component selects the most relevant passages mentioning the character based on semantic similarity. Using the character’s name and aliases, all named mentions in the novel were first located with regular expressions, and 200-word windows were then extracted around each mention. We set this number to 200 experimentally, as it provides a good compromise between capturing relevant information and keeping the context size within the models’ limits. We then compute

passage embeddings with E5-large⁹ (Wang et al., 2024b), selected for its lightweight architecture and its query-based sentence embeddings, as well as its recent application in literary analysis, where it has proven effective in literary contexts (Hatzel and Biemann, 2024). To derive attribute-specific relevant excerpts, we perform the following steps:

- **Instruction:** We define a general instruction prompt, common to all attributes query
- **Attribute Query:** For each target attribute, a specific query is created (e.g., for Age: “Is the character a child, a teenager, an adult, or a senior?”). All queries are available in Appendix C.
- **Scoring:** We encode all queries and passages representations, then compute cosine similarities between both.
- **Filtering:** We select the ten most similar passages per attributes, and when multiple attributes are requested, duplicates across top-10 lists are removed. This list will be used as the input to the LLM-based Attribute Inference step.

6.2. Attribute Inference

Using the attribute-relevant excerpts derived from the retrieval step, we prompt an LLM to infer the requested character attributes. A single prompt is constructed for each character, combining all requested attributes in one query. In this work, we use Qwen3-8B (Yang et al., 2025); and Phi-4 14B (Abdin et al., 2024) as LLM backbones. Both models are quantized to 4 bits (Q4_K_M) to optimize

⁹<https://huggingface.co/intfloat/multilingual-e5-large-instruct>

		Qwen3-8b	Phi4-14b	Baseline
Weighted F1	Gender	0.993	0.997	0.561
	Age	0.783	0.800	0.660
	Type (h./n-h.)	0.969	0.971	0.910
Soft F1	Age	0.927	0.941	0.908
Micro F1	Spoken-Languages	0.712	0.749	0.576

Table 5: Results for close and semi-closed class attributes, with their dedicated metric. Type (h./n-h.) corresponds to the binary classification of Type Human vs Non-Human.

memory usage and inference speed, and we kept default generation parameters.

We observed that the model outputs were often noisy and partially malformed (e.g., containing reasoning traces, Markdown artifacts, or broken JSON structures). Thus, we apply a post-processing step to clean the predictions before converting them into structured character attributes (see Appendix D for details).

7. Results

We report the results of the RAG-based pipeline on S-VoCAL, assessed through the proposed evaluation framework.

7.1. Closed and Semi-Closed Class Results

For closed and semi closed class attributes, we compare LLM performance against a majority-value baseline, which always predicts the most frequent value for each attribute in S-VoCAL. The baseline provides a minimal reference point to contextualize the results obtained through the RAG-based pipeline.

We report results in Table 5. For closed-class attributes, both RAG-based pipelines perform relatively well, and significantly better than the majority-value baseline. When using Phi-4 14B as the LLM backbone, we observe slightly higher scores than when using Qwen3-8B. Both systems obtain very high Weighted F1 scores for *Gender* and *Type*, while the results for *Age* are slightly lower. We observe improved results when evaluating the *Age* attribute with soft-F1 which indicates that most errors occurred between adjacent classes, a behavior also reported in the manual annotations.

For the semi-closed attribute Spoken Languages, both RAG-based pipelines achieve considerable

Attribute	mHAS \uparrow	
	Qwen3-8b	Phi4-14b
Occupation	0.522	0.486
Residence	0.498	0.502
Origin	0.420	0.422
Type (n-h.)	0.785	0.781
Physical health	0.141	0.165

Table 6: Mean Human-Aligned Score (mHAS), for open attributes. For most attributes, values of 0.5 indicates a partially correct prediction, with lacking information. Type (n-h.) = Type (non-human).

improvements over the baseline, with scores indicating a majority of correctly inferred languages. However, the gap to the maximum score remains significant, suggesting occasional omissions or errors in the set of inferred spoken languages.

7.2. Open Class Results

We report in Table 6 the mean Human-Aligned Scores (mHAS), derived from the similarity metric presented in Section 5.2. We chose to report mHAS rather than raw similarity scores to facilitate interpretation, and refer the reader to Appendix E if interested in raw similarity scores. mHAS ranges from 0 to 1, where a value of 0.5 often indicates a partial valid answer. We refer the reader to Appendix B for detailed interpretation values for each attribute.

Type reaches the highest mHAS (≈ 0.78), corresponding to mostly correct predictions, but lacking precision in some cases (e.g., predicting “bird” instead of “crow”). *Occupation* (≈ 0.48 – 0.52) shows partially correct inference, with still some of the character’s occupations missing from the prediction. *Residence* (≈ 0.5) corresponds to predictions close to the expected value at a broader level, such as correct country but wrong region, or partially correct prediction, with part of the places to predict missing. *Origin* (≈ 0.42) falls between high-level correctness (right continent but wrong country) and partial overlap (some correct places but also some omissions). *Physical Health* (≈ 0.15) falls below the first interpretive level (0.33 = vague but correct distinction between poor and good health), implying that systems often fail to detect whether the character has a health condition or not. These results suggest that while *Type* is well retrieved, *Occupation*, *Residence* and *Origin* are often partially correct, and *Physical Health* remains challenging

Between the two LLMs pipelines, there is no clear winner for retrieving open-class attributes: we observe only minor fluctuations across attributes, except for *Occupation* which appears slightly better inferred by Qwen3-8b. Interestingly, we observed

no clear linear relationship between Qwen3 and Phi4 performance: if a model performs well on a specific instance, it does not necessarily imply that the other model also perform well on this instance, suggesting that both models have independent predictive behavior. While the reasons behind this independence remains unclear, we hypothesize that the two models might exhibit different book memorization patterns (Chang et al., 2023), which could lead to a model performing well on instances of a particular book that the other model has not memorized, suggesting worse performance on this book. We further detail this point in Appendix F.

8. Conclusion

This work introduces S-VoCAL, a dataset and evaluation framework designed to assess how well character attributes relevant to the design of a character’s voices can be inferred from novels. S-VoCAL entails a broad range of attributes grounded in sociophonetic studies, and covers 192 books and 952 characters. Given the attribute heterogeneity, we designed an extensive evaluation framework that includes standard metrics for closed and semi-closed cases, and a novel semantic similarity metric based on recent Large Language Models for open-class attributes. S-VoCAL thus offers a standardized benchmark for evaluating systems designed to infer voice-relevant character attributes from novels. We show the applicability of S-VoCAL by applying a simple RAG pipeline that infers voice-related attributes for a given character from entire novels. Our results suggest that this pipeline handles closed-class attributes (e.g., *Age*, *Gender*, *Type*) reliably, but remain limited in handling open attributes such as *Physical Health* or *Origin*. S-VoCAL contributes to the creation of immersive synthetic book narration, where each character is assigned a voice consistent with their identity.

9. Limitations

While this work presents an effort towards the creation of a dataset and evaluation framework for inferring voice-related attributes of fictional characters, it bears several limitations.

First, S-VoCAL only contains publicly available books drawn from Project Gutenberg. These books often entails canonical literature, and are very likely memorized by most Large Language Models (LLMs). This limitation is crucial when evaluating LLM-based attribute inference, as their reported performance on S-VoCAL does not necessarily imply that they will perform well on recent books, unseen during pre-training. However, we believe that S-VoCAL structured evaluation framework remains

useful when evaluating on such recent books. Alternatively, book memorization and the extended knowledge of LLMs on books seen during pretraining – which positively impact the accuracy of predicted attributes – might play a crucial role in the faithful voice design of widely known characters.

Besides, our manual annotations only covered a subset of the *Age* attribute, as we believe it was the most crucial attribute to design faithful character voices. However, other attributes present in S-VoCAL remain largely underrepresented, and extending this annotation to these attributes could be very beneficial to improve coverage and reliability of the evaluation across all attribute types.

In addition, S-VoCAL provides 8 static character attributes derived from Wikidata. However, many of the selected attributes are likely to evolve over time, such as *Age*, *Physical Health* or *Spoken Languages*. We found that recovering these dynamic patterns from any sources, either Wikidata or analysis websites, was particularly challenging, and decided instead to opt for static attributes that serve as a minimal standpoint to evaluate attribute inference systems. However, we believe that the dynamic modeling of attributes is of crucial importance in the faithful representation of a character’s voice throughout the story, and hope that S-VoCAL will serve as a first step towards this goal.

Finally, other character properties than those selected may have an influence on the character’s voices. However, such characteristics are either highly contextual (e.g., their emotions) or subject to interpretative judgments of the reader, and are therefore less consistently represented in structured resources. We believe that S-VoCAL provides a first step towards general character voice design, which can further be contextually rendered through additional inference layers that we leave as promising directions for future work.

10. Ethical Considerations

The dataset introduced in this work may carry structural biases, as Project Gutenberg, and especially books in English, mainly includes Western literary classics and Wikidata available information remains uneven across entities. Such imbalances can influence the types of characters and attributes that are most frequently represented and, consequently, evaluated.

Moreover, future work on synthetic voice generation should ensure that the use of the defined attributes remains responsible, by creating clear ethical guidelines and monitoring practices, to avoid unintended stereotypical associations between character traits and speaking styles.

Finally, because S-VoCAL dataset was built exclusively from publicly available books, it remains

imbalanced with a predominance of adult, male, and human characters. This distribution does not fully reflect the diversity of characters in literature, and including a more diverse set of books in future work could help address this limitation.

11. Bibliographical References

- Marah Abidin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Xin Wang, Rachel Ward, Yue Wu, Dingli Yu, Cyril Zhang, and Yi Zhang. 2024. [Phi-4 technical report](#).
- David Abercrombie. 1967. *Elements of General Phonetics*. Edinburgh University Press, Edinburgh.
- David Bamman, Ted Underwood, and Noah A Smith. 2014. A bayesian mixed effects model of literary character. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 370–379.
- Sabyasachee Baruah and Shrikanth Narayanan. 2024. [Character attribute extraction from movie scripts using llms](#). In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8270–8275.
- Faeze Brahman, Meng Huang, Oyvind Tafjord, Chao Zhao, Mrinmaya Sachan, and Snigdha Chaturvedi. 2021. "let your characters tell their story": A dataset for character-centric narrative understanding. *arXiv preprint arXiv:2109.05438*.
- James Brown. 2021. The Art of Storytelling – How Voice Over Brings Audiobooks to Life. <https://jamesbrownvoice.com/voice-over-bringing-audiobooks-to-life>.
- Gerald J. Canter. 1963a. Speech characteristics of patients with parkinson's disease. i. intensity, pitch and duration. *Journal of Speech and Hearing Disorders*, 28:221–229.
- Gerald J. Canter. 1963b. [Speech characteristics of patients with parkinson's disease: I. intensity, pitch, and duration](#). *Journal of Speech and Hearing Disorders*, 28(3):221–229.
- Kent Chang, Mackenzie Cramer, Sandeep Soni, and David Bamman. 2023. [Speak, memory: An archaeology of books known to ChatGPT/GPT-4](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7312–7327, Singapore. Association for Computational Linguistics.
- Ella Creamer. 2025. UK audiobook revenue up by almost a third last year . <https://shorturl.at/5lRYk>.
- Alan Cruttenden. 1986. *Intonation*. Cambridge University Press, Cambridge.
- Albert Di Cristo. 2013. *La prosodie de la parole*. De Boeck Superieur.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. [Summeval: Re-evaluating summarization evaluation](#).
- Sascha Frühholz and Pascal Belin. 2018. *The Oxford Handbook of Voice Perception*. Oxford University Press.
- Pablo Gervás and Gonzalo Méndez. 2024. Tagging narrative with propp's character functions using large language models. In *Text2Story@ ECIR*, pages 137–148.
- Brittany Goodwin. 2024. So You Want To Narrate Audiobooks. <https://www.brittanygoodwin.com/blog/2019/12/8/so-you-want-to-narrate-audiobooks>.
- Alexander Gurung and Mirella Lapata. 2024. [Chiron: Rich character representations in long-form narratives](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, page 8523–8547. Association for Computational Linguistics.
- Hans Ole Hatzel and Chris Biemann. 2024. [Story embeddings — narrative-focused representations of fictional stories](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5931–5943, Miami, Florida, USA. Association for Computational Linguistics.

- Harry Hollien and Thomas Shipp. 1972a. [Speaking fundamental frequency and chronologic age in males](#). *Journal of Speech and Hearing Research*, 15(1):155–159.
- Harry Hollien and Thomas F. Shipp. 1972b. [Speaking fundamental frequency and chronological age in males](#). *Journal of Speech and Hearing Research*, 15:155–159.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, Yang Zhang, and Boris Ginsburg. 2024. [Ruler: What’s the real context size of your long-context language models?](#)
- Elias Iosif and Taniya Mishra. 2014. [From speaker identification to affective analysis: A multi-step system for analyzing children’s stories](#). In *Proceedings of the 3rd Workshop on Computational Linguistics for Literature (CLFL)*, pages 40–49, Gothenburg, Sweden. Association for Computational Linguistics.
- Brandon Jaipersaud, Zining Zhu, Frank Rudzicz, and Elliot Creager. 2024. [Show, don’t tell: Uncovering implicit character portrayal using llms](#).
- Klaus Krippendorff. 2011. [Computing krippendorff’s alpha-reliability](#).
- William Labov. 1973. *Sociolinguistic patterns*. 4. University of Pennsylvania press.
- John Laver and Peter Trudgill. 1979. [Phonetic and linguistic markers in speech](#). *Social markers in speech*, 1:32.
- John Laver and Peter Trudgill. 1991. [Phonetic and linguistic markers in speech](#). In John Laver, editor, *The Gift of Speech*, pages 235–264. Edinburgh University Press, Edinburgh.
- Louis Mahon and Mirella Lapata. 2024. [A modular approach for multimodal summarization of tv shows](#).
- Gaspard Michel, Elena V. Epure, and Christophe Cerisara. 2025. [Libriquote: A speech dataset of fictional character utterances for expressive zero-shot speech synthesis](#).
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FACTScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Francis Nolan. 1987. [Alan cruttenden, intonation](#). cambridge: Cambridge university press, 1986. pp. xiv + 214. *Journal of Linguistics*, 23(2):455–458.
- Argyrios Papoudakis, Mirella Lapata, and Frank Keller. 2024. [BookWorm: A dataset for character description and analysis](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4471–4500, Miami, Florida, USA. Association for Computational Linguistics.
- Kyeongman Park, Seongho Joo, and Kyomin Jung. 2025. [Multiactor-audiobook: Zero-shot audiobook generation with faces and voices of multiple speakers](#).
- Herman Parret. 2002. [Chapitre 2. La qualité de la voix](#), Le Point philosophique, pages 39–66. De Boeck Supérieur, Louvain-la-Neuve.
- Charuta Pethe, Bach Pham, Felix D Childress, Yunting Yin, and Steven Skiena. 2025. [Prosody analysis of audiobooks](#). In *2025 19th International Conference on Semantic Computing (ICSC)*, page 217–221. IEEE.
- Emma Rodero and Ignacio Lucas. 2023. [Synthetic versus human voices in audiobooks: The human emotional intimacy effect](#). *New Media & Society*, 25(7):1746–1764.
- Nilanjana Roy. 2024. [Do audiobooks count as reading?](#) *Financial Times*.
- Piotr Sorokowski, Katarzyna Pisanski, Tomasz Frąckowiak, Aleksander Kobylarek, and Agata Groyecka-Bernard. 2024. [Voice-based judgments of sex, height, weight, attractiveness, health, and psychological traits based on free speech versus scripted speech](#). *Psychonomic Bulletin & Review*, pages 1–10.
- Dominik Stammbach, Maria Antoniak, and Elliott Ash. 2022. [Heroes, villains, and victims, and gpt-3: Automated extraction of character roles without training data](#).
- Statista. 2025a. [Audiobooks - China](#). <https://www.statista.com/outlook/amo/media/books/audiobooks/china>.
- Statista. 2025b. [Audiobooks - Europe](#). <https://www.statista.com/outlook/amo/media/books/audiobooks/europe>.
- Statista. 2025c. [Audiobooks - United States](#). <https://www.statista.com/outlook/amo/media/books/audiobooks/united-states>.

- Kobie Van Krieken, Hans Hoeken, and José Sanders. 2017. Evoking and measuring identification with narrative characters—a linguistic cues framework. *Frontiers in psychology*, 8:1190.
- Krishnapriya Vishnubhotla, Adam Hammond, Graeme Hirst, and Saif Mohammad. 2024. [The emotion dynamics of literary novels](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2557–2574, Bangkok, Thailand. Association for Computational Linguistics.
- Brendan Walsh, Mark Hamilton, Greg Newby, Xi Wang, Serena Ruan, Sheng Zhao, Lei He, Shaofei Zhang, Eric Dettinger, William T. Freeman, and Markus Weimer. 2023. [Large-scale automatic audiobook creation](#).
- Chonghua Wang, Haodong Duan, Songyang Zhang, Dahua Lin, and Kai Chen. 2024a. [Ada-level: Evaluating long-context llms with length-adaptable benchmarks](#).
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2024b. [Text embeddings by weakly-supervised contrastive pre-training](#).
- Alissa P Wolters, Young-Suk Grace Kim, and John William Szura. 2022. Is reading prosody related to reading comprehension? a meta-analysis. *Scientific Studies of Reading*, 26(1):1–20.
- Detai Xin, Sharath Adavanne, Federico Ang, Ashish Kulkarni, Shinnosuke Takamichi, and Hiroshi Saruwatari. 2022. [Improving speech prosody of audiobook text-to-speech synthesis with acoustic and textual contexts](#).
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chu-jie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. [Qwen3 technical report](#).
- Xinfeng Yuan, Siyu Yuan, Yuhan Cui, Tianhe Lin, Xintao Wang, Rui Xu, Jiangjie Chen, and Deqing Yang. 2024. [Evaluating character understanding of large language models via character profiling from fictional works](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8015–8036, Miami, Florida, USA. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. [Bertscore: Evaluating text generation with BERT](#). *CoRR*, abs/1904.09675.
- Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025. [Qwen3 embedding: Advancing text embedding and reranking through foundation models](#).

Appendices

A. Annotation Guidelines for Character Age

We describe below the guidelines given to human participants to annotate character age. Characters are grouped into four age categories:

- **Child:** From 0 to 12 year old
- **Teenage:** From 13 to 17 year old
- **Adult:** From 18 to 59 year old
- **Senior:** From 60 year old and beyond

A.1. Sources

Characters must be looked up online using their full name along with the book title, so their age can be inferred from web sources such as well-known websites Wikipedia, SparkNotes, Grade-Saver, LitCharts, and others. This type of website usually includes plot summaries of the book and character descriptions or analyses, from which character's age can be inferred. As the goal is to get trustworthy information, cross-sourcing the annotation is encouraged. If multiple sources provide conflicting information, choose the most consistent description or leave the annotation blank if uncertainty remains.

A.2. General Principle

The age category should be assigned using the following priority order:

1. Explicit age stated in the text.
2. Relative age information (e.g., "younger than his 18-year-old sister").
3. Direct descriptive terms or keywords (e.g., "boy", "old man", "teenage girl").
4. Behavior, role, or life situation (e.g., attending school, having a stable job, being married or a parent).

A.3. Steps of Decision for Annotation

1. **Is a numerical age explicitly stated?**
→ Assign the category using the table above.
2. **If not, is the age given relative to another character?**
Examples: "younger than his 18-year-old sister", a person described as older than someone who's already annotated as a senior.
→ Infer the most probable category.

Category	Frequent Keywords / Expressions
Child	baby, infant, toddler, boy, girl, little, kid, schoolboy, schoolgirl
Teenager	teen, adolescent, high school student, schoolboy (high school), young man, young woman (could also be "adult" depending on the context)
Adult	man, woman, husband, wife, worker, soldier, parent, teacher, merchant
Senior	old man, old woman, elderly, grandfather, grandmother, retired, frail, aged

Table 7: Frequent keywords used to infer age category.

3. **If not, are there clear descriptive terms?**
Examples: "child", "woman", "old man".
→ Assign the category accordingly.
4. **If not, can behavior or social role indicate age?**
 - Attending elementary school → Child.
 - In high school → Teenager.
 - Having a stable job or being married → Adult.
 - Being described as "elderly" or "frail" → Senior.
5. **If no reasonable conclusion can be made:**
Leave the annotation blank.

4. Special Rules for Non-Human Entities

- If the age is stated and realistic (e.g., an elf "20 years old"), apply the same rules as for humans.
- If the age is unrealistic (e.g., a "200-year-old vampire"), assign the category based on the character's behavior and appearance.
- If no age is stated, apply the same decision steps as for humans.

5. Examples of Common Keywords

We provide in Table 7 a list of common keywords used by age-category.

B. Guidelines for Assessing Model Predictions

We describe below the guidelines given to human participants to assess model predictions in open-class attributes. These guidelines include examples per category that can be used to interpret our mean-Human-Aligned-Scores.

B.1. General Evaluation Setup

For each attribute, three elements are defined:

- {attribute}_pred: the model's prediction.
- {attribute}_gold: the gold reference (correct answer).
- {attribute}_agreement: the agreement score assigned by the annotator.

Only the agreement score is annotated. Scores are chosen from a predefined set of values specific to each attribute.

B.2. Origin and Residence

Origin and Residence are evaluated as list attributes using the following agreement scale:

- **0.0**: Prediction is completely wrong (e.g., unrelated to the question, wrong continent).
- **0.25**: Prediction is partially related at a high level (e.g., correct continent but wrong country).
- **0.5**: Prediction is partially correct at a lower level (correct country, but wrong region or sub-area) OR partially correct but missing places.
- **0.75**: Prediction is correct but less specific than the gold, or overly specific and not fully verifiable.
- **1.0**: Prediction is fully correct and fully correspond to the gold, synonyms included. In the case of cities, accept cities which are very close to each other. If gold is a country, accept capital city or very big city in the country.

Examples

- 0.0: gold = ["France"], pred = ["America"]
- 0.5: gold = ["France", "Bretagne"], pred = ["France", "Marseille"]
- 1.0: gold = ["France"], pred = ["France", "Paris"]

B.3. Physical Health

- **0.0**: Prediction not related or incorrect good/poor status.
- **0.33**: Vague prediction, but correctly identifies good/poor health.
- **0.66**: Correct but less specific than gold.
- **1.0**: Fully correct or synonymous prediction.

Examples

- 0.0: gold = ["Missing leg"], pred = ["good"]
- 0.33: gold = ["Cancer"], pred = ["ill"]
- 1.0: gold = ["Deaf"], pred = ["Deafness"]

B.4. Occupation

- **0.0**: Completely wrong occupation.
- **0.25**: Same general domain but incorrect.
- **0.5**: Prediction partially correct but lacks some occupations not related at all to those found.
- **0.75**: Prediction in the same domain but not as precise as gold or prediction correct + unrelated occupation not mentioned in gold.
- **1.0**: Correct prediction or synonym, can be more precise than gold.

Examples

- 0.0: gold = ["writer"], pred = ["farmer"]
- 0.25: gold=["writer"] pred = ["editor"]
- 0.5: gold = ["novel author", "nurse", "mother"], pred = ["writer"]
- 1.0: gold = ["soldier", "captain"], pred = ["military sailor"]

B.5. Type

- **0.0**: Completely wrong or human/non-human confusion.
- **0.33**: Wrong entity but correct human/non-human distinction.
- **0.66**: Prediction more general than gold.
- **1.0**: Correct prediction or synonym.

Examples

- 0.0: gold = ["human"], pred = ["cat"]
- 0.33: gold = ["cat"], pred = ["dog"]
- 0.66: gold = ["horse"], pred = ["equine"]

C. Attribute Query Templates

For each character and each attribute, the following query templates are used for retrieval of the relevant passages with E5-large model:

- **Age:** "Is {character_name} a child, a teenager, an adult, or a senior?"
- **Gender:** "Is {character_name} male or female?"
- **Origin:** "Where is {character_name} from?"
- **Residence:** "Where does {character_name} live?"
- **Spoken Languages:** "What languages does {character_name} speak?"
- **Type:** "What type of entity is {character_name}?"
- **Occupation:** "What is {character_name}'s occupation?"
- **Physical Health:** "How is {character_name}'s health condition?"

D. Post-processing of Model Outputs

We summarize in Table 9 the main post-processing operations applied to raw LLM outputs. Code for post-processing is available on the associated repository at <https://github.com/AbigailBerthe/S-VoCAL>.

D.1. Example

We present below some examples of the post-processing operation applied to the LLM outputs

Before post-processing (simplified):

```
<think> (...) </think>
{
  "origin": ["Europe", "Germany", "Parchim"],
  "residence": ["Germany", "Hamburg"],
  "spoken_languages": [
    "German", "English", 'null']
}
```

Attribute	Qwen3	Phi4	Baseline
Origin	0.481	0.446	0.512
Residence	0.426	0.476	0.366
Type (non-human)	0.869	0.810	0.569
Occupation	0.561	0.558	0.403
Physical Health	0.301	0.329	0.317

Table 8: Mean cosine similarity scores per open attributes

After post-processing:

```
{
  "origin": [
    "Europe", "Germany", "Parchim"],
  "residence": ["Germany", "Hamburg"],
  "spoken_languages": [
    "German", "English"]
}
```

E. Raw Similarity Scores

Table 8 reports the mean cosine similarity scores for each open attribute.

F. Distribution of Phi-4 and Qwen3 mHAS

Figure 3 presents an instance-level comparison of the mHAS scores assigned to outputs produced by Phi-4 and Qwen3, across the open-class attributes. We see no linear relationship between model's performance, indicating that models tend to make errors on different instances, although provided with the same input. We hypothesize that this behavior might be due to different book-memorization per model.

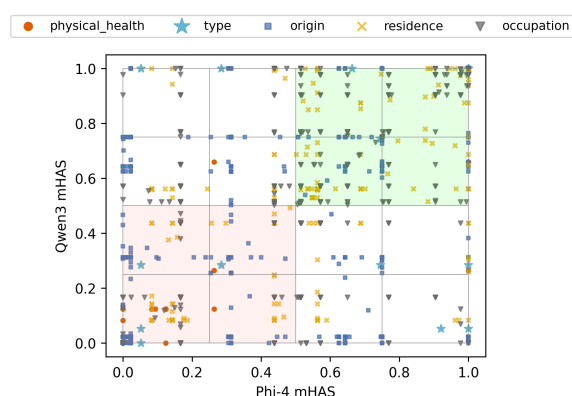


Figure 3: Comparison of Phi-4 and Qwen3 mHAS scores across attributes.

Input problem (example)	Cleaning operation	Expected result
Reasoning traces <think>...</think>	Remove reasoning segment	Traces removed
Markdown fences ``json ... ``	Remove formatting markers	Keep only JSON
Invalid values : None	Normalize invalid value	Valid JSON
Multiple JSON blocks	Select final JSON block	Single target block
Unquoted list items [Paris, London]	Add missing quotes	["Paris", "London"]
Extra brackets or commas]],	Repair unbalanced separators	Balanced brackets
Missing commas between ob- jects	Repair list separators	Properly separated elements
Nested malformed lists [["Paris"]]	Normalize list structure	["Paris"]
Literal 'null' inside lists	Remove placeholder value	Cleaned list
origin as key-value pairs ("ori- gin": 'country': "France", "city": "Paris")	Flatten into list of values	"origin": ["France", "Paris"]
Non-string input when string ex- pected	Discard attribute	<i>null</i> as safe default

Table 9: Cleaning rules applied to model outputs