



HAL
open science

Age Recommendation from Texts and Sentences for Children

Rashedur Rahman, Gwénolé Lecorvé, Nicolas Béchet

► **To cite this version:**

Rashedur Rahman, Gwénolé Lecorvé, Nicolas Béchet. Age Recommendation from Texts and Sentences for Children. 2024. hal-04756683

HAL Id: hal-04756683

<https://hal.science/hal-04756683v1>

Preprint submitted on 28 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Age Recommendation from Texts and Sentences for Children

Rashedur Rahman

Univ Rennes, CNRS, IRISA
Lannion, France
md-rashedur.rahman@irisa.fr

Gwénoél Lecorvé

Orange
Lannion, France
gwenole.lecorve@orange.com

Nicolas Béchet

Univ. Bretagne Sud, CNRS, IRISA
Vannes, France
nicolas.bechet@irisa.fr

Abstract

Children have less text understanding capability than adults. Moreover, this capability differs among the children of different ages. Hence, automatically predicting a recommended age based on texts or sentences would be a great benefit to propose adequate texts to children and to help authors writing in the most appropriate way. This paper presents our recent advances on the age recommendation task. We consider age recommendation as a regression task, and discuss the need for appropriate evaluation metrics, study the use of state-of-the-art machine learning model, namely Transformers, and compare it to different models coming from the literature. Our results are also compared with recommendations made by experts. Further, this paper deals with preliminary explainability of the age prediction model by analyzing various linguistic features. We conduct the experiments on a dataset of 3, 673 French texts (132K sentences, 2.5M words). To recommend age at the text level and sentence level, our best models achieve MAE scores of 0.98 and 1.83 respectively on the test set. Also, compared to the recommendations made by experts, our sentence-level recommendation model gets a similar score to the experts, while the text-level recommendation model outperforms the experts by an MAE score of 1.48.

1 Introduction

Nowadays, children spend considerable time on the Internet. The primary concern is to ensure them to accessing adequate online content depending on their age. Therefore, in recent years, safe Internet for children has gained interest in many research domains (Tomczyk and Kopecký, 2016; Byrne and Burton, 2017; Livingstone, 2019). Most studies focus on filtering abusive texts containing hate, violence, pornography, etc. (Liu and Forss, 2015; Suvorov et al., 2013). However, the adequacy also relates to the capacity of a child to understand the textual content (blogs, newspapers, etc.). Beyond the issue of safely browsing the web, this question of comprehensibility more widely concerns all types of documents that could be submitted or narrated to children, like books, schoolbook, etc. Thus, besides the impact on the user experience (children’s and parent’s ones), recommending and understanding the linguistic adequacy of a textual content for children of a given age is of primary interest to help children build up their knowledge.

Overall, text understanding by children is a well known topic in psycho-linguistics and cognitive sciences (Gathercole, 1999; Tartas, 2010; Mouw et al., 2019). In particular, key findings have shown the impact of memory (Gathercole, 1999), temporality (Tartas, 2010; Hickmann, 2012), and emotions (Davidson, 2006; Mouw et al., 2019). On the contrary, this is still poorly studied in Natural Language Processing (NLP). Yet, the automatic recommendation of a target age for a given text could be benefit to various applications. For instance, this could help search engines to provide custom search settings by indexing a recommended age along with each document/web page. Computer-assisted writing tools could also be proposed to help authors or editors to propose suitable texts.

Toward this objective, this paper studies the task of automatically predicting a recommended age from textual productions. This task is addressed as a machine learning regression problem where multiple dimensions are defined and studied. The idea is to provide a solid and detailed frame for further progress in this area, both regarding methodological aspects and experimental results. This includes the following contributions:

1. The task of age recommendation from texts is defined and discussed to clarify difficulties or shortcomings possibly induced by some key choices, especially when defining the the fuzzy notion of "recommended" age. This mixes questions about the quality/reliability of the data annotations, the biases that specific text genres/styles can introduce or the evaluation of the results.
2. The paper proposed a dataset for textual age recommendation. This dataset is composed of French texts of various genres (novels, newspapers, and encyclopedias) for a wide range of ages. Moreover, the approach for the construction of this dataset can be easily transposed to other languages.
3. A particular attention is paid to the potentially relevant metrics for the task through the proposition of various metrics. Their respective behavior is studied through an artificial study case, and the best metrics are then used all along the paper.
4. With existing machine learning techniques, various recommendation models are proposed and compared. They cover various settings such that age can be recommended from texts or isolated sentences; using word/sentence embeddings, linguistic expert features or both; and with more or less complex models (transformers, feed-forward neural network or random forests). The results are compared against naive, historical and expert baselines.
5. Since explainability is required for the perspective of human interactive applications, the paper proposes an opening where various linguistic features are studied to understand what makes a text easy or difficult to understand.

In the remainder, Section 2 discusses the related work, Section 3 precisely defines the age recommendation task, and Section 4 details the dataset used for the experiments. An exhaustive study of the evaluation metrics is provided in Section 5 which is followed by the recommendation models and experimental results in Section 6. Finally, Section 7 investigates which linguistic features contribute most for age recommendation.

2 Related Work

Age recommendation from text is related to text understanding/readability analysis, and it has been studied in the fields of human and social sciences (Section 2.1), and NLP (Section 2.2).

2.1 Human and Social Sciences

In the literature of psycho-linguistics and cognitive sciences, text comprehension by children has been studied well. Children have less text understanding capabilities than adults as their brain is still developing. The brain activity of a child peaks at 4 years. For language acquisition, the activity can be equivalent to 150% of an adult (Gathercole, 1999). Short-term memory affects language understanding or accomplishment of complex tasks, and such memory is developed particularly between two to eight years (Gathercole, 1999). Acquisition of temporal notions is crucial for children to understand calendar-time and chronological orders (Tartas, 2010; Hickmann, 2012). Emotions are also important factors to establishing and maintaining the coherency of facts in a text (Mouw et al., 2019), and the basic emotions (joy, anger, sadness, fear, etc.) are acquired around the age of 10 (Davidson, 2006). When children start learning to read, (Frith, 1985) argued that reading is acquired through three main stages: logographic, alphabetical, and orthographic stage. The logographic stage refers to the faculty of recognizing the drawing of a word rather than deciphering it, and is developed between the ages of 5 to 6 years. At later ages, the second stage (alphabetical), splitting a word into simpler graphical units (graphemes) and phonological units (phonemes), and the third stage (orthographic), linguistic ability to break down a word into meaningful units (morphemes) are developed.

However, one can relate age recommendation to the task of text readability, where various scores have historically been proposed to quantify the difficulty of a text, like the Flesch-Kincaid score (Kincaid et al., 1975), Dale-Chall formula (Dale and Chall, 1948; Chall and Dale, 1995) or the Gunning fog index (Robert, 1968). The Flesch-Kincaid score computes frequency ratios on syllables, while the Dale-Chall formula

additionally counts difficult words. In contrast, the Gunning fog index considers the ratio of word to sentence counts, and the ratio of complex words to total words counts in a text. More recently, (Todirascu et al., 2013) studied the use of coherence and cohesion properties to assess the readability of texts, and they found some good correlations between such property and readability. (Islam and Rahman, 2014) proposed to predict texts into four classes (very easy, easy, medium, and difficult) where the readability classes corresponded to some child age ranges. The approach inherits the readability index features along with the lexical features. In a very recent work, (Wilkens et al., 2022) developed a readability assessment toolkit for french texts that aggregates various linguistic variables to train a classifier and a regression model. Also, this toolkit makes a ranking of the different features by measuring their correlations to the 9 readability levels.

2.2 Natural Language Processing

In spite of these psycho-linguistic studies, the adequacy between ages and linguistic or cognitive skills has been less studied in the field of NLP. The first step comes from (Schwarm and Ostendorf, 2005) where the authors explored how to automatically predict from which US school grade a child could read news articles. This problem was considered as a classification task among four classes. The model was a support vector machine with word n-gram, lexical and syntactic features.

The traditional study of text readability (Dell’Orletta et al., 2011; Islam et al., 2012) or age recommendation (Schwarm and Ostendorf, 2005) mostly relies on the hand-crafted linguistic features which may become very difficult to extract for low-resource languages. In recent years, using neural networks for text readability classification (Mesgar and Strube, 2018; Balyan et al., 2020; Martinc et al., 2021; Feng et al., 2022) and age prediction from text (Bayot and Gonçalves, 2017; Chen et al., 2019) has risen great interest in the field of computational linguistics. One major strength of the current approaches is to delegate the step of feature extraction to the neural network model based on the sole word embeddings of the input texts. Hence, such methods do not anymore necessarily require hand-crafted features for the training purpose. For instance, Blandin et al. (2020) explored feed-forward neural networks to recommend the age corresponding to a text. The authors showed, among a large set of various linguistic features, word embedding features were the most contributory. However, these models still suffer to achieve very good scores, and lack explainability of the predictions, but they show how difficult the task is.

Apart from age recommendation or text readability, one popular related task is to characterize authors of textual contents. In Nguyen et al. (2011), linear regression was used with gender and different textual features to predict the age of the authors of blogs, telephone conversations, and online forum posts. Similarly, Chen et al. (2019) proposed LSTM-based regression and classification models to guess adult ages (from 14 to 34 years) of blog authors. The same task is performed by Bayot and Gonçalves (2017) on tweets using a Convolutional Neural Network (CNN). Finally, an active learning based approach has also been explored in Chen et al. (2016) as a regression task for age prediction from social media texts by using textual and social features. The task of author’s age prediction from text is also a difficult one as the prediction performances are not so good yet. However, these studies show prospects of using regression models for age recommendation task.

Finally, in a very broad scope, in the last years, transformers (Wolf et al., 2020) have become a very attractive neural architecture for many NLP tasks. Especially, encoder models like BERT (Devlin et al., 2019) for many regression and classification tasks. Generally, such a model is pre-trained with a very large dataset and then it is used to perform on some downstream tasks by fine-tuning it with a smaller dataset. This architectural improvement motivates us to explore the advanced models for our age recommendation task. Nonetheless, it is interesting to highlight that these models are slower and more energy-consuming than former machine learning models. Hence, their benefit in terms of performance for the task should be balanced with the relative benefit compared to simpler approaches.

Yet, to our knowledge, no research studied any pre-trained BERT type model for age recommendation from text. In this paper, we propose to use CamemBERT (Martin et al., 2020), a pre-trained French language model built on the RoBERTa (Liu et al., 2019) architecture. We consider age recommendation as a regression task. To compare with our proposed CamemBERT approaches, this paper explores also

Feed-forward and a Random Forest models. Moreover, we introduce two metrics (by tackling some special cases) for evaluating age recommendation, and provide a deeper analysis on these metrics. Finally, to explain the recommendations or predictions (as a preliminary explainability), this paper presents different feature ranking approaches to show which linguistic features are most contributory for age recommendation.

3 Task Definition

We aim to assist children in their reading activities (online, leisure, or in schools) by suggesting them comprehensible texts according to their ages, and to facilitate the authors in writing the appropriate texts for the children. Therefore, the goal is to predict the target age from a given text.

3.1 Challenges

Text understanding capability differs to different children based on their ages, literacy, vocabulary etc. The children of same age may not have the same capacity of comprehending a piece of text. Some related studies worked on recommending texts to different school-grade levels (Schwarm and Ostendorf, 2005), and on text-complexity classification for different groups of school-grade levels (Islam et al., 2014). However, the grade levels corresponding to child-ages vary in different cultures and/or countries. When some grade levels or child-ages are merged into a group (or class), another question arises how large the interval of grade levels (or child-ages) is to be allowed. On the contrary, if one considers text recommendation for each individual child-age (or grade level), it lacks sufficient annotated data (to train a recommendation model) corresponding to the ages. Thus, from the computational linguistic point of view, defining the target age for automatically recommending age is very difficult. It requires to do some trade-offs among the different issues.

3.2 Target Age

In age recommendation task, the target age can be defined in several ways. First, one can think of the minimum age at which a text is comprehensible. However, in real world, suggesting a text to a child based on the minimum age may not be very practical as he/she cannot find it very interesting after certain age. For example, in general, a child of 12 years does not like the text (poem, fiction etc.) written for the kids of 4 years although he/she understands it very well. Such issues can be tackled by defining the target age as an interval (or range) of lower bound and upper bound, $[a, b]$ meaning that a text, T is comprehensible by the children of ages between a and b years. This interval further can be deduced to a mean-age, $\mu = \frac{a+b}{2}$ with a considerable deviation. For example, if a text corresponds to the age-range of $[4 - 8]$, the target age is 6 years with a deviation of 2 years. In this paper, the target age is mainly defined as an age-range, but it studies the mean-age(μ) as well.

3.3 Recommendation Method

Regarding the definition of the target age, we consider age recommendation as a regression task to predict the age as a numeric value from an input. The input of age recommendation can be either a full text or a sentence. The text level and sentence level recommendations allow us to deeper analyze different aspects and to serve different objectives of this task. In text level age recommendation method, the target age is predicted by analyzing the full text as a whole. However, for sentence level recommendation, each sentence is analyzed separately. Further, the sentence level predictions can be aggregated (e.g., mean aggregation) to make a text level global recommendation. Thus, the text level recommendation method can serve only the readers while the sentence level recommendation method can be useful for an author as well to incrementally verify (while writing a text) if the text is suitable for the readers of a certain age.

In age recommendation, an input text/sentence is represented with several features which are computed by various linguistic analysis. The feature-values are fed to a regression model in a supervised manner. Such a model is able to predict an age-range, an individual bound (min or max), and a mean age depending on the training setup.

	Genre	# texts	# sentences	Age range (avg.)	Mean age (avg.)
Train	Encyclopedia	738	21,960	[11.33 , 15.71]	13.52
	Newspaper	793	18,245	[09.44 , 13.85]	11.64
	Fiction	966	47,277	[09.06 , 11.85]	10.45
	Overall	2,509	87,882	[09.73 , 13.26]	11.49
Validation	Encyclopedia	201	6,495	[10.62 , 15.13]	12.87
	Newspaper	147	3,224	[09.38 , 13.78]	11.58
	Fiction	255	12,471	[07.86 , 10.90]	09.38
	Overall	605	22,261	[08.90 , 12.57]	10.74
Test	Encyclopedia	135	4,325	[11.52 , 15.86]	13.69
	Newspaper	189	4,118	[09.49 , 13.86]	11.67
	Fiction	231	13,186	[08.66 , 11.83]	10.24
	Overall	559	21,701	[09.41 , 13.03]	11.22
Total	Encyclopedia	1074	32,780	[11.21 , 15.61]	13.41
	Newspaper	1129	25,587	[9.44 , 13.84]	11.64
	Fiction	1452	72,934	[8.78 , 11.68]	10.23
	Overall	3,673	131,844	[9.53 , 13.11]	11.32

Table 1: Sizes and ages for the whole, training, validation and tests datasets. Statistics are given overall and for each text genre.

4 Data

This section presents the corpus of 3,673 French texts (132K sentences, 2.5M words) that we have constituted to perform age recommendation. The main statistics of this corpus are given in Table 1 and are described in the remainder. This corpus is intended to be publicly available¹.

4.1 Age Annotations

All texts are associated with an age range. As previously discussed, this range is interpreted as the rough interval of ages from which the text is mostly adequate and can be fully understood. Age ranges come from indications given by the authors or editors in the case of children-dedicated texts, whereas they are arbitrary set to [14 , 18] for adults. It is important to highlight that these annotations are probably imperfect since editors and authors are not psycho-linguists². Indeed, age ranges are usually empirically defined, based on the feeling and experience of the authors and editors, as well as marketing considerations³.

When working on isolated sentences later on in the paper, please note that all sentences will be associated with age range of the text from which they come. For example, if a text annotated with an age range of $[a, b]$, all the sentences in the text are considered with the same age range. This is an inaccurate assumption since it is clear that all sentences of a difficult text may not be difficult when taken in isolation. Maybe only few of them are difficult or even the difficulty may come from the reasoning linking the sentences. However, this is the best approximation that one could reasonably expect as annotating thousands of sentences would be difficult and there would not be any guarantee that these annotations are consistent with the text level ones (again, the logics of the editors and authors is unknown).

¹Due to the legal issue, now we are currently able to publicly release 25% of this dataset. The complete dataset will be released very soon.

²Experiments of Section 6.3 provide an idea of the differences between these reference age ranges and those recommended by experts.

³For instance, age ranges for novels/stories are usually associated as a whole to collections, rather than individually to each book. Furthermore, some strategies may incite to lower the age range if this can help differentiating a product from its concurrent. Furthermore, age ranges tend to be large to attract more parents.

4.2 Balancing and partitioning

To avoid linguistic biases due to specific genres, the corpus mixes encyclopedic, newspaper and fictional texts. Encyclopedic texts comes from Vikidia⁴ and Wikimini⁵ for children, and Wikipedia article on difficult topics regarding the adult age range. Newspaper texts come from various newspapers and magazines like *Le P'titLibé*⁶ or *Albert*⁷ for children, whereas they come from usual newspapers for adults (*L'Humanité*, *Le Monde diplomatique*, etc.). Finally, fictional texts come from stories or novels with various authors.

The corpus is balanced in terms of text genres and age ranges. However, no specific attention has been paid on balancing the topics across the age ranges and genres. This is simply because multiplying the constraints makes it impossible to build a corpus large enough to perform machine learning⁸.

Text lengths are very different according to their genre and target age, especially regarding novels which can be very long for adults. The information about length could be taken by the models as a clue in favor of high ages. To avoid this bias, all texts larger than 10,000 characters have been segmented on paragraph boundaries⁹ with a segment length of 5,000 characters approximately. This segmentation avoids the biases due to the genre and target age too. Moreover, it is consistent with one of our goals to index online texts (which are rarely presented as one very long page but rather paginated) or to assist writers (they may probably not want to study their whole text but rather portions of it).

The whole dataset is partitioned into train/validation/test sets with respective proportions of 68.3/16.5/15.2% in terms of texts, and 66.7/16.9/16.5% regarding the sentences. All texts segments coming from a same original text are gathered in the same subset to guarantee that the validation and test data are completely unseen data after training. In all subsets, fictions count more texts and sentences than two other genres of encyclopedia and newspaper. Moreover, the average age ranges and mean ages per subset and per genre vary slightly. Especially, it can be noticed that the ages from the validation set differ more to the training set than the test set does.

4.3 Statistics on Ages

The distributions of the texts and sentences per age are shown in Figure 1. One text or sentence with age range $[a, b]$ is counted up for an age x as soon as $x \in [a, b]$. This figure shows that the ages from 8 to 12 years are slightly more represented, especially regarding sentences. This corresponds to the period where children can read but are still considered as children (as opposed to teenagers). In contrast, few texts/sentences correspond to the ages between 0 to 3 years. This is normal as the children of these ages are too young to read and their attention time is smaller. Hence, these texts are supposed to be narrated, and they are short.

The detailed statistics of the dataset per age range are shown in Table 2. The dataset contains texts for 32 different age ranges. However, the train, validation and test sets individually does not cover the same age ranges. The three sets contain 29, 19, and 24 age ranges, respectively where 18 age ranges are in common.

5 Evaluation metrics

This section presents and studies different metrics to compute the error between a reference interval $r = [a, b]$ provided by the annotations, and a hypothesis interval $h = [c, d]$. While this study does not pretend to be exhaustive, the objective is to highlight the main questions for the age recommendation task and to come up with one or few metrics for the remainder of the paper, hopefully future work in the community as well.

⁴fr.vikidia.org

⁵fr.wikimini.org

⁶<https://ptitlibe.liberation.fr/>

⁷<https://www.journal-albert.fr/>

⁸In this regard, the corpus could have been much large if text genres were not balanced since encyclopedic data for children is easy to massively collect.

⁹In practice, these boundaries are detected by blank lines instead of simply line returns since the latter strategy would badly segment dialogues.

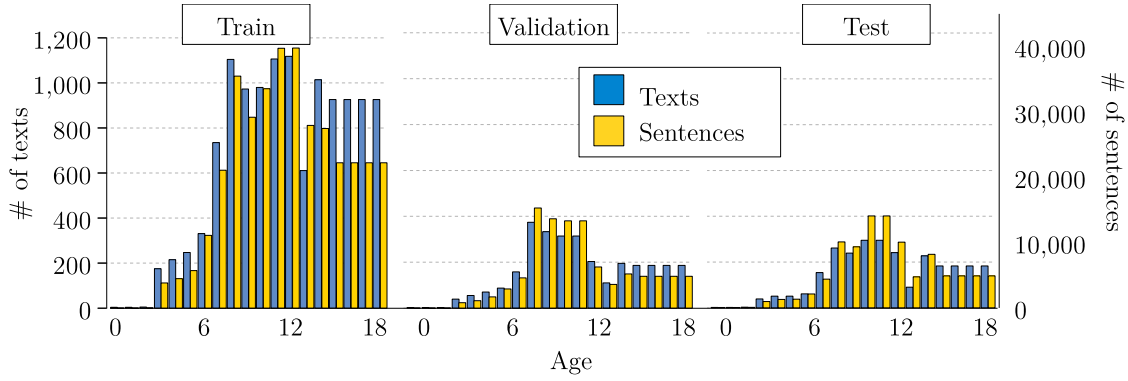


Figure 1: Distribution of the texts and sentences over the different ages in the train, validation, and test sets

Age		Train			Validation			Test		
Range	Mean	Texts	Sent.	Tokens	Texts	Sent.	Tokens	Texts	Sent.	Tokens
[0 , 2]	1.0	1	6	73	–	–	–	–	–	–
[0 , 3]	1.5	3	29	255	2	33	277	2	60	686
[2 , 4]	3.0	–	–	–	–	–	–	1	17	258
[2 , 7]	4.5	1	20	486	–	–	–	–	–	–
[3 , 5]	4.0	28	1,116	13,318	7	241	3,029	8	206	2,402
[3 , 6]	4.5	79	1,552	25,660	11	197	3,215	16	360	5,143
[3 , 7]	5.0	15	321	5,034	7	104	1,669	4	88	1,225
[3 , 8]	5.5	29	683	12,761	11	262	4,371	5	107	1,925
[3 , 9]	6.0	20	459	7,918	1	23	356	3	132	2,229
[4 , 6]	5.0	–	–	–	1	58	650	1	121	1,121
[4 , 7]	5.5	1	20	336	–	–	–	1	37	948
[4 , 8]	6.0	31	494	6,917	14	219	3,318	10	166	2,393
[4 , 9]	6.5	11	217	4,614	3	69	1,358	2	59	1,254
[5 , 7]	6.0	5	275	4,301	–	–	–	–	–	–
[5 , 8]	6.5	1	38	623	–	–	–	–	–	–
[5 , 9]	7.0	25	989	17,420	15	592	11,493	1	48	948
[5 , 12]	8.5	1	27	632	–	–	–	–	–	–
[6 , 8]	7.0	70	5,571	69,513	16	1,208	15,494	7	479	6,311
[6 , 9]	7.5	42	1,365	24,396	8	276	5,414	8	361	7,219
[6 , 12]	9.0	–	–	–	–	–	–	3	167	2,879
[7 , 11]	9.0	17	1,096	13,608	–	–	–	4	279	3,306
[7 , 12]	9.5	466	11,231	232,898	83	1,989	41,918	106	2,477	51,352
[8 , 10]	9.0	57	4,915	58,352	–	–	–	–	–	–
[8 , 11]	9.5	29	1,706	27,799	120	7,562	104,882	62	4,232	56,106
[8 , 12]	10.0	28	765	13,507	2	80	1,390	4	129	2,283
[8 , 13]	10.5	277	8,781	223,954	101	3,356	79,749	46	1,456	36,687
[10 , 12]	11.0	12	751	9,707	8	645	7,105	38	2,533	32,724
[10 , 13]	11.5	63	4,169	57,710	–	–	–	–	–	–
[10 , 14]	12.0	30	2,815	28,214	–	–	–	32	2,790	27,946
[11 , 13]	12.0	183	11,593	175,219	–	–	–	–	–	–
[12 , 14]	13.0	58	2,844	43,870	9	365	6,243	12	503	8,401
[14 , 18]	16.0	926	24,034	654,148	186	4,982	129,408	183	4,894	125,796
Total		2,509	87,882	1,733,243	605	22,261	421,339	559	21,701	381,542

Table 2: Number of texts, sentences and tokens for each age range in the train, validation and test sets. Age ranges are sorted according to their lower bound first. The sign – is used when no data is available.

Overall, most approaches vary according to the meaning behind the notion of interval. On the one hand, one can assume that an interval $[x, y]$ is exact, that is all children whose age is in this interval will understand the text. Besides, uncertainty can be taken into account to reflect the fact that the real (reference) range/value is unknown and the only available information is that this is somewhere within $[x, y]$. As previously discussed, this uncertainty assumption seems particularly reasonable for our data since their underlying annotation process by the authors and editors is not perfectly understood.

Below, the section presents several approaches and their most relevant metrics¹⁰. Finally, these metrics are compared and discussed on a study case.

5.1 Usual vector metrics

Considering ranges as 2D vectors $\vec{r} = \begin{pmatrix} a \\ b \end{pmatrix}$ and $\vec{h} = \begin{pmatrix} c \\ d \end{pmatrix}$, usual metrics are the L_1 and L_2 distances (“absolute Error” and “root squared error”), that is:

$$L_1(\vec{h} - \vec{r}) = |c - a| + |d - b|, \quad (1)$$

$$\text{and } L_2(\vec{h} - \vec{r}) = \sqrt{(c - a)^2 + (d - b)^2}. \quad (2)$$

Intuitively, one could argue that it is better when the hypothesis interval is included in the reference one, rather than the contrary. This can be taken into consideration by measuring the angle θ between the vector $\vec{h} - \vec{r} = \begin{pmatrix} c - a \\ d - b \end{pmatrix}$ and $\vec{i} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ ¹¹. Mathematically, this is computed as:

$$\cos(\theta) = \cos(\vec{h} - \vec{r}, \vec{i}) = \begin{cases} 1 & \text{if } \vec{r} = \vec{h}, \\ \frac{c - a - d + b}{\sqrt{2} L_2(\vec{h} - \vec{r})} & \text{otherwise.} \end{cases} \quad (3)$$

L_2 and $\cos(\theta)$ can then be linearly interpolated as follows to provide a new metric:

$$\theta\text{-}L_2(\vec{r}, \vec{h}) = L_2(\vec{h} - \vec{r}) + \alpha(1 - \cos(\theta)), \quad (4)$$

where α is an empirically tunable weight, and $1 - \cos(\theta)$ ranges in $[0, 2]$.

The main drawback of all these metrics is that the lower and upper bounds of the intervals are processed independently, whereas it is clear that they are not independent.

5.2 Jaccard distance

Jaccard distance, which is usually applied on sets, can be extended to intervals and can be used to evaluate age ranges. The definition is as follows:

$$J([a, b], [c, d]) = 1 - \frac{\| [a, b] \cap [c, d] \|}{\| [a, b] \cup [c, d] \|}, \quad (5)$$

where $\|\cdot\|$ denotes the norm of an interval. The resulting measure ranges in $[0, 1]$ where 0 means that the 2 intervals do not overlap at all, whereas 1 is returned if they are equal.

The Jaccard distance is debatable for various reasons. First, it does not consider the size of the gap when the intervals do not overlap. Second, J is symmetric, whereas the intuition would argue for a prevalence of the reference interval r over the hypothesized one h . Finally, $J(r, h)$ normalizes the error according to the size of the intervals. Hence, big errors on large intervals are equivalent to small errors on small intervals. This last argument can be overwhelmed by scaling J according to the size of the intervals, for instance, using the average interval size:

$$J_{\text{year}}([a, b], [c, d]) = \frac{|b - a| + |d - c|}{2} \times J([a, b], [c, d]). \quad (6)$$

¹⁰Several other variants have been studied but are not reported here since they do not bring interesting theoretical or experimental results.

¹¹An illustration of this principle is provided in Appendix A.

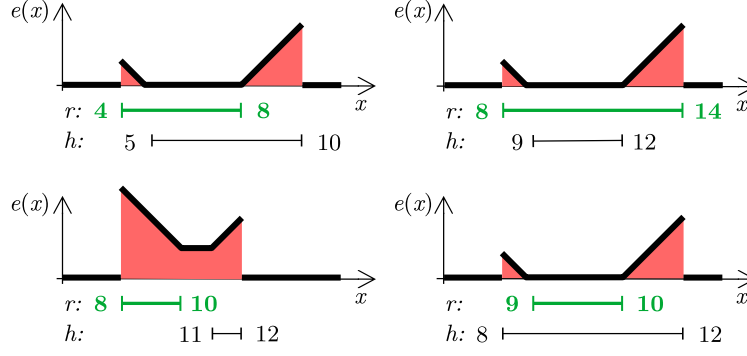


Figure 2: Evolution of the local error e and the resulting global error IE (filled areas below bold lines) for prototypical situations of the reference (r) and hypothesis (h) ranges.

5.3 Mean-based metrics

Assuming that the bounds of a given age range are somewhat uncertain, one may argue that the interval should be reduced to its center/mean value. When applying this principle on both intervals, this defines a mean-to-mean error, denoted as follows:

$$\mu E_{a,b}(c, d) = | \mu(a, b) - \mu(c, d) | , \quad (7)$$

where $\mu(x, y)$ is the mean of x and y . While this measure may look oversimplistic, we believe that it is interesting because it is easily interpretable.

This metric can be relaxed further by considering that a recommendation is acceptable (i.e., no error) as soon as the mean of the hypothesis interval lies in the reference interval. This is defined as the Bound Error (BE):

$$BE_{a,b}(c, d) = \begin{cases} 0 & \text{if } \mu(c, d) \in [a, b] , \\ a - \mu(c, d) & \text{if } \mu(c, d) < a , \\ \mu(c, d) - b & \text{if } \mu(c, d) > b . \end{cases} \quad (8)$$

5.4 Integral of local errors

In a last approach, the error between two age ranges can be seen as the sum of a local errors $e(x)$ for age x over the recommendation domain $[0, 18]$. In practice, this domain can be reduced to the smallest segment which both includes the reference and hypothesis intervals because there is no error to be counted outside of it. Hence, we define the Integral Error (IE) as:

$$IE_{a,b}(c, d) = \int_{\min(a,c)}^{\max(b,d)} e_{a,b,c,d}(x) dx . \quad (9)$$

IE is interesting since it enables a large panel of logics through the definition of different functions e .

A straightforward definition for $e(x)$ consists in measuring the distance between x and r , the distance between x and h , and to sum the two distances. The distance between x and a given interval $[s, t]$ is null if $x \in [s, t]$, and the absolute distance to the closest bound otherwise. This can be written as follows:

$$e_{a,b,c,d}(x) = \begin{cases} 0 & \text{if } x \in [a, b] \text{ and } x \in [c, d] , \\ \min(|x - c|, |x - d|) & \text{if } x \in [a, b] \text{ and } x \notin [c, d] , \\ \min(|x - a|, |x - b|) & \text{if } x \notin [a, b] \text{ and } x \in [c, d] , \\ \min(|x - c|, |x - d|) + \min(|x - a|, |x - b|) & \text{if } x \notin [a, b] \text{ and } x \notin [c, d] . \end{cases} \quad (10)$$

After integrating and applying square root to express the result as “years” (instead of “squared years”), this brings to the following Symmetric Integral Error (Sym-IE):

$$Sym-IE_{a,b}(c, d) = \sqrt{\frac{(a-c)^2 + (d-b)^2}{2}} = \frac{L_2(\vec{h} - \vec{r})}{\sqrt{2}} . \quad (11)$$

Figure 2 illustrates the behavior of $e(x)$ (solid line) and $Sym-IE$ (hatched area) on different cases.

The frame of local errors can be easily extended. For instance, Sym-IE can be generalized by weighting differently the local error due to r and due to h with a factor $\beta \in [0, 1]$:

$$e_{a,b,c,d}(x) = \begin{cases} 0 & \text{if } x \in [a, b] \text{ and } x \in [c, d], \\ \beta \min(|x - c|, |x - d|) & \text{if } x \in [a, b] \text{ and } x \notin [c, d], \\ (1 - \beta) \min(|x - a|, |x - b|) & \text{if } x \notin [a, b] \text{ and } x \in [c, d], \\ \beta \min(|x - c|, |x - d|) + (1 - \beta) \min(|x - a|, |x - b|) & \text{if } x \notin [a, b] \text{ and } x \notin [c, d], \end{cases} \quad (12)$$

A global error, called β -IE, is then defined as:

$$\beta-IE_{a,b}(c, d) = \sqrt{\beta \times (\max(0, c - a)^2 + \max(0, b - d)^2) + (1 - \beta) \times (\max(0, a - c)^2 + \max(0, d - b)^2)}. \quad (13)$$

The setting $\beta = 0.5$ corresponds to Sym-IE, while $\beta < 0.5$ gives more importance to fitting the reference interval than the hypothesis.

Other extensions could be imagined, which we believe makes this approach better than the previous ones on a theoretical point of view. For instance, one could integrate the fact that errors on early ages are more impacting than those on higher ages, since the cognitive development of children slows down as they grow up.

5.5 Study case and discussion

Among all the previously defined metrics, a good metric would behave in the same way as a human would intuitively do. This means that, given a reference age range, a good metric would rank a set of hypothesis ranges in the same order as a human. To do so, we designed a set of $n = 20$ hypothesis intervals (h_i , $1 \leq i \leq n$) with respect to a reference $r = [8, 12]$, and ranked them in an oracle way. These hypotheses reflect typical situations that one would want the metric to distinguish appropriately. Given a metric d , all hypotheses h_i can be sorted in the ascending order considering $d(r, h_i)$, and associated to their rank $\rho_d(i)$. Then, the quality of d is measured by comparing the ranking ρ_d with the oracle ranking ρ^* using the average Spearman's footrule distance S with the oracle ranking (Spearman, 1904) (the lower the better):

$$S(\rho^*, \rho_d) = \frac{1}{n} \sum_{i=1}^n |\rho^*(i) - \rho_d(i)|. \quad (14)$$

Table 3 shows the rankings for each metric¹² along with the resulting value for S . The metrics under investigation are compared to what would do a random metric (i.e., $d(r, h_i)$ is a random positive real number)¹³. Values of the parameters $\alpha = 0.5$ and $\beta = \frac{1}{3}$ for the parametric metrics are those leading to the best results for S .¹⁴

Overall, it appears that the values for S are very varied. Some metrics lead to high values. For J , this is consistent with the drawbacks previously highlighted. For BE , it appears that this comes from the fact that many hypotheses cannot be (purposely) distinguished. As shown in the tables, a linear interpolation with μE can easily alleviate, even for very small values of α . The same phenomenon is observed with Asym-IE. At the opposite, the best metrics reach very low values, which denote a strong correlation with the human judgement (β -IE and θ -L₂).

Even if the hypothesis ranges used in this experiment do not perfectly fit the expected distribution of the hypotheses on the real data, we believe that this study provides a valuable basis to rank the

¹²The corresponding values for $d(r, h_i)$ are provided in Appendix A.

¹³Let one note that the expected Spearman's footrule distance S for this random metric is an empirical approximation since no analytical form for the expectation of Spearman's footrule distance S is known in the literature.

¹⁴Eventhough this would be a methodological issue for some experiments, we judged that this optimization process is acceptable since the objective is not to provide the best parameters but to understand the way that the metrics behave.

Hypothesis	Oracle	Vector			Jaccard		Mean			Local error		Random
		L_1	L_2	$\theta\text{-}L_2$ ($\alpha=0.5$)	J	J_{year}	μE	BE	$\alpha\mu\text{E}+(1-\alpha)\text{BE}$ ($\alpha = 0.5$)	Sym-IE	$\beta\text{-IE}$ ($\beta=1/2$)	
$[c, d]$	σ	π (Δ)	π (Δ)	π (Δ)	π (Δ)	π (Δ)	π (Δ)	π (Δ)	π (Δ)	π (Δ)	π (Δ)	
8 12	1	1 (0)	1 (0)	1 (0)	1 (0)	1 (0)	1 (0)	1 (0)	1 (0)	1 (0)	1 (0)	/
8.1 12	2	2 (0)	2 (0)	2 (0)	3 (-1)	2 (0)	1 (1)	1 (1)	1 (1)	2 (0)	2 (0)	/
7.9 12	3	2 (1)	2 (1)	4 (-1)	2 (1)	3 (0)	1 (2)	1 (2)	1 (2)	2 (1)	3 (0)	/
8.5 12	4	4 (0)	4 (0)	3 (1)	6 (-2)	4 (0)	1 (3)	1 (3)	1 (3)	4 (0)	4 (0)	/
9 12	5	4 (1)	6 (1)	5 (0)	6 (-1)	4 (1)	10 (-5)	1 (4)	10 (-5)	6 (-1)	5 (0)	/
9 11	6	8 (-2)	8 (2)	6 (0)	11 (-5)	8 (-2)	1 (5)	1 (5)	1 (5)	8 (-2)	7 (-1)	/
7.5 13	6	4 (2)	4 (2)	7 (-1)	4 (2)	6 (0)	1 (5)	1 (5)	1 (5)	4 (2)	6 (0)	/
7 12	8	4 (4)	6 (2)	8 (0)	4 (4)	6 (2)	10 (-2)	1 (7)	10 (-2)	6 (2)	8 (0)	/
10 10	9	12 (-3)	12 (3)	11 (-2)	14 (-5)	12 (-3)	1 (8)	1 (8)	1 (8)	12 (-3)	11 (-2)	/
9 13	9	8 (1)	8 (1)	9 (0)	10 (-1)	9 (0)	12 (-3)	1 (8)	12 (-3)	8 (1)	9 (0)	/
7 13	11	8 (3)	8 (3)	10 (1)	8 (3)	10 (1)	1 (10)	1 (10)	1 (10)	8 (3)	10 (1)	/
6 12	12	8 (4)	11 (1)	12 (0)	8 (4)	10 (2)	12 (0)	1 (11)	12 (0)	11 (1)	12 (0)	/
8 8	13	12 (1)	14 (1)	14 (-1)	14 (-1)	12 (1)	15 (-2)	1 (12)	15 (-2)	14 (-1)	13 (0)	/
12 12	13	12 (1)	14 (1)	14 (-1)	14 (-1)	12 (1)	15 (-2)	1 (12)	15 (-2)	14 (-1)	13 (0)	/
6 14	13	12 (1)	12 (1)	13 (0)	11 (2)	17 (-4)	1 (12)	1 (12)	1 (12)	12 (1)	15 (-2)	/
7 8	16	16 (0)	16 (0)	16 (0)	14 (2)	16 (0)	17 (-1)	17 (-1)	17 (-1)	16 (0)	16 (0)	/
6 8	17	17 (0)	17 (0)	17 (0)	14 (3)	17 (0)	18 (-1)	18 (-1)	18 (-1)	17 (0)	17 (0)	/
5 7	18	18 (0)	18 (0)	18 (0)	14 (4)	17 (1)	19 (-1)	19 (-1)	19 (-1)	18 (0)	18 (0)	/
4 7	19	19 (0)	19 (0)	19 (0)	14 (5)	20 (-1)	20 (-1)	20 (-1)	20 (-1)	19 (0)	19 (0)	/
0 18	20	20 (0)	20 (0)	20 (0)	13 (7)	21 (-1)	12 (8)	1 (19)	12 (8)	20 (0)	20 (0)	/
0 0	21	21 (0)	21 (0)	21 (0)	14 (7)	12 (9)	21 (0)	21 (0)	21 (0)	21 (0)	21 (0)	/
$S \rightarrow$	0	1.1	0.9	0.4	2.9	1.4	3.4	5.9	3.4	0.9	0.3	9.7

Table 3: Ranks given by the different metrics for various hypothesis intervals $[c, d]$ according to the reference interval $[a, b] = [8, 12]$. The rank difference with the Oracle is given in parenthesis (Δ), and Spearman’s footrule distances S against the Oracle are given by the last row.

models experimented in this paper. As such, the results will be provided in terms of $\beta\text{-IE}_{\beta=0.4}$, $\theta\text{-}L_{2,\beta=0.4}$. Additionally, the results will be reported in terms of μE (error between the means) because this is easily interpretable and corresponds to the situation where only one number is expected from age recommendation, instead of a range.

6 Age Recommendation Experiments

We consider age recommendation (prediction of a target age) as a regression task. This section details the different approaches and machine learning models for age prediction from a text/sentence, and shows the experimental results.

6.1 Experiment Protocols

The age recommendation approaches studied here are (a) OneShot text-level (by considering a text as a whole), (b) sentence-level, and (c) aggregation text-level. In aggregation approach, the sentence-level recommendations are aggregated to make a text-level global recommendation. Moreover, in each approach, a text/sentence is represented with a vector of several features (as real numbers). This paper studies mainly two kinds of features: word embedding features and linguistic expert features. For the embedding features, we use two pretrained embeddings, FastText and CamemBERT. FastText is a CBOW-based word embeddings (dimension of 300) which is pre-trained on French Wikipedia using fastText (Grave et al., 2018). Unlike FastText, CamemBERT (Martin et al., 2020) is a BERT-based contextual embeddings (dimension of 768), trained on the French subcorpus of a multilingual corpus named OSCAR. However,

Feature Category	# Features
Lexicon	5
Graphemes	6
Morphosyntax	7
Verbal Tenses	24
Person/Number	5
Syntactic Dependencies	8
Logical Connectors	16
Phonetics	9
Sentiments	27
Overall	107

Table 4: Linguistic expert features of 9 categories inherited from the related study

Data	Representation with Features	Expression
Sentence	1. Avg. of token/word-vectors from FastText	$S_{FT} = \overline{V_{token}(FastText)}$
	2. Avg. of token/word-vectors from CamemBERT	$S_{CB} = \overline{V_{token}(CamemBERT)}$
	3. Sentence-level linguistic expert features	S_{Expert}
Text	1. Avg. of sentence-vectors from FastText	$T_{FT} = \overline{S_{FT}}$
	2. Sequence of sentence-vectors from CamemBERT	$T_{CB} = \{S_{CB}^1, S_{CB}^2, \dots, S_{CB}^N\}$
	3. Avg. of sentence-level linguistic expert features	$T_{Expert} = \overline{S_{Expert}}$

Table 5: Different representations of a sentence and a text to train an age-recommendation model

the linguistic expert features are a set of 107 features of 9 categories adapted from the literature of related studies. Table 4 presents the 9 categories with their feature counts¹⁵.

With the word embeddings and linguistic expert features, different representations of a sentence and a text are detailed in Table 5. Several regression models are trained with these representations of features to predict the target age from a text/sentence. Here we study the following regression models for age recommendation.

6.1.1 Regression Models

CamemBERT, a French language model relies on a RoBERTa architecture (Liu et al., 2019), which includes positional encoding and multi-head masking with attention mechanism to learn the language representation. It bidirectionally learns contextual relations from words in all positions unlike the sequential learning in RNN architecture. A pretrained CamemBERT provides also the contextual vector of a word or sentence (mean of the word-vectors in a sentence). CamemBERT offers performing single-output regression task. For age-range recommendation, given a sentence as input, two CamemBERT models are employed to recommend the lower bound and upper bound separately.

A RNN model e.g., GRU (Cho et al., 2014) are able to learn one-way long-term dependencies of a sequence. Such a model is highly relevant to text-based recommendations as a text holds the necessary properties of sequence. Although CamemBERT allows regression on long texts (maximum 512 tokens), many texts in our dataset are much larger than this limit. Therefore, for OneShot text-level age recommendation, this paper uses a GRU network fed with T_{CB} . We studied RNN models also for the sentence-level age recommendation but the results were not so good to report here.

A **Feed-forward** is a multilayer perceptron network with an input layer, hidden layer(s), and an output layer. With less parameters, it is comparatively a lighter deep learning architecture than the advanced ones e.g., BERT, transformers etc. In age recommendation task, the Feed-forward model is fed with the feature-vectors of size N representing a text/sentence, and directly outputs the two bounds of an age-range.

A **Random Forest** is an ensemble learning model to perform classification or regression task which

¹⁵Appendix B describes the feature categories in detail.

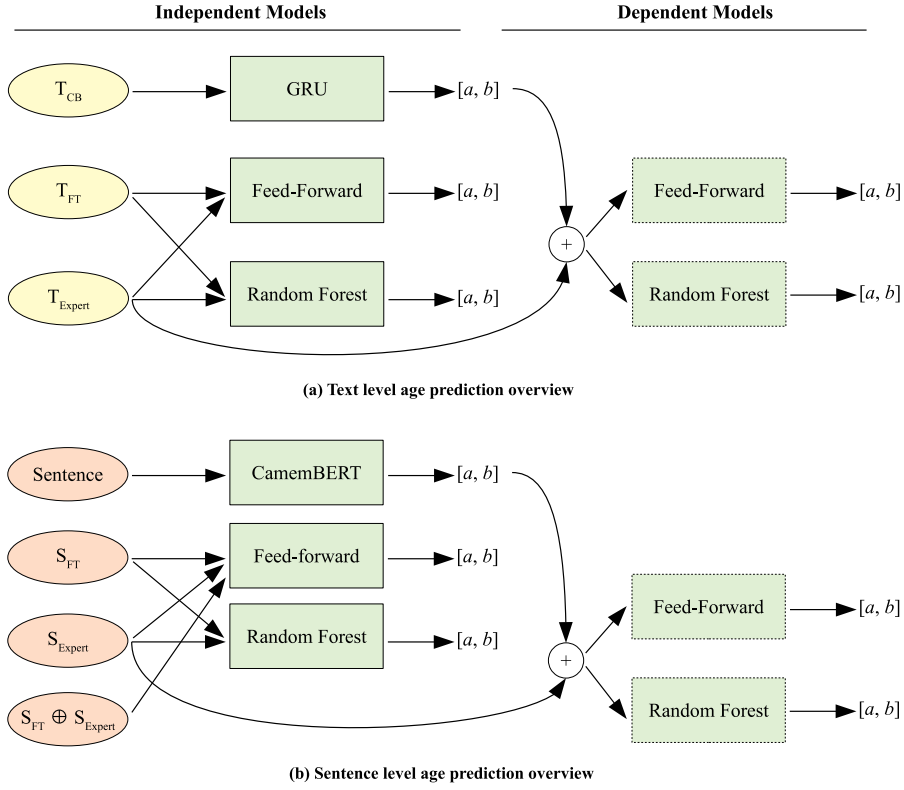


Figure 3: Overview of the text-level and sentence-level age recommendation models

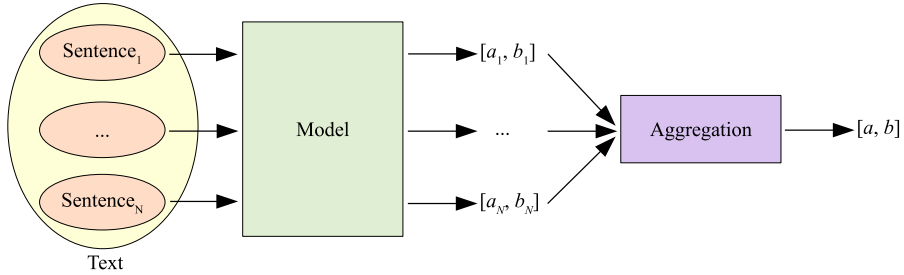


Figure 4: Text-level age recommendation by aggregating sentence-level recommendations

combines multiple decision trees to make the final prediction. A Random Forest reduces the overfitting and gets rid of the limitations of a decision tree algorithm. It offers explainability of the recommendations. We employ Random Forest regressor to directly predict the lower bound and upper bound of an age-range like the RNN and Feed-forward ones.

Combinations of the different features and regression models The regression models are fed with different types of features regarding either age recommendation from text or sentence. Figure 3 shows the overview of different regression models and their associated inputs. A model that takes embedding features and expert features as input is named as an *independent model*. In contrast, we name a model as *dependent model* when the model's input is the output of another model (in particular, GRU and CamemBERT), concatenated with the expert-features. Thus Feed-forward and Random Forest act as both independent and dependent models based on the input. However, the GRU (for text-level recommendation) and CamemBERT (for sentence-level recommendation) models are used only as the independent models.

$$grade\ level = 0.39 \left(\frac{total\ words}{total\ sentences} \right) + 11.8 \left(\frac{total\ syllables}{total\ words} \right) - 15.59 \quad (15)$$

Model	Text Level Prediction	Sentence Level Prediction
Feed-forward	hidden_layers = 6, hidden_units = 200, activation = ReLU, epoch = 500	hidden_layers = 6, hidden_units = 200, activation = ReLU, epoch = 500
Random Forest	nb_estimators = 100, criterion = 'gini', nb_leaf_node = ∞	nb_estimators = 100, criterion = 'gini', nb_leaf_node = ∞
GRU	units = 128, activation = tanh, batch_size = 128, drop_out = 0.20, epoch = 50	Not Used
CamemBERT	Not Used	max_sentence_len = 100, lr = 1e-5, batch_size = 32, epoch = 3

Table 6: Hyper-parameters used in different models for fine-tuning purpose

Sentence aggregation, Grade-level, and Naive approach We also study aggregation models to recommend the age for a text from its sentence level recommendations as shown in Figure 4. This study explores different aggregation approaches e.g., mean, median, and RNN but finds mean aggregation as the most robust and reliable. Therefore, this paper reports only the mean aggregation, $[a, b] = \frac{\sum_{i=1}^N [a_i, b_i]}{N}$, where $[a_i, b_i]$ s are recommended ages at the sentence level, and N is the number of sentences in a text. This paper further uses *Flesch-Kincaid grade-level score* (Kincaid et al., 1975) which computes a grade-level (Eq. 15) for an input text. It is a very simple model but widely studied baseline in text-readability analysis. For a text, to derive the recommended age from Flesch-Kincaid score, we use a base-age of 5.5, and add it to the grade-level. Thus, for example, regarding a text, if the grade-level is 4.5, the recommended age becomes $4.5 + 5.5 = 10$ years. As Flesch-Kincaid score returns a single grade-level, this paper considers the derived recommended age as a mean-age (μ), and it is comparable with the recommendations of other models by μ -score only. Finally we use a *Naive* model, for any text or sentence as input, it outputs the mean values of age-ranges $[9.86, 13.78]$ for text-level and $[9.73, 13.26]$ for sentence-level) computed from the training dataset. This naive approach helps us to justify the other models, either they learn well from the data or not.

6.1.2 Training and Parameter Tuning

The age recommendation models are built on the train dataset. Further, to confirm the stability, each model is fine-tuned with several hyper-parameters on the validation dataset. Table 6 details the different hyper-parameters calibrated in the models.

6.2 Age Recommendation Scores

Here we present the scores obtained from the different age recommendation models. Basically these scores refer to the errors made by the different recommendation models. To confirm the stability of the age recommendation models, this section reports the μ -score of each model on both validation and test datasets. However, the θ - L_2 and β - IE scores of each model are presented for test dataset only.

6.2.1 OneShot text-level recommendation

In OneShot text-level age recommendation, we train GRU, Feed-forward and Random Forest models with the different types of features as shown in Fig. 3(a). Table 7 shows the μE scores of different models on the validation and test datasets. All the models except Flesch-Kincaid significantly outperform the Naive approach. Flesch-Kincaid method gets the worst μE score and it is not very surprising as this method is very simple. The Naive approach obtains better score than Flesch-Kincaid method because the age-range of a large portion of the dataset is approximately same as the Naive approach computes. However, as an independent model, GRU with T_{CB} features results the best scores for the three genres Encyclopedia (E), Newspaper (N) and Fiction (F) on both validation and test datasets. With T_{FT} features, the Feed-forward model performs the second best. Moreover, Random Forest performs better than the Feed-forward model when both are fed with T_{Expert} features. Interestingly, as a dependent model, Random Forest with the combination of GRU recommended min_age, max_age, and T_{Expert} features obtains a little better score than the GRU model. Also in terms of the proposed metrics, θ - L_2 and β - IE , GRU with T_{CB} features outperforms other models as shown in Table 8. These results signify the effectiveness of T_{CB} for OneShot

Model/Features	Validation				Test			
	E	N	F	All	E	N	F	All
Naive	2.80	2.94	3.78	3.25	3.24	2.94	3.12	3.10
Flesch-Kincaid	6.22	4.97	5.88	5.76	6.92	4.61	4.78	5.23
GRU/T _{CB}	1.42	0.30	1.97	1.39	1.18	0.35	1.40	0.98
Feed-forward/T _{FT}	1.70	1.10	2.61	1.96	1.72	1.01	1.79	1.52
Feed-forward/T _{Expert}	2.31	1.95	2.71	2.42	2.15	1.90	2.51	2.22
Random Forest/T _{FT}	2.08	1.33	2.84	2.23	1.80	1.21	2.19	1.78
Random Forest/T _{Expert}	2.04	1.40	2.42	2.06	1.70	1.35	2.18	1.78
Feed-forward/(GRU_pred.+T _{Expert})*	1.72	0.79	2.26	1.72	1.42	0.85	1.70	1.34
Random Forest/(GRU_pred.+T _{Expert})*	1.33	0.23	1.95	1.34	1.05	0.26	1.37	0.91

Table 7: OneShot text-level μE scores (* indicates dependent model)

Model/Features	$\theta-L_2$				$\beta-IE$			
	E	N	F	All	E	N	F	All
Naive	5.07	4.69	5.13	4.98	3.21	2.90	3.31	3.16
GRU/T _{CB}	2.09	1.05	2.66	1.96	1.17	0.39	1.47	1.02
Feed-forward/T _{FT}	2.95	1.96	3.13	2.70	1.71	1.03	1.84	1.55
Feed-forward/T _{Expert}	3.58	3.29	4.21	3.75	2.16	1.95	2.63	2.28
Random Forest/T _{FT}	3.04	2.21	3.70	3.06	1.80	1.25	2.23	1.81
Random Forest/T _{Expert}	2.90	2.40	3.73	3.07	1.70	1.37	2.25	1.81
Feed-forward/(GRU_pred.+T _{Expert})*	2.52	1.76	3.06	2.48	1.42	0.90	1.8	1.40
Random Forest/(GRU_pred.+T _{Expert})*	1.71	0.60	2.50	1.66	1.04	0.28	1.43	0.94

Table 8: OneShot text-level $\theta-L_2$ and $\beta-IE$ scores on test dataset (* indicates dependent model)

Model/Features	Validation				Test			
	E	N	F	All	E	N	F	All
Naive	2.67	2.61	2.94	2.82	3.26	2.63	2.20	2.50
CamemBERT/-	1.93	1.3	2.21	2.00	2.02	1.28	1.94	1.83
Feed-forward/S _{FT}	2.31	1.65	2.86	2.53	2.41	1.61	2.38	2.24
Feed-forward/S _{Expert}	2.75	2.60	2.94	2.84	2.75	2.51	2.68	2.66
Feed-forward/(S _{FT} ⊕S _{Expert})	2.19	1.60	2.87	2.49	2.31	1.55	2.39	2.21
Random Forest/S _{FT}	2.42	2.15	2.43	2.39	2.68	2.02	2.01	2.15
Random Forest/S _{Expert}	2.41	2.23	2.41	2.39	2.62	2.17	2.04	2.19
Feed-forward/(CamemBERT_pred.⊕S _{Expert})*	2.08	1.38	2.71	2.34	2.08	1.34	2.33	2.09
Random Forest/(CamemBERT_pred.⊕S _{Expert})*	1.97	1.22	2.42	2.12	2.00	1.18	1.98	1.83

Table 9: Sentence level μE scores (* indicates dependent model)

text-level age recommendation. It is as expected because CamemBERT holds the contextual information of the words in a sentence, and such information helps a neural network model to extract useful hidden features.

6.2.2 Sentence-level recommendations

For the sentence level experiments, we employ a pretrained CamemBERT as a regression model by fine-tuning it with our training data. CamemBERT does not allow multi-value regression. Therefore, this paper trains two CamemBERT models separately to recommend the lower bound (minimum age) and the upper bound (maximum age) of an age range for sentence level age recommendation. Later the mean age is derived from the recommended bounds. However, Feed-forward and Random Forest models are trained to directly recommend the age range like the OneShot age recommendation.

Table 9 details the sentence level μE scores for the three genres on the validation and test datasets. In this case, Feed-forward with S_{Expert} features performs worst having a μE score higher than the

Model/Features	$\theta-L_2$				$\beta-IE$			
	E	N	F	All	E	N	F	All
Naive	5.11	4.28	3.83	4.18	3.22	2.59	2.35	2.58
CamemBERT/-	3.49	2.44	3.41	3.25	2.13	1.38	2.03	1.93
Feed-forward/ S_{FT}	3.95	2.81	4.00	3.77	2.40	1.66	2.47	2.30
Feed-forward/ S_{Expert}	4.42	4.11	4.46	4.39	2.73	2.54	2.79	2.73
Feed-forward/ $(S_{FT} \oplus S_{Expert})$	3.80	2.72	4.01	3.72	2.31	1.60	2.48	2.28
Random Forest/ S_{FT}	4.29	3.47	3.48	3.65	2.66	2.08	2.09	2.20
Random Forest/ S_{Expert}	4.22	3.67	3.56	3.72	2.62	2.21	2.14	2.25
Feed-forward/ $(CamemBERT_pred. \oplus S_{Expert})^*$	3.42	2.42	3.93	3.54	2.08	1.39	2.42	2.16
Random Forest/ $(CamemBERT_pred. \oplus S_{Expert})^*$	3.28	2.16	3.41	3.15	1.99	1.25	2.05	1.89

Table 10: Sentence-level $\theta-L_2$ and $\beta-IE$ scores on test dataset (* indicates dependent model)

Model/Features	Validation				Test			
	E	N	F	All	E	N	F	All
Naive	2.80	2.94	3.78	3.25	3.24	2.94	3.12	3.10
CamemBERT/-	1.86	1.12	1.94	1.72	1.92	1.16	1.73	1.59
Feed-forward/ S_{FT}	2.11	1.33	2.48	2.09	2.21	1.36	1.91	1.80
Feed-forward/ S_{Expert}	2.31	2.12	2.83	2.50	2.46	2.08	2.41	2.32
Feed-forward/ $(S_{FT} \oplus S_{Expert})$	2.01	1.29	2.47	2.04	2.12	1.30	1.93	1.76
Random Forest/ S_{FT}	2.45	2.14	2.84	2.55	2.67	2.07	2.42	2.37
Random Forest/ S_{Expert}	2.44	2.23	2.93	2.60	2.60	2.21	2.53	2.44

Table 11: Text level μE scores by Mean-aggregation

Naive approach. All other models perform better than the Naive one. The CamemBERT model with its contextual word embeddings achieves the best scores for each genre on both datasets. Interestingly, on Feed-forward model, S_{FT} features perform better than the S_{Expert} features. The S_{FT} and S_{Expert} features jointly improves a little over S_{Expert} features but seems not significant. Like the OneShot text-level, Random Forest with S_{Expert} features gets better score than the Feed-forward one. However, as dependent models (CamemBERT recommended min-age and max-age, concatenated with S_{Expert} features), the Feed-forward and Random Forest do not improve the scores over the CamemBERT.

We also observe the $\theta-L_2$ $\beta-IE$ scores (Table 10) for the sentence level age recommendation on the test dataset and notice the coherency of the performances by different models as seen for μE scores. These scores show that CamemBERT model significantly outperforms the other models without requiring any explicit feature engineering. However, it is not surprising that the sentence level scores are worse than the OneShot text-level scores as the training sentences are not annotated by the experts rather it is done based on the text annotations which leads to inaccurate sentence annotation.

6.2.3 From sentences to text-level recommendations

The sentence level scores still look very decent and it motivates us to explore aggregation of sentence level recommendations to make a text-level global recommendation. Here we study the mean aggregation which derives the mean value from the recommended ages of the sentences in a text.

As shown in Table 11, by significantly outperforming the Naive approach, the mean-aggregation (especially for the CamemBERT model) gives very good μE scores to recommend age for a text. CamemBERT obtains μE scores of 1.72 and 1.59 on the validation and test datasets which are 0.33 and 0.61 points worse than the best (GRU/ T_{CB}) OneShot model. Although, Feed-forward performs better than Random Forest, both get some good scores on the validation and test datasets. Like the OneShot text-level and sentence-level, here again we notice that on the Feed-forward model, S_{FT} features performs better than the S_{Expert} features. Jointly the S_{FT} and S_{Expert} features gets insignificantly a little better score over S_{FT} features alone. Similar behaviour is noticed with the $\theta-L_2$ and $\beta-IE$ scores on the test dataset as shown in Table 12.

Model/Features	$\theta-L_2$				$\beta-IE$			
	E	N	F	All	E	N	F	All
Naive	5.07	4.69	5.13	4.98	3.21	2.90	3.31	3.16
CamemBERT/-	3.27	2.18	3.07	2.83	1.97	1.21	1.78	1.64
Feed-forward/S _{FT}	3.61	2.42	3.35	3.10	2.19	1.37	1.96	1.82
Feed-forward/Expert	3.97	3.48	4.08	3.86	2.42	2.06	2.49	2.34
Feed-forward/(S _{FT} ⊕S _{Expert})	3.49	2.33	3.37	3.05	2.10	1.31	1.98	1.78
Random Forest/S _{FT}	4.26	3.46	4.08	3.92	2.63	2.04	2.48	2.38
Random Forest/S _{Expert}	4.17	3.65	4.28	4.05	2.56	2.18	2.63	2.46

Table 12: Text level $\theta-L_2$ and $\beta-IE$ scores by mean-aggregation on test dataset

Sentence Level Prediction				Text Level Prediction			
Model	μE	$\theta-L_2$	$\beta-IE$	Model	μE	$\theta-L_2$	$\beta-IE$
Naive	5.02	7.70	5.14	Naive	4.55	7.11	4.74
CamemBERT/-	3.43	5.54	3.21	Mean Aggr. (CamemBERT/-)	2.95	4.83	2.76
				OneShot (GRU/T _{CB})	1.47	2.72	1.43
Expert 1	3.28	5.23	3.27	Expert 1	2.60	4.23	2.60
Expert 2	3.25	5.28	3.33	Expert 2	3.50	5.76	3.63
Expert 3	3.18	5.28	3.26	Expert 3	2.72	4.47	2.70
Expert mean	2.90	4.76	2.93	Expert mean	2.95	4.81	2.95

Table 13: Age prediction scores (μE , $\theta-L_2$ and $\beta-IE$) for the expert annotated 80 sentences and 24 texts from the test set.

Overall, the scores by mean-aggregation are very promising and this aggregation method provides significant insights on the age-recommendation particularly for the authors view-point. This approach can be very effective for an author to incrementally verify a text, either it is suitable for a target age.

6.2.4 Comparison of recommendations per age and age-range

Up to now, [GRU/T_{CB}, Feed-forward/T_{FT}] and [CamemBERT/-, Feed-forward/S_{FT}] are the best performing two pairs of models for the OneShot text-level and sentence-level age recommendation. Here we compare their recommendation performances for each age and age-range.

Fig. 5 shows the μE scores per age (a and b), and per age-range (c and d) at the text level (a and c) and sentence level (b and d) on test set. Almost in all cases (except text level scores per age-range), the GRU/T_{CB} (text level) and CamemBERT/- (sentence level) outperform the Feed-forward/S_{FT} (text level) and Feed-forward/S_{FT} (sentence level) models.

For OneShot text-level age recommendation, GRU/T_{CB} performs better for recommending ages of 8 to 12 years (older children) and 14 to 18 years (adults). For other ages, the two models perform quite similar. However, in sentence level age recommendation, CamemBERT/- gets much better scores for child-ages than adult-ages. Overall, CamemBERT has outstanding contributions in age recommendation task.

6.3 Comparison with Experts

This section compares the recommendations from our best models with those of three psycho-linguist experts on sentences and texts. To do so, the experts were requested to annotate 20 texts from the test set and 80 sentences randomly selected from these 20 texts. The different models are applied on the same sentences and texts. Table 13 compares μE , $\theta-L_2$ and $\beta-IE$ scores measured for the Naive approach, CamemBERT and GRU models, every expert, and when averaging the expert’s recommendations.

The CamemBERT and GRU models significantly outperform the Naive approach for the sentence-level and text-level age recommendation. However, CamemBERT/- model gets high μE , $\theta-L_2$ and $\beta-IE$ scores at the sentence level but still performs quite similar to the individual experts. At the text level, CamemBERT/- with mean-aggregation obtains the same μE as that of the expert-mean. As an OneShot model, using sentence-vectors from CamemBERT, GRU/T_{CB} significantly outperforms the experts. These

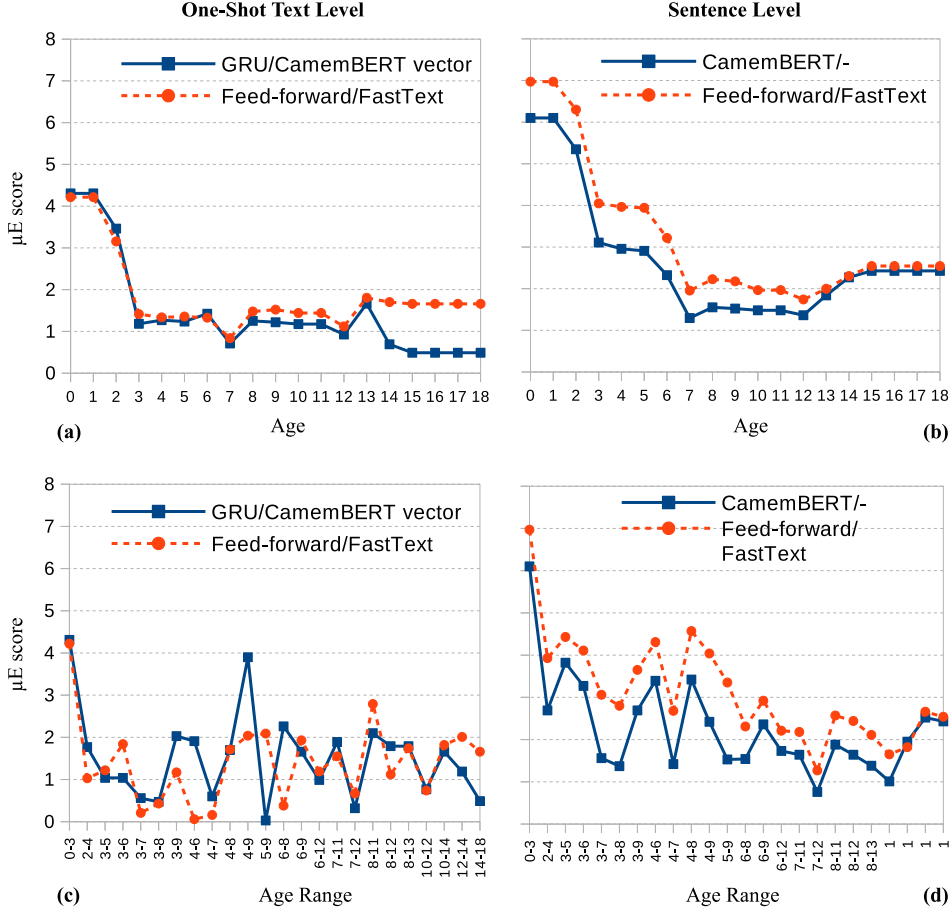


Figure 5: Per age (a and b), and per age-range (c and d) μE scores at the OneShot text level (a and c) and sentence level (b and d) on test set

Sentence	Actual Age Range / Mean	CamemBERT Range / Mean	μE	$\theta-L_2$	$\beta-IE$
1. Yum, here is a Mosquito!	[4, 8] / 6.0	[4.73, 7.39] / 6.06	0.06	0.95	0.55
2. This afternoon, Tremolo was to play at a funeral.	[8, 11] / 9.5	[8.98, 11.27] / 10.13	0.63	1.27	0.60
3. If we wait later, it is more complicated.	[12, 14] / 13.0	[7.80, 12.04] / 9.92	3.08	5.30	3.62
4. By then, you will have had time to comfort a little girl who has already shed tears on her fault.	[14, 18] / 16.0	[7.35, 11.50] / 9.43	6.57	9.80	6.60

Table 14: Examples of age recommendation by the CamemBERT/- model for some expert-annotated sentences. Here the sentences are translated in English while the French version is available in Table 20, Appendix C.

results indicate that the models built on CamemBERT perform better than human annotators. Table 14 presents some sentence level age recommendations where the sentences are annotated by experts. The CamemBERT/- model gives very good recommendation for the first two sentences while it fails very badly for the last two instances.

7 Feature Analysis for Explainability

Overall the age recommendation models obtain very decent scores. Although the *Expert* feature set (with any regressor) does not achieve the best score, it allows us to analyze which linguistic properties have higher contributions in age recommendation. As an opening of explainability of the age prediction models, this section explores feature ranking by different approaches. Before detail analysis, Table 15 describes

Feature Name	Description
SentenceLength	Number of words in a sentence.
WordLengthStd	Std. deviation of the length of each word in a sentence.
LemmaDiversity	Number of different lemmas in a sentence.
StopwordsProportion	Proportion of stop words to the total number of words in a sentence.
PhonemeNumberSentence	Number of phonemes in the sentence (phonetic length).
PhonemeDiversitySentence	Number of different phonemes in a sentence.
PhonemeNumberAvg	Average number of phonemes per word.
PhonemeNumberStd	Std. deviation of the number of phonemes per word.
PhonemeProbStd	Std. deviation of the average probability of the phonemes in a sentence.
PhonemeDiversityStd	Std. deviation of the number of different phonemes in each word.
PolarityScore	Polarity score of a sentence [-1, 1].

Table 15: Description of some the linguistic expert features

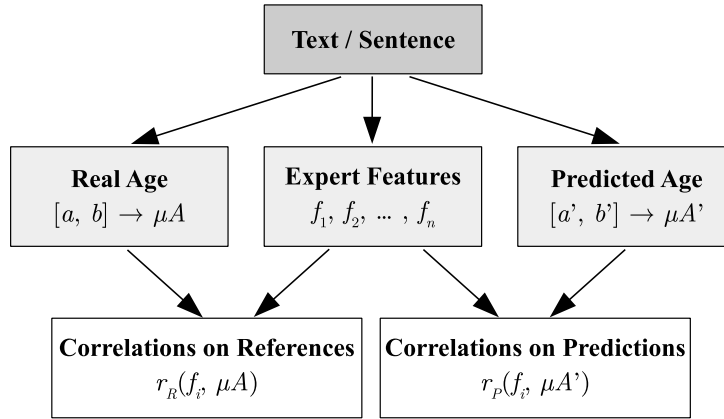


Figure 6: Correlation between expert feature and real age, and expert feature and predicted age

Features	Real Age			Predicted Age		
	Text	Sentence	Avg.*	Text (GRU/T _{CB})	Sentence (CamemBERT/-)	Avg.
PhonemeNumberAvg	0.56 / 2	0.25 / 1	0.40 / 1	0.48 / 3	0.27 / 1	0.38 / 2
PhonemeNumberStd	0.58 / 1	0.22 / 2	0.40 / 2	0.52 / 1	0.26 / 2	0.39 / 1
WordLengthStd	0.54 / 3	0.19 / 7	0.37 / 3	0.50 / 2	0.25 / 4	0.38 / 3
LemmaDiversity	0.50 / 4	0.22 / 4	0.36 / 4	0.48 / 6	0.24 / 5	0.36 / 5
SentenceLength	0.49 / 5	0.22 / 3	0.36 / 5	0.48 / 4	0.25 / 3	0.37 / 4
PhonemeNumberSentence	-0.34 / 2	-0.15 / 1	-0.25 / 1	-0.31 / 3	-0.16 / 1	-0.24 / 2
PhonemeDiversitySentence	-0.34 / 3	-0.15 / 2	-0.25 / 2	-0.31 / 4	-0.16 / 2	-0.24 / 3
PhonemeProbStd	-0.37 / 1	-0.10 / 4	-0.24 / 3	-0.39 / 1	-0.15 / 3	-0.27 / 1
PhonemesDiversityStd	-0.32 / 6	-0.14 / 3	-0.23 / 4	-0.25 / 6	-0.15 / 4	-0.20 / 4
StopWordsProportion	-0.33 / 5	-0.08 / 5	-0.20 / 5	-0.31 / 5	-0.09 / 6	-0.20 / 5

Table 16: Correlation (correlation score / feature rank) of *Expert* features to the real and predicted mean-ages at the text and sentence levels

some linguistic expert features which are presented in this section¹⁶.

7.1 Correlation Between Feature and Age

At this point, we are interested to measure the correlation (Pearson correlation coefficient) between the features and ages (on test set) to identify the most contributory features.

¹⁶Appendix B describes all the expert features.

Sentence Level		Text Level	
Feature	Score	Feature	Score
PhonemeNumberAvg	0.1021	PhonemeNumberAvg	0.1596
PhonemeNumberStd	0.0383	SentenceLength	0.1425
SentenceLength	0.0368	PhonemeNumberStd	0.0847
PhonemeProbStd	0.0132	PhonemeProbStd	0.0245
StopWordsProportion	0.0130	PolarityScore	0.0213

Table 17: Top 5 *Expert* features (with rank-scores) by permutation feature ranking at the sentence and text levels

Features	Sentence	Text
Expert	2.66	2.22
Expert - Phonetics	2.77 / +0.11	2.33 / +0.11
Expert - Dependencies	2.71 / +0.05	2.27 / +0.05
Expert - Morphosyntax	2.71 / +0.05	2.26 / +0.04
Expert - Lexicon	2.69 / +0.03	2.29 / +0.07
Expert - Sentiments	2.69 / +0.03	2.18 / -0.04
Expert - Person/Number	2.68 / +0.02	2.16 / -0.06
Expert - Connectors	2.67 / +0.01	2.21 / -0.01
Expert - VerbalTenses	2.66 / +0.00	2.19 / -0.03
Expert - Graphemes	2.65 / -0.01	2.18 / -0.04

Table 18: Ablation test ($\mu E / \Delta$) of the *Expert* features with Feed-forward model at the sentence and text levels

First, correlation is measured between each expert-feature and the real ages at the sentence and text levels, and then done similar with the predicted ages from two models, GRU/ T_{CB} (text level) and CamemBERT/- (sentence level) as shown in Fig. 6. We hypothesize that the top ranked correlated features will be similar for the real age and recommended age for a good age recommendation model. Table 16 shows the top 5 positively correlated (upper part) and top 5 negatively correlated (bottom part) features for text level and sentence level predictions. Each correlation score is presented with its rank among the features. For example, in second column (Real Age/Text), the correlation score for *PhonemeNumberStd* (0.58 / 1) is 0.58 with the rank of 1. From this table, it is very clear that the top ranked positively and negatively correlated features are common in all settings although their ranks displace a little. From the correlation based feature ranking, one can notice that among different types of features the *phonetic* features obtains the best rank scores.

7.2 Permutation Feature Ranking

The idea of permutation feature ranking is to randomly shuffle an individual feature value to break the relationship between the feature and the target value. Thus it measures the drop of performance due to discarding the feature, and finally ranks the features. Here we do it for the expert features (T_{Expert} and S_{Expert}) at the sentence level and text level on the test set by employing Random Forest regressor, a built-in library in scikit-learn.

Table 17 enlists the top 5 features for different feature combinations at the sentence and text levels. For ranking on the *Expert* features, most of the top ranked features are related to *Phonetics*. Moreover, the *SentenceLength* and *StopWordsProportion* are found very important features for age recommendation at the sentence and text levels. These features are also seen in the list of top ranked features by correlation based feature ranking in Table 16.

7.3 Ablation Test

Finally we do an ablation test with different categories of the *Expert* features on the test set. The idea of ablation test is like the permutation feature ranking where an individual feature is shuffled and its impact on the performance is measured. Here ablation test is done with the Feed-forward model, but instead of shuffling a single feature, all the features of same category are discarded.

Table 18 presents the μE scores of the ablation test where almost all the feature categories except *VerbalTenses* and *Graphemes* show positive impacts at sentence level. Interestingly, the *Phonetics*,

Dependencies, *Morphosyntax* and *Lexicon* features have positive impacts at both sentence and text levels. All the feature ranking approaches studied in this paper find the phonetics features the most common and contributory in any settings. Although these feature categories individually do not have very high impact, they altogether perform quite well in age recommendation task.

8 Conclusion and Future Work

To recommend adequate texts to the children of different ages, this paper presented a detail work on age prediction from texts/sentences. This study discussed several aspects about the definition of target age, evaluation metrics, and explored various machine learning models for age recommendation. Almost all the models perform very decently while particularly two models built on top of CamemBERT achieve the best scores, even by outperforming the human experts. The results at the text level are very good and suggest the use of the best model for real applications. On sentences, the error rates are higher, probably due to less reliable references (derived from the texts). Our preliminary study on the explainability of age recommendation models finds that the linguistic phonetic features are most prominent regarding the different difficulty levels of the texts/sentences. The dependency, morphosyntax and lexicon features also show their importance in some extents.

As an explainability task, to explain text-difficulty for different child-ages, this paper looks forward to identifying the potential words and phrases locally (instead of global context) which impact the difficulty levels in a text/sentence. In future, we also aim to work on text simplification (by paraphrasing) to offer a better reading experience to the children.

Acknowledgement

This work has been partially funded by the French National Research Agency (ANR) as part of the TextToKids project (ANR-19-CE38-0014).

References

- Renu Balyan, Kathryn S McCarthy, and Danielle S McNamara. 2020. Applying natural language processing and hierarchical machine learning approaches to text difficulty classification. *International Journal of Artificial Intelligence in Education*, pages 1–34.
- Roy Khristopher Bayot and Teresa Gonçalves. 2017. Age and gender classification of tweets using convolutional neural networks. In *International Workshop on Machine Learning, Optimization, and Big Data*, pages 337–348. Springer.
- Alexis Blandin, Gwénoél Lecorvé, Delphine Battistelli, and Aline Étienne. 2020. Age recommendation for texts. In *Proceedings of the Language Resources and Evaluation Conference (LREC)*, pages 1431–1439.
- Jasmina Byrne and Patrick Burton. 2017. Children as internet users: how can evidence better inform policy debate? *Journal of Cyber Policy*, 2(1):39–52.
- Marie Candito, Joakim Nivre, Pascal Denis, and Enrique Henestroza Anguiano. 2010. [Benchmarking of statistical dependency parsers for french](http://alpage.inria.fr/statgram/frdep/fr_stat_dep_parsing.html). In *Proceedings of the International Conference on Computational Linguistics*. http://alpage.inria.fr/statgram/frdep/fr_stat_dep_parsing.html.
- Jeanne Sternlicht Chall and Edgar Dale. 1995. *Readability revisited: The new Dale-Chall readability formula*. Brookline Books.
- Jing Chen, Long Cheng, Xi Yang, Jun Liang, Bing Quan, and Shoushan Li. 2019. Joint learning with both classification and regression models for age prediction. In *Journal of Physics: Conference Series*, volume 1168.
- Jing Chen, Shoushan Li, Bin Dai, and Guodong Zhou. 2016. Active learning for age regression in social media. In *Chinese Computational Linguistics and Natural Language Processing Based on Naturally Annotated Big Data*, pages 351–362. Springer.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](https://arxiv.org/abs/1409.1259). pages 1724–1734.

- Edgar Dale and Jeanne S Chall. 1948. A formula for predicting readability: Instructions. *Educational research bulletin*, pages 37–54.
- Denise Davidson. 2006. The role of basic, self-conscious and self-conscious evaluative emotions in children's memory and understanding of emotion. *Motivation and Emotion*, 30(3).
- Felice Dell'Orletta, Simonetta Montemagni, and Giulia Venturi. 2011. Read-it: Assessing readability of italian texts with a view to text simplification. In *Proceedings of the second workshop on speech and language processing for assistive technologies*, pages 73–83.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Duddington, Jonathan. 2014. *espeak*.
- Han Feng, Sizai Hou, Le-Yin Wei, and Ding-Xuan Zhou. 2022. Cnn models for readability of chinese texts. *Mathematical Foundations of Computing*, 5(4):351.
- Uta Frith. 1985. Beneath the surface of developmental dyslexia. In *K. E. Patterson, J. C. Marshall, & M. Coltheart (Eds.), Surface Dyslexia: Neuropsychological and Cognitive Studies of Phonological Reading*.
- Susan Gathercole. 1999. [Cognitive approaches to the development of short-term memory](#). *Trends in cognitive sciences*, 3.
- L. H. Geyer. 1977. [Recognition and confusion of the lowercase alphabet](#). *Perception & Psychophysics*, 22(5).
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- D. Gromer and M. Weiss. 1990. *Lire, tome 1 : apprendre à lire*. Armand Colin.
- Maya Hickmann. 2012. Diversité des langues et acquisition du langage: espace et temporalité chez l'enfant. *Langages*, (4).
- Zahurul Islam, Alexander Mehler, and Rashedur Rahman. 2012. Text readability classification of textbooks of a low-resource language. In *Proceedings of the Pacific Asia Conference on Language, Information, and Computation*, pages 545–553.
- Zahurul Islam, Md Rashedur Rahman, and Alexander Mehler. 2014. Readability classification of bangla texts. In *Proceedings of the International Conference on Intelligent Text Processing and Computational Linguistics (CICLing)*, pages 507–518. Springer.
- Zahurul Islam and Rashedur Rahman. 2014. Readability of bangla news articles for children. In *Proceedings of the Pacific Asia Conference on Language, Information and Computing*, pages 309–317.
- J. Peter Kincaid, Robert P. Jr Fishburne, Richard L. Rogers, and Brad S Chissom. 1975. "derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel". *Institute for Simulation and Training*.
- Shuhua Liu and Thomas Forss. 2015. Text classification models for web content filtering and online safety. In *Proceedings of the IEEE International Conference on Data Mining Workshop (ICDMW)*, pages 961–968. IEEE.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Sonia Livingstone. 2019. Eu kids online. *The international encyclopedia of media literacy*, pages 1–17.
- Steven Loria. 2018. *Textblob*.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.

- Matej Martinc, Senja Pollak, and Marko Robnik-Šikonja. 2021. Supervised and unsupervised neural approaches to text readability. *Computational Linguistics*, 47(1):141–179.
- Mohsen Mesgar and Michael Strube. 2018. A neural local coherence model for text quality assessment. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4328–4339.
- Jolien M Mouw, Linda Van Leijenhorst, Nadira Saab, Marleen S Danel, and Paul van den Broek. 2019. Contributions of emotion understanding to narrative comprehension in children and adults. *European Journal of Developmental Psychology*, 16(1).
- Dong Nguyen, Noah A Smith, and Carolyn Rose. 2011. Author age prediction from text using linear regression. In *Proceedings of the 5th ACL-HLT workshop on language technology for cultural heritage, social sciences, and humanities*, pages 115–123.
- Annie Piolat and Rachid Bannour. 2009. An example of text analysis software (emotaix-tropes) use: The influence of anxiety on expressive writing. *Current psychology letters. Behaviour, brain & cognition*, 25(2), 2009).
- NL Ranks. 2019. [List of french stop words](#).
- Gunning Robert. 1968. The technique of clear writing. *Revised Edition*. New York: McGraw Hill.
- Sarah Schwarm and Mari Ostendorf. 2005. [Reading level assessment using support vector machines and statistical language models](#).
- Charles Spearman. 1904. The proof and measurement of association between two things. *The American Journal of Psychology*, 1(15):72–101.
- Roman Suvorov, Ilya Sochenkov, and Ilya Tikhomirov. 2013. Method for pornography filtering in the web based on automatic classification and natural language processing. In *Proceedings of the International Conference on Speech and Computer*, pages 233–240. Springer.
- Valérie Tartas. 2010. Le développement de notions temporelles par l’enfant. *Développements*, 4.
- Amalia Todirascu, Thomas François, Nuria Gala, Cédric Fairon, Anne-Laure Ligozat, and Delphine Bernhard. 2013. Coherence and cohesion for the assessment of text readability. In *Proceedings of 10th International Workshop on Natural Language Processing and Cognitive Science (NLPCS 2013)*, pages 11–19.
- Łukasz Tomczyk and Kamil Kopecký. 2016. Children and youth safety on the internet: Experiences from czech republic and poland. *Telematics and Informatics*, 33(3):822–833.
- Rodrigo Wilkens, David Alfter, Xiaou Wang, Alice Pintard, Anaïs Tack, Kevin P Yancey, and Thomas François. 2022. Fabra: French aggregator-based readability assessment toolkit. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1217–1233.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). pages 38–45.

Appendix

A Details of the Evaluation Metrics

A.1 Combining distance and angle

As depicted in Figure 7, different hypotheses can lead to a same error, whereas one may expect to distinguish them. Especially when introducing the uncertainty of the reference interval, one may consider that inner errors are less severe than outer ones. For instance, based on the reference $[8, 12]$, $h_3 = [10, 10]$ is a better recommendation than $h_1 = [6, 14]$. Considering the cosine of the angle θ between of the vector $\vec{h} - \vec{r} = \begin{pmatrix} c - a \\ d - b \end{pmatrix}$ and the inner direction $\vec{v} = \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ alleviates this situation. The vector \vec{v} symbolizes the fact that it is fine if the hypothesized lower bound is greater than the one of the reference, or if the hypothesized upper bound is lower than the one of the reference. An ideal value for θ is 0 (cosine is 1) while the worst situation is $\theta = \pi$ or $-\pi$ (cosine is -1).

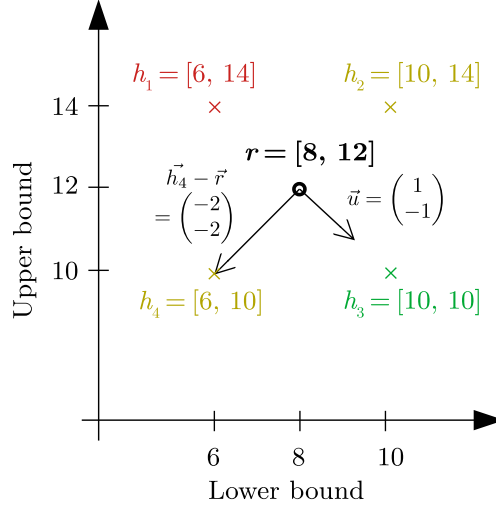


Figure 7: Examples of 4 ranges with equal L_2 errors.

Hyp.		Vector			Jaccard		Mean			Local error	
[c, d]		L_1	L_2	$\theta\text{-}L_2$ ($\alpha = 0.5$)	J	J_{year}	μE	BE	$\alpha \cdot \mu\text{E} + (1-\alpha) \cdot \text{BE}$ ($\alpha = 0.5$)	Sym-IE	$\beta\text{-IE}$ ($\beta = 1/2$)
8	12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
8.1	11.9	0.10	0.14	0.14	0.05	0.19	0.00	0.00	0.00	0.10	0.08
7.9	12.1	0.10	0.14	1.14	0.05	0.20	0.00	0.00	0.00	0.10	0.12
8.5	11.5	0.50	0.71	0.71	0.25	0.88	0.00	0.00	0.00	0.50	0.41
9	12	0.50	1.00	1.15	0.25	0.88	0.50	0.00	0.25	0.71	0.57
9	11	1.00	1.41	1.41	0.50	1.50	0.00	0.00	0.00	1.00	0.81
7.5	12.5	0.50	0.71	1.71	0.20	0.90	0.00	0.00	0.00	0.50	0.58
7	12	0.50	1.00	1.85	0.20	0.90	0.50	0.00	0.25	0.71	0.82
10	10	2.00	2.83	2.83	1.00	2.00	0.00	0.00	0.00	2.00	1.62
9	13	1.00	1.41	1.91	0.40	1.60	1.00	0.00	0.50	1.00	1.00
7	13	1.00	1.41	2.41	0.33	1.67	0.00	0.00	0.00	1.00	1.16
6	12	1.00	2.00	2.85	0.33	1.67	1.00	0.00	0.50	1.41	1.64
8	8	2.00	4.00	4.15	1.00	2.00	2.00	0.00	1.00	2.83	2.30
12	12	2.00	4.00	4.15	1.00	2.00	2.00	0.00	1.00	2.83	2.30
6	14	2.00	2.83	3.83	0.50	3.00	0.00	0.00	0.00	2.00	2.32
7	8	2.50	4.12	4.37	1.00	2.50	2.50	0.50	1.50	2.92	2.44
6	8	3.00	4.47	4.81	1.00	3.00	3.00	1.00	2.00	3.16	2.82
5	7	4.00	5.83	6.21	1.00	3.00	4.00	2.00	3.00	4.12	3.78
4	7	4.50	6.40	6.85	1.00	3.50	4.50	2.50	3.50	4.53	4.36
0	18	7.00	10.00	10.99	0.78	8.56	1.00	0.00	0.50	7.07	8.19
0	0	10.00	14.42	14.82	1.00	2.00	10.00	8.00	9.00	10.20	9.51

Table 19: Values returned by the metrics for various hypothesis intervals $[c, d]$ according to the reference interval $[a, b] = [8, 12]$.

A.2 Error values for each metric

Table 19 reports the error given by various metrics over the set of 20 samples considered in Section 5.5, sorted according to their ascending Oracle rank.

B Expert Features

The state of the art leads us to consider a list of 38 linguistic aspects for age recommendation task. As detailed below, these aspects are gathered in 9 categories, totaling a feature vector of 107 real values for each input text/sentences.

1. Lexical information (5 features)

- Mean and standard deviation of log probability of the words in French. Log probabilities have been derived from the language model for French for the speech recognition, trained on types of various types.
- Diversity of lemmas.
- Mean and standard deviation over the frequencies of the words.

2. Graphy/typography (6 features)

- Mean and standard deviation over the graphical confusability score of the words. To do so, we consider a graphical confusion score $c(x, y)$ between two graphemes x and y . Then, given a word $w = [w_1, w_2 \dots w_N]$, the confusability score is computed as the cumulative confusion between each pair of consecutive graphemes in the word, that is: $\sum_{i=1}^{N-1} c(w_i, w_{i+1})$. In practice, the confusion score c is taken from (Geyer, 1977).
- Mean and standard deviation over of the length of the words.
- Ratio of characters (including punctuation marks) against the number of words.
- Ratio of punctuation marks against the number of words.

3. Morphosyntax (7 features)

- Proportion of the following grammatical classes: verbs, state verbs, names, adjectives, clitics, temporal adverbs. Part of speech tags are generated using Bonsai (Candito et al., 2010).
- Proportion of stop words in a list of 114 words from (Ranks, 2019).

4. Verbal tenses (24 features)

- Number of different verbal tenses
- Proportions of 7 so-called simple tenses: present, past simple, future, imperfect, subjunctive present, conditional present, infinitive.
- Proportions of 7 composed tenses: compound past (passé composé), past past (passé antérieur), future past (futur antérieur), more than perfect (plus que parfait), subjunctive past, past conditional, past infinitive.
- Number of different temporal systems: past, present, future.
- Proportions of conjugated verbs for each of the 3 temporal systems: past, present, future.
- Proportion of compound tenses.
- Proportion of simple tenses.
- Proportion of each mode: infinitive, indicative, subjunctive.

5. Genders and numbers (5 features)

- Proportion of conjugated verbs in the first person.
- Proportion of conjugated verbs in the second person.
- Proportion of conjugated verbs in the third person.
- Proportion of conjugated verbs in the singular form.
- Proportion of conjugated verbs in the plural form.

6. Syntactic dependencies (8 features)

- Number of words per sentence.
- Average distances (word count) between a word and its dependencies. Dependency parsing is achieved using Bonsai (Candito et al., 2010).
- Maximum distances (word count) between a word and its dependencies.
- Mean and standard deviation of dependencies per word (words that points to a given word).
- Mean and standard deviation of the distances between each word and the words to which it points.
- Depth of the dependency tree.

7. Logical connectors (16 features)

- Proportion of logical connectors for each of the following types: addition; time; goal; cause; comparison; concession; conclusion; condition; consequence; enumeration; explanation; illustration; justification; opposition; restriction; exclusion. Since the way to gather connectors in categories, varies across papers, the categorization used is a consensus of all of them.

8. Phonetics (9 features)

- Number of phonemes in the sentence, as generated using the grapheme-to-phoneme convertor of eSpeak (Duddington, Jonathan, 2014).
- Number of different phonemes in the sentence.
- Frequency of the phonemes over the whole sentence.
- Mean and variance of the phonetic ordinariness scores of the words. The ordinariness score is computed as the average probability of appearance of each phoneme in French, as given in (Gromer and Weiss, 1990).
- Mean and variance of the word-based diversity of the phonemes.
- Mean and variance of the number of phonemes per word.

9. Sentiments/emotions (27 features)

- Score of subjectivity as used in the sentiment classifier TextBlob (Loria, 2018).
- Score of polarity (still using TextBlob).
- Proportion of words identified as trigger for a predefined set of 25 emotions: neutral, admiration, love, appeasement, daring, anger, behavior, guilt, disgust, displeasure, desire, embarrassment, empathy, pride, impassibility, inhumanity, jealousy, joy, contempt, unspecified, pride, fear, resentment, surprise, sadness.

This list is a refinement of the EMOTAIX dictionary (Piolat and Bannour, 2009).

C Results

Sentence	Actual Age Range / Mean	CamemBERT Range / Mean	μE	$\theta-L_2$	$\beta-IE$
1. Miam, voilà un moucheron!	[4, 8] / 6.0	[4.73, 7.39] / 6.06	0.06	0.95	0.55
2. Cet après-midi-là, Trémolo devait jouer à un enterrement.	[8, 11] / 9.5	[8.98, 11.27] / 10.13	0.62	1.27	0.60
3. Si on attend plus tard, c'est plus compliqué.	[12, 14] / 13.0	[7.80, 12.04] / 9.92	3.08	5.30	3.62
4. D'ici là, vous aurez eu le temps de rassurer une petite fille qui a déjà versé des larmes sur sa faute.	[14, 18] / 16.0	[7.35, 11.50] / 9.43	6.57	9.80	6.60

Table 20: Examples of age recommendation by the CamemBERT/- model for some expert-annotated sentences in French.