

GePpeTto Carves Italian into a Language Model

Lorenzo De Mattei^{•◊*†}, Michele Cafagna[†], Felice Dell’Orletta^{*}, Malvina Nissim[◊], Marco Guerini[‡]

[•]Department of Computer Science, University of Pisa, Italy

[◊]Center for Language and Cognition Groningen, University of Groningen, The Netherlands

^{*}ItaliaNLP Lab, Istituto di Linguistica Computazionale “Antonio Zampolli”, Pisa, Italy

[†]Aptus.AI, Pisa, Italy

[‡]Fondazione Bruno Kessler, Trento, Italy

lorenzo.demattei@di.unipi.it, michele@aptus.ai,

felice.dellorletta@ilc.cnr.it, m.nissim@rug.nl, guerini@fbk.eu

Abstract

In the last few years, pre-trained neural architectures have provided impressive improvements across several NLP tasks. Still, generative language models are available mainly for English. We develop GePpeTto, the first generative language model for Italian, built using the GPT-2 architecture. We provide a thorough analysis of GePpeTto’s quality by means of both an automatic and a human-based evaluation. The automatic assessment consists in (i) calculating perplexity across different genres and (ii) a profiling analysis over GePpeTto’s writing characteristics. We find that GePpeTto’s production is a sort of *bon-sai* version of human production, with shorter but yet complex sentences. Human evaluation is performed over a sentence completion task, where GePpeTto’s output is judged as natural more often than not, and much closer to the original human texts than to a simpler language model which we take as baseline.

1 Introduction

Language Models (LMs) based on pre-trained architectures such as BERT (Devlin et al., 2019) and GPT-2 (Radford et al., 2019) have provided impressive improvements across several NLP tasks. While for BERT-based architectures several monolingual models other than English have been developed, language-specific implementations of generative pre-trained transformer based models, such as GPT-2, are not widely available yet. As a contribution to fill this gap, we developed GePpeTto, the first generative language model for Italian, using the original GPT-2 as a blueprint.

Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

The evaluation of generated text is known to be intrinsically difficult (Gatt and Krahmer, 2018); we adopt here an encompassing approach, performing both automatic and human-based evaluations. The automatic assessment consists in two strategies: the first involves calculating perplexity across different language models trained on various datasets representing different genres. This serves to understand how good GePpeTto is as a language model, and how much it captures the various genres. The second one is a profiling analysis where, by means of a series of linguistic features, we capture some of GePpeTto’s writing characteristics, and compare them to those of the data it was trained on. Finally, the human evaluation is performed over a sentence completion task where GePpeTto is evaluated against gold standard sentences as well as a simple Markov-based baseline.

We make the model available to the community:

<https://github.com/LoreDema/GePpeTto>.

2 GePpeTto

GePpeTto was trained using the original settings of GPT-2 on a collection of Italian texts amounting to almost 13GB. Details on data and model’s parameters are provided in the following sections.

2.1 Data

The training set comprises two main sources. The first one is a dump of Italian Wikipedia (November 2019), consisting of 2.8GB of text. The content was extracted using the Wikiextractor tool (Attardi, 2012). The second one is the ItWac corpus (Baroni et al., 2009), which amounts to 11GB of web texts. This collection provides a mix of standard and less standard Italian, on a rather wide chronological span, with older texts than the Wikipedia

dump (the latter stretches only to the late 2000s).

Minimal processing was applied to the texts. All Wikipedia documents were prefixed by the token “Wikipedia” followed by the page’s title words. All ItWac texts were introduced by the token “Links” followed by the webpage address the text was coming from. For all texts in both collections, end of document was marked with the string `<|endoftext|>`, as done for the original GPT-2 training set (Radford et al., 2019).

2.2 Model

GePpeTto was trained on 4 GPUs Tesla T4, using the PyTorch implementation available on huggingface.¹ The model’s size corresponds to GPT-2 small, with 12 layers and 117M parameters (Radford et al., 2019). The vocabulary size is 30,000, the batch size is 32 and the block size is 100. We use 10,000 warmup steps, the adam optimiser, and an initial learning rate of $5e-5$. We train the model for 620,000 steps.

2.3 Examples

Through some selected examples, we illustrate here GePpeTto’s generation behaviour. We prompt it with different inputs, and show how the generated text is indeed style-compliant with the prompts. In the first example, we start the prompt with the ‘keyword’ “Wikipedia”, priming GePpeTto to reproduce text in the style and format of Wikipedia pages (recall that the tag “Wikipedia” was preposed to all Wikipedia training data).

prompt: “Wikipedia La Repubblica del Burmini”

“Wikipedia La Repubblica del Burmini \n\nLa Repubblica del Burmini (in francese “La Repubblica del Burginai”), originariamente repubblica di Francia, venne creata nel 1859 insieme al Regno d’Italia, con capitale Bourges, per promuovere l’unificazione dei due territori della Francia. Le città della Francia si unirono nel 1877 ed i cinque stati facenti parte dell’impero erano Francia, Italia e Savoia.”

In the second example, we prompt GePpeTto with the same text “*La Repubblica del Burmini*” but preposing to it the ‘keyword’ “Links” followed by the web address of an Italian newspaper, priming GePpeTto to

¹<https://huggingface.co>.

reproduce a text in the style of a news article. Interestingly the output is in news style but the prompt is interpreted as the header of the news, i.e., ‘place and date’.

prompt: “Links <http://www.corriere.it> La Repubblica del Burmini”

“Links <http://www.corriere.it> La Repubblica del Burmini - 14 gennaio 2003. Nota: Per gli ambiti regionali, le Regioni dovranno determinare in un solo anno le modalità di trasporto, tra gli altri, degli studenti del secondo ciclo, e le modalità per pagare eventuali supplementi a carico del servizio sanitario nazionale. Si tratta, in altre parole, di un decreto legislativo con un semplice riferimento alla Regione Lombardia.

In the following three examples we can see that even without using any specific keywords, GePpeTto’s is easily style-primed by clear prompts. In the first case, we use as input a string prompting forum-like language, in the second one, we use legal terms in the prompt, and in the last one, we prompt GePpeTto with more sophisticated language, which the model appears to readily pick up in its generated text.

prompt: “Piero DM wrote:”

“Piero DM wrote: >Una domanda sul lavoro a tempo pieno dovrebbe far sorridere.....<< Scusate, ma vorrei avere una vostra opinione.”

prompt: “Visto l’articolo 4 del dpcm:”

“Visto l’articolo 4 del dpcm 29 dicembre 1999, con il quale viene disposto l’inquadramento tra il personale interessato alle operazioni e le RSU per l’applicazione del quarto comma del predetto articolo 4, paragrafo 1, lettera a) del presente accordo,”

prompt: “l’intricato esistenzialismo”

“l’intricato esistenzialismo diventa il motore di quella riflessione metafisico-filosofica di cui il primo pensatore si avvale. L’altruismo di fondo appartiene quasi completamente alla filosofia; l’acuto verso la razionalità senza limite.

3 Automatic Evaluation

GePpeTto is trained as a language model for Italian. To assess its closeness to actual Ital-

ian texts, we calculate perplexity on a variety of sources, including a small leave out test set (1%) of GePpeTto’s training corpus (Section 3.1). In addition, we explore GePpeTto’s linguistic profile by comparing its production with human-written texts along a series of linguistic features (Section 3.2).

3.1 Perplexity

As a first evaluation, we are interested in understanding the quality of GePpeTto as a language model in its own training domain. As a second evaluation we want test its performance at zero-shot domain transfer (i.e. language modeling of a different domain). We use perplexity as a measure of language modelling performance. The different domains we consider, and the relative corpora we use, are as follows:

- own domains: Wikipedia and ItWac;
- legal domain: a corpus of Italian laws scraped from EUR-Lex² (tables excluded);
- news: a corpus of articles from the online versions of two newspapers, i.e., *la Repubblica*³ and *Il Giornale*⁴ (De Mattei et al., 2020);
- social media: a corpus of forum comments (Maslennikova et al., 2019).

To compute the perplexity scores (Table 1) we used a random sample of 4M tokens for each corpus. As expected, GePpeTto performs better on its own domains. Although ItWac is four times bigger than Wikipedia, the lower performance on the former might be due to ItWac being open domain with a large diversity of styles, while Wikipedia is more ‘standardised’. Consistently with this hypothesis, we observe a similar trend in ‘out-of-domain’ testing, where GePpeTto performs better on domains with a well coded style, namely legal documents. On domains with less coded styles, such as news and especially forum comments, we observe a performance drop.

If we compare perplexity scores with the original English GPT-2 small model, we see that GePpeTto’s results are slightly worse on the own domain corpora, which could be due to the smaller size of the training set. Out-of-domain perplexity scores are comparable between the two models.

3.2 Linguistic Profiling

For our second evaluation, we used Profiling-UD (Brunato et al., 2020), a tool for the automatic

²<https://eur-lex.europa.eu/>

³<https://www.repubblica.it>

⁴<https://www.ilgiornale.it/>

DOMAIN	PERPLEXITY
Wikipedia	26.4910
ItWac	30.9698
Legal	39.6087
News	48.3468
Social Media	131.3812

Table 1: Perplexity of GePpeTto over several in-domain and out-of-domain corpora.

Feature	Original		GePpeTto	
	μ	std	μ	std
CPT	4.809	0.959	4.750	1.127
TPS	32.302	28.322	20.382	11.127
TPC	12.393	11.504	10.711	8.529
LL _{max}	13.290	13.370	8.922	6.112
LL _{avg}	2.555	1.002	2.373	0.676

Table 2: Main linguistic features considered in our analysis. CPT = chars per token, TPS = token per sentence, TPC = tokens per clause, LL = links length.

analysis of texts that extracts several linguistic features of varying complexity. These features range from raw text properties, such as average length of words and sentences, to lexical, morpho-syntactic, and syntactic properties, such as part-of-speech (POS) distribution and inflectional properties of verbs. More complex aspects of sentence structure are derived from syntactic annotation, and model global and local properties of parsed tree structure, such as the order of subjects/objects with respect to the verb, the distribution of syntactic relations, and the use of subordination.

In our analysis we focus on two macro aspects of GePpeTto’s output, namely lexical complexity and syntactic complexity, and compare them to human productions. To do so, we rely on a selection of Profiling-UD’s features which we use as proxies for the macro-aspects that we consider.

We run the profiling analysis on a sample of both gold and generated texts. For gold, we randomly sample the test set for a total of about 19k sentences. For GePpeTto, we pick the first token from each of the 19k gold sentences, and use it as a prompt to the model. We profile these generated texts.

Lexical complexity. We proxy lexical complexity with the number of characters per word, overall frequency of tokens, also with reference to an ex-

ternal dictionary, and POS distribution.

The number of characters per token (CPT), which indicates whether shorter (usually more common) or longer (usually more complex/specialised) words are used, is completely comparable across the original (4.80, std=0.96) and GePpeTto’s (4.75, std=1.13) language models – see Table 2. This suggests that the complexity of the used vocabulary is not that different.

We compute a reference dictionary of token frequency on ItWac (≈ 1.5 billion tokens), and compare observed token frequency in both gold and generated text to this reference. We observe that in gold sentences, each token has a probability of 0.912 to be in the top 5‰ of most frequent tokens. In the generated sentences, the probability grows to 0.935, suggesting that GePpeTto is more likely to use more frequent words rather than rarer ones. This observation is in line with previous research which showed that for Nucleus Sampled texts, such as those produced by GPT-2, all tokens come from the top-p%, since the long tail is cut off, while for human produced texts, the probability of all tokens being drawn from the top-p% of the language distribution goes to zero as document length increases (Gehrmann et al., 2019; Zellers et al., 2019).

Regarding POS distribution, we observe that while for most POS tags usage is comparable, for a few others the two language models differ. The latter are, specifically, auxiliaries and proper nouns, which GePpeTto tends to overgenerate in comparison to the original model, and adjectives, which GePpeTto instead uses less than in the original texts. This is seen also for nouns and verbs, but the differences are relatively minimal. Conjunctions are also overall less frequent in GePpeTto. A detailed table will be included in the final version.

Syntactic complexity. At the level of syntax, we proxy complexity by the number of tokens per sentence, and the number of tokens per clause. We also look at the length of a dependency link, that is calculated as the number of words occurring linearly between the syntactic head and its dependent (excluding punctuation dependencies). The value associated with this feature corresponds to the average value extracted for all dependencies in a text. This information is complemented with the feature *Maximum dependency link* corresponding to the longest dependency link for each sentence.

When comparing the number of tokens per sentence (TPS, Table 2), we see that it’s much lower

for GePpeTto’s production rather than for human texts (20.4 tokens per sentence on average for GePpeTto vs 32.3 for gold texts), indicating that GePpeTto generates shorter sentences. Contextually, we also observe that GePpeTto’s generated sentences exhibit less variation in length (smaller STD) than human sentences (larger STD).

The difference in number of tokens at the clause level is relatively smaller, with clauses of length 12.4 in human texts vs 10.7 in GePpeTto (TPC, see Table 2). Considering that a clause is proxied by the presence of a verbal/copular head, it seems that sentences produced by GePpeTto, though shorter, are similar in complexity given the proportional distribution of verbal heads.

The above values taken together might suggest that while complexity at the macro level (sentence length) is higher for natural sentences, at the micro level (clause length) complexity of GePpeTto’s generations and human texts is more similar. While this intuition will require further linguistic analysis, observing the length of syntactic links seems to support it. This feature proxies quite well syntactic complexity, since it indicates how maximally far (and how far on average) a dependent and its head are within a sentence. Both the maximum length and the average length are higher for human texts (LL_{max} and LL_{avg} , see Table 2). However, if we look at them proportionally to sentence length, we find that they are comparable: normalising the longest link by the number of tokens per sentence (LL_{max}/TPS), we obtain similar values for gold (0.411) and for GePpeTto (0.438). This suggests that GePpeTto produces somewhat shorter sentences, but their internal complexity relatively corresponds to the internal complexity of the longer sentences produced by humans.

4 Human evaluation

We also test GePpeTto’s ability to generate Italian texts through a sentence completion task. The automatically generated sentences are presented to human subjects for evaluation on perceived naturalness and compared to gold ones and to a baseline.

While the original (gold) texts represent an upperbound for GePpeTto, we do not actually have a lowerbound against which the quality of GePpeTto can be assessed. To provide a comparison, we train a simple Markov model that would be able to generate text and use it as our baseline. Since the size of a Markov model dra-

matically grows with its vocabulary size, we use 1 million randomly sampled sentences from the same training-set used for GePpeTto. We train a Markov chain generator using the `markovify`⁵ implementation with state size 2, then we generate synthetic texts starting from the last 2 tokens of same prompts used for GePpeTto.

4.1 Tasks

Human subjects are asked to perform two evaluation tasks. One is a comparative ranking task, where subjects are asked to rank three portions of text (produced by gold, GePpeTto, baseline) according to perceived naturalness. The other is a classification task, where subjects are asked to tell, according to their intuition, if a portion of text, seen in isolation, is automatically generated (*yes*, *no*, *can't tell*).

Experimental design. The experiment includes 12 conditions of the stimulus material in a 4x3 design. One level (A) with three conditions is given by {gold, GePpeTto, baseline}. The second level (B) is the `prompt+completion` combination that results in 4 conditions {5+5, 5+10, 10+5, 10+10}. We use 100 different prompts (randomly selected gold sentences truncated at 5 and 10 tokens). Each of the 100 prompts enters each of the 12 conditions of the 4x3 design, for a total of 12 different stimuli. Basically, each 5 or 10 tokens prompt is completed with 5 or 10 tokens coming either from gold, GePpeTto, or the baseline model. Table 3 shows an example of all the stimuli deriving from the same 5- or 10-token prompt.

Each subject is assigned either to the ranking or to the classification task.

In ranking, we opt for a between subject evaluation set up by assigning each subject to one of the (B) conditions and offer the three versions of (A) to be ranked. For example, one subject is asked to evaluate all the 100 prompts in the 5+5 configuration (dimension B) for the three realisations, i.e., gold, GePpeTto, and baseline (dimension A).

For the classification experiments, we again opt for a between subject evaluation set up, this time by assigning each subject to one of the 12 conditions, randomly picked up for each prompt. In other words, we make sure that each subject is exposed to only one completion per prompt, randomising prompt order. By seeing only one (out of 12) realisation per prompt, each subject sees a

given prompt only once and we can therefore avoid cross-comparison effects of different completions of the same prompt, which could otherwise potentially lead again to an implicit ranking task.

Material. The materials are prepared as follows: we have selected 100 random documents/sentences and have cut them at their 5 first tokens and also their 10 first tokens. Each 5-token and 10-token prompt was given to GePpeTto and baseline so that the models could continue the text.

For each prompt, we obtain one single generated text by the two automatic models and chop them at 5 or at 10 tokens. In other words, each chopped version is derived from the same generated output which is just cut at different lengths.

We cut the sentences (including the original one) to control for the effect of text length. Indeed, we observed in Section 3.2 that GePpeTto generates shorter sentences than humans, which could represent a strong bias in evaluation. In Table 3, we show examples of all the possible stimulus material configurations according to the `prompt+completion` conditions of level (B).

Instructions and subjects. For both the ranking and classification experiments, subjects were told that they will have to evaluate excerpts of text along a ‘more natural vs. more artificial’ dimension. All stimuli used in both scenarios are the same.

For the ranking scenario, subjects were asked to “*rank the given examples from the most natural to the most artificial*”, where the inputs are three texts (gold, GePpeTto, baseline), all starting with the same prompt, thus the same five or ten tokens.

For the classification scenario, subjects saw instead the portions of text in isolation, and could answer *yes*, *no*, or *can't tell* to the question “*according to your intuition is this sentence written by an artificial intelligence?*”.

A total of 24 unique subjects (12 females) carried out the tasks using Google Forms. Twelve subjects (6 females) were assigned to Task 1 and the others to Task 2. Each subject evaluated 100 cases, and each case was evaluated by three different subjects.

4.2 Results

First, we discuss the results of our human evaluation separately, with observations related to the ranking task and observations related to the classification task. Subsequently, we knit together the two outcomes to draw a wider picture of how humans assess the quality of GePpeTto’s output.

⁵<https://github.com/jsvine/markovify>.

5 token prompt: Mentre per quanto riguarda gli 10 token prompt: Mentre per quanto riguarda gli accordi per la fornitura di	
Gold	
5+5	Mentre per quanto riguarda gli accordi per la fornitura di
5+10	Mentre per quanto riguarda gli accordi per la fornitura di latte, in scadenza questa
10+5	Mentre per quanto riguarda gli accordi per la fornitura di latte, in scadenza questa
10+10	Mentre per quanto riguarda gli accordi per la fornitura di latte, in scadenza questa settimana, Alemanno ha detto
GePpeTto	
5+5	Mentre per quanto riguarda gli emendamenti, fa presente che il
5+10	Mentre per quanto riguarda gli emendamenti, fa presente che il suo gruppo non ha sottoscritto
10+5	Mentre per quanto riguarda gli accordi per la fornitura di beni e servizi, i fatti
10+10	Mentre per quanto riguarda gli accordi per la fornitura di beni e servizi, i fatti in suo possesso hanno come
Markov-based baseline	
5+5	Mentre per quanto riguarda gli aspetti più significativi del mondo
5+10	Mentre per quanto riguarda gli aspetti più significativi del mondo editoriali, con priorità di sviluppo
10+5	Mentre per quanto riguarda gli accordi per la fornitura di biciclette elettriche a 48 bit
10+10	Mentre per quanto riguarda gli accordi per la fornitura di biciclette elettriche a 48 bit (281,5 trilioni di operazioni e

Table 3: Example outputs (stimuli) for different prompt lengths of the same original sentence.

Ranking Overall, results show that the most frequently chosen completion is the gold one, followed by GePpeTto and then the Markov baseline, but the baseline is far more distant from GePpeTto than GePpeTto from gold (Figure 1). If we look at results in more detail (see Table 4), based on the variable that we have considered in the experimental set up, namely length of input and continuation as well as overall sentence length, we observe that the order of preference for gold is 10+10, then 5+10, then 10+5, and lastly 5+5, while for the automatic models the order is 5+5, 10+5, 5+10, and then 10+10, suggesting the following.

First, the shortest the sentence, the hardest it is to discriminate between gold and generated text; indeed, the 5+5 condition is the one that results best for the two models and worst for gold.

Second, when the sentence is the longest (10+10), it is easiest for the subjects to discriminate the gold from the generated sentences. It is also interesting to note that in this condition we observe the largest gap between the two generation models, with GePpeTto getting ranked higher than Markov more than in the other conditions.




Figure 1: Ranking results for the three models

Third, at equal sentence length (15 tokens) the situation is a bit more fuzzy, but we can observe a slight tendency where it is easier to spot as automatically generated the 5+10 rather than 10+5 cases. This, in combination with the previous observation, seems to imply that the longer the generated text, the easier it is to figure out which texts are automatically produced, which makes sense, since there is more ‘space’ for the models to make mistakes.

model	5+5			5+10			10+5			10+10		
	1 st	2 nd	3 rd	1 st	2 nd	3 rd	1 st	2 nd	3 rd	1 st	2 nd	3 rd
Gold	54	30	16	62	31	7	60	27	13	70	21	9
GePpeTto	34	43	23	30	46	24	33	43	24	23	59	18
Markov	12	27	61	8	23	69	7	30	63	7	20	73

Table 4: Percentages of ranking results according to the various stimulus material conditions.




Figure 2: Classification results for the three models

Classification Overall, results show that across all conditions, gold sentences are most often rightly identified as not automatically generated (68% of “no” to the question whether the output was produced by an artificial intelligence), followed by GePpeTto (54%), and lastly by the Markov baseline (26%), indicating, as expected, that the latter produces the least natural outputs. Figure 2 reports the distribution over the various answers. Also in this case the distance between GePpeTto and gold is lower than GePpeTto and the baseline (double in percentage points), indicating that the production of GePpeTto is approaching natural language. It is also interesting to see that the highest percentage of “can’t tell” is recorded for GePpeTto, meaning that for this model it was harder than for baseline and gold to decide whether the text was automatic or not.

Let us look at results in more detail (Table 5), focusing again on length of input and continuation. Regarding continuation, we observe that *+5 conditions are better than *+10 conditions for both automatic models, indicating that the least generated text, the more natural the fragment is perceived.

Regarding input length, we see that for GePpeTto a longer prompt yields better results (10+5 is better than 5+5, and 10+10 is better than

5+10). With 10-token prompts, GePpeTto generates text that is (i) assessed as natural as much as the original text when completed with 5 tokens (62% GePpeTto, 63% original), and (ii) judged as natural 50% of the times when completed with 10 tokens. This seems to suggest that a longer input context is beneficial to GePpeTto when completion size is kept constant. However, we may wonder whether GePpeTto is evaluated as more natural because the generated text is actually better given the more context to start with, or simply because there is more gold text in the stimulus. If it were just for the contribution of a longer gold portion in the stimulus, we should see a similar behaviour for the baseline. Instead, we see that prompt size doesn’t matter for the baseline, at least for the 5 token completion case (33% in both 5+5 and 10+5). In the 10-completions (5+10 and 10+10), the larger amount of gold data in the stimulus probably does alleviate a little the very low naturalness induced by the generated text. While we can tentatively postulate that GePpeTto generates better text when more input is provided, further investigation is required to provide more solid evidence.

Summary of Results. Intersecting the observations from the two experimental setups provides us with a complete picture. In ranking (thus when the models are directly compared), both GePpeTto and the baseline perform best in the 5+5 and 10+5 conditions, suggesting that automatic generation can easily be spotted when compared side by side with human text. In other words, the least generated material, the better.

However, looking at classification, where each textual material is evaluated in isolation, we see that the two models behave in fact very differently. First, there is a much larger proportion of cases produced by GePpeTto that are deemed “natural” (54%) compared to Markov (26%). Second, the margin of uncertainty when judging GePpeTto is higher than for the baseline and

model	5+5			5+10			10+5			10+10		
	yes	no	ct	yes	no	ct	yes	no	ct	yes	no	ct
Gold	26	66	8	27	68	5	32	63	5	28	71	1
GePpeTto	32	55	13	48	46	6	32	62	6	42	50	8
Markov	62	33	5	80	13	7	61	33	6	71	19	10

Table 5: Percentages of classification results according to the various stimulus material conditions. Is the text automatically generated? {yes, no, can’t tell (ct)}.

for original text. Lastly, given the same completion size, GePpeTto performs better when its prompt is longer. Whether this is an effect of a larger proportion of gold data in the stimulus or it has to do with providing the model with a larger input context is left to future investigation.

5 Conclusion

GePpeTto is the first GPT-2-based language model for Italian. Through both automatic and manual evaluation we assessed its quality on a variety of texts and in comparison to gold data as well as another statistical generation model. Results show that GePpeTto is able to produce text which is much closer to human quality rather than to the text generated by the other generation model we have used. Linguistic analysis also highlights that GePpeTto’s production is quite similar to human production, though in a sort of bonsai version, since its sentences are on average shorter than the original texts, but with similar complexity.

The availability of GePpeTto opens up substantial possibilities. In the same way that GPT-2 is changing the approach to several NLP English tasks, we can expect GePpeTto to serve a similar purpose in Italian language processing.

References

Giuseppe Attardi. 2012. Wikiextractor. <http://attardi.github.io/wikiextractor>.

Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. 2009. The WaCky wide web: a collection of very large linguistically processed web-crawled corpora. *Language resources and evaluation*, 43(3):209–226.

Dominique Brunato, Andrea Cimino, Felice Dell’Orletta, Simonetta Montemagni, and Giulia Venturi. 2020. Profiling-UD: a Tool for Linguistic Profiling of Texts. In *Proceedings of the Twelfth International Conference on Language*

Resources and Evaluation (LREC 2020), Marseille, France. European Language Resources Association (ELRA).

Lorenzo De Mattei, Michele Cafagna, Felice Dell’Orletta, and Malvina Nissim. 2020. Invisible to People but not to Machines: Evaluation of Style-aware Headline Generation in Absence of Reliable Human Judgment. In *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC 2020)*, Marseille, France. European Language Resources Association (ELRA).

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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Albert Gatt and Emiel Krahmer. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61:65–170.

Sebastian Gehrmann, Hendrik Strobelt, and Alexander M Rush. 2019. Gltr: Statistical detection and visualization of generated text. *arXiv preprint arXiv:1906.04043*.

Aleksandra Maslennikova, Paolo Labruna, Andrea Cimino, and Felice Dell’Orletta. 2019. Quanti anni hai? Age Identification for Italian. In *Proceedings of the Sixth Italian Conference on Computational Linguistics (CLiC-it 2019)*, Bari, Italy. CEUR Proceedings 2481.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. In *Advances in Neural Information Processing Systems*, pages 9051–9062.