

# Camoscio: an Italian Instruction-tuned LLaMA

Andrea Santilli<sup>1</sup>, Emanuele Rodolà<sup>1</sup>

<sup>1</sup>*Sapienza University of Rome - Computer Science Department*

## Abstract

In recent years Large Language Models have improved the state of the art on several natural language processing tasks. However, their availability is frequently restricted to paid API services, posing challenges for researchers in conducting extensive investigations. On the other hand, while some open-source models have been proposed by the community, they are typically English-centric or multilingual without a specific adaptation for the Italian language. In an effort to democratize the available and open resources for the Italian language, in this paper we introduce Camoscio: a language model specifically tuned to follow users' prompts in Italian. Specifically, we finetuned the smallest variant of LLaMA (7b) with LoRA on a corpus of instruction prompts translated to Italian via ChatGPT. Results indicate that the model's zero-shot performance on various downstream tasks in Italian competes favorably with existing models specifically finetuned for those tasks. All the artifacts (code, dataset, model) are released to the community at the following url: <https://github.com/teelinsan/camoscio>

## Keywords

Large Language Models, Instruction-tuned Models, Resources for the Italian Language

## 1. Introduction

In recent years, Large Language Models (LLMs) have made remarkable advancements in the field of natural language processing, demonstrating state-of-the-art performance on various tasks [1, 2, 3]. However, the majority of these models are typically controlled by for-profit organizations that release just a paid API for receiving responses based on input textual prompts. This severely constrains researchers from conducting comprehensive and meaningful research, as they lack access to both the model's weights and the training data regime. This limitation is particularly relevant for privacy-sensitive applications (e.g., medical domain) where data cannot be shared with external providers.

On the other hand, several open-source models<sup>1</sup> have been proposed as an alternative to closed models [4, 5, 6]. However, most of these models are English-centric or multilingual, albeit with performance that lags behind their monolingual counterparts. Furthermore, in these latter models, support for the Italian language is usually poor. For example, BLOOM – the largest open multilingual model available up to date – has not been trained on any Italian data, while LLaMA has only a small percentage of training data in the Italian language<sup>2</sup>. In addition to this, most of these models are only trained with the standard language modeling objective (i.e., predict the next token given the previous ones) on corpora

of raw textual data, while it has been shown that a second training step of instruction-tuning is crucial to increase downstream performance [7, 8, 9]. Recently, a step in this direction has been made by Taori et al. [10] with the release of Stanford Alpaca, an instruction-tuned version of LLaMA for the English language. Following this approach, in this paper we propose Camoscio as an instruction-tuned version of LLaMA for the Italian language by translating to Italian the instruction-tuning dataset of Stanford Alpaca. In particular, we finetuned the smallest version of LLaMA (7 billion parameters) with LoRA [11], a parameter-efficient finetuning technique that allows to train larger models on standard desktop hardware.

Our contributions are the following:

- We introduce an instruction-tuning dataset for the Italian language, stemming from the Stanford Alpaca [10] dataset, translating it to Italian.
- We train Camoscio on this dataset and evaluate its zero-shot performance on several downstream tasks for the Italian language (NewsSum-IT, SQuAD-IT, XFORMAL IT).
- We release all the artifacts (code, dataset, model checkpoints) to the community.

## 2. Background

Large language models have emerged as a general class of models capable of performing a wide range of tasks without explicit finetuning by just leveraging in-context examples [12]. They've garnered popularity not only in the natural language processing domain but also across audio, image, and multimodal domains [13, 14, 15], with most of the approaches scaling or optimizing their performance [2, 16].

*CLiC-it 2023: 9th Italian Conference on Computational Linguistics, Nov 30 – Dec 02, 2023, Venice, Italy*

✉ [santilli@di.uniroma1.it](mailto:santilli@di.uniroma1.it) (A. Santilli); [rodola@di.uniroma1.it](mailto:rodola@di.uniroma1.it) (E. Rodolà)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).  
CEUR Workshop Proceedings (CEUR-WS.org)

<sup>1</sup>Actual openness depends on the model license.

<sup>2</sup>Less than 4.5% of training data comes from Wikipedia in 20 different languages, including Italian.



**Table 1**

Results on SQuAD-IT. All the models are trained on the SQuAD-IT training set, except for Camoscio which is evaluated in a zero-shot fashion. The additional evaluation metric *Exact Match via ChatGPT* is highlighted in grey. The scores F1 and EM for competitor models are reported from their respective papers.

	SQuAD-IT						
	F1	EM	EM-GPT	R1	R2	RL	BS
DrQA-IT [32]	.659	.561	-	-	-	-	-
mBERT [35]	.760	.650	-	-	-	-	-
BERT <sup>3</sup> [18]	.753	.638	-	-	-	-	-
MiniLM [36]	.720	.577	-	-	-	-	-
MiniLM <sub>+st</sub> [36]	.745	.620	-	-	-	-	-
XLM-R Large <sub>+st</sub> [36]	.804	.676	-	-	-	-	-
mT5 Small [22]	.660	.560	.684	.617	.347	.617	.712
mT5 Base [22]	.757	.663	.745	.709	.396	.708	.770
IT5 Small [22]	.716	.619	.602	.671	.372	.671	.743
IT5 Base [22]	.761	.663	.600	.712	.406	.712	.770
IT5 Large [22]	.780	.691	.641	.730	.412	.729	.784
Camoscio-7b (0-shot)	.270	.077	.576	.242	.133	.241	.237

At inference time, the same prompt is used to generate the answer. Only the text generated after “[...] ### Risposta:” is used as final output. We sample from the model using *top-p* sampling [37] with a temperature of 0.2,  $p = 0.75$ ,  $k = 40$ , and beam search with 4 beams.

We refer to Appendix A for the additional implementation details.

## 4. Experiments

Currently, there is a very limited availability of datasets for a solid evaluation of the broad capabilities these general-purpose models possess. This is true for English but especially for the Italian language, although the community is moving towards this direction [38]. To evaluate our model we decided to follow the same evaluation protocol proposed in Sarti and Nissim [22]. Compared to their approach, we do not perform any training on the downstream tasks, i.e., we perform just the evaluation on the test set in a zero-shot fashion by providing to the model a textual description of the task (e.g., “*Riassumi il seguente articolo*”). We compared the performance of our model on standard Italian benchmarks for summarization (NewsSum-IT), question answering (SQuAD-IT), and style transfer (XFORMAL IT).

Compared to Sarti and Nissim [22], we do not include the Wikipedia for Italian Text Summarization (WITS) corpus [39] since Wikipedia is included in the original training corpus of LLaMA [6]. We also omitted the news style transfer task between “Il Giornale” to “La Repubblica” (and vice-versa) based on CHANGE-IT [40], since Camoscio has no concepts of “Il Giornale” or “La Repubblica” styles (i.e., it was never exposed during training

or finetuning to this kind of articles, although we recognize it might be interesting to analyze this in a few-shot setting). We describe in the next paragraphs the three datasets used for the evaluation.

**News Summarization.** We evaluate the news article summarization capabilities of Camoscio using the dataset NewsSum-IT proposed by Sarti and Nissim [22]. This dataset is obtained by merging two newspaper sources (“Fanpage.it” and “Il Post”) scraped by the Applied Recognition Technology Laboratory<sup>4</sup> and available on the Hugging Face Hub [41]. We used only the test split for the zero-shot evaluation and asked the model to generate an answer given the instruction “*Dopo aver letto il testo qui sotto, riassumilo adeguatamente.*” provided in the textual prompt and the news text provided as input (complete prompt as explained in §3.2). We use the same evaluation metrics of Sarti and Nissim [22] and report the average across the two newspapers as in their work.

**Question Answering.** To assess the model performance on extractive question answering, we used the SQuAD-IT dataset [32]. This dataset is composed of sets of paragraphs, questions, and answers derived from the original SQuAD dataset [42] via machine translation and subsequent filtering of problematic instances. As for the previous datasets, we used just the test split for zero-shot evaluation. The model is asked to generate an answer given the instruction “*Dopo aver letto il paragrafo qui sotto, rispondi correttamente alla successiva domanda.*” We

<sup>3</sup><https://huggingface.co/antoniocappiello/bert-base-italian-uncased-squad-it>

<sup>4</sup><https://huggingface.co/ARTELab>

**Table 2**

Results on formality style transfer (XFORMAL IT) for the formal-to-informal (F  $\rightarrow$  I) and informal-to-formal (I  $\rightarrow$  F) directions. Competitors’ scores reported from Sarti and Nissim (2022).

	XFORMAL (IT) F $\rightarrow$ I				XFORMAL (IT) I $\rightarrow$ F			
	R1	R2	RL	BS	R1	R2	RL	BS
mT5 Small	.651	.450	.631	.666	.638	.446	.620	.684
mT5 Base	.653	.449	.632	.667	.661	.471	.642	.712
IT5 Small	.650	.450	.631	.663	.646	.451	.628	.702
IT5 Base	.652	.446	.632	.665	.583	.403	.561	.641
IT5 Large	.611	.409	.586	.613	.663	.477	.645	.714
Camoscio-7b (0-shot)	.645	.436	.623	.651	.622	.428	.600	.667

evaluated the generated answers using the script from Sarti and Nissim [22]. Furthermore, we also used an additional metric “ChatGPT Exact Match” to better assess the performance. We explain this metric in the following subsection “Evaluation Metrics”.

**Formality Style Transfer.** We assess the style transfer capabilities of Camoscio using the Italian subset of the XFORMAL dataset [43], hereafter referred to as XFORMAL-IT. The dataset consists of forum messages from the GYAFC corpus [44] automatically translated covering several topics (entertainment, music, family, and relationships). The test set is constructed by using crowdworkers via Amazon Mechanical Turk to collect formal-informal pairs directly in Italian. The model is evaluated in both style transfer directions (Formal to Informal and Informal to Formal). We use only the test split for the zero-shot evaluation and ask the model to generate an answer given the instruction “Dato il seguente testo scritto in modo formale, riscrivilo in modo informale.” and vice versa according to the style transfer direction.

#### 4.1. Evaluation Metrics

We use the same evaluation protocol and scripts of Sarti and Nissim [22]. Specifically, for evaluating lexical matches, we rely on the language-independent ROUGE metric proposed by Lin [45] in the variants unigram (R1), bigram (R2), and Longest Common Subsequence (RL). To gauge semantic correspondence, we employ the trained BERTScore metric [46] with a widely used BERT model pre-trained on Italian<sup>5</sup> and the same baseline scores as Sarti and Nissim [22]. Following previous works, for evaluating the Question-Answering task we employ exact-match (EM) and F1-score (F1). However, since Camoscio is not trained on the output distribution of the question-answering dataset, these metrics will fail to assess the correctness of the output since the EM will count as zero

<sup>5</sup>dbmdz/bert-base-italian-xxl-uncased

even with a correct output but different wording. To account for these variations, we used an approach similar to Zheng et al. [47] that leverages an external LM (in our case *gpt-3.5-turbo*) to judge whether the answer provided by a model is correct (1) or not (0) given the question and the ground-truth answer. We refer to this metric as Exact Match via ChatGPT (EM-GPT) and explain it with additional details in Appendix B.

#### 4.2. Results and Discussion

**Question Answering.** Table 1 shows the results of Camoscio compared to other methods used in the literature. We observe that the metrics commonly used for the task (Exact Match and F1) are very low compared to all the other models. Although this is generally expected since we are comparing trained models with an untrained one, the exact match score is suspiciously low. Looking at the output responses, we noted that Camoscio produces correct but wordy answers (e.g., “La crisi petrolifera del 1973 è iniziata nell’ottobre 1973.” instead of “ottobre 1973”) making the system to perform bad on this score despite the fact that it produces correct answers. Since all the other systems are trained on the datasets, they are aligned with the expected target distribution and the exact match metric is an effective choice. Nevertheless, when it comes to the zero-shot configuration in Camoscio, this conventional metric fails to accurately capture the true performance of the task.

To this end, we evaluated the model also with standard evaluation metrics for generative models (R1, R2, RL, BS). However, we also observe in this case low scores despite the fact that a qualitative examination of the provided answers suggests an overall higher quality. This is possibly due to the different lengths between the produced answers (long) and the ground truth (short) and reinforces the necessity of developing a more precise metric to accurately gauge task performance.

For this purpose, we used instead the metric *Exact*

**Table 3**  
Results on NewSum-IT

	NewsSum-IT			
	R1	R2	RL	BS
mBART Large <sup>6,7</sup>	.377	.194	.291	-
mT5 Small	.323	.150	.248	.375
mT5 Base	.340	.161	.262	.393
IT5 Small	.330	.155	.258	.386
IT5 Base	.339	.160	.263	.044
IT5 Large	.251	.101	.195	.315
Camoscio-7b (0-shot)	.250	.104	.174	.190

*Match via ChatGPT* explained in §4.1. This metric shows that the actual zero-shot performance of Camoscio is in line with the other trained models (.576) and it is also way higher compared to the original EM metric (.077), confirming the need for another type of metric to evaluate the task in the zero-shot setting. Results also show that the EM-GPT metric of trained models correlates well with the existing EM metric, even though with a little marginal difference. This suggests that this metric could serve as an approximate estimation of the model’s actual performance, although it might be subject to bias according to the model used for estimation.

**Style Transfer & Summarization.** Tables 2 and 3 show results respectively for the formality style transfer and news summarization task. We can observe that the zero-shot performance of Camoscio in both tasks is competitive with trained models. According to the model and training dataset, these latter might achieve slightly better scores at the expense of a less generalist model. Looking at the qualitative results, we note however that the summarization task on “Il Post” and “Fanpage” is affected by some common failure cases.

**Failure Cases.** The most common failure case consists of the model not producing an answer at all after the input prompt (4.93% of cases on “Il Post” and 21.16% cases on “Fanpage”). We think that it might be due to the input document of these examples being too long and out of distribution compared to the training documents seen in the instruction-tuning dataset (max length 256 tokens). This might confuse the model and trigger the generation of the end-of-sentence token. Other failure cases include the model responding with a copy of the instruction prompt. For example, we found that in 0.75% of the cases in the “Il Post” split of NewSum-IT, when asked to summarize the article the model responded with

<sup>6</sup><https://huggingface.co/ARTELab/mbart-summarization-ilpost>

<sup>7</sup><https://huggingface.co/ARTELab/mbart-summarization-fanpage>

“Questo articolo non è più commentabile. Abbonati al Post per commentare le altre notizie.” We found this behavior peculiar, considering that the input prompt never mentions that articles were collected from “Il Post”. Looking carefully at the dataset, we found that these documents were not correctly scraped from the web and the model was simply verbatim copying the input error message above (i.e., the message is already summarized and the model is proposing it as-is).

**Open-ended generation.** We show here some qualitative results on the task of open-ended generation. Compared to the other tasks above, in this task we provide the model with just the instruction without any specific input. As a tribute to GePpeTto [20], the first generative language model for Italian, we pose to the model the following instruction: “*Riassumi la storia di Pinocchio*”. The model answered: “*La storia di Pinocchio racconta di un bambino di nome Pinocchio che vive con i suoi genitori in un piccolo villaggio. Un giorno, Pinocchio incontra un burattino magico che gli dona la possibilità di diventare un vero bambino. Pinocchio inizia una serie di avventure mentre cerca di diventare un vero bambino. Durante il suo viaggio, Pinocchio fa amicizia con molti personaggi e impara molte lezioni di vita. Alla fine, Pinocchio raggiunge il suo obiettivo e diventa un vero bambino*”. At first glance the answer might seem correct, it is grammatically accurate and fluent in Italian. However, knowing the story, the model makes several factual errors like characterizing Pinocchio as a kid (instead of a wooden puppet) living with parents (instead of Geppetto) and introducing a nonexistent magic puppet. Overall this example highlights that, although promising, these models have well-known limitations like hallucinations, factual errors, and several kinds of biases [48, 49, 50]. Consequently, it is essential to exercise caution when utilizing them, keeping these limitations in mind.

## 5. Conclusion

In this paper, we introduced Camoscio, a 7 billion instruction-tuned model for the Italian language, together with its Italian instruction-tuning dataset. Results show that the zero-shot performance of Camoscio on several downstream tasks in Italian is competitive with existing models specifically finetuned for those tasks. Despite the known limitations of these kinds of models, this is a first step towards a generalist model capable of performing a wide range of tasks in Italian without explicit finetuning. This is particularly relevant especially in several domains where data is scarce or not available (e.g., medical domain). In an effort to democratize the available and open resources for the Italian language, we release all the artifacts (code, dataset, model) to the community.

## 6. Limitations

Results shown in the paper highlight zero-shot performance competitive with existing finetuned models on three different tasks: summarization (NewsSumIT), question answering (SQuAD-IT), and style transfer (XFORMAL IT). However, it is unclear whether this is true also for other tasks, especially those out of training distribution of the instruction-tuning dataset (see Figure 1). Evaluating and thoroughly assessing the performance of these kinds of models is still an open research question. In addition to this, as already mentioned, the model suffers from common problems that affect language models such as hallucinations, factual errors, and several kinds of biases.

## Acknowledgments

We thank Danilo Croce for pointing out existing implementation issues with the tokenization and the training objective in the *alpaca-lora* repository and Gabriele Sarti for sharing datasets and evaluation protocols used in IT5. We thank all the anonymous reviewers at *CLiC-it 2023* for the feedback provided.

## References

- [1] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *Advances in neural information processing systems* 33 (2020) 1877–1901.
- [2] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al., Palm: Scaling language modeling with pathways, *arXiv preprint arXiv:2204.02311* (2022).
- [3] OpenAI, Gpt-4 technical report, 2023. [arXiv:2303.08774](https://arxiv.org/abs/2303.08774).
- [4] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, et al., Opt: Open pre-trained transformer language models, *arXiv preprint arXiv:2205.01068* (2022).
- [5] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé, et al., Bloom: A 176b-parameter open-access multilingual language model, *arXiv preprint arXiv:2211.05100* (2022).
- [6] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al., Llama: Open and efficient foundation language models, *arXiv preprint arXiv:2302.13971* (2023).
- [7] V. Sanh, A. Webson, C. Raffel, S. Bach, L. Sutawika, Z. Alyafeai, A. Chaffin, A. Stiegler, A. Raja, M. Dey, M. S. Bari, C. Xu, U. Thakker, S. S. Sharma, E. Szczechla, T. Kim, G. Chhablani, N. Nayak, D. Datta, J. Chang, M. T.-J. Jiang, H. Wang, M. Manica, S. Shen, Z. X. Yong, H. Pandey, R. Bawden, T. Wang, T. Neeraj, J. Rozen, A. Sharma, A. Santilli, T. Fevry, J. A. Fries, R. Teehan, T. L. Scao, S. Biderman, L. Gao, T. Wolf, A. M. Rush, Multitask prompted training enables zero-shot task generalization, in: *International Conference on Learning Representations*, 2022. URL: <https://openreview.net/forum?id=9Vrb9D0WI4>.
- [8] J. Wei, M. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, Q. V. Le, Finetuned language models are zero-shot learners, in: *International Conference on Learning Representations*, 2021.
- [9] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, E. Li, X. Wang, M. Dehghani, S. Brahma, et al., Scaling instruction-finetuned language models, *arXiv preprint arXiv:2210.11416* (2022).
- [10] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, T. B. Hashimoto, Stanford alpaca: An instruction-following llama model, [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023.
- [11] E. J. Hu, yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, LoRA: Low-rank adaptation of large language models, in: *International Conference on Learning Representations*, 2022. URL: <https://openreview.net/forum?id=nZeVKeeFYf9>.
- [12] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, et al., On the opportunities and risks of foundation models, *arXiv preprint arXiv:2108.07258* (2021).
- [13] E. Postolache, G. Mariani, M. Mancusi, A. Santilli, L. Cosmo, E. Rodolà, Latent autoregressive source separation, *Proceedings of the AAAI Conference on Artificial Intelligence* 37 (2023) 9444–9452. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/26131>. doi:10.1609/aaai.v37i8.26131.
- [14] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, An image is worth 16x16 words: Transformers for image recognition at scale, in: *International Conference on Learning Representations*, 2021. URL: <https://openreview.net/forum?id=YicbFdNTTy>.
- [15] G. Trappolini, A. Santilli, E. Rodolà, A. Halevy, F. Silvestri, Multimodal neural databases, in: *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Infor-*

- mation Retrieval, SIGIR '23, Association for Computing Machinery, New York, NY, USA, 2023, p. 2619–2628. URL: <https://doi.org/10.1145/3539618.3591930>. doi:10.1145/3539618.3591930.
- [16] A. Santilli, S. Severino, E. Postolache, V. Maiorca, M. Mancusi, R. Marin, E. Rodola, Accelerating transformer inference for translation via parallel decoding, in: Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Toronto, Canada, 2023, pp. 12336–12355. URL: <https://aclanthology.org/2023.acl-long.689>.
- [17] M. Polignano, P. Basile, M. De Gemmis, G. Semeraro, V. Basile, et al., Alberto: Italian bert language understanding model for nlp challenging tasks based on tweets, in: CEUR Workshop Proceedings, volume 2481, CEUR, 2019, pp. 1–6.
- [18] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, 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), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186. URL: <https://aclanthology.org/N19-1423>. doi:10.18653/v1/N19-1423.
- [19] V. Basile, M. Lai, M. Sanguinetti, et al., Long-term social media data collection at the university of turin, in: Proceedings of the Fifth Italian Conference on Computational Linguistics (CLiC-it 2018), CEUR-WS, 2018, pp. 1–6.
- [20] L. D. Mattei, M. Cafagna, F. Dell’Orletta, M. Nissim, M. Guerini, Geppetto carves italian into a language model, in: J. Monti, F. Dell’Orletta, F. Tamburini (Eds.), Proceedings of the Seventh Italian Conference on Computational Linguistics, CLiC-it 2020, Bologna, Italy, March 1-3, 2021, volume 2769 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020. URL: [https://ceur-ws.org/Vol-2769/paper\\_46.pdf](https://ceur-ws.org/Vol-2769/paper_46.pdf).
- [21] M. Baroni, S. Bernardini, A. Ferraresi, E. Zanchetta, The wacky wide web: a collection of very large linguistically processed web-crawled corpora, *Language resources and evaluation* 43 (2009) 209–226.
- [22] G. Sarti, M. Nissim, It5: Large-scale text-to-text pretraining for italian language understanding and generation, arXiv preprint arXiv:2203.03759 (2022).
- [23] L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, C. Raffel, mT5: A massively multilingual pre-trained text-to-text transformer, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Online, 2021, pp. 483–498. URL: <https://aclanthology.org/2021.naacl-main.41>. doi:10.18653/v1/2021.naacl-main.41.
- [24] M. La Quatra, L. Cagliero, Bart-it: An efficient sequence-to-sequence model for italian text summarization, *Future Internet* 15 (2023). URL: <https://www.mdpi.com/1999-5903/15/1/15>. doi:10.3390/fi15010015.
- [25] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, L. Zettlemoyer, BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Online, 2020, pp. 7871–7880. URL: <https://aclanthology.org/2020.acl-main.703>. doi:10.18653/v1/2020.acl-main.703.
- [26] A. Bacciu, G. Trappolini, A. Santilli, E. Rodolà, F. Silvestri, Fauno: The italian large language model that will leave you senza parole!, in: F. M. Nardini, N. Tonello, G. Faggioli, A. Ferrara (Eds.), Proceedings of the 13th Italian Information Retrieval Workshop (IIR 2023), Pisa, Italy, June 8-9, 2023, volume 3448 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023, pp. 9–17. URL: <https://ceur-ws.org/Vol-3448/paper-24.pdf>.
- [27] C. Xu, D. Guo, N. Duan, J. McAuley, Baize: An open-source chat model with parameter-efficient tuning on self-chat data, 2023. arXiv:2304.01196.
- [28] Michael, Stambecco: Italian instruction-following llama model, <https://github.com/mchl-labs/stambecco>, 2023.
- [29] B. Peng, C. Li, P. He, M. Galley, J. Gao, Instruction tuning with gpt-4, arXiv preprint arXiv:2304.03277 (2023).
- [30] E. J. Wang, Alpaca-lora, <https://github.com/tloen/alpaca-lora>, 2023.
- [31] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N. A. Smith, D. Khashabi, H. Hajishirzi, Self-instruct: Aligning language models with self-generated instructions, in: Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Toronto, Canada, 2023, pp. 13484–13508. URL: <https://aclanthology.org/2023.acl-long.754>.
- [32] D. Croce, A. Zelenanska, R. Basili, Neural learning for question answering in italian, in: *AI\* IA 2018—Advances in Artificial Intelligence: XVIIth International Conference of the Italian Association for Artificial Intelligence*, Trento, Italy, November 20–23, 2018, Proceedings 17, Springer, 2018, pp. 389–402.
- [33] A. Scaiella, D. Croce, R. Basili, Large scale datasets for image and video captioning in italian, *Italian*

- Journal of Computational Linguistics 2 (2019) 49–60. URL: [http://www.ai-ic.it/IJCoL/v5n2/IJCOL\\_5\\_2\\_3\\_\\_scaiella\\_et\\_al.pdf](http://www.ai-ic.it/IJCoL/v5n2/IJCOL_5_2_3__scaiella_et_al.pdf).
- [34] C. Larcher, M. Piau, P. Finardi, P. Gengo, P. Esposito, V. Caridà, Cabrita: closing the gap for foreign languages, 2023. [arXiv:2308.11878](https://arxiv.org/abs/2308.11878).
- [35] D. Croce, G. Brandi, R. Basili, Deep bidirectional transformers for italian question answering, in: R. Bernardi, R. Navigli, G. Semeraro (Eds.), Proceedings of the Sixth Italian Conference on Computational Linguistics, Bari, Italy, November 13-15, 2019, volume 2481 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2019. URL: <https://ceur-ws.org/Vol-2481/paper25.pdf>.
- [36] A. Riabi, T. Scialom, R. Keraron, B. Sagot, D. Seddah, J. Staiano, Synthetic data augmentation for zero-shot cross-lingual question answering, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021, pp. 7016–7030.
- [37] A. Holtzman, J. Buys, L. Du, M. Forbes, Y. Choi, The curious case of neural text degeneration, in: International Conference on Learning Representations, 2020. URL: <https://openreview.net/forum?id=rygGQyrFvH>.
- [38] V. Basile, L. Bioglio, A. Bosca, C. Bosco, V. Patti, UINAUIL: A unified benchmark for Italian natural language understanding, in: Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations), Association for Computational Linguistics, Toronto, Canada, 2023, pp. 348–356. URL: <https://aclanthology.org/2023.acl-demo.33>. doi:10.18653/v1/2023.acl-demo.33.
- [39] S. Casola, A. Lavelli, WITS: wikipedia for italian text summarization, in: E. Fersini, M. Passarotti, V. Patti (Eds.), Proceedings of the Eighth Italian Conference on Computational Linguistics, CLiC-it 2021, Milan, Italy, January 26-28, 2022, volume 3033 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2021. URL: <https://ceur-ws.org/Vol-3033/paper65.pdf>.
- [40] L. De Mattei, M. Cafagna, A. AI, F. Dell’Orletta, M. Nissim, A. Gatt, Change-it@ evalita 2020: Change headlines, adapt news, generate, Proceedings of the Seventh Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020) 2765 (2020).
- [41] Q. Lhoest, A. Villanova del Moral, Y. Jernite, A. Thakur, P. von Platen, S. Patil, J. Chaumond, M. Drame, J. Plu, L. Tunstall, J. Davison, M. Šaško, G. Chhablani, B. Malik, S. Brandeis, T. Le Scao, V. Sanh, C. Xu, N. Patry, A. McMillan-Major, P. Schmid, S. Gugger, C. Delangue, T. Matussière, L. Debut, S. Bekman, P. Cistac, T. Goehringer, V. Mustar, F. Lagunas, A. Rush, T. Wolf, Datasets: A community library for natural language processing, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 175–184. URL: <https://aclanthology.org/2021.emnlp-demo.21>. [arXiv:2109.02846](https://arxiv.org/abs/2109.02846).
- [42] P. Rajpurkar, J. Zhang, K. Lopyrev, P. Liang, Squad: 100,000+ questions for machine comprehension of text, in: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, 2016, pp. 2383–2392.
- [43] E. Briakou, D. Lu, K. Zhang, J. Tetreault, Olá, bonjour, salve! xformal: A benchmark for multilingual formality style transfer, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2021, pp. 3199–3216.
- [44] S. Rao, J. Tetreault, Dear sir or madam, may i introduce the gyafc dataset: Corpus, benchmarks and metrics for formality style transfer, in: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), 2018, pp. 129–140.
- [45] C.-Y. Lin, ROUGE: A package for automatic evaluation of summaries, in: Text Summarization Branches Out, Association for Computational Linguistics, Barcelona, Spain, 2004, pp. 74–81. URL: <https://aclanthology.org/W04-1013>.
- [46] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, Y. Artzi, Bertscore: Evaluating text generation with bert, in: International Conference on Learning Representations, 2019.
- [47] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. E. Gonzalez, I. Stoica, Judging llm-as-a-judge with mt-bench and chatbot arena, 2023. [arXiv:2306.05685](https://arxiv.org/abs/2306.05685).
- [48] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, P. Fung, Survey of hallucination in natural language generation, *ACM Comput. Surv.* 55 (2023). URL: <https://doi.org/10.1145/3571730>. doi:10.1145/3571730.
- [49] E. M. Bender, T. Gebru, A. McMillan-Major, S. Shmitchell, On the dangers of stochastic parrots: Can language models be too big?, in: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’21, Association for Computing Machinery, New York, NY, USA, 2021, p. 610–623. URL: <https://doi.org/10.1145/3442188.3445922>. doi:10.1145/3442188.3445922.
- [50] E. Sheng, K.-W. Chang, P. Natarajan, N. Peng, The woman worked as a babysitter: On biases in lan-



guage generation, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Association for Computational Linguistics, Hong Kong, China, 2019, pp. 3407–3412. URL: <https://aclanthology.org/D19-1339>. doi:10.18653/v1/D19-1339.

- [51] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, A. M. Rush, Transformers: State-of-the-Art Natural Language Processing. Association for Computational Linguistics, 2020, pp. 38–45. URL: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- [52] S. Mangrulkar, S. Gugger, L. Debut, Y. Belkada, S. Paul, Peft: State-of-the-art parameter-efficient fine-tuning methods, <https://github.com/huggingface/peft>, 2022.
- [53] T. Dettmers, M. Lewis, Y. Belkada, L. Zettlemoyer, Llm.int8(): 8-bit matrix multiplication for transformers at scale, in: S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), Advances in Neural Information Processing Systems, volume 35, Curran Associates, Inc., 2022, pp. 30318–30332.
- [54] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, in: International Conference on Learning Representations, 2018.

## A. Implementation Details

The model was trained with the LoRA Parameter-efficient Finetuning technique [11], using the Hugging Face Transformers, PEFT, Datasets libraries [51, 52, 41] and the library Alpaca-LoRA [30]. Specifically, it was trained for 3 epochs with int8 quantization [53] on a standard desktop GPU Nvidia 3090 on a machine with Ubuntu 20.04.4 LTS, AMD Ryzen 9 3900X 12-Core Processor and 32GB of RAM. The model was trained with batches of dimension 4 and gradient accumulation to obtain a final “virtual batch” of 128. The maximum length used for training is 256 tokens. The learning rate is set to  $3 \times 10^{-4}$  with AdamW [54] and a total of 100 warmup steps are performed. We used a *lora\_r* (i.e., the dimensionality of the low-rank update of the matrices) equals to 8, *lora\_alpha* equals to 16 and *lora\_dropout* equals to 0.05. We used LoRA adapters just for the matrices *Query* and *Value* in all the attention layers in the LLaMA model, following the original LoRA paper. We used the LLaMA 7 billion checkpoint by loading it from the Hugging Face Hub repository “*decapoda-research/llama-7b-hf*”.

## B. Exact Match via ChatGPT

*Exact Match via ChatGPT* is a metric we introduced to evaluate the performance of Camoscio in the zero-shot setting on the question-answering task. This metric assesses whether the answer provided by a model is correct or not, compared to a ground-truth answer, without the need to have an exact string match (Exact Match). Specifically, we used an external LM (in our case *gpt-3.5-turbo*) that acts as a judge with the scope of verifying the correctness of the answer. We used a prompt similar to the following to compute this metric<sup>8</sup>:

“Given the context below and the corresponding question, please indicate whether the answer is correct (1) or not (0). Use a dict format in the response.

```
Context: {Context}
Question: {Question}
Correct gold answer: {Correct_answer}
Answer: {Answer}”
```

This prompt returns 1 if the answer is correct or zero otherwise. These values are then used to compute the *Exact Match via ChatGPT* with the same procedure of the Exact Match score.

## C. Generation Examples

We propose here some generation examples. We also provide in the GitHub repository the generated text for all the examples in the test datasets, together with other open-ended generation examples.

### C.1. NewSum-IT (“Il Post”)

*Input text:*

L’esame di maturità è sempre un momento importante, quando lo si vive, anche se poi forse a distanza di anni lo si ricorderà con distacco e pensando che non era poi così difficile, o memorabile. Per molti continuerà a farsi vivo di notte, nei sogni. Quello di quest’anno, che inizia oggi, di certo sarà un po’ diverso: senza tema e dizionari, senza attese di gruppo fuori dalla scuola prima di iniziare, ma con mascherine e gel igienizzanti. Agli studenti della Scuola Holden, la scuola di scrittura di Torino fondata da Alessandro Baricco, è stato chiesto di scrivere una lettera ai maturandi del 2020, per incoraggiarli e dare loro sostegno. Quella scritta da Nausica Anerdi, che ha frequentato il primo anno di Academy, il corso di laurea triennale in scrittura della Holden, è stata pubblicata oggi sul Post e su Repubblica, La Stampa e sulla pagina di Torino del Corriere della Sera. Tutte le altre potete

<sup>8</sup>Evaluation script available here

leggerle qui.

*Camoscio generated text:*

Nausica Anerdi, che ha frequentato il primo anno di Academy, il corso di laurea triennale in scrittura della Holden, ha scritto una lettera ai maturandi del 2020, per incoraggiarli e dare loro sostegno.

*mT5-base generated text:*

La lettera ai maturandi del 2020. È stata scritta da Nausica Anerdi, che ha frequentato il primo anno di Academy, la scuola di scrittura di Torino fondata da Alessandro Baricco.

*Ground Truth:*

Una lettera a chi sta per fare la maturità. Da parte di una studentessa della Scuola Holden, del corso di laurea triennale in scrittura.

---

## C.2. NewSum-IT (“Fanpage.it”)

*Input text:*

Continuano ad aumentare i nuovi casi di coronavirus nel nostro Paese. Sono stati 2.800 i contagi registrati ieri: numeri che preoccupano il governo e che ricordano quelli delle fasi più critiche dell'emergenza. Domani l'esecutivo si riunirà e valuterà se sia il caso di rendere più severe le norme anti-contagio attualmente in vigore. Entro la prossima settimana si attende il nuovo Dpcm contenente le misure di contrasto all'epidemia, mentre si valuta la proroga dello stato di emergenza fino al prossimo 31 gennaio 2021. Ma vediamo quindi quali sono queste nuove regole che il governo sta pensando di introdurre per frenare la curva dei contagi. L'obbligo di portare la mascherina all'aperto, già introdotto nei giorni scorsi in alcune zone, sarà esteso a tutto il territorio nazionale. Oltre quindi a confermare la necessità di indossare sempre il dispositivo di protezione nei luoghi chiusi, di igienizzare frequentemente le mani e di rispettare le distanze di sicurezza e il divieto di assembramento, il governo studia se rendere alcune misure più stringenti. In particolare, saranno potenziati i controlli nei luoghi della movida o dove è più facile che si vadano a costituire affollamenti. Le operazioni di vigilanza saranno affidate anche ai militari impegnati nel progetto “Strade secure”. Il ministro della Salute, Roberto Speranza, si sarebbe detto favorevole all'estensione dell'obbligo di mascherina all'aperto a tutto il Paese, mantenendosi per quella linea di prudenza sostenuta fin dallo scoppio dell'emergenza. La stessa considerazione, d'altronde, è condivisa anche dalla autorità locali, dal momento che diversi governatori hanno già emesso delle ordinanze che impongono di coprire naso e bocca all'aperto, con molte salate per

chi non rispetta le nuove regole. Il governo continua ad escludere la possibilità di tornare a una situazione di lockdown generale. Ma se i contagi continuassero a crescere a questi ritmi, potrebbero essere imposte alcune limitazioni in specifiche zone. In questo senso si sta anche considerando di introdurre un chiusura per bar e locali alle 22 o alle 23: una chiusura anticipata affinché non si ripetano situazioni di movida, come abbiamo visto anche quest'estate, in cui si favorisce la diffusione dell'infezione. All'aperto non sarà inoltre consentito riunire più di mille persone alla volta (ovviamente sempre secondo le regole del distanziamento): questa regola sarà valida sia per le manifestazioni che per gli eventi sportivi. Anche il limite per gli stadi resta a mille spettatori. Negli spazi chiusi, invece, il tetto rimane a 200 persone: questo vale sia per cinema, teatri e sale da concerto, ma anche per le celebrazioni private come i matrimoni. Le discoteche resteranno sicuramente chiuse, nei ristoranti si dovrà continuare a rispettare il distanziamento e anche nei negozi gli ingressi saranno contingentati.

*Camoscio generated text:*

Il governo sta valutando l'estensione dell'obbligo di mascherina all'aperto a tutto il Paese.

*mT5-base generated text:*

Il governo sta pensando di introdurre nuove regole anti-contagio per frenare la curva dei contagi. Il provvedimento è stato esteso a tutto il territorio nazionale, mentre si valuta la proroga dello stato di emergenza fino al prossimo 31 gennaio 2021. Vediamo quindi quali sono le nuove regole che il governo sta pensando di introdurre per frenare la curva dei contagi.

*Ground Truth:*

Con i nuovi casi di coronavirus in continuo aumento, il governo pensa a inasprire le regole anti-contagio per contrastare la diffusione dell'infezione. Si va verso l'obbligo di indossare la mascherina anche all'aperto e la chiusura a bar e ristoranti alle 22 o alle 23. Anche l'esercito sarà impiegato nei controlli.

---

## C.3. XFORMAL-IT (F→I)

*Input text:*

Buongiorno, quella su E-Bay è molto bella.

*Camoscio generated text:*

Ciao, quella su E-Bay è davvero bella!

*mT5-base generated text:*

Buongiorno, quella su E-Bay è molto bella.

*Ground Truth:*

!ciaooooooooquella su e-bay è troppo carina!!

---

#### **C.4. XFORMAL-IT (I→F)**

*Input text:*

no mai , però ora per recuperare vado subito a lavare i piatti!

*Camoscio generated text:*

È ora per lavare i piatti.

*mT5-base generated text:*

No, ma ora per recuperare, vado subito a lavare i piatti.

*Ground Truth:*

Giammai! Ora però per recuperare vado subito a lavare i piatti

---

#### **C.5. SQuAD-IT**

*Input text:*

La crisi petrolifera del 1973 iniziò nell' ottobre 1973 quando i membri dell' Organizzazione dei Paesi esportatori di petrolio arabo (OAPEC, composta dai membri arabi dell' OPEC più Egitto e Siria) proclamarono un embargo petrolifero. Alla fine dell' embargo, nel marzo 1974, il prezzo del petrolio era salito da 3 dollari al barile a quasi 12 dollari a livello mondiale; i prezzi americani erano notevolmente più elevati. L' embargo ha causato una crisi petrolifera, o "shock", con molti effetti a breve e lungo termine sulla politica globale e sull' economia globale. Più tardi fu chiamato il "primo shock petrolifero", seguito dalla crisi petrolifera del 1979, definita il "secondo shock petrolifero". Domanda: Quando è stata la seconda crisi petrolifera?

*Camoscio generated text:*

La seconda crisi petrolifera è stata nel 1979.

*mT5-base generated text:*

1979

*Ground Truth:*

1979