

Received June 25, 2021, accepted August 13, 2021, date of publication August 16, 2021, date of current version August 25, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3105278

ELECTRA for Neural Coreference Resolution in Italian

RAFFAELE GUARASCI¹, ANIELLO MINUTOLO¹, EMANUELE DAMIANO¹, GIUSEPPE DE PIETRO¹, HAMIDO FUJITA^{2,3,4}, (Senior Member, IEEE), AND MASSIMO ESPOSITO¹

¹Institute for High Performance Computing and Networking (ICAR), National Research Council (CNR), 80131 Naples, Italy

²Faculty of Information Technology, Ho Chi Minh City University of Technology (HUTECH), Ho Chi Minh City 70000, Vietnam

³Andalusian Research Institute in Data Science and Computational Intelligence (DaSCI), University of Granada, 18011 Granada, Spain

⁴Faculty of Software and Information Science, Iwate Prefectural University, Iwate 15252, Japan

Corresponding author: Aniello Minutolo (aniello.minutolo@cnr.it)

ABSTRACT In recent years, the impact of Neural Language Models has changed every field of Natural Language Processing. In this scenario, coreference resolution has been among the least considered task, especially in language other than English. This work proposes a coreference resolution system for Italian, based on a neural end-to-end architecture integrating ELECTRA language model and trained on OntoCor-efIT, a novel Italian dataset built starting from OntoNotes. Even if some approaches for Italian have been proposed in the last decade, to the best of our knowledge, this is the first neural coreference resolver aimed specifically to Italian. The performance of the system is evaluated with respect to three different metrics and also assessed by replacing ELECTRA with the widely-used BERT language model, since its usage has proven to be effective in the coreference resolution task in English. A qualitative analysis has also been conducted, showing how different grammatical categories affect performance in an inflectional and morphological-rich language like Italian. The overall results have shown the effectiveness of the proposed solution, providing a baseline for future developments of this line of research in Italian.

INDEX TERMS Coreference resolution, ELECTRA, Italian dataset, deep learning, natural language processing.

I. INTRODUCTION

Coreference resolution is one of the tasks that has always fascinated scholars of Natural Language Processing (NLP). Nevertheless, it has a troubled history in the field and it still cannot be considered as a solved task [1]–[4]. Many aspects concerning the nature itself of the coreference are still debated both in Computational [5] and Theoretical Linguistics [6]. A large number of proposed approaches have followed one another over the years, from early rule-based and statistical [7] to machine and deep learning ones [5].

Recent years with the rise of neural language models (NLM) have led to a change. Approaches based on the latest language models like ELMo (*Embeddings from Language Models*) [8], [9] and BERT (*Bidirectional Encoder Representations from Transformers*) [10]–[12] have been proposed and they have brought significant performance gains over all previous approaches [13].

The associate editor coordinating the review of this manuscript and approving it for publication was Ali Shariq Imran^{1b}.

This has opened the way to interesting research perspectives, but this new line of research has so far been almost exclusively the domain of the English language. The rapid growth of approaches available for English does not have an equivalent in other languages. This is partly due to the scarcity of existing resources. English, indeed, can rely on a large number of resources of various domains and sizes, the same cannot be said for the other languages.

Trying to address this shortcoming, this paper proposes a NLM-based system addressing the coreference resolution task in the Italian language, using a novel dataset built starting from OntoNotes [14].

The system end-to-end architecture is derived from [8], but, differently, it exploits a novel language model named ELECTRA (*Efficiently Learning an Encoder that Classifies Token Replacements Accurately*) [15]. ELECTRA has shown a better ability to capture contextual word representations, substantially outperforming, in its downstream performance, other models, like BERT, given the same model size, data, and compute [16]. Moreover, for ELECTRA, a pre-trained

model also exists for the Italian language, that is trained on a standard-Italian written corpora and, thus, results adequate for the coreference purposes.

Concerning the coreference dataset used in this work, the obligatory choice has fallen on OntoCoreFIT,¹ a novel corpus collecting almost 55k utterances in Italian created with a cross-lingual approach starting from OntoNotes. This corpus ensures a good coverage due to its size and domain-independence, in addition collected utterances are compliant with Italian grammar being subjected to linguistic refinement step and to quantitative and qualitative evaluation.

Very few coreference resources have been developed for Italian, the majority of which are outdated and small in size. This is one of the main causes of very slow progress in the field in recent years. Moreover, currently, there are no studies that make a comprehensive analysis of different models to evaluate performances in specific NLP tasks, like coreference resolution, in Italian language. As far as known, this is the first neural coreference resolution system specifically focused on Italian.

Summarizing, this work focuses on these main objectives: i) exploiting a very promising NLM, i.e. ELECTRA, and a novel Italian corpus to make a robust and efficient coreference resolution system for the Italian language, ii) replacing ELECTRA with BERT in order to compare their performance and assess the validity of the proposed system, iii) providing a baseline for future developments of this line of research.

The contribution is structured as follows. Section 2 reviews the state-of-the art of approaches for coreference resolution starting from early rule-based approaches to modern deep learning ones. In Sections 3 and 4 the experimental assessment is outlined, including the details of the neural language model and the dataset used. The Section 5 focuses on the presentation and the discussion of the obtained results, from both a quantitative and qualitative point of view. Finally, Section 6 reports conclusion and future works.

II. BACKGROUND AND RELATED WORK

A. RULE-BASED APPROACHES

Reference resolution task in NLP has been widely considered as a task which inevitably depends on some hand-crafted rules. These rules are based on syntactic and semantic features of the text under consideration. First algorithms [17], [18] were based on syntax or discourse structure, they used rule-based approaches and a manual evaluation. Another similar work [19] added to previous ones the use of knowledge resources and morpho-semantic information. The main limitation of these early approaches was the total lack of a standard for the evaluation phase [20], especially with the advent of shared evaluation campaigns starting from MUC [21], ACE [22] and CoNLL [23], [24].

A rule-based approach with a more robust evaluation was proposed by [25]. This simple modular approach relied only

¹OntoCoreFIT is made available upon request at: <https://nlpit.na.icar.cnr.it/nlp4it/#/datasets/coref>

on syntax and was evaluated on multiple standard datasets. Further improved modified versions were created from this work [26], [27].

B. MACHINE-LEARNING APPROACHES

The increased availability of annotated corpora like MUC or ACE led to the development of machine learning approaches to coreference resolution.

The most fruitful family of machine-learning approaches dealt with coreference as a set of pairwise connections. It used a classifier to decide if two proper nouns (NP) were co-referent. First, these models created training instances to reduce the imbalance between co-referent and non-coreferent entities in the training samples. Types of classifiers proposed ranged from statistical learners [28] to random forests [29]. Finally, a NP partition was generated, in order to test the trained model on a test set to obtain the coreference chains. To handle this task, several clustering techniques were developed, i.e. best-first clustering [30], closest-first clustering [31], correlational clustering [32], Bell Tree beam search [33] and graph partitioning algorithms [34]. Other studies proposed to combine the classifier with effective partitioning using Integer Linear Programming [35] or to completely eliminate the classification phase [36].

Other types of proposed approaches were represented by entity-mention model, that focused on a single underlying entity of each referent in discourse [37], and by mention-pair model, which used a binary classifier to decide whether an antecedent was coreferent with the mention [38], [39].

C. DEEP LEARNING APPROACHES

As happened in all fields of NLP, the advent of Deep Learning rapidly subverted existing approaches. The possibility of representing words using vectors that embed their semantic relationships and the lowest number of required features drastically reduced the need to rely on manually-created features. These approaches captured dependencies between mentions using Recurrent Neural Networks (RNN) or Long short-term memory (LSTM). Limitations of deep learning approaches lie in their poor adaptability to the domain. They frequently need domain-specific adjustment before being used with satisfying results.

The first neural model proposed [1] focused on two critical aspects of coreference resolution: the identification of non-anaphoric references in texts and the ability to distinguish mentions from non-mentions. The model, trained on features extracted from BASIC [40], performed better than all existing approaches. Later developed models [2], [3] incorporated entity-level information produced by a RNN in order to exploit a global features about entity clusters. These models relied on including features defined on mention-pair clusters.

Another approach [41] used the neural mention ranking model [3] in order to replace the heuristic loss functions with reinforced-learning based policy gradient algorithm. Currently, the state-of-the-art approach [42], is based on end-to-end neural model with the construction of

high-dimensional word embeddings to represent words of annotated documents. Although difficult to maintain because of its high-dimensionality, this system, based on LSTM, has as its strengths the ability to capture long term dependencies.

More recent approaches use BERT model to word representation [8], [11] or its modified version SpanBERT to create span representations and increase BERT's maximum segment limitations [12].

D. MULTILINGUAL APPROACHES

Even if the vast majority of coreference resolution approaches was focused on English language — as in many other NLP research areas — individual language-specific approaches was also developed. Dedicated approaches for major European languages were proposed: German [43]–[45], Spanish [46], Portuguese [47] Czech [48], French [49] and many others. Concerning Italian, very few systems for coreference resolution were proposed so far [50], [51]

Evaluation campaigns like Semeval 2010 [52] and Conll 2012 [24] and the development of multilingual resources like Ontonotes [53] or ParCor [54] shifted the focus to systems that can be language-independent or able to be adapted to several languages simultaneously. However, current language-independent systems were based only on shallow approaches exploiting universal part-of-speech tagset [55] and universal dependencies [56]. Many issues are still unsolved and no system is yet able to handle language-specific features [57].

A recent approach that looks promising, in particular in case of low-resource languages, uses cross-lingual methodologies exploiting existing resources developed in other languages. In particular, for low-resource languages, projection-based techniques have been proposed [58], [59]. Projection is a technique consisting in automatically transferring annotations from a resource-rich language to a low-resource language across parallel corpora. In particular, these approaches have used English as source language, and they have been tested for Spanish and Italian [58], Portuguese and Spanish [60], German and Russian [61]. Other studies have tested a direct transfer learning between languages by using multilingual word embeddings, using a model trained on a language for other languages that share a common semantic space [62]: experiments have been carried out on Chinese, Spanish, Portuguese and English [47], [63].

III. MATERIAL AND METHODS

Figure 1 shows the working process behind the proposed coreference resolution system. First, documents composing the OntoCorefIT dataset are given as input. Secondly, the end-to-end *c2f-coref* architecture proposed by [8], [42] and powered by the usage of ELECTRA to represent input tokens, is leveraged to calculate the coreference predictions. A detailed description of the coreference data set realized for the Italian language and the system architecture is provided below.




FIGURE 1. Working process of the proposed system.

A. DATASET

The dataset used here is OntoCorefIT, currently the largest dataset for coreference in the Italian language built from OntoNotes.

Even though other corpora exist for coreference in Italian, none of them is suitable for the purpose of this work. The first reason is given by the size. Corpora currently available for the Italian language are quite small (see Table 1), therefore they should not particularly suitable to be used as training set for a neural language model. On the contrary, OntoCorefIT consists of almost 55k sentences for a total of 704k tokens. A second reason relies in the domain-independence. OntoCorefIT corpus is based on the English version of OntoNotes, so it collects data from a set of heterogeneous domains. By contrast, other Italian existing resources are focused on specific domains: VENEX data come from financial newspapers [64] and spoken task-oriented dialogues [65]; I-CAB and LiveMemories are limited to a restricted material related to the Italian region of Trentino-Alto Adige/Südtirol.

TABLE 1. Size comparison of Italian coreference datasets.

Corpus	Size (words)	Annotation Scheme
OntoCorefIT	708k	CoNLL
Venex (2004)	40k	MATE/MMAX
i-Cab (2006)	250k	ACE
LiveMemories (2010)	250k	ARRAU

OntoCorefIT has been automatically generated from OntoNotes dataset by using a cross-lingual approach, i.e. English utterances are automatically translated and refined in order to obtain utterances respecting Italian grammar and contextually, preserving original coreference and mentions. OntoCorefIT shares the same CoNLL-like annotation schema of original English Ontonotes, ensuring robustness in the evaluation being the *de facto* standard since CoNLL2012 to most recent state-of-the-art systems [42].

In more detail, starting from OntoNotes dataset, firstly, a multi-level translation has been performed to translate utterances from English to Italian, trying to preserve original mentions without losing in the translation the tokens composing the mentions, their positions, and the verbal agreements involving them.

Secondly, the translated utterances are linguistically refined in order to improve their readability and, thus, produce an output text as close as possible to the Italian grammar. In particular, language-specific rules derived from the

Theoretical linguistics have been introduced, covering most frequent phenomena in Italian sentences, i.e. gender and number agreement, subject pronoun deletion or inflection phenomena.

In particular, due the pro-drop nature of the Italian language, the explicit subject pronoun can always be removed. Thus, in case of a pronoun subject tagged as a mention, the tag has been shifted to the verb referred by the pronoun, otherwise it has been simply removed from the utterance.

In line with Ontonotes, all utterances have been successively annotated with Part-of-Speech (POS) tags in order to obtain a CoNLL-like format. Morphosyntactic analysis with POS-tagging has been automatically performed for the utterances in Italian. Thus, the lack of any supervision may have introduced some errors in tag assignment, in particular concerning cases presenting greater ambiguity in certain categories (e.g. pronouns and determiners). Italian, due to its specific features, is considered a more complex language than English [66]–[68] and, therefore, it exhibits greater criticalities in performing pos-tagging and parsing tasks [69], [70].

Table 2 reports an overview of the OntoCorefIT dataset, showing the total number of utterances (*utts*), coreferences (*coarefs*) and tokens for the three partitions into which the dataset is divided, namely *Train*, *Test* and *Dev*.

TABLE 2. Overview of the OntoCorefIT dataset.

	Train	Test	Dev
utts	44,073	5,415	5,363
corefs	40,648	5,039	4,372
tokens	568,641	71,293	68,796

The subsets *Train*, *Test* and *Dev* are arranged into sets of documents each of which is composed of an ordered list of non-overlapping partitions of ordered utterances.

B. SYSTEM ARCHITECTURE

The neural architecture of the proposed system is inspired by the end-to-end *c2f-coref* system proposed by [8], [42]. That end-to-end neural model has been very successful in the literature and it is the current state of the art for the English OntoNotes dataset [24].

In accordance with [8], [42], the task of coreference resolution is defined as a set of antecedent assignments y_i for each span i , with $1 \leq i \leq N$, belonging to a given document D that contains T tokens and $N = \frac{T(T+1)}{2}$ possible text spans.

In particular, all spans are considered as potential mentions and, for each span i , the set of antecedent assignments y_i , i.e. mentions preceding the span under examination and referring to the same entity, is calculated. The set of possible assignments for each y_i is $Y(i) = \{\epsilon, 1, \dots, i - 1\}$, which includes the dummy antecedent ϵ and all preceding spans. The dummy antecedent is used to cover the case when a span is not an entity mention, as well as the case when a span is an entity mention but is not coreferent with other spans. Grouping all spans connected by a set of antecedent predictions allows to define a final clustering.

To realize this task, the model proposed in [8] and [42] learns a conditional probability distribution $P(y_1, \dots, y_N|D)$ whose most likely configuration corresponds to the correct clustering. This distribution is calculated as the product of multinomials for each span:

$$P(y_1, \dots, y_N|D) = \prod_{i=1}^N P(y_i|D) = \prod_{i=1}^N \frac{\exp(s(i, y_i))}{\sum_{y' \in Y(i)} \exp(s(i, y'))} \quad (1)$$

where $s(i, j)$ is a pairwise score for a coreference link between span i and span j in document D . This coreference score is computed as follows:

$$s(i, j) = \begin{cases} 0 & j = \epsilon \\ s_m(i) + s_m(j) + s_a(i, j) & j \neq \epsilon \end{cases} \quad (2)$$

It is equal to 0 in case of dummy antecedent, otherwise it is the sum of three terms, namely $s_m(i)$ and $s_m(j)$ are the scores indicating that the spans i and j are mentions, and $s_a(i, j)$ is the score indicating that the span j is an antecedent for the span i .

The model predicts the best antecedent score if all non-dummy scores are positive, otherwise it vanishes.

Differently from [8], [42], each span i is given an embedding representation h_i by using ELECTRA as shown in Figure 2 and described in the next subsection.




FIGURE 2. Span representation using ELECTRA and attention-based head words.

Given these span representations, the scoring functions s_m and s_a are calculated, as shown in Figure 3, via feed-forward neural networks $FFNN_m$ and $FFNN_a$ as follows:

$$s_m(i) = w_m \cdot FFNN_m(h_i) \quad (3)$$

$$s_a(i, j) = w_a \cdot FFNN_a([h_i, h_j, h_i \circ h_j, \phi(i, j)]) \quad (4)$$

where \cdot is the dot product; \circ is element-wise multiplication; $FFNN$ is a feed forward neural network calculating a non-linear mapping from input to output; $s_a(i, j)$ includes explicit element-wise similarity of each span e_i and a feature vector $\phi(i, j)$ containing information about speaker, genre and other syntactic metadata.

For the training, the marginal log-likelihood of all correct antecedents implied by the gold clustering is optimized:

$$\log \prod_{i=1}^N \sum_{\hat{y} \in Y(i) \cap GOLD(i)} P(\hat{y}) \quad (5)$$

where $GOLD(i)$ is the set of spans in the gold cluster containing the span i .

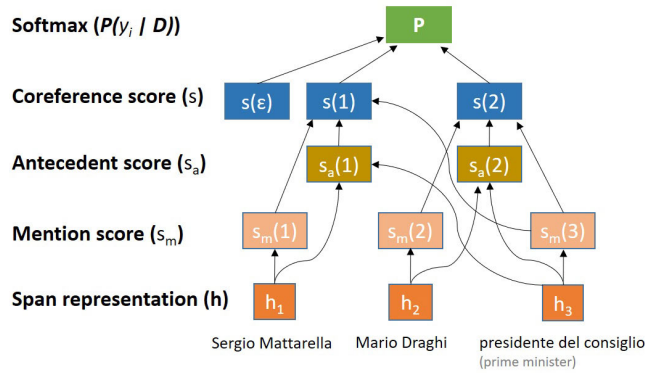


FIGURE 3. Calculation of mention and antecedent scoring functions.

C. ELECTRA

Unlike previous work based on LSTM or BERT [4], [8], [11], the proposed system lies on a new NLM inspired by GAN networks [71], ELECTRA [15]. ELECTRA, has shown to be more compute-efficient than BERT, allowing to achieve better performance keeping the dimensions of the model unchanged.

In more detail, the pre-training step in ELECTRA is performed exploiting the masked language model (MLM) in a more efficient way than BERT. It uses two Transformer models, that share the same word embedding, namely a generator G and a discriminator D , and it is based on training D to distinguish “fake” or replaced input tokens produced by G in the sequence. This approach, called replaced token detection (RTD), allows to use a minor number of examples without losing in performance.

In particular, for a given input sequence, where some tokens are randomly replaced with a [MASK] token, G is trained to predict the original tokens for all masked ones. On the other hand, G is given input sequences built by replacing [MASK] tokens with “fake” ones produced by G and it is trained to predict whether they are original or “fake”.

More formally, given an input sentence s of raw text χ , composed by a sequence of tokens $s = w_1, w_2, \dots, w_n$ where w_t ($1 \leq t \leq n$) represents the generic token (e.g. word, subword or character), both G and D firstly encode s into a sequence of contextualised vector representations $h(s) = h_1, h_2, \dots, h_n$.

Then, for a given position t so that the corresponding $w_t = [MASK]$, the generator outputs the probability to generate a particular token w_t , with a softmax layer:

$$p_G(w_t | s) = \frac{e(w_t)^T h_G(s)_t}{\sum_{w'} \exp(e(w')^T h_G(s)_t)} \quad (6)$$

where $e(\cdot)$ represents the embedding function.

On the other hand, the discriminator predicts whether w_t is the original or “fake”, using a sigmoid layer:

$$D(s, t) = \text{sigmoid}(e(w_t)^T h_D(s)_t) \quad (7)$$

During the pre-training, G employs the following loss function:

$$\mathcal{L}_{Gen} = \mathcal{L}_{MLM} = \mathbb{E} \left(\sum_{i \in m} -\log p_G(w_i | s^{masked}) \right) \quad (8)$$

where $m = m_1, m_2, \dots, m_k$ are k random selected words and s^{masked} is the sentence with the masked words.

On the other hand, D uses the following loss function:

$$\mathcal{L}_{Dis} = \mathbb{E} \left(\sum_{t=1}^n -\mathbb{I}(w_t^{corrupt} = x_t) \log D(s^{corrupt}, t) + -\mathbb{I}(w_t^{corrupt} \neq x_t) \log D(s^{corrupt}, t) \right) \quad (9)$$

where $w_t^{corrupt}$ is the corrupted word within the corrupted sentence $s^{corrupt}$.

Finally, the following combined loss is minimised:

$$\min_{\theta_G, \theta_D} \sum_{s \in \chi} \mathcal{L}_{Gen}(s, \theta_G) + \lambda \mathcal{L}_{Dis}(s, \theta_D) \quad (10)$$

At the end of the pre-training, G is discarded and only the discriminator model D is used.

The main reason that improve ELECTRA efficiency respect to MLM BERT-like models is that predictions are calculated not only over masked tokens, but also for each token and the discriminator loss can be calculated over all input tokens.

In the proposed system, ELECTRA is used to represent each span by considering the embeddings for its boundary tokens as well as for a head token calculated over all its words. In more detail, for each span i , its representation h_i obtained by using ELECTRA is given by:

$$h_i = [x_{START(i)}^*, x_{END(i)}^*, \hat{x}_i, \phi(i)] \quad (11)$$

where $x_{START(i)}^*$ and $x_{END(i)}^*$ are the embedding representations of the boundary tokens, \hat{x}_i represents a soft head word calculated over all the span tokens and $\phi(i)$ is a vector feature which encodes the span size.

The soft head word \hat{x}_i is calculated by using an attention mechanism on the words of each span [72], as follows:

$$\alpha_t = w_\alpha \cdot FFNN_\alpha(x_t) \quad (12)$$

$$a_{i,t} = \frac{\exp(\alpha_t)}{\sum_{k=START(i)}^{END(i)} \exp(\alpha_k)} \quad (13)$$

$$\hat{x}_i = \sum_{t=START(i)}^{END(i)} a_{i,t} \cdot x_t \quad (14)$$

where \hat{x}_i is the weighted sum of the word vectors x_t belonging to the span i , and $a_{i,t}$ are automatically learned weights.

IV. EXPERIMENTAL SETUP AND METRICS

Hereafter the experimental setup and the evaluation metrics are described in Section IV-A and IV-B respectively.




FIGURE 4. ELECTRA overview with replaced token detection. The generator G is usually a MLM trained with the discriminator D but it may virtually be any model producing an output distribution over tokens.

A. EXPERIMENTAL SETUP

To realize the proposed system, the implementation² of the coreference model proposed by [42] has been used, by exploiting ELECTRA embeddings to calculate span representations. ELECTRA model in its base (cased) version³ has been tested, which is made available by Hugging Face Transformers⁴ framework. This framework provides state-of-the-art Transformer-based architectures with thousands of pre-trained models in over a hundred languages for NLP tasks. In particular, this specific ELECTRA model has been pretrained on a corpus of 81GB, made of a recent Wikipedia dump, various texts from the OPUS [73] corpora collections and the Italian part of the OSCAR corpus [74].

In order to assess the effectiveness of the coreference resolution model integrating ELECTRA, further experiments have been arranged, where ELECTRA is replaced by BERT also considering its base (cased) version.⁵ To the best of our knowledge, no other available implementations exist for the particular coreference resolution task in Italian. The choice of using BERT-base in the coreference resolution model as a valid option for comparison is justified also by the fact that it has proven to be effective in the coreference resolution task in English, as shown in [11]–[13].

In detail, the architectures of ELECTRA and BERT are characterised by 12 encoder layers, known as Transformers Blocks, and 12 attention heads (or Self-Attention as introduced in [75]), hence feed forward networks with a hidden size of 768. Each training session has been fixed of 24 epochs, with a learning rate of 0.1. More architectural details and training hyper-parameters are reported in Table 3. All experiments have been performed on a deep learning workstation, with 40 Intel(R) Xeon(R) CPUs E5-2630 v4 @ 2.20GHz, 256 GB of RAM and 4 GPUs GeForce GTX 1080 Ti. The operating system is Ubuntu Linux 16.04.7 LTS.

Using the division of OntoCorefIT in the training, validation and testing datasets shown in Table 2, the results have been derived by averaging the performance of the coreference system integrating ELECTRA or BERT over five repetitions

TABLE 3. Hyper-parameters.

Hyperparameter	Value
Epochs	24
Dropout	0.3
Learning rate	from 0.1 up to 0.00001
Loss	marginalized
Feature Embedding size	20
Max Span Width	30
Max training Sentences	6
Max segment Length	256
Dimensions Hidden State	256
Number of Attention Heads	12
Number of Hidden Layers	12
Hidden size	768
Number of Hidden Layers	12
Parameters	110M
Vocabulary Size	32102

and finally reporting the arithmetic mean of the results, rounded to the second decimal place.

B. EVALUATION METRICS

Concerning evaluation metrics, there are still open issues in the literature and several metrics have been proposed, each of which tries to address biases of the earlier ones. For the purpose of this work, official metrics provided by the Conll 2012 shared task have been taken into account. In particular, the MELA metric has been adopted [76], which combines three different metrics addressing different dimensions: MUC [77], B^3 [78] and $CEAF_e$ [33]. These metrics consider the true set of entities K (named key or key partition) obtained through manual annotation of the entities, and the predicted (or response) set of entities R , i.e. answer partition produced by the system. In particular, they are defined as follows:

- MUC is a link-based metric. It compares the entities defined by the links in the key and the response. MUC considers a cluster of references as linked references, each reference is linked to one or more references. MUC metric primarily measures the number of link modifications required to make the result entity set R identical to the key entity set K . Precision is calculated as follows:

$$Precision_{MUC} = \frac{\sum_{r_j \in R} \frac{|r_j| - |P(r_j)|}{|r_j|}}{\sum_{r_j \in R} (|r_j| - 1)} \quad (15)$$

while recall is equal to:

$$Recall_{MUC} = \frac{\sum_{k_i \in K} \frac{|k_i| - |P(k_i)|}{|k_i|}}{\sum_{k_i \in K} (|k_i| - 1)} \quad (16)$$

- B^3 is a mention-based metric. It first computes precision and recall for each mention, and then calculates the weighted average of these individual precision and recall scores to obtain global precision and recall. In particular, for each mention m of K , the Recall is computed by considering the fraction of the correct mentions included in the predicted entity that contains m . On the other hand, the precision is computed by exchanging the gold

²<https://github.com/lxucs/coref-hoi>

³<https://huggingface.co/dbmdz/electra-base-italian-xxl-cased-discriminator>

⁴<https://github.com/huggingface/transformers>

⁵<https://huggingface.co/dbmdz/bert-base-italian-xxl-cased>

TABLE 4. Comparative results achieved with the coreference resolution system using ELECTRA and BERT models.

Model	run	Identification of Mentions			Coreference									
					MUC			B^3			$CEAF_e$			avg F1
		R	P	F1	R	P	F1	R	P	F1				
ELECTRA cased-xxl	AVG	83.21	89.48	86.23	77.34	80.57	78.72	68.08	71.45	69.71	62.60	73.91	67.78	72.13
BERT cased-xxl	AVG	80.00	89.33	84.39	73.44	79.56	76.38	64.19	70.83	67.34	59.25	72.24	65.10	69.60

entities with the predicted ones. If K is the key entity containing mention m , and R is the response entity containing mention m , precision and recall for the mention m are calculated as:

$$Precision_{B^3} = \frac{\sum_{k_i \in K} \sum_{r_j \in R} \frac{|k_i \cap r_j|}{|k_i|}}{\sum_{r_j \in R} |r_j|} \quad (17)$$

$$Recall_{B^3} = \frac{\sum_{k_i \in K} \sum_{r_j \in R} \frac{|k_i \cap r_j|}{|k_i|}}{\sum_{k_i \in K} |k_i|} \quad (18)$$

- $CEAF_e$ is an entity-based metric. It is a particular instance of CEAF, a metric based on the assumption that each entity of K is mapped to a single entity of R , and vice versa. It uses a similarity measure to find the best one-to-one mapping between entities in K and entities in R . The best mapping is the one that maximizes the overall similarity of the entities, ϕ . In the case of $CEAF_e$, ϕ is given by the following equation:

$$\phi(k_i, r_i) = \frac{2|k_i \cap r_i|}{|k_i| + |r_i|} \quad (19)$$

Recall is equal to the total similarity divided by the number of mentions in K :

$$Recall_{CEAF} = \frac{\sum_{k_i \in K^*} \phi(k_i, g^*(k_i))}{\sum_{k_i \in K} \phi(k_i, k_i)} \quad (20)$$

where g^* is a function that associates to every entity of K an entity of R , whereas K^* is the set of key entities included in the optimal mapping.

Precision is the total similarity divided by the number of mentions in R :

$$Precision_{CEAF} = \frac{\sum_{k_i \in K^*} \phi(k_i, g^*(k_i))}{\sum_{r_i \in R} \phi(r_i, r_i)} \quad (21)$$

According to [24], [79], the combination of these three metrics allows for a good compromise, balancing the limits of the various measures. It is worth noting that an evaluation metric must possess two properties, interpretability, which indicates the goodness of the detected entities and coreference, and discriminability, which allows to distinguish between good and not good decisions. B^3 and $CEAF_e$ provide the property of discriminability, but not interpretability. On the other hand MUC benefits from interpretability, but it is the least discriminating among the three metrics. Since none of the three is reliable if taken individually, it is common practice to use the average of the three as the overall metric.

V. RESULTS AND DISCUSSION

This section presents and discusses the results achieved from both a quantitative and qualitative perspective.

A. QUANTITATIVE ANALYSIS

Table 4 reports results obtained using ELECTRA and BERT coupled with the coreference resolution system with respect to the three metrics described above. In particular, results are detailed with respect to two sub-tasks: *Identification of mentions* and *Coreference*, according to the criteria proposed in Conll 2012 shared task.

The first sub-task calculates the correctness of the mentions that are produced, without considering the coreference link structure, in other words, it is not verified that mentions refer to the right entity; the second one performs this check, evaluating the correctness of the expected coreference links between the key and answer mentions.

Concerning *Identification of mentions* task, the impact of ELECTRA can be seen in the average improvement of about 2 percentage points. Results of *Coreference* task according to three different metrics (columns MUC, B^3 and $CEAF_e$ of the table) show a slight drop in performance with respect to the first task in both models. This result is not surprising since the higher complexity of the coreference task justifies lower results.

With respect to the metrics used, MUC, being link-based, achieves better performance on precision and recall, respectively; while $CEAF_e$ has the lowest scores, especially with regard to recall, which does not even reach 63% using the system with ELECTRA and is even lower than 60% when BERT is used in place of ELECTRA. B^3 average scores are quite similar to those obtained with $CEAF_e$, even if there is a change in behaviour: the recall is higher than $CEAF_e$, to the detriment of precision which achieves the values of 71,45% and 70,83% when the system uses ELECTRA and BERT respectively.

Even there is still no agreement in the literature on the reliability of coreference resolution metrics, the best result obtained with MUC is consistent with other work on the English language [11]. One possible explanation for this result is that MUC has been proven to be robust despite being the least discriminating metric [80]. A further reason is that MUC is particularly suitable on data with many links per mention, as in the case of OntoCoreFIT.

B. ERROR ANALYSIS

A further error analysis is reported, which assesses the performance of the coreference resolution system integrating

ELECTRA model with respect to the different Parts-Of-Speech (POS) each mention belongs to.

Table 5 reports the frequency of occurrence of single-token coreferences occurring in the OntoCorefIT dataset, grouped in relation to their POS ordered in a decreasing manner.

TABLE 5. Distribution of most frequent single-token coreferences grouped with respect to their POS in the OntoCorefIT dataset.

Part-of-Speech	Tag	Percentage
Pronouns	PRON	33.6%
Proper Nouns	PROPN	10.6%
Verbs	VERB	8.6%
Determiners	DET	6.8%
Nouns	NOUN	1.08%
Adverbs	ADV	1.04%

Numbers in the table confirm findings in other studies [81], [82] that indicate pronouns as the primary POS used to co-refer to an entity.

Table 6 outlines the numbers of single-token mentions that are predicted (*p*) or not predicted (*np*) correctly by the system using ELECTRA or BERT, grouped with reference to the three most frequent POSs associated, namely pronouns (PRON), proper nouns (PROPN) and verbs (VERB).

TABLE 6. Predicted (*p*) and not-predicted (*np*) mentions for each POS.

POS	tot	ELECTRA		BERT	
		<i>p</i>	<i>np</i>	<i>p</i>	<i>np</i>
PRON	1,053	991 (94.1%)	62 (5.8%)	983 (93.3%)	70 (6.6%)
<i>personal</i>	553	531 (96.0%)	22 (3.9%)	523 (94.5%)	30 (5.4%)
<i>possessive</i>	362	354 (97.7%)	8 (2.2%)	352 (97.2%)	10 (2.7%)
<i>demonstrative</i>	138	106 (76.8%)	32 (23.1%)	108 (78.2%)	30 (21.7%)
PROPN	639	605 (94.6%)	34 (5.3%)	597 (93.4%)	42 (6.5%)
VERB	525	392 (74.6%)	133 (25.3%)	372 (70.8%)	153 (29.1%)

Table 6 shows discrepancies in the effectiveness of predictions in relation to different part-of-speech categories for the coreference resolution system using ELECTRA or BERT models. Concerning most frequent categories, namely pronouns and proper nouns, the system integrating ELECTRA model reaches a percentage of wrong predictions that is lower than 5.8% and 5.3% respectively. On the other hand, the system integrating BERT model obtains scores that are slightly lower, with 6.6% of not predicted pronouns and 6.5% of not predicted proper nouns. Concerning verbs, the system has an error rate of more than 25% when using both the models, but the impact on the overall results is low due to the fact that verbs account for only 8.6% of the total POS categories occurring in the dataset (see Table 5). This result can be explained also considering the criteria used to construct the OntoCorefIT dataset. Indeed, mentions have been frequently shifted to verbs — following the approach already proposed in [84] and [85] — when the subject could be omitted, as in the case of personal pronouns with subject role (*pronoun-dropping*), adding a higher ambiguity with respect to final identification of mentions belonging to the verb POS.

A more detailed analysis of pronouns shows a different trend between correctly predicted mentions and not. Indeed, the pronouns category is the most representative (33,6% of single-token coreferences) and it includes various types of pronouns with different functions and syntactic structures,

ranging from subject and object pronouns to possessives and demonstratives.

Substantial differences can be noted between the predicted and erroneous pronoun types (as shown in Table 6). More than 80% of correctly predicted pronouns belongs to possessives and personal pronouns: respectively 535 and 362 out of a total of 1053 for the system using ELECTRA while 983 and 523 for the system adopting BERT. In particular, when ELECTRA is used instead of BERT, the system predicts both types of pronouns with a greater accuracy, with an increase of two percentage points in the case of personal pronouns.

Mention predictions are much more inaccurate in case of demonstratives; the system using both the models achieves even 80% of correct predictions. This is also the only case where the system using BERT performs better than when ELECTRA is adopted, achieving 108 of the 138 correctly predicted mentions (78,2%), while stopping with ELECTRA at 106 (76,8%). This is probably explained by the complexity of this pronominal category in the Italian language. Demonstratives present in fact several difficulties in Italian, both at morphological and a syntactic level. They are characterized — like possessives — by several inflected form, depending on gender and number and on the initial letter of the word they refer to. In addition, they are usually involved in long-distance dependencies and subject-object clause construction, i.e. “Ed è **questo** che ho deciso di chiedere” (*And that is what I decided to ask* in English). Finally, a particular case within the group of demonstratives is the pronoun “ciò” (“that” in English) that occupies 15% of the wrong prediction. “Ciò” is invariant neuter pronoun and its lack of inflection produces ambiguities and leads to a number of challenges in some situation, i.e. subject-verb agreement.

C. QUALITATIVE ANALYSIS

To deepen the analysis on the typology of errors, a qualitative analysis of the performance of the coreference system comparing the results achieved by using ELECTRA and BERT language models has been carried out with reference to OntoCorefIT dataset.

Table 7 shows a snippet containing an example related to most frequent POS in the dataset for each language model used. **Bold** text indicates correct prediction with respect to underlined mentions, while incorrect assignments are shown using *italic* text. Square brackets indicate dropped subject pronoun in Italian.

Concerning pronouns, the system using ELECTRA shows a correct prediction in a interrogative sentence with a relative clause. The singular feminine third-person personal pronoun “lei” (*her*) acting as a mention, has the role of subject into the relative clause “pensi che lei lo sta salvando” (*You think she is saving him*) which refers to the noun “presidente” (*president*) in the main clause. But the second sentence contains an error “Non ha previsto il risultato di lui . . . Ma ha rifiutato ciò” (*He did not anticipate his findings, But he refused it*), even if it has a simple syntax with no subordinates. The sentence has

TABLE 7. Examples of correct (bold) and incorrect (italic) predictions with respect to single-token mentions (underline) for most frequent part-of-speech categories.

PoS	ELECTRA	BERT
PRON	Uno dei suoi compiti come presidente ... Pensi che <u>lei</u> lo sta salvando per il libro? (One of her tasks as president... You think <u>she</u> is saving him for the book?) Non ha previsto il risultato di lui . Ma ha rifiutato <u>ciò</u> (He did not anticipate his findings. But he refused <u>it</u>)	<u>Questi</u> hanno scritto oggi... Lei non ha condiviso le note con loro (<u>They</u> wrote today... She did not share the notes with them) Ho pensato a lungo a questo... Molti criticano <u>ciò</u> (I thought about <u>this</u> for a long time... Many people criticise <u>this</u>)
PROPN	Credo che la <u>Cina</u> sia un concorrente... La nazione controlla un sacco del debito (I believe that China it is a competitor. The country controls a lot of the debt) <u>Elia</u> è venuto da noi. <i>Ha dato</i> un sacchetto... (<u>Elijah</u> came to us. [He] <i>gave</i> a bag...)	L'ex avvocato di <u>Clinton</u> ... Sono mosse accuse contro di lui (<u>Clinton's</u> former lawyer... Allegations are made against him) <u>Iraq</u> ha trionfato sul male dell'Occidente ...il Paese crede questo (<u>Iraq</u> triumphed over the evil of the West ... the Country believes this)
VERB	Abbiamo visto <u>Leo</u> ... Pensava che potesse dividere gli Stati Uniti. (We saw <u>Leo</u> ... He thought it could divide the United States) [Lei] Non <u>ha visto</u> alcuna luce... [Lei] <i>Diventerà</i> la tutrice legale (She <u>saw</u> no light... She <i>will become</i> the legal guardian)	Non <u>aveva</u> Ken Starr a confutare... molti hanno seguito il processo (He <u>didn't have</u> Ken Starr to refute...many followed the trial) Ci <u>riferivamo</u> a esso... è sempre difficile <i>farlo</i> (We were referring to it... it is always difficult <i>to do that</i>)

a negative construction and the mention is the neuter pronoun “ciò” (*it*) in an object position. A similar behaviour can be observed for the same POS category for the system using BERT model (third column of the Table 7). The correctly predicted mention occurs as subject “Questi hanno scritto oggi” (*They wrote today*), while the system with BERT fails the correct assignment when the mention occurs as indirect object introduced by a preposition “a questo” (*about this*).

A similar behaviour can be observed for the proper nouns POS. The system using ELECTRA is able to correctly predict the proper noun with object function inside the relative clause introduced by the preposition “Credo che la Cina” (*I believe that China*), while proper noun “Elia” (*Elijah*) is not correctly referred by verb “[Egli] Ha dato” (*He gave*). Maybe the error in an elementary sentence such as this one composed of the simple Subject-Verb-Object (SVO) order is given by the dropping of the subject, which remains implicit in Italian, by making the correct prediction more difficult. The system’s behaviour with BERT for this POS category is more ambiguous, as shown with two example utterances presenting a similar syntax. In the first example, the proper noun acting as subject is into a prepositional phrase “L’ex avvocato di Clinton” (*Clinton’s former lawyer*) and it is correctly predicted. The second utterance has a construction even simpler with the mention as subject in preverbal position at the beginning of the sentence “Iraq”. Despite this, it is not correctly predicted. However, it is important to note that the error rate for proper nouns is the lowest of all the POSs (5.3% and 6.6% respectively), therefore the typology of not predicted sentences is not very representative.

The scenario is different concerning the verb POS, the grammatical category for which there is a smaller gap between correct predictions and incorrectly predicted relationships (see Table 6). Even more than in the previous PoS, two factors seem to influence whether or not the prediction is correct: position of the mention (subject/object) and subject expressed or omitted. For instance, in the example utterance “Abbiamo visto Leo . . . Pensava che potesse dividere gli Stati Uniti” (*We saw Leo . . . He thought it could divide the United States*), the system with ELECTRA correctly associates the verb mention holding the relative clause “**Pensava** che . . .” (*He thought*), to the noun “Leo”. This sentence may appear more articulated, since it does not have a linear syntax and it presents a subordinate clause. However, it should be noted

that the coreference “Leo” is in the object position, therefore the prediction must not refer to an implicit element that has been shifted to the adjacent verb, as in the next example. For the sentence “Diventerà la tutrice legale” (*She will become the legal guardian*) the system with ELECTRA generates a wrong prediction. The verb “diventerà” (*will become*) is tagged as mention, since the subject pronoun [lei] (*She*) is not expressed in the Italian sentence. In addition, this sentence presents two omitted subject pronoun, both for mention and for coreference, thus increasing the difficulty in making a correct prediction. By contrast, the system’s behaviour with BERT is slightly better in this category. An example of correctly predicted mention for the verbs category is represented by a simple SVO utterance with a negative construction “Non aveva Ken Starr . . .” (*He didn’t have Ken Starr. . .*). The last utterance “Ci riferivamo a esso . . . è sempre difficile farlo” (*We were referring to it . . . it is always difficult to do that*) shows an error of the system using BERT since a particular case of verb PoS is occurred. Indeed, the utterance contains a dative construction with a clitic pronoun “Ci” (literally *us*) preceding the mention “riferivamo” (*were referring*) and an enclitic form merged with the verb in the form of suffix -lo for the coreference “farlo” (*to do that*).

VI. CONCLUSION AND FUTURE WORK

In this paper a coreference resolution system for the Italian language has been presented. The system is based on an end-to-end architecture and it uses ELECTRA as a language model to represent input tokens. For the training step, OntoCorefIT dataset has been used. It is the biggest coreference resolution dataset for the Italian language, built using a cross-lingual methodology starting from OntoNotes. As far as is known, there are currently no existing works for Italian based on neural networks for coreference, nor applications that use ELECTRA as language model.

Therefore, the performance of the system integrating ELECTRA have not been assessed with a direct comparison with other existing solutions for the Italian Language, but they have been evaluated by considering BERT in place of ELECTRA in its base (cased) version. Experiments with the system using ELECTRA have achieved better results than using BERT. Scores have been calculated according to *de facto* standard metrics proposed in Conll2012 task, with respect to the sub-tasks *Identification of mentions* and *Coreference*,

showing an increase of at least two percentage points with the system with ELECTRA with respect to BERT. In detail, the system with ELECTRA or with BERT achieves a F_1 score of 72,13% and 69,60% for *Identification of mentions* task and of 86,23% and 84,39% for *Coreference* task, respectively.

According to the results, an error analysis has been carried out aiming at evaluating which grammatical category (PoS) is mostly error-prone for single-token coreferences. Worse performances have been shown for the VERB category, since in OntoCoreFIT dataset, mentions have been frequently shifted to verbs, in case the subject could be omitted, to accommodate the syntactic constructions of Italian. Hence, this has added a higher ambiguity in the final recognition. A deeper analysis has been conducted with respect to mentions belonging to the pronouns category, showing better performances with possessives and personal pronouns with respect to demonstratives, due to a high variability related to morphological and a syntactic aspects. Finally, a qualitative analysis has been conducted to compare the behaviour of the system when ELECTRA or BERT is used as language model, also reporting examples of correct and incorrect predicted mentions.


Further studies could be conducted on the possibility of using other language models specifically for the coreference resolution task. Furthermore, from a strictly linguistic point of view, a deeper analysis not only limited to grammatical categories but also to other linguistic features affecting performances could be carried out as future work.

REFERENCES

- [1] S. Wiseman, A. M. Rush, S. Shieber, and J. Weston, "Learning anaphoricity and antecedent ranking features for coreference resolution," in *Proc. 53rd Annu. Meeting Assoc. Comput. Linguistics, 7th Int. Joint Conf. Natural Lang. Process.*, vol. 1, 2015, pp. 1416–1426.
- [2] S. Wiseman, A. M. Rush, and S. M. Shieber, "Learning global features for coreference resolution," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2016, pp. 994–1004.
- [3] K. Clark and C. D. Manning, "Deep reinforcement learning for mention-ranking coreference models," 2016, *arXiv:1609.08667*. [Online]. Available: <http://arxiv.org/abs/1609.08667>
- [4] C. Lee, S. Jung, and C.-E. Park, "Anaphora resolution with pointer networks," *Pattern Recognit. Lett.*, vol. 95, pp. 1–7, Aug. 2017.
- [5] R. Sukthankar, S. Poria, E. Cambria, and R. Thirunavukarasu, "Anaphora and coreference resolution: A review," *Inf. Fusion*, vol. 59, pp. 139–162, Jul. 2020.
- [6] S. E. Blackwell, "Testing the neo-Gricean pragmatic theory of anaphora: The influence of consistency constraints on interpretations of coreference in Spanish," *J. Pragmatics*, vol. 33, no. 6, pp. 901–941, Jun. 2001.
- [7] V. Ng, "Supervised noun phrase coreference research: The first fifteen years," in *Proc. 48th Annu. Meeting Assoc. Comput. Linguistics*, 2010, pp. 1396–1411.
- [8] K. Lee, L. He, and L. Zettlemoyer, "Higher-order coreference resolution with coarse-to-fine inference," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, vol. 2, 2018, pp. 687–692.
- [9] H. Fei, X. Li, D. Li, and P. Li, "End-to-end deep reinforcement learning based coreference resolution," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, Florence, Italy, Jul. 2019, pp. 660–665.
- [10] B. Kantord and A. Globerson, "Coreference resolution with entity equalization," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, Florence, Italy, Jul. 2019, pp. 673–677.
- [11] M. Joshi, O. Levy, L. Zettlemoyer, and D. Weld, "BERT for coreference resolution: Baselines and analysis," in *Proc. Conf. Empirical Methods Natural Lang. Process., 9th Int. Joint Conf. Natural Lang. Process. (EMNLP-IJCNLP)*, Hong Kong, Nov. 2019, pp. 5803–5808.
- [12] M. Joshi, D. Chen, Y. Liu, D. S. Weld, L. Zettlemoyer, and O. Levy, "SpanBERT: Improving pre-training by representing and predicting spans," *Trans. Assoc. Comput. Linguistics*, vol. 8, pp. 64–77, Dec. 2020.
- [13] L. Xu and J. D. Choi, "Revealing the myth of higher-order inference in coreference resolution," 2020, *arXiv:2009.12013*. [Online]. Available: <http://arxiv.org/abs/2009.12013>
- [14] E. Hovy, M. Marcus, M. Palmer, L. Ramshaw, and R. Weischedel, "OntoNotes: The 90% solution," in *Proc. Hum. Lang. Technol. Conf. NAACL, Companion, Short Papers (XX NAACL)*, 2006, pp. 57–60.
- [15] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning, "ELECTRA: Pre-training text encoders as discriminators rather than generators," 2020, *arXiv:2003.10555*. [Online]. Available: <http://arxiv.org/abs/2003.10555>
- [16] A. Rogers, O. Kovaleva, and A. Rumshisky, "A primer in BERTology: What we know about how BERT works," *Trans. Assoc. Comput. Linguistics*, vol. 8, pp. 842–866, Dec. 2020.
- [17] J. R. Hobbs, "Resolving pronoun references," *Lingua*, vol. 44, no. 4, pp. 311–338, Apr. 1978.
- [18] J. B. Grosz, K. A. Joshi, and S. Weinstein, "Centering: A framework for modelling the local coherence of discourse," Inst. Res. Cogn. Sci., Univ. Pennsylvania, Philadelphia, PA, USA, Tech. Rep. IRCS-95-01, Jan. 1995.
- [19] S. Lappin and H. J. Leass, "An algorithm for pronominal anaphora resolution," *Comput. Linguistics*, vol. 20, no. 4, pp. 535–561, 1994.
- [20] R. Mitkov, "Towards a more consistent and comprehensive evaluation of anaphora resolution algorithms and systems," *Appl. Artif. Intell.*, vol. 15, no. 3, pp. 253–276, Mar. 2001.
- [21] R. Grishman and B. Sundheim, "Message understanding conference-6: A brief history," in *Proc. 16th Conf. Comput. Linguistics (COLING)*, vol. 1, 1996, pp. 466–471.
- [22] G. R. Doddington, A. Mitchell, M. A. Przybocki, L. A. Ramshaw, S. M. Strassel, and R. M. Weischedel, "The automatic content extraction (ACE) program-tasks, data, and evaluation," in *Proc. Lrec*, Lisbon, Portugal, vol. 2, 2004, pp. 837–840.
- [23] S. Pradhan, L. Ramshaw, M. Marcus, M. Palmer, R. Weischedel, and N. Xue, "Conll-2011 shared task: Modeling unrestricted coreference in ontototes," in *Proc. 15th Conf. Comput. Natural Lang. Learn., Shared Task*, 2011, pp. 1–27.
- [24] S. Pradhan, A. Moschitti, N. Xue, O. Uryupina, and Y. Zhang, "CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes," in *Proc. Joint Conf. EMNLP CoNLL-Shared Task*, 2012, pp. 1–40.
- [25] A. Haghighi and D. Klein, "Simple coreference resolution with rich syntactic and semantic features," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2009, pp. 1152–1161.
- [26] K. Raghunathan, H. Lee, S. Rangarajan, N. Chambers, M. Surdeanu, D. Jurafsky, and C. D. Manning, "A multi-pass sieve for coreference resolution," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2010, pp. 492–501.
- [27] H. Lee, A. Chang, Y. Peirsman, N. Chambers, M. Surdeanu, and D. Jurafsky, "Deterministic coreference resolution based on entity-centric, precision-ranked rules," *Comput. Linguistics*, vol. 39, no. 4, pp. 885–916, 2013.
- [28] N. Ge, J. Hale, and E. Charniak, "A statistical approach to anaphora resolution," in *Proc. 6th Workshop Very Large Corpora*, 1998, pp. 161–170.
- [29] H. Lee, M. Surdeanu, and D. Jurafsky, "A scaffolding approach to coreference resolution integrating statistical and rule-based models," *Natural Lang. Eng.*, vol. 23, no. 5, pp. 1–30, 2017.
- [30] V. Ng and C. Cardie, "Improving machine learning approaches to coreference resolution," in *Proc. 40th Annu. Meeting Assoc. Comput. Linguistics (ACL)*, 2001, pp. 104–111.
- [31] W. M. Soon, H. T. Ng, and D. C. Y. Lim, "A machine learning approach to coreference resolution of noun phrases," *Comput. Linguistics*, vol. 27, no. 4, pp. 521–544, Dec. 2001.
- [32] A. McCallum and B. Wellner, "Conditional models of identity uncertainty with application to noun coreference," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 17, 2004, pp. 905–912.
- [33] X. Luo, "On coreference resolution performance metrics," in *Proc. Conf. Hum. Lang. Technol. Empirical Methods Natural Lang. Process. (HLT)*, 2005, pp. 25–32.
- [34] C. Nicolae and G. Nicolae, "BestCut: A graph algorithm for coreference resolution," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2006, pp. 275–283.
- [35] J. R. Finkel and C. D. Manning, "Enforcing transitivity in coreference resolution," in *Proc. 46th Annu. Meeting Assoc. Comput. Linguistics Hum. Lang. Technol. Short Papers (HLT)*, 2008, pp. 45–48.

- [36] E. Fernandes, C. D. Santos, and R. L. Milidiú, "Latent structure perceptron with feature induction for unrestricted coreference resolution," in *Proc. Joint Conf. EMNLP CoNLL-Shared Task*, 2012, pp. 41–48.
- [37] X. Yang, G. Zhou, J. Su, and C. L. Tan, *Improving Noun Phrase Coreference Resolution by Matching Strings* (Lecture Notes in Artificial Intelligence), vol. 3248. Berlin, Germany: Springer, 2005, pp. 22–31.
- [38] X. Yang, J. Su, and C. L. Tan, "A twin-candidate model for learning-based anaphora resolution," *Comput. Linguistics*, vol. 34, no. 3, pp. 327–356, Sep. 2008.
- [39] P. Denis and J. Baldridge, "Specialized models and ranking for coreference resolution," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2008, pp. 660–669.
- [40] G. Durrett and D. Klein, "Easy victories and uphill battles in coreference resolution," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2013, pp. 1971–1982.
- [41] K. Clark and C. D. Manning, "Improving coreference resolution by learning entity-level distributed representations," in *Proc. 54th Annu. Meeting Assoc. Comput. Linguistics*, vol. 1, 2016, pp. 643–653.
- [42] K. Lee, L. He, M. Lewis, and L. Zettlemoyer, "End-to-end neural coreference resolution," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2017, pp. 188–197.
- [43] Y. Versley, *Using Web to Resolve Coreferent Bridging German Newspaper Text*. Princeton, NJ, USA: Citeseer, 2007.
- [44] M. Strube, S. Rapp, and C. Müller, "The influence of minimum edit distance on reference resolution," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2002, pp. 312–319.
- [45] D. Tuggener, "Incremental coreference resolution for German," Ph.D. dissertation, Fac. Arts Social Sci., Univ. Zurich, Zürich, Switzerland, 2016.
- [46] F. Acerenza, M. Rabosto, M. Zubizarreta, A. Rosá, and D. Wonsever, "Coreference resolution between sources of opinions in Spanish texts," in *Proc. 38th Conferencia Latinoamericana En Inf. (CLEI)*, Oct. 2012, pp. 1–8.
- [47] A. F. Cruz, G. Rocha, and H. L. Cardoso, "Exploring Spanish corpora for Portuguese coreference resolution," in *Proc. 5th Int. Conf. Social Netw. Anal., Manage. Secur. (SNAMS)*, Oct. 2018, pp. 290–295.
- [48] A. Nedoluzhko and J. Mirovský, "How dependency trees and tectogramatics help annotating coreference and bridging relations in Prague Dependency Treebank," in *Proc. 2nd Int. Conf. Dependency Linguistics (DepLing)*, 2013, pp. 244–251.
- [49] E. Lassalle and P. Denis, "Leveraging different meronymy discovery methods for bridging resolution in French," in *Discourse Anaphora and Anaphor Resolution Colloquium*. Berlin, Germany: Springer, 2011, pp. 35–46.
- [50] S. Broscheit, M. Poesio, S. P. Ponzetto, K. J. Rodriguez, L. Romano, O. Uryupina, Y. Versley, and R. Zanoli, "BART: A multilingual anaphora resolution system," in *Proc. 5th Int. Workshop Semantic Eval.*, 2010, pp. 104–107.
- [51] M. Poesio, O. Uryupina, and Y. Versley, "Creating a coreference resolution system for Italian," in *Proc. LREC*, 2010, pp. 713–716.
- [52] M. Recasens, L. Márquez, M. A. Martí, M. Taulé, V. Hoste, M. Poesio, and Y. Versley, "SemEval-2010 task 1: Coreference resolution in multiple languages," in *Proc. Workshop Semantic Evaluations, Recent Achievements Future Directions (DEW)*, 2009, pp. 1–8.
- [53] R. Weischedel, M. Palmer, M. Marcus, E. Hovy, S. Pradhan, L. Ramshaw, N. Xue, A. Taylor, J. Kaufman, and M. Franchini, "OntoNotes release 5.0," Linguistic Data Consortium, Philadelphia, PA, USA, Tech. Rep. LDC2013T19, 2013.
- [54] L. Guillou, C. Hardmeier, A. Smith, J. Tiedemann, and B. Webber, "ParCor 1.0: A parallel pronoun-coreference corpus to support statistical MT," in *Proc. 9th Int. Conf. Lang. Resour. Eval. (LREC)*, Reykjavík, Iceland, May 2014, pp. 3191–3198.
- [55] S. Petrov, D. Das, and R. McDonald, "A universal part-of-speech tagset," in *Proc. 8th Int. Conf. Lang. Resour. Eval. (LREC)*, 2012, pp. 2089–2096.
- [56] J. Nivre, "Towards a universal grammar for natural language processing," in *Proc. Int. Conf. Intell. Text Process. Comput. Linguistics*. Cham, Switzerland: Springer, 2015, pp. 3–16.
- [57] S. Kübler and D. Zhekova, "Multilingual coreference resolution," *Lang. Linguistics Compass*, vol. 10, no. 11, pp. 614–631, Nov. 2016.
- [58] A. Rahman and V. Ng, "Translation-based projection for multilingual coreference resolution," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2012, pp. 720–730.
- [59] Y. Grishina and M. Stede, "Knowledge-lean projection of coreference chains across languages," in *Proc. 8th Workshop Building Using Comparable Corpora*, 2015, pp. 14–22.
- [60] A. F. T. Martins, "Transferring coreference resolvers with posterior regularization," in *Proc. 53rd Annu. Meeting Assoc. Comput. Linguistics, 7th Int. Joint Conf. Natural Lang. Process.*, vol. 1, 2015, pp. 1427–1437.
- [61] M. Novák, A. Nedoluzhko, and Z. Zabokrtský, "Projection-based coreference resolution using deep syntax," in *Proc. 2nd Workshop Coreference Resolution Beyond OntoNotes (CORBON)*, 2017, pp. 56–64.
- [62] G. Kundu, A. Sil, R. Florian, and W. Hamza, "Neural cross-lingual coreference resolution and its application to entity linking," 2018, *arXiv:1806.10201*. [Online]. Available: <http://arxiv.org/abs/1806.10201>
- [63] H. Hajishirzi, L. Zilles, D. S. Weld, and L. Zettlemoyer, "Joint coreference resolution and named-entity linking with multi-pass sieves," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2013, pp. 289–299.
- [64] S. Montemagni, F. Barsotti, M. Battista, N. Calzolari, O. Corazzari, A. Zampolli, F. Fanciulli, M. Massetani, R. Raffaelli, R. Basili, M. T. Pazienza, D. Saracino, F. Zanzotto, N. Mana, F. Pianesi, and R. Delmonte, "The Italian syntactic-semantic treebank: Architecture, annotation, tools and evaluation," in *Proc. Workshop Linguistically Interpreted Corpora (COLING)*. Luxembourg: Centre Universitaire, Aug. 2000, pp. 18–27. [Online]. Available: <https://aclanthology.org/W00-1903>
- [65] A. Bristot, L. Chiran, and R. Delmonte, "Verso un'annotazione XML di dialoghi spontanei per l'analisi sintattico-semanticca," in *Proc. 11th Giornate di Studio GFS, Multimodalità e Multimedialità Nella Comunicazione*, 2000, pp. 42–50.
- [66] J. H. McWhorter, "The worlds simplest grammars are creole grammars," *Linguistic Typol.*, vol. 5, nos. 2–3, pp. 125–166, Jan. 2001.
- [67] C. A. Ferguson, "Simplified registers and linguistic theory," in *Exceptional Language and Linguistics*. New York, NY, USA: Academic, 1982, pp. 49–66.
- [68] D. Brunato and F. Dell'Orletta, "On the order of words in Italian: A study on genre vs complexity," in *Proc. 4th Int. Conf. Dependency Linguistics (DepLing)*, 2017, pp. 25–31.
- [69] D. Seddah, R. Tsarfaty, S. Kübler, M. Candito, J. Choi, R. Farkas, J. Foster, I. Goenaga, K. Gojenola, Y. Goldberg, and S. Green, "Overview of the SPMRL 2013 shared task: Cross-framework evaluation of parsing morphologically rich languages," in *Proc. 4th Workshop Stat. Parsing Morphol.-Rich Lang.*, Seattle, WA, USA, Oct. 2013, pp. 146–182.
- [70] R. Tsarfaty, J. Nivre, and E. Andersson, "Cross-framework evaluation for statistical parsing," in *Proc. 13th Conf. Eur. Chapter Assoc. Comput. Linguistics*, 2012, pp. 44–54.
- [71] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," 2014, *arXiv:1406.2661*. [Online]. Available: <http://arxiv.org/abs/1406.2661>
- [72] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," 2014, *arXiv:1409.0473*. [Online]. Available: <http://arxiv.org/abs/1409.0473>
- [73] J. Tiedemann, "Opus-parallel corpora for everyone," in *Proc. 19th Annu. Conf. Eur. Assoc. Mach. Transl., Projects/Products*, 2016, pp. 384–384.
- [74] P. J. O. Suárez, B. Šagot, and L. Romary, "Asynchronous pipelines for processing huge corpora on medium to low resource infrastructures," in *Proc. Workshop Challenges Manage. Large Corpora (CMLC)*, Cardiff, Wales, Jul. 2019, pp. 9–16.
- [75] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, N. A. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst., Annu. Conf. Neural Inf. Process. Syst.*, I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, Eds., Long Beach, CA, USA, Dec. 2017, pp. 5998–6008.
- [76] P. Denis and J. Baldridge, "Global joint models for coreference resolution and named entity classification," *Procesamiento del Lenguaje Natural*, vol. 42, pp. 87–96, 2009.
- [77] M. Vilain, J. Burger, J. Aberdeen, D. Connolly, and L. Hirschman, "A model-theoretic coreference scoring scheme," in *Proc. 6th Conf. Message Understand. (MUC)*, Nov. 1995, pp. 45–52.
- [78] A. Bagga, "Algorithms for scoring coreference chains," in *Proc. Linguistic Coreference Workshop 1st Conf. Lang. Resour. Eval. (LREC)*, Granada, Spain, May 1998, pp. 563–566.
- [79] L. Márquez, M. Recasens, and E. Sapena, "Coreference resolution: An empirical study based on SemEval-2010 shared task 1," *Lang. Resour. Eval.*, vol. 47, no. 3, pp. 661–694, Sep. 2013.
- [80] N. S. Moosavi and M. Strube, "Which coreference evaluation metric do you trust? A proposal for a link-based entity aware metric," in *Proc. 54th Annu. Meeting Assoc. Comput. Linguistics*, vol. 1, 2016, pp. 632–642.

- [81] M.-C. Müller, "Fully automatic resolution of 'it', 'this' and 'that' in unrestricted multi-party dialog," Univ. Tübingen, Tübingen, Germany, 2008.
- [82] M. A. Kabadjov, "A comprehensive evaluation of anaphora resolution and discourse-new classification," Ph.D. dissertation, Dept. Comput. Sci., Univ. Essex, Colchester, U.K., 2007.
- [83] D. McKelvie, A. Isard, A. Mengel, M. B. Möller, M. Grosse, and M. Klein, "The MATE workbench—An annotation tool for XML coded speech corpora," *Speech Commun.*, vol. 33, nos. 1–2, pp. 97–112, Jan. 2001.
- [84] K. J. Rodriguez, F. Delogu, Y. Versley, E. W. Stemle, and M. Poesio, "Anaphoric annotation of Wikipedia and blogs in the live memories corpus," in *Proc. LREC*, 2010, pp. 157–163.



RAFFAELE GUARASCI received the M.Sc. degree in computational linguistics from the University of Pisa and the Ph.D. degree in computational linguistics from the University of Salerno. He has been an Adjunct Professor in the Ph.D. program on computational linguistics and social media with the University of Salerno and on the program committee of national and international conferences. He is currently a Researcher with the Institute for High Performance Computing and Networking, National Research Council (ICAR-CNR), Italy, where he is a member of the Cognitive Systems Laboratory. His current research interests include interaction between neural language models and linguistics theoretical and cognitive issues to the application of artificial intelligence methods based on machine/deep learning to natural language processing tasks.



ANIELLO MINUTOLO received the M.Sc. degree in computer science engineering from the University of Naples Federico II and the Ph.D. degree in information technology engineering from the University of Naples Parthenope. Since 2018, he has been a Contract Professor of informatics with the Faculty of Engineering, University of Naples Federico II. He is currently a Researcher with the Institute for High Performance Computing and Networking, National Research Council (ICAR-CNR), Italy. His current research interests include artificial intelligence, decision support systems, dialog systems, and knowledge management, and modeling and reasoning. He has been involved in different national and European projects and on the program committee of some international conferences and workshops. Moreover, he is also a member of the editorial board of some international journals.



EMANUELE DAMIANO received the M.Sc. degree in computer science from the University of Naples Federico II. He is currently a Graduate Research Fellow with the Institute for High Performance Computing and Networking, National Research Council (ICAR-CNR), Italy, where he is a member of the Cognitive Systems Laboratory and has been involved in different national projects. His research interests include artificial intelligence approaches based on machine/deep learning applied to natural language processing and question answering.



GIUSEPPE DE PIETRO is currently the Director of the Institute for High Performance Computing and Networking, CNR-ICAR, and an Adjunct Professor with the College of Science and Technology, Temple University, Philadelphia. He has been actively involved in many European and national projects, with industrial co-operations. He is the author of over 200 scientific articles published in international journals and conferences. His current research interests include cognitive computing, clinical decision support systems, and software architectures for e-health. He is also a KES International Member. He is also involved in many program committees and journal editorial boards.



HAMIDO FUJITA (Senior Member, IEEE) received the B.S. degree in electrical engineering from the University of Manchester, Manchester, U.K., in 1979, the master's and Ph.D. degrees in information engineering from Tohoku University, Sendai, Japan, in 1985 and 1988, respectively, the Doctor Honoris Causa degrees from Óbuda University, Budapest, Hungary, and Timisoara Technical University, Timisoara, Romania, in 2013 and 2018, respectively. He won the title of Honorary Professor from Óbuda University, in 2011. He is currently a Research Professor with the Faculty of Information Technology, Ho Chi Minh City University of Technology, Ho Chi Minh City, Vietnam. He is also a Professor in artificial intelligence with Iwate Prefectural University, Takizawa, Japan, and the Director of Intelligent Software Systems. He is also an Adjunct Professor of computer science and artificial intelligence with Stockholm University, Stockholm, Sweden; the University of Technology Sydney, Ultimo, NSW, Australia; and the National Taiwan Ocean University, Keelung, Taiwan. He has supervised the Ph.D. students jointly with the University of Laval, Quebec City, QC, Canada; the University of Technology Sydney, Oregon State University, Corvallis, OR, USA; the University of Paris 1 Pantheon-Sorbonne, Paris, France; and the University of Genoa, Genoa, Italy. He has four international patents in software systems and several research projects with Japanese industry and partners. He was a recipient of the Honorary Scholar Award from the University of Technology Sydney, in 2012. He is the Editor-in-Chief of *Knowledge-Based Systems*. He is also the Vice President of the International Society of Applied Intelligence and the Editor-in-Chief of *Applied Intelligence* (Springer). He is also a Highly Cited Researcher in cross-field for the year 2019 and in computer science for the year 2020, respectively (Clarivate Analytics). He has given many keynotes in many prestigious international conferences on intelligent systems and subjective intelligence. He headed a number of projects including Intelligent HCI, a project related to mental cloning for healthcare systems as an intelligent user interface between human users and computers, and the SCOPE project on virtual doctor systems for medical applications.



MASSIMO ESPOSITO received the M.Sc. degree (*cum laude*) in computer science engineering from the University of Naples Federico II, in March 2004, and the Ph.D. degree in information technology engineering from the University of Naples Parthenope, in April 2011. Since 2012, he has been a Contract Professor in informatics with the Faculty of Engineering, University of Naples Federico II. Since 2016, he has been responsible for the Cognitive Systems Laboratory, Institute for High Performance Computing and Networking, National Research Council (ICAR-CNR), Italy, where he is currently a Researcher. He is the author of over 100 peer-reviewed papers on international journals and conference proceedings. His current research interests include artificial intelligence (AI) and are focused on AI algorithms and techniques, mixing deep learning and knowledge-based technologies, for building intelligent systems to converse, understand natural language, and answer to questions, with an emphasis on the distributional neural representation of words and sentences, and on specific natural language tasks, such as part of speech tagging, sentence classification, and open information extraction. He has been involved in different national and European projects. He has been on the program committee of many international conferences and workshops and, moreover, is also a member of the editorial board of some international journals. He received the university's first-level master's degree, namely, European Master on Critical Networked Systems, in December 2007.

...