SAPIENZA UNIVERSITÀ DI ROMA

COGNITIVE PHILOLOGY
No 18 (2025)
ISSN 2035-391X

Old Texts, New Tools: A Data Science Approach to Authorship Attribution for Early Italian Poetry

Aris Anagnostopoulos[*], Domenico Chicca[*], Francesca Cupelloni[**]

[*]Department of Computer, Control, and Management Engineering, Sapienza University of Rome;
[**]Department of Letters and Modern Cultures, Sapienza University of Rome

1. Introduction

Some years ago, from a statistical standpoint, Arjuna Tuzzi (2015) observed that, when it comes to selecting a computational method to detect an author's stylistic hand, «nella letteratura scientifica si trovano centinaia di proposte diverse, per esempio la distanza del coseno, la Delta di Burrows, la distanza intertestuale di Labbé, e nessuna si può considerare migliore in assoluto». Aware of the lack of shared protocols for the so-called non-traditional attribution methodologies, in this study we present the results of an experimental application of a data science approach – in terms to be clarified below – to early Italian texts of uncertain authorship. We have developed an inclusive machine-learning system that makes recommendations for candidate authors together with explanations for these recommendations, while allowing the philologist to overview the entire building choices by constantly consulting the output of the model. We have chosen a sample consisting of seven unknown texts for which the author is unknown but the name of the Florentine poet Antonio Pucci (1309-1390 approx.) has been proposed by scholars as a possible candidate. This project stems from interdisciplinary research and aims to use the selected corpus merely as a testing ground from which to draw additional insights. Disregarding issues of authorship, this investigation has been conducted by considering texts

iuxta propria principia, not only or not so much in relation to their possible authors. Our primary interest lies in refining a method, rather than proposing new – or reviving old – authorship hypotheses[1].

1.1 Attributive Philology and Artificial Intelligence

Attributive philology, a growing important field of humanistic research, aims at investigating a set of possible candidate authors of a text of uncertain authorship on the basis of historical, paleographical, bibliographical, stylistic and linguistic elements[2]. Medieval studies are a particularly fertile ground for this kind of investigation: as it is well known, many medieval texts are anonymous or attributed by manuscripts to different authors, especially to major ones. One of the most interesting exceptions in medieval Italian literature is the Florentine poet Antonio Pucci, whose name, although he is a minor author, often appears in manuscripts alongside works of doubtful attribution. His poetry, in particular, forms such a vast and unbounded corpus that it has discouraged so far the preparation of a critical edition. For this reason, Pucci seemed an ideal test case for this *in corpore vili* experimentation: quantitatively consistent, inherently intertextual, and especially relevant given our previous research[3].

In this work we present a new framework for attributive philology and train it on a dataset of early Italian texts specific to our task in order to assess the accuracy of results on a corpus of doubtful rhymes attributed to Pucci by scholars and/or manuscripts. Our framework is based on a document classifier, which, after being trained on texts of the candidate authors, proposes candidate authors for the selected unknown texts. Equally important, they provide justification for their output, allowing the philologist to understand and weigh the various hypotheses. We show that the combination of (1), an ensemble NN-based classification system, (2), an explanation producing component, and (3), human expertise, can circumvent a lot of the difficulties mentioned in the previous section. Our analysis confirms some of the existing theories for some texts and proposes some new plausible ones for others (see Section 6.3). Whereas the approach was guided by the particular texts, the overall methodology could be used for other areas in philology. It showcases that AI techniques developed for big data applications can be of high value in literature and that human guidance can mitigate the scarcity of data. However, a critical factor that is still challenging to bypass is the lack of large data amounts. For this reason, when applied to early Italian texts, the term "Data Science" itself is nearly oxymoronic: neural network-based approaches typically require

[1] A recent subject-specific survey of published research using machine learning for the study of ancient texts in Sommerschield *et al*. 2023. As far as authorship attribution is concerned, cf. Miranda-García, Calle-Martín 2008, Seroussi, Zukerman, and Bohnert 2014, Cupelloni 2021.

[2] Cf. Stoppelli 2020, p. 20.

[3] Cf. Cupelloni 2022.

hundreds of thousands of training samples to function effectively. For a given author, despite using all the texts that we know that he has written at the time of the study, we often have only a handful of texts (see Section 4). It is a challenge to surpass this limitation and make use of the very successful modern techniques with limited amounts of data. Some inherent methodological challenges in studying this ancient language, such as its polymorphic nature, its semantic values—often very different from modern Italian, and its dimensionally reduced dataset, must also be considered. For this reason, even though technically the problem can be considered as a *text/document-classification problem*, it has various particularities requiring the strong-interaction between the expert philologist and the AI system; see Section 5.

1.2 Roadmap

In Section 2 we present some related work on authorship attribution via computational means and in Section 3 we survey the work on attributive philology in OI. In Section 4 we describe how we built our corpus, and we provide some background information on the authors that we study and on early Italian texts in general. Section 5 presents our approach for authorship attribution. In Section 6 we present our results, first on text in which the authors are known, and then in unknown texts, allowing us to compare our findings with the existing philological theories. Finally, in Section 7 we discuss the value and the limitation of our approach and we conclude. We designed our approach for study on OI, but the approach may be of interest for general authorship attribution. Readers who are not interested in the particular use on early Italian may skip Section 3 and skim over Section 4.

2. Related Work on Authorship Attribution

There is a very limited number of studies that use statistical techniques for author attribution on early Italian (see Section 3); much more data are available from authorship attribution's studies in other ancient languages, such us Arabic.

Ouamour and Sayoud (2013) studied what types of features work most effectively on very short ancient Arabic texts that vary in length between 209 and 800 words. In their experiments they tested characters, character n-gram, words, and word n-ngram for the task of authorship attribution with a support vector machine (SVM) classifier. Despite the brevity of their texts, they obtained 80% of accuracy with character bi-gram, tri-gram, and rare words. In a more recent work, Salami and Momtazi (2020) propose the use of a neural model for poet identification. They stacked a convolutional layer on top of a bidirectional recurrent layer to grasp semantic features from the input text. They performed several experiments by varying the number of authors and tested their classifier on both entire poems and individual lines in an attempt to establish the minimum usable text length. In the hardest setup they achieved 78% f-

measure.

Ivanov (2019) experimented on more recent texts: on the *historical document corpus*, a corpus containing 224 English-language texts written by 38 authors in the 18th century. He investigated the effectiveness of assonance as a stylistic marker for authorship attribution. Assonance is a form of imperfect rhyme, consisting in having two or more successive verses ending in syllables that contain different consonants but the same vowel, producing a similar, but different, sound. He conducted several experiments, by comparing or combining part-of-speech tags, alliteration, and assonance stylistic features. He used SVM classifiers, with sequential minimal optimization and multilayer perceptron. The main claim of the article is that, when the number of candidate authors is large, assonance is not a useful feature when used along with more traditional ones; however, for a small set of candidate authors it can help in the authorship attribution task.

The area of computer forensics studies the problem of authorship attribution, but it is focused mostly on modern texts. The *PAN@CLEF* is a series of scientific events focused on digital forensics and *stylometry*, the discipline whose goal is the analysis of authors' writing style through the use of statistical methods. Custodio and Paraboni (2018) designed an ensemble classifier (a classifier that is composed of multiple subclassifiers each of which uses a different approach to analyze the input text), which was able to provide superior results compared to the other approaches presented at the PAN@CLEF 2018 event[4]. Similar findings presented Muttenthaler *et al.* (2019) a year later. One of our classification approaches (Section 5.15.1) uses an ensemble classifier as well.

Some work on stylometry studies what features should be extracted from the input text and used for classification. Stamatatos (2017) studies how to extract text features that are able to capture an author's writing style but not the content; this may be useful when an author writes texts that cover a wide range of topics yet uses a similar style in his writing. The question of style vs. content is also the topic of a study by Sari *et al.* (2018), who show the strong dependence of the effectiveness of the features on the input dataset: content-based features are more effective if the input data contain texts of high topical variety, whereas style-based features are preferred for more uniform datasets (e.g., movie reviews). Our application scenario of early Italian is different than this area of research, yet some of the intuition can be carried to our case as well: we will see that the fact that the authors of our study have written texts of many different styles is a factor that we take into account when choosing how to train our models.

Neural network models, which have been proven successful in the last years in many fields, have found their way also into authorship attribution. One of their advantages is the ability to extract features automatically. Shresta *et al*. (2017) employed character N-gram convolutional neural networks, which

[4] https://pan.webis.de/

allowed them to work with very short text samples (e.g., tweets). Their study uses temporal convolution to build a feature map out of variable length character N-grams. Jafariakinabad *et al.* (2019) proposed a hierarchical attention model. They show that syntax and structure are features powerful enough to shape the author's style. First, they learn the syntactic representation of sentences in a given text using part-of-speech, then they aggregate them in the representation of the whole text using recurrent neural networks and in particular bidirectional-LSTM. Their work makes explicit effort to incorporate syntactic information into the model. Another example is the work of Ferracane *et al.* (2017). It extends the work of Shresta et al. by adding features capturing syntactic information and they improve the results. Based on this knowledge, our ensemble model incorporates some of these features as they prove to be informative also for early Italian (see Section 5.1).

3. A Survey of Attributive Philology and Digital Analysis of Early Italian Texts

In Italy, attributive philology was defined for the first time by Gianfranco Contini in the *Enciclopedia del Novecento* (Contini 1977), also referring to Roberto Longhi's attributionism and Leo Spitzer's *Stilkritik*. The Italian philologist underlines here the great caution required in drawing definitive conclusions only on a stylistic basis, that is, determining once and for all whether a certain stylistic system or system of stylistic elements can really be considered an internal signature[5]. At Contini's time, Erdman and Fogel (1966) discussed problems of attribution from Shakespeare to Keats with a similar focus on the limitations and the efficacy of internal evidence.

Since the assembly of various electronic corpora of texts ("and in some cases even their extinction, due to the obsolescence of the computer technologies with which they were produced") [6], such as, limiting to early Italian texts, the *Tesoro della Lingua Italiana delle Origini* (*TLIO*) and the *Letteratura Italiana Zanichelli* (*LIZ*), the *modus operandi* has changed significantly in terms of the possibility to verify philologists' initial intuition: corpus interrogation is a crucial point to bring further empirical basis to the analysis of a text whose attribution is uncertain. Many case studies have benefited greatly by the correct use of Italian text databases (for a proper approach to them, see, for example, Kennedy and Miceli 2001): in some of the most recent ones in Italian medieval literature (Cupelloni 2018; Ruggiero 2020; Stoppelli 2020, 2021), the use of concordances and frequency lists has provided valuable insights into potential candidates for authorship. In addition to the scrutiny of electronic indexes, stylometric methods can play a decisive role in the attributive methodology of medieval texts. As recently emphasized from the viewpoint of stylometry (Kestemont 2012; Canettieri 2013), medieval philology is

[5] On this topic see also Cortelazzo and Tuzzi 2017, pages 12–13.
[6] Luvisotto 2020.

a field devoted to the study of a large number of cases of unresolved authorship and therefore a privileged context for computational authorship attribution techniques. Nevertheless, these methods do not offer conclusive proof and need to be combined with qualitative analyses (e.g., stylistic and philological considerations) to work synergistically.

This necessity for interdisciplinary approaches has encouraged some interesting collaborations in recent years; for example, between philologists and physicists (e.g., Canettieri *et al.* 2008) or between linguists and statisticians (e.g., Cortelazzo *et al*. 2013, 2016; Cortelazzo and Tuzzi 2017, 2018; for a review of Italian multidisciplinary projects in Digital Humanities, see Ciotti 2014). In most cases, scientific literature has shown a growing interest in modern and contemporary Italian literature, whereas quantitative linguistic approaches to medieval Italian texts are still underrepresented in academic studies. To the best of our knowledge, they are limited to those of Barber (1985, 1990), Robey (1997), Canettieri (2011, 2013, 2016), Canettieri and Italia (2013), and Corbera *et al.* (2019).

Corbera et al. (2019) is worthy of a particular mention. The study is focused on the application of statistical techniques for *author verification*, namely, to verify whether the famous *Epistle to Cangrande* was written by Dante Alighieri. The text, one of the most controversial among the works of the Italian poet (along with the *Fiore*; see Section 4.1), is considered a forgery by some scholars. This opinion has been reinforced by the results of the classifier proposed in Corbera et al. (2019), whose predictions output permit to assert the *Epistle* is the product of an unnamed forger. However, as the same scholars underline, the results of the study should be carefully interpreted; only further well-controlled research will answer many open questions about the text.

4. The Corpus

We now present the corpus that we use in this work.

4.1 Models and Criteria

The reference corpus for the texts used in our model is the OVI corpus, created by the Italian National Research Council (CNR-Opera del Vocabolario Italiano, located in Florence)[7]. It is the largest database available today of early Italian literature, which contains, as of June 16, 2025, 3,725 texts, from the first text that can be said to be Italian (Placito di Capua, 960 CE) to the end of the 14th century (1375, Giovanni Boccaccio's death, is a symbolic cut-off point). In this study, we concentrate on a much smaller corpus composed exclusively of literary works ascribable to the period of 1300

[7] http://gattoweb.ovi.cnr.it/. The corresponding lemmatized database used to compile the early Italian dictionary, proposed by the same CNR-Opera del Vocabolario Italiano (Florence), is the Tesoro della Lingua Italiana delle Origini (http://tlio.ovi.cnr.it/TLIO/).

CE–1400 CE; among them, there are texts whose origins we do not know and are hard to attribute. The unknown texts considered in this study are all related to Antonio Pucci: we select texts from his corpus of works (3 out of 12)[8], along with five texts for whom the name of the poet has been mentioned as one of the most likely authors.

In more detail, Table 1 show, for each text, the candidate author (one or more) whom scholars consider most probable; without any pretense of being exhaustive, we consider in this table only the main scholars' proposals, sometimes followed by the indication of the manuscripts on which the proposal is built.

Text	Candidate author	Candidate author	Candidate author
Andandosi la volpe cum la gata (caudate sonnet)	Antonio Pucci (Contini 1938)	A poet's of Pucci's circle (Contini 1938)	
Bel Gherardino (*ottava rima*)	Antonio Pucci (Zambrini 1867)		
Cantari del Carduino (ottava rima)	Antonio Pucci (Rajna 1873)		
Ciascun faccia per sé (ballad)	Antonio Pucci Corsi 1969; Firenze, Biblioteca Nazionale Centrale, II.IV.61)	Niccolò Soldanieri (Trucchi 1847)	
Detto d'Amore (seven-syllable lines)	Dante Alighieri (Contini 1984)		
Fiore[9] (series of sonnets)	Dante Alighieri (Mazzoni 1923, Rajna, Contini 1984, etc.)	Antonio Pucci (Fasani 1973, 1974-1975)	Dante da Maiano (Stoppelli 2011)
Quando 'l consiglio degli ucce' si tenne (caudate sonnet)	Dante Alighieri (De Robertis 2002; Firenze, Biblioteca Laurenziana, Laurenziano Rediano 184; Firenze, Biblioteca Nazionale Centrale, II IV 114)	Antonio Pucci (Contini 1995, Marti 2002, Giunta 2014; Cupelloni 2019)	

Table 1. Proposed authorships in the literature for our unknown texts.

[8] For the complete list see Mirabile's archive (http://www.mirabileweb.it/author-rom/antonio-pucci--author/LIO_229674).

[9] For the text we consider only some of the scholars' hypotheses proposed so far.

A few words about some of these texts. Among them, the most well-known case is that of the *Fiore*. Since the discovery of the Montpellier manuscript (École de Médecine, H 438) by the French philologist Ferdinand Castets (1881), who identified the self-named "ser Durante" as Dante Alighieri, the poem inspired much debate in the Italian academic world (since Contini 1984)[10]. The attributive question, still not resolved, has so far produced different answers (the work has been assigned to various authors: Dante Alighieri, Dante da Maiano, Antonio Pucci, Immanuel Romano, etc.) mainly on the basis of stylistic and linguistic elements, sometimes supported by means of computer-aided research (Canettieri 2011a, 2011b).

Another interesting open question is related to the sonnet *Quando 'l consiglio degli ucce' si tenne*. In this case, two manuscripts attribute the sonnet to Dante; in other two, the text is anonymous. According to Domenico De Robertis (2002), this situation implies Dante's authorship; for Contini (1995), Marti (2002) and Giunta (2014), the silence of two out of four codexes could be meaningful and suggest that it is not he the author of the text.

The question is different for the *Bel Gherardino* and *Cantari del Carduino*. In these cases, it is not a question of identifying the author: anonymity is the rule rather than otherwise in this kind of poetry, even if there are some interesting proposals of attribution in literature (the most remarkable one, about Pucci, could be read in Rajna 1873). Pucci is the only well-known author of the genre of *cantari* whose name we know. He usually signs his work as follows: "Al vostro onor questo fe' Antonio Pucci". We have decided, therefore, to include the two texts in our list of unknown texts to verify whether the model effectively recognizes the presence of a common language between them and Pucci's work.

4.2 Authors of Early Italian Texts

The authors included in our corpus-based analysis are the following:

(1) Dante Alighieri (1265-1321);

(2) Giovanni Boccaccio (1313-1375);

(3) Antonio da Ferrara (1315-1374 approx.);

(4) Dante da Maiano (second half of the 13th century?);

(5) Fazio degli Uberti (1301?-1367?);

(6) Francesco Petrarca (1304-1374);

(7) Antonio Pucci (1309 approx.-1390?);

(8) Franco Sacchetti (1330 approx.- 1440 approx.);

[10] For an overview of the discussion and a survey of arguments both for and against Dante's authorship, see Stoppelli 2011.

(9) Simone Serdini (1360 approx.-1419 or 1420);

(10) Niccolò Soldanieri (second half of the 14th century);

(11) Giovanni Villani (1280 approx.-1348).

To enhance, from a multidisciplinary perspective, the philological and critical approach in terms of investigation both of the manuscript tradition and of some previous scholars' hypotheses, we considered, firstly, authors explicitly documented by manuscripts or who have been previously named by scholars as the possible candidates of one or more of our seven unknown texts. Secondly, to increase linguistic information for training the model, we added to the corpus the entire work of other seven authors of the same geographical area (Tuscany) and chronological period (13th and 14th centuries), such as Giovanni Boccaccio and Francesco Petrarca. In Figure 1 we characterize the literature work of each author with respect to the textual genre. The number of texts included in the corpus is to be understood as merely illustrative; only a representative sample, reflecting the variety of the Tuscan medieval panorama, is covered in this study, providing, as we wished, a contribution for future insights and endeavours.

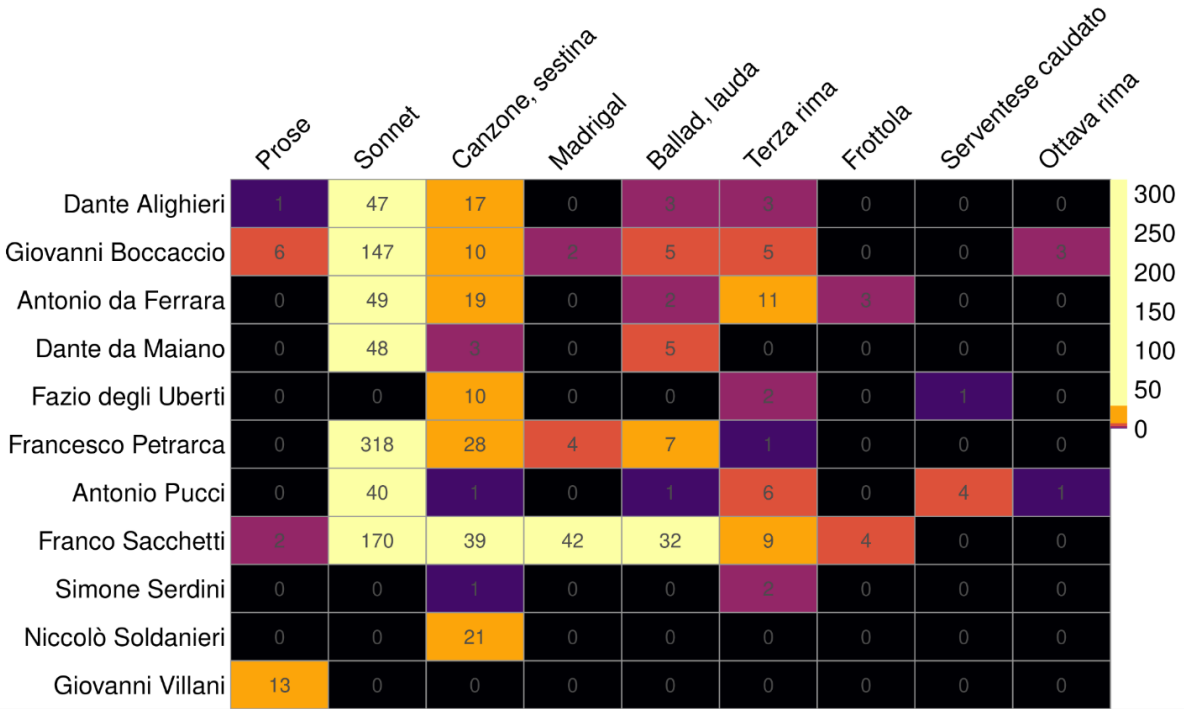

Figure 1. Number of texts in each style written by the authors.

4.3 Corpus Criteria

For a correct interpretation of the results (see Section 6), it must be considered that our dataset is:

1. heterogeneous in terms of genres (prose/poetry): there are four out of eleven authors whose selected texts are both in verse and in prose (Dante Alighieri, Giovanni Boccaccio, Antonio Pucci, Franco Sacchetti); for the other ones, the works included for analysis and scoring are only poems (Antonio da Ferrara, Dante da Maiano, Fazio degli Uberti, Francesco Petrarca, Simone Serdini, Niccolò Soldanieri) or a chronical prose (Villani).
2. diachronically motivated: each author of the corpus lived and worked during the 13th and/or 14th centuries (except for Simone Serdini, who died in 1420);
3. diatopically representative: all texts can be classified as written in early Tuscan.

The analysis includes a total of 1,148 texts and 329,078 tokens: a not large but representative specimen of the general panorama of medieval Italian literature selected for testing the accuracy of our model. In this study we consider only a selection of the literary works included in the OVI corpus. Poetry forms the main component of our genre-diversified corpus of known texts: 819 sonnets; 139 canzoni; 48 madrigals; 55 ballads; 39 poetic works in *terza rima*; 10 *sestine*; 7 *frottole*; 5 *serventesi caudati*; 4 poetic works in *ottava rima* (or *cantari*). As far as our corpus's works in prose are concerned, they are ten, five of which written by Giovanni Boccaccio (*Decameron*, *Corbaccio*, *Elegia di Madonna Fiammetta*, *Filocolo* and *Trattatello in laude di Dante*), two by Franco Sacchetti's (*Trecentonovelle* and *Signor mio, ben che io sia grave di corpo*), and one by Giovanni Villani (*Nuova Cronica*). For each author we have reported in Figure 2 the total number of words.

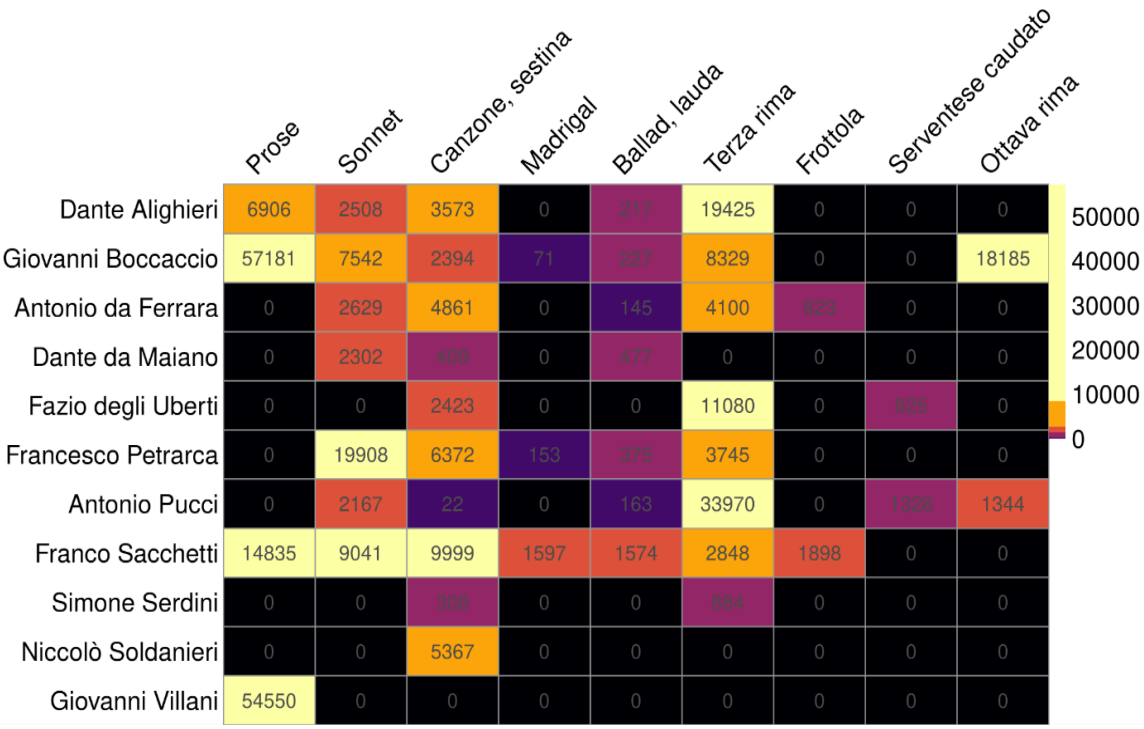

Figure 2. Number of words in each genre written by the authors.

5. Our Approach

Our approach has two modules: (1) A neural-network based text classifier, and (2) a module providing explanations for the choices made by the classifier. Those two modules interact with the human expert (philologist) who studies the results, interprets the explanations and guides the adjustments to the system, as we discuss in Section 7. See Figure 3. Note that the nature of the problem makes indispensable the interaction of the three components. See Section 7 for more details. We now proceed into explaining the two modules.

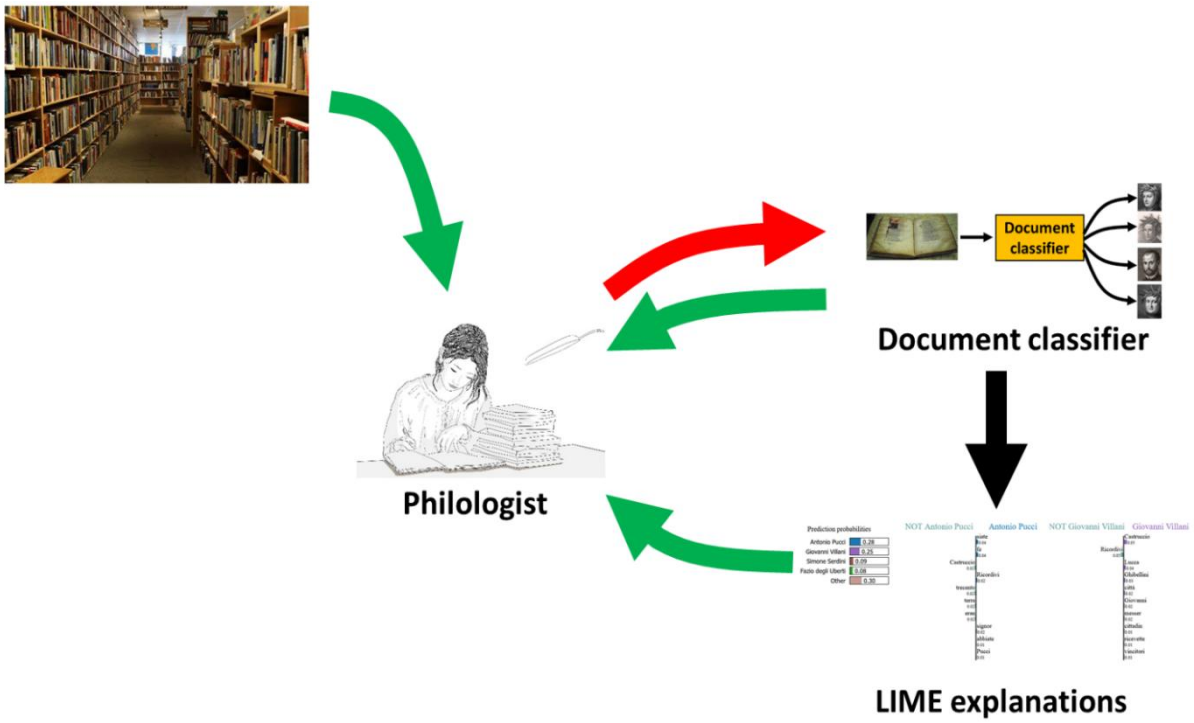

Figure 3. Our approach: The philologist consults the output of the classifier, and the explanations provided to guide the design parameters and interpret the results.

5.1 Document Classification

In this module we treat the problem of authorship attribution as a *multiclass classification machine learning task*: For each author under consideration we are provided with a collection of texts written by him and the goal is to train a machine-learning model (aka a classifier) that, when presented with an unknown text, for each of the authors computes a score that is higher when the author is more likely to have written the text under consideration. We use two main classifiers: an ensemble neural-network (NN) classifier and a support vector machine (SVM) (Cortes and Vapnik 1995).

*Ensemble NN classifier***.** The classifier is an ensemble model of four neural-networks (NN), each of which processes a different set of features. See Figure 4 for a global overview of the architecture.

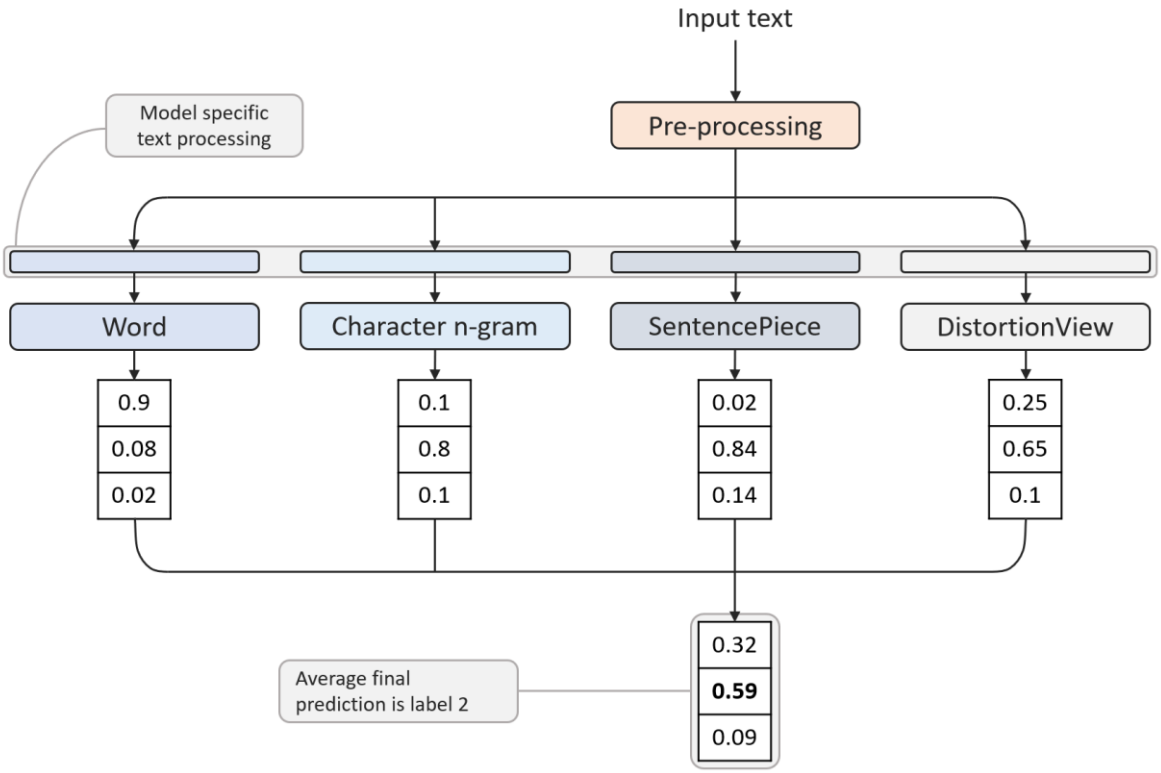

Figure 4. Ensemble classifier scheme.

The first step consists in preprocessing and is common for all four models. Following the literature, and after experimentation we followed the following steps: (1) Removed stopwords and words that we considered not useful for the attribution task (this may appear bizarre for an authorship attribution task, yet in our case it is necessary; see the discussion in Section 7. (2) Retained case folding, accents, and punctuation symbols. (3) We normalized some symbols that used different characters (e.g., double quotation marks). (4) We normalized white spaces into a single white-space character. (5) We separated punctuation symbols from words, except for the periods appearing in abbreviations. Finally, (6) we broke long text sequences into sequences of length at most L=256 *tokens,* where a token is a word or a punctuation symbol.

After the common preprocessing phase, each of the four NN classifiers processes the input and extracts its own set of features. We call the four models *Word*, *Character n-gram* (*Char*), *SentencePiece* (*Piece*), and *DistortionView* (*Dist*). In Figure 5 we can see an example of how each model extracts features from the input phrase. For the *Word* model each word and each punctuation mark (except for abbreviations) is a single feature. The *Char* model receives as features the ones of *Word* and in addition all sequences of n characters (n=3), with the beginning and end of words being indicated with two special characters. The *Piece* model introduced by Kudo and Richardson (2018), is a tokenization method, which selects subword segments in an unsupervised way. Finally, *Dist* (cf.

Stamatatos 2017) is a model that replaces with wildcards words that appear a limited number of times; the goal is to capture some stylistic elements.

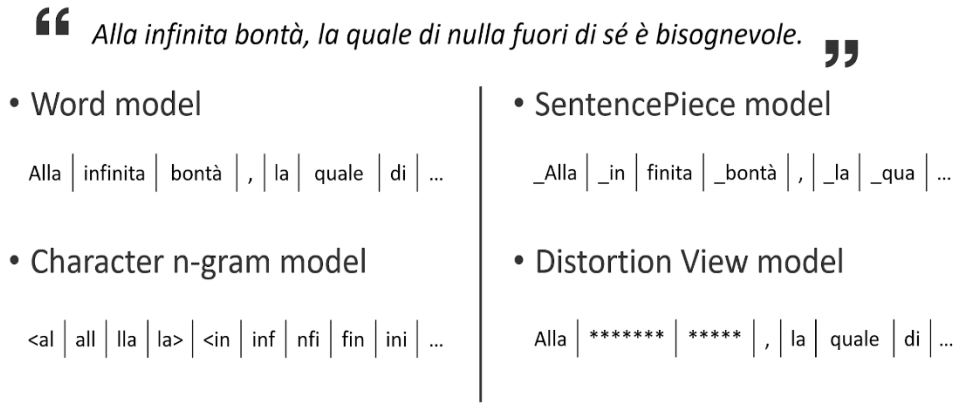

Figure 5. Example of input feature preparation.

For each of the four models we split the input into segments of length L (L=256); if the final length of the vector is less than *L* we pad it with a special character, if it is larger (it can happen in the SentencePiece model) we truncate it. The input is fed to a NN using a one-hot representation. Each of the four NNs has the same architecture, which is based on the Transformer model (Vaswani *et al*. 2017). In the supplementary material we provide the precise description. Thus, for each input segment we obtain four *C*-dimensional vectors (C=11). A final weighted average over the four models, gives a unique (*C*-dimensional) probability distribution over the *C* candidate authors. By averaging the scores of the segments for the entire text, we obtain a unique probability distribution (i.e., a *C*-dimensional vector) for a text over the *C* candidate authors. The hyperparameter values that we used in the model are available in the supplementary material.

*SVM classifier***.** The second classifier is a standard SVM classifier; as we can see in Section 2 SVMs are used frequently for the problem of authorship attribution, We use as input features the entire set of features (i.e., the four classes of features) that we use for the ensemble classifier that we just described, using tf-idf weighting. We refer to this model as *SVM rich*.

5.2 Explainability

The second module provides clues to the philologist about why the model attributed a document to a given author. We decided to use the popular LIME framework (cf. Ribeiro *et al*. 2016) as it is versatile and it is able to provide simple and effective explanations to users who are not machine-learning experts. LIME, which stands for *local interpretable model-agnostic explainer*, is a method designed to explain the output of a classifier. In particular, it attempts to explain the output

of a black-box classifier for a given input by using a second *explainable model*[11] and training it to have the same output as the original model for the instance under examination and instances near it. In our case, as an explainable model we use ridge regression using a bag-of-words representation[12], and to create new input instances that are close to the one under consideration it removes random words.

To summarize, given a particular input text and the output of our ensemble model (a score for each candidate author reflecting the model's estimate of the likelihood that he authored the input text), the explainability module provides a potential explanation of the words that weigh towards each author. We can see an example of the output in Figure 6. To understand it, it shows that the model predicts Pucci as the most probable author (score=0.28) of the text *O lucchesi pregiati* followed closely by Villani (score=0.25). Focusing on Pucci, the words *siate* and *fe'* are the main words that appear in the text and increase the score towards the author, whereas the word *Castruccio* is the main word that decreases the score towards him. This output can be invaluable as it allows the user (the philologist) to use her expertise and evaluate whether the model is accurate.

5.3 Use of the System

The two modules that we described are the technical components of our system. They provide new tools to the philologist for analyzing old texts and making more informed decisions for their provenance. As we illustrate in Figure 3, the philologist interacts with these modules in multiple ways during her analysis task:

(1) She looks at the output of the classifier and evaluates how the results compare with the current theories.

(2) She observes justification that the LIME explainability model for the classifier's output

(3) With her knowledge in the field evaluates whether the justifications are sound. This allows us to tweak some of the parameters of the system. For example:

- Consider as stopwords words that should have no attributive value (but for instance may have been used frequently by a particular copyist).
- Add/remove some set of features to the classifier.
- Detect potential omissions or rooms for improvement.

In Section 7 we discuss the use of this process for our study in OI.

[11] Interpretable models are models that provide interpretations for their decisions. One common way that this is done is by assessing the importance of each input feature for the output value. Examples of interpretable models are linear regression, logistic regression, and decision trees.

[12] https://github.com/marcotcr/lime/

5.4 Experimental Setting

Our experimental setting consists of two sets of experiments. First, we used only the text with known authors to evaluate and compare our models and set the main hyperparameters. After deciding on the model, we use it on the unknown texts. We now describe these two phases.

Known texts. The first step of our approach is to use the known texts for training some classification models for authorship attribution. This process allows us to decide on the use of the ensemble model of Section 5.1 and on the hyperparameters of the model (how many levels, etc.). To do that we performed a stratified 5-fold cross validation: for each author under consideration, we created five folds making sure that for each author and each style (see the supplementary material) each fold contains the same number of input words to the largest extent possible.

Unknown texts. The second step is our ultimate goal, the assessment of authorhood likelihood for the unknown texts. For this we used the ensemble model with the hyperparameter settings that we selected from the experiments with the known texts. However, to set the rest of the weights we use the entire dataset of known texts.

Baselines. To decide on a model, in the measurement with the known text we performed a series of comparisons with two baseline models; the comparisons are both both quantitative and qualitative. The first baseline model is an SVM classifier, where we use standard single words as input features and tf-idf weighting; we call this model *SVM simple*. The second model is a logistic regression model, again using single words as input features and tf-idf weighting.

Evaluation measures. To evaluate the performance of the classifiers we use the *precision*, *recall*, *F1 measure* or simply *F1*, and *accuracy* measures, which are used typically in information retrievar. In the supplementary material you can find the precise formulae.

Qualitative analysis. This is another crucial step, where we applied the domain knowledge of a philologist, expert in the area (one of the authors) of OI. Whereas pure numerical scores are definitely indicative of the system's performance, they do not tell the entire story, especially when the amount of data available is limited. Then multiple problems may appear, such as overfitting. We therefore evaluated the results by observing the most important features that could lead to the classifiers' output, and in base of the findings we adjusted the parameters of our system. The explainability module (Section 5.2) is indispensable for this task. We discuss more about this in Section 7.

6. Results

We now present our experimental findings. We start by with the known texts, and we proceed with the unknown ones.

6.1 Known Texts

First, we present the numerical results of the known texts. We have summarized the scores of the four classifiers (presented in Sections 5.1 and 5.4) in Table 2.

Author	Logistic regression			SVM simple			Ensemble NN			SVM rich		
	pre	rec	F1	pre	rec	F1	pre	rec	F1	pre	rec	F1
Antonio Pucci	0.95	0.32	0.48	0.62	0.65	0.63	0.74	0.54	0.62	0.77	0.74	0.76
Antonio da Ferrara	0.94	0.92	0.93	0.91	0.89	0.90	0.88	0.96	0.92	0.88	0.87	0.87
Dante Alighieri	0.70	0.50	0.58	0.73	0.68	0.70	0.80	0.72	0.76	0.82	0.84	0.83
Dante da Maiano	0.92	0.95	0.93	0.96	0.85	0.90	0.82	0.91	0.86	0.95	0.88	0.91
Fazio degli Uberti	1.00	0.05	0.10	0.86	0.30	0.44	0.67	0.20	0.31	0.35	0.55	0.43
Francesco Petrarca	0.75	1.00	0.86	0.88	0.97	0.92	0.85	1.00	0.92	0.98	0.97	0.97
Franco Sacchetti	0.74	0.96	0.84	0.88	0.89	0.89	0.87	0.93	0.90	0.93	0.86	0.89
Giovanni Boccaccio	0.70	0.21	0.33	0.72	0.61	0.66	0.84	0.53	0.65	0.77	0.85	0.81
Giovanni Villani	1.00	1.00	1.00	1.00	1.00	1.00	0.87	1.00	0.93	0.61	1.00	0.75
Niccolò Soldanieri	0.95	0.63	0.76	0.75	0.90	0.82	0.78	0.83	0.81	0.96	0.90	0.93
Simone Serdini	0.00	0.00	0.00	0.67	0.80	0.73	1.00	0.20	0.33	0.00	0.00	0.00
Accuracy		0.77			0.84			0.85			0.88	

Table 2. Results on known texts.

Apart from the logistic model, the results of the other classifiers are similar. Looking at the SVM model, we can see that the richer feature set is able to improve the results. Restricting our attention to the classifiers with the richer set, we can see that depending on the author one may perform better than the other, whereas in the overall accuracy the SVM classifier has a slightly higher performance. Given that there is not a clear winner between the two models, we use both of them in the rest of our study.

6.2 Understanding the Results

Before applying the models on the unknown corpus, we give an example of the LIME explanations given for some texts. We have chosen an example from the ensemble NN classifier, which is able to showcase how the classifier is able to select some key features. It is a ballad by Antonio Pucci, *O lucchesi pregiati* (1370 approx.), which summarizes some of the main events occured in Lucca from 1313 to 1335. The model effectively recognizes Pucci's authorship and, at the same

time, identifies his source, Giovanni Villani's *Nuova Cronica*, by pinpointing not only grammatical and thematic words, but also anthroponyms (*Castruccio*) and toponyms (*Lucca*, *Pisa* etc.):

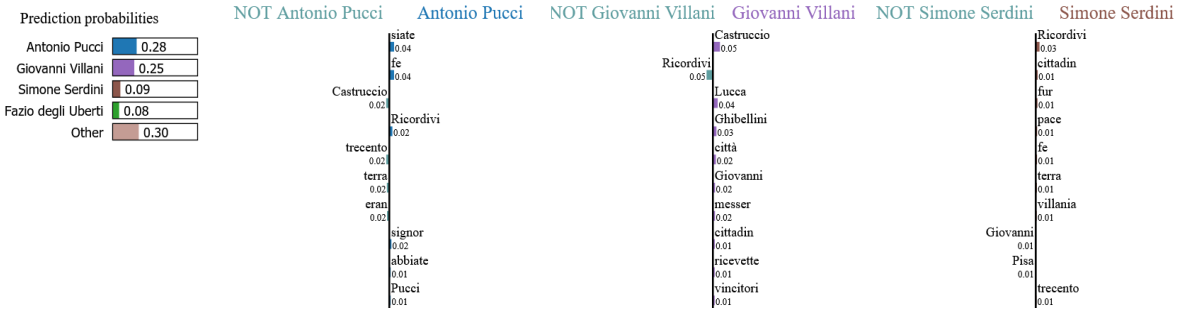

Figure 6. LIME output of Pucci's "O lucchesi pregiati."

In the supplementary material you can find two more examples that showcase the model's performance.

6.3 Unknown Texts

We are now ready to present the results for the seven unknown texts; we recall that Section 4.1 shows the most prevalent hypotheses for authorship. We display the results in Table 3.

Text	AP	AdF	DA	DdM	FdU	FP	FS	GB	GV	NS	SS
Andandosi la volpe	**0.182**	**0.141**	**0.189**	0.086	0.014	0.062	**0.248**	0.031	0.006	0.027	0.013
	0.591	0.010	0.047	0.005	0.005	0.007	**0.164**	**0.161**	0.004	0.003	0.002
Bel Gherardino	**0.257**	0.014	0.093	0.030	**0.215**	0.007	0.080	**0.237**	0.031	0.010	0.027
	0.202	0.030	**0.111**	0.017	**0.255**	0.023	0.040	**0.202**	0.054	0.037	0.037
Cantari Carduino	**0.117**	0.069	**0.217**	0.102	**0.174**	0.007	0.046	0.104	0.052	0.028	0.085
	0.070	0.011	0.041	0.005	**0.118**	0.007	0.019	**0.608**	0.083	0.016	0.021
Ciascun faccia per sé	**0.132**	**0.129**	0.021	0.090	0.005	0.097	**0.123**	0.012	0.013	**0.353**	0.025
	0.013	0.066	0.006	0.002	0.001	0.005	**0.887**	0.014	0.001	0.005	0.000
Detto d'Amore	**0.119**	0.042	**0.230**	**0.140**	**0.163**	0.056	0.125	0.025	0.004	0.083	0.014
	0.129	0.013	**0.505**	0.012	**0.151**	0.011	0.024	0.061	0.010	0.069	0.014
Fiore	**0.147**	0.121	**0.242**	0.077	0.047	0.015	0.084	**0.164**	0.028	0.051	0.023
	0.118	0.032	**0.284**	0.014	**0.159**	0.015	0.026	0.094	**0.110**	**0.121**	0.024
Quando il consiglio	**0.316**	0.043	**0.238**	0.043	**0.101**	0.5059	0.040	0.071	0.010	0.061	0.018
	0.381	0.015	**0.137**	0.007	0.008	0.092	**0.174**	**0.149**	0.011	0.019	0.008

Table 3. Output score for given for each text by our classifier to each author. For each text, the first line is the output of the ensemble NN classifier, and the second one the output of the SVM rich classifier. Each column corresponds to a candidate author: Antonio Pucci (AP), Antonio da Ferrara (AdF), Dante Alighieri (DA), Dante da Maiano (DdM), Fazio degli Uberti (FdU), Franco Sacchetti (FS), Francesco Petrarca (FP), Giovanni Boccaccio (GB), Giovanni Villani (GV), Niccolò Soldanieri (NS), Simone Serdini (SS).

In some cases the results corroborate some previous hypotheses in literature, especially about the *Detto d'Amore* and the *Fiore*; in other ones, they provide new investigative leads worthy of

consideration, as in the case of *Andandosi la volpe*, attributed to Sacchetti, and *Quando 'l consiglio*, attributed to Pucci. However, regardless of the authors' names provided (for some texts, as we said, the real authors may remain anonymous forever, if other manuscripts do not give any definitive proof), the model seems to work reasonably well, especially for the correct identification not only of the textual typology but also of the historical and cultural environment in which the unknown texts may have been conceived and produced. A good example of the first type (textual typology) are our analysis results of *Bel Gherardino* (text 2) that show interesting connections between the text and the work of the only two writers of *ottave* in the corpus: Giovanni Boccaccio and Antonio Pucci. An example of the second type (historical and cultural environment) is the outcome for *Ciascun faccia per sé* (text 4), for whom the analysis significantly identifies the same potential suspects selected by scholars.

Andandosi la volpe cum la gata (text 1): the caudate sonnet is likely to be written by Antonio Pucci or by one of Pucci's circle (Contini, 1938, pp. 290–291; see Section 4.1). The ensemble NN model assigns it to Franco Sacchetti, with the highest score (0.248), followed by Dante (0.189) and Pucci (0.182), whereas the SVM model attributes it more clearly to Pucci (0.591). Sacchetti belongs to Pucci's same cultural *côté*: he is one of the most representative authors of the Florentine expressive and playful poetics. Therefore, both results are coherent with Contini's observations ("Pucci or one of Pucci's circle") awaiting, of course, further and essential philological verification.

Bel Gherardino (text 2): although our features do not capture style, both the ensemble model and the SVM one identify effectively the only two authors of *cantari* included in the corpus: Pucci (0.257; 0.202) and Boccaccio (0.237; 0.202)[13]. In this case, as we said (see Section 4.1), it is not important the author identification, even if the results for *Bel Gherardino* allow us to reconsider an "old-fashioned" opinion by the Italianist Alessandro D'Ancona, who suggested that Pucci was the most likely author of the text (see Zambrini 1867, p. 17): It is not out of the question that the real author is not included in the corpus and that he will remain anonymous forever. Rather, it is important that the quantitative methods adopted seem to manage to efficiently detect Ol texts' lexical profile by pinpointing the most important grammatical and lexical features (e.g., the rich chivalric lexicon: *cortese* 'courtly,' *signor* 'sir,' *cavaliere* 'knight,' *donzella* 'amsel,' and so on).

Cantari di Carduino (text 3): as far as the *Cantari di Carduino* are concerned, our ensemble model

[13] Instead, Fazio's authorship (0.215) is likely to be excluded: at the present state of knowledge, he wasn't an author of *cantari*. Also in this case, a high score is assigned to an author who showed many intertextual connections with Dante, Boccaccio and Pucci himself.

catches the linguistic similarity between the text and the work of: Dante (0.217), Fazio degli Uberti (0.174) and Pucci (0.117), whereas the SVM model recognizes the important presence of Boccaccio's language (0.608). Despite more fragmented results, the ensemble effectively proposes a set of plausible authors too, such as Pucci. Moreover, although Dante and Fazio – the poets with the highest score – are not writers of *cantari*, their presence is explainable by the high number of intertextual connections between them (Lorenzi 2010). Hence, the triangle Dante-Fazio-Pucci identifies some of the most salient stylistic references of the text among those included in the corpus and gives some new suggestions for further linguistic research.

Ciascun faccia per sé (text 4): from the ensemble model's viewpoint, the analysis significantly confirms the rather high probability of Niccolò Soldanieri's authorship (0.353; see Section 4.3); the model assigns the ballad also to Pucci and Antonio da Ferrara, with lower scores (0.132; 0.129). It is especially interesting that the ensemble model's two candidate authors, Soldanieri and Pucci, are the most likely ones for the scholars too, in some cases without any documentary basis (the name of the author appears in none of the manuscripts: see Trucchi 1846), in other ones, instead, basing upon the testimony of manuscripts (cf. Corsi 1969; more recent edition of the text in Calvia 2017, page XXIII). As far as SVM model's results are concerned, also in this case they are more evident when compared to ensemble model's ones: Franco Sacchetti stands out as the most probable candidate for authorship (0.887).

Detto d'Amore and *Fiore* (texts 5-6): even though, in this case, scientific literature is very rich in proposals and there is still not a unanimous conclusion (see Section 4.3), both the ensemble model's predictions and the SVM model's ones significantly agree with the most accredited theory, which argues Dante's authorship (0.242 and 0.284 for the *Fiore*; 0.230 and 0.505 for the *Detto*). Among the other candidates selected by both models (for the *Fiore*: Pucci, Boccaccio, Fazio and Soldanieri; for the *Detto*: Fazio and, effectively, the minor poet Dante da Maiano), Dante da Maiano deserves a particular mention. As Stoppelli 2020 pointed out, who wrote the *Fiore* could be someone else than Dante, maybe a poet not far from the comic-realistic tradition as the same Dante da Maiano, "not so much for the compatibility of the name Durante (whose Dante is the hypocoristic form), [...] but more because of the quantity and quality of the exclusive references between the *Fiore* and Dante da Maiano's rhymes, as well as for the common centonistic technique of making verses."

And what happens if we try to analyze both texts together? We report here the results that confirm Dante's first place (0.249), followed by Pucci (0.143) and Boccaccio (0.139). In this case the model catches some very interesting and rarely documented words, such as *ipocristo* 'hypocrite or

Antichrist' (from the evangelical expression "hypocritae tristes" of *Matt*. 6, 16)[14].

Quando 'l consiglio (text 7): both the ensemble model and the SVM one assign the caudate sonnet to Pucci (0.316 and 0.381). The difference is instead in the "position" of Dante, who is the second most probable candidate for the ensemble model (0.238), but only the fourth one for the svm model (0.137), overcome by Sacchetti (0.174) and Boccaccio (0.149). As in the case of *Ciascun faccia per sé*, the group of potential suspects identified by scholars and that one defined by the two classifiers significantly coincide. In more detail, this result allows us to add some new elements to the recent hypothesis of Pucci's authorship (cf. Cupelloni 2018), proposed on the basis of important previous observations (see Contini 1995, page 278; Marti 2002, pages 516-518; Giunta 2014, page 572), with the recognition of some key words such as *cornacchia* 'crow', *gonnella* 'skirt', *dischiuda* 'it discloses', etc.:

7. Conclusion

Computer-aided authorship attribution is one of the newest and most interesting fields for 21th medieval philologists. In this article we presented a methodological proposal that integrates a new quantitative approach into a traditional qualitative analysis, leaving the philologist overview the entire building choices by consulting the output of the system. We have applied our framework on Italian literature of the 13th and 14th century, choosing a sample of seven unknown texts for which the author is unknown, but the name of Antonio Pucci has been proposed in literature. As one of the prime examples of disciplinary intersections in authorship attribution of early Italian texts, this kind of analysis confirms to be a quite reliable and promising tool, but only with the recommended Continian "great caution" (cf. Contini 1977). Indeed, the findings presented in Section 6 show that computational approaches can be a very useful tool in the arsenal of philologists but only when critically applied. Such approaches can capture subtle relationships between the terms or the word structures that are used by the various authors and propose answers for candidate authors. For example, the analysis of *Bel Gherardino*, written in *ottava rima* gave as candidate authors the two authors in our corpus who have effectively practiced this metrical form. Whereas such a deduction may be rather immediate for a philologist, this is by no means immediate for our classifiers: input features fed to the classifiers (Section 5.1) capture only the words (or parts of the words) and are completely agnostic on the style; yet the output shows that there are correlations between the style of a text and the terms used by the author which the classifiers are able to detect. Such findings are not obvious *a priori* and can provide new hints for authorship attribution. Here we want to reaffirm the importance of the triangle Philologist–

[14] For this word, documented only in the *Fiore*, see TLIO, s.v. *ipocristo*.

Classifier–Explainability (Section 5.3). Our system is the product of an iterative process in which we would observe the classifier and the explanation results and perform adjustments. The presence of the expert who overviews the whole process is crucial for a number of reasons:

(1) When a large volume of data is available, then a sophisticated classifier is able to learn what features are important, how they may interact with each other, and so on. For Early Italian texts the amount of text available is very limited (see Section 1.1). In such cases, feature selection is indispensable (Duboue 2020) and can significantly assist the classification task.

(2) The main hypothesis in machine-learning tasks such as classification is that the distribution of the test set and of the test set is the same: this ensures that the classifier learnt using the test set can provide a reasonable output when applied on the test set. To the extent possible we tried to satisfy this hypothesis for the known texts: for each author we tried to create the 5 folds (Section 5.4) in a way that each fold contains a similar number of texts and of words of each style. (Note that avoiding doing this gave us inferior results). However, when dealing with the unknown texts, the situation is different. We used the entire corpus of known texts for training, but we have no guarantees that the unknown texts can be considered coming from the same distribution. The small number of texts makes the variance high; even worse, there may be confounding factors: consider a text written by an author X (who may belong in our corpus, but we have not identified). The fact that the identity of X is unknown is an indication that several signals that are present in the known texts may not be present in the text under consideration. Some of these signals may be related to the document content, breaking the statistical hypothesis of machine learning that we mentioned previously. For this reason, it is important for the expert to observe the classifier output and the explanations and to assess the results.

Regarding the classifiers, we maintained those two that gave the best results: SVM and Ensemble NN. Linear SVMs have a long history in text classification because of their high performance (Anagnostopoulos *et al.* 2009), demonstrated also in the authorship attribution for Italian texts (Canettieri 2013; Italia and Canettieri 2013). In the last 3–4 years, the transformer neural network architecture has become one of the methods of choice for text-related tasks. Our ensemble classifier is based on the transformer architecture; however, given that training a transformer requires orders of magnitude more data, we reduced it in complexity to reduce the number of model parameters and lower the chance for overfitting.

The application of both classifiers allows us to ponder on the most appropriate methodologies for early Italian texts. For example, if we make a comparison of the two models based on the results for one of the unknown sonnets examined here, *Quando 'l consiglio*, SVM and ensemble NN recognize both the

fingerprints of the poet Antonio Pucci, the same candidate author proposed by some scholars on the basis of strictly philogical, stylistic and linguistic observations (Contini 1995, page 278; Marti 2002, pages 516–518; Giunta 2014, page 572). But things change if we consider the other candidates identified, for whom the scores are significantly different (Dante for the Ensemble model; Boccaccio and Sacchetti for SVM). The results for this case study show clearly that the Ensemble NN classifier can be considered more valuable than SVM. In other words, it makes more innovative recommendations by recognizing the hand of an author with less high but more accurate scores as compared to SVM's ones. In any case, the reflections on benefits and shortcomings of two major classifiers represent a step forward in enhancing statistical tools able to support and verify the philologist's hypothesis, combining with the already well experimented "corpus interrogation" (Section 3) so as to bring further empirical basis to the analysis of an unknown text. For future work, we would like to incorporate as input features some structures that are known to be important for attribution purposes. Our current approach uses as features mainly the words that appear in the texts and manipulations of them; instead, we intend to include features that capture the style of a text and the rhyme pattern of the words. Even though just by using content features the system was able in some cases to detect the writing style of the authors, if we explicitly provide such information to the classifiers they may make more informed decisions. Data science has transformed many scientific disciplines, and we believe that data-science frameworks such as this one may offer new directions to philology as well, but only if critically applied as a support and an assisting tool and not as a replacement of refined methods.

References

Anagnostopoulos et al. 2009
Aris Anagnostopoulos - Andrei Broder - Kunal Punera, *Effective and efficient classification on a search-engine model*, «Knowledge and Information Systems», 16.2 (2009), pp. 129–154.

Barber 1985
Joseph A. Barber, *Prospettive per un'analisi statistica del Fiore*, «Revue des études italiennes», 32 (1985), pp. 5–24.

Barber 1990
Joseph A. Barber, *A statistical analysis of the Fiore*, «Lectura Dantis», 6 (1990), pp. 100–122.

Barolini 2006
Teodolinda Barolini, *Dante and the Origins of Italian Literary Culture*, New York: Fordham University Press, 2006.

Beltrami 2002
Pietro G. Beltrami, *Gli strumenti della poesia*, Bologna: il Mulino, 2002.

Bendinelli Predelli 2017
Antonio Pucci, *Cantari della guerra di Pisa*, a cura di Maria Bendinelli Predelli, Firenze: Società Editrice Fiorentina, 2017.

Böhning 1992
Dankmar Böhning, *Multinomial logistic regression algorithm*, «Annals of the Institute of Statistical Mathematics», 44.1 (1992), pp. 197–200.

Brown 2020
Josh Brown, *Language history from below: Standardization and Koineization in Renaissance Italy*, «Journal of Historical Sociolinguistics», 6.1 (2020), online. [DA VERIFICARE pp./doi o e-location]

Calvia 2017
Nicolò del Preposto, *Opera Completa. Edizione critica commentata dei testi intonati e delle musiche*, a cura di Antonio Calvia, Firenze: Edizioni del Galluzzo per la Fondazione Ezio Franceschini, 2017.

Canettieri 2011
Paolo Canettieri, *Il Fiore e il Detto d'Amore*, «Critica del Testo», 14.1 (2011), pp. 519–530.

Canettieri 2013
Paolo Canettieri, *Le impronte digitali dell'autore. Un metodo di attribuzione automatizzata per i testi delle letterature romanze*, «Le forme e la storia», 6.2 (2013), pp. 229–243.

Canettieri 2016
Paolo Canettieri, *Chi non ha scritto il Fiore*, in *Sulle tracce del "Fiore"*, a cura di N. Tonelli, Firenze: Le Lettere, 2016, pp. 121–134.

Canettieri et al. 2008
Paolo Canettieri - Vittorio Loreto - Marta Rovetta - Giovanna Santini, *Philology and Information Theory*, «Cognitive Philology», 1 (2008), online.

Castellani 1952
Arrigo Castellani, *Nuovi testi fiorentini del Dugento*, vol. I, Firenze: Sansoni, 1952.

Castets 1881
Ferdinand Castets, *Il "Fiore", poème italien du XIIIe siècle, en CCXXXII sonnets, imité du "Roman de la Rose" par Durante. Texte inédit publié avec facsimile*, Montpellier-Paris: Société pour l'étude des langues romanes - Maisonneuve, 1881.

Ciociola 2011
Claudio Ciociola, *Sestina*, in *Enciclopedia dell'Italiano*, Roma: Istituto della Enciclopedia Italiana Treccani, 2011, online, https://www.treccani.it/enciclopedia/sestina_%28Enciclopedia-dell%27Italiano%29/.

Ciotti 2014
Digital Humanities: progetti italiani ed esperienze di convergenza, a cura di Fabio Ciotti, Roma: Sapienza Università Editrice (per AIUCD), 2014.

Contini 1938
Gianfranco Contini, *Un manoscritto ferrarese quattrocentesco di scritture popolareggianti*, «Archivum Romanicum», 22 (1938), pp. 281–319.

Contini 1977
Gianfranco Contini, *Filologia*, in *Enciclopedia del Novecento*, Roma: Istituto della Enciclopedia Italiana Treccani, 1977, online, https://www.treccani.it/enciclopedia/filologia_%28Enciclopedia-del-Novecento%29/.

Contini 1984
Il "Fiore" e il "Detto d'amore" attribuibili a Dante Alighieri, a cura di Gianfranco Contini, Milano: Mondadori, 1984.

Contini 1995
Dante Alighieri, *Rime*, a cura di Gianfranco Contini, Torino: Einaudi, 1995.

Corbara et al. 2019
Silvia Corbara - Alejandro Moreo - Fabrizio Sebastiani - Mirko Tavoni, *The Epistle to Cangrande Through the Lens of Computational Authorship Verification*, in *New Trends in Image Analysis and Processing – ICIAP 2019*, edited by M. Cristani et al., Switzerland: Springer Nature, 2019, pp. 148–158.

Corsi 1969
Rimatori del Trecento, a cura di Giuseppe Corsi, Torino: UTET, 1969.

Cortelazzo et al. 2013
Michele A. Cortelazzo - Paolo Nadalutti - Arjuna Tuzzi, *Improving Labbé's intertextual distance: testing a revised version on a large corpus of Italian literature*, «Journal of Quantitative Linguistics», 20.2 (2013), pp. 125–152.

Cortelazzo et al. 2016
Michele A. Cortelazzo - Paolo Nadalutti - Stefano Ondelli - Arjuna Tuzzi, *Authorship Attribution and Text Clustering for Contemporary Italian Novels*, in *QUALICO 2016 book of abstracts*, Germany: University of Trier, 2016.

Cortelazzo - Tuzzi 2017
Michele A. Cortelazzo e Arjuna Tuzzi, *Sulle tracce di Elena Ferrante: questioni di metodo e primi risultati*, in *Testi, corpora, confronti interlinguistici: approcci qualitativi e quantitativi*, a cura di G. Palumbo, Trieste: EUT Edizioni Università di Trieste, 2017, pp. 11–24.

Cortelazzo - Tuzzi 2018
Michele A. Cortelazzo e Arjuna Tuzzi, *What is Elena Ferrante? A comparative analysis of a secretive bestselling Italian writer*, «Digital Scholarship in the Humanities», 33.3 (2018), pp. 685–702.

Cortes - Vapnik 1995
Corinna Cortes e Vladimir Vapnik, *Support vector machine*, «Machine Learning», 20.3 (1995), pp. 273–297.

Cupelloni 2018
Francesca Cupelloni, *Sull'attribuzione del sonetto Quando il consiglio*, «Rivista di studi danteschi», 18.1 (2018), pp. 184–200.

Cupelloni 2019
Francesca Cupelloni, *Sul lessico del Centiloquio. Primi assaggi*, «Bollettino di italianistica», 2 (2019), pp. 38–52.

Cupelloni 2021
Francesca Cupelloni, *Metodi non tradizionali di filologia attributiva. Bilanci e prospettive di ricerca*, «Ecdotica», 18 (2021), pp. 81–102.

Cupelloni 2022
Francesca Cupelloni, *La lingua di Antonio Pucci. Indagini su lessico, sintassi e testualità*, Firenze: Cesati, 2022.

Custódio - Paraboni 2018
José Eleandro Custódio e Ivandré Paraboni, *EACH-USP ensemble cross-domain authorship attribution*, in *Working Notes of CLEF 2018 – Conference and Labs of the Evaluation Forum*, 2018.

De Giorgi 2013
Fulvio De Giorgi, *Contini nel mondo rosminiano*, «Annali della Scuola Normale Superiore di Pisa. Classe di Lettere e Filosofia», 5.2 (2013), pp. 599–688.

Duboue 2020
Pablo Duboue, *The Art of Feature Engineering: Essentials for Machine Learning*, Cambridge: Cambridge University Press, 2020.

Erdman - Fogel 1966
David V. Erdman e Ephim G. Fogel, *Evidence for Authorship: Essay on Problems of Attribution*, Ithaca: Cornell University Press, 1966.

Fasani 1973
Remo Fasani, *Ancora per l'attribuzione del Fiore al Pucci*, «Studi e problemi di critica testuale», 6 (1973), pp. 22–68.

Fasani 1974–1975
Remo Fasani, *Il Fiore e la poesia del Pucci*, «Deutsches Dante-Jahrbuch», 49–50 (1974–1975), pp. 81–141.

Fasani 2008
Remo Fasani, *Il "Fiore" e il "Detto d'amore" attribuiti a Immanuel Romano*, Ravenna: Longo, 2008.

Ferracane - Wang - Mooney 2017
Elisa Ferracane - Su Wang - Raymond Mooney, *Leveraging discourse information effectively for authorship*

attribution, in *Proceedings of the Eighth International Joint Conference on Natural Language Processing (IJCNLP 2017)*, 2017, pp. 584–593. [DA VERIFICARE: luogo/editore]

Ferri 1909
Antonio Pucci, *La poesia popolare di Antonio Pucci*, a cura di Ferruccio Ferri, Bologna: Libreria L. Beltrami, 1909.

Frati 1913
Le "Rime" del Codice Isoldiano (Bologn. Univ. 1739), a cura di Lodovico Frati, Bologna: Romagnoli-Dall'Acqua, 1913.

Formisano 2012
Opere di dubbia attribuzione. Il Fiore e il Detto d'Amore, a cura di Luciano Formisano, Roma: Salerno Editrice, 2012 (Nuova Edizione Commentata delle Opere di Dante, 7).

Giunta 2004
Claudio Giunta, *Sul rapporto tra prosa e poesia nel Medioevo e sulla frottola*, in *Storia della lingua e filologia. Per Alfredo Stussi nel suo sessantacinquesimo compleanno*, a cura di M. Zaccarello e L. Tomasin, Firenze: Edizioni del Galluzzo, 2004, pp. 35–72.

Giunta 2014
Dante Alighieri, *Rime*, a cura di Claudio Giunta, Milano: Mondadori, 2014.

Heijkant 1997
Marie-José Heijkant, *"De re Artu parlar cio che si scrisse". Orality in the Breton "cantari"*, «Amsterdamer Beiträge zur Älteren Germanistik», 48 (1997), pp. 57–70.

Italia - Canettieri 2013
Paola Italia e Paolo Canettieri, *Un caso di attribuzionismo novecentesco: il Diario Postumo di Montale*, «Cognitive Philology», 6 (2013), online.

Ivanov 2019
Lubomir Ivanov, *Learning Patterns of Assonance for Authorship Attribution of Historical Texts*, in *Proceedings of the Thirty-Second International FLAIRS Conference (FLAIRS-32)*, 2019, pp. 191–196.

Jafariakinabad - Tarnpradab - Hua 2019
Fereshteh Jafariakinabad - Sansiri Tarnpradab - Kien A. Hua, *Syntactic recurrent neural network for authorship attribution*, «arXiv», arXiv:1902.09723 (2019).

Kennedy - Miceli 2001
Claire Kennedy e Tinizia Miceli, *An evolution of intermediate students' approaches to corpus investigation*, «Language Learning & Technology», 5.3 (2001), pp. 77–90.

Kestemont 2012
Mike Kestemont, *Stylometry for Medieval Authorship Studies: An Application to Rhyme Words*, «Digital Philology: A Journal of Medieval Cultures», 1.1 (2012), pp. 42–72.

Kingma - Ba 2015
Diederik P. Kingma e Jimmy Ba, *Adam: A method for stochastic optimization*, in *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*, 2015, arXiv:1412.6980.

Kleinhenz 2000
Cristopher Kleinhenz (ed.), *Medieval Italy: An Encyclopedia*, New York: Routledge, 2000.

Kudo - Richardson 2018
Taku Kudo e John Richardson, *SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing*, in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, arXiv:1808.06226.

LIRIO = Leonardi, Lino *et alii*, editors. 2013. *Corpus della poesia lirica italiana delle origini. Dagli inizi al 1400*, Edizioni del Galluzzo per la Fondazione Ezio Franceschini, Firenze, http://lirioweb.ovi.cnr.it.

LIZ = Stoppelli, Pasquale, and Eugenio Picchi. *LIZ. Letteratura Italiana Zanichelli, cd-rom dei testi della letteratura italiana*, Zanichelli, Bologna.

Loomis 1925
Roger Sherman Loomis, *Medieval Iconography and the Question of Arthurian Origins*, «Modern Language Notes», 40.2 (1925), pp. 65–70.

Lorenzi 2010
Cristiano Lorenzi, *Echi danteschi (e stilnovistici) nelle Rime di Fazio degli Uberti*, «Italianistica», 39.2 (2010), pp. 95–113.

Luvisotto 2020
Alessia Luvisotto, *Chronicle of the seminar "Attributive philology and digital analysis of literary texts"*, 2020, online (www.humanitiesforchange.org).

Marti 2002
Mario Marti, *L'edizione nazionale delle "Rime" di Dante Alighieri*, «Giornale Storico della Letteratura Italiana», 179.588 (2002), pp. 511–524.

Mazzoni 1923
Guido Mazzoni, *Il "Fiore" e il "Detto d'Amore" attribuiti a Dante Alighieri. Testo del XIII secolo*, Firenze: Alinari, 1923.

Miranda-García - Calle-Martín 2008
Antonio Miranda-García e Javier Calle-Martín, *A Survey of Non-traditional Authorship Attribution Studies*, «Ecdotica», 5 (2008), pp. 47–168.

Miranda-García - Calle-Martín 2012
Antonio Miranda-García e Javier Calle-Martín, *Stylometry and Authorship Attribution: Introduction to the Special Issue*, «English Studies», 93.3 (2012), pp. 251–258.

Moretti 2013
Franco Moretti, *Distant Reading*, London: Verso, 2013.

Muttenthaler - Lucas - Amann 2019
Lukas Muttenthaler - Gordon Lucas - Janek Amann, *Authorship Attribution in Fan-Fictional Texts given variable length Character and Word N-Grams*, in *Working Notes of CLEF 2019*, edited by L. Cappellato - N. Ferro - David E. Losada - Henning Müller, Lugano, Switzerland, 2019.

Ouamour - Sayoud 2012
Siham Ouamour e Halim Sayoud, *Authorship attribution of ancient texts written by ten Arabic travelers using a SMO-SVM classifier*, in *Proceedings of the 2012 International Conference on Communications and Information Technology (ICCIT)*, 2012, pp. 44–47.

Ribeiro et al. 2016
Marco Tulio Ribeiro - Sameer Singh - Carlos Guestrin, *Why should I trust you?: Explaining the predictions of any classifier*, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Diego, California, 2016, pp. 97–101.

Robey 1997
David Robey, *The Fiore and the Comedy: some Computerized Comparisons*, in *The Fiore in Context. Dante, France, Tuscany*, edited by Z. G. Baranski e P. Boyde, Notre Dame-London: University of Notre Dame Press, 1997, pp. 109–134.

Roelli 2020
Philipp Roelli, *Handbook of Stemmatology: History, Methodology, Digital Approaches*, Berlin/Boston: De Gruyter, 2020.

Ruggiero 2020
Fabio Ruggiero, *A sonnet to be re-attributed to Cino?*, «Nuova rivista di letteratura italiana», 23.1 (2020), pp. 11–31.

Salami - Momtazi 2020
Dariush Salami e Saeedehand Momtazi, *Recurrent convolutional neural networks for poet identification*, «Digital Scholarship in the Humanities», 36.2 (2020), pp. 472–481.

Sapegno 1952
Poeti minori del Trecento, a cura di Natalino Sapegno, in *Letteratura Italiana Ricciardi*, Milano-Napoli: Ricciardi, 1952.

Sari et al. 2018
Yunita Sari - Mark Stevenson - Andreas Vlachos, *Topic or style? exploring the most useful features for authorship attribution*, in *Proceedings of the 27th International Conference on Computational Linguistics*, Santa Fe, New Mexico, USA, 2018, pp. 343–353.

Seroussi et al. 2014
Yaniri Seroussi - Ingrid Zukerman - Fabian Bohnert, *Authorship Attribution with Topic Models*, «Computational Linguistics», 40.2 (2014), pp. 269–310.

Shrestha et al. 2017
Prasha Shrestha - Sebastian Sierra - Fabio González - Manuel Montes - Paolo Rosso - Thamar Solorio,

Convolutional neural networks for authorship attribution of short texts, in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics. Volume 2: Short Papers*, 2017, pp. 669–674.

Solimena 1980
Adriana Solimena, *Repertorio metrico dello Stilnovo*, Roma: Società Filologica Romana, 1980.

Sommerschield et al. 2023
Thea Sommerschield - Yannis Assael - John Pavlopoulos - Vanessa Stefanak - Andrew Senior - Chris Dyer - John Bodel - Jonathan Prag - Ion Androutsopoulos - Nando de Freitas, *Machine Learning for Ancient Languages: A Survey*, «Computational Linguistics», 49.3 (2023), pp. 703–747.

Stamatatos 2017
Efstathios Stamatatos, *Authorship attribution using text distortion*, in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics. Volume 1: Long Papers*, 2017, pp. 1138–1149.

Stanghellini 2009
Cecco Angiolieri, *Il "Fiore". Nuove congetture testuali e interpretazioni*, a cura di Menotti Stanghellini, Monteriggioni: Il Leccio, 2009.

Stoppelli 2011
Pasquale Stoppelli, *Dante e la paternità del "Fiore"*, Roma: Salerno Editrice, 2011.

Stoppelli 2017
Pasquale Stoppelli, *Per un nuovo profilo di Dante da Maiano*, in *La poesia in Italia prima di Dante. Atti del Colloquio internazionale di Italianistica (Roma, 10–12 giugno 2015)*, a cura di F. Suitner, Ravenna: Longo, 2017, pp. 65–74.

Stoppelli 2020
Pasquale Stoppelli, *L'equivoco del nome. Rime incerte fra Dante Alighieri e Dante da Maiano*, Roma: Salerno Editrice, 2020.

TLIO = Beltrami, Pietro G., director, *TLIO. Tesoro della Lingua Italiana delle Origini*. Firenze: CNR-Opera del Vocabolario Italiano, http://tlio.ovi.cnr.it/TLIO.

Trucchi 1847
Francesco Trucchi, *Poesie italiane inedite di dugento autori dall'origine della lingua infino al secolo XVII*, Prato: Ranieri Guasti, 1847.

Vaswani et al. 2017
Ashish Vaswani - Noam Shazeer - Niki Parmar - Jakob Uszkoreit - Llion Jones - Aidan N. Gomez - Lukasz Kaiser - Illia Polosukhin, *Attention is all you need*, in *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS 2017)*, 2017, pp. 6000–6010.

Volpi 1907
Rime di trecentisti minori, a cura di Guglielmo Volpi, Firenze: Sansoni, 1907.

Zambrini 1867
Antonio Pucci, *Cantare del bel Gherardino: novella cavalleresca in ottava rima del secolo XIV*, a cura di Francesco Zambrini, Bologna: Gaetano Romagnoli, 1867.

Supplementary Material

Details on the NN Classifier (Supplementary to Section 5.1)

Here we present details on the four NN models of our classifier. Each of the four NNs has the same architecture, which is based on the Transformer model (Vaswani et al. 2017); see Figure 2. The first layer is an embedding layer, which, for each input token it creates an E-dimensional embedding (E=300). Next the input, represented as an $L\times E$-dimensional matrix, enters into a Transformer encoder with six heads. The output of the (single) encoder, an $L\times E$-dimensional matrix, is then averaged over all the tokens, giving an E-dimensional vector representing the embedding of the entire initial sequence. This vector goes through a fully-connected network with a SoftMax activation function, which gives a probability distribution over the C=11 potential authors of the segment. Thus, for each input segment we obtain four C-dimensional vectors.

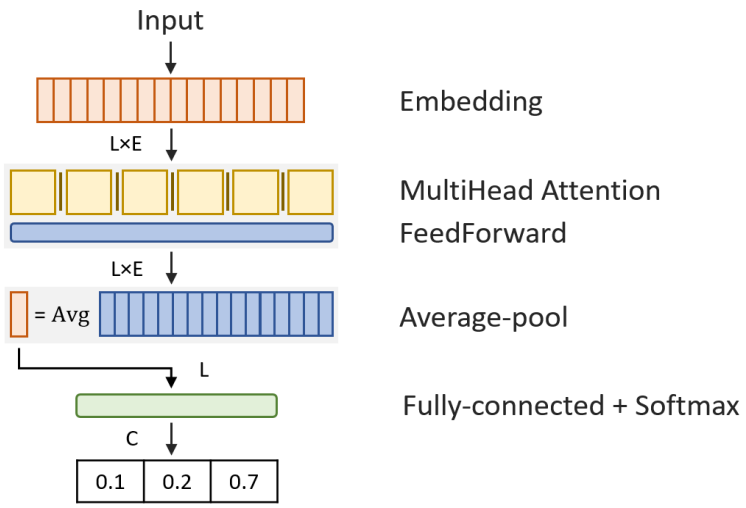

Figure 2. Classifier architecture.

Table 4 contains the combination of hyper-parameters that yielded the best validation score. Because we have a small dataset we added dropout layers with rate 0.4 for regularization. During training we used mini-batches of size B=512, using Adam optimizer (Kingma and Ba 2015) with learning rate 2e-05. We trained each model for a maximum of 50 epochs with early stopping and chose the models that maximized the sum of the scores of the outputs of the four models.

Hyper-parameter	Optimal value
Sequence length (L)	256
Vocabulary size	All tokens
Embedding size (E)	300
Number of heads (H)	6
Hidden size	300
Batch size (B)	512
Dropout	0.4
Learning rate	2e-5
Max. epochs	50

Table 4. The hyperparameters of our model.

Evaluation Measures (Supplementary to Section 5.4 Errore. L'origine riferimento non è stata trovata.**)**

To evaluate the performance of the classifiers we use the *precision*, *recall*, *F1 measure* or simply *F1*, and *accuracy* measures, which we now define. For a given author X, we first define:

TP_X: True positive: number of documents written by author X, for which the classifier gives to X the highest score

FP_X: False positive: number of documents not written by author X, for which the classifier does gives to X the highest score

TN_X: True negative: number of documents not written by author X, for which the classifier does not give to X the highest score

TN_X: False negative: number of documents written by author X, for which the classifier does not give to X the highest score

We also define

$$Precision_X = \frac{TP_X}{TP_X + FP_X},$$

$$Recall_X = \frac{TP_X}{TP_X + FN_X},$$

and

$$F1_X = \frac{2 \cdot Precision_X \cdot Recall_X}{Precision_X + Precision_X}.$$

Finally, we define the overall *accuracy* of a classifier to be equal to the sum, over all the eleven authors, of the true positives, divided by the total number, N, of documents in the test set:

$$Accuracy = \frac{1}{N} \cdot \sum_X TP_X$$

Overview of Early Italian Styles

With the risk of oversimplifying a complex topic but for the benefit of the readers who are not familiar with early-Italian writing, we provide a brief description of the styles found in our corpus[15].

1. *Sonnet*, *sonetto ritornellato*, *caudato*, *doppio*, and *rinterzato*: a sonnet usually consists of an octave (8 lines) and a sestet (6 lines) of hendecasyllables, variously rhyming (e.g., the most ancient scheme for the octave is ABABABAB and the most recurring one in the 14th century is ABBAABBA). The four variations included in our corpus, *ritornellato, caudato, doppio*, and *rinterzato*, constitute an extension of the sonnet's traditional form: the first type (*ritornellato*) displays a concluding couplet (*ritornello doppio*); the second one (*caudato*) a three-line *coda* (a seven-syllable line and a rhyming couplet); in the third one (*doppio*) a seven-syllable line is added after the odd lines of the quatrains and after the middle line of the triplets (as in Dante's *Morte villana*: AaBBbA AaBBbA CDdC CDdC), whereas in the *rinterzato* a seven-syllable line usually follows the first verse of each triplet (but there are different cases, as in the twenty-four-verse type of sonetto rinterzato *Quando 'l consiglio degli ucce' si tenne*, one of the unknown texts: AaBBbA AaBBbA CcDDd CCcD dEE)[16];

2. *Canzone* and *sestina*: The traditional *canzone* or *canzone petrarchesca*, the most prestigious meter in medieval Italian poetry, has a variable number of stanzas; the conclusion is usually relied on a shorter stanza called *congedo*. The sestina, used most notably by Dante and Petrarca, is a type of *canzone* made up of six stanzas of six hendecasyllables: the six rhyming words of the first stanza always and necessarily return in the five subsequent stanzas, according to a principle of rotation that takes the name of *retrogradatio cruciata* (see Ciociola 2011);

3. *Madrigal*: The madrigal is a short lyric form that consists of a group of 2 to 5 tercets followed by a couplet; Petrarca's four madrigals are among the finest of the 14th century;

4. *Ballad* and *lauda*: Originally intended for music performance, ancient ballads show a variable structure (*zagialesca*, Italian ballad, etc.) whose essential feature is the *ripresa*, that is, a refrain that constantly precedes the text. The *lauda* is a type of popular devotional ballad with a religious topic, as in Pucci's *Veggendo ber Gesù aceto e fele*;

5. *Terza rima*: The rhyme scheme of the so called *terza rima* is ABA BCB CDC…, normally all hendecasyllables; the widespread use of this style can be explained by the enthusiastic reception of Dante's *Commedia* (e.g., in Petrarca's *Trionfi*) and it can also appear in the form of a historical-chronachistic chapter with an informative and propagandistic intent, as in Pucci's *Centiloquio*.

6. *Frottola*: The *frottola* is a lyrical form that, differently from the other styles we present here, does not have a strict structure: its only structural principle is given by the rhyme, iterated for two, three or more times (see Giunta 2004), as in Franco Sacchetti's *Chi drieto va* and *La lingua nova*;

7. *Serventese caudato*: a popular metric form with stanzas composed of a series of "long lines" (hendecasyllables, octosyllables, etc.), rhymed with each other, and a shorter concluding verse; the most recurring form is AA...b BB...c CC...d;

---

[15] The main point of reference for the descriptions is Beltrami 2002.

[16] The interpretation of this sonnet's rhyme scheme is still in doubt (see Solimena 1980, page 14; Cupelloni 2018, pages 187–188).

8. *Ottava rima*: a stanza of all rhymed hendecasyllables (ABABABCC) used for *cantari* and epic poems; one of the oldest examples of this poetic form included in the corpus is Boccaccio's *Filostrato*.

Additional Examples of LIME Explanations (Supplementary to Section 6.2)

A second text is a poem by Niccolò Soldanieri, *Dato che fu a questo mondo il lume* (second half of the 14th century), a didactic *canzone* rich in moral precepts. Also in this case, the model identifies the right author with a very high score, followed by two "minor poets" (see Section XXX) belonging to his same cultural background (the two "Antonios", Pucci and da Ferrara):

Figure 3. LIME output of Soldanieri's "Dato che fu a questo mondo il lume."

The last example concerns Dante's rhymes, considered as a whole. In this case the model not only identifies Alighieri's authorship with a relevant score, but seems to catch the linguistic difference between thirteenth-century Florentine (the language of Dante Alighieri and Dante da Maiano) and fourteenth-century Florentine (the most represented authors of this century are, significantly, the other two "crowns of Florence", Boccaccio and Petrarca):

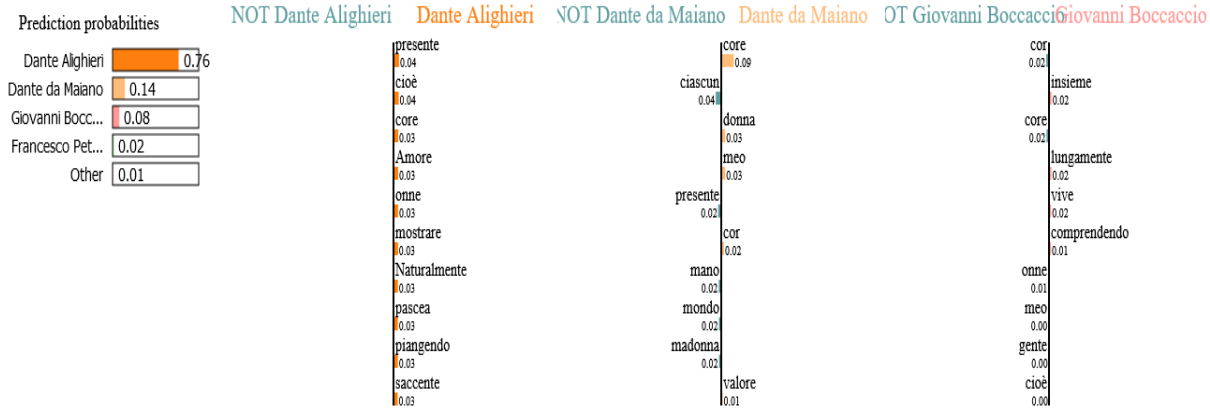

Figure 4. LIME output of Dante's rhymes.

LIME Explanations for the Unknown Texts (Supplementary to Section 6.3)

Andandosi la volpe cum la gata (text 1):

Bel Gherardino (text 2):

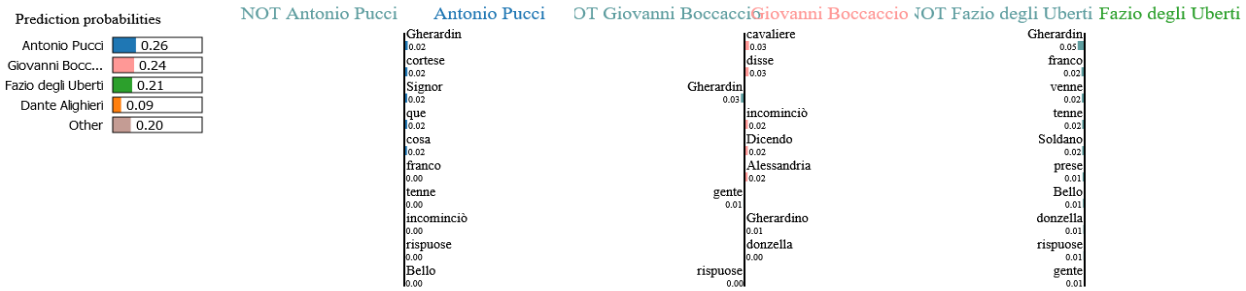

Cantari di Carduino (text 3):

Ciascun faccia per sé (text 4):

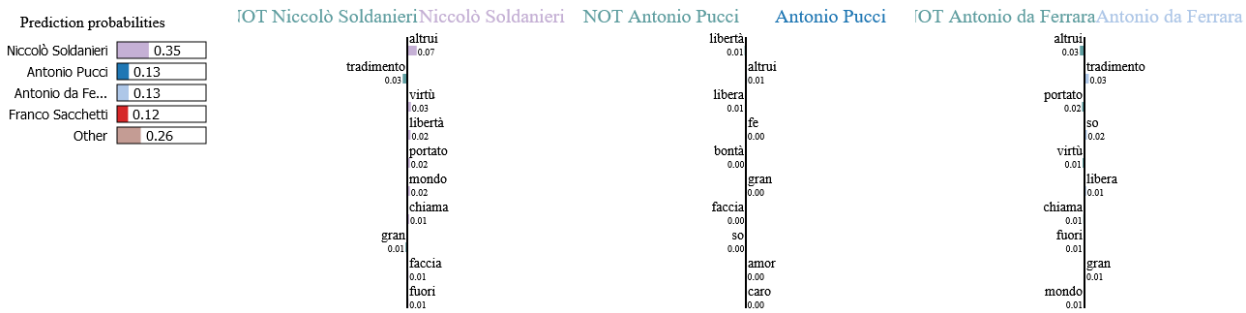

Detto d'Amore (text 5):

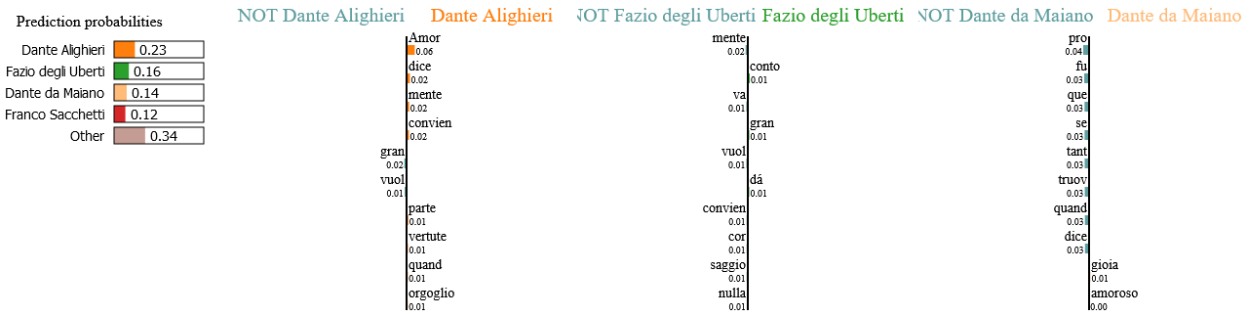

Fiore (texts 6):

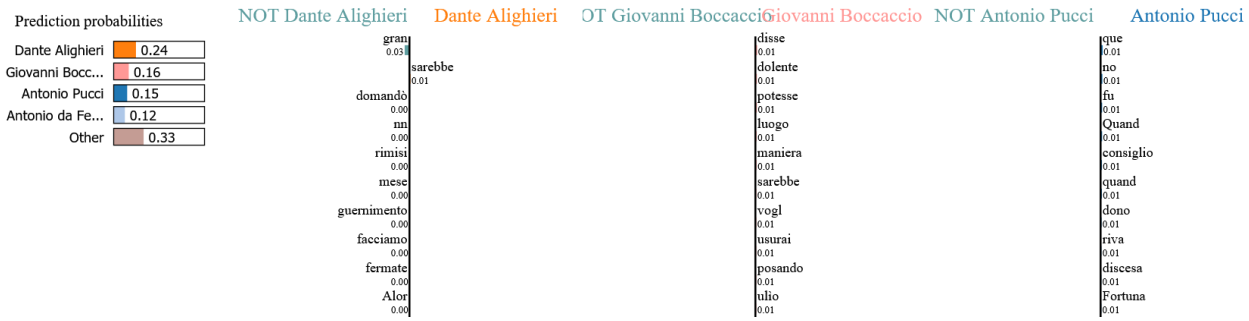

Detto d'Amore and *Fiore* combined (texts 5-6):

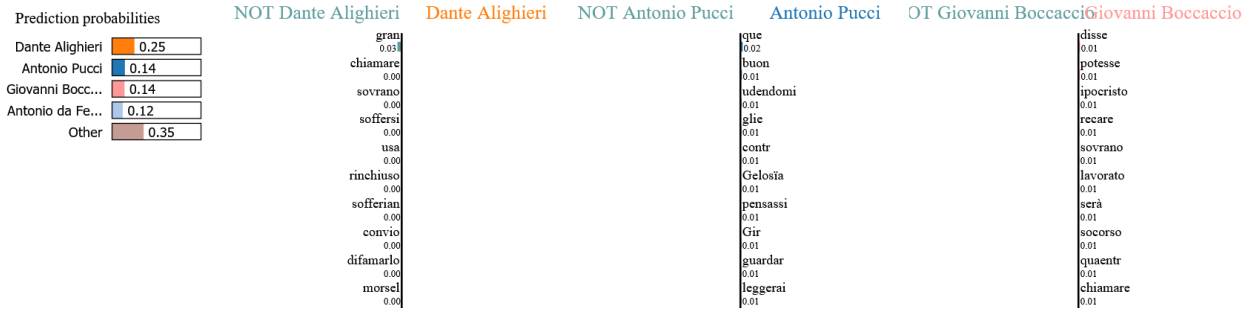

Quando 'l consiglio (text 7):