

(Stem and Word) Predictability in Italian verb paradigms: An Entropy-Based Study Exploiting the New Resource *LeFFI*

Matteo Pellegrini¹, Alessandra Teresa Cignarella^{2,3}

1. Liceo Statale “Augusto Monti” di Chieri, Italy

2. Dipartimento di Informatica, Università degli Studi di Torino, Italy

3. PRHLT Research Center, Universitat Politècnica de València, Spain

pellegrini.matteo@liceomonti.edu.it, cigna@di.unito.it

Abstract

English. In this paper we present *LeFFI*, an inflected lexicon of Italian listing all the available wordforms of 2,053 verbs. We then use this resource to perform an entropy-based analysis of the mutual predictability of wordforms within Italian verb paradigms, and compare our findings to the ones of previous work on stem predictability in Italian verb inflection.

1 Introduction

The pioneering work of Aronoff (1994) has inspired an influential line of research where predictability within inflectional paradigms is modelled by resorting to the notion of morphomic stems – i.e., stems that cannot be considered as bearing any meaning, as they appear in groups of cells that do not share a fixed morphosyntactic content. In this perspective, every lexeme is seen as equipped with a set of indexed stems, that only for regular lexemes are mutually predictable, while for irregular verbs they need to be independently stored. From each of these stems, a fixed set of wordforms can be obtained by adding the appropriate inflectional endings. An analysis relying on these assumptions was proposed by Maiden (1992) and subsequent work – see Maiden (2018) for a recent survey – to account for the patterns of stem allomorphy that are found in the verbal inflection of Romance languages in general. More detailed implementations of these ideas have then been provided for individual languages, among them Italian (Pirrelli and Battista, 2000; Montermini and Boyé, 2012; Montermini and Bonami, 2013). Another possibility that has been explored in more recent times is tackling the issue of inflectional predictability in terms of predictions of wordforms from one another, without

assuming a given segmentation in stems vs. endings, in a fully word-based, abstractive (Blevins, 2016) approach. Within this framework, Ackerman et al. (2009) propose to estimate the reliability of inflectional predictions by means of the information-theoretic notion of conditional entropy. Building on this work, Bonami and Boyé (2014) outline a procedure that allows to compute entropy values estimating the uncertainty in predicting one cell from another one directly from a lexicon of fully inflected wordforms in phonological transcription, using the type frequency of different inflectional patterns to estimate their probability of application. This method has been applied to French by Bonami and Boyé (2014), to Latin by Pellegrini (2020), and it has been used for typological comparison on a small sample of languages by Beniamine (2018), who also provides a freely available toolkit (Qumin) allowing to perform this computation automatically for any language.

A similar entropy-based analysis has not been proposed for Italian yet. To be able to use the Qumin toolkit to perform it, it is necessary to have an inflected lexicon listing all the wordforms of a representative number of lexemes in phonological transcription, like e.g. Flexique for French (Bonami et al., 2014) or LatInflexi for Latin (Pellegrini and Passarotti, 2018). Looking for such a resource for Italian, we can see that in most lexicons wordforms are given in orthographic transcription – see e.g. Morph-it! (Zanchetta and Baroni, 2005) and CoLFIS (Bertinetto et al., 2005). On the other hand, in PhonItalia (Goslin et al., 2014) there are phonological transcriptions, but not all the inflected wordforms of each lexeme are listed. To the best of our knowledge, the only resource providing phonological transcriptions of the full paradigm of lexemes is GLAFF-IT (Calderone et al., 2017), but due to the way in which it was created, it proves to be too noisy to be used for entropy computations as such.

In this paper, we describe the work that was done to obtain a smaller, but cleaner version of GLAFF-

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

IT. We then use this resource to perform an entropy-based analysis of predictability in Italian verb inflection. After briefly describing the methodology, we present our results comparing them with the findings of previous stem-based analyses.

2 The Resource

In order to build *LeFFI (Lessico delle Forme Flesse dell'Italiano)*, we have firstly consulted GLAFF-IT, a free machine-readable dictionary based on *Wikizionario*, the Italian language edition of Wiktionary. It is a morphophonological Italian lexicon which contains a total of 485,135 word-forms among verbs, nouns, adjectives and adverbs, in both orthographic and phonological IPA transcription. Since our interest for the present research lies only in verbs, in this step a total of 411,770 verbal forms in phonological transcription have been extracted from GLAFF-IT, together with the citation form (the infinitive) of the lexeme they belong to, thus resulting in a list of the complete paradigms of 7,552 verbs. To indicate the morphosyntactic properties expressed by each wordform, we use the notation of the Leipzig Glossing Rules (Comrie et al., 2008), both in our resource and in the examples shown in this paper.

Due to the large amount of manual work needed in order to obtain our resource, for the time being we have decided to focus only on a fraction of this list. So as not to lose quantitatively relevant data, our selection was based on the frequency of lexemes, as reported in the CoLFIS frequency lexicon. We have thus crossed the list of 7,552 verbs extracted from GLAFF-IT with the 5,193 verbal lexemes contained in CoLFIS, and kept only the ones with a frequency higher than 10. The resulting dataset, listing the 53 available, non-periphrastic cells of 2,053 verbs, is still large enough to allow for reasonably safe generalizations on Italian verb inflection.

After these automatic steps, several manual changes have been made in order to obtain the current version of our resource. Firstly, it should be noticed that many of the phonological transcriptions provided by GLAFF-IT are obtained automatically from the orthographic form. In some cases, however, it is not possible to infer a precise phonological transcription from orthography alone, because some graphemes can correspond to different phonemes. In such cases, the phonological transcriptions provided by GLAFF-IT are underspec-

ified: for instance, the symbol E is used for the grapheme ⟨e⟩, that can correspond to /e/ or /ɛ/, and similarly O for ⟨o⟩ (/o/ or /ɔ/), S for ⟨s⟩ (/s/ or /z/), Z for ⟨z⟩ (/ts/ or /dz/). While we have manually reconducted ⟨s⟩, ⟨z⟩ and a few other marginal ambiguous graphemes to the actual phonemes they correspond to, for ⟨e⟩ and ⟨o⟩ we have decided to keep the same neutralization as in GLAFF-IT. This choice is due to the fact that manually disambiguating all cases to reflect the actual pronunciation in the standard variety of Italian would have been very time consuming, but it is also justified by the fact that in many varieties (including the northern ones of the authors) these distinctions are not made.

Another systematic correction concerns the placement of stress, that for many wordforms have been obtained automatically in GLAFF-IT, and sometimes turns out not to be in the right place: for instance, in many third-plural forms, the stress is incorrectly placed on the penultimate (e.g. PRS.IND.3SG /divent'ano/ 'they become', /okkup'ano/ 'they occupy'), while in our resource we move it to the (pre)antepenultimate (e.g. /div'entano/, /'okkupano/). While in other cases it was possible to correct stress position in an automatic way, by moving the stress to the syllable where it is systematically placed (e.g. the antepenultimate in forms like PRET.IND.3SG /f'etfero/ 'they did'), in this case, since there are two alternatives, the changes had to be done semi-automatically, by automatically moving the stress to the antepenultimate, and then manually moving it to the preantepenultimate whenever needed.

In cases of cells containing more than one wordform, we keep only one of the cell-mates. Wherever it was possible, we have used Thornton (2008)'s description of overabundance in Italian verb inflection to select the less marginal variant (e.g., keeping /d'evo/ rather than /d'ebbo/ in the PRS.IND.1SG of *DOVERE* 'must').

Several other punctual corrections were manually made on the data of GLAFF-IT, yielding the current version of our resource, that is clean enough to be able to perform an entropy-based analysis shedding light on the patterns of interpredictability between wordforms in Italian verb paradigms.

3 The Method

The Qumin toolkit computes implicative entropy values estimating the uncertainty in predicting each paradigm cell assuming knowledge of one (or more

than one) wordform, following the procedure described in Beniamine (2018). Here, we illustrate the methodology using the data given in Table 1.

lexeme	conj.	GER	PRS.IND.2PL
AMARE ‘love’	1 st	/am'ando/	/am'ate/
VEDERE ‘see’	2 nd	/ved'endo/	/ved'ete/
SENTIRE ‘hear’	3 rd	/sent'endo/	/sent'ite/

Table 1: Italian verbs of different conjugations.

The first step of the procedure consists in classifying verbs according to the patterns of formal alternation between wordforms, and the phonological context in which such alternations are attested. As is shown in the second column of Table 2, 1st and 2nd conjugation verbs display the same pattern (1), while 3rd conjugation verbs use another pattern (2). The second step is another classification based on the patterns that can potentially be applied to GER to obtain PRS.IND.2PL. As can be seen in the third column of Table 2, verbs of the 2nd and 3rd conjugation are in the same class (B), because patterns 1 and 2 can potentially be applied to a GER ending in /endo/, while only pattern 1 can be applied to 1st conjugation verbs with GER in /ando/. Entropy is then computed for each of the classes of this second classification, weighing the probability of application of different patterns by means of their type frequency in the data, i.e., the number of verbs in which they are attested: here, data from *LeFFI* are given in the last column of Table 2.

lexeme	pattern/context (1SG ↔ 3SG)	applicable patterns	n. verbs
AMARE	1 (_ndo ↔ _te / V.#)	A (1)	1,505
VEDERE	1 (_ndo ↔ _te / V.#)	B (1,2)	320
SENTIRE	2 (_endo ↔ _ite / C.#)	B (1,2)	215

Table 2: Information used to compute the entropy of predicting PRS.IND.2PL from GER.

$$\begin{aligned}
& H(\text{PRS.IND.2PL}|\text{GER}) \\
&= - \left(\frac{1,505}{2,040} \times \left(\frac{1,505}{1,505} \times \log_2 \frac{1,505}{1,505} \right) + \right. \\
&\quad \left. + \frac{535}{2,040} \times \left(\frac{320}{535} \times \log_2 \frac{320}{535} + \frac{215}{535} \times \log_2 \frac{215}{535} \right) \right) = \\
&= - \left(\frac{1,505}{2,040} \times 0 + \frac{535}{2,040} \times -0.972 \right) = \\
&= 0.255
\end{aligned}
\tag{1}$$

As is shown in Equation 1, there is no uncertainty in class A: given a GER in /ando/, PRS.IND.2PL cannot but be in /ate/. On the other hand, given a GER in /endo/, PRS.IND.2PL can be in in /ete/ (applying pattern 1) or in /ite/ (applying pattern 2). As a consequence, there is some uncertainty in this case. The entropy values of different classes are then summed and weighed – again on the basis of type frequency – in a single entropy value, that estimates the overall uncertainty in predicting PRS.IND.2PL from GER in Italian verbs.

4 Results

Giving the data of *LeFFI* as input to the Qumin toolkit, the output is an entropy-based distance matrix of all the cells of Italian verb paradigms. We do not show it here for reasons of space as it comprises 53 columns and rows, but we use its values to draw a mapping of the paradigm in zones of full interpretability, where two cells A, B are conflated in the same zone if they can be predicted from one another with no uncertainty, i.e. if $H(A|B) = H(B|A) = 0$. The outcome of this grouping is given in Table 3.

	1SG	2SG	3SG	1PL	2PL	3PL
FUT.IND	Z6	Z6	Z6	Z6	Z6	Z6
PRS.COND	Z6	Z6	Z6	Z6	Z6	Z6
PRS.SBJV.	Z2	Z2	Z2	Z4	Z4	Z2
PRS.IND	Z12	Z14	Z15	Z4	Z10	Z13
IPRF.IND	Z9	Z9	Z9	Z9	Z9	Z9
IPRF.SBJV	Z1	Z1	Z1	Z1	Z1	Z1
PRET.IND	Z5	Z1	Z5	Z1	Z1	Z5
IMP		Z3		Z4	Z10	

	M.SG	F.SG	M.PL	F.PL
PST.PTCP	Z7	Z7	Z7	Z7
PRS.PTCP	Z11		Z11	
GER	Z11			
INF	Z8			

Table 3: Zones of interpretability in Italian verb paradigms: verbal forms.

Given this mapping, we can obtain what Stump and Finkel (2013) call a distillation of the paradigm by keeping only one cell for each zone, since all the other cells are trivially predictable. In Table 4, we show entropy values for a distillation of the Italian verbal paradigm.

It is interesting to compare the mapping of Table 3, based on the predictability of wordforms, to the one based on the predictability of stems that has been proposed in the works cited in Section 1, where cells are grouped together if they

	Z1	Z2	Z3	Z4	Z5	Z6	Z7	Z8	Z9	Z10	Z11	Z12	Z13	Z14	Z15
Z1 (IPRF.SBJV.3SG)		.431	.364	.416	.045	.029	.029	.091	.009	.019	.009	.385	.347	.444	.357
Z2 (PRS.SBJV.3SG)	.405		.213	.474	.394	.423	.393	.436	.413	.405	.342	.213	.213	0	.213
Z3 (IMP.2SG)	.269	.006		.708	.226	.273	.235	.310	.273	.268	.172	.005	.005	0	.002
Z4 (PRS.IND.1PL)	1.239	1.238	1.419		.852	.743	.893	1.215	1.230	1.228	1.033	.965	1.322	.775	1.394
Z5 (PRET.IND.3SG)	.015	.443	.374	.451		.034	.056	.035	.044	.015	.044	.397	.359	.457	.370
Z6 (FUT.IND.3SG)	.495	.864	.856	.527	.231		.187	.526	.498	.503	.487	.460	.760	.466	.853
Z7 (PST.PTCP.F.SG)	.013	.435	.378	.426	.005	.026		.027	.006	.006	.006	.387	.352	.458	.376
Z8 (INF)	.032	.435	.366	.524	.045	.160	.033		.030	0	.026	.386	.358	.450	.377
Z9 (IPRF.IND.3SG)	.011	.429	.367	.417	.044	.458	.031	.086		.010	0	.380	.346	.442	.357
Z10 (PRS.IND.2PL)	.041	.435	.366	.428	.053	.494	.033	.083	.031		.026	.386	.358	.451	.367
Z11 (GER)	.265	.557	.503	.417	.128	.262	.141	.325	.254	.256		.474	.472	.582	.502
Z12 (PRS.IND.1SG)	.731	.830	.567	.703	.442	.366	.429	.712	.731	.727	.682		.830	.118	.572
Z13 (PRS.IND.3PL)	.248	0	0	.620	.229	.243	.228	.278	.248	.248	.176	0		0	0
Z14 (PRS.IND.2SG)	.991	.559	1.033	.544	.622	.502	.627	.976	.991	.986	.909	.329	.744		1.038
Z15 (PRS.IND.3SG)	.269	.003	0	.717	.233	.268	.239	.312	.270	.268	.173	.003	.003	0	

Table 4: Entropy values for a distillation of the Italian verb paradigm.

are based on the same stem. For this comparison, we refer to Montermini and Bonami (2013), where the most recent version of the stem-based mapping is provided. In their description, 8 stems are identified, while our word-based mapping is composed of 15 zones. In particular, Z1-9-10-11 of our mapping correspond to the zones including cells that are based on the same stem S1 in Montermini and Bonami (2013)’s analysis: this is why they are all colored with different shades of red in Table 3. Similarly, our Z2-12-13 (different shades of blue) include cells based on Montermini and Bonami (2013)’s S2 and our Z3-14-15 (different shades of green) include cells based on Montermini and Bonami (2013)’s S3. As for the other zones of our mapping, there is a one-to-one correspondence with the stems identified by Montermini and Bonami (2013).

The discrepancies between the two approaches are mostly due to two different reasons: (i) the presence of a few, highly irregular verbs¹ that are not accounted for by Montermini and Bonami (2013)’s analysis, but are included in our dataset, and, therefore, in our entropy-based analysis; (ii) more systematic opacities of some wordforms, that are poorly informative on the conjugation of lexemes.

As an example of case (i), PRS.IND.2PL and IPRF.IND.3SG can almost always be predicted from one another by replacing the final segments /te/ with /va/, or *vice versa*: e.g. AMARE (PRS.IND.2PL /am'ate/, IPRF.IND.3SG /am'ava/) and SENTIRE (PRS.IND.2PL /sent'ite/, IPRF.IND.3SG /sent'iva/).

¹Namely: ANDARE ‘to go’, AVERE ‘to have’, DARE ‘to give’, DIRE ‘to say’, ESSERE ‘to be’, FARE ‘to do’, SAPERE ‘to know’, and STARE ‘to stay’.

However, this generalization does not hold for a handful of highly irregular verbs, as is exemplified by DIRE ‘say’, with PRS.IND.2PL /d'ite/ but IPRF.IND.3SG /dit'eva/. Of course, the picture is different depending on the presence of such irregular verbs in the data. If they are excluded, as in Montermini and Bonami (2013), the two cells can be considered as based on the same stem (S1) and, thus, as being fully interpredictable. If they are included, as happens in our data, the two cells have to be assigned to different zones, since there is some uncertainty in predicting the cells from one another. However, entropy is very low in such cases, thanks to the weighing based on type frequency (see the corresponding values in Table 4). It should be noticed that the lexemes that are not considered by Montermini and Bonami (2013) because of their irregularity are among the verbs with higher token frequency in Italian (all ranking among the first 13 positions in COLFIS). This makes their exclusion less worrisome, as the irregular formal patterns they display can plausibly be considered as being learned by rote. Nevertheless, our entropy-based picture can be considered as achieving a higher level of granularity in the description.

As an example of case (ii), PRS.IND.2SG and PRS.IND.3SG are in the same zone in Montermini and Bonami (2013), because they are both considered as obtained from S3: in particular, PRS.IND.3SG is identical to S3, while to obtain PRS.IND.2SG the final vowel of S3 has to be replaced by /a/. In both cases, knowing the shape of S3 is sufficient to infer the cell without any uncertainty. However, in our word-based perspective there is uncertainty when guessing PRS.IND.3SG from PRS.IND.2SG: the latter al-

ways ends in /i/ (e.g. AMARE /'ami/, VEDERE /v'edi/), neutralizing the distinction between verbs of different conjugations, and, thus, not allowing to discriminate between 1st conjugation verbs with S3 and PRS.IND.3SG in /a/ (e.g. AMARE /'ama/) and 2nd and 3rd conjugation verbs with S3 and PRS.IND.3SG in /e/ (e.g. VEDERE /v'ede/).

These examples show that our method allows to identify sources of uncertainty that are downplayed in the stem-based picture, either because of their quantitative marginality – case (i) – or because they are obscured by the use of an abstract stem, that however is not always inferrable by the shape of the single wordform used as predictor – case (ii).

However, it should be noticed that at least the possible availability of more exhaustive stem spaces accounting for all the formal variation of Italian verb inflection, without excluding highly irregular verbs – thus corresponding to our case (i) – was already acknowledged in the works cited above: see e.g. Pirrelli and Battista (2000, Footnote 16) and Montermini and Bonami (2013, Footnote 9). Indeed, there is of course a trade-off between the number of zones in which the paradigm is split on the one hand, and the coverage of the identified zones with respect to the whole lexicon on the other hand. In the stem-based mapping, the choice is not to make the number of zones too high, at the (minimal) cost of not accounting for a handful of irregular verbs. Conversely, in the word-based mapping that we adopt in the present paper, the higher number of zones is compensated by a complete coverage of the whole lexicon. Now, how many of the zones are actually identified and learned by speakers is an empirical matter that should be tackled by means of psycholinguistic experiments. However, what is important to keep in mind is that this gap between the two approaches can be filled, either by drawing the stem space in such a way that it covers also for irregular verbs, or by reducing the number of zones in the word-based analysis gradually collapsing zones of interpredictability for increasing values of implicative entropy. For instance, if the criterion for two cells to be assigned to the same zone is for them to be predictable from one another with an implicative entropy value lower than 0.01, rather than 0, then Z3,13,15 can be merged in a same zone. If the threshold is set at 0.02, also Z1 and Z9 can be conflated in the same zone, to which also Z7 can be added with threshold set at 0.03.

On the other hand, the discrepancy between the

two approaches generated by more systematic, but unidirectional opacities such as the one described above in (ii) could be avoided if in the entropy-based mapping we decided that having null entropy in one direction would be a sufficient criterion for two cells to be assigned to the same zone – i.e., two cells belong to the same zone if either $H(A|B)$ or $H(B|A) = 0$.

5 Conclusions

In this paper, we have presented the inflected lexicon of Italian verbs *LeFFI*. We have then exploited it to investigate predictability in Italian verb inflection, using implicative entropy to estimate the uncertainty in predicting wordforms from one another. The results have been used to obtain a mapping of the paradigm in zones of interpredictability, that we have compared to the mapping of stems proposed in previous work, showing that our word-based procedure is capable of capturing aspects that are downplayed, if not ignored in the stem-based approach.

Besides their theoretical interest, both the resource and the information-theoretic approach potentially have more practical applications, for instance in the field of psycholinguistics. The resource provides a very clean but sufficiently large dataset of forms that can be used as a source of input for fine-grained experiments. In such experiments, it would be possible to test if the different levels of predictability between cells identified by different values of implicative entropy find a correspondence in the process of acquisition of inflectional morphology by L1 and L2 speakers – i.e., if the pairs of cells between which there are higher implicative entropy values are indeed the ones on which learners are more uncertain. More generally, our entropy-based evaluation of uncertainty in inflectional predictions can be considered as a measure of (at least one aspect of) morphological complexity, that can be used also in other areas, for instance to assess text readability.

6 Availability of Data and Tools

The data and tools used in this study are freely available online, allowing for an easy replication of the presented results. *LeFFI* can be found in the following repository: <https://github.com/matteo-pellegrini/LeFFI>. The Qumin toolkit that was used to automatically perform entropy computations can be freely downloaded at: <https://github.com/XachaB/Qumin>.

References

- Farrell Ackerman, James P. Blevins, and Robert Malouf. 2009. Parts and wholes: Patterns of relatedness in complex morphological systems and why they matter. In James P. Blevins and Juliette Blevins, editors, *Analogy in grammar: Form and acquisition*, pages 54–82. Oxford University Press, Oxford.
- Mark Aronoff. 1994. *Morphology by itself: Stems and inflectional classes*. MIT press, Cambridge.
- Sacha Beniamine. 2018. *Classifications flexionnelles. Étude quantitative des structures de paradigmes*. Ph.D. thesis, Université Sorbonne Paris Cité-Université Paris Diderot.
- Pier Marco Bertinetto, Cristina Burani, Alessandro Laudanna, Lucia Marconi, Daniela Ratti, Claudia Rolando, and Anna Maria Thornton. 2005. CoLFIS (Corpus e Lessico di Frequenza dell’Italiano Scritto).
- James P. Blevins. 2016. *Word and paradigm morphology*. Oxford University Press, Oxford.
- Olivier Bonami and Gilles Boyé. 2014. De formes en thèmes. In Florence Villoing, Sarah Leroy, and Sophie David, editors, *Foisonnements morphologiques. Études en hommage à Françoise Kerleroux*, pages 17–45. Presses Universitaires de Paris-Ouest, Paris.
- Olivier Bonami, Gauthier Caron, and Clément Plancq. 2014. Construction d’un lexique flexionnel phonétisé libre du français. In *Congrès Mondial de Linguistique Française – CMLF 2014*, volume 8, pages 2583–2596. EDP Sciences.
- Basilio Calderone, Matteo Pascoli, Franck Sajous, and Nabil Hathout. 2017. Hybrid method for stress prediction applied to GLAFF-IT, a large-scale Italian lexicon. In *International Conference on Language, Data and Knowledge*, pages 26–41, Cham. Springer.
- Bernard Comrie, Martin Haspelmath, and Balthasar Bickel. 2008. The Leipzig Glossing Rules: Conventions for interlinear morpheme-by-morpheme glosses. *Department of Linguistics of the Max Planck Institute for Evolutionary Anthropology and the Department of Linguistics of the University of Leipzig*.
- Jeremy Goslin, Claudia Galluzzi, and Cristina Romani. 2014. PhonItalia: a phonological lexicon for Italian. *Behavior research methods*, 46(3):872–886.
- Martin Maiden. 1992. Irregularity as a determinant of morphological change. *Journal of linguistics*, 28(2):285–312.
- Martin Maiden. 2018. *The Romance verb: Morphomic structure and diachrony*. Oxford University Press, Oxford.
- Fabio Montermini and Olivier Bonami. 2013. Stem spaces and predictability in verbal inflection. *Lingue e linguaggio*, 12(2):171–190.
- Fabio Montermini and Gilles Boyé. 2012. Stem relations and inflection class assignment in Italian. *Word Structure*, 5(1):69–87.
- Matteo Pellegrini and Marco Passarotti. 2018. LatInfLexi: an inflected lexicon of Latin verbs. In *Proceedings of the Fifth Italian Conference on Computational Linguistics (CLiC-it 2018)*.
- Matteo Pellegrini. 2020. Using LatInfLexi for an Entropy-Based Assessment of Predictability in Latin Inflection. In *Proceedings of LT4HALA 2020-1st Workshop on Language Technologies for Historical and Ancient Languages*, pages 37–46.
- Vito Pirrelli and Marco Battista. 2000. The paradigmatic dimension of stem allomorphy in Italian verb inflection: 2628. *Italian Journal of Linguistics*, 12(2):307–380.
- Gregory Stump and Raphael A. Finkel. 2013. *Morphological typology: From word to paradigm*. Cambridge University Press, Cambridge.
- Anna M. Thornton. 2008. A non-canonical phenomenon in Italian verb morphology: double forms realizing the same cell. *Poster presented at OxMorph1-Oxford*.
- Eros Zanchetta and Marco Baroni. 2005. Morph-it!: A free corpus-based morphological resource for the Italian language. *Proceedings of corpus linguistics*.