

# Almawave-SLU: a New Dataset for SLU in Italian

Valentina Bellomaria, Giuseppe Castellucci, Andrea Favalli and Raniero Romagnoli

Language Technology Lab  
Almawave srl  
[first name initial].[last name]@almawave.it

## Abstract

The widespread use of conversational and question answering systems made it necessary improve the performances of speaker intent detection and understanding of related semantic slots, i.e., Spoken Language Understanding (SLU). Often, these tasks are approached with supervised learning methods, which needs considerable labeled datasets. This paper<sup>1</sup> presents the first Italian dataset for SLU in voice assistants scenario. It is the product of a semi-automatic procedure and is used as a benchmark of various open source and commercial systems.

## 1 Introduction

Conversational interfaces, e.g., Google’s Home or Amazon’s Alexa, are becoming pervasive in daily life. As an important part of any conversation, language understanding aims at extracting the meaning a partner is trying to convey. Spoken Language Understanding (SLU) plays a fundamental role in such a scenario. Generally speaking, in SLU a spoken utterance is first transcribed, then semantic information is extracted. Language understanding, i.e., extracting a semantic “frame” from a transcribed user utterance, typically involves: i) Intent Detection (ID) and ii) Slot Filling (SF) (Tur et al., 2010). The former makes the classification of a user utterance into an intent, i.e., the purpose of the user. The latter finds what are the “arguments” of such intent. As an example, let us consider Figure 1, where the user asks for playing a song (Intent=PlayMusic) (*with or without you*, Slot=song) of an artist (*U2*, Slot=artist). Usually, supervised learning methods are adopted

play	with	or	without	you	by	U2
riproduci	with	or	without	you	degli	U2
O	B-song	I-song	I-song	I-song	O	B-artist

Figure 1: An example of Slot Filling in IOB format for a sentence with intent *PlayMusic*.

for SLU. Their efficacy strongly depends on the availability of labeled data. There are various approaches to the production of labeled data, depending on the intricacy of the problem, on the characteristics of the data, and on the available resources (e.g., annotators, time and budget). When the reuse of existing public data is not feasible, manual labeling should be accomplished, eventually by automating part of the labeling process.

In this work, we present the first public dataset for the Italian language for SLU. It is generated by a semi-automatic procedure from an existing English dataset annotated with intents and slots. We have translated the sentences into Italian and reported the annotations based on a token span algorithm. Then, the translation, spans and consistency of the entities in Italian have been manually validated. Finally, the dataset is used as benchmark for NLU systems. In particular, we will compare a recent state-of-the-art (SOTA) approach (Castellucci et al., 2019) with Rasa (ras, 2019) taken from the open source world, IBM Watson Assistant (wat, 2019), Google DialogFlow (dia, 2019) and, finally, Microsoft LUIS (msl, 2019), some commercial solutions in use.

Following, in section 2 related works will be discussed; In section 3 the dataset generation will be discussed. Section 4 we will present the experiments. Finally, in section 5 we will draw the conclusions.

## 2 Related Work

SLU has been addressed in the Natural Language Processing community mainly in the English lan-

<sup>1</sup>Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

guage. A well-known dataset used to demonstrate and benchmark various NLU algorithms is Airline Travel Information System (ATIS) (Hemphill et al., 1990) dataset, which consists of spoken queries on flight related information. In (Braun et al., 2017) three dataset for Intent classification task were presented. *AskUbuntu Corpus* and *Web Application Corpus* were extracted from Stack-Exchange and the third one, i.e., *Chatbot Corpus*, was originated from a Telegram chatbot. The newer multi-intent dataset SNIPS (Coucke et al., 2018) is the starting point for the work presented in this paper. An alternative approach to manual or semi-automatic labeling is the one proposed by the data scientists of the Snorkel project with Snorkel Drybell (Bach et al., 2018) that aims at automating the labeling through the use of data programming. Other works have explored the possibility of creating datasets in a language starting from datasets in other languages, such as (Jabaian et al., 2010) and (Stepanov et al., 2013). Regarding the Italian language two main works can be pointed out (Raymond et al., 2008; Vanzo et al., 2016). Our work differs mainly in the application domain (i.e., we focus on the voice assistants scenario). In particular, (Raymond et al., 2008) mainly focuses on dialogues in a customer service scenario; (Vanzo et al., 2016) focuses on Human-Robot interaction.

### 3 Almawave-SLU: A new dataset for Italian SLU

We created the new dataset <sup>2</sup> starting from the SNIPS dataset (Coucke et al., 2018), which is in English. It contains 14,484 annotated examples<sup>3</sup> with respect to 7 intents and 39 slots. In table 1 an excerpt of the dataset is shown. We started from this dataset as: i) it contains a reasonable amount of examples; ii) it is multi-domain; iii) we believe it could represent a more realistic setting in today’s voice assistants scenario.

We performed a semi-automatic procedure consisting of two phases: an automatic translation with contextual alignment of intents and slots; a manual validation of the translations and annotations. The resulting dataset, i.e., Almawave-SLU, has fewer training examples, a total of 7,142 and the same number of validation and test examples of the original dataset. Again, 7

<sup>2</sup>The Almawave-SLU dataset is available for download. To obtain it, please send an e-mail to the authors.

<sup>3</sup>There are 13084, 700 and 700 for training, validation and test, respectively.

intents and 39 slots have been annotated. Table 2 shows the distribution of examples for each intent.

#### 3.1 Translation and Annotation

In a first phase, we translated each English example in Italian by using the Translator Text API: part of the Microsoft Azure Cognitive Services. In order to create a more valuable resource in Italian, we also performed an automatic substitution of the names of movies, movie theatres, books, restaurants and of the locations with some Italian counterpart. First, we collected from the Web a set  $E$  of about 20,000 Italian versions of such entities; then, we substituted each entity in the sentences of the dataset with one randomly chosen from  $E$ .

After the translation, an automatic annotation was performed. The intent associated with the English sentence has been copied to its Italian counterpart. Slots have been transferred by aligning the source and target tokens<sup>4</sup> and by copying the corresponding slot annotation. In case of exceptions, e.g., multiple alignments on the same token or missing alignment, we left the token without annotation.

#### 3.2 Human Revision

In a second phase, the dataset was divided into 6 different sets, each containing about 1,190 sentences. Each set was assigned to 2 annotators<sup>5</sup>, and each was asked to review the translation from English to Italian and the reliability of the automatic annotation. The guideline was to consider a valid annotation when both the alignment and the semantic slots were correct. Moreover, also a semantic consistency check was performed: e.g., served dish and restaurant type or city and region or song and singer. The 2 annotators have been used to cross-check the annotations, in order to provide more reliable revisions. When the 2 annotators disagreed, the annotations have been validated by a third different annotator.

During the validation phase some interesting phenomena emerged. <sup>6</sup> For example, there have been cases of inconsistency between the restaurant name and the type of served dish when the name of the restaurant mentioned the kind of food served, e.g., "*Prenota un tavolo da Pizza Party per mangiare noodles*". There were also wrong associations between the type of restaurant and service

<sup>4</sup>The alignment was provided by the Translator API.

<sup>5</sup>A total of 6 annotators were available.

<sup>6</sup>Some inconsistencies were in the original dataset

AddToPlaylist	Add the song virales de siempre by the cary brothers to my gym playlist.
BookRestaurant	I want to book a top-rated brasserie for 7 people.
GetWeather	What kind of weather will be in Ukraine one minute from now?
PlayMusic	Play Subconscious Lobotomy from Jennifer Paull.
RateBook	Rate The children of Niobe 1 out of 6 points.
SearchCreativeWork	Looking for a creative work called Plant Ecology
SearchScreeningEvent	Is Bartok the Magnificent playing at seven AM?

Table 1: Examples from the SNIPS dataset. The first column indicates the intent, the second columns contains an example.

requested, e.g. *"Prenota nell'area piscina per 4 persone in un camion-ristorante"*. A truck restaurant is actually a van equipped for fast-food in the street. Again, among the cases of unlikely associations resulting from automatic replacement, the inconsistency between temperatures and cities is mentioned, in cases like "snow in the Sahara". Another type of problem occurred when the same slot was used to identify very different objects. For example, for the intent *SearchCreativeWork*, the slot *object\_name* was used for paintings, games, movies, etc... We can observe and analyze a couple of examples for this intent: *Can you find me the work, The Curse of Oak Island ?* and *Can you find me, Hey Man ?*. The first example contains *The Curse of Oak Island*, that is a television series and the second refers to *Hey Man* that is a music album, but both are labeled as *object\_name*, where the *object\_type* are different and not specified. In all these cases, the annotators were asked to correct the sentences and the annotations, accordingly. Again, in the case of *BookRestaurant* intent a manual revision was made when in the same sentence the city and state coexist: to make the data more relevant to the Italian language, the region relative to the city is changed, e.g. *"I need a table for 5 at a highly rated gastropub in Saint Paul, MN"* is translated and adapted for Italian in *"Vorrei prenotare un tavolo per 5 in un gastropub molto apprezzato a Biella, Piemonte"*.

	Train	Train-R	Valid	Test
AddToPlayList	744	185	100	124
BookRestaurant	967	250	100	92
GetWeather	791	195	100	104
PlayMusic	972	240	100	86
RateBook	765	181	100	80
SearchCreativeWork	752	172	100	107
SearchScreeningEvent	751	202	100	107

Table 2: Almayave-SLU Datasets statistics. Train-R is the reduced training set.

### 3.3 Automatic Translation Analysis

In many cases, machine translation lacked context awareness: this isn't an easy task due to phenomena as polysemy, homonymy, metaphors and idioms. There can be problems of lexical ambiguities when a word has more than one meaning and can produce wrong interpretations. For example, the verb "to play" can mean "spend time doing enjoyable things", such as "using toys and taking part in games", "perform music" or "perform the part of a character".

Human intervention occurred to maintain the meaning of the text dependent on cultural and situational contexts. Different translation errors were modified by the annotators. For example, the automatic translation of the sentence *Play Have You Met Miss Jones by Nicole from Google Music.* was *Gioca hai incontrato Miss Jones di Nicole da Google Music.*, but the correct Italian version is *Riproduci Have You Met Miss Jones di Nicole da Google Music.* In this case the wrong translation of the verb *play* causes a meaningless sentence.

Often, translation errors are due to the presence of prepositions, that have the same function in Italian as they do in English. Unfortunately, these cannot be directly translated. Each preposition is represented by a group of related senses, some of which are very close and similar while others are rather weak and distant. For example, the Italian preposition "di" can have six different English counterparts – of, by, about, from, at, and than. For example, in the SNIPS dataset the sentence *I need a table for 2 on feb. 18 at Main Deli Steak House* was translated as *Ho bisogno di un tavolo per 2 su Feb. 18 presso Main Deli Steak House.* Here, the translation of "on" is wrong: the correct Italian version should translate it as "il". Another example with wrong preposition translation is the sentence *"What will the weather be one month from now in Chad ?"*, the automatic translation of "one month from now" is "un mese da ora" but the correct translation is "tra un mese".

Common errors were in the translation of temporal expression, that are different between Italian and English. For example the translation of the sentence “*Book a table in Fiji for zero a.m*” was “*Prenotare un tavolo in Fiji per zero a.m*” but in Italian “zero a.m” is “mezzanotte”.

Other errors were specific of some intents, as they tend to have more slangs. For example, the translation of *GetWeather*’s sentences was problematic because the main verb is often misinterpreted, while in the sentences related to the intent *BookRestaurant* a frequent failure occurred on the interpretation of prepositions. For example, the sentence “*Will it get chilly in North Creek Forest?*” was translated as “*Otterrà freddo in North Creek Forest?*”, while the correct translation is “*Farà freddo a North CreekForest?*”. In this case, the system misinterpreted the context, assigning to “get” the wrong meaning.

## 4 Benchmarking SLU Systems

Nowadays, there are several human-machine interacting platforms, commercial and open source. Machine learning algorithms enable these systems to understand natural language utterances, match them to intents, and extract structured data. We decided to use the Alwave-SLU dataset with the following SLU systems.

### 4.1 SLU Systems

**RASA.** RASA (ras, 2019) is an open source alternative to popular NLP tools for the classification of intentions and the extraction of entities. Rasa contains a set of high-level APIs to produce a language parser through the use of NLP and ML libraries, via the configuration of the pipeline and embeddings. It seems to be very fast to train, does not require great computing power and, despite this, it seems to get excellent results.

**LUIS.** Language Understanding service (msl, 2019) allows the construction of applications that can receive input in natural language and extract the meaning from it through the use of Machine Learning algorithms. LUIS was chosen as it provides also an easy-to-use graphical interface dedicated to less experienced users. For this system the computation is completely done remotely and no configuration is needed.

**Watson Assistant.** IBM’s Watson Assistant (wat, 2019) is a white label cloud service that al-

lows software developers to embed a virtual assistant, that use Watson AI machine learning and NLU, in their software. Watson Assistant allows customers to protect information gathered through user interaction in a private cloud. It was chosen because it was conceived for an industrial market and for its long tradition in this task.

**DialogFlow.** Dialogflow (dia, 2019) is a Google service to build engaging voice and text-based conversational interfaces, powered by a natural language understanding (NLU) engine. Dialogflow makes it easy to connect the bot service to a number of channels and runs on Google Cloud Platform, so it can scale to hundreds of millions of users. DialogFlow was chosen due to its wide distribution and ease of use of the interface.

**Bert-Joint.** It is a SOTA approach to SLU adopting a joint Deep Learning architecture in an attention-based recurrent frameworks (Castellucci et al., 2019). It exploits the successful Bidirectional Encoder Representations from Transformers (BERT) model to pre-train language representations. In (Castellucci et al., 2019), the authors extend the BERT model in order to perform the two tasks of ID and SF jointly. In particular, two classifiers are trained jointly on top of the BERT representations by means of a specific loss function.

### 4.2 Experimental Setup

Alwave-SLU has been used for training and evaluation of Rasa, Luis, Watson Assistant, DialogFlow and Bert-Joint. Another evaluation is made on 3 different training datasets, i.e Train-R, of reduced dimensions with respect to the Alwave-SLU, each about 1,400 sentences equally distributed on intent.

The train/validation/test split used for the evaluations is 5,742 (1,400 for Train-R), 700 and 700, respectively. Regarding Rasa, we used version 1.0.7, and we adopted the standard “supervised embeddings” pipeline, since it is recommended in the official documentation. This pipeline consists of a *WhiteSpaceTokenizer*, that was modified to avoid the filter of punctuation tokens, a *Regex Featurizer*, a *Conditional Random Field* to extract entities, a *Bag-of-words Featurizer* and an *Intent Classifier*. LUIS was tested against the api v2.0, and the loading of data to train the system with LUIS APP VERSION 0.1. Unfortunately Watson

System	Eval-1 with Train set			Eval-2 with Train-R set		
	Intent	Slot	Sentence	Intent	Slot	Sentence
<b>Rasa</b>	96.42	85.40	65.76	93.84	78.58	52.25
<b>LUIS</b>	95.99	79.47	50.57	94.46	72.51	35.53
<b>Watson Assistant</b>	96.56	-	-	95.03	-	-
<b>Dialogflow</b>	95.56	74.62	46.16	93.60	65.23	36.68
<b>Bert-Joint</b>	<b>97.6</b>	<b>90.0</b>	<b>77.1</b>	<b>96.13</b>	<b>83.04</b>	<b>65.23</b>

Table 3: Overall scores for Intent and Slot

Assistant supports only English models for the annotations of contextual entities, i.e. slots; therefore, we have only measured the intents<sup>7</sup>. Regarding DialogFlow, a “Standard” (free) utility has been created with API version 2; the python library “dialogflow” has been used for the predictions.<sup>8</sup> DialogFlow allows the choice between pure ML mode (“ML only”) and hybrid rule-based and ML mode (“match mode”). We chosen ML mode. Regarding the BERT-Joint system, a pre-trained BERT model is adopted, which is available on the BERT authors website<sup>9</sup>. This model is composed of 12-layer and the size of the hidden state is 768. The multi-head self-attention is composed of 12 heads for a total of 110M parameters. As suggested in (Castellucci et al., 2019), we adopted a dropout strategy applied to the final hidden states before the intent/slot classifiers. We tuned the following hyper-parameters over the validation set: (i) number of epochs among (5, 10, 20, 50); (ii) Dropout keep probability among (0.5, 0.7 and 0.9). We adopted the Adam optimizer (Kingma and Ba, 2015) with parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , L2 weight decay 0.01 and learning rate  $2e-5$  over batches of size 64.

### 4.3 Experimental Results

In table 3 the performances of the systems are shown. The SF performance is the F1 while the ID and Sentence performances are measured with the accuracy. We also show an evaluation carried out with models trained on three different split of reduced size derived from the whole dataset. The reported value is the average of measurements obtained separately on the entire test dataset.

<sup>7</sup>Refer to Table 3. Entity feature support details at <https://cloud.ibm.com/docs/services/assistant?topic=assistant-language-support>

<sup>8</sup><https://cloud.google.com/dialogflow/docs/reference/rest/v2/projects.agent.intents#Part>

<sup>9</sup>[https://storage.googleapis.com/bert\\_models/2018\\_11\\_23/multi\\_cased\\_L-12\\_H-768\\_A-12.zip](https://storage.googleapis.com/bert_models/2018_11_23/multi_cased_L-12_H-768_A-12.zip)

Regarding the ID task, all models are performing similarly, but Bert-Joint F1 score is slightly higher than others. For SF task, notice that there are significant differences between LUIS, DialogFlow and Rasa performances.

Finally, Bert-Joint achieved the top score on joint classification, in the assessments with the two different sizes of the dataset. The adaptation of nominal entities in Italian may have amplified the problem for the other models.

## 5 Conclusion

The contributions of this work are two-fold: first, we presented and released the first Italian SLU dataset (Almawave-SLU) in the voice assistants context. It is composed of 7,142 sentences annotated with respect to intents and slots, almost equally distributed on the 7 different intents. The effort spent on the construction of this new resource, according to the semi-automatic procedure described, is about 24 FTE<sup>10</sup>, with an average production of about 300 examples per day. We consider this effort lower than typical efforts to create linguistic resources from scratch.

Second, we compared some of the most popular NLU services with this data. The results show they all have similar features and performances. However, compared to another specific architecture for SLU, i.e., Bert-Joint, they perform worse. It was expected and it demonstrates the Almawave-SLU can be a valuable dataset to train and test SLU systems on the Italian language. In future, we hope to continuously improve the data and to extend the dataset.

## 6 Acknowledgment

The authors would like to thank to David Alessandrini, Silvana De Benedictis, Raffaele Mazzocca, Roberto Pellegrini and Federico Wolenski for the support in the annotation, revision and evaluation phases.

<sup>10</sup>Full Time Equivalent

## References

- Stephen H. Bach, Daniel Rodriguez, Yintao Liu, Chong Luo, Haidong Shao, Cassandra Xia, Souvik Sen, Alexander J. Ratner, Braden Hancock, Houman Alborzi, Rahul Kuchhal, Christopher Re, and Rob Malkin. 2018. Snorkel drybell: A case study in deploying weak supervision at industrial scale. *CoRR*, abs/1812.00417.
- Daniel Braun, Adrian Hernandez-Mendez, Florian Matthes, and Manfred Langen. 2017. Evaluating natural language understanding services for conversational question answering systems. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 174–185, Saarbrücken, Germany, August. Association for Computational Linguistics.
- Giuseppe Castellucci, Valentina Bellomaria, Andrea Favalli, and Raniero Romagnoli. 2019. Multilingual intent detection and slot filling in a joint bert-based model. *CoRR*, abs/1907.02884.
- Alice Coucke, Alaa Saade, Adrien Ball, Theodore Bluche, Alexandre Caulier, David Leroy, Clement Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, Mael Primet, and Joseph Dureau. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *CoRR*, abs/1805.10190.
2019. Google dialogflow. <https://dialogflow.com>.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. The atis spoken language systems pilot corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley*.
- Bassam Jabaian, Laurent Besacier, and Fabrice Lefèvre. 2010. Investigating multiple approaches for slu portability to a new language. In *INTER-SPEECH*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
2019. Microsoft luis on azure. <https://azure.microsoft.com/it-it/services/cognitive-services/language-understanding-intelligent-service/>.
2019. Rasa: Open source conversational ai. <https://rasa.com/>.
- Christian Raymond, Kepa Joseba Rodriguez, and Giuseppe Riccardi. 2008. Active annotation in the LUNA Italian corpus of spontaneous dialogues. In *LREC 2008*.
- Evgeny Stepanov, Ilya Kashkarev, Ali Orkan Bayer, Giuseppe Riccardi, and Arindam Ghosh. 2013. Language style and domain adaptation for cross-language slu porting. pages 144–149, 12.
- G. Tur, D. Hakkani-Tur, and L. Heck. 2010. What is left to be understood in atis? In *2010 IEEE Spoken Language Technology Workshop*, pages 19–24, Dec.
- Andrea Vanzo, Danilo Croce, Giuseppe Castellucci, Roberto Basili, and Daniele Nardi. 2016. Spoken language understanding for service robotics in italian. In Giovanni Adorni, Stefano Cagnoni, Marco Gori, and Marco Maratea, editors, *AI\*IA 2016 Advances in Artificial Intelligence*, pages 477–489, Cham. Springer International Publishing.
2019. Ibm watson assistant v1. <https://cloud.ibm.com/apidocs/assistant>.