

# FreeLing 2.1: Five years of open-source language processing tools

Lluís Padró\*, Miquel Collado\*, Samuel Reese<sup>◇</sup>, Marina Lloberes<sup>†</sup>, Irene Castellón<sup>†</sup>

\*Software Department – TALP Research Center  
Universitat Politècnica de Catalunya  
padro@lsi.upc.edu, miquel.collado@gmail.com

<sup>◇</sup>ISAE – Supaero  
Université Paul Sabatier  
samuel.reese@supaero.org

<sup>†</sup>GRIAL research group  
Universitat de Barcelona  
{marina.lloberes, icastellon}@ub.edu

## Abstract

FreeLing is an open-source multilingual language processing library providing a wide range of language analyzers for several languages. It offers text processing and language annotation facilities to natural language processing application developers, simplifying the task of building those applications. FreeLing is customizable and extensible. Developers can use the default linguistic resources (dictionaries, lexicons, grammars, etc.) directly, or extend them, adapt them to specific domains, or even develop new ones for specific languages.

This paper overviews the recent history of this tool, summarizes the improvements and extensions incorporated in the latest version, and depicts the architecture of the library. Special focus is brought to the fact and consequences of the library being open-source: After five years and over 35,000 downloads, a growing user community has extended the initial three languages (English, Spanish and Catalan) to eight (adding Galician, Italian, Welsh, Portuguese, and Asturian), proving that the collaborative open model is a productive approach for the development of NLP tools and resources.

## 1. Introduction

Basic language processing tasks such as tokenizing, morphological analysis, lemmatizing, part-of-speech tagging, word sense disambiguation (WSD), dependency parsing, etc. are a need for most natural language processing (NLP) applications such as Machine Translation, Summarization, Dialogue systems, Text mining, etc.

This makes language analyzers very valuable resources for researchers and developers in NLP. Also, the lack of out-of-the-box state-of-the-art systems is a severe bottleneck for faster progress in the area, both in research and development.

FreeLing was undertaken with the believe that steps should be taken towards general availability of basic NLP tools and resources, which may be used without restrictions. Thus, to enable faster advances and more portable systems in our area, an open-source model was chosen.

After five years (first version was released on 2004), over 35,000 downloads, and a growing user community which has extended the initial three languages (English, Spanish and Catalan) to eight (adding Galician, Italian, Welsh, Portuguese, and Asturian) prove

that the collaborative open model is a productive approach for the development of NLP tools and resources.

In the next section we review FreeLing evolution in its five-year life. Section 3 describes the most recent enhancements to the library, and Section 4 presents a shallow technical overview of the library. Section 5 outlines some conclusions and further work.

## 2. FreeLing project evolution

The first version of FreeLing (1.0) was released in 2003, and presented to the community in LREC 2004 (Carreras et al., 2004). That first version was a re-implementation of several tools (Atserias et al., 1998; Carreras and Padró, 2002) developed at the TALP research center<sup>1</sup> in UPC<sup>2</sup>.

The new version, was rewritten from scratch with the following goals:

- Build a fast state-of-the-art analyzer that could be used to process large amounts of text and to develop NLP applications.

<sup>1</sup><http://www.talp.upc.edu>

<sup>2</sup><http://www.upc.edu>

- Disseminate the results of the NLP research at TALP center, and make it available to the community
- Set a collaborative environment for the development of language analyzers, where contributions from any NLP community member can fit and revert in a global benefit.

To achieve these goals, the tool was structured as a highly modular C++ library consisting of several classes, each able to perform a different kind of analysis. Under this approach, new classes or services can be added to the library, and applications can select which services will be used, how, and when.

The goals related to dissemination and community collaboration were achieved releasing the software under an open-source GPL license, which has made it possible to enrich the project with many contributions, and make the tool useful for a larger number of people, as the increasing number of downloads suggests.

Currently, the number of downloads of the library is over 35,000 since it was first released, showing a steady growth at each release, and a large expansion at the latest 2.1 version (see Table 1).

Versions		total #downl	monthly average #downl
1.0, 1.0.1, 1.1	Oct.03–Sep.04	608	54
1.2, 1.3, 1.4, 1.5b1	Oct.04–Sep.06	2,250	96
1.5	Oct.06–Jan.08	2,451	161
2.0, 2.1a1, 2.1b1,	Feb.08–Sep.09	3,297	163
2.1a1, 2.1	Oct.09–Mar.10	26,713	4,452

Table 1: Evolution of FreeLing downloads

The original three supported languages (Catalan, Spanish, and English) have been brought up to eight by this community, with the inclusion of Italian, Galician, Portuguese, Welsh, and Asturian. Also, the original size-limited dictionaries for Catalan and Spanish have been replaced by much larger state-of-the-art resources.

Finally, the increasing availability of other open-source language analysis tools has enabled us to extend the library capabilities with long-awaited services such as Word Sense Disambiguation (WSD), obtained thanks to the inclusion of the UKB disambiguator developed by (Agirre and Soroa, 2009). See (Padró et al., 2010) for details on its integration in FreeLing.

We regard this evolution as very satisfactory, since it is fulfilling our goals of creating a useful tool which can

be enlarged by the community, used by application developers, and become a flexible platform enabling the integration of new resources and services.

### 3. Recent extensions

The most recent extensions included in FreeLing library can be classified in three types: extensions for linguistic resources for existing languages, support for new languages, and processing modules offering new services.

#### 3.1. Improving existing languages

As in most open-source projects, a large part of FreeLing contributions are related to the enhancement of already existing features. In our case, noticeable contributions include:

- Development of a number recognition module for Italian, which was not present in previous versions, thanks to Vitalie Scurtu<sup>3</sup>.
- Enlargement of Spanish dictionary from 6,000 to 75,000 lemmas, thanks to the Spanish Resource Grammar (SRG) project (Marimon et al., 2007).
- Enlargement of the Catalan dictionary from 7,000 to 71,000 lemmas, thanks to the Catalan speller project *El Corrector*<sup>4</sup>.
- Enlargement of the Galician dictionary from 7,000 to 50,000 lemmas, thanks to the PLN research group at Universidade de Santiago de Compostela<sup>5</sup>.
- Development and improvement of Spanish, Catalan, and English chunking and dependency grammars by the GRIAL research group<sup>6</sup> (Carrera et al., 2008), and thanks to EuroOpenTrad and KNOW projects, funded by Spanish Government

#### 3.2. Support for new languages

Another important bulk of contributions are those devoted to include support for new languages in FreeLing. This task requires on the one hand, the compilation of a morphological dictionary which provides possible lemmas and PoS tags for each form, and on the other hand, the development of a hand tagged corpus (consistent with the dictionary) to enable the training of the PoS taggers. In addition, other minor adjustments have to be done, such as adapting the tokenizer rules to the particularities of the language, writing appropriate rules to deal with suffixation, etc.

Languages that have been recently included are Welsh, Asturian, and Portuguese. All of them offer a chain of processing from plain text up to PoS tagging.

<sup>3</sup><http://scurtu.sitonline.it>

<sup>4</sup><http://www.elcorrector.cat>

<sup>5</sup><http://gramatica.usc.es>

<sup>6</sup><http://grial.uab.es>







