The DELPH-IN Grammary: a Curated Repository of Grammars and Treebanks

Francis Bond♠ **and Dan Flickinger**♢
♠ Palacký University
fcbond.github.io/
♢ Independent Researcher

bond@ieee.org, danflick@alumni.stanford.edu

Abstract

Precision computational grammars encode detailed linguistic analyses and compositional semantics that support rigorous investigation of grammatical phenomena, but their development requires substantial expertise and maintenance. To ensure long-term sustainability and accessibility of these resources, we present the DELPH-IN Grammary, a curated collection of twenty three HPSG grammars spanning seventeen languages and eight language families. The repository includes mature broad-coverage grammars (English, German, Japanese, Norwegian, Spanish) with associated treebanks, as well as grammars for typologically diverse and less-resourced languages. Each grammar is standardized with metadata, compiled using the ACE parser/generator, and loaded into the Linguistic Type Database for detailed inspection. Following FAIR principles, all resources are version-controlled and archived on Zenodo with annual releases synchronized to community development cycles. The Grammary enables reproducible grammar research, cross-linguistic typological studies, semantic parsing development, and grammar engineering pedagogy, providing the depth and theoretical grounding that complements data-driven approaches in computational linguistics. Our goal is to establish a sustainable model for preserving these valuable resources which bridge the critical gap between theoretical linguistics and empirical corpus based research. The Grammary is available at `https://github.com/delph-in/grammary` (Zenodo doi: zenodo.18945956).

Keywords: computational grammars, HPSG, DELPH-IN, language resources, FAIR principles, precision parsing, compositional semantics, treebanks, grammar engineering, multilingual NLP

1. Introduction

Grammars of natural languages are inherently labour-intensive to produce, and computational grammars add to the effort required by demanding of the grammarian a level of consistency, explicitness, and internal coherence not always found in grammar texts. Implemented grammars can take years to develop, a time span that is often at odds with the pace of development and maintenance of associated software for analyzing and processing them, and also at odds with cycles of funding to support continuing development of these grammars. In order to preserve and enhance the accessibility of such valuable and relatively scarce resources, we have developed a central open-source repository which gathers up-to-date versions of one group of grammars and provides a common viewing platform for inspecting, downloading, and experimenting with them. We call this repository the DELPH-IN Grammary,[1] and in this paper describe the grammars in the collection along with the processing and inspection tools available to enable further study, development, and evaluation.

[1] `grammary, gramarye, gramary` (noun) 1. (obsolete) Grammar; learning. 2. (archaic) Mystical learning; the occult, magic, sorcery. (Wiktionary, accessed 2026-03-08)

2. Background

The grammars we have collected here are from the Deep Linguistic Processing with HPSG Initiative (DELPH-IN: Uszkoreit et al., 2000; Uszkoreit, 2002), a long-running collaboration on linguistically precise, semantics-bearing grammars and tools grounded in HPSG and Minimal Recursion Semantics (MRS).[2] Its flagship ERG (English Resource Grammar) has served as a workhorse for research on broad-coverage parsing/generation and for interfaces where fine-grained meaning representations matter (Flickinger, 2000). DELPH-IN curates an ecosystem (LKB, PET/ACE, LOGON pipelines, ERS documentation) and sister grammars such as Jacy (Japanese) and NorSource (Norwegian), as well as the LinGO Grammar Matrix for rapid prototyping of precision HPSG grammars (Copestake, 2002; Bender et al., 2002). The motivation is explicit: deliver *precision* and *compositional semantics* that support both parsing and generation.

In terms of scale, DELPH-IN's resources cover multiple languages with varying maturity; ERG is the most developed, supported by decades of engineering, regression tests, and semantic documentation (Oepen and Flickinger, 2019). The Grammar Matrix has spawned dozens of research grammars

[2] delph-in.github.io

Proceedings of the Fifteenth Language Resources and Evaluation Conference (LREC 2026), pages 11776–11786
11–16 May 2026. ©ELRA Language Resources Association (ELRA), 2026

and field-work fragments, lowering the barrier for typologically diverse languages to enter the precision-grammar space (Bender et al., 2002). While the raw number of languages is smaller than in treebank-centric projects, the depth per language is unusually high, including underspecification, multi-word expressions, robust rules, and MRS-based logical forms (Copestake et al., 2005). The DELPH-IN approach borrows from software engineering in encouraging frequent testing: the grammar is developed, used to parse a corpus, treebanked, revised and treebanked again ensuring that it maintains its coverage of the corpus (Oepen et al., 2004a).

The richness of the grammars comes at a cost: it is hard to train grammarians, and due to the complexity of the grammars it is hard to understand the internal workings. For that reason, there are various tools for inspecting the grammars from within the grammatical engines (ACE and the LKB) and through the annotated corpora via the linguistic type database (LTDB: Hashimoto et al., 2007).

2.1. Formalism

DELPH-IN grammars are developed within the linguistic framework of Head-driven Phrase Structure Grammar: (HPSG: Pollard and Sag, 1994), a highly lexicalized, constraint-based approach to defining the syntactic and semantic properties of human languages. These grammars each include a rich lexicon and a set of lexical and syntactic rules whose properties are organized in a hierarchy of types, each of which defines constraints on the well-formedness of words or phrases of the language. These constraints are specified in a formal language called the Type Description Language (TDL: Krieger and Schafer, 1994), implemented in several unification-based parsers and generators. The reference semantic formalism used in all DELPH-IN grammars is Minimal Recursion Semantics (MRS: Copestake et al., 2005), an underspecified, flat representation of the compositional semantics of a sentence, systematically constructed via unification from the semantics of its constituents.

2.2. Tools

Among the processing engines developed within the DELPH-IN ecosystem is an efficient parser/generator called ACE (Answer Constraint Engine: Crysmann and Packard, 2012), which can apply each DELPH-IN grammar to sentences of that language when parsing text, and which can in principle also generate sentences of that language when presented with valid MRSs as input. When parsing, ACE exhaustively computes all valid analyses for the input string which are licensed by the grammar, and can rank these analyses using a statistical model trained on a manually-constructed treebank when available.

The DELPH-IN inventory of tools also includes a full-forest treebanker (FFTB: Packard, 2015), which provides a graphical interface for a discriminant-based approach to manually annotating the intended analysis of each sentence in a corpus. An important property of this method is the recording of the binary choices made by the annotator in selecting the intended analysis, so that later updates to the grammar can be accommodated in the FFTB to efficiently update the associated treebank, keeping the two resources in sync over time. For some DELPH-IN grammars, the treebanks are maintained by the grammarian for reference and evaluation as the grammar develops, while for some larger grammars, the associated treebank is also used to train a statistical model that can be applied during parsing to rank the candidate analyses licensed by the grammar (Oepen et al., 2004b).

An important tool for examining the contents of these treebanks and their associated grammars is the Linguistic Type Database (LTDB: Hashimoto et al., 2007), which can import the TDL files defining a grammar together with its treebank, to produce a fine-grained and comprehensive view of these complex systems. The LTDB also makes use of metadata and individual type-level documentation strings supplied by each grammarian, and further computes salient summary properties of each grammar to construct an informative overview. Figure 1 shows an example of a lexical entry from the Mandarin Chinese Grammar, Zhong, in the LTDB. It shows the name of the lexical type, a brief human readable description, some examples of words from this lexical type and sentences using them. Each sentence can show the parse derivation tree, and semantics (as avm, dependency (shown) or text (where it can be input to the generator)). Grammarians can search by type, lemma or predicate. In addition, they can examine a complete list of rules (including inflectional rules and constructions) and lexical types, with information on frequency in the corpus and lexicon. The goal is to make this information available both for linguists interested in the grammar of the language, and to grammar engineers who want to study the analyses or improve the grammar. As this is a web interface, the grammarian does not need to install any of the DELPH-IN infrastructure to explore the grammars and accompanying treebanks.

We used `pydelphin` (Goodman, 2019) to extend the LTDB so that it also provides an interactive demo — you can parse sentences and see the results as syntactic trees or semantic dependencies, with all types (rules, lexical entries and predicates) linked to their entry in the grammar. If the grammar

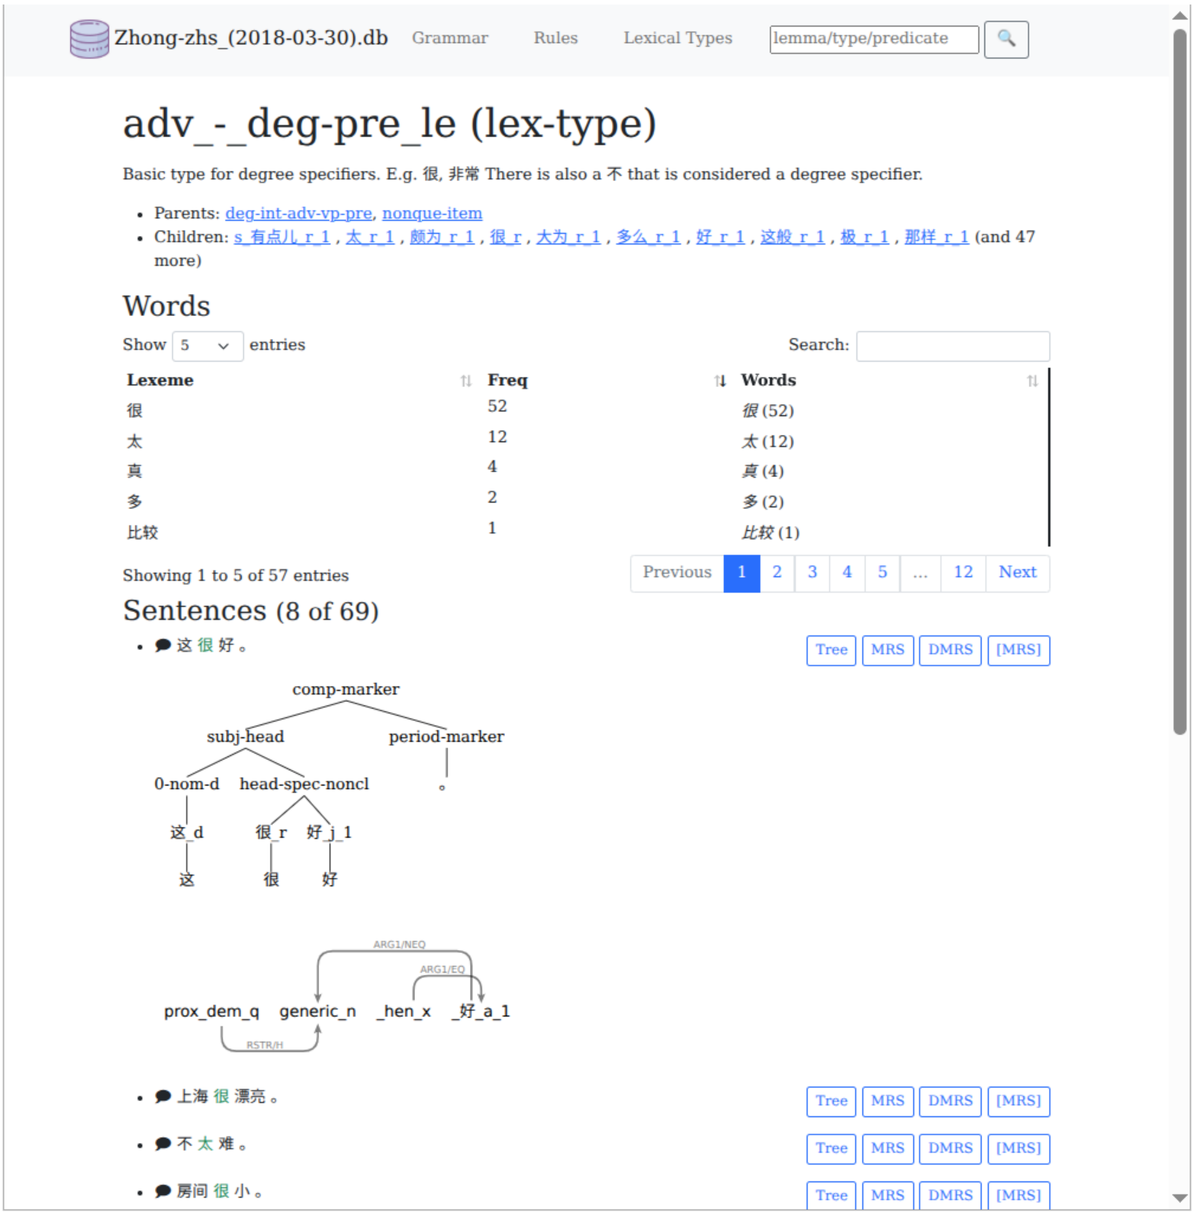

Figure 1: A Lexical entry from Zhong

supports generation, you can also select a parse result and then generate from it.

3. The Grammar

Because grammar engineering is such a labour-intensive and expert-driven process, its results should be made as broadly reusable as possible. To this end, we have gathered all known grammars that run under the DELPH-IN framework to ensure that they are Findable, Accessible, Interoperable, and Reusable (FAIR: Wilkinson et al., 2016). DELPH-IN has a long tradition of openness: the LOGON project, for instance, distributed grammars and tools through an open Subversion server (Lønning et al., 2004). However, that infrastructure is no longer maintained and many newer grammars are not included. Our goal is therefore to extend this tradition using contemporary open-science practices and to create a sustainable, FAIR-compliant collection of precision grammars.

To build the Grammar, we first surveyed the literature and repositories to locate all released grammars and to identify their most recent versions. Some exist only as legacy tar archives, while more recent ones are hosted on GitHub. For each grammar we record its canonical URL and, where necessary, retrieval instructions (e.g. via `git` or `svn`). This creates a living index that links linguistic resources to their maintainers and version histories.

Next, we enriched each grammar with minimal metadata required by the Linguistic Type Database:

Grammar	License	Lexicon	Rules	Trees
Abui	MIT	83	62	0
AraGram - an HPSG grammar of Modern Arabic	MIT	150	61	0
Bulgarian Resource Grammar	MIT	183	280	0
Cantonese HPSG	MIT	123	26	0
English Resource Grammar	MIT	44,580	392	97,650
English Resource Grammar (for Singlish)	MIT	44,649	408	44
English Resource Grammar (for Dictionaries)	MIT	44,580	396	1,508
English Resource Grammar (for Education)	MIT	46,733	471	0
German Grammar	LGPLLR	64,017	1,022	102
Hausa Grammar	LGPLLR	498	132	0
HeGram - an HPSG grammar of Modern Hebrew	MIT	432,196	64	0
Indonesian Resource Grammar	MIT	23,972	85	2,861
Jacy Japanese Grammar	MIT	56,866	118	8,797
Korean Resource Grammar	MIT	39,942	352	22
Zhong (Mandarin Chinese Traditional)	MIT	283	56	0
Norwegian HPSG grammar	LGPLLR	74,805	449	1,202
Norwegian syntax-based grammar	MIT	70,092	123	0
PorGram: A Portuguese HPSG Grammar	MIT	1,539	160	0
Russian Resource Grammar	LGPLLR	401	66	0
Spanish Resource Grammar	MIT	54,623	655	2,677
Wambaya aux+verb cluster grammar	MIT	1,528	239	727
Wambaya argument composition grammar	MIT	1,528	239	726
Zhong (Mandarin Chinese Simplified)	MIT	38,907	61	681

Table 1: Summary of the Grammars Collected in the Grammary

a short identifier, full name, and path to its configuration file. Where available, we also specified the associated treebank directory and the subset of treebanks to include. Additional metadata (such as license, maintainers, project website and more) is automatically parsed and displayed when present, and we have added it to most grammars. Up-to-date metadata is important for findability and usability; by making it more visible, we ensure that it is kept current.

A number of grammars required minor repair or adaptation to compile successfully under current toolchains. When grammars are still actively maintained, we contributed patches upstream; otherwise, we preserve the fixes locally and apply them automatically after download. This step, while technically demanding, is crucial: it transforms archived grammars into immediately executable and testable resources. Repairs were required, for example, for BURGER (Bulgarian), GG (German), HAG (Hausa), HeGram/AraGram (Hebrew/Arabic), the Russian Resource Grammar, and the Wambaya grammars. Such maintenance demands detailed knowledge of the formalism and compiler behaviour, but it yields resources that can be run consistently across platforms. This standardisation step allows the grammars to be uniformly accessed and processed by DELPH-IN tools. Finally, in order to make the demo easier to use, we added sample sentences to each grammar, to show what can be parsed, and how the input should look (i.e. how it should be transcribed or tokenized). We show an example of a sentence being parsed and then generated in Figure 2.

Finally, all grammars are loaded into the LTDB and compiled with ACE, ensuring they are immediately searchable and usable for parsing and generation. The resulting collection is compressed, archived, and deposited in Zenodo so that the data is persistently available for both human researchers and automated workflows.[3] The Grammary thus provides a stable reference snapshot of all compiled grammars, enabling reproducible experiments, cross-framework evaluation, and long-term preservation independent of particular software versions. Crucially, the grammars are then accessible through the LTDB web interface with no local installation required and no need to understand the full details of the DELPH-IN formalism. We plan to issue annual releases in coordination with the DELPH-IN Summit, supporting ongoing community development and versioned benchmarking.

The LTDB provides several layers of interactivity designed for both linguists and grammar engineers. Each grammar's rules and lexical types are browseable as dedicated pages, showing treebank frequencies alongside example sentences drawn directly from the treebank, with the words instantiating the relevant type highlighted in context. Within the TDL definition of any type, all type

[3] zenodo.18945956

names are rendered as hyperlinks, enabling engineers to navigate the full inheritance hierarchy by clicking through parent and daughter types — a feature particularly valuable when tracing the source of a constraint or understanding how an analysis is constructed. Types show in which file and from which line they are defined. For grammars with associated treebanks, frequency information is available at every level: per rule, per lexical type, and per lexeme. Finally, the parse demo goes beyond simple parsing: for grammars that support generation (most notably the ERG), a user can select any parse result, inspect its MRS semantic representation, and then generate surface strings from it — a vivid illustration of the bidirectional competence of these grammars and the semantic representations they produce.

4. The Grammars

Table 1 summarizes the grammars compiled in the Grammary. It gives the name (hyperlinked to the website if any), the license, number of lexical items, number of rules and number of treebanked sentences. Much more information is available for each grammar in the Grammary; this is just a short summary.

Grammars in the Grammary cover 8 major language families with 17 distinct languages modelled in 24 grammar implementations (Table 2).

4.1. Abui

The Abui grammar is a small research grammar for Abui, a language spoken on the island of Alor in Indonesia (Kratochvíl, 2026). It is an active–stative language: Agents are typically realized as NPs or free pronouns, while Patients are either NPs or pronominal prefixes on verbs, very different from any of the other languages in the collection. Abui has no written tradition, so all the data is collected and transcribed by field linguists with the help of the local language community. The grammar was first developed with the grammar Matrix, and then extended, but still has only modest coverage.

4.2. Arabic

The grammar of Modern Standard Arabic is one of a closely related pair together with the grammar of Hebrew described below (Arad Greshler et al., 2015). Developed in parallel with common constraints defined in a shared set of types, this small grammar provides illustrations of linguistic phenomena that illuminate contrasts with Hebrew.

4.3. Bulgarian

BURGER is a relatively small grammar of Bulgarian that includes a sampling of lexical entries but a rather richer inventory of rules which encode analyses of a wide range of syntactic phenomena in the language. The grammar does not include a corresponding treebank, although it was developed alongside a large-scale treebank for Bulgarian (BulTreebank (Simov et al., 2004).

4.4. Cantonese (Yue)

The Cantonese grammar (yue) started as a Matrix grammar. It is being extended to cover common phenomena, and a few interesting phenomena where it differs from Mandarin Chinese, such as aspectual particles or classifier use, have been developed in depth (Sio and Song, 2015; Bond and Sio, 2025). It is a good example of grammar reuse, extensively borrowing analyses from the Mandarin Chinese grammar (Zhong).

4.5. English

The English Resource Grammar (ERG) is a broad-coverage grammar enabling analyses of more than 95% of sentences comprising corpora in several genres including newspaper text, Wikipedia, the Brown Corpus, transcriptions of spoken dialogues, and e-commerce customer email (Flickinger, 2000). The implementation, under active development, includes three notable variants to analyze particular variants from standard English, including one for the Singaporean dialect Singlish (Chow and Bond, 2022), one for dictionary definitions (with e.g. object drop), and one that accommodates mildly ungrammatical sentences to enable applications for teaching English grammar (Suppes et al., 2014). The large Redwoods treebank (Oepen et al., 2004b) accompanies each release of the ERG.

4.6. German

The grammar of German (GG: Crysmann, 2005) saw its beginnings in the Verbmobil project of the 1990s (Wahlster, 2000), and was under continuous development for more than a decade with several authors, resulting in a broad-coverage grammar with a large lexicon. There is a small treebank illustrating some of the many linguistic phenomena whose analyses are implemented in the grammar.

4.7. Hausa

The Hausa grammar (HaG: Crysmann, 2012) includes a relatively small illustrative lexicon but implements a significant range of linguistic phenomena well documented in a series of publications.

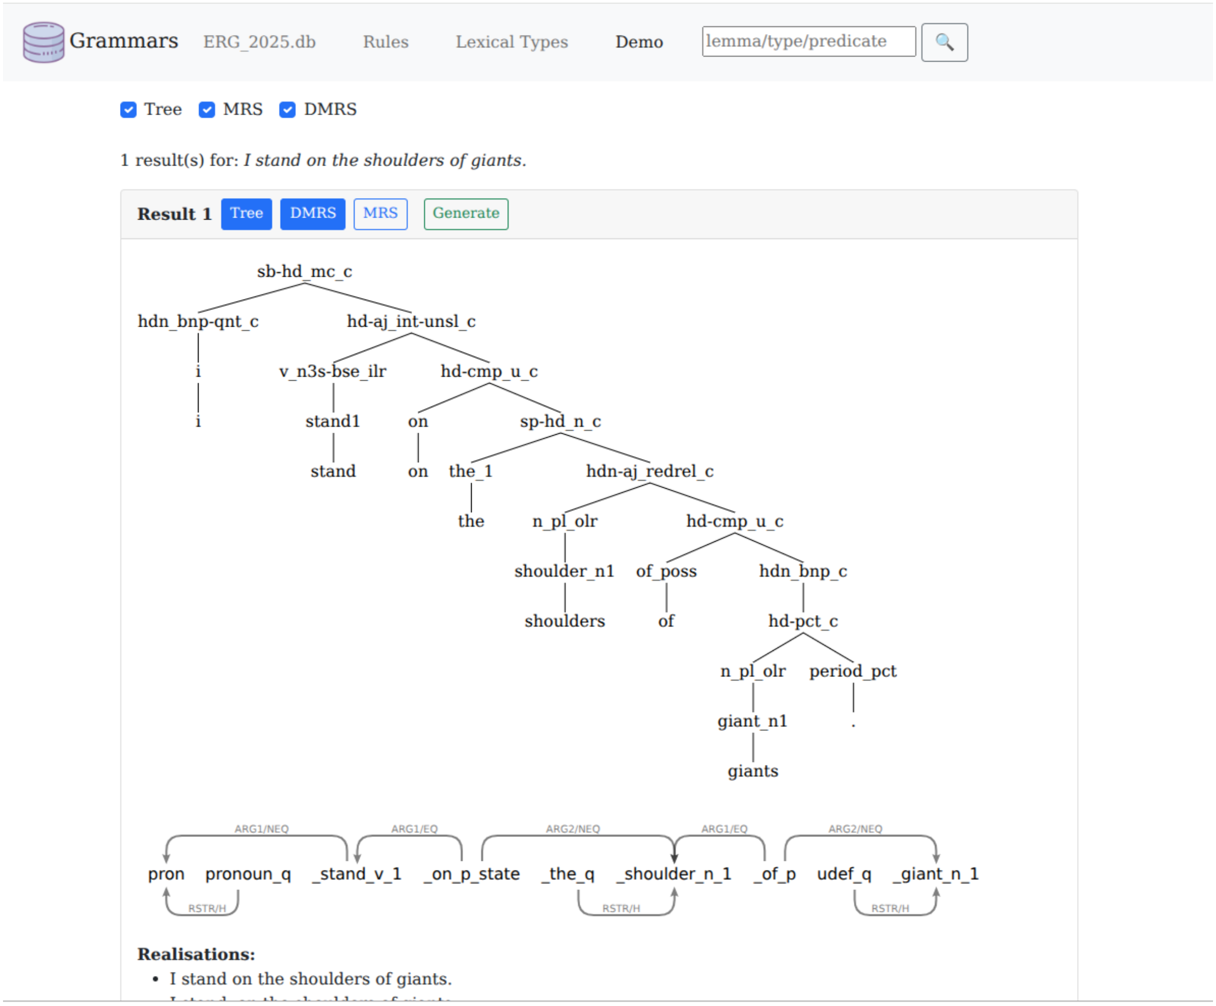

Figure 2: A sentence being parsed in the demo

Family	Languages (subfamily)
Afro-Asiatic	Hausa (Chadic); Arabic, Hebrew (Semitic)
Austronesian	Indonesian
Indo-European	English, German, Norwegian (Germanic); Portuguese, Spanish (Romance); Bulgarian, Russian (Slavic)
Japonic	Japanese
Koreanic	Korean
Mirndi (Australian)	Wambaya
Sino-Tibetan	Mandarin Chinese, Cantonese
Timor–Alor–Pantar	Abui

Table 2: Languages and language families represented in the Grammary.

No treebank is yet available for the grammar, which continues in active development.

4.8. Hebrew

The grammar of Hebrew HeGram: (HeGram: Arad Greshler et al., 2015), developed in parallel with the Arabic grammar described above, includes by far the largest lexicon of the DELPH-IN grammars, as it expands out many morphological patterns. It has an implementation of a range of linguistic phenomena, although there is no accompanying treebank.

4.9. Indonesian (Indra)

Indra started as a Matrix grammar but was substantially extended both in size and to cover phenomena such as valency changing with different affixes and reduplication (Moeljadi et al., 2015). It has two treebanks, JATI — built from dictionary definition sentences and Cendana — built from an online chat corpus (Moeljadi, 2017; Moeljadi et al., 2019).

4.10. Japanese (Jacy)

Jacy is one of the oldest grammars, being originally developed for the Verbmobil project (Siegel et al., 2016). It was one of the grammars used to develop the Grammar Matrix. It has a large lexicon, good coverage, and a treebank large enough to train statistical models (Bond et al., 2008).

4.11. Korean (KRG)

The Korean resource grammar is a medium-size grammar that has gone through two rounds of grammar development (Kim and Yang, 2003; Kim et al., 2011). In the first, the main emphasis was on syntactically interesting phenomena; in the second, there was more emphasis on improving the coverage. It does not have an extensive treebank.

4.12. Mandarin Chinese (Zhong)

Zhong was also developed in two phases, the first ManGO (Yang, 2011) focused on a variety of linguistic phenomena, while in the second there was more emphasis on coverage (Fan et al., 2015). It was subsequently extended to cover ungrammatical phenomena so that it could be used in language teaching, inspired by the ERG's mal-rules (Morgado da Costa and Bond, 2022), and is under continuing development.

4.13. Norwegian

Two substantial grammars of Norwegian are included in the repository: Norsource (Hellan and Haugereid, 2003), developed at NTNU with a large lexicon and broad syntactic coverage; and Norsyg (Haugereid, 2009), developed at the Western Norway University, also with a large lexicon, implementation of a wide range of syntactic phenomena, and a focus on support for incremental left-to-right parsing. The Norsource grammar is accompanied by a moderately sized treebank.

4.14. Portuguese

The grammar of Portuguese (PorGram: `https://github.com/LR-POR/PorGram`) is a relatively young Matrix-derived implementation of Brazilian Portuguese with an illustrative lexicon and a growing set of rules, as yet with no treebank available.

4.15. Russian

Originally developed as part of an ambitious design to produce a set of grammars for several Slavic languages, this grammar of Russian (RRG: Avgustinova and Zhang, 2009) implements a modest set of linguistic phenomena using a small lexicon, with no treebank.

4.16. Spanish

The Spanish Resource Grammar (SRG: Marimon, 2010) is one of the more well developed implementations in this set, with a large lexicon, broad coverage of linguistic phenomena over a range of corpora, and a substantial treebank. The grammar has undergone two phases of development, with the first significant effort taking place in Barcelona, and the more recent second phase in A Coruna (Zamaraeva et al., 2024).

4.17. Wambaya

The grammar of Wambaya (Nordlinger, 1998; Bender, 2010) includes implemented analyses and a lexicon sufficient to account for all of the recorded language data available. The repository includes two versions of the grammar, one with a focus on the aux-verb cluster, and the other on argument composition. Both are accompanied by a treebank with analyses of the available data.

5. Related Work

Several projects share DELPH-IN's goal of developing broad-coverage computational grammars grounded in linguistic theory.

CoreGram (Müller, 2015; Penn, 2004) is an HPSG-based collection running on TRALE,[4] a Prolog-based system for typed feature logic. CoreGram emphasizes a shared *core* of cross-linguistic constraints with language-specific extensions, providing a laboratory for testing theoretical hypotheses about universal versus language-particular properties. It covers several typologically diverse languages (German, Mandarin, Danish, Maltese, Persian, English, Yiddish) with detailed analyses and substantial documentation. The grammars all have test-suites associated with them and can be downloaded. There is also a system Grammix, which allows you to download and run them in a virtual machine. However, they cannot be explored online, and the test-suites are not released with associated parses.

ParGram (Parallel Grammar Project) (Butt et al., 1999, 2002) builds computational grammars in Lexical Functional Grammar using the XLE platform. ParGram aligns cross-linguistic analyses to enable contrastive study while maintaining practical robustness for applications. Its "feature space" discipline—documented conventions for features and

[4]TRALE

mappings—enables parallel grammar development across languages (English, German, French, Norwegian, Japanese, Urdu) (Butt et al., 2003; Dalrymple et al., 2020). While XLE and extensive documentation are publicly available, the grammars themselves are not openly released.

Grammatical Framework (GF) (Ranta, 2011, 2020) is a functional programming language for multilingual grammar applications, combining abstract syntax with language-specific concrete implementations. GF's Resource Grammar Library covers diverse languages (from major European languages to Amharic, Somali, Nepali, Urdu) and is open source under GPL/LGPL. GF excels at multilingual generation and controlled natural language interfaces, though developing a new language to full RGL status requires 3-9 months of expert effort.

Universal Dependencies (UD) (Nivre et al., 2020) takes a fundamentally different approach: rather than executable grammars, UD provides a cross-linguistically consistent annotation scheme for treebanks across 150+ languages. UD's scale—hundreds of treebanks with regular releases—has transformed multilingual parsing and cross-lingual NLP. However, UD focuses on surface syntactic structures without model-theoretic constraints or compositional semantics. In practice, UD and precision grammar collections are complementary: UD offers breadth and statistical learnability; precision grammars provide depth and semantic representations.

6. Discussion and Future Work

The Grammary represents the first FAIR-compliant release of precision grammars with their treebanks, making these resources as accessible and reproducible as UD treebanks. By standardizing access to executable grammars that provide fine-grained linguistic analyses and compositional semantics, the Grammary enables new research directions: comparative evaluation across formalisms, hybrid approaches combining symbolic and statistical methods, and cross-linguistic typological studies with rich semantic annotations.

We plan three main improvements. First, we will compile all grammars with the LKB development environment in addition to ACE, maximizing compatibility and potentially revealing latent bugs. Second, we will establish annual releases synchronized with the DELPH-IN Summit, ensuring the Grammary reflects ongoing community development while providing stable versioned snapshots for reproducible research. Third, we will use the search tools developed for UD to make the syntax trees and semantic dependencies graph-searchable.

By extending the openness and reproducibility principles that have transformed corpus linguistics into the domain of grammar engineering, the Grammary serves as a foundation for sustained research in precision grammar development, semantic parsing, and theoretically grounded computational linguistics.

7. Conclusion

The DELPH-IN Grammary ensures the continued availability of a significant collection of implemented HPSG grammars in widely varying stages of development, many with associated treebanks which help to illustrate the linguistic analyses intended by their authors. By design, the Grammary will continue to reflect the most current version of each grammar and treebank over time, enabling sustained experimentation and evaluation of linguistic analyses across families of languages, all within a common framework and using common reference formalisms for the grammar type definitions in TDL and for the corresponding semantic representations in MRS.

8. Ethics Statement

All resources in the Grammary are distributed under open-source licenses with appropriate attribution to original developers. For grammars of endangered or minority languages, we respect community intellectual property and follow original developers' licensing preferences. By making these resources FAIR-compliant, we aim to democratize access to precision grammar technology for researchers and educators globally.

We acknowledge that NLP technologies can be misused. However, these grammars are primarily research tools requiring expert knowledge, reducing risks of harmful deployment.

9. Limitations

The Grammary has several important limitations: (1) **Coverage:** Grammars typically achieve lower coverage on unseen text than statistical parsers, with substantial variation across the collection. (2) **Diversity:** Limited to 23 grammars across 17 languages, with gaps in African, indigenous American, and Southeast Asian languages. (3) **Expertise:** Effective extension of the grammars requires substantial background in HPSG, MRS, and DELPH-IN tooling, limiting accessibility. (4) **Treebanks:** Not all grammars have treebanks, and sizes vary dramatically (hundreds to tens of thousands of sentences).

Despite these limitations, the Grammary provides valuable resources for research requiring explicit linguistic analysis and compositional semantics.

10. Bibliographical References

Tali Arad Greshler, Livnat Herzig Sheinfux, Nurit Melnik, and Shuly Wintner. 2015. Development of maximally reusable grammars: Parallel development of Hebrew and Arabic grammars. In *Proceedings of the 22nd International Conference on Head-Driven Phrase Structure Grammar, Nanyang Technological University (NTU), Singapore*, pages 27–40, Stanford, CA. CSLI Publications.

Tania Avgustinova and Yi Zhang. 2009. Exploiting the Russian national corpus in the development of a Russian Resource Grammar. In Núria Bel, Erhard Hinrichs, Kiril Simov, and Petya Osenova, editors, *Adaptation of Language Resources and Technology to New Domains at the RANLP 2009 Conference, Borovets, Bulgaria*, pages 1–11. INCOMA Ltd., Shoumen, Bulgaria.

Emily M. Bender. 2010. Reweaving a grammar for Wambaya: A case study in grammar engineering for linguistic hypothesis testing. *Linguistic Issues in Language Technology*, 3.

Emily M. Bender, Dan Flickinger, and Stephan Oepen. 2002. The grammar matrix: An open-source starter-kit for the rapid development of cross-linguistically consistent broad-coverage precision grammars. In *Proceedings of the Workshop on Grammar Engineering and Evaluation at the 19th International Conference on Computational Linguistics*, pages 8–14, Taipei, Taiwan.

Francis Bond, Sanae Fujita, and Takaaki Tanaka. 2008. The Hinoki syntactic and semantic treebank of Japanese. *Language Resources and Evaluation*, 42(2):243–251. (Re-issue of DOI 10.1007/s10579-007-9036-6 as Springer lost the Japanese text).

Francis Bond and Joanna Ut Seong Sio. 2025. Inner and outer aspect in Cantonese. In *Papers Presented at the 32nd International Conference on Head-Driven Phrase Structure Grammar (HPSG 2025)*.

Miriam Butt, Helge Dyvik, Tracy Holloway King, Hiroshi Masuichi, and Christian Rohrer. 2002. The parallel grammar project. In *Proceedings of COLING 2002 Workshop on Grammar Engineering and Evaluation*, pages 1–7.

Miriam Butt, Martin Forst, Tracy H. King, and Jonas Kuhn. 2003. The feature space in parallel grammar writing. In *Proceedings of the ESSLLI 2003 Workshop on Ideas and Strategies for Multilingual Grammar Development*, Vienna, Austria.

Miriam Butt, Tracy Holloway King, María-Eugenia Niño, and Frédérique Segond. 1999. *A Grammar Writer's Cookbook*. CSLI publications.

Siew Yeng Chow and Francis Bond. 2022. Singlish where got rules one? constructing a computational grammar for Singlish. In *13th International Conference on Language Resources and Evaluation (LREC 2022)*.

Ann Copestake. 2002. *Implementing Typed Feature Structure Grammars*. CSLI Publications.

Ann Copestake, Daniel P. Flickinger, Carl J. Pollard, and Ivan A. Sag. 2005. Minimal Recursion Semantics: An introduction. *Research on Language and Computation*, 3:281–332.

Berthold Crysmann. 2005. Relative clause extraposition in German: An efficient and portable implementation. *Research on Language and Computation*, 1(3):61–82.

Berthold Crysmann. 2012. HaG: A computational grammar of Hausa. In *Selected Proceedings of the 42nd Annual Conference on African Linguistics (ACAL 42)*, pages 321–337, Somerville, MA. Cascadilla Press.

Berthold Crysmann and Woodley Packard. 2012. Towards efficient HPSG generation for German, a non-configurational language. In *COLING*, page 695–710.

Mary Dalrymple, Agnieszka Patejuk, and Mark-Matthias Zymla. 2020. Xle+glue: A new tool for integrating semantic analysis in xle. In *Proceedings of LFG 2020*, pages 89–108.

Zhenzhen Fan, Sanghoun Song, and Francis Bond. 2015. Building Zhong [|], a Chinese HPSG meta-grammar. In *Proceedings of the 22nd International Conference on Head-Driven Phrase Structure Grammar (HPSG 2015)*, pages 97–110.

Dan Flickinger. 2000. On building a more efficient grammar by exploiting types. *Natural Language Engineering*, 6(1):15–28. (Special Issue on Efficient Processing with HPSG).

Michael Wayne Goodman. 2019. A python library for deep linguistic resources. In *2019 Pacific Neighborhood Consortium Annual Conference and Joint Meetings (PNC)*, Singapore.

Chikara Hashimoto, Francis Bond, and Dan Flickinger. 2007. The lextype DB: A web-based framework for collaborative multilingual grammar and treebank development. In *First International Workshop on Intercultural Collaboration (IWIC-2007)*, pages 44–58.

Petter Haugereid. 2009. *Phrasal Subconstructions: A Constructionalist Grammar Design, Exemplified with Norwegian and English*. Ph.D. thesis, Norwegian University of Science and Technology.

Lars Hellan and Petter Haugereid. 2003. Norsource – An excercise in the Matrix Grammar building design. In *Proceedings of the ESSLLI 2003 Workshop "Ideas and Strategies for Multilingual Grammar Development"*, Vienna, Austria.

Jong-Bok Kim and Jaehyung Yang. 2003. Korean Phrase Structure Grammar and Its Implementations into the LKB System. In *Proceedings of the 17th Pacific Asia Conference on Language, Information and Computation*.

Jong-Bok Kim, Jaehyung Yang, Sanghoun Song, and Francis Bond. 2011. Processing of Korean and the development of the Korean resource grammar. *Linguistic Research*, 28(3):635–672.

František Kratochvíl. 2026. Abui: A grammatical sketch. In Nicholas Evans and Sebastian Fedden, editors, *The Oxford Guide to the Papuan Languages*, chapter 24. Oxford University Press. Prepublication version available at https://www.academia.edu/121496041/Abui_a_grammatical_sketch_for_the_Oxford_Guide_to_the_Papuan_Languages_.

Hans-Ulrich Krieger and Ulrich Schafer. 1994. TDL- a type description language for constraint-based grammars. In *COLING 1994 Volume 2: The 15th International Conference on Computational Linguistics*.

Jan Tore Lønning, Stephan Oepen, Dorothee Beermann, Lars Hellan, John Carroll, Helge Dyvik, Dan Flickinger, Janne Bondi Johannessen, Paul Meurer, Torbjørn Nordgård, Victoria Rosén, and Erik Velldal. 2004. LOGON. A Norwegian MT effort. In *Proceedings of the Workshop in Recent Advances in Scandinavian Machine Translation*, Uppsala, Sweden.

Montserrat Marimon. 2010. The Spanish resource grammar. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*.

David Moeljadi. 2017. Building JATI: A treebank for indonesian. In *Proceedings of the Atma Jaya Conference on Corpus Studies (ConCorps 2017)*, Jakarta, Indonesia.

David Moeljadi, Francis Bond, and Sanghoun Song. 2015. Building an HPSG-based Indonesian resource grammar (INDRA). In *ACL 2015 Workshop on Grammar Engineering Across Frameworks (GEAF 2015)*.

David Moeljadi, Aditya Kurniawan, and Debaditya Goswami. 2019. Building cendana: a treebank for informal indonesian. In *Proceedings of the 33rd Pacific Asia Conference on Language, Information and Computation (PACLIC 33)*, pages 156–164, Hakodate, Japan.

Luís Morgado da Costa and Francis Bond. 2022. Grammatical error detection using HPSG grammars: Diagnosing common Mandarin Chinese grammatical errors. In *Proceedings of the 29th International Conference on Head-Driven Phrase Structure Grammar (HPSG 2022)*, Online: Nagoya/Tokyo.

Stefan Müller. 2015. The CoreGram project: Theoretical linguistics, theory development and verification. *Journal of Language Modelling*, 3(1):21–86.

Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. Universal Dependencies v2: An evergrowing multilingual treebank collection. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.

Rachel Nordlinger. 1998. *A Grammar of Wambaya, Northern Australia*. Pacific Linguistics, Canberra.

Stephan Oepen and Dan Flickinger. 2019. The ERG at MRP 2019: Radically compositional semantic dependencies. In *Proceedings of the Shared Task on Cross-Framework Meaning Representation Parsing at the 2019 Conference on Natural Language Learning, CoNLL 2019, Hong Kong, Novemer 3, 2019*, pages 40–44. Association for Computational Linguistics.

Stephan Oepen, Dan Flickinger, and Francis Bond. 2004a. Towards holistic grammar engineering and testing — grafting treebank maintenance into the grammar revision cycle. In *Beyond Shallow Analyses — Formalisms and Statistical Modelling for Deep Analysis (Workshop at IJCNLP-2004)*, Hainan Island.

Stephan Oepen, Daniel P. Flickinger, Kristina Toutanova, and Christopher D. Manning. 2004b. LinGO redwoods: A rich and dynamic treebank for HPSG. *Research on Language and Computation*, 2(4):575–596.

Woodley Packard. 2015. Full forest treebanking. Master's thesis, University of Washington.

Gerald Penn. 2004. Balancing clarity and efficiency in typed feature logic through delaying. In *Proceedings of the Association for Computational Linguistics 2004*, pages 239–246.

Carl J. Pollard and Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. Studies in Contemporary Linguistics. The University of Chicago Press, Chicago.

Aarne Ranta. 2011. *Grammatical Framework: Programming with Multilingual Grammars*. CSLI Publications, Stanford.

Aarne Ranta. 2020. Scaling up the grammatical framework from controlled languages to robust pipelines. *Computational Linguistics*, 46(3):529–562.

Melanie Siegel, Emily M. Bender, and Francis Bond. 2016. *Jacy: An Implemented Grammar of Japanese*. CSLI Studies in Computational Linguistics. CSLI Publications.

Kiril Simov, Petya Osenova, Alexander Simov, and Milen Kouylekov. 2004. Design and implementation of the Bulgarian HPSG-based treebank. *Research on Language and Computation*, 2(4):495–522.

Joanna Ut-Seong Sio and Sanghoun Song. 2015. Divergence in expressing definiteness between Mandarin and Cantonese. In *Proceedings of the 22nd International Conference on Head-Driven Phrase Structure Grammar (HPSG 2015)*, pages 178–195.

Pat Suppes, T. Liang, E. E. Macken, and Dan Flickinger. 2014. Positive technological and negative pre-test-score effects in a four-year assessment of low socioeconomic status K-8 student learning in computer-based math and language arts courses. *Computers & Education*, 71:23–32.

Hans Uszkoreit. 2002. New chances for deep linguistic processing. *Proceedings of the 19th International Conference on Computational Linguistics (COLING 2002)*.

Hans Uszkoreit, Daniel Flickinger, Walter Kasper, and Ivan A. Sag. 2000. Deep analysis with HPSG. In Wolfgang Wahlster, editor, *Verbmobil. Foundations of Speech-to-Speech Translation*, artificial intelligence edition, pages 321 – 330. Springer, Berlin, Germany.

Wolfgang Wahlster, editor. 2000. *Verbmobil: Foundations of Speech-to-Speech Translation*. Artificial Intelligence. Springer Verlag, Berlin/""Heidelberg/""New York, NY.

Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al. 2016. The fair guiding principles for scientific data management and stewardship. *Scientific data*, 3.

Justin Chunlei Yang. 2011. Mandarin grammar online: ManGO. Presentation at the 2011 DELPH-IN Summit.

Olga Zamaraeva, Lorena S. Allegue, and Carlos Gómez-Rodríguez. 2024. Spanish resource grammar version 2023. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15093–15104, Torino, Italia. ELRA and ICCL.