

Grammatist at WMT 2026 General MT Task: An Agentic LLM Framework for Machine Translation with Explicit Linguistic Knowledge

Samuel Frontull^{1,*}, Jean Maillard^{2,*},
Christopher Haberland^{3,2}, Ruth Videsott⁴, Stefano Lusito²

¹University of Innsbruck, Austria, ²Council for Ligurian Linguistic Heritage, Italy,

³University of Washington, USA, ⁴Free University of Bozen–Bolzano, Italy

* Corresponding authors, equal contribution: samuel.frontull@uibk.ac.at, jean@maillard.it

Abstract

We describe Grammatist,¹ our submission to the unconstrained track of the WMT 2026 General Machine Translation Task, developed in direct collaboration with community members of two regional languages of Italy, Ladin and Ligurian. Ladin and Ligurian represent challenging scenarios for machine translation (MT), as they have limited resources for conventional data-driven approaches. However, this scarcity of large-scale parallel data does not reflect a lack of linguistic documentation: both languages are extensively documented in dictionaries and grammar books, which contain valuable knowledge but are not readily exploitable by conventional MT approaches. Grammatist addresses this gap by making this information available to a large language model (LLM) at inference time. The framework is both agentic and model-agnostic: it dynamically consults dictionaries and grammar books during translation to retrieve the information most relevant to the input. Any LLM can serve as the underlying base model, enabling Grammatist to benefit from advances in increasingly capable models. For the WMT 2026 task, we apply Grammatist to structured document-level translation from English into Ligurian and Ladin.

1 Introduction

Modern approaches to low-resource machine translation commonly address limited supervision by adapting pre-trained multilingual models (fine-tuning) and by augmenting available parallel data (back-translation). While effective in many settings (Haddow et al., 2022), these approaches remain fundamentally dependent on the information that can be learned from the available parallel data, where rare words and linguistic patterns may receive too little evidence.

However, limited parallel data does not necessarily imply a lack of linguistic documentation. Many

languages are well documented with carefully curated grammars, dictionaries, and other linguistic resources and Ligurian and Ladin are representative examples of such languages. Approaches for MT that rely on parallel corpora, however, provide no direct mechanism for exploiting the information contained in other forms of linguistic documentation. As a result, substantial linguistic knowledge may remain inaccessible to the translation system.

Large language models (LLMs) offer an opportunity to approach this setting from a broader perspective. Their ability to perform in-context learning, interact with tools, and take actions based on the information they retrieve opens a different paradigm for (low-resource) MT that makes it possible to consider linguistic resources beyond parallel corpora (Zhang et al., 2024; Pei et al., 2025; Marmonier et al., 2025). This is the idea behind *Grammatist*, an agentic framework for machine translation. Grammatist combines LLMs with linguistic resources through prompting and tool use. The system was previously introduced and evaluated for sentence-level Italian–Ligurian and Italian–Ladin translation (Maillard et al., 2026). In this work:

- we apply Grammatist to English-to-Ligurian and English-to-Ladin (document-level) translation;
- we show how LLM reasoning capabilities can be used to dynamically retrieve and exploit linguistic knowledge beyond the language pair directly represented in available dictionaries;
- we evaluate the approach in the WMT26 General Machine Translation Task and analyse the effectiveness and limitations of tool-assisted translation for under-resourced languages.

With this approach, we propose a flexible framework for incorporating external linguistic knowledge into LLM-based machine translation at inference time.

¹<https://github.com/ConseggioLigure/grammatist>




Figure 1: Overview of the Grammatist agentic translation framework, from (Maillard et al., 2026). While the framework is able to also leverage parallel data, no bitext was used in the experiments described here.

2 System Description

Grammatist is an agent-based machine translation tool that provides LLMs with access to dictionaries, enabling them to autonomously consult and leverage these resources to generate translations. The overall architecture of the system is sketched in Figure 1. In this section, we first describe the linguistic resources used for Ladin and Ligurian translation. We then detail how we implement the agentic translation framework and how the LLM interacts with these resources during translation.

2.1 Linguistic Resources

Dictionaries We use two bilingual dictionaries as the primary lexicographical resources for the regional languages considered in this work: the Italian–Ligurian dictionary (Maillard and Lusito, 2021) and the Italian–Ladin (Val Badia) dictionary (Moling, 2016). Table 1 summarises their structural properties, highlighting the richness of the linguistic information they encode.

Both resources provide substantial lexical coverage, with 13,704 Italian source lemmas in the Ligurian dictionary and 29,737 in the Ladin dictionary. Beyond word-to-word mappings, they include multi-word expressions (16.6% and 20.8% of lemmas, respectively) and distinguish homonyms or part-of-speech-specific entries, preserving important lexical and semantic distinctions.

The dictionaries also capture substantial translation ambiguity, containing 18,598 and 34,982 word senses and associating many senses with multiple target-language equivalents. In particular, 34.8% of Ligurian senses and 47.8% of Ladin senses have more than one translation.

Beyond lexical correspondences, both resources provide inflection tables and usage examples that support context-sensitive interpretation and gener-

ation of morphologically valid target forms. Although these resources are indexed from Italian rather than English, they remain useful for WMT26 because the LLM can bridge languages at inference time.

However, effectively exploiting the information depends on the reasoning capabilities of the underlying model: the system must identify the relevant sense, relate retrieved Italian-based evidence to the English source context, and select an appropriate target-language realisation. This makes the resources particularly suitable for an agentic framework, where retrieval and reasoning are combined during translation.

The dictionaries are stored as .jsonl files. Each line contains a headword and one or more entries, allowing distinct lexical entries, including homonymous and part-of-speech-specific entries, to remain separate. An entry can contain definitions, one or more target-language equivalents, bilingual usage examples, multiword expressions, and inflectional tables.

Grammar Rules Besides dictionary evidence, Grammatist incorporates general grammatical rules. The grammar rules provide the model with explicit information about morphosyntactic phenomena that are not reliably captured by transfer from Italian or by general multilingual pretraining. Both languages exhibit grammatical structures that differ substantially from Italian, such as obligatory subject clitics in Ligurian (e.g., *O l’à*) and verb-particle constructions in Ladin (e.g., *surantó...fora*). These language-specific phenomena require explicit linguistic guidance to produce idiomatic translations.

Grammar guidance is directly added to the system prompt. The grammar tips used for the submission are given in Appendix B. More details are provided in Maillard et al. (2026).

	Italian–Ligurian Maillard and Lusito (2021)	Italian–Ladin (Val Badia) Moling (2016)
1. Source Language Coverage (Italian Input)		
Unique source lemmas	13,704	29,737
<i>e.g., mela, mangiare</i>	13,704	29,737
Multiword expression lemmas	2,278 (16.6%)	6,197 (20.8%)
<i>e.g., per favore, prima donna</i>	2,278 (16.6%)	6,197 (20.8%)
Lemmas with multiple entries (POS/homonym splits)		
<i>e.g., piano</i> → [Entry 1: <i>noun</i>] vs. [Entry 2: <i>adverb</i>]	201 (1.5%)	2,008 (6.8%)
2. Dictionary Structure & Macrostructure		
Total entry blocks	13,917	31,979
Word senses (semantic definitions)		
<i>e.g., capo</i> → <i>Sense 1: "boss"</i> vs. <i>Sense 2: "head"</i>	18,598	34,982
Average entries per lemma	1.016	1.075
Average senses per entry	1.336	1.094
3. Target Language Output & Ambiguity (Ligurian / Ladin)		
Total translation equivalents	28,108	70,496
Unique target surface forms	15,434	31,359
Senses with multiple target translations (synonymy)	6,468 (34.8%)	16,738 (47.8%)
Average translations per sense	1.511	2.015
4. Richness & Attestation (Examples & Grammar)		
Total example usage sentences	13,747	18,556
Senses accompanied by examples	11,454 (61.6%)	10,204 (29.2%)
Average examples per sense	0.739	0.530
Entries with inflection tables	7,180 (51.6%)	19,693 (61.6%)

Table 1: Detailed statistical profile of the bilingual lexicographical resources.

2.2 System Architecture

Chunking and Prompting. The input text is first divided into smaller chunks. The chunking algorithm is breakpoint-oriented, with minimum and maximum chunk lengths of 400 and 800 characters. It searches for a boundary in the following priority order: (i) whitespace or a newline immediately following a sentence-ending period, exclamation mark, or question mark; (ii) a colon or semicolon; (iii) a comma; and (iv) any whitespace. Breakpoints before the 400-character lower bound are skipped, and the algorithm then searches for a breakpoint up to 800 characters, considering these types in the order above.²

The chunks are then independently processed. For each chunk, the model is given a system prompt, a task-specific domain instruction, and a native function calling-tool (`dictionary_lookups`) for dictionary lookups.

The system prompt instructs the model to identify every source word, including function words, lemmatise inflected forms, and request the relevant verb tenses needed to formulate dictionary

lookup queries.³ The model is then instructed to use the retrieved dictionary evidence together with the specified grammar rules to produce the final translation. The prompt also defines the tool-use policy and the conditions under which dictionary queries should be issued. In particular, it states: “*If the source text is not Italian, translate it to Italian and then proceed as above.*” Thus, for English input, the model first translates the source into Italian and uses the resulting Italian forms as headwords for subsequent Italian–Ligurian or Italian–Ladin dictionary lookups. The complete system prompt is provided in Appendix A.

Figure 2 illustrates the interface between the model and the dictionary lookup tool. The left panel shows the strict JSON schema registered with the model API, which specifies the structure of the tool arguments and the grammatical tenses that can be requested. The right panel shows an example of a batched lookup request generated by GPT-5.4 for the source phrase “*I want to eat an apple now.*”

Tool and Dictionary Evidence. Given the Italian phrase “*mangio una mela*”, the model could emit the structured request shown in Figure 2b. The verbal lemma receives a tense request, while the non-

²If no breakpoint of one type is found, it moves to the next; if none is found, the chunk is cut at 800 characters. A short final segment is merged with the preceding chunk when their combined length remains within the 800-character maximum.

³It also permits synonym and fallback queries when a preceding lookup fails.

```

{
  "type": "function",
  "name": "dictionary_lookups",
  "description": "Look up dictionary words.",
  "parameters": {
    "type": "object",
    "additionalProperties": false,
    "required": ["dictionary_lookups"],
    "properties": {
      "dictionary_lookups": {
        "type": "array",
        "items": {
          "type": "object",
          "additionalProperties": false,
          "required": ["lemma", "tenses"],
          "properties": {
            "lemma": {"type": "string"},
            "tenses": {
              "type": "array",
              "items": {"type": "string",
                "enum": [
                  "IndPres", "IndImp", "IndFut", "ConzPres",
                  "ConzImp", "Cnd", "Imp", "Inf", "Ger", "PartPast"
                ]
              }
            }
          }
        }
      }
    }
  },
  "strict": true
}

```

(a) Tool API schema definition.

```

{
  "dictionary_lookups": [
    {
      "lemma": "io",
      "tenses": []
    },
    {
      "lemma": "volere",
      "tenses": [
        "IndPres",
        "Inf"
      ]
    },
    {
      "lemma": "mangiare",
      "tenses": [
        "Inf"
      ]
    },
    {
      "lemma": "uno",
      "tenses": []
    },
    {
      "lemma": "mela",
      "tenses": []
    },
    {
      "lemma": "adesso",
      "tenses": []
    }
  ]
}

```

(b) Model-generated batched lookup requests.

Figure 2: Integration of the dictionary lookup tool. The subfigure 2a shows the strict JSON schema defining the tool’s parameters and allowed grammatical tenses. The subfigure 2b displays the actual batched lookup request generated by GPT-5.4 for the source phrase “*I want to eat an apple now.*” Notice that the model correctly translates the English input internally to query the corresponding Italian lemmas.

verbal entries receive empty lists. The local tool performs exact, lower-cased lemma lookup, merges duplicate requests, and returns a condensed representation rendered as Markdown, as illustrated in Appendix D.

Although the general framework also supports parallel-data retrieval, the WMT26 system described here uses only dictionaries and grammar rules; apart from example sentences included in the dictionaries, no additional parallel corpus is used.

Agent Loop and Generation. The agent loop is designed to keep tool use efficient and bounded while preserving the necessary context across calls. The Responses API maintains the conversation through `previous_response_id`, so after a tool call the client sends only the new tool result and user input instead of replaying the entire conversation history. In the WMT26 configuration, each chunk therefore uses at most one tool-use round and one batched tool invocation. The translated chunks are then concatenated in their original order to produce the final translation.

3 Validating the Approach

Before applying Grammarist to the WMT26 blind test set, we first validate the approach on development benchmarks. We evaluate both Ladin and Ligurian on FLORES+ (Costa-jussà et al., 2024; Frontull et al., 2025) and Ligurian additionally on BOUQuET (Andrews et al., 2025).⁴

Table 2 compares Grammarist with a zero-shot baseline using the same GPT-5.4 model (gpt-5.4-2026-03-05). The baseline was prompted to translate source texts using only core task instructions (target language specification and output format). Grammarist uses the same model as its generation backend but additionally provides the external information and tools described in Section 2.2.⁵ Grammarist substantially improves translation quality across all three evaluation settings. The observed gains provide evidence that the additional

⁴The Ladin subset of the BOUQuET benchmark was not publicly available during system development.

⁵The two systems differ in both the information available to the model and the resulting prompt structure, so the results should not be interpreted as an isolated ablation of the dictionary component.

Direction	Dataset	Baseline	Grammatist	Tokens (M) In / Out	Inference Avg. (s)	Dictionary lookups		
		BLEU / chrF	BLEU / chrF			Found	Missing	Success
eng→lij	BOUQuET	12.41 / 40.30	28.96 / 53.94	6.37 / 1.39	29.68	9,472	885	91.46%
eng→lij	FLORES+	14.62 / 46.85	29.89 / 59.01	9.96 / 2.48	40.50	17,875	2,809	86.42%
eng→lld	FLORES+	8.17 / 38.31	26.53 / 54.48	9.08 / 2.80	40.90	18,833	2,041	90.22%

Table 2: BLEU / chrF scores for baseline translation and Grammatist. The baseline employs pure GPT-5.4 (gpt-5.4-2026-03-05) prompted zero-shot to translate source texts with only the core task instruction and target language specification, completely omitting external context, grammar tips, or dictionary tools. Grammatist, for which we also report token usage, inference time, and dictionary lookup statistics, utilises this same model as backend.

linguistic information and structured tool use are useful. Importantly, the intermediate English-to-Italian translation does not appear to be a significant bottleneck: the approach performs consistently well on both Ligurian and Ladin. Table 2 also reports inference and retrieval statistics for Grammatist. The system processes between 6.37M and 9.96M input tokens across the evaluation sets, with average inference times of approximately 30–41 seconds per sentence. Dictionary lookup success remains high, ranging from 86.42% to 91.46% across the three settings showing most lexical candidates proposed during translation can be resolved directly through the dictionary. Although additional retrieval iterations could potentially recover some of these candidates, doing so would increase both inference time and token usage.

These development results validate the effectiveness and practical feasibility of the core Grammatist approach. Across both languages and evaluation benchmarks, the system achieves translation quality beyond previously reported results for these language pairs (Maillard et al., 2026; Frontull and Ströhle, 2026). The development experiments therefore provided sufficient evidence that the approach was robust enough to scale to the WMT26 blind test set, where we adapted the system only with some domain-specific instructions.

4 WMT 2026 General MT Shared Task

In previous work, Grammatist was evaluated on the FLORES+ (Dale et al., 2025) and BOUQuET (Andrews et al., 2025) benchmarks. Here, we study its performance in the document-level setting of WMT, which introduces different challenges related to longer contexts, document coherence, and structured input preservation. These characteristics provide a natural setting for investigating the capabilities and practical costs of LLM-based translation systems.

Translation of the WMT26 Testset Translations of the WMT26 test set are generated using a dedicated inference pipeline. All experiments use GPT-5.4 (openai/gpt-5.4-2026-03-05) with medium reasoning effort, allowing a single tool call per chunk translation request. Text documents are translated in a single model invocation using the task-specific instruction provided by the benchmark. The same procedure is used for the *software* domain: despite the additional requirement to preserve valid JSON structure, no domain-specific processing or post-processing is required beyond including the provided software-domain instructions in the prompt. The framework produces structurally valid translated JSON directly, preserving the top-level type, object keys, and list lengths. This capability depends in part on the instruction-following and structured-output abilities of the underlying LLM.

All generated translations undergo automatic validation to guarantee compatibility with the WMT evaluation pipeline. Specifically, the system enforces non-empty outputs, exact line-count parity with the source document, and preservation of blank-line alignment. Translations that violate these constraints are automatically regenerated, up to three additional attempts, ensuring that formatting errors do not affect evaluation.

Resource Usage We report the average inference time, number of LLM (tool) calls, and token consumption per document in the WMT test set in Table 3. The number of calls depends primarily on document length, as documents are divided into independently translated chunks (only a single tool call allowed for each chunk). Grammatist requires relatively small chunks because retrieved dictionary entries and other linguistic evidence substantially expand the model context during translation (as can be seen in the examples in Appendix D). Longer documents therefore require more calls and

Table 3: Inference cost analysis per document for translating the WMT test set across domains for Ligurian and Ladin. N indicates the number of evaluated document files per domain. All values are averages per document.

Domain	N	Ligurian (lij_Latn)			Ladin (lld_Latn)		
		Time (s)	LLM Calls	Tokens (In/Out)	Time (s)	LLM Calls	Tokens (In/Out)
news	14	240.64	11.00	197,532.8 / 24,052.7	268.03	11.00	144,619.1 / 25,678.1
social	44	143.32	5.77	117,180.1 / 14,657.8	156.09	5.77	81,253.9 / 15,440.1
software	7	222.47	8.57	137,597.7 / 21,692.3	226.36	8.57	100,895.3 / 22,127.9
speech	133	44.18	2.00	34,192.7 / 4,134.3	44.13	2.00	23,461.4 / 4,152.9

consume more tokens.

Allowing more extensive tool use could increase these requirements further. For example, on the news test set for English–Ligurian translation, the system requires an average of 240.64 seconds and 197.5K input tokens per document, while shorter documents in the speech domain require 44.18 seconds and 34.2K input tokens. Similar trends are observed across the English–Ladin test sets. These results quantify the resource requirements of the submitted system and provide insight into the practical cost of applying LLM-based translation systems to WMT-scale evaluation scenarios.

5 Conclusion

We presented Grammarist, an agentic machine translation framework applied to English–Ligurian and English–Ladin translation submitted to the WMT 2026 General Machine Translation Task. The system combines the multilingual capabilities of a general-purpose LLM with explicit linguistic knowledge retrieved dynamically from Italian–Ligurian and Italian–Ladin dictionaries for English-to-target translation. We discuss the observed limitations and directions for future work.

Limitations and Future Work Grammarist has substantially higher inference requirements than direct LLM translation, partly because retrieved linguistic resources expand the context and require relatively small input chunks. To limit resource consumption in our WMT submission, we allowed only a single tool-use round per chunk. This restriction may limit the amount of linguistic evidence the model can retrieve for difficult translations; allowing additional tool use could improve translation quality, by allowing e.g. the system to recover from cases where a queried lemma is missing from a lexical resource, but would further increase token consumption, latency, and cost. Future work should therefore investigate more selective retrieval and more compact representations of retrieved re-

sources.

The chunk-based processing introduces some limitations at the document level. Chunks are non-overlapping and translated independently, with no source or translation context, tool-call history, or document-level memory carried from one chunk to the next. As a result, document-level coherence is not explicitly enforced across chunks and may be affected, particularly when resolving context-dependent expressions or maintaining consistent terminology. A further limitation comes from restricting the system to a single tool call per chunk. The LLM may not always know which words or expressions exist, and a single lookup call also prevents it from performing a follow-up search for the lookups that fail. A more detailed analysis of retrieval failures and of the effect of these limitations on translation quality is left for future work.

Acknowledgments

We thank the *Istitut Ladin Micurá de Rù* and the *Council for Ligurian Linguistic Heritage* for their support and for providing lexical and grammatical resources. We thank the organizers of the WMT 2026 General Machine Translation Shared Task for their work in organizing the task and supporting participants throughout the submission and evaluation process.

References

- Pierre Andrews, Mikel Artetxe, Mariano Coria Meglioli, Marta R. Costa-jussà, Joe Chuang, David Dale, Mark Duppenthaler, Nathaniel Paul Ekberg, Cynthia Gao, Daniel Edward Licht, Jean Maillard, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Eduardo Sánchez, Ioannis Tsiamas, Arina Turkatenco, Albert Ventayol-Boada, and Shireen Yates. 2025. [BOUQuET : dataset, benchmark and open initiative for universal quality evaluation in translation](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 27515–27535, Suzhou, China. Association for Computational Linguistics.

- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Mailard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, and 20 others. 2024. [Scaling neural machine translation to 200 languages](#). *Nature*, 630(8018):841–846.
- David Dale, Laurie Burchell, Jean Maillard, Idris Abdulmumin, Antonios Anastasopoulos, Isaac Caswell, and Philipp Koehn. 2025. [Findings of the WMT 2025 shared task of the open language data initiative](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 495–502, Suzhou, China. Association for Computational Linguistics.
- Samuel Frontull, Thomas Ströhle, Carlo Zoli, Werner Pescosta, Ulrike Frenademez, Matteo Ruggeri, Daria Valentin, Karin Comploj, Gabriel Perathoner, Silvia Liotto, and Paolo Anvidalfarei. 2025. [Bringing Ladin to FLORES+](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 1061–1071, Suzhou, China. Association for Computational Linguistics.
- Samuel Frontull and Thomas Ströhle. 2026. [Every word presented in context: Syntactic coverage as objective for low-resource machine translation with large language models](#). In *Proceedings of the Fifteenth Language Resources and Evaluation Conference (LREC 2026)*, pages 8824–8837, Palma, Mallorca, Spain. European Language Resources Association (ELRA).
- Barry Haddow, Rachel Bawden, Antonio Valerio Miceli Barone, Jindřich Helcl, and Alexandra Birch. 2022. [Survey of low-resource machine translation](#). *Computational Linguistics*, 48(3):673–732.
- Jean Maillard, Samuel Frontull, Christopher Haberland, Ruth Videsott, and Stefano Lusito. 2026. How far can an agentic framework close the gap between llms and specialised mt for low-resource languages? In *Proceedings of the Twelfth Italian Conference on Computational Linguistics (CLiC-it 2026)*, Palermo, Italy. CEUR Workshop Proceedings. To appear.
- Jean Maillard and Stefano Lusito. 2021. DEIZE: Diçionäio eletrônico italian-zeneise. <https://conseggio-ligure.org/it/dizionario/deize>. Online dictionary, accessed 4 June 2026.
- Malik Marmonier, Rachel Bawden, and Benoît Sagot. 2025. [Explicit learning and the LLM in machine translation](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 31372–31422, Suzhou, China. Association for Computational Linguistics.
- Sara Moling, editor. 2016. *Dizionario Italiano-Ladino Val Badia / Dizionar Ladin Val Badia-Talian*. Istitut Ladin “Micurà de Rü”, San Martin de Tor (BZ).
- Renhao Pei, Yihong Liu, Peiqin Lin, François Yvon, and Hinrich Schuetze. 2025. [Understanding in-context machine translation for low-resource languages: A case study on Manchu](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8767–8788, Vienna, Austria. Association for Computational Linguistics.
- Kexun Zhang, Yee Man Choi, Zhenqiao Song, Taiqi He, William Yang Wang, and Lei Li. 2024. [Hire a linguist!: Learning endangered languages in LLMs with in-context linguistic descriptions](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15654–15669, Bangkok, Thailand. Association for Computational Linguistics.

A System Prompt

The text below is the system prompt used in Grammatist. The WMT26 submission allowed one tool-use round per chunk.

```
You are a professional translator tasked with
translating Italian text into {Ligurian (Genoese)|Ladin
(Val Badia)}. Given some Italian source text, you must:
1. Use the relevant tools that are provided to you to
look up essential information.
2. Use the output of the tool calls and the grammar tips
below to produce a translation.
```

Step 1: Using the tools

```
Identify every single word in the source text, including
function words and trivial words, and look them up by
making a native tool call with the "dictionary_lookups"
tool.
```

```
Provide a list of relevant tenses for the verb lookups.
Shape: `{"lemma":"...", "tenses":[...]}`.
```

```
For verbs, select only the strictly needed tenses you'd
want to look up in "tenses" amongst: IndPres (present
indicative), IndImp (imperfect), IndFut (future),
ConzPres (present subjunctive), ConzImp (imperfect
subj.), Cnd (conditional), Imp (imperative), Inf, Ger,
PartPast.
```

```
For compound Italian tenses, request the needed
{Ligurian|Ladin} finite auxiliary tense plus PartPast.
Present participles should be looked up individually as
if they were nouns or adjectives. For non-verbs,
"tenses" should be an empty list.
```

```
For pronominal verbs, request the base verb plus any
particles when useful.
```

```
You can make multiple tool calls in subsequent rounds. If
a word isn't found in the dictionary, look up a synonym.
If the dictionary entry lists a useful multi-word
expression, look it up. You are allowed a limited number
of rounds of tool calling. You will be told how many
rounds you have left, after which you will be asked to
translate.
```

IMPORTANT:

1. You MUST look up EVERY SINGLE WORD, including trivial prepositions like "in" and "a", conjunctions like "e" and "o", determiners, etc. Failure to look up even a single word will be considered a critical error and will lead to a poor translation.
2. You MUST preserve all capitalization, white space (including trailing spaces), and markup in the translation.

Step 2: Output the translation

```
Translate the source text using the grammar tips and the
output from the dictionary calls, which will contain
usage examples, relevant conjugation tables and
declension tables. Output only the translated text,
preserve capitalization, white space and markup.
```

B Grammar Tips

Key grammar tips

In Ligurian, subject clitics are mandatory before verbs in the 2nd and 3rd person singular.
2nd singular: *_ti_* (or *_t'_* before vowels): *mangi* → *ti mangi*, *guardi* → *t'ammi*.
3rd singular: *_o_* (m) / *_a_* (f), with *_l'_* before vowels: *il cane mangia* → *o can o mangia*.
Subject clitics are not used in impersonal constructions (*oggi piove* → *ancheu ceuve*), in nominalized infinitives (*mi piace giocare* → *me piaxe zugà*), or with perception verbs (*mi sembra di esserci stato* → *me pà d'èsighe stàto*).

Ligurian has two polite pronouns, *_voscià_* and *_voi_*.
Voscià is used with interlocutors of higher or upper social class; it conjugates like the 3rd person singular and requires the subject clitic *_scià_*: *lei lavora* → *voscià scià travaggia*; *lei non è d'accordo* → *voscià scià no l'è d'accòrdio*.

Voi is used among people of the same lower-middle class or with people of lower social status and conjugates like the 2nd person plural: *lei lavora* → *voi travaggiæ*; *lei non è d'accordo* → *voi no sei d'accòrdio*.
When in doubt, use *_voscià_*. Never translate the Italian polite pronoun *_lei_* with *_lê_*.

The modal verbs *_dovei_* (must), *_poei_* (can), and *_voei_* (want) always take the auxiliary *_avei_* (to have), regardless of the verb they govern: *me ne sono dovuto andare* → *ò dovuo anàmene*; *non è potuta uscire* → *a no l'à posciuo sciortì*.

When in Ligurian the verb precedes the subject, the verb is always in the singular, without a clitic: *giocano i bambini* → *zeuga i figgeu*; *cadono le foglie* → *cazze e feugge*; *mi piacciono i ravioli* → *me piaxe i ravieu*.
This rule also applies to compound tenses, in which case the past participle is also masculine singular: *sono venute le zie* → *l'è vegnuo e lalle*.
Excluded from this rule are cases where the verb has a presentational function (generally the verb *_èse_*): *sono bambini* → *en figgeu*; *erano giochi popolari a quei tempi* → *ean zeughi popolari à quelli tempi*.

To form dates in Ligurian:

- before the day number, use the articulated preposition *_a-i_* (but for the first of the month, use the singular *_a-o_* with the ordinal indicator: *_a-o l'z_*);
- before the month name, use the simple preposition *_de_* (elided to *_d'_* if the month begins with a vowel: *d'agosto*, *d'ottobre*);
- before the year number, use the articulated preposition *_do_*.

Example: *la festa è martedì 1 febbraio 2000* → *a festa a l'è mätēsdi a-o l'z de frevâ do 2000*.
Before the day number, it is also possible to use simply the article *_i_* (or *_o_* for the first of the month).
Using only the article is mandatory when the date is part of a nominal predicate: *domani è il 13 settembre 2004* → *doman l'è i 13 de dexembre do 2004*.

Morphological adaptation rules for words not present in the dictionary: *-ale* → *-ale*; *-are* → *-are*; *-aria/-ario* → *-äia/-äio*; *-ese* → *-eise*; *-ezza* → *-essa*; *-ia/-ie* → *-ia/-ie*; *-ile* → *-i*; *-ino/-ina* → *-in/-iña*; *-iolo/-iola* → *-ieu/-ieua*; *-ismo* → *-iximo/-ismo*; *-issimo* → *-iscimo*; *-ito/-ita* → *-io/-ia*; *-tà* → *-tæ*; *-ura* → *-ua*; *-uto/-uta* → *-uo/-ua*; *-zione* → *-çion*.

Key grammar tips

In the Ladin of Val Badia, expressing the subject is mandatory. When the subject is represented by a personal pronoun (whether stressed or clitic), the pronoun is always mandatory.
The Ladin pronominal system has two series of pronouns: stressed (tonic) pronouns and unstressed (atonic) pronouns (which are clitics).
In most cases, clitics are used. These have a fixed position: either directly before the verb or attached directly to the verb in an enclitic position (if an inversion occurs).

Examples with clitic pronouns: *_Al sà da cujiné_* (He knows how to cook.); *_I sun gnüs adalerch cun la ferata_* (I arrived here by train.)

Stressed pronouns are used only in specific contexts: when two subjects are linked by a conjunction and the pronoun does not hold a fixed position before the verb (*_Èra y sü colaboradus á ciafé le pest por proiec inovatifs tl artejanat_*); when a new referent appears in the text (*_Francësch ê bele n pü' ti agn. Èra, süa fomena, ê indere bindebó jona y plëna d'energia_*); when there is a need to emphasize the subject because it holds important informative value (*_La maiú pert dla iuria é dla minunga che ël sides le miú y toches da ciafé le post pro l'orchestra_*); when there is a contrast between two elements (*_Os m'ëis dit che al foss da mudé auto, iô foss bëgn ciamó stada contënta cun le vedl_*).

In the Ladin of Val Badia, the finite verb has a fixed position in declarative sentences: the second position. When the first position is occupied not by the subject, but by an object, an adverbial construction, or a subordinate clause, the verb and the subject must invert: the verb remains in the second position, while the subject shifts behind the verb.

Sentence with an initial subject: *_Maria sona la chitara_*

Sentences with adverbial constructions in the initial position: *_Sëgn sona Maria la chitara_*; *_Cun ligrëza sona Maria la chitara_*

Sentence with an initial object: *_Les audanzas á le diretur ciafé de porsona_*

Verb-subject inversion is also mandatory for interrogative sentences: *_Tan dí á pa le paura restelé?_*

When the verb and the subject are inverted and the subject is a subject clitic, the clitic is attached directly to the verb ending.

Sentence with an initial subject: *_I sun stá a Aunejia_*; *_Te manges ion jopa da cern_*

Sentence with initial adverbial constructions (and enclitic position of the clitic): *_Inier sunse stá a Aunejia_*; *_Canche t'es püre mangeste ion jopa da cern_*.
The enclitic form does not fully correspond to the proclitic form for every person.

100% correspondence between enclitic and proclitic forms: 1st person singular (*_i_* → *_i_*): in most cases, the verb ending and the clitic overlap (*_i vá, i jará_* → *spo vái, spo jarái_*; *_i mangi, i mangiará_* → *spo mangi, spo mangiarái_*); 2nd person singular (*_te_* → *_te_*): *_te vas, te jaras_* → *spo vaste, spo jaraste_*; *_te manges, te mangiaras_* → *spo mangeste, spo mangiaraste_*.

Partial correspondence between enclitic and proclitic forms: 3rd person singular (*_al_* vs. *_l/el_*; *_ara_* vs. *_ra/era_*): *_al/ara vá, al/ara jará_* → *spo vái / vára, spo jarál / jarára_*; *_al/ara mangia, al/ara mangiará_* → *spo mangel / mangera, spo mangiarál / mangiarára_*; 3rd person plural (*_ai_* vs. *_i_*; *_ares_* vs. *_res/-eres_*): *_ai/ares vá, ai/ares jará_* → *spo vái / váres, spo jarái / jaráres_*; *_ai/ares mangia, ai/ares mangiará_* → *spo mangi / mergeres, spo mangiarái / mangiaráres_*.

No correspondence between enclitic and proclitic forms: 1st person plural (*_i_* → *_es_*): *_i jun, i jarun_* → *spo junse, spo jarunse_*; *_i mangiun, i mangiarun_* → *spo mangiunse, spo mangiarunse_*; 2nd person plural (*_i_* → *_e_*): *_i jëis, i jarëis_* → *spo jëise, spo jarëise_*; *_i mangëis, i mangiarëis_* → *spo mangëise, spo mangiarëise_*.

In Ladin, the prepositions *_a, da, de, te, sön_* bind directly to the definite article.

a (to/at) + feminine singular *_la/l'_* → *_ala/al'_* (*_ala_* lá before consonant, *_al'_* incuntada before vowel); + masculine singular *_le/l'_* → *_al_* (*_al_* mutur, al pröm); + feminine plural *_les_* → *_ales_* (*_ales_* atres, ales maestres); + masculine plural *_li/i_* → *_al/ai_* (*_ai_* ostis, ai membri).

da (from) + feminine singular *_la/l'_* → *_dala/dal'_* (*_dala_* scora, dal'atra pert before vowel); + masculine singular *_le/l'_* → *_dal_* (*_dal_* ombolt, dal mëteman); + feminine plural *_les_* → *_dales_* (*_dales_* idees, dales cater); + masculine plural *_le/i_* → *_dai_* (*_dai_* interesc, dai comuns).

de (of) + feminine singular *_la/l'_* → *_dla/dl'_* (*_dla_* comunità, dl'umanité before vowel); + masculine singular *_le/l'_* → *_di/dl_* (*_di_* auto, dl laur); + feminine plural *_les_* → *_dles_* (*_dles_* operadesses, dles regoles); + masculine plural *_li/i_* → *_di_* (*_di_* afars, di seturs).

te (in) + feminine singular _la/l'_ → _tla/tl'_ (_tla banca, tl'agricoltora_ before vowel); + masculine singular _le/l'_ → _tl_ (_tl interes, tl paise_); + feminine plural _les_ → _tles_ (_tles intenziuns, tles cités_); + masculine plural _li/i_ → _ti_ (_ti agn, ti mëise_).

sön (on) + feminine singular _la/l'_ → _söla/söl'_ (_söla tofla, söl'incrujada_ before vowel); + masculine singular _le/l'_ → _söl_ (_söl termo, söl armé_); + feminine plural _les_ → _söles_ (_söles isoles, söles massëdles_); + masculine plural _li/i_ → _söi_ (_söi odli, söi libri_).

Before masculine nouns starting with a vowel, the article combined with the preposition is not apostrophized: _tl_ auto, _al_ ater, _dal_ idraulich.

These same prepositions do not bind to the indefinite article; they remain completely separate: te n auto, a n'atra, da n idraulich.

In the Ladin of Val Badia, there is no distinct conditional form used in the apodosis (the main clause of a conditional sentence expressing the result). Instead, the subjunctive mood is employed in both clauses of the conditional sentence.

Italian (conditional in apodosis): _Se Paolo avesse superato l'esame della patente, sua nonna gli avrebbe regalato una macchina._

English translation: _If Paolo had passed his driving test, his grandmother would have given him a car._

Ladin of Val Badia (subjunctive in both clauses): _Sce Paul ess passé l'ejam dla patent, ti ess süa lâ scinché n auto._

In the Ladin of Val Badia, the imperfect tense is not formed using an intervocalic _-v-_. Instead, it is formed by lengthening the vowel in the verb ending, which is represented graphically by a circumflex accent.

The final vowel of the infinitive ending matches the vowel of the imperfect tense only and exclusively in the category of verbs ending in "-i" (where _-i_ = -î). In all other verb conjugations, the final vowels change:

Verbs ending in _-é_ (changes to _-â_): present _te manges_ → imperfect _te mangiâs_.

Verbs ending in _-e_ (changes to _-ô_): present _i perdi_ → imperfect _i pordô_.

Verbs ending in _-ëi_ (changes to _-ô_): present _al plej_ → imperfect _al playô_.

Verbs ending in _-i_ (retains matching vowel _-î_): present _ara dorm_ → imperfect _ara dormî_.

You are a professional {target_lang} translator, tasked with providing translations suitable for use in {target_lang}. Your goal is to accurately convey the meaning and nuances of the original text while adhering to {target_lang} grammar and vocabulary and ensuring that the translation is natural. The original text is a news article. Ensure the translation is formal and consistent with journalistic standards. Produce only the {target_lang} translation, without any additional explanations or commentary. Maintain the HTML formatting of the original source text. Please translate the following text into {target_lang}:

C.2 Social Media Domain

For user-generated social media content, the instructions specified an informal, conversational style while handling platform-specific elements such as hashtags, URLs, and expressive punctuation.

You are a professional {target_lang} translator, tasked with providing translations suitable for use in {target_lang}. Your goal is to accurately convey the meaning and nuances of the original text while adhering to {target_lang} grammar and vocabulary and ensuring that the translation is natural. The original text is user-generated content from a social media platform. Do not reproduce spelling mistakes. Reproduce marks of expressiveness that communicate meaningful intent (e.g. enthusiasm through capitalisation or elongation) in a way that is natural in {target_lang}. Copy URLs and user handles directly rather than translating them. However, translate hashtags as appropriate for the translation to be natural for social media text. Follow the punctuation of the source text as best as possible. Additional punctuation should be added only if not doing so would seriously alter the comprehension of the text. Translate text in an informal style, like close friends talking, even if it changes the original tone. Produce only the {target_lang} translation, without any additional explanations or commentary. Maintain the HTML formatting of the original source text. Please translate the following text into {target_lang}:

C Domain-Specific Instructions

To guide the translation process according to the specific characteristics and constraints of each domain, we applied targeted domain instructions during generation. In the prompts below, the placeholder {target_lang} is dynamically instantiated according to the target language and variety descriptor—specifically, *Ladin, Val Badia, Italy* (lld_Latn) for Ladin or *Ligurian, Italy* (lij_Latn) for Ligurian.

C.1 News Domain

For journalistic content, the system prompt enforced formal register and preservation of source HTML formatting.

C.3 Spoken Language Domain

For automatically transcribed speech, the instructions prioritized maintaining colloquial flow and handling transcription artifacts like interrupted words and non-linguistic sounds.

You are a professional {target_lang} translator, tasked with providing translations suitable for use in {target_lang}. Your goal is to accurately convey the meaning and nuances of the original text while adhering to {target_lang} grammar and vocabulary and ensuring that the translation is natural. The original text is automatically transcribed from spoken language and can contain errors. Maintain the flow and colloquial style of the speaker in the translation. Do not include non-linguistic sounds (e.g. laughter, groans, hesitation sounds, etc.), but do include interjections. If a word is interrupted, either guess the full word if possible or otherwise omit it. Keep foreign words as they are when translating. Produce only the {target_lang} translation, without any additional explanations or commentary. Output the text such that each sentence is on a separate line. Please translate the following text into {target_lang}:

C.4 Software/JSON Domain

For localization of software interface strings, the prompt mandated strict preservation of JSON structure, object keys, placeholders, and escape sequences.

```
Translate the JSON string values into {target_lang}.
Preserve the exact JSON fragment syntax, structure,
object keys, array order, item counts, placeholders,
escape sequences, numbers, booleans, and nulls from the
input. Do not add, remove, or repair braces, brackets,
quotes, commas, or other delimiters. Translate only
natural-language string values; copy keys, placeholders,
product names, code identifiers, URLs, and non-linguistic
tokens unless translation is clearly required inside
user-facing text. Output only the translated JSON text,
with no commentary or Markdown code fences.
```

D Dictionary Evidence Examples

The following outputs demonstrate the rendered Markdown evidence returned by the dictionary lookup tool, for the lookup requests:

```
{
  "dictionary_lookups": [
    {
      "lemma": "volere", "tenses": ["IndPres", "Inf"]},
    {
      "lemma": "mangiare", "tenses": ["Inf"]}
  ]
}
```

Dictionary evidence for Ligurian (Genoese)

****volere** **v_tr**** (avere l'intenzione) _voei_: volevo fargli una sorpresa _voeiva fâghe unna sorpreisa_; mi scusi, non volevo disturbarla _scià me scuse, no voeiva desturbâla_. ◊ (desiderare) _voei_, _à qcn.) èse cao_: voglio sapere cosa vi ha detto la vicina _veuggio savei còs'a va dito a vexiña_; non sa nemmeno lui ciò che vuole _o no ò sa manco lè quello ch'o veu_; vuoi che ti accompagni dal medico? _ti veu che t'accompagne da-o mego?_ vorrei diventasse un po' più indipendente _me saieiva cao ch'a vegnisse un pò ciù indipendente_. ◊ (con valore modale) _voei_: quando gliel'ho detto, non ci voleva credere _quande ghe l'ò dito, o no voeiva creddighe_; voglio proprio vedere come se la caverà questa volta _veuggio pròpio vedde comm'o ne sciortia stavòtta_; vuoi farmi il favore di smetterla? _ti veu fâme o piaxe de ciantâla?_ volete abbassare un po' il volume della musica, che è tardi? _voei ammermâ un pò o volumme da muxica, che l'é tardi?_. ◊ (esigere) _voei_: vogliono essere rispettati _veuan vegni respettâ_; voglio una spiegazione, e la voglio adesso! _veuggio unna spiegaçion, e â veuggio oua!_. ◊ (concedere) _lasciâ_: non posso uscire stasera, i miei non vogliono _staseia no pòsso sciorti, i mæ no me lascian_; se mia moglie vuole, domani vengo a pescare con voi _se mæ moggê a me lascia, doman vègno â pescâ con viatri_. ****Nota**** Con i verbi modali _dovei_ "dovere", _poei_ "potere" e _voei_ "volere" seguiti da un infinito, in genovese si usa solo l'ausiliare _avei_, mentre in italiano si ricorre ad _avere_ o a _essere_ a seconda dell'ausiliare richiesto dal verbo all'infinito: _ò dovuo corri_ "sono dovuto correre", _no ò posciuo anâghe_ "non sono potuto andarci", _an vosciuo parti_ "sono voluti partire". Fanno eccezione le forme riflessive intransitive, dove l'uso dell'ausiliare è parallelo a quello dell'italiano: _a se n'é dovua anâ_ "se n'è dovuta andare", _an vosciuo anâsene_ "sono voluti andarsene". ♦ ****Multi-word expressions**** ****chi troppo vuole nulla stringe****, ****non c'è peggior sordo di chi non vuole sentire****, ****non volerne sapere più nulla****, ****volente o nolente****, ****volere bene****, ****volere dire**** ♦ Conjugations: ****èse****: ind. pres.: son, t'è, o/a l'é, semmo, sei, en/son / ****lasciâ****: ind. pres.: lascio, ti lasci, o/a lascia, lascemmo, lascia, lascian / ****voei****: ind. pres.: veuggio, ti veu, o/a veu, voemmo/voemo, voei, veuan

****volere** **s_m**** (volontà, desiderio) _voentâ_ (fem), _voei_: questo è il volere del popolo _sta chi a l'é a voentâ do pòpolo_; il testamento rispetta il volere della nonna _o testamento o rispetta a voentâ da madonnâ_. ♦ Declensions: voei (m. s.), voei (m. p.)

****mangiare** **v_tr**** (ingerire del cibo) _mangiâ_, _pittâ_, _ciccâ_: non vedo l'ora di mangiare quella pizza! _no me pâ mai de tardi de mangiâ quella pissa!_ Stefano ha mangiato un panino al volo prima della riunione _o Steva o l'à mangiou un panetto a-a spedia prima da reunion_. ◊ (fare un pasto) _mangiâ_: di solito mangiamo verso le otto di sera _pe-o sòlito mangemmo pe eutt'oe da seia_; vieni a mangiare da noi domenica sera? _ti vègni â mangiâ da niatri domenega de poidisnâ?_. ◊ (guadagnare illecitamente) _giasciâ_, _mangiâ adòsso_ (à qcs.): hanno mangiato tutti i fondi destinati alla ricostruzione _an giasciou tutti i fondi destinâ a-a reconstruçion_. ****Nota**** L'espressione _mette i pè sott'a-a tōa_, letteralmente "mettere i piedi sotto al tavolo", ha il significato di "mangiare serviti". Viene usata, spesso con tono ironico, da chi è abitualmente incaricato di cucinare o servire per sottolineare che, in una data occasione, ha potuto semplicemente sedersi a tavola e mangiare senza dover fare nulla. ♦ ****Nota**** Si segnala anche _mangiâ into papè come i gatti_, letteralmente "mangiare nella carta come i gatti", che descrive un modo informale e scomposto di mangiare, senza piatto né posate; e _mangiâ de strangoggion_ "mangiare strangolandosi" o _mangiâ co-o tortaiue_ "mangiare con l'imbutto", che indicano il mangiare in fretta, ingoiando quasi senza aver masticato. ♦ ****s_m**** (ciò che si mangia) _mangiâ_, _çibbo_, _pitto_ (colloquiale): qui il mangiare è ottimo e abbondante _chi o mangiâ o l'é milia e in abundansa_. ♦ ****Multi-word expressions**** ****da mangiare****, ****la bellezza non si mangia****, ****mangiare a quattro palmenti****, ****mangiare a secco****, ****mangiare fino a scoppiare****, ****mangiare il pane a tradimento****, ****mangiare la foglia****, ****mangiarsi il fegato**** ♦ Declensions: çibbo (m. s.), çibbi (f. s.) / pitto (m. s.), pittì (m. p.)

Dictionary evidence for Ladin (Val Badia)

****volere** **s_m**** (varie traduzioni possibili) _deider_, _orëi_, _orenté_: rispettare il volere di qcn. _respeté la orenté de zacai_. Declensions: d/e/j/i/d/e/r (m. s.), d/e/j/i/d/e/r/s (m. p.) / o/r/ë/i (m. s.) / o/r/e/n/t/é (m. s.)

****volere** **v_tr**** (tradotto come 'orëi') _orëi_: «Andiamo al cinema?» «Come vuoi.» «Junse al chino?» «Sciöche t'os.»; che cosa vuoi che ti dica? _ci oste pa ch'i te dijés?_ cosa vuoi ancora da me? _ci oste pa ciamó da me?_ Dio ha voluto così _Dì â orù insciö_; il mio mal di denti non vuole passare _mi me de dënz n'ó nia passé_. Declensions: o/r/ü/d/a (m. s.), o/r/ü/d/e/s (m. p.)

****mangiare** **s_m**** (varie traduzioni possibili) _mangé_, _mangiaria_, _pasté_, _patjù_ da mangé_: il mangiare è pronto _le mangé é arjigné_; il mangiare era buono _le mangé é bun_; un mangiare leggero _n mangé lisier_; un mangiare pesante _n mangé pesocho_; un mangiare sano _n mangé sann_. Declensions: m/a/n/g/é (m. s.) / m/a/n/g/i/a/r/i/a (m. s.), m/a/n/g/i/a/r/i/e/s (m. p.) / p/a/s/t (m. s.), p/a/s/ć (m. p.)

****mangiare** **v_tr**** (varie traduzioni possibili) _ciablé_, _gosté_, _mangé_ sò, _se mangé ite_, _cacé jö_: dar da mangiare agli affamati _ti dé da mangé a chi ch'â fan_; dar da mangiare agli affamati _ti dé da mangé ai afamâ_; dar da mangiare ai poveri _ti dé da mangé ai pürî_; la ruggine mangia il ferro _la rüja mangia sò le fer_; mangia e bevi _mangia y bër_. Declensions: c/i/a/b/l/é (m. s.) / g/o/s/t/é (m. s.) / m/a/n/g/i/a/d/a (m. s.), m/a/n/g/i/a/d/e/s (m. p.)