

# A Savoyard Version of the FLORES+ Corpus

Malik Marmonier<sup>†</sup> Alain Favre<sup>‡</sup> Benoît Sagot<sup>†</sup> Rachel Bawden<sup>†</sup>

<sup>†</sup>Inria, Paris, France <sup>‡</sup>Institut de la langue savoyarde

{firstname.lastname}@inria.fr<sup>†</sup> {firstname.lastname}@langue-savoyarde.com<sup>‡</sup>

## Abstract

We present a Savoyard version of the FLORES+ corpus, a contribution to the Open Language Data Initiative (OLDI) shared task at WMT 2026. Savoyard, an endangered Arpitan dialect of the Gallo-Romance language family, is currently undergoing active revitalization. We give an overview of its linguistic history, describe our data collection process, and report baseline scores for a range of systems on the resulting French↔Savoyard evaluation set as a point of comparison for future work. We hope that this dataset will stimulate further data collection efforts for Savoyard and the broader Arpitan dialect continuum, supporting the development of NLP resources for these endangered languages and contributing to the preservation of humanity’s linguistic heritage.

## 1 Introduction

Machine Translation (MT) has advanced considerably over the past four decades (Brown et al., 1990; Koehn et al., 2003; Sutskever et al., 2014; Bahdanau et al., 2015; Vaswani et al., 2017). This progress owes much to the field’s departure from evaluation practices based on nebulous, self-reported notions of linguistic coverage in favor of shared, empirical evaluation frameworks (Liberman, 2010). Yet these advances have primarily benefited a handful of dominant languages, leaving lower-resourced varieties largely excluded (Joshi et al., 2020; Haddow et al., 2022). Addressing this disparity requires extending the same empirical evaluation rigor that drove earlier advances to low-resource settings.

It is explicitly to this end that the FLORES+ dataset was initially compiled (Goyal et al., 2022, as Flores-101) and continuously expanded (NLLB Team, 2024; Gordeev et al., 2024; Kuzhuget et al., 2024; Ali et al., 2024; Abdulmumin et al., 2024; Perez-Ortiz et al., 2024; Mæhlum et al., 2025; Okttem et al., 2025; Frontull et al., 2025; Lou et al.,




Figure 1: The Arpitan dialect continuum, with the Savoyard domain in a darker green. Outer boundaries are schematic renderings based on (Kristol, 2016) and French Wikipedia maps; political borders and hydrography are drawn from Natural Earth.

2025; Samil and Adelani, 2026, *inter alia*), establishing a standardized evaluation framework for more than 200 languages and script variants, with a deliberate focus on otherwise underrepresented languages.

As a direct extension of these past efforts, our work presents a version of the FLORES+ corpus in the Savoyard language, an endangered Arpitan dialect of the Gallo-Romance family currently the object of active revitalization efforts. In the following sections, we provide an overview of the linguistic history of Savoyard, present our data collection process, and report various baseline performance scores on the resulting evaluation set.

We hope that this dataset will stimulate further data collection efforts for Savoyard and the broader Arpitan dialect continuum, ultimately supporting the development of NLP resources to aid in the preservation of these endangered languages.



Figure 2: Lake Annecy, in the historical heart of the Savoyard domain. Photo by Alex Brown, CC BY 2.0.

## 2 Language Overview

### 2.1 What’s in a name: Arpitan, Franco-Provençal, or Francoprovençal?

Latin, implanted across Gaul in the wake of the Roman conquest, gave rise to a family of Romance vernaculars spanning modern France and Belgium, together with parts of Switzerland and northern Italy. These “Gallo-Romance” languages were long divided into two branches along a north–south axis, named for their medieval reflexes of Latin *hoc* (‘yes’): the *langues d’oïl* of the north (Smith, 2016), chief among them French, and the *langues d’oc*, or Occitan, of the south (Olivieri and Sauzet, 2016)—echoing a classification of the Romance vernaculars famously articulated by Dante in the early 14th century (Alighieri, 1996). A third grouping went unrecognized until Ascoli (1878) demonstrated, on phonetic grounds, that the dialects of east-central France, French-speaking Switzerland, and the Aosta Valley of Italy form a coherent group, neither varieties of an Oc or Oïl language nor a mixture of both, but constituting a distinct linguistic development characterized by its own treatment of the Latin substratum.

Ascoli baptized this new grouping *franco-provençale* (‘Franco-Provençal’), a hyphenated name that—in yoking together the labels of an Oïl and an Oc language—seemed to some to announce precisely the hybrid status his analysis had refuted, and the label has therefore been faulted ever since (Tuaillon, 1983). In an attempt to mitigate this perceived misnomer, later scholars suppressed the hyphen, *Francoprovençal*, recasting the two elements as a unified linguistic sign (Tuaillon, 1983)—a spelling now standard in the scientific literature (Kristol, 2016). Nonetheless, this solution was deemed unsatisfactory by language activists who promoted the neo-endonym “Arpitan”

as a replacement (from a local root *arp-* ‘mountain pasture’).<sup>1</sup> Once contentious for its separatist and ethno-nationalist overtones in the Italian context that birthed it (Bichurina, 2019), the term “Arpitan” has since been largely depoliticized and is seeing growing institutional adoption. It has now supplanted *Franco-Provençal* and *Francoprovençal* as the reference name of the language under SIL.<sup>2</sup> It is also the term preferred by the *Institut de la langue savoyarde*—whose translation services produced the FLORES+ partition this paper presents—to refer to the broader language grouping to which Savoyard belongs.

### 2.2 Savoyard

Savoyard<sup>3</sup> is a subset of the Arpitan dialect continuum. It denotes the Arpitan varieties historically spoken in the County of Savoy—a feudal state of the Holy Roman Empire founded in 1003 and raised to a duchy in 1416, which straddled the modern French–Italian–Swiss alpine borders, with its Arpitan-speaking core comprising the present-day French *départements* of *Savoie* and *Haute-Savoie*, and most of the Swiss Canton of Geneva.<sup>4</sup>

As Latin receded from administrative and literary use in the late European Middle Ages, Romance vernaculars—such as French or the various dialects of the Italian city-states—moved into the registers it vacated. Arpitan did so outside Savoy, to a limited extent.<sup>5</sup> Savoyard, however, was deprived

<sup>1</sup>As is often the case with stigmatized vernaculars in French-speaking diglossic contexts, the actual endonym is the pejorative term “*patois/patouè*”.

<sup>2</sup><https://iso639-3.sil.org/code/frp>; we note that ISO 639-3 code is still frp.

<sup>3</sup>Glottocode *savo1253*. We follow Glottolog in using “Savoyard,” though we note that the Institut de la langue savoyarde prefers “Savoyan” as the English name of the language.

<sup>4</sup>Excluding the city of Geneva itself.

<sup>5</sup>For instance, in the Lyonnais, in the devotional writings of

of such an opportunity due to the Duchy's early adoption of French as the language of its court and chancery.<sup>6</sup> It was therefore left without the administrative or literary functions through which other vernaculars acquired prestige.

The prestige of Savoyard nonetheless improved markedly in the 19th century, when an abundant literature written in Savoyard began to emerge, prompted by Romantic interest in folklore and popular tradition (Bichurina and Dunoyer, 2021), with poet Amélie Gex (1835–1883) as its most accomplished representative. This development, however, was cut short by the annexation of Savoy to France in 1860. Aggressive French nation-building policies had long identified linguistic diversity with anti-republican feudalism (Grégoire, 1794), and the vernaculars—now uniformly denigrated as *patois*—were driven out of public life across France. In Savoy, this process was particularly virulent, compressed as it was into a shorter time span due to the region's more recent inclusion into the French polity, and the school system was its principal instrument: children were forbidden to speak Savoyard, humiliated and punished when they did, translating political assimilation into practices of *de facto* linguistic violence (Skutnabb-Kangas, 2000).

Revitalization efforts began in earnest in the second half of the 20th century, as the vocabulary of decolonization was taken up by minority-language communities in the First World itself, and gathered pace around the turn of the millennium, when the framing of language loss in terms borrowed from environmental conservation entered French public debate and inflected public policy (Bichurina, 2019).

A late-1980s estimate placed the number of Savoyard speakers at approximately 30,000 (Tuailon, 1989), though reliable counts remain difficult to establish (Boutillier, 2022). Indeed, recent sociolinguistic work—the FORA survey (Bert and Costa, 2009) and subsequent fieldwork by Bichurina and Dunoyer (2021)—reveals a complex linguistic landscape characterized by four main categories of speakers. The “native” or traditional speakers, generally aged between 80 and 90 years old, acquired the language in early childhood but

their use of Savoyard often remains invisible in the public sphere due to past stigmatization. In contrast, “late” speakers—predominantly educated retirees in their 60s and 70s who had limited exposure to Savoyard in their childhood but re-engaged with the language in adulthood—are highly visible and drive preservation efforts through dedicated associations. Alongside these, a small but growing group of “new” speakers, primarily teenagers and young adults, approach Savoyard as a learned language, driven by broader cultural motives. Finally, the landscape includes a vast pool of “passive speakers” who comprehend the language but struggle to express themselves in it. Because these individuals can rapidly transition into active speakers within supportive environments, their fluctuating proficiency complicates attempts to establish static speaker counts.

While a comprehensive linguistic description of Savoyard falls outside the scope of this paper,<sup>7</sup> its defining characteristic—marked dialectal fragmentation—is highly relevant to corpus building. As an Arpitan variety, Savoyard exhibits significant micro-linguistic diversity, exacerbated over centuries by the geographic isolation of alpine valleys. This diversity extends beyond lexical variations to encompass morphological and phonetic divergences, where the evolution of shared Latin roots often differs from one village to the next. In a computational context, this significant local variance presents a unique modeling challenge, underscoring the necessity of a supra-dialectal orthography capable of unifying these diverse variants.

### 2.3 Revitalization and Written Standards

Savoyard revitalization efforts rest mainly on networks of private non-profit associations: the local *groupes patoisants* and, since 2002, the *Institut de la langue savoyarde*.<sup>8</sup> A persistent barrier to the revitalization of Savoyard—and Arpitan more broadly—through French public education stemmed from the absence of a unified orthography. Despite a softening stance toward regional languages, the French state's ideological framework still equates linguistic legitimacy with a codified written standard (Bichurina and Dunoyer, 2021). This text-centric ideology historically underpinned the *loi Deixonne* (1951)—the foundational law for the teaching of regional languages in

the mystic Marguerite d'Oingt (d. 1310), and in the Dauphiné, where Justinian's *Institutes* were rendered into the dialect of Grenoble (Bichurina and Dunoyer, 2021).

<sup>6</sup>An early adoption that predated even that of the French Kingdom. The neighboring Republic of Geneva replaced Latin with French on a similar timeline, and to the equal detriment of the local Arpitan vernacular.

<sup>7</sup>See (Désormaux, 1981) and (Viret, 2015) for detailed accounts and additional references.

<sup>8</sup>In Savoyard: *Enstitut de la Lengoua Savoyârda*.

France—which systematically excluded uncoded vernaculars.

Defining a written standard for Arpitan thus became a strategic imperative for this linguistic community, albeit a deeply divisive one. For decades, the dominant written form in Savoy was the *Graphie de Conflans* (Tuailon, 1984). Developed in the 1980s by linguists and native speakers, this highly phonetic system was designed to faithfully transcribe the exact pronunciation of local village dialects. While it preserved the region’s immense micro-linguistic diversity, its hyper-local fidelity severely hindered broader intercomprehension and official recognition. To bridge this gap, linguist Dominique Stich introduced the *Orthographe de Référence B* (ORB) in the early 2000s (Stich, 2001). As a supra-dialectal, etymological system, ORB aims to abstract local phonetic and morphological variations into a single graphemic standard for the entire Arpitan domain.

While ORB initially faced fierce resistance from traditional speakers—who found it harder to master and feared it erased their specific local heritage—it was championed by “new” speakers and revitalization activists for its unifying character and political leverage (Bichurina, 2019; Dunoyer, 2019). Ultimately, the strategic necessity of a unified standard appears to have prevailed, and it is in great part thanks to the growing acceptance of the principle of a supra-dialectal written standard that the French Ministry of Education officially authorized the teaching of Arpitan in French public schools in 2021 (Ministère de l’Éducation nationale, de la Jeunesse et des Sports, 2021), with the active collaboration of the *Institut de la langue savoyarde*.

### 3 Data Collection

This section outlines the data collection process for the Savoyard partitions of the FLORES+ dev and devtest sets, which were translated from French by the translation service of the *Institut de la langue savoyarde* (ILS).

#### 3.1 Target Linguistic Variant and Orthography

The Savoyard variant selected as the target language for the partition is termed *Savoyard médian* (Median Savoyard) by the ILS. It is the emergent standard that the ILS teaches through MOOCs and online seminars. It is closest to the dialect of Thônes, a town located near the geographic center

of Savoy, and therefore broadly understandable by the greatest number of speakers, though its emergence as a written standard is mostly tied to lexicographical efforts by Favre and Stich (2005), recently expanded by Favre (2026). The orthographic norm selected for the partition is the aforementioned supra-dialectal ORB spelling as documented in (Stich, 2003).

#### 3.2 Translation Process

The translation workload, initially shared between two translators, was ultimately borne almost entirely by Alain Favre, a proficient late speaker and lexicographer who contributed in the past to the development of a rule-based Apertium system for the French-Arpitan language pair (Forcada and Tyers, 2016; Alòs, 2020).

Translators were instructed to follow the official OLDI guidelines,<sup>9</sup> which were presented to them during an online meeting. They worked within a customized spreadsheet interface (see Figure 3), with one sheet for each of the dev and devtest partitions. A source text column displayed the French segments of FLORES+, originally extracted from English Wikinews, Wikijunior, and Wikivoyage documents and translated into French by the original compilers of the corpus (Goyal et al., 2022). A column displaying the URL of the source documents allowed the translators to refer to them, but since these documents are in English, a column displaying the Italian translation of the segments proved more useful for resolving ambiguities in the French text. While the translators were instructed to translate the devtest partition strictly from scratch, for the dev partition they were provided with translation hypotheses generated by the rule-based French-Arpitan Apertium system (Forcada and Tyers, 2016) for inspiration (column not shown in Figure 3). This design choice was motivated by the desire to collect in-context-learning or finetuning data for the automatic post-editing of Apertium output into Savoyard.<sup>10</sup> Such data would in turn facilitate large-scale data augmentation.

The translators were given three and a half months to complete the assigned task and were compensated at a fair rate. During their work, they faced diverse challenges, both stylistic (the formal “press release” style typical of Wikinews source texts, for instance, does not historically exist in

<sup>9</sup><https://oldi.org/translation-guidelines.pdf>

<sup>10</sup>This design choice was discussed with and approved by the OLDI shared task organizers beforehand.

The image shows the FLORES+ translation interface. At the top, there's a menu bar with options like 'Fichier', 'Édition', 'Affichage', 'Insertion', 'Format', 'Données', 'Outils', 'Géométrie', 'Extensions', and 'Aide'. Below the menu is a toolbar with various icons for text editing and formatting. The main area is a table with columns for source and target text, and a status column. The table contains multiple rows of text, mostly in French and Italian, with some English text interspersed. The interface is designed for managing and translating large volumes of text.

Figure 3: Translation interface.

Savoyard) and terminological. When the translation of FLORES+ surfaced terminological gaps in the existing Savoyard lexicography (Favre and Stich, 2005; Viret, 2019; Favre, 2026), neologisms were coined by prioritizing the intrinsic lexical resources of Arpitan over direct French borrowings.

The second translator—a linguistics student—contributed to the early stages of the project, providing a small batch of translations and reviewing the first several hundred segments before having to unexpectedly withdraw from the project; their involvement nonetheless yielded valuable stylistic and grammatical discussions. Following their departure, the primary translator established a formalized lexical and syntactic review grid, based on various reference texts, to ensure consistency across the entire dataset. All drafted translations were subsequently proofread and standardized against this grid.

Upon receipt of the translations, length-ratio heuristics were used as a first validation step to check for alignment problems. More substantial validation was provided by the baseline process presented in the next section and by the accompanying human annotation campaign; a random sample of the resulting dataset is displayed in Appendix D.

## 4 Baselines

To contextualize the difficulty of the French↔Savoyard translation task, and to provide reference points for future work on this dataset, we report the performance of a range of baseline approaches, spanning naïve copy of the source segment, rule-based systems, vanilla finetuning of neural MT (NMT) systems, sequence-level distillation of rule-based and LLM<sup>11</sup>-post-edited MT output into NMT systems, and LLM-based approaches. All systems are evaluated on the devtest partition of the Savoyard FLORES+ corpus we compiled, in both translation directions (FR→FRP and FRP→FR),<sup>12</sup> using the BLEU (Papineni et al., 2002), chrF++ (Popović, 2015, 2017), and TER (Snover et al., 2006) metrics.<sup>13</sup> Results are given in Table 1.

**Naïve copy** As a lower-bound reference, we report the performance of a trivial copy-source baseline, which simply returns the source seg-

<sup>11</sup>Large Language Model

<sup>12</sup>We rely on the ISO 639-3 FRP code for the Arpitan language as a shorthand for Savoyard in this section.

<sup>13</sup>Metrics were computed using sacreBLEU (Post, 2018). BLEU : nrefs:1lcase:mixedlfff:noltok:13alsmooth:explversion:2.6.0; chrF++ : nrefs:1lcase:mixedlfff:yeslnc:6lnw:2lpspace:nolversion:2.6.0; TER : nrefs:1lcase:lcltok:tercomlnorm:nolpunct:yeslasian:nolversion:2.6.0

| System                | Modality                      | FR → FRP |        |          |       | FRP → FR |        |          |       |
|-----------------------|-------------------------------|----------|--------|----------|-------|----------|--------|----------|-------|
|                       |                               | Extr. %  | BLEU ↑ | chrF++ ↑ | TER ↓ | Extr. %  | BLEU ↑ | chrF++ ↑ | TER ↓ |
| Copy Source Baseline  | -                             | 100.00   | 10.5   | 40.2     | 70.1  | 100.00   | 10.5   | 39.3     | 70.2  |
| Apertium (Rule-Based) | Standard Generation           | 100.00   | 74.7   | 86.0     | 15.3  | 100.00   | 68.8   | 81.4     | 19.8  |
| NLLB-200 (600M)       | Finetuned Baseline (Gold)     | 100.00   | 35.0   | 62.9     | 39.5  | 100.00   | 63.2   | 78.3     | 23.8  |
|                       | Distilled (Gold+Silver)       | 100.00   | 69.6   | 83.3     | 18.1  | 100.00   | 72.4   | 84.2     | 17.1  |
| Gemma-4-31B           | Zero-Shot   Direct (No-Think) | 100.00   | 12.3   | 42.6     | 65.4  | 100.00   | 71.4   | 82.9     | 18.6  |
|                       | Zero-Shot   Reasoning         | 90.32    | 11.7   | 39.6     | 67.5  | 98.72    | 71.9   | 82.8     | 19.0  |
|                       | Few-Shot   Direct (No-Think)  | 100.00   | 58.4   | 77.5     | 24.5  | 100.00   | 77.3   | 86.6     | 14.4  |
|                       | Few-Shot   Reasoning          | 82.61    | 45.7   | 64.8     | 39.2  | 89.92    | 69.6   | 79.7     | 22.7  |
| Gemma-4-E4B           | Zero-Shot   Direct (No-Think) | 100.00   | 9.8    | 39.4     | 70.8  | 100.00   | 61.2   | 76.1     | 26.3  |
|                       | Zero-Shot   Reasoning         | 100.00   | 8.8    | 37.4     | 72.8  | 100.00   | 58.3   | 74.5     | 29.9  |
|                       | Few-Shot   Direct (No-Think)  | 100.00   | 40.2   | 66.0     | 38.1  | 100.00   | 65.5   | 79.0     | 23.2  |
|                       | Few-Shot   Reasoning          | 100.00   | 33.4   | 60.8     | 43.6  | 100.00   | 62.5   | 77.1     | 26.1  |
| Gemma-4-E2B           | Zero-Shot   Direct (No-Think) | 100.00   | 8.9    | 38.5     | 72.7  | 100.00   | 54.9   | 72.2     | 30.8  |
|                       | Zero-Shot   Reasoning         | 100.00   | 9.0    | 37.2     | 74.9  | 100.00   | 33.7   | 67.0     | 63.1  |
|                       | Few-Shot   Direct (No-Think)  | 100.00   | 27.1   | 55.4     | 50.3  | 100.00   | 53.0   | 69.3     | 33.9  |
|                       | Few-Shot   Reasoning          | 100.00   | 4.2    | 38.5     | 161.5 | 100.00   | 29.2   | 64.5     | 68.4  |

Table 1: Evaluation metrics for both translation directions (FR → FRP and FRP → FR).

ment (French in the FR→FRP, Savoyard in the FRP→FR) unmodified as the translation hypothesis under evaluation. Despite the close phylogenetic relationship and ongoing diglossic contact between French and Savoyard, this baseline performs poorly in both directions (BLEU 10.5, chrF++ 40.2/39.3), indicating that the two languages display only marginal overlap in their surface forms and that the FLORES+ test partition meaningfully discriminates between systems with genuine translation ability and mere copying.

**Rule-based system** We next evaluate the French–Arpitan Apertium rule-based system (Forcada and Tyers, 2016; Alòs, 2020),<sup>14</sup> which performs strongly in both directions, and in fact constitutes the best-performing system for FR→FRP (BLEU 74.7, chrF++ 86.0, TER 15.3). While the Arpitan Apertium system was not used by the translators to translate the devtest partition of FLORES+, one of the translators was instrumental in its development, as previously mentioned, and its output is therefore likely to closely mirror this translator’s idiolect. Apertium performance is, however, asymmetric: it is comparatively weaker for FRP→FR (BLEU 68.8, chrF++ 81.4, TER 19.8), a pattern we return to below.

**NMT finetuning** We finetune NLLB-200-600M (NLLB Team, 2024) directly on the dev partition of our Savoyard FLORES+ corpus (with 20% held out for validation), using a learning rate of 2e-5, an effective batch size of 16, weight decay of 0.01, mixed-precision training, and a best-epoch-saving strategy over 10 epochs. Given the limited size of this gold-reference training signal,

performance is considerably lower than that of Apertium, and strikingly asymmetric across directions: FRP→FR reaches BLEU 63.2 (chrF++ 78.3), whereas FR→FRP falls to BLEU 35.0 (chrF++ 62.9). This gap suggests that, under a small-data regime, generation into the lower-resourced Savoyard is substantially harder than generation into French, for which NLLB’s pretrained capacity can offset the scarcity of in-domain supervision.

**NMT distillation-finetuning** Our second NMT configuration, Distilled (Gold+Silver), augments the gold development data with silver parallel data generated through a two-step translation pipeline. First, we translate the French source text from the OLDI Seed (Maillard et al., 2023; NLLB Team, 2024; Marmonier et al., 2025) and SMOLSENT (Caswell et al., 2025)<sup>15</sup> corpora into Arpitan using Apertium. Second, we use Gemma-4-31B to post-edit this Arpitan draft into Savoyard. To guide the LLM, we prompt it with few-shot triplet demonstrations consisting of the French source segments from the FLORES+ dev set, the Apertium Arpitan draft translations presented to the human translators during the annotation process, and the gold Savoyard reference from the compiled Savoyard dev partition (see Appendix A). We then finetune NLLB-200 on the concatenated gold and silver datasets. This approach substantially closes the gap with Apertium and reverses the direction-wise asymmetry observed for rule-based translation, consistent with the pattern already seen for the NLLB Baseline model above: for FRP→FR, the distilled model (BLEU 72.4, chrF++ 84.2) outperforms Apertium, while for FR→FRP it remains

<sup>14</sup><https://github.com/apertium/apertium-fra-frp>

<sup>15</sup>We compiled a French version of SMOLSENT, which we contribute to the same 2026 OLDI Shared Task.

slightly behind (BLEU 69.6 vs. 74.7).

**LLMs** We also evaluate three models from the Gemma-4 family (31B, E4B, and E2B; [Google DeepMind 2026](#)) across two binary axes. The first axis, Zero-Shot vs. Few-Shot, varies the level of in-context supervision. Zero-shot models receive only a brief prompt instructing them to translate the segment based on prior knowledge, while few-shot models are supplied with the entire *dev* partition structured as ordered bitexts within their prompt. The second axis, Direct (No-Think) vs. Reasoning, controls the generation modality. Direct models must emit the translation immediately within XML tags, whereas Reasoning models are permitted an explicit internal monologue before producing the tagged output.<sup>16</sup> See Appendix B for prompt examples.

Two patterns in LLM performance emerge from Table 1.

First, few-shot demonstrations provide substantial and consistent gains across all model sizes and translation directions. For instance, under Direct generation, Gemma-4-31B improves from 12.3 to 58.4 BLEU for FR→FRP, and from 71.4 to 77.3 BLEU for FRP→FR; identical trends hold for E4B and E2B. This indicates that the *dev* partition of FLORES+, despite its modest size, offers a potent in-context-learning signal for a genuinely low-resource variety that is likely absent from pretraining data while closely related to a high-resource language.

Second, while standard model ranking remains constant across modalities (Gemma-4-31B > E4B > E2B), allowing models to reason before translating is universally harmful, with the performance penalty exhibiting a non-monotonic relationship to model size. Performance degrades most sharply at both ends of the model-size spectrum, though through distinct failure mechanisms: the 31B variant falls into infinite self-checking loops that exhaust the token limit (dropping extraction rates to 82.61% for few-shot FR→FRP), whereas the smaller E2B model frequently emits its chain-of-thought directly within the output tags (driving TER scores above 160). The intermediate

<sup>16</sup>We use a regular expression to extract the string within the <translation>...</translation> tags. Extraction failures (reported as *Extr. %* in Table 1) are penalized as empty strings. Inference is powered by vLLM ([Kwon et al., 2023](#)) and Hugging Face transformers ([Wolf et al., 2020](#)). Max output tokens are capped at 1,024 for Direct and 16,384 for Reasoning; vLLM’s max-model-len is set to 240,000 for the few-shot settings.

E4B model incurs a milder penalty, only suffering from reasoning-induced linguistic drift, occasionally yielding Italian-inflected or French-lexified pseudo-translations (e.g., “*L’armonia al luo de travai...*”) adequately penalized by the metrics. Ultimately, no model, at any scale, appears to successfully leverage an explicit internal monologue productively for this translation task.<sup>17</sup>

## 5 Human Validation

To validate the automatic metrics and incidentally the quality of the human translations, we conducted a blind pairwise human evaluation campaign. The annotation was performed by Alain Favre, who also served as the primary translator for the evaluation set. To preserve statistical power while remaining mindful of the annotator’s time, we restricted the experiment to a subset of the most promising systems: the rule-based Apertium platform, the distilled NLLB-200 (600M) model, and Gemma-4-31B (Few-Shot | Direct). The annotator was presented with 100 round-robin matchups per translation direction and asked to indicate his preference between two anonymized system outputs, with the option of declaring a tie, and with the gold human reference injected as a hidden system. The resulting pairwise judgments were aggregated into a global ranking using the TrueSkill algorithm ([Herbrich et al., 2007](#)), with default priors ( $\mu = 25.0$ ,  $\sigma = 8.333$ ) and ties resolved via TrueSkill’s native draw mechanism. For each system, we report in Table 2 a conservative score ( $\mu - 3\sigma$ ) and, for adjacent pairs in the ranking, the probability that the higher-ranked system’s true skill exceeds the lower-ranked system’s, computed in closed form from the difference of their posterior Gaussians.

These results show that human judgments closely mirror the trends observed in the automatic metrics, highlighting a stark asymmetry depending on the translation direction. In the FR → FRP di-

<sup>17</sup>To test whether these performance differences are statistically significant, we performed paired bootstrap resampling ([Koehn, 2004](#)). The Direct-over-Reasoning advantage is significant ( $p < 0.05$ ) in ten of the twelve model/regime/direction comparisons. Two comparisons deviate: zero-shot Gemma-4-31B in FRP→FR, where the two modalities are statistically indistinguishable (+0.5 chrF++, n.s.), and few-shot Gemma-4-E2B in FRP→FR, where Reasoning significantly outperforms Direct (−2.0 chrF++,  $p < 0.05$ ). Full pairwise significance matrices are given in Appendix C. We note that the significance matrices report *segment-level* mean chrF++ differences, whereas Table 1 reports corpus-level scores; the two diverge for the degenerate E2B Reasoning configurations, whose spurious chain-of-thought output is penalized far more heavily under corpus-level aggregation.

| Direction | System                    | Cons. Score | $\mu$ | $\sigma$ | Prob. > Next | Significance           |
|-----------|---------------------------|-------------|-------|----------|--------------|------------------------|
| FR → FRP  | Hidden Human Reference    | 27.32       | 31.41 | 1.36     | >99.9%       | $p < 0.001$            |
|           | Apertium (Rule-Based)     | 20.53       | 23.94 | 1.14     | 67.3%        | n.s.                   |
|           | NLLB-200 Distilled (600M) | 19.79       | 23.22 | 1.14     | 81.3%        | n.s.                   |
|           | Gemma-4-31B (Few-Shot)    | 18.29       | 21.77 | 1.16     | —            | —                      |
| FRP → FR  | Hidden Human Reference    | 25.26       | 28.92 | 1.22     | 88.1%        | n.s. ( $p \geq 0.05$ ) |
|           | Gemma-4-31B (Few-Shot)    | 23.51       | 26.94 | 1.14     | 97.5%        | $p < 0.05$             |
|           | NLLB-200 Distilled (600M) | 20.30       | 23.75 | 1.15     | 81.4%        | n.s.                   |
|           | Apertium (Rule-Based)     | 18.67       | 22.27 | 1.20     | —            | —                      |

Table 2: TrueSkill human evaluation leaderboards for both translation directions. The Conservative Score represents the lower 99% credible bound of the system’s estimated true skill ( $\mu - 3\sigma$ ). The “Prob. > Next” column indicates the probability of the system’s true skill being strictly greater than that of the system ranked immediately below it, alongside its corresponding statistical significance threshold (n.s. = not significant).

rection, the hidden human reference ( $\mu = 31.41$ ) significantly outperformed all machine systems ( $p < 0.001$ ). Among the automated approaches, the hand-crafted rules of Apertium ( $\mu = 23.94$ ) and the Distilled NLLB model ( $\mu = 23.22$ ) performed comparably, effectively tying for first place. Notably, the Distilled NLLB model outperformed Gemma-4-31B ( $\mu = 21.77$ ). Conversely, in the FRP → FR direction, Gemma-4-31B Few-Shot Direct ( $\mu = 26.94$ ) demonstrated overwhelming superiority over both NLLB and Apertium. Remarkably, the probability of superiority between the hidden human reference and Gemma-4-31B was only 88.14%, indicating that the LLM achieved near-human parity ( $p \geq 0.05$ ) when decoding into standard French.

The automatic and human evaluations therefore converge, validating both the automatic metrics for this language pair and the quality of the reference translation, to an extent that is further discussed in the Limitations section. The resulting picture is asymmetric: translation into French is approaching human parity, while into Savoyard the human reference remains significantly preferred over all systems and the best baseline system is rule-based.

## 6 Related Work

The first MT evaluation dataset bearing the name “FLORES” was released in 2019 and targeted two resource-constrained Indo-Aryan languages: Nepali and Sinhala (Guzmán et al., 2019). Its source segments were extracted from Wikipedia articles, and professionally translated, with stringent iterative rounds of quality assurance. In order to mitigate the effects of translationese on MT evaluation (Zhang and Toral, 2019; Graham et al., 2020), source segments were also sampled from Nepali and Sinhala Wikipedia, rather than English alone.

The subsequent FLORES-101 dataset (Goyal et al., 2022) forms the actual basis of FLORES+. While it retained the core methodology of professional translation and rigorous quality assurance protocols, its 2,009 source segments bear no relation to those of the original FLORES, having been drawn instead from English-language Wikimedia projects (Wikitravel, Wikinews, Wikijunior) and translated into 101 target languages. This English-centric design made possible a fully aligned many-to-many evaluation across 10,100 language pairs, allowing researchers to assess regionally relevant, non-English translation directions.

FLORES-101 was further expanded into FLORES-200 in the context of the NLLB-200 project (NLLB Team, 2024), covering 200 languages as the name implies. The corpus, now curated by the Open Language Data Initiative (Burchell et al., 2024; Dale et al., 2025) under the name FLORES+, is regularly extended into new low-resource languages (Gordeev et al., 2024; Kuzhuget et al., 2024; Ali et al., 2024; Perez-Ortiz et al., 2024; Frontull et al., 2025; Lou et al., 2025; Samil and Adelani, 2026), with important contributions also including corrections to previously released partitions (Abdulmumin et al., 2024; Perez-Ortiz et al., 2024; Mæhlum et al., 2025; Oktem et al., 2025).

Other multi-parallel evaluation corpora include TICO-19 (Anastasopoulos et al., 2020), which targets the medical domain across 38 languages, and NTREX-128 (Federmann et al., 2022), which translates 1,997 WMT19 news segments (Barrault et al., 2019) into 128 languages.<sup>18</sup> Similarly, WMT24++ (Deutsch et al., 2025) extends the WMT24 test set

<sup>18</sup>NTREX-128’s exposure to web crawlers on GitHub appears to have contaminated LLM pre-training data, limiting its use as a strict evaluation benchmark but making it a valuable few-shot training resource.

(Kocmi et al., 2024) to 55 languages. Despite potential GDPR constraints regarding its social media data, it uniquely provides both initial translations and final post-edits, making it ideal for studying translation dynamics (cf. Marmonier et al. 2026). Moving away from English-centric models, the BOUQuET dataset (Andrews et al., 2025) uses original texts handcrafted across eight pivot languages. Spanning 275 variants, it covers diverse registers—such as fiction and dialogues—while maintaining low structural complexity to encourage crowd-sourced extensions.

To the best of our knowledge, no machine translation evaluation corpora for Savoyard or the broader Arpitan dialect continuum currently exist, making the present dataset the first of its kind for this language family.

## 7 Conclusion

We have presented a Savoyard partition of the FLORES+ corpus, together with a range of baseline scores intended to give future work a point of comparison. Savoyard is but one variety within a much larger and diverse whole: the Arpitan dialect continuum extends across the French departments of eastern and central France (including varieties such as Bressan, Lyonnais, Forézien, and Dauphinois), French-speaking Switzerland (encompassing dialects like Vaudois, Genevois, Neuchâtelois, Fribourgeois, and Valaisan), and the Aosta Valley of Italy. All of these are low-resource varieties from the point of view of natural language processing, and their revitalization would likely benefit greatly from translation technology. We hope that the present work will stimulate similar efforts for these dialects elsewhere, contributing in turn to the preservation of humanity’s linguistic diversity.

## Limitations

The main limitations of this work stem from the very scarcity of speakers that motivates it.

**A single-translator corpus.** The workload, initially to be shared between two translators, was ultimately borne almost entirely by one, the pool of expert translators being extremely small. This does not bear on the quality of the resulting corpus, as it was produced by a lexicographer of Savoyard working under the auspices of the *Institut de la langue savoyarde*, the body whose lexicographical and normative work defines the Median Savoyard

standard the partition targets. It does, however, bound what the corpus attests.

**Human evaluation.** The annotation campaign was conducted by the translator himself. The design was blind: outputs were anonymized and the annotator was not informed that the gold references had been injected as a hidden system. Some residual self-preference along stylistic lines cannot be excluded, since the norms against which the systems were judged are those the annotator/translator himself applies. The campaign is best read, accordingly, as establishing agreement with the ILS standard—which is the property our references are meant to have—rather than as an independent measure of acceptability across the entire Savoyard dialect continuum.

**Confounds in the baselines.** The translator contributed in the past to the French–Arpitan Apertium system. Although the devtest was translated from scratch, Apertium’s FR→FRP scores are plausibly inflated by this shared provenance.

## Acknowledgments

This work was funded by the French *Agence Nationale de la Recherche* (ANR) under the project TraLaLaM (“ANR-23-IAS1-0006”), as well as by Inria under the “Défi”-type project COLaF. It was also partly funded by Benoît Sagot’s chair in the PRAIRIE-PSAI institute, funded by the ANR as part of the “France 2030” strategy under the reference ANR-23-IACL-0008.

This work was also granted access to the HPC resources of IDRIS under allocations 2025-AD011015117R2, 2025-AD011012254R5 and 2025-AD011017228 made by GENCI. The authors acknowledge the use of LLMs as coding assistants during the implementation of this work.

We are grateful to François Yvon for his help with the administrative aspects of this project. We also thank Franck Monod, Operations Manager at the *Institut de la langue savoyarde*, and Amandine Pinget, Data Scientist at the *Institut de la langue savoyarde*, for their assistance and support in coordinating this collaboration.

## References

Idris Abdulmumin, Sthembiso Mkhwanazi, Mahlatse Mbooi, Shamsuddeen Hassan Muhammad, Ibrahim Said Ahmad, Neo Putini, Miehleketo Mathabela, Matimba Shingange, Tajuddeen Gwadabe,

- and Vukosi Marivate. 2024. [Correcting FLORES Evaluation Dataset for Four African Languages](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 570–578, Miami, Florida, USA. Association for Computational Linguistics.
- Felermينو Dario Mario Ali, Henrique Lopes Cardoso, and Rui Sousa-Silva. 2024. [Expanding FLORES+ Benchmark for More Low-Resource Settings: Portuguese-Emakhuwa Machine Translation Evaluation](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 579–592, Miami, Florida, USA. Association for Computational Linguistics.
- Dante Alighieri. 1996. *De vulgari eloquentia*. Number 5 in Cambridge Medieval Classics. Cambridge University Press, Cambridge.
- Hèctor Alòs. 2020. [GSOC 2020 rapport final: Traducteur bidirectionnel français-arpitan](#). Apertium Wiki. Accessed: 2026-08-03.
- Antonios Anastasopoulos, Alessandro Cattelan, Zi-Yi Dou, Marcello Federico, Christian Federmann, Dmitriy Genzel, Francisco Guzmán, Junjie Hu, Macduff Hughes, Philipp Koehn, Rosie Lazar, Will Lewis, Graham Neubig, Mengmeng Niu, Alp Öktem, Eric Paquin, Grace Tang, and Sylwia Tur. 2020. [TICO-19: the Translation Initiative for COvid-19](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- 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.
- Graziadio Isaia Ascoli. 1878. Schizzi franco-provenzali. *Archivio Glottologico Italiano*, 3:61–120.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. [Findings of the 2019 Conference on Machine Translation \(WMT19\)](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.
- Michel Bert and James Costa. 2009. [Etude fora : Francoprovençal et occitan en rhône-alpes](#). Rapport de recherche, L’Institut Pierre Gardette.
- Natalia Bichurina. 2019. *L’émergence du franco-provençal : langue minoritaire et communauté autour du Mont-Blanc*. Maison des sciences de l’homme d’Aquitaine. Préface de Manuel Meune.
- Natalia Bichurina and Christiane Dunoyer. 2021. [Le francoprovençal en savoie : Histoire et pratiques contemporaines](#). *Les Dossiers du Musée Savoisien : Revue numérique*, 7. [En ligne].
- L. Boutillier. 2022. [Le terme savoyard « monchu » intègre le dictionnaire](#). *La Savoie*.
- Peter F. Brown, John Cocke, Stephen A. Della Pietra, Vincent J. Della Pietra, Fredrick Jelinek, John D. Lafferty, Robert L. Mercer, and Paul S. Roossin. 1990. A Statistical Approach to Machine Translation. *Comput. Linguist.*, 16(2):79–85.
- Laurie Burchell, Jean Maillard, Antonios Anastasopoulos, Christian Federmann, Philipp Koehn, and Skyler Wang. 2024. [Findings of the WMT 2024 Shared Task of the Open Language Data Initiative](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 110–117, Miami, Florida, USA. Association for Computational Linguistics.
- Isaac Caswell, Elizabeth Nielsen, Jiaming Luo, Colin Cherry, Geza Kovacs, Hadar Shemtov, Partha Talukdar, Dinesh Tewari, Baba Mamadi Diane, Djibrila Diane, Solo Farabado Cissé, Koulako Moussa Doumbouya, Edoardo Ferrante, Alessandro Guasoni, Christopher Homan, Mamadou K. Keita, Sudhamoy DebBarma, Ali Kuzhuget, David Anugraha, and 5 others. 2025. [SMOL: Professionally Translated Parallel Data for 115 Under-represented Languages](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 1103–1123, Suzhou, China. Association for Computational Linguistics.
- David Dale, Laurie Burchell, Jean Maillard, Idris Abdumumin, 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.
- Joseph Désormaux. 1981. *Bibliographie méthodique des parlers de Savoie: langue et littérature*. Laffitte, Marseille, France. Facsimile reprint of the original 1922 edition published in Annecy by Imprimerie commerciale.
- Daniel Deutsch, Eleftheria Briakou, Isaac Rayburn Caswell, Mara Finkelstein, Rebecca Galor, Juraj Juraska, Geza Kovacs, Alison Lui, Ricardo Rei, Jason Riesa, Shruti Rijhwani, Parker Riley, Elizabeth Salesky, Firas Trabelsi, Stephanie Winkler, Biao Zhang, and Markus Freitag. 2025. [WMT24++: Expanding the Language Coverage of WMT24 to 55](#)

- [Languages & Dialects](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 12257–12284, Vienna, Austria. Association for Computational Linguistics.
- Christiane Dunoyer, editor. 2019. *Regards croisés sur la standardisation du francoprovençal: Actes de la conférence annuelle sur l'activité scientifique du Centre d'Études francoprovençales (Saint-Nicolas, 11 novembre 2017)*. Région autonome Vallée d'Aoste, Aoste.
- Alain Favre. 2026. *Diccionèro de fata: Savoyârd – Francès / Français – Savoyard*. Yoran Embanner, Spézet.
- Alain Favre and Dominique Stich. 2005. *Diccionèro de fata / Minidico: Savoyârd / Francès – Français / Savoyard*. Yoran Embanner, Spézet.
- Christian Federmann, Tom Kocmi, and Ying Xin. 2022. [NTREX-128 – News Test References for MT Evaluation of 128 Languages](#). In *Proceedings of the First Workshop on Scaling Up Multilingual Evaluation*, pages 21–24, Online. Association for Computational Linguistics.
- Mikel L. Forcada and Francis M. Tyers. 2016. [Aperitium: a free/open source platform for machine translation and basic language technology](#). In *Proceedings of the 19th Annual Conference of the European Association for Machine Translation: Projects/Products*, Riga, Latvia. Baltic Journal of Modern Computing.
- 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.
- Google DeepMind. 2026. Gemma 4 model card. [https://ai.google.dev/gemma/docs/core/model\\_card\\_4](https://ai.google.dev/gemma/docs/core/model_card_4).
- Isai Gordeev, Sergey Kuldin, and David Dale. 2024. [FLORES+ Translation and Machine Translation Evaluation for the Erzya Language](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 614–623, Miami, Florida, USA. Association for Computational Linguistics.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The Flores-101 Evaluation Benchmark for Low-Resource and Multilingual Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 10:522–538.
- Yvette Graham, Barry Haddow, and Philipp Koehn. 2020. [Statistical Power and Translationese in Machine Translation Evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 72–81, Online. Association for Computational Linguistics.
- Henri Grégoire. 1794. [Rapport sur la nécessité et les moyens d’anéantir les patois et d’universaliser l’usage de la langue française](#). Historical document, Convention nationale. Available on Wikisource.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. [The FLORES Evaluation Datasets for Low-Resource Machine Translation: Nepali–English and Sinhala–English](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6098–6111, Hong Kong, China. Association for Computational Linguistics.
- 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.
- Ralf Herbrich, Tom Minka, and Thore Graepel. 2007. [Trueskill™: A bayesian skill rating system](#). In Bernhard Schölkopf, John Platt, and Thomas Hofmann, editors, *Advances in Neural Information Processing Systems 19: Proceedings of the 2006 Conference*. The MIT Press.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The State and Fate of Linguistic Diversity and Inclusion in the NLP World](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Marzena Karpinska, Philipp Koehn, Benjamin Marie, Christof Monz, Kenton Murray, Masaaki Nagata, Martin Popel, Maja Popović, and 3 others. 2024. [Findings of the WMT24 General Machine Translation Shared Task: The LLM Era Is Here but MT Is Not Solved Yet](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 1–46, Miami, Florida, USA. Association for Computational Linguistics.
- Philipp Koehn. 2004. [Statistical Significance Tests for Machine Translation Evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Philipp Koehn, Franz Josef Och, and Daniel Marcu. 2003. [Statistical phrase-based translation](#). In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - Volume 1*, NAACL ’03, page 48–54, USA. Association for Computational Linguistics.

- Andres Kristol. 2016. [Francoprovençal](#). In Adam Ledgeway and Martin Maiden, editors, *The Oxford Guide to the Romance Languages*. Oxford University Press.
- Ali Kuzhuget, Airana Mongush, and Nachyn-Enkhedorzhu Oorzhak. 2024. [Enhancing Tuvan Language Resources through the FLORES Dataset](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 593–599, Miami, Florida, USA. Association for Computational Linguistics.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Mark Liberman. 2010. [Obituary: Fred Jelinek](#). *Computational Linguistics*, 36(4):595–599.
- Andrés Lou, Juan Antonio Pérez-Ortiz, Felipe Sánchez-Martínez, Miquel Esplà-Gomis, and Víctor M. Sánchez-Cartagena. 2025. [FLORES+ Mayas: Generating Textual Resources to Foster the Development of Language Technologies for Mayan Languages](#). In *Proceedings of Machine Translation Summit XX: Volume 2*, pages 89–90, Geneva, Switzerland. European Association for Machine Translation.
- Petter Mæhlum, Anders Næss Evensen, and Yves Scherrer. 2025. [Improved Norwegian Bokmål Translations for FLORES](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 1124–1132, Suzhou, China. Association for Computational Linguistics.
- Jean Maillard, Cynthia Gao, Elahe Kalbassi, Kaushik Ram Sadagopan, Vedanuj Goswami, Philipp Koehn, Angela Fan, and Francisco Guzman. 2023. [Small Data, Big Impact: Leveraging Minimal Data for Effective Machine Translation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2740–2756, Toronto, Canada. Association for Computational Linguistics.
- Malik Marmonier, Rachel Bawden, and Benoît Sagot. 2026. [Testing the Deliteralization Hypothesis in Human and Machine Translation](#). *Preprint*, arXiv:2605.25686.
- Malik Marmonier, Benoît Sagot, and Rachel Bawden. 2025. [A French Version of the OLDI Seed Corpus](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 1048–1060, Suzhou, China. Association for Computational Linguistics.
- Ministère de l'Éducation nationale, de la Jeunesse et des Sports. 2021. Bulletin officiel n° 47 du 16 décembre 2021. [https://www.education.gouv.fr/sites/default/files/document/BOENJS\\_47\\_1423493.pdf-311421.pdf](https://www.education.gouv.fr/sites/default/files/document/BOENJS_47_1423493.pdf-311421.pdf). Accessed: July 9, 2026.
- NLLB Team. 2024. [Scaling neural machine translation to 200 languages](#). *Nature*, 630(8018):841–846.
- Alp Oktem, Mohamed Aymane Farhi, Brahim Essaidi, Naceur Jabouja, and Farida Boudichat. 2025. [Correcting the Tamazight Portions of FLORES+ and OLDI Seed Datasets](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 1072–1080, Suzhou, China. Association for Computational Linguistics.
- Michèle Oliuéri and Patrick Sauzet. 2016. [Southern Gallo-Romance \(Occitan\)](#). In Adam Ledgeway and Martin Maiden, editors, *The Oxford Guide to the Romance Languages*. Oxford University Press.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a Method for Automatic Evaluation of Machine Translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Juan Antonio Perez-Ortiz, Felipe Sánchez-Martínez, Víctor M. Sánchez-Cartagena, Miquel Esplà-Gomis, Aaron Galiano Jimenez, Antoni Oliver, Claudi Aventín-Boya, Alejandro Pardos, Cristina Valdés, Jusèp Loís Sans Socasau, and Juan Pablo Martínez. 2024. [Expanding the FLORES+ Multilingual Benchmark with Translations for Aragonese, Aranese, Asturian, and Valencian](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 547–555, Miami, Florida, USA. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- Matt Post. 2018. [A Call for Clarity in Reporting BLEU Scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- République française. 1951. Loi n° 51-46 du 11 janvier 1951 relative à l'enseignement des langues et dialectes locaux. <https://www.legifrance.gouv.fr/loda/id/JORFTEXT000000886638>. Journal officiel de la République française, 13 janvier 1951, p. 483. Accessed: 2026-06-05.
- Hadia Mohmmadosman Ahmed Samil and David Ifeoluwa Adelani. 2026. [Sudanese-Flores: Extending FLORES+ to Sudanese Arabic Dialect](#). In *Proceedings of the 7th Workshop on African Natural Language Processing (AfricaNLP 2026)*, pages 243–247,

- Rabat, Morocco. Association for Computational Linguistics.
- Tove Skutnabb-Kangas. 2000. *Linguistic Genocide in Education—Or Worldwide Diversity and Human Rights?* Lawrence Erlbaum Associates, Mahwah, NJ. Includes bibliographical references and index.
- John Charles Smith. 2016. [French and northern Gallo-Romance](#). In Adam Ledgeway and Martin Maiden, editors, *The Oxford Guide to the Romance Languages*. Oxford University Press.
- Matthew Snover, Bonnie Dorr, Rich Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A Study of Translation Edit Rate with Targeted Human Annotation](#). In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA. Association for Machine Translation in the Americas.
- Dominique Stich. 2001. *Francoprovençal : proposition d'une orthographe supra-dialectale standardisée*. Thèse de doctorat, Université Paris 5. Moteur de recherche des thèses françaises: theses.fr.
- Dominique Stich. 2003. *Dictionnaire francoprovençal / français, français / francoprovençal*. Éditions Le Carré, Thonon-les-Bains.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, page 3104–3112, Cambridge, MA, USA. MIT Press.
- Gaston Tuaillon. 1983. *Le Francoprovençal: Progrès d'une définition*. Centre d'Études Francoprovençales « René Willien », Saint-Nicolas (Aoste).
- Gaston Tuaillon. 1984. La graphie de Conflans pour le patois savoyard. *Les Nouvelles du Centre d'Études francoprovençales « René Willien »*, page 5.
- Gaston Tuaillon. 1989. Le francoprovençal, langue oubliée. In Geneviève Vernes, editor, *Vingt-cinq communautés linguistiques de la France*, volume 1, page 204. L'Harmattan, Paris.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NeurIPS'17, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Roger Viret. 2015. *Graméra savoyârda: Arbané*. Institut de la langue savoyarde, Challes-les-Eaux, France.
- Roger Viret. 2019. *Dictionnaire français–savoyard*. Troisième édition revue et augmentée.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-Art Natural Language Processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Mike Zhang and Antonio Toral. 2019. [The Effect of Translationese in Machine Translation Test Sets](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 73–81, Florence, Italy. Association for Computational Linguistics.

## A Data Distillation prompt

This appendix provides an example of the template used for the *Distilled (Gold+Silver)* dataset generation, where Gemma-4-31B was provided with French–Arpitan–Savoyard few-shot triplet demonstrations to post-edit the rule-based Apertium drafts.

### Gemma-4-31B Distillation Prompt Template

Savoyard (endonym: *savoyârd*) is a Franco-Provençal dialect of the Gallo-Romance family. It is spoken in some territories of the historical Duchy of Savoy, nowadays a geographic area spanning Savoie and Haute-Savoie, France, and the Canton of Geneva, Switzerland.

You are to translate a French segment of text into Savoyard. To assist you in this task, we will provide you with translation hypotheses produced by Apertium, a rule-based machine translation platform. It will provide a rough candidate translation from French into an Arpitan macrolanguage. Your task is to post-edit and refine this candidate to produce a fluent, orthographically correct, and natural Savoyard translation.

To guide your style and vocabulary, here is a parallel corpus of expert translations demonstrating the task at hand:

<parallel\_corpus>

**French:** *Des scientifiques de l'école de*

médecine de l'université de Stanford ont annoncé ce lundi la création d'un nouvel outil de diagnostic, qui permettrait de différencier les cellules en fonction de leur type. Il s'agit d'une petite puce imprimable, qui peut être produite au moyen d'une imprimante à jet d'encre standard, pour un coût d'environ un cent de dollar pièce.

**Apertium Candidate:** Des científicos de l'ècoula de mèdecina de l'univèrsitàt de Stanford ant anonciè ceti delon la crèacion d'un novèl outil de diagnostic, que pèmetrèt de difèrenciér les cèlules en fonccion de lor tipo. Il s'agèt d'una petit puge imprimable, que pôt être fêta u moyen d'una èstampiosa a jèt d'encro estandârd, por un côte d'enveron un cent de dolâr pièce.

**Savoyard:** Des científicos de l'ècoula de mèdecina de l'univèrsitàt de Stanford ant anonciè ceti delon la crèacion d'un novèl outil de diagnostic, que pèmetrèt de difèrenciér les cèlules en fonccion de lor tipo. O est una petiôta puge emprimâbla, que pôt être fêta avouéc una èstampiosa a jèt d'encro estandârd, por un côte d'enveron un cent de dolâr pièce.

...

**French:** Les compagnies aériennes qui les proposent sont notamment Air Canada, Delta Air Lines, Lufthansa pour les vols en provenance des États-Unis ou du Canada, et WestJet.

**Apertium Candidate:** Les compagnies d'aviacion que los propôsont sont spécialement Air Canadâ, Dèlta Air Lines, Lufthansa por los vòls en provegnance des Ètats-Unis ou du Canadâ, et WestJet.

**Savoyard:** Les compagnies d'aviacion que los propôsont sont spécialement Air Canada, Dèlta Air Lines, Lufthansa por los vòls que vegnont des Ètats-Unis ou du Canadâ, et WestJet.

**French:** Dans tous les cas, vous devez réserver par téléphone directement

auprès de la compagnie aérienne.

**Apertium Candidate:** Dens tôs los câs, vos dête resèrvar per tèlèfono dirèctament uprès de la compagnie d'aviacion.

**Savoyard:** Diens tôs los câs, vos dêde resèrvar per tèlèfono dirèctament uprès de la compagnie d'aviacion.

</parallel\_corpus>

You must only output your final translation between the XML tags and absolutely nothing else. Do not include any reasoning, notes, or introductions. The expected XML format is: <translation>your translation here</translation>

Here is the segment you must translate:

**French:** La première avancée majeure en termes d'abstraction fut l'utilisation de chiffres pour représenter les nombres.

**Apertium Candidate:** La première avancîe majore en tèrmenos d'abstraccion at étâ l'usâjo de chifros por représentar los nombros.

**Savoyard:**

## B LLM Translation Prompt Templates

This appendix contains the prompts used for the evaluation of the Gemma-4 models (31B, E4B, and E2B) described in Section 4. We provide examples in the FR→FRP translation direction; the FRP→FR differs minimally, for instance in the order in which dev set bitexts are presented.

### B.1 Zero-Shot Templates

#### B.1.1 Direct (No-Think) Modality

The following template was used for the zero-shot direct generation condition. The model was instructed to output the final translation immediately within XML tags.

#### Zero-Shot Direct Translation Prompt

Savoyard (endonym: *savoyârd*) is a Franco-Provençal dialect of the Gallo-Romance family. It is spoken in some territories of the historical Duchy of Savoy, nowadays a geographic area spanning Savoie and

Haute-Savoie, France and the Canton of Geneva, Switzerland.

You are to translate the following French text segment into Savoyard. **You must only output the translation between the XML tags and absolutely nothing else. Do not include any reasoning, notes, or introductions.** The expected XML format is: `<translation>your translation here</translation>`

Segment to translate:

*Les travailleurs doivent souvent obtenir l'accord de leurs supérieurs pour toute décision qu'ils prennent et sont censés obéir aux instructions de leurs supérieurs sans poser de questions.*

### B.1.2 Reasoning Modality

The following template was used for the zero-shot reasoning condition, explicitly allowing the model to utilize an internal monologue prior to emitting the translation.

#### Zero-Shot Reasoning Translation Prompt

Savoyard (endonym: *savoyârd*) is a Franco-Provençal dialect of the Gallo-Romance family. It is spoken in some territories of the historical Duchy of Savoy, nowadays a geographic area spanning Savoie and Haute-Savoie, France and the Canton of Geneva, Switzerland.

You are to translate the following French text segment into Savoyard. **You can reflect on the task at hand prior to producing your final translation.** You must output your final translation between XML tags thusly: `<translation>your translation here</translation>`

Segment to translate:

*Les travailleurs doivent souvent obtenir l'accord de leurs supérieurs pour toute décision qu'ils prennent et sont censés obéir aux instructions de leurs supérieurs sans poser de questions.*

## B.2 Few-Shot Templates

### B.2.1 Direct (No-Think) Modality

For the few-shot direct condition, the prompt was appended with the entire structured *dev* partition bitext as in-context demonstrations, immediately followed by the target source query.

#### Few-Shot Direct Translation Prompt

Savoyard (endonym: *savoyârd*) is a Franco-Provençal dialect of the Gallo-Romance family. It is spoken in some territories of the historical Duchy of Savoy, nowadays a geographic area spanning Savoie and Haute-Savoie, France and the Canton of Geneva, Switzerland.

You will be asked to translate a French text segment into Savoyard. **You must only output the translation between the XML tags and absolutely nothing else. Do not include any reasoning, notes, or introductions.** The expected XML format is: `<translation>your translation here</translation>`

To assist you in this task, here is a small parallel corpus. It should dispel doubts you might have regarding the orthography and grammar of this romance dialect:

`<parallel_corpus>`

**French:** *Des scientifiques de l'école de médecine de l'université de Stanford ont annoncé ce lundi la création d'un nouvel outil de diagnostic, qui permettrait de différencier les cellules en fonction de leur type. Il s'agit d'une petite puce imprimable, qui peut être produite au moyen d'une imprimante à jet d'encre standard, pour un coût d'environ un cent de dollar pièce.*

**Savoyard:** *Des científicos de l'ècoula de mèdecina de l'univèrsitàt de Stanford ant anonciè ceti delon la crèacion d'un novél outil de diagnostic, que pèmetrèt de difèrenciér les cèlules en fonccion de lor tipo. O est una petiòta puge emprimâbla, que pôèt ètre fèta avouéc una èstampiosa a jèt d'encro estandârd, por un côte d'enveron un cent de dolâr piéce.*

**French:** *Selon les chercheurs principaux, cela pourrait permettre une détection précoce du cancer, de la tuberculose, du VIH et du paludisme chez les patients des pays à faible revenu, où le taux de survie à des maladies telles que le cancer du sein peut être deux fois moins élevé que dans les pays riches.*

**Savoyard:** *D'après los chërchiors principaux, cen porrêt pèrmetre una dètèccion couètosa du chancro, de la tubèrculosa, du VIH et du paludismo chiéz los malâdos des pays a fêbla renta, yô que lo tôx de survia a des maladies coment lo chancro du nènè pôût être doux côps muens hiôt que diens los pays rechos.*

...

**French:** *Dans tous les cas, vous devez réserver par téléphone directement auprès de la compagnie aérienne.*

**Savoyard:** *Diens tôs los câs, vos dêde resèrvan per tèlèfono dirèctament uprès de la compagnie d'aviacion.*  
</parallel\_corpus>

Here is the French segment you must translate into **Savoyard**:

*Les travailleurs doivent souvent obtenir l'accord de leurs supérieurs pour toute décision qu'ils prennent et sont censés obéir aux instructions de leurs supérieurs sans poser de questions.*

**A reminder that you must only output the translation between the XML tags and absolutely nothing else.** Do not include any reasoning, notes, or introductions. Format: <translation>your translation here</translation>

### B.2.2 Reasoning Modality

For the few-shot reasoning condition, the system prompt instructed the model to reason before outputting the final tags, combined with the structural sequence of the *dev* bitext.

### Few-Shot Reasoning Translation Prompt

Savoyard (endonym: *savoyârd*) is a Franco-Provençal dialect of the Gallo-Romance family. It is spoken in some territories of the historical Duchy of Savoy, nowadays a geographic area spanning Savoie and Haute-Savoie, France and the Canton of Geneva, Switzerland.

You will be asked to translate a French text segment into Savoyard. **You can reflect on the task at hand prior to producing your final translation.** You must output your final translation between XML tags thusly: <translation>your translation here</translation>

To assist you in this task, here is a small parallel corpus. It should dispel doubts you might have regarding the orthography and grammar of this romance dialect:

<parallel\_corpus>

**French:** *Des scientifiques de l'école de médecine de l'université de Stanford ont annoncé ce lundi la création d'un nouvel outil de diagnostic, qui permettrait de différencier les cellules en fonction de leur type. Il s'agit d'une petite puce imprimable, qui peut être produite au moyen d'une imprimante à jet d'encre standard, pour un coût d'environ un cent de dollar pièce.*

**Savoyard:** *Des scientifiques de l'ècoula de mèdecina de l'univèrsitàt de Stanford ant anonciê ceti delon la crèacion d'un novél outil de diagnostico, que pèrmetrèt de difèrenciér les cèlules en fonccion de lor tipo. O est una petiôta puge emprimâbla, que pôût être fêta avouéc una èstampiosa a jèt d'encro estandârd, por un côte d'enveron un cent de dolâr pièce.*

**French:** *Selon les chercheurs principaux, cela pourrait permettre une détection précoce du cancer, de la tuberculose, du VIH et du paludisme chez les patients des pays à faible revenu, où le taux de survie à des maladies telles que le cancer du sein peut être deux fois moins élevé que dans les*

*pays riches.*

**Savoyard:** *D'après los chèrchiors principâls, cen porrêt pèrmètre una dètèccion couètosa du chancro, de la tubèrculosa, du VIH et du paludismo chiéz los malâdos des payis a fêbla renta, yô que lo tôx de survia a des maladies coment lo chancro du nè nè pôt être doux côps muens hiôt que diens los payis rechos.*

...

**French:** *Dans tous les cas, vous devez réserver par téléphone directement auprès de la compagnie aérienne.*

**Savoyard:** *Diens tôs los câs, vos dêde resèrvar per tèlèfono dirèctament uprés de la compagnie d'aviacion.*

</parallel\_corpus>

Here is the French segment you must translate into **Savoyard**:

*Les travailleurs doivent souvent obtenir l'accord de leurs supérieurs pour toute décision qu'ils prennent et sont censés obéir aux instructions de leurs supérieurs sans poser de questions.*

**A reminder that you can reflect on the task at hand prior to producing your final translation** but must output your final translation between XML tags thusly: <translation>your translation here</translation>

## C Statistical Significance Matrices

To assess whether the observed performance margins between systems are statistically meaningful, we conduct paired bootstrap resampling with 1,000 iterations. Figures 4 and 5 illustrate the mean chrF++ differences between all evaluated configurations for both translation directions.

Cells marked with an asterisk (\*) indicate that the performance delta between the row system and column system is statistically significant at  $p < 0.05$ .




Figure 4: Statistical Significance Heatmap: FR → FRP




Figure 5: Statistical Significance Heatmap: FRP → FR

## **D Data Sample**

We provide a random sample of the compiled dataset in Table 3.

| ID  | English Source                                                                                                                                                                                                                    | French Translation                                                                                                                                                                                                                                                                                                  | Savoyard Translation                                                                                                                                                                                                                                                          |
|-----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 992 | The tourist season for the hill stations generally peaks during the Indian summer.                                                                                                                                                | <i>La saison touristique des stations de montagne culmine généralement pendant l'été indien.</i>                                                                                                                                                                                                                    | <i>La sêson touristica des stacions de montagne avente son maximon en gènèral pendent lo chôd-temps indien.</i>                                                                                                                                                               |
| 553 | Middle distance running is a relatively inexpensive sport; however, there are many misconceptions regarding the few pieces of equipment required to participate.                                                                  | <i>Le demi-fond est un sport relativement peu coûteux. L'idée qu'il nécessiterait un équipement basique est aussi fausse que la plupart des mythes qui entourent ce sport.</i>                                                                                                                                      | <i>Lo demi-fond est un sport .que côte pas tant. L'idê qu'il arêt fôta d'una forniture basica est asse fôssa coment la plepârt des mitos qu'entouront ceti sport.</i>                                                                                                         |
| 152 | Virgin have only purchased the 'good bank' of Northern Rock, not the asset management company.                                                                                                                                    | <i>Virgin n'a acheté que la «banque saine» de Northern Rock, pas la société de gestion d'actifs.</i>                                                                                                                                                                                                                | <i>Virgin at achetâ solament la «banca sana» de Northern Rock, pas la sociêtât de maneyence d'actifs.</i>                                                                                                                                                                     |
| 189 | The typhoon is estimated to be moving toward China at eleven kph.                                                                                                                                                                 | <i>On estime que le typhon se dirige vers la Chine à une vitesse de 11 km/h.</i>                                                                                                                                                                                                                                    | <i>On èstime que lo tifon s'aduit vers la China a una vitèsse de 11 km/h.</i>                                                                                                                                                                                                 |
| 463 | Unfortunately, when you make a DVD, it's borders will most likely be cut too, and if the video had subtitles too close to the bottom, they won't be fully shown.                                                                  | <i>Malheureusement, lorsque vous réalisez un DVD, ses bords seront probablement coupés aussi, et si la vidéo comportait des sous-titres trop proches du bas de l'écran, ils n'apparaîtront pas.</i>                                                                                                                 | <i>Mâlherosament, quand vos fabrecâd un DVD, sos bôrds seront probâblement copâs asse-ben, et se la vidèò at des sot-titros trop prôches du bâs de l'ècran, ils sè vèrront pas.</i>                                                                                           |
| 953 | Airlines such as Emirates, Etihad Airways, Qatar Airways & Turkish Airlines have greatly expanded their services to Africa, and offer connections to many major African cities at competitive rates than other European airlines. | <i>Les compagnies aériennes telles que Airlines such as Emirates, Etihad Airways, Qatar Airways et Turkish Airlines ont grandement développé leur offres en Afrique, et offrent des liaisons vers les villes africaines majeures à des tarifs compétitifs comparés aux autres compagnies aériennes européennes.</i> | <i>Les compagnies d'aviacion coment Emirates, Etihad Airways, Qatar Airways et Turkish Airlines ant rudament dêvelopâ lor ofres en Africa, et ofront des liêsons vers les veles afriquènes mès grantes a des prix compétitifs comparâ ux compagnies d'aviacion eropèènes.</i> |
| 48  | Mr Reid managed to drive the New Zealand's A1GP car, Black Beauty at speeds over 160km/h seven times over the bridge.                                                                                                             | <i>M. Reid est parvenu au volant de la voiture néo-zélandaise de l'A1 GP dénommée Black Beauty, à dépasser à 7 reprises une vitesse de 160 km/h au dessus du pont.</i>                                                                                                                                              | <i>M. Reid est aventâ u volant de la vouètura nèò-zèlandèsa de l'A1 GP apelâye Black Beauty, a dêpassar a 7 reprêses una vitèsse de 160 km/h u dessus du pont.</i>                                                                                                            |
| 406 | Yes! King Tutankhamun, sometimes referred to as "King Tut" or "The Boy King", is one of the most well known ancient Egyptian kings in modern times.                                                                               | <i>Oui! Le roi Toutankhamon, parfois appelé «Roi Tout» ou «L'Enfant roi», est l'un des anciens monarques égyptiens les plus connus des temps modernes.</i>                                                                                                                                                          | <i>Ouè! Lo rê Toutânkhamon, des côps apelâ «Rê Tot» ou «L'Enfant rê», est yon des ancians monarcos ègipcians mès famox des temps modèrns.</i>                                                                                                                                 |
| 937 | In some federal countries, such as the United States and Canada, income tax is levied both at the federal level and at the local level, so the rates and brackets can vary from region to region.                                 | <i>Dans certains pays fédéraux tels que les États-Unis et le Canada, l'impôt sur le revenu est prélevé à la fois au niveau fédéral et au niveau local, de sorte que les taux et les tranches d'imposition peuvent varier d'une région à l'autre.</i>                                                                | <i>Diens des payis fèdèrâls coment los États-Unis et lo Canadâ, l'emposicion sur la renta est prèlevâye a côp u nivél fèdèrâl et u nivél local, de façon que los tôx et les trenches d'emposicion pouessont vareyer d'una règion a l'ôtro.</i>                                |
| 86  | The amount and thickness of the pack ice, according to Pittman, is the worst it has been for sealers in the past 15 years.                                                                                                        | <i>Selon M. Pittman, la quantité et l'épaisseur de la banquise sont les pires qu'aient connues les chasseurs de phoques au cours des 15 dernières années.</i>                                                                                                                                                       | <i>D'après M. Pittman, la quantitât et l'èpèssior de la banquisa sont les pires qu'èyont cognèssues los chaciors de foques diens los 15 dêrrièrs ans.</i>                                                                                                                     |

Table 3: Random sample of ten sentence pairs from the *dev* partition of our Savoyard version of the FLORES+ corpus.