

# MedInjection-FR: Exploring the Role of Native, Synthetic, and Translated Data in Biomedical Instruction Tuning

**Ikram Belmadani<sup>1,2</sup>, Oumaima El Khettari<sup>2</sup>,  
Pacôme Constant dit Beauvils<sup>3,4</sup>, Benoit Favre<sup>1,5</sup>, Richard Dufour<sup>2</sup>**

<sup>1</sup>Aix-Marseille Univ., CNRS, LIS UMR 7020, 13000 Marseille, France,

<sup>2</sup>Nantes Univ., École Centrale Nantes, CNRS, LS2N, UMR 6004, 44000 Nantes, France,

<sup>3</sup>Nantes Univ., CHU Nantes, PHU 11: Santé Publique,

Clinique des données, INSERM, CIC 1413, 44000 Nantes, France ,

<sup>4</sup> Nantes Univ., CNRS, INSERM, L'institut du thorax, 44000 Nantes, France,

<sup>5</sup> Université Grenoble Alpes, CNRS, INRIA, Grenoble INP,

{first.last}@{univ-amu.fr, univ-nantes.fr, chu-nantes.fr}

## Abstract

Instruction tuning has become essential for adapting large language models (LLMs) to follow domain-specific prompts. Yet, in specialized fields such as medicine, the scarcity of high-quality French instruction data limits effective supervision. To address this gap, we introduce MedInjection-FR, a large-scale French biomedical instruction dataset comprising 571K instruction-response pairs drawn from three complementary sources: native, synthetic, and translated data. We design a controlled experimental framework to systematically assess how data provenance affects instruction tuning, using Qwen-4B-Instruct fine-tuned across seven configurations combining these sources. Results show that native data yield the strongest performance, while mixed setups, particularly native and translated, provide complementary benefits. Synthetic data alone remains less effective but contributes positively when balanced with native supervision. Evaluation on open-ended QA combines automatic metrics, LLM-as-a-judge assessment, and human expert review; although LLM-based judgments correlate best with human ratings, they show sensitivity to verbosity. These findings highlight that data authenticity and diversity jointly shape downstream adaptation and that heterogeneous supervision can mitigate the scarcity of native French medical instructions.

**Keywords:** French biomedical NLP, instruction tuning, data provenance, large language models

## 1. Introduction

Supervised fine-tuning (SFT) has become the standard approach to adapting LLMs to follow instructions (Shengyu et al., 2023). However, adaptation in specialized domains, such as medicine, cannot rely solely on continued pretraining over domain-specific text. While domain-adaptive pretraining (DAPT) or continual pretraining (CPT) can enrich models with in-domain vocabulary and factual knowledge, these methods do not necessarily improve, and may even erode, instruction-following behavior (Christophe et al., 2024). In practice, models exposed to extended domain pretraining sometimes “forget” how to respond coherently to prompts or adhere to task formats learned during earlier instruction tuning (Kotha et al., 2024). These limitations make instruction-oriented supervision indispensable for aligning models with clinical reasoning and communicative expectations.

Yet, this process relies on a crucial resource: an **instruction dataset**. Such datasets are not always easy to obtain, especially in expert domains like the medical field, where creating high-quality (instruction, response) pairs requires substantial domain knowledge. Constructing these datasets often demands manual labor, either for curating the pairs themselves or for evaluating automatically gener-

ated ones to ensure factual and clinical accuracy.

Moreover, most existing biomedical instruction datasets are centered on English, with only a few exceptions for the French language, such as FrenchMedMCQA (Labrak et al., 2023a) and MedQAI (Bazoge, 2025). This leaves a significant gap for other languages, particularly French, where linguistic, legal, and ethical limitations often constrain access to medical data. As a result, the scarcity of large, high-quality native resources raises the question of whether alternative data sources, such as automatically generated or translated instructions, can effectively complement or partially replace native supervision.

While synthetic and translated instructions can potentially increase data diversity and expand topic coverage, their true impact on downstream adaptation remains insufficiently understood. Do they actually enhance instruction-following and domain reasoning, or do they introduce stylistic and factual inconsistencies that hinder alignment? Addressing this uncertainty requires a systematic comparison of data sources under controlled conditions.

To this end, we introduce MedInjection-FR, a new French biomedical instruction dataset designed to inject medical knowledge into LLMs through diverse forms of supervision. MedInjection-FR is composed of three complementary subsets: (i) **Na-**

NATIVE	TRANSLATED	SYNTHETIC
<p><b>Instruction</b></p> <p>Un patient souffre d'une compression médullaire. Testant sa flexion de hanche, il décolle la cuisse du lit mais ne résiste pas. Il étend la jambe dans le lit, mais pas contre pesantier. Il remonteorteils et pied gauche contre vous, plus faiblement qu'à droite. Cotation MRC ?</p> <p><b>Options</b></p> <p>A: Muscle psoas iliaque 2/5, B: Muscle psoas iliaque 3/5, C: Muscle quadriceps 2/5, D: Muscle quadriceps 3/5, E: Releveurs du pied 4/5 à gauche</p> <p><b>Answer</b></p> <p><b>B C E</b></p>	<p><b>Instruction</b></p> <p>Un psychologue souligne que les parents devraient utiliser des conséquences appropriées pour gérer les mauvais comportements. Le but peut être de confirmer un handicap supposé, attirer l'attention, obtenir du pouvoir ou se venger. Quelle orientation théorique ?</p> <p><b>Options</b></p> <p>A: Adlérienne, B: Existentielle, C: Cognitivo-comportementale, D: Systèmes, E: Gestalt, F: Analyse transactionnelle, G: Psychodynamique, H: Humaniste</p> <p><b>Answer</b></p> <p><b>A</b></p>	<p><b>Context</b></p> <p>Le cancer anaplasique de la thyroïde (CAT) est parmi les cancers les plus agressifs avec une survie médiane de 4 mois. Orientation urgente vers centre multidisciplinaire essentiel. Les analyses moléculaires et nouvelles thérapies ciblées améliorent les résultats. Le CAT avec mutation BRAF V600E bénéficie de dabrafenib et trametinib en première ligne.</p> <p><b>Instruction</b></p> <p>Quels sont les facteurs pronostiques associés au cancer anaplasique de la thyroïde?</p> <p><b>Answer</b></p> <p>Les facteurs pronostiques incluent la présence de mutations BRAF V600E, le stade de la maladie, l'âge du patient, et la réponse au traitement.</p>

Figure 1: Sample instances from the three MedInjection-FR components. Each includes the task type (MCQ, MCQU, or OEQ) and, when applicable, its supporting context. The figure illustrates the diversity of medical reasoning tasks and prompt styles across data sources.

**Native** instructions derived directly from real French medical texts and resources; (ii) **Synthetic** instructions automatically generated from French biomedical abstracts and clinical cases using large generative models; and (iii) **Translated** instructions adapted from existing English biomedical instruction datasets.

This composition enables a systematic investigation of how the nature and provenance of training data influence model adaptation. We hypothesize that native data, being linguistically and culturally aligned, offer the most direct benefit for medical reasoning in French, but that combining heterogeneous sources may further enhance robustness and generalization particularly when native data are scarce.

To test this hypothesis, we conduct a series of experiments in which instruction-tuned LLMs are trained on each subset individually and on different combinations thereof. The resulting models are evaluated on multiple biomedical question answering and reasoning benchmarks. This study allows us to quantify the contribution of each data type and analyze the trade-offs between authenticity, diversity, and linguistic adaptation in instruction tuning. Our contributions are threefold:

- We release MedInjection-FR, the first open large-scale French biomedical instruction dataset, comprising 571 436 (instruction, response) pairs that integrate native, synthetic, and translated supervision sources: <https://github.com/ikram28/MedInjection-FR>.
- We provide an experimental framework to assess the effect of data origin on the performance of instruction-tuned models.
- We offer empirical insights into the relative benefits of native versus augmented data for medical domain adaptation in French.

## 2. Related Work

Instruction tuning has emerged as a key paradigm for aligning LLMs with human intent (Lou et al., 2024; Han et al., 2025; Liang et al., 2025). It allows the model to generalize to unseen tasks through supervised fine-tuning or instruction-response pairs. This paradigm has also been extended to the biomedical domain, with adaptations of general instruction-tuning datasets resulting in AlpaCare (Zhang et al., 2023), BioInstruct (Tran et al., 2024), and MedAlpaca (Han et al., 2023), which provide diverse medical instruction-response pairs spanning a diversity of tasks. Other works, such as MedMCQA (Pal et al., 2022), MedQA (Jin et al., 2021), and PubMedQA (Jin et al., 2019), provide large-scale benchmarks for evaluating medical reasoning and factuality, but they remain English-centric.

French biomedical datasets are scarce, and notable examples include FrenchMedMCQA (Labrak et al., 2023a) and MediQAI (Bazoge, 2025). To overcome annotation costs and privacy constraints, recent work leverages LLMs to generate synthetic instructions (Shashidhar et al., 2025; Wang et al., 2023; Zeng et al., 2024) or translate existing datasets (Kim and Park, 2024). The trade-off between native, synthetic, and translated data, however, remains an open question. MedInjection-FR addresses this by providing a French biomedical instruction dataset combining all three sources, enabling systematic study of data provenance in instruction-tuned models.

## 3. Dataset Construction

### 3.1. Overview

The MedInjection-FR dataset comprises 571 436 instruction-response pairs distributed across three complementary components: native, synthetic, and translated. As shown in Table 1, the translated subset constitutes the largest portion of the dataset

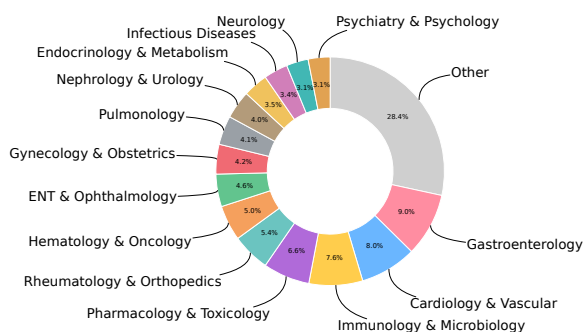


Figure 2: Distribution of medical specialties in MedInjection-FR (native + synthetic components). Proportions are normalized across fourteen major medical categories derived from specialty meta-data.

Component	Train	Validation	Test	Total
Native	57 563	5055	14 629	77 247
Generated	76 506	—	—	76 506
Translated	366 370	38 011	13 293	417 674
<b>Total</b>	<b>500 439</b>	<b>43 066</b>	<b>27 931</b>	<b>571 436</b>

Table 1: Number of instruction–response pairs in each component (native, synthetic, and translated) and their respective splits into training, validation, and test sets.

(417 K pairs), followed by the native (77 K) and synthetic (76 K) subsets. Examples from each component are illustrated in Figure 1. Each subset is divided into training, validation, and test splits.

The dataset includes three task formats: open-ended question answering (OEQ), multiple-choice question answering with multiple correct answers (MCQ), and multiple-choice question answering with a unique correct answer (MCQU). Each of these may appear either as a context-based task, in which a medical scenario is described and followed by a question-answer pair derived from that context, or a standalone instruction. Their respective distributions are reported in Table 2.

Specialty metadata were available for the native and synthetic components, and grouped into fourteen broad medical categories following the standard French taxonomy. As shown in Figure 2, MedInjection-FR spans a wide range of clinical and theoretical domains, with the most represented fields being gastroenterology (9%), cardiology (8%), immunology and microbiology (7.6%), and pharmacology (6.6%), supporting a general-purpose use for French medical instruction tuning.

### 3.2. Native Data

The native component combines curated datasets and web-scraped French medical resources to re-

QA Type	OEQ	MCQ	MCQU
<b>Count</b>	57 509	59 592	454 335

Table 2: Counts of OEQ, MCQ and MCQU question types within MedInjection-FR.

Model	BLEU	COMET
<b>GPT-4o-mini</b>	51.01	0.8751
<b>Gemini 2.0 Flash</b>	53.72	0.8783
<b>WMT 2024 best system</b>	53.54	0.8760

Table 3: BLEU and COMET scores obtained by Gemini 2.0 Flash-001 and GPT-4o-mini on biomedical translation benchmarks, compared to the best WMT 2024 system.

flect authentic domain knowledge. It includes 526 question-answer pairs collected from S-Editions (S-E), a French platform for educational resources for medical students, and 32 603 items from MediQA (Bazoge, 2025), which are sourced from national medical examinations covering 41 medical specialties. FrenchMedMCQA (Labrak et al., 2023a) contributes 3 105 pharmacy-focused multiple-choice questions, while two repositories by mlabonne (medical-cases-fr (Mlabonne, 2024a) and medical-mcqa-fr (Mlabonne, 2024b)) add 12 194 examples originating from French medical exam databases. The FrBMedQA (Kaddari and Bouchentouf, 2022) dataset complements these sources with 19 836 questions derived from French biomedical Wikipedia articles spanning eight UMLS semantic groups (chemicals and drugs, anatomy, physiology, disorders, phenomena, procedures, genes and molecular sequences, and devices). These questions were originally formulated in a closed format and were converted into multiple-choice questions using GPT-4o-mini (Hurst et al., 2024) to standardize the structure across sources.

Finally, the native subset integrates parallel biomedical translation corpora derived from the WMT challenge repositories (bio). These bilingual data were reformulated as instruction-response pairs, where each instruction requested the French translation of an English biomedical passage.

### 3.3. Translated Data

To expand data coverage, a large collection of English biomedical instruction datasets was translated into French using Gemini 2.0 Flash (Goo) and GPT-4o-mini (Hurst et al., 2024). The translated component comprises 417 674 instruction-response pairs derived from several established English resources, including MedQA (Hendrycks et al., 2020), PubMedQA (Jin et al., 2019), MedMCQA (Pal et al., 2022), six medical categories of MMLU (Hendrycks et al., 2020), K-QA (Manes et al., 2024), MMLU-

PRO (Wang et al., 2024, psychology, biology, and health), and MedXpertQA (Zuo et al., 2025).

To assess translation quality, we evaluated the outputs on the WMT 2024 Biomedical Translation Task corpus (Neves et al., 2024) using BLEU and COMET metrics. As shown in Table 3, both models achieved performance comparable to the best WMT 2024 system, suggesting that the translated subset maintains high semantic fidelity and linguistic quality.

### 3.4. Synthetic Data

The synthetic component was generated exclusively for the training set using GPT-4o (Hurst et al., 2024) as the generation model. Source contexts included clinical cases from the DEFT-2021 (Poulain and Connes, 2021) and DIAMED (Labrak et al., 2024) datasets, and biomedical abstracts from the MORFITT (Labrak et al., 2023b) corpus. From these materials, GPT-4o produced 33 493 MCQ and 43 013 OEQ. The complete generation prompts for each corpus and task type are reported in Appendix 13.1.

Each clinical case served as the basis for eight distinct instructional tasks: clinical note summarization, factual medical question answering, differential diagnosis support, treatment plan suggestion, laboratory result interpretation, patient instruction generation, drug interaction checking, and medical specialty classification. Abstracts were used to generate a smaller set of tasks centered on summarization, factual question answering, and specialty classification.

To assess the quality of the generated data, we employed four LLMs as evaluators: GPT-4.1-mini (OpenAI, 2025a), Gemini 2.0 Flash (Goo), MedGemma-27B (Sellergren et al., 2025), and Qwen3-Next-80B-A3B-Instruct (Team, 2025). These models were selected to provide diverse perspectives on evaluation, covering variations in domain specialization (general-purpose vs. medical-specific), model size (medium vs. large), architectural design (dense vs. mixture-of-experts), and licensing (open-source vs. proprietary).

Each model rated the generated question-answer pairs according to predefined criteria. For MCQ data, scores ranged from 1 (irrelevant or incorrect) to 3 (accurate and contextually coherent), the prompt is specified in Figure 13. For OEQ data, a five-point scale was used, where 1 denoted irrelevance and 5 indicated a fully accurate and comprehensive response. The extended scale for OEQ was adopted to reflect the varying degrees of partial correctness typical of open-ended responses. The automatic evaluation prompts used for synthetic data validation are detailed in Appendix 13.2.

As shown in Figure 3, all evaluator models rated the majority of MCQ generations as high-quality

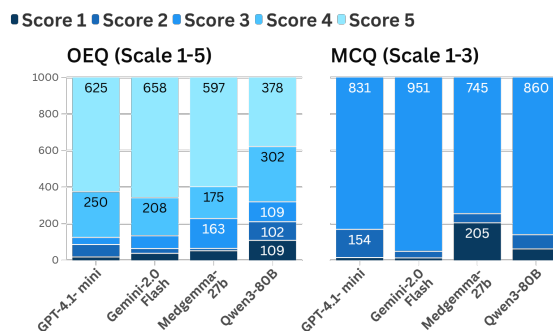


Figure 3: Quality evaluation of synthetic data by four LLM judges. The left panel reports OEQ scores on a 1-5 scale, and the right panel reports MCQ scores on a 1-3 scale.

(score 3), demonstrating strong contextual alignment. For OEQ, a larger spread was observed across the 3-5 range, with Gemini 2.0 Flash and GPT-4.1-mini producing the most favorable judgments overall. The results suggest that the generated subset maintains reasonable factual and contextual quality, with some variability across tasks and response types.

## 4. Experimental Setup

To investigate how training data provenance impacts model adaptation, we conduct a controlled comparison of multiple training configurations. Specifically, we address the following research question:

*How do different sources of French medical instruction data (native, synthetic, translated) impact LLM performance, and do combining these sources yield complementary benefits compared to relying on a single source?*

### 4.1. Training Configurations

To systematically evaluate this question, we constructed seven dataset configurations isolating the contribution of each data origin: NAT (native only), TRAD (translated only), SYN (synthetic only), NAT-TRAD, NAT-SYN, TRAD-SYN, and ALL (all sources combined). Each configuration contains 33 493 examples, corresponding to the smallest common sample size across all sources to ensure a fair comparison. The datasets consist exclusively of multiple-choice tasks, as open-ended QA data were mainly reserved for evaluation. For mixed configurations, we maintained equal proportions of examples from each contributing subset to avoid bias toward a specific source.

To examine whether the subsets differ substantially in their semantic content, we visualized in-

struction embeddings from each data source using t-SNE (Maaten and Hinton, 2008). The resulting projection (Figure 4) provides a qualitative view of their distributional overlap. Native and translated instructions exhibit strong local mixing, suggesting similar semantic neighborhoods, while synthetic data display slightly broader dispersion, illustrating greater topical variability. Overall, the subsets appear moderately distinct yet sufficiently aligned to support joint fine-tuning without major domain shift.

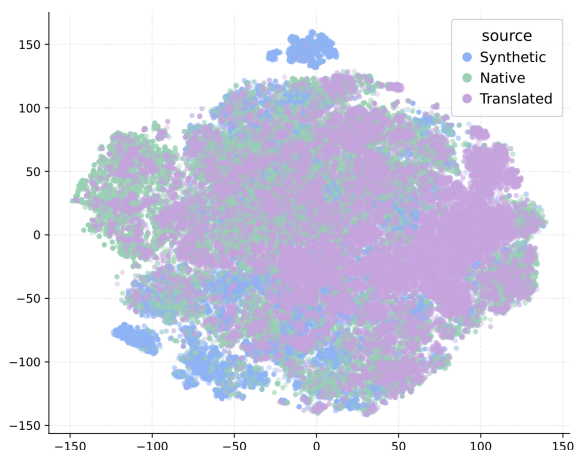


Figure 4: t-SNE visualization of instruction embeddings across the three data sources (native, translated, synthetic).

This design allows for a controlled ablation study, where performance differences can be directly attributed to the inclusion or exclusion of particular data types rather than to scale effects or task imbalances.

## 4.2. Base Model and Training Procedure

We used Qwen-4B-Instruct as the base model for all experiments. This model was selected for its multilingual pretraining, compact 4B-parameter architecture, and extended context capacity that accommodates our longest inputs ( $\approx 6\,300$  tokens) without truncation. Such coverage is particularly important for processing detailed clinical cases and multi-sentence biomedical instructions, which often require long contextual reasoning. The 4B scale offers a practical balance between representation capacity and computational cost.

We applied SFT using the DoRA (Liu et al., 2024, Weight Decomposed Low-Rank Adaptation) technique, an extension of LoRA (Hu et al., 2022) that decomposes pretrained weights into magnitude and direction components.

DoRA was chosen after preliminary comparisons with LoRA, during which it demonstrated superior stability and efficiency for small-scale biomedical fine-tuning.

Each run was trained for 10 epochs with a batch size of 12, a learning rate of  $1 \times 10^{-4}$ , and gradient accumulation over 8 steps, using a cosine learning rate scheduler with a warmup ratio of 0.05. For the DoRA configuration, the rank parameter was set to 16, the scaling factor (lora alpha) to 16, and the dropout probability to 0.05. These adapters were applied to the attention and feed-forward projection layers, specifically the following modules: *q\_proj*, *k\_proj*, *v\_proj*, *o\_proj*, *gate\_proj*, *up\_proj*, and *down\_proj*. All fine-tuning runs were conducted under identical training conditions, with the dataset configuration being the only varying factor. The hyperparameters were not explicitly tuned in this study but were chosen following standard practices for adapter-based fine-tuning in instruction-tuning setups. Preliminary monitoring of training dynamics, including loss curves and gradient norms, confirmed stable optimization without signs of divergence or overfitting, supporting the suitability of these settings for our experimental framework.

## 4.3. Evaluation Protocol

Evaluation was carried out deterministically using greedy decoding with a temperature of 0 across all QA types (MCQ, MCQU, and OEQ). For MCQ and MCQU, performance was assessed using the Exact Match (EM) accuracy, computed under two decoding modes: (1) standard greedy decoding and (2) constrained decoding, where we constrain the token vocabulary to be one of the answer choice letters (e.g., one of [“A”, “B”, “C”, “D”] for a four-choice QA dataset) and treat the answer choice with the highest token probability as the model’s prediction.

To improve robustness and account for potential position bias in multiple-choice evaluation (Pezeshkpour and Hruschka, 2024), we additionally performed three runs per configuration with randomized answer orderings. Reported scores correspond to the average performance across these runs.

For MCQ tasks that include multiple correct answers, we computed the Hamming score to capture partial correctness.

The *Exact Match (EM)* metric measures the proportion of predictions that exactly match the gold answer set:

$$EM = \frac{1}{N} \sum_{i=1}^N [\hat{y}_i = y_i]$$

where  $N$  is the total number of questions,  $y_i$  the gold answer (or answer set), and  $\hat{y}_i$  the model prediction. The *Hamming score* quantifies partial agreement between predicted and reference label

sets:

$$\text{Hamming Score} = \frac{1}{N} \sum_{i=1}^N \frac{|y_i \cap \hat{y}_i|}{|y_i \cup \hat{y}_i|}$$

Evaluation for OEQ required a distinct approach, as the outputs are free-form text rather than discrete labels. To capture both linguistic similarity and factual correctness, we employed a complementary evaluation protocol combining: (1) automatic text-based metrics to quantify lexical and semantic overlap, and (2) a human and LLM-based evaluation to assess factual accuracy and clinical soundness. Sections 4.3.1 and 4.3.2 describe each component in detail.

#### 4.3.1. Automatic Evaluation

We established automatic baselines using standard text similarity metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Banerjee and Lavie, 2005), and BERTScore (Zhang\* et al., 2020). These metrics quantify lexical and semantic overlap between model outputs and gold references. However, they remain limited in assessing factual correctness and reasoning adequacy, especially in the biomedical domain, where clinically valid answers may differ linguistically from references (Yim et al., 2025).

#### 4.3.2. Human and LLM-as-a-Judge Evaluation

To overcome the limitations of automatic metrics, we explored an LLM-as-a-judge approach to automatically assess the quality of generated answers.

To identify the most reliable model for this purpose, we conducted a meta-evaluation study on a sample of 100 OEQ pairs drawn from the test set. Each candidate LLM judge was prompted with the same evaluation instructions as those provided to a human medical expert. The full evaluation prompt provided to the LLM judges is reproduced in Appendix 13.3.

The study involved five LLMs chosen to represent a diversity of characteristics: three general-purpose models (GPT-5 OpenAI, 2025b, GPT-4.1-mini OpenAI, 2025a and Gemini 2.0 Flash Goo) and two biomedical models (MedGemma-27B Sellergren et al., 2025 and HuatuoGPT-o1-72B Chen et al., 2025).

A licensed physician independently rated the same 100 samples, assigning binary scores (1 for equivalent, 0 for non-equivalent) based on whether the generated response conveyed the essential medical information required by the reference. Minor lexical or stylistic differences were considered acceptable as long as the medical content was accurate and complete. We then computed Pearson correlations between human ratings and the scores

		pearson_r	p_value
LLMs	MedGemma-27B	0.61	$1.70 \times 10^{-11}$
	GEMINI-FLASH-2.0	0.57	$7.33 \times 10^{-10}$
	GPT-4.1-mini	0.49	$1.80 \times 10^{-7}$
	HuatuoGPT-o1-72B	0.47	$7.78 \times 10^{-7}$
	GPT-5	0.38	$9.56 \times 10^{-5}$
Metrics	ROUGE-2	0.36	$2.22 \times 10^{-4}$
	ROUGE-L	0.30	$2.39 \times 10^{-3}$
	ROUGE-1	0.29	$3.48 \times 10^{-3}$
	METEOR	0.23	$2.09 \times 10^{-2}$
	BERTScore-F1	0.17	$9.49 \times 10^{-2}$
	BLEU	0.02	$8.06 \times 10^{-1}$

Table 4: Pearson correlation coefficients ( $r$ ) and associated significance levels ( $p$ ) between human expert judgments and automated evaluation metrics on 100 sampled OEQ responses.

produced by each LLM judge and by the automatic metrics.

As shown in Table 4, MedGemma-27B, the smallest of the biomedical models, exhibited the strongest alignment with expert judgments ( $r = 0.61$ ,  $p = 1.74 \times 10^{-11}$ ), outperforming both larger general models and the larger medical LLM. This suggests that domain specialization plays a more decisive role than model scale in capturing factual correctness.

Traditional automatic metrics, by contrast, showed only weak to moderate correlations with human evaluation, with ROUGE-2 reaching the highest value among them ( $r \approx 0.36$ ).

These findings underscore the limited reliability of lexical and embedding-based similarity measures for biomedical QA, where valid responses often employ paraphrase or alternate formulations that reduce surface overlap, while linguistically similar but clinically incorrect outputs may receive inflated scores.

## 5. Results

In this section, we present the results obtained from the experimental configurations described in Section 4. For multiple-choice tasks (MCQ and MCQU), we report the scores from a single evaluation run. Results averaged over the three randomized runs (used to account for potential position bias) are provided in Appendix 13.4.

### 5.1. Multiple-Choice Question Answering

**General trends** Table 5 summarizes the performance of all configurations on MCQ and MCQU tasks, under both greedy and constrained decoding. Overall, the results reveal consistent behavior across decoding strategies for all fine-tuned

Model	MCQ				MCQU		Aggregation	
	Greedy		Constrained		Greedy	Constrained	Greedy	Constrained
	EM	Hamming	EM	Hamming	EM	EM	EM	EM
QWEN-4B (base)	15.92	43.78	18.44	53.42	53.14	56.17*	34.53	37.31
QWEN-4B-NAT	<b>27.5</b>	<b>59.73</b>	<b>27.5</b>	<b>59.82</b>	53.49	53.67	40.5	40.59
QWEN-4B-TRAD	18.56	53.08	18.0	52.8	54.65	54.88	36.6	36.44
QWEN-4B-SYN	13.3	49.79	12.73	49.94	46.46	46.73	29.88	29.73
QWEN-4B-NAT-TRAD	25.65	59.21*	25.65	59.33*	<b>56.94</b>	<b>57.1</b>	<b>41.29</b>	<b>41.37</b>
QWEN-4B-NAT-SYN	24.31	58.82	24.31	58.96	53.96	54.19	39.14	39.25
QWEN-4B-TRAD-SYN	18.69	52.98	18.38	52.95	53.63	53.99	36.16	36.18
QWEN-4B-ALL	25.95*	58.59	25.98*	58.77	55.78*	55.95	40.86*	40.97*

Table 5: Performance of all fine-tuning configurations on MCQ and MCQU under greedy and constrained decoding. Results are reported as Exact Match (EM) and Hamming scores. The aggregated columns report the average of EM scores across MCQ and MCQU formats. Since both settings correspond to MCQA and differ only in the number of correct options (single-answer vs. multiple-answer), we treat them as complementary variants of the same task. The aggregation therefore provides a compact summary of overall MCQA performance. **Bold** values indicate the best-performing models, while asterisks (\*) denote the second-best results within each column.

models, except for the base model. The approximately three-point gap between greedy (34.53) and constrained (37.31) decoding observed in QWEN-4B suggests that, although the base model was instruction-tuned, it was not explicitly optimized to output valid choice labels. After fine-tuning, however, this discrepancy nearly disappears, indicating that exposure to task-specific supervision helped models produce the expected answer format.

**Task-level comparison** Performance differences between MCQ and MCQU also reflect the relative difficulty of the tasks. Scores on MCQ are systematically lower than those on MCQU, as answering a single-correct question requires identifying one exact label, whereas MCQ allows partial correctness. This is further supported by the higher Hamming scores compared to EM values, demonstrating that the models often predict partially correct answer sets even when the EM criterion is not satisfied.

**Effect of individual data sources** When training on individual data sources, clear distinctions emerge. Models fine-tuned on native data (QWEN-4B-NAT) achieve the largest improvement in constrained decoding relative to the base model (40.59 EM vs. 37.31), underscoring the importance of linguistic alignment and exposure to authentic French medical discourse. In contrast, the synthetic-only configuration (QWEN-4B-SYN) yields the weakest results (29.73 EM), confirming that automatically generated data, though diverse, introduce stylistic variability and factual noise. The translated-only model (QWEN-4B-TRAD) shows a modest apparent gain under greedy decoding (36.6 EM) compared with the base model (34.53), yet this advan-

tage disappears once decoding is constrained to valid answer labels. This pattern indicates that the observed improvement results from the decoding strategy difference rather than genuine progress in answer accuracy.

**Impact of mixed data configurations** Combining multiple data types consistently improves performance compared to single-source setups. The NAT-TRAD configuration achieves the highest overall accuracy (41.29 greedy, 41.37 constrained), confirming the strong complementarity between linguistically aligned native supervision and the conceptual diversity introduced by translated English datasets. The fully mixed ALL model performs comparably (40.86 greedy, 40.97 constrained), suggesting that heterogeneous supervision enhances robustness without sacrificing linguistic coherence. The NAT-SYN configuration also yields competitive results (39.25 constrained), indicating that augmenting native data with synthetic examples can inject stylistic and task diversity while preserving factual grounding. In contrast, the TRAD-SYN configuration provides only marginal improvements over TRAD alone (36.55 vs. 36.44 constrained), suggesting limited synergy between two non-native sources. Overall, these trends highlight that native supervision plays a stabilizing role, anchoring learning when paired with noisier or cross-lingual data, and remains essential for achieving consistent gains in French medical instruction tuning.

## 5.2. Open-Ended Question Answering

Table 6 presents the results for OEQ, evaluated using both automatic metrics and the LLM-as-a-judge framework. Overall, traditional text similarity met-

Model	OEQ				
	BLEU	ROUGE-2	METEOR	BERTScore_F1	LLM-as-a-Judge
<b>QWEN-4B (base)</b>	12.08	20.82	<b>30.78</b>	27.22	<b>0.36</b>
<b>QWEN-4B-NAT</b>	31.85*	22.67	25.58	33.44	0.24
<b>QWEN-4B-TRAD</b>	29.48	21.42	26.49	27.39	0.22
<b>QWEN-4B-SYN</b>	26.73	20.81	24.79	27.54	0.31*
<b>QWEN-4B-NAT-TRAD</b>	<b>32.09</b>	23.05*	25.90	35.60*	0.23
<b>QWEN-4B-NAT-SYN</b>	30.63	22.36	23.73	21.65	0.34
<b>QWEN-4B-TRAD-SYN</b>	28.07	21.94	25.77	34.98	0.23
<b>QWEN-ALL</b>	31.74	<b>24.74</b>	29.36*	<b>39.67</b>	0.25

Table 6: Results on OEQ evaluated with automatic text similarity metrics and the LLM-as-a-judge (MedGemma-27B) framework. **Bold** numbers mark the best results, and asterisks (\*) indicate the second-best within each column.

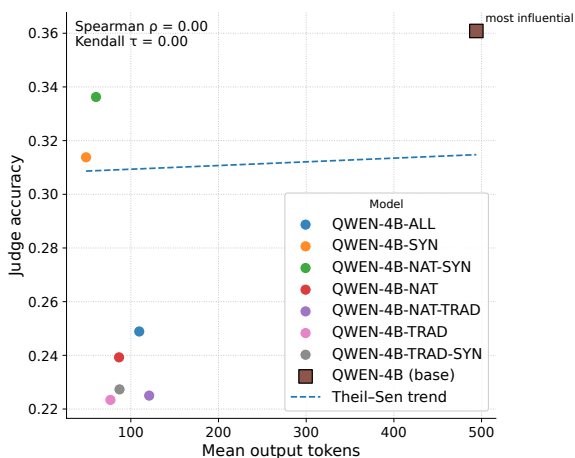


Figure 5: Relationship between mean output length and LLM-judge accuracy across fine-tuned models. Each point represents one configuration, with the dashed line indicating the best-fit linear trend.

rics (BLEU, ROUGE, METEOR, BERTScore) exhibit moderate sensitivity to fine-tuning, with BLEU exhibiting the largest improvement (+20 points) and other metrics such as ROUGE-2 and BERTScore rising more modestly (+3-8 points). However, as discussed in Section 4.3, these metrics are restricted to surface-level overlap.

A notable observation from the LLM-as-a-judge evaluation is that the base model achieves the highest mean score (0.36), surpassing all fine-tuned variants. Because the base model also produces substantially longer responses, we examined whether this effect could reflect a potential verbosity bias, i.e., the tendency of evaluators to favor longer outputs (Ye et al., 2025; Saito et al., 2023).

Figure 5 presents the association between mean output length (log scale) and judge-assigned accuracy. Using non-parametric correlation measures, both Spearman's  $\rho$  and Kendall's  $\tau$  equal 0.00 ( $p = 1$ ), indicating no consistent monotonic rela-

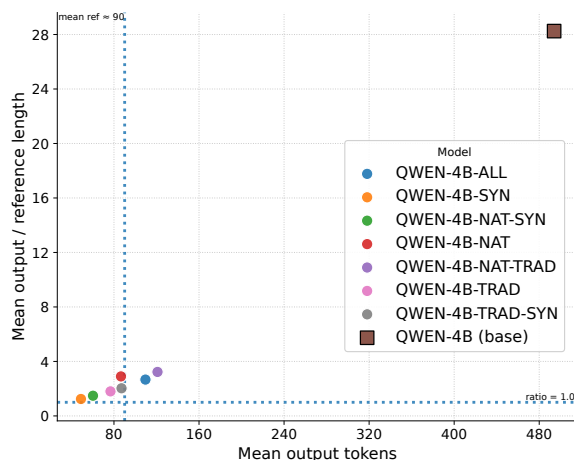


Figure 6: Comparison between average output length and reference length for each model. The vertical dashed line marks the mean reference length ( $\approx 90$  tokens), and the horizontal line represents a perfect length ratio of 1.

tionship across configurations. A leave-one-out sensitivity analysis further shows that the apparent association is unstable and largely influenced by a single configuration.

Figure 6 complements this analysis by comparing the ratio between model output length and reference length. The base model stands out clearly, producing responses approximately 28 times longer than the reference answers, whereas fine-tuned models remain much closer to the target length (generally between 1 and 3.5 times the reference length). Although this visual contrast highlights a substantial verbosity difference, it does not translate into a systematic rank-based correlation across models.

Taken together, these findings suggest that while the base model's higher score coincides with markedly longer outputs, the overall evidence for a

generalized verbosity bias remains limited. The observed pattern appears configuration-specific rather than indicative of a robust relationship between response length and evaluation score.

This pattern suggests that instruction tuning on MedInjection-FR encourages models to produce more concise outputs compared to the base model. From a linguistic perspective, models trained on native or mixed data display stronger lexical alignment and semantic adequacy, as evidenced by their higher ROUGE-2 and BERTScore values. In contrast, purely synthetic or translated data tend to underperform, likely due to minor stylistic inconsistencies or factual drift introduced during generation and translation.

## 6. Discussion

The results presented in Section 5 provide insights into how the nature and provenance of instruction data affect the adaptation of French medical LLMs. Our discussion primarily draws on the MCQ and MCQU results.

**Complementarity of data sources.** The comparison across training configurations confirms that native French data remain the most effective for improving model adaptation, particularly under constrained decoding. However, when native data are limited, combining them with other sources, translated or synthetic, proves particularly valuable. Despite containing only half or one-third of the native examples, mixed configurations such as NAT-TRAD, NAT-SYN, and ALL achieve comparable or even slightly higher aggregated scores. This suggests that heterogeneous supervision may introduce complementary information that enhances generalization. When used alone, both synthetic and translated datasets remain less effective than native data, probably due to noise, translation artifacts, or stylistic inconsistencies. Nevertheless, these data types likely contribute additional variation and conceptual diversity that help models better generalize to unseen formulations. Native supervision, however, appears essential to anchor learning and mitigate noise from these less controlled sources. Such findings are especially relevant in low-resource contexts where large-scale native instruction data are difficult to obtain.

**Insights from open-ended QA.** The OEQ evaluation highlights several limitations of current automatic evaluation methods. Standard overlap-based metrics (e.g., BLEU, ROUGE, BERTScore) exhibit only moderate sensitivity to fine-tuning, reflecting their limited ability to capture factual or clinical correctness in biomedical settings.

Regarding potential verbosity effects, our non-parametric analysis did not reveal a consistent monotonic relationship between response length and judge-assigned accuracy. Although the base model generates substantially longer outputs and receives higher scores, this pattern does not generalize across fine-tuned configurations, and rank-based correlation measures (Spearman's  $\rho$  and Kendall's  $\tau$ ) do not indicate a robust association.

Taken together, these findings suggest two key takeaways: (1) automatic and LLM-based metrics may not reliably reflect factual precision in biomedical reasoning, and (2) evaluation of medical QA tasks benefits from domain-aware or human-in-the-loop assessment to ensure that judgments are grounded in clinical relevance rather than surface similarity or isolated configuration effects.

## 7. Conclusion

This work introduces MedInjection-FR, a comprehensive French biomedical instruction dataset designed to investigate how the provenance of supervision data affects LLM adaptation. Through a controlled series of fine-tuning experiments, we showed that while native data provide the most reliable linguistic and contextual alignment with French medical reasoning, combining them with translated or synthetic sources can achieve comparable and sometimes superior performance, even with a smaller share of native examples. This finding suggests that heterogeneous training data can represent a viable alternative when authentic French resources are scarce. Moreover, our evaluation highlights that automatic and LLM-based metrics, although most correlated with human judgments, can still be influenced by verbosity bias. Together, these results deepen our understanding of how data composition shapes instruction-tuned biomedical LLMs and underscore the need for more robust, domain-aware evaluation methods in future multilingual adaptation work.

## 8. Limitations

While the proposed study isolates the effects of data provenance under controlled conditions, several limitations remain.

First, thematic coverage and task difficulty may vary across data sources, potentially influencing performance independently of linguistic origin. Although we control for sample size and task format, fully disentangling content effects from provenance effects remains challenging.

Second, synthetic data quality was assessed using a numerical scoring scheme (Section 3.4) rather than qualitative or adjectival judgments (e.g., Likert-style descriptors). While numerical scores

facilitate aggregation and comparison across evaluators, they may compress nuanced distinctions in answer quality. Moreover, although we report aggregate quality statistics, synthetic instances were not filtered based on these scores prior to fine-tuning. As a result, residual quality variation may have influenced downstream performance. A more detailed qualitative error analysis or selective filtering strategy could provide further insight into the relationship between generation quality and model adaptation.

Third, the evaluation relies partly on automatic and LLM-based scoring, both of which can misrepresent factual correctness or penalize concise yet accurate outputs. Although we calibrate the LLM-as-a-judge using human annotations, such approaches remain imperfect proxies for expert evaluation.

Finally, our experiments are limited to a single backbone model (Qwen-4B-Instruct). Scaling to larger architectures or models with different pre-training distributions may reveal distinct adaptation dynamics or different sensitivity to data provenance.

## 9. Ethical Considerations

This work introduces MedInjection-FR, a French biomedical instruction dataset composed of native, translated, and synthetically generated data. All source materials were drawn from publicly available biomedical datasets and educational resources. No personally identifiable health information (PHI) was used, and no patient-level private data were collected or generated.

The released dataset and fine-tuned models are provided solely for research purposes in natural language processing and machine learning.

Finally, the observed sensitivity to positional bias in multiple-choice evaluation (Appendix 13.4) highlights broader methodological concerns in benchmarking medical LLMs. Future work should prioritize evaluation protocols that reduce unintended biases and better reflect clinically meaningful reasoning.

## 10. Acknowledgements

This work was financially supported by ANR MALADES (ANR-23-IAS1-0005). It was provided with computing HPC and storage resources by GENCI at IDRIS thanks to the grants 2025-AD011015256R1 and 2025-AD011016540 on the supercomputer Jean Zay's H100 partition.

## 11. Bibliographical References

Gemini 2.0 flash – generative ai on vertex ai. <https://cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/2-0-flash>.

Satanjeev Banerjee and Alon Lavie. 2005. **ME-TEOR: An automatic metric for MT evaluation with improved correlation with human judgments**. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.

Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, and Benyou Wang. 2025. **Towards medical complex reasoning with LLMs through medical verifiable problems**. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 14552–14573, Vienna, Austria. Association for Computational Linguistics.

Clement Christophe, Tathagata Raha, Svetlana Maslenkova, Muhammad Umar Salman, Praveenkumar Kanithi, Marco AF Pimentel, and Shadab Khan. 2024. **Beyond fine-tuning: Unleashing the potential of continuous pretraining for clinical LLMs**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10549–10561, Miami, Florida, USA. Association for Computational Linguistics.

Xudong Han, Junjie Yang, Tianyang Wang, Ziqian Bi, Junfeng Hao, and Junhao Song. 2025. **Towards alignment-centric paradigm: A survey of instruction tuning in large language models**. *arXiv preprint arXiv:2508.17184*.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. **Lora: Low-rank adaptation of large language models**. *ICLR*, 1(2):3.

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. **Gpt-4o system card**. *arXiv preprint arXiv:2410.21276*.

Yungi Kim and Chanjun Park. 2024. **Instatrans: An instruction-aware translation framework for non-english instruction datasets**. *arXiv preprint arXiv:2410.01512*.

Suhas Kotha, Jacob Mitchell Springer, and Aditi Raghunathan. 2024. **Understanding catastrophic forgetting in language models via implicit inference**. In *The Twelfth International Conference on Learning Representations*.

- Yiming Liang, Tianyu Zheng, Xinrun Du, Ge Zhang, Jiaheng Liu, Xingwei Qu, Wenqiang Zu, Xingrun Xing, Chujie Zheng, Lei Ma, et al. 2025. Aligning instruction tuning with pre-training. *arXiv preprint arXiv:2501.09368*.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. 2024. Dora: Weight-decomposed low-rank adaptation. In *Forty-first International Conference on Machine Learning*.
- Renze Lou, Kai Zhang, and Wenpeng Yin. 2024. Large language model instruction following: A survey of progresses and challenges. *Computational Linguistics*, 50(3):1053–1095.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- Mariana Neves, Cristian Grozea, Philippe Thomas, Roland Roller, Rachel Bawden, Aurélie Névéol, Steffen Castle, Vanessa Bonato, Giorgio Maria Di Nunzio, Federica Vezzani, Maika Vicente Navarro, Lana Yeganova, and Antonio Jimeno Yepes. 2024. **Findings of the WMT 2024 biomedical translation shared task: Test sets on abstract level**. In *Proceedings of the Ninth Conference on Machine Translation*, pages 124–138, Miami, Florida, USA. Association for Computational Linguistics.
- OpenAI. 2025a. Introducing gpt-4.1 in the api. <https://openai.com/index/gpt-4-1/>.
- OpenAI. 2025b. Introducing gpt-5. <https://openai.com/index/introducing-gpt-5/>.
- 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.
- Pouya Pezeshkpour and Estevam Hruschka. 2024. **Large language models sensitivity to the order of options in multiple-choice questions**. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2006–2017, Mexico City, Mexico. Association for Computational Linguistics.
- Keita Saito, Akifumi Wachi, Koki Wataoka, and Youhei Akimoto. 2023. **Verbosity bias in preference labeling by large language models**. *ArXiv*, abs/2310.10076.
- Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo Kohlberger, Shawn Xu, Fayaz Jamil, Cían Hughes, Charles Lau, et al. 2025. Medgemma technical report. *arXiv preprint arXiv:2507.05201*.
- Sumuk Shashidhar, Clémentine Fourier, Alina Lozovskaya, Thomas Wolf, Gokhan Tur, and Dilek Hakkani-Tür. 2025. **Yourbench: Dynamic evaluation set generation with LLMs**. In *Second Conference on Language Modeling*.
- Zhang Shengyu, Dong Linfeng, Li Xiaoya, Zhang Sen, Sun Xiaofei, Wang Shuhe, Li Jiwei, Runyi Hu, Zhang Tianwei, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*.
- Qwen Team. 2025. **Qwen3 technical report**.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. **Self-instruct: Aligning language models with self-generated instructions**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh V Chawla, and Xiangliang Zhang. 2025. **Justice or prejudice? quantifying biases in LLM-as-a-judge**. In *The Thirteenth International Conference on Learning Representations*.
- Wen-wai Yim, Asma Ben Abacha, Zixuan Yu, Robert Doerning, Fei Xia, and Meliha Yetisgen. 2025. Morqa: Benchmarking evaluation metrics for medical open-ended question answering. *arXiv preprint arXiv:2509.12405*.
- Weihao Zeng, Can Xu, Yingxiu Zhao, Jian-Guang Lou, and Weizhu Chen. 2024. **Automatic instruction evolving for large language models**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6998–7018, Miami, Florida, USA. Association for Computational Linguistics.
- Tianyi Zhang\*, Varsha Kishore\*, Felix Wu\*, Kilian Q. Weinberger, and Yoav Artzi. 2020. **Bertscore: Evaluating text generation with bert**. In *International Conference on Learning Representations*.

Xinlu Zhang, Chenxin Tian, Xianjun Yang, Lichang Chen, Zekun Li, and Linda Ruth Petzold. 2023. Alpacare: Instruction-tuned large language models for medical application. *arXiv preprint arXiv:2310.14558*.

## 12. Language Resource References

corpora: Parallel corpora for the biomedical domain. <https://github.com/biomedical-translation-corpora/corpora>.

S-editions. <https://s-editions.fr/>.

Adrien Bazoge. 2025. Mediqal: A french medical question answering dataset for knowledge and reasoning evaluation. *arXiv preprint arXiv:2507.20917*.

Tianyu Han, Lisa C Adams, Jens-Michalis Pappoannou, Paul Grundmann, Tom Oberhauser, Alexander Löser, Daniel Truhn, and Keno K Bressemer. 2023. Medalpaca—an open-source collection of medical conversational ai models and training data. *arXiv preprint arXiv:2304.08247*.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multi-task language understanding. *arXiv preprint arXiv:2009.03300*.

Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.

Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. PubMedQA: A dataset for biomedical research question answering. 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 2567–2577, Hong Kong, China. Association for Computational Linguistics.

Zakaria Kaddari and Toumi Bouchentouf. 2022. Frbmedqa: the first french biomedical question answering dataset. *IAES International Journal of Artificial Intelligence*, 11(4):1588.

Yanis Labrak, Adrien Bazoge, Richard Dufour, Mickael Rouvier, Emmanuel Morin, Béatrice Daille, and Pierre-Antoine Gourraud. 2023a.

Frenchmedmcqa: A french multiple-choice question answering dataset for medical domain. *arXiv preprint arXiv:2304.04280*.

Yanis Labrak, Adrien Bazoge, Oumaima El Khettari, Mickael Rouvier, Pacome Constant Dit Beaufils, Natalia Grabar, Béatrice Daille, Solen Quiniou, Emmanuel Morin, Pierre-Antoine Gourraud, and Richard Dufour. 2024. DrBenchmark: A large language understanding evaluation benchmark for French biomedical domain. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5376–5390, Torino, Italia. ELRA and ICCL.

Yanis Labrak, Mickael Rouvier, and Richard Dufour. 2023b. MORFITT : Un corpus multi-labels d'articles scientifiques français dans le domaine biomédical. In *Actes de CORIA-TALN 2023. Actes de l'atelier "Analyse et Recherche de Textes Scientifiques" (ARTS)@TALN 2023*, pages 66–70, Paris, France. ATALA.

Itay Manes, Naama Ronn, David Cohen, Ran Ilan Ber, Zehavi Horowitz-Kugler, and Gabriel Stanovsky. 2024. K-QA: A real-world medical Q&A benchmark. In *Proceedings of the 23rd Workshop on Biomedical Natural Language Processing*, pages 277–294, Bangkok, Thailand. Association for Computational Linguistics.

Mlabonne. 2024a. medical-cases-fr: Dataset by mlabonne. <https://huggingface.co/datasets/mlabonne/medical-cases-fr>.

Mlabonne. 2024b. medical-mqca-fr: Dataset by mlabonne. <https://huggingface.co/datasets/mlabonne/medical-mqca-fr>.

Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 248–260. PMLR.

Timotheé Poulain and Victor Connes. 2021. DEFT 2021: Évaluation automatique de réponses courtes, une approche basée sur la sélection de traits lexicaux et augmentation de données (DEFT 2021 : Automatic short answer grading, a lexical features selection and data augmentation based approach). In *Actes de la 28e Conférence sur le Traitement Automatique des Langues Naturelles. Atelier DÉfi Fouille de Textes (DEFT)*, pages 31–40, Lille, France. ATALA.

Hieu Tran, Zhichao Yang, Zonghai Yao, and Hong Yu. 2024. Bioinstruct: instruction tuning of

large language models for biomedical natural language processing. *Journal of the American Medical Informatics Association*, 31(9):1821–1832.

Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290.

Yuxin Zuo, Shang Qu, Yifei Li, Zhangren Chen, Xuekai Zhu, Ermo Hua, Kaiyan Zhang, Ning Ding, and Bowen Zhou. 2025. Medxpertqa: Benchmarking expert-level medical reasoning and understanding. *arXiv preprint arXiv:2501.18362*.

## 13. Appendices

### 13.1. Synthetic Data Generation

This appendix reports the exact system prompts used to generate the synthetic instruction–response pairs. We provide separate prompts for each source corpus (MORFITT abstracts, DEFT clinical cases, and DiaMed cases) and for each task format (OEQA and MCQA), in order to ensure full transparency and reproducibility of the generation process. The OEQA generation prompts are provided in Figures 7–9, and the MCQA generation prompts are provided in Figures 10 and 11.

### 13.2. Synthetic Data Evaluation

Here we report the exact evaluation prompts used to assess the quality of the synthetic data. We provide separate prompts for OEQA and MCQA formats, detailing the scoring criteria given to the evaluator models. The OEQA and MCQA evaluation prompts are shown in Figures 12 and 13, respectively.

### 13.3. Evaluation Prompt provided to LLM Judges

Figure 14 presents the prompt provided to LLM judges to assess equivalence between candidate answers and the ground truth.

### 13.4. Position Bias in MCQ Evaluation

To assess the impact of potential position bias in multiple-choice evaluation (Pezeshkpour and Hruschka, 2024), we conducted three additional evaluation runs for each configuration, randomly permuting the answer choice ordering at each run. Reported values in Table 7 correspond to the macro-average performance across datasets and across

the three permutations, while standard deviations quantify variability across runs.

**General Trends** Compared to the single-run results reported in Table 5, overall EM scores decrease across all models and decoding modes when answer positions are randomized. For example, the base model’s aggregated constrained EM drops from 37.31 (single run) to 23.20 under permutation, and similar reductions are observed for all fine-tuned variants. This indicates that fixed answer ordering can inflate performance estimates, likely due to positional priors learned during instruction tuning.

Importantly, however, the relative ranking of configurations remains largely stable. Fine-tuned models continue to outperform the base model, and configurations including native supervision consistently rank among the strongest setups.

**Task-Level Comparison** As in the main analysis in Section 5, MCQ scores remain systematically lower than MCQU scores. The higher Hamming scores relative to Exact Match persist across permutations, indicating that models frequently predict partially correct answer sets even when failing the strict EM criterion. Randomizing answer order does not alter this structural pattern.

**Effect of Individual Data Sources** Under permutation, native-only fine-tuning (QWEN-4B-NAT) remains the strongest single-source configuration. Synthetic-only models (QWEN-4B-SYN) continue to yield the lowest EM scores across both MCQ and MCQU. Translated-only setups (QWEN-4B-TRAD) show intermediate performance. Although absolute scores decrease compared to the single-run setting, the hierarchy among individual data sources remains consistent.

**Impact of Mixed Data Configurations** Mixed configurations (NAT-TRAD, NAT-SYN, and ALL) continue to outperform most single-source models across randomized runs. NAT-TRAD and ALL remain among the top-performing configurations in aggregated EM, indicating that their advantage is not an artifact of fixed answer ordering.

**Variance Across Runs** Standard deviations across the three permutations are substantial, indicating that answer reordering can markedly affect absolute EM scores. This confirms the presence of non-negligible position sensitivity in multiple-choice evaluation. Importantly, however, variability remains comparable across configurations, and no model exhibits disproportionately higher instability relative to others. Thus, while position bias significantly impacts absolute accuracy, it does not

### System prompt (Morfit Corpus)

You are a medical expert tasked with generating complex instruction-output pairs from a French biomedical abstract. The pairs should reflect hospital-relevant tasks and be sophisticated enough that only a doctor could answer them accurately, avoiding simple or direct questions (e.g., about sample size or obvious facts).

From this abstract, generate 6 instruction-output pairs for the following 3 tasks:

1. Clinical Note Summarization (task 1): 1 pair, based strictly on the abstract.
2. Factual Medical Question Answering (task 2): 4 pairs, questions must be standalone, general medical queries inspired by the abstract's content but phrased without any reference to the study itself; answers must reflect knowledge from the abstract.
3. Specialty Classification (task 3): 1 pair, instructing to classify the text, with the output listing the specialties provided.

Follow these instructions:

1. Format each pair as a JSON object with three fields:
  - "task": An integer (1-3) corresponding to the task number above.
  - "instruction": The instruction or question in French.
  - "output": A concise, factual answer based solely on the text (no speculation beyond what's implied).
3. Ensure questions require clinical reasoning, synthesis, or expertise (e.g., differential diagnosis, treatment rationale), not basic recall.
4. Return the pairs in a JSON list, e.g.:
 

```
[
  { "task": , "instruction": "", "output": "" },
]
```

Figure 7: System prompt used to generate OEQA synthetic instruction-response pairs from the MORFIT abstract corpus.

Model	QCM				QCMU		Aggregation	
	Greedy		Constrained		Greedy	Constrained	Greedy	Constrained
	EM	Hamming	EM	Hamming	EM	EM	EM	EM
QWEN-4B (base)	10.76 <sup>4.43</sup>	37.09 <sup>5.84</sup>	12.72 <sup>5.10</sup>	45.10 <sup>7.25</sup>	30.98 <sup>19.18</sup>	33.68 <sup>19.64</sup>	20.87 <sup>13.92</sup>	23.20 <sup>14.35</sup>
QWEN-4B-NAT	18.27 <sup>7.95</sup>	49.20 <sup>9.27</sup>	18.27 <sup>7.95</sup>	49.23 <sup>9.32</sup>	32.80 <sup>18.03</sup>	32.86 <sup>18.14</sup>	25.53 <sup>13.93</sup>	25.56 <sup>14.00</sup>
QWEN-4B-TRAD	12.48 <sup>5.63</sup>	45.41 <sup>6.83</sup>	11.91 <sup>5.56</sup>	44.96 <sup>6.86</sup>	33.45 <sup>18.54</sup>	33.54 <sup>18.70</sup>	22.97 <sup>13.70</sup>	22.73 <sup>13.79</sup>
QWEN-4B-SYN	9.90 <sup>2.88</sup>	44.47 <sup>4.55</sup>	9.29 <sup>2.91</sup>	44.62 <sup>4.63</sup>	31.18 <sup>13.28</sup>	31.27 <sup>13.43</sup>	20.54 <sup>9.61</sup>	20.28 <sup>9.72</sup>
QWEN-4B-NAT-TRAD	16.67 <sup>7.78</sup>	49.08 <sup>8.78</sup>	16.67 <sup>7.78</sup>	49.14 <sup>8.83</sup>	33.80 <sup>20.04</sup>	33.86 <sup>20.13</sup>	25.23 <sup>15.20</sup>	25.27 <sup>15.26</sup>
QWEN-4B-NAT-SYN	15.81 <sup>7.48</sup>	48.44 <sup>8.92</sup>	15.81 <sup>7.48</sup>	48.49 <sup>9.00</sup>	32.70 <sup>18.37</sup>	32.80 <sup>18.53</sup>	24.26 <sup>14.03</sup>	24.30 <sup>14.13</sup>
QWEN-4B-TRAD-SYN	12.38 <sup>5.47</sup>	45.57 <sup>6.42</sup>	12.01 <sup>5.52</sup>	45.45 <sup>6.50</sup>	33.18 <sup>17.71</sup>	33.30 <sup>17.92</sup>	22.78 <sup>13.11</sup>	22.65 <sup>13.26</sup>
QWEN-4B-ALL	16.96 <sup>7.56</sup>	48.84 <sup>8.59</sup>	16.98 <sup>7.58</sup>	48.90 <sup>8.69</sup>	33.57 <sup>19.15</sup>	33.63 <sup>19.25</sup>	25.27 <sup>14.56</sup>	25.30 <sup>14.63</sup>

Table 7: Performance of all configurations on QCM and QCMU under three answer-order permutations to assess positional bias. Results are macro-averaged across datasets and reported as Exact Match (EM) and Hamming scores. Standard deviations (superscript) correspond to the sample standard deviation across the three answer-order runs, reflecting sensitivity to option reordering. Aggregated columns report the average EM across QCM and QCMU. Bold values indicate the best-performing models, and asterisks (\*) indicate the second-best within each column.

fundamentally alter the comparative ranking of data configurations.

invariant scoring protocols, or calibration strategies explicitly designed to reduce positional priors.

Overall, the permutation analysis confirms that the main findings reported in Section 5 are robust to answer choice randomization at the comparative level. Nevertheless, the magnitude of the observed variance highlights the need for future work to systematically account for and mitigate positional artifacts in MCQ evaluation. Potential directions include answer-order averaging, permutation-

### System prompt (DEFT Corpus)

You are a medical expert tasked with generating complex instruction-output pairs from a French clinical case report. The pairs should reflect hospital-relevant tasks and be sophisticated enough that only a doctor could answer them accurately, avoiding simple or direct questions (e.g., about age, sex, or obvious facts).

From the text, generate 15 instruction-output pairs for the following 8 tasks:

1. Clinical Note Summarization (task 1) : 2 pairs
2. Factual Medical Question Answering (task 2) : 2 pairs
3. Differential Diagnosis Support (task 3) : 2 pairs
4. Treatment Plan Suggestion (task 4): 2 pairs
5. Lab Result Interpretation (task 5): 2 pairs
6. Patient Instruction Generation (task 6): 2 pairs
7. Drug Interaction Check (task 7): 2 pairs
8. Specialty Classification (task 8): 1 pair, instructing to classify the clinical case into the relevant medical specialty (or specialties) it pertains to. Return the identified specialty or specialties in French as a list (e.g., ["Cardiologie", "Oncologie"]).

Follow these instructions:

1. Format each pair as a JSON object with three fields:
  - "task": An integer (1-8) corresponding to the task number above.
  - "instruction": The instruction or question in French.
  - "output": A concise, factual answer in French based solely on the text (no speculation beyond what's implied).
3. Ensure questions require clinical reasoning, synthesis, or expertise (e.g., differential diagnosis, treatment rationale), not basic recall.
4. Return the pairs in a JSON list, e.g.:

```
[  
  {"task": , "instruction": "", "output": ""},  
]
```

Figure 8: System prompt used to generate OEQA synthetic instruction-response pairs from DEFT clinical cases.

### System prompt (DiaMed Corpus)

You are a medical expert tasked with generating complex instruction-output pairs from a French clinical case report for fine-tuning a language model. The pairs should reflect hospital-relevant tasks and be sophisticated enough that only a doctor could answer them accurately, avoiding simple or direct questions (e.g., about age, sex, or obvious facts).

From the text, generate 16 instruction-output pairs (2 pairs per task) for the following 8 tasks:

1. Clinical Note Summarization (task 1)
2. Factual Medical Question Answering (task 2)
3. Differential Diagnosis Support (task 3)
4. Treatment Plan Suggestion (task 4)
5. Lab Result Interpretation (task 5)
6. Patient Instruction Generation (task 6)
7. Billing/ICD-10 Code Suggestion (task 7)
8. Drug Interaction Check (task 8)

Follow these instructions:

1. Format each pair as a JSON object with three fields:
  - "task": An integer (1-8) corresponding to the task number above.
  - "instruction": The instruction or question in French.
  - "output": A concise, factual answer based solely on the text (no speculation beyond what's implied).
3. Ensure questions require clinical reasoning, synthesis, or expertise (e.g., differential diagnosis, treatment rationale), not basic recall.
4. Return the pairs in a JSON list, e.g.:

```
[  
  {"task": , "instruction": "", "output": ""},  
]
```

Figure 9: System prompt used to generate OEQA synthetic instruction–response pairs from DiaMed clinical cases.

### System prompt (Morfitt Corpus)

You are a medical expert tasked with generating standalone multiple-choice questions (MCQAs) from a French biomedical abstract. These MCQAs will be used to fine-tune a language model for general medical knowledge reasoning.

#### Task:

Generate 5 MCQA pairs for Factual Medical Question Answering task.

#### Guidelines

- Each question must be standalone - do not reference the abstract, article, or any specific study.
- Questions must reflect general medical knowledge that can be inferred or derived from the abstract's content.
- Focus on factual, clinically relevant knowledge: definitions, mechanisms, risk factors, complications, disease characteristics, etc.
- Avoid trivial or surface-level recall; require some synthesis or application.
- Do not copy phrases or wording directly from the abstract.
- Use French for all content.

#### Output format:

Return a JSON list of 5 MCQA items. Each item must follow this schema:

```
[
  {
    "question": "",
    "options": {
      "A": "",
      "B": "",
      "C": "",
      "D": ""
    },
    "answer": <Letter(s) corresponding to the correct option choice>
  }
]
```

Figure 10: System prompt used to generate MCQA items from the MORFITT abstract corpus.

### System prompt (DEFT & DiaMed Corpus)

You are a medical expert tasked with generating complex, clinically relevant MCQAs from a French clinical case report. These MCQAs will be used to fine-tune a language model and must reflect real hospital tasks.

#### Requirements:

- Questions must require clinical reasoning, decision-making, or synthesis of information.
- Avoid trivial questions (e.g., patient age, gender, or obvious facts stated directly).
- Design questions only a trained medical doctor could answer accurately.
- Some questions may have more than one correct answer.

#### Tasks:

Generate 14 MCQA pairs (2 per task) across the following 7 categories:

1. Factual Medical Question Answering - questions about pathophysiology, definitions, or clinical facts requiring reasoning.
2. Differential Diagnosis Support - identify the most probable cause given a set of symptoms.
3. Treatment Plan Suggestion - suggest the best next step or therapeutic intervention.
4. Lab Result Interpretation - interpret bloodwork, imaging, or test outcomes in context.
5. Patient Instruction Generation - what a clinician would explain or recommend to the patient.
6. Billing/ICD-10 Code Suggestion - infer the most appropriate billing or diagnostic code.
7. Drug Interaction Check - identify potential interactions or contraindications in a treatment plan.

#### Output format:

Return a JSON list of 14 questions. Each item should follow this schema:

```
[
  {
    "task": ,
    "question": "",
    "options": {
      "A": "",
      "B": "",
      "C": "",
      "D": ""
    },
    "answer": <Letter(s) corresponding to the correct option choice>
  }
]
```

Figure 11: System prompt used to generate MCQA items from DEFT and DiaMed clinical cases.

#### System prompt

You are a medical evaluator tasked with assessing question-answer pairs within a given context. Provide a score from 1 to 5 based on the provided score criteria.

Do not include any other opening, closing, or explanations.

Score criteria:

Score 1: The question-answer pair is completely irrelevant or incorrect given the context. The answer has major factual errors.

Score 2: The question is somewhat relevant but the answer has significant inaccuracies or lacks important details from the context.

Score 3: The question is relevant and the answer is mostly accurate but contains some minor factual errors or omissions.

Score 4: The question is clear and relevant, and the answer is accurate based on the context with only very minor omissions.

Score 5: The question is clear, relevant, and the answer is completely accurate and comprehensive based on the given context.

Remember, your score should consider both the relevance of the context to the medical domain and the accuracy of the question-answer pair. Return only the score no other text is allowed.

Figure 12: System prompt used to automatically evaluate OEQA synthetic question-answer pairs.

#### System prompt

You are a medical evaluator tasked with assessing the correctness of multiple-choice question-answer pairs within a given context. Provide a score from 1 to 5 based on the criteria below. Do not include any other opening, closing, or explanations.

Score criteria:

Score 1: The question-answer pair is completely irrelevant or incorrect given the context. The answer is incorrect.

Score 2: The question is relevant but the answer is incorrect given the context.

Score 3: The question is clear, relevant, and the answer is completely accurate based on the given context.

Return only the score, no other text is allowed.

Figure 13: System prompt used to automatically evaluate MCQA synthetic question-answer items.

#### System prompt

You are a medical evaluator tasked with assessing whether a candidate answer is equivalent to the ground truth.

Assign a score strictly according to the criteria below. Do not include any explanations, comments, or extra text in your response.

Scoring criteria:

Score 0: Not equivalent

Score 1: Equivalent

Definition of equivalence:

Two answers are considered equivalent if the essential expected information is covered. Minor differences in wording, or additional or missing details, are acceptable as long as the candidate answer would be considered an acceptable response to the question.

Return only the score (0 or 1), nothing else.

Figure 14: System prompt provided to LLM judges to score answer equivalence against the ground truth.