

PROPELLA-1: MULTI-PROPERTY DOCUMENT ANNOTATION FOR LLM DATA CURATION AT SCALE

Maximilian Idahl^{1,2} Benedikt Droste¹ Björn Plüster¹ Jan Philipp Harries¹

¹ellamind ²Leibniz University Hannover

ABSTRACT

Since FineWeb-Edu, data curation for LLM pretraining has predominantly relied on single scalar quality scores produced by small classifiers. A single score conflates multiple quality dimensions, prevents flexible filtering, and offers no interpretability. We introduce propella-1, a family of small multilingual LLMs (0.6B, 1.7B, 4B parameters) that annotate text documents across 18 properties organized into six categories: core content, classification, quality and value, audience and purpose, safety and compliance, and geographic relevance. The models support 57 languages and produce structured JSON annotations conforming to a predefined schema. Evaluated against a frontier commercial LLM as a reference annotator, the 4B model achieves higher agreement than much larger general-purpose models. We release propella-annotations, a dataset of over three billion document annotations covering major pretraining corpora including data from FineWeb-2, FinePDFs, HPLT 3.0, and Nemotron-CC. Using these annotations, we present a multi-dimensional compositional analysis of widely used pretraining datasets, revealing substantial differences in quality, reasoning depth, and content composition that single-score approaches cannot capture. All model weights and annotations are released under permissive, commercial-use licenses.¹

1 INTRODUCTION

Data quality is among the highest-leverage factors in LLM pretraining. FineWeb-Edu matches the next-best open dataset on MMLU with approximately $8\times$ fewer training tokens through educational quality filtering (Penedo et al., 2024), OLMo 3 matches Qwen-3 32B with roughly $6\times$ fewer tokens through rigorous data curation (Team OLMo et al., 2025), and BeyondWeb reports $7.7\times$ faster training through systematic document rephrasing (Maini et al., 2025). These results establish that investing in data quality yields multiplicative, not marginal, improvements in training efficiency.

The dominant paradigm for model-based data quality filtering since 2024 has been the FineWeb-Edu pipeline: use a large LLM to annotate a sample of documents with a single scalar “educational value” score, distill those labels into a small classifier (FastText or a linear head on encoder embeddings), and threshold at scale (Penedo et al., 2024; Li et al., 2024). This approach has two fundamental limitations.

First, a single scalar score conflates multiple quality dimensions. Educational value is not the only relevant dimension for pretraining data. Content with high reasoning depth but low pedagogical structure (e.g., legal analysis, technical specifications) may receive low educational scores despite being highly valuable for training. Conversely, a document may score well on educational value while containing heavy commercial bias, outdated information, or personally identifiable information. FineWeb-Edu’s own analysis reveals significant topic bias: filtering for educational value systematically favors certain topics (education, history, science) and removes others (entertainment, business, travel) (Penedo et al., 2024).

Second, most model-based quality filtering has been developed for English only. FineWeb-Edu, DCLM, and Nemotron-CC are English-only systems (Penedo et al., 2024; Li et al., 2024; Su et al.,

¹Model weights under Apache 2.0, annotations under CC-BY-4.0. All artifacts are available at <https://hf.co/collections/ellamind/propella-1>.

2025). While recent work has extended quality filtering to multiple languages through per-language FastText classifiers (Messmer et al., 2025) or multilingual LLM-as-judge approaches (Ali et al., 2025), these still rely on single-score quality measures.

We introduce propella-1, a family of small multilingual LLMs that move beyond single-score filtering to structured multi-property annotation. Our contributions are as follows:

1. **Annotation framework.** We define 18 properties across six categories that capture complementary quality dimensions, enabling flexible, composable filtering strategies for LLM training data curation (§3).
2. **propella-1 models.** We train three decoder-only models (0.6B, 1.7B, 4B) based on Qwen-3, that annotate documents in 57 languages, achieving high agreement with frontier models at orders-of-magnitude lower cost (§4, §5).
3. **propella-annotations.** We release over three billion document annotations covering major pretraining corpora, the largest open release of multi-property annotations for LLM training data to date (§6).
4. **Dataset insights.** We demonstrate the analytical value of multi-property annotation through three case studies that reveal compositional differences across datasets, quality tiers, and languages that single-score approaches cannot capture (§7).

2 RELATED WORK

Heuristic filtering. Rule-based filtering has been a cornerstone of web-scale pretraining data curation (Ortiz Suárez et al., 2019; Raffel et al., 2020). Common heuristics include length filters, character ratio thresholds, repetition removal, and blocklists. FineWeb’s ablation study reveals that many heuristics have low or even negative impact, and that some C4 rules hurt downstream performance (Penedo et al., 2024). Moreover, Su et al. (2025) find that standard heuristic filtering removes 18% of high-quality tokens, and that disabling heuristic filters for high-quality documents improves downstream accuracy. Heuristic approaches also carry language-specific assumptions: rules requiring terminal punctuation fail for languages without sentence-final punctuation, and word-length heuristics do not apply to languages without spaces.

Perplexity-based filtering. CCNet introduced perplexity filtering using n-gram language models trained on Wikipedia (Wenzek et al., 2020). While effective for removing disfluent text across languages, perplexity filtering is biased toward Wikipedia-like content and cannot capture semantic quality: fluent nonsense scores well, while valuable but unconventional content (code, conversational text, technical writing) is penalized (Rae et al., 2021; Penedo et al., 2024).

Classifier-based quality scoring. Two paradigms have emerged for training quality classifiers. The first uses curated positive examples (Wikipedia, textbooks) paired with random web samples to train binary classifiers (Li et al., 2024; Joulin et al., 2017). The second uses LLM-as-judge annotations to train regressors on encoder embeddings, as pioneered by FineWeb-Edu (Penedo et al., 2024). Both produce single scalar quality scores.

Multilingual quality filtering. FineWeb-2-HQ trains separate FastText and MLP scorers per language using diverse multilingual training sources (Messmer et al., 2025). JQL evaluates LLM judges for educational quality across 35 languages and distills their scores into lightweight regressors (Ali et al., 2025). These extend single-score filtering to non-English languages but do not address the single-score limitation itself.

Register-based annotation. Recent work on web register classification annotates documents with hierarchical register labels (e.g., narrative, informational, instructional) rather than a single quality score (Myntti et al., 2024; 2025). HPLT 3.0 (Oepen et al., 2025) applies register annotations at web scale alongside quality and PII labels, representing the broadest annotation effort among these systems. Myntti et al. (2025) show that register-based data selection can improve or harm LLM performance, indicating that text type matters independently of quality.

Multi-dimensional quality assessment. Several recent systems assess data quality along multiple dimensions rather than a single score. QuRating (Wettig et al., 2024) annotates a 260B token corpus with scalar ratings across four quality dimensions using pairwise LLM judgments. DataMan (Peng et al., 2025) derives 14 quality criteria and annotates 447B tokens, while Meta-rater (Zhuang et al., 2025) integrates four quality dimensions with existing signals through learned weightings, demonstrating that multi-dimensional selection can double convergence speed. RedPajama V2 (Weber et al., 2024) accompanies over 100 trillion tokens of multilingual web text with quality signals and meta-data. These approaches produce scalar ratings or heuristic signals for data selection; none combines structured categorical annotations, broad property coverage (including safety, geographic relevance, and content classification), multilingual support, and a public release of large-scale annotations.

propella-1 addresses this gap with structured categorical annotation across 18 properties for 57 languages, at billion-document scale, with all annotations publicly released.

3 ANNOTATION FRAMEWORK

3.1 DESIGN PHILOSOPHY

Our annotation framework is motivated by the observation that data quality is inherently multi-dimensional. A document’s value for LLM pretraining depends on its content integrity, topical relevance, reasoning depth, audience level, safety profile, and many other factors. Reducing this to a single scalar discards information that practitioners need for flexible, task-specific curation.

We designed the framework to satisfy three requirements: (1) each property should enable a concrete filtering use case; (2) properties should be complementary, not redundant; (3) annotation values should use enumerated categories with clear rubric definitions to ensure consistency.

3.2 PROPERTIES

propella-1 annotates documents across 18 properties organized into six categories, summarized in Table 1.

Each property was chosen because it enables a concrete curation scenario. Educational value supports curriculum-style training schedules. Reasoning indicators allow selecting content rich in analytical and explanatory reasoning, relevant for training reasoning models. Commercial bias enables filtering out marketing-dominated content that may degrade model objectivity. Time-sensitivity distinguishes evergreen reference material from ephemeral news, supporting freshness-aware data selection. Content type and business sector enable domain-targeted dataset construction. Safety and PII properties support compliance-aware filtering.

Because annotations are structured as typed JSON fields, practitioners can compose arbitrary filtering predicates, for example selecting documents that simultaneously satisfy high reasoning depth, low commercial bias, and no PII. A representative annotation output is shown in Figure 6 (Appendix B).

3.3 RUBRIC DEVELOPMENT

Consistent annotation at scale requires precise operational definitions. We developed the annotation rubric through a systematic iterative process over approximately two weeks. Starting from a large set of candidate properties drawn from data curation practice, we consolidated overlapping dimensions and removed properties that did not enable distinct filtering use cases, arriving at a core set of 17. Through subsequent pilot annotation rounds, we refined definitions, adjusted category boundaries, and resolved ambiguities. The one-sentence description property was added late in the process to enable English-speaking practitioners to inspect and navigate multilingual datasets without requiring knowledge of the source language, yielding the final set of 18.

The complete rubric spans approximately 8,000 words and constitutes a comprehensive annotation manual. For each property, the rubric specifies all permissible values with detailed selection criteria, concrete positive and negative examples, and key decision points that disambiguate between adjacent categories. For ordinal properties such as content quality and content integrity, percentage-based thresholds and graduated severity levels provide annotators (and models) with clear boundaries.

Table 1: The 18 annotation properties of propella-1, organized by category. $|V|$ denotes the number of enumerated values; – indicates open vocabulary (free-text or ISO-3166 country names). Property type indicates the evaluation metric used: ordinal properties are evaluated with Quadratic Weighted Kappa (QWK), the binary property with F1, multi-select properties with Intersection-over-Union (IoU, also known as the Jaccard index), and the free-text property is excluded from quantitative evaluation.

Category	Property	Description	Type	$ V $
Core Content	Content Integrity	Completeness and technical quality	Ordinal	4
	Content Ratio	Content vs. navigation/UI ratio	Ordinal	5
	Content Length	Amount of substantive content	Ordinal	4
Classification	One-Sentence Desc.	Neutral summary of the document	Free-text	–
	Content Type	Functional structure and purpose	Multi-select	18
	Business Sector	Industry domain relevance	Multi-select	37
	Technical Content	Specialized knowledge type	Multi-select	7
Quality & Value	Content Quality	Writing and presentation quality	Ordinal	5
	Information Density	Signal vs. redundancy ratio	Ordinal	5
	Educational Value	Teaching and learning potential	Ordinal	5
	Reasoning Indicators	Logical reasoning depth	Ordinal	5
Audience & Purpose	Audience Level	Target sophistication level	Ordinal	6
	Commercial Bias	Commercial influence on objectivity	Ordinal	5
	Time-Sensitivity	Temporal value decay	Ordinal	4
Safety	Content Safety	Inappropriate or harmful content	Ordinal	5
	PII Presence	Personally identifiable information	Binary	2
Geographic	Regional Relevance	Regional/cultural context	Multi-select	14
	Country Relevance	Specific country relevance	Multi-select	–

Multi-select properties such as content type (18 possible values), business sector (37 sectors), and technical content (7 categories) include multi-label combination examples that illustrate when and how to assign multiple labels simultaneously.

Several properties receive particularly detailed treatment. Educational value and reasoning indicators include structured decision trees, automation heuristics (e.g., lexical and structural cue counts per thousand tokens), disambiguation rules that distinguish related concepts (e.g., explanatory vs. analytical reasoning, procedural steps vs. logical argumentation), and annotated borderline examples. The information density property defines common repetition patterns to watch for, such as copy-paste artifacts and circular reasoning. Safety and PII properties include explicit guidelines for distinguishing contextual mentions (e.g., public figures, academic authors) from genuine privacy concerns.

The rubric also addresses challenges specific to multilingual annotation. Language-specific quality indicators account for proper character encoding in non-Latin scripts, adjusted word-density expectations for agglutinative languages such as Turkish and Finnish, and appropriate register assessment for languages with formalized politeness systems such as Japanese and Korean. Regional and country relevance properties include explicit guidelines clarifying that language does not determine geographic relevance (e.g., Spanish-language content about Asian markets should not be classified as Latin American).

The full rubric is provided in Appendix F and released alongside the model weights.

4 PROPELLA-1 MODELS

Architecture. The propella-1 family consists of three models based on the Qwen-3 architecture (Yang et al., 2025): 0.6B, 1.7B, and 4B parameters. We chose decoder-only models over encoder-based classifiers for three reasons: (1) they handle long documents natively (up to 64K context), (2) they produce all 18 annotations as structured JSON output removing the need for per-property classifier training, and (3) through fine-tuning, they internalize detailed annotation guidelines and achieve strong performance with a short, efficient prompt.

Training data. Training data was obtained by annotating a diverse document sample using multiple frontier LLMs, prompted with the full annotation rubric and strict response schema. The training set covers 57 languages with approximately 35% English and the remainder spanning European languages, Arabic, Chinese, Japanese, Korean, Thai, and others. Dedicated support for code, mathematical content, and post-training data (instruction-following, conversations) is included. Documents were sampled from diverse sources including web crawls, PDFs, and curated datasets to ensure the models generalize across content types. Detailed breakdowns of the language and source distributions are provided in Appendix A. To address limitations of API content filters, which sometimes refused to annotate problematic documents, a small subset of documents was annotated manually.

Training setup. Models were fine-tuned with a 64K context length using fp8 mixed-precision training on $4 \times$ H100 GPUs within a few hours per model variant. The target output format is a JSON object conforming to a predefined schema with enumerated values for all categorical properties. No whitespace is included in the JSON output, to reduce the number of output tokens during inference.

Prompt design. The propella-1 models use a compact system prompt (~ 800 tokens) that lists all properties with their enumerated values and brief descriptions (see Appendix C). This prompt is significantly shorter than the full annotation prompt ($\sim 14\text{K}$ tokens) used during training data creation.

5 EVALUATION

5.1 SETUP

We evaluate on a test set of 3,000 documents. For these documents, we obtain annotations from Gemini-3-Pro with reasoning effort set to “high”. We use these as reference labels, treating Gemini-3-Pro as a high-quality proxy for this structured annotation task. We discuss the implications and limitations of this choice in §8. The one-sentence description property is excluded from quantitative evaluation.

5.2 METRICS

Properties are grouped by type, each evaluated with an appropriate metric:

- **Ordinal properties** (11 properties): Quadratic Weighted Kappa (QWK), which measures agreement while accounting for the ordinal nature of labels, penalizing larger disagreements more heavily.
- **Binary property** (1 property: PII presence): F1 score.
- **Multi-select properties** (5 properties): Intersection-over-Union (IoU), averaged across samples.

We compute an overall score as a weighted average reflecting the proportion of each property type:

$$\text{overall} = \frac{11}{17} \times \overline{\text{QWK}} + \frac{1}{17} \times \text{F1} + \frac{5}{17} \times \overline{\text{IoU}} \quad (1)$$

5.3 BASELINES

We compare against several general-purpose LLMs, including models from the Gemini (Comanici et al., 2025), Qwen (Yang et al., 2025), and Gemma (Team, 2025) families, as well as Mistral-Small (Mistral AI, 2025) and SmoLLM3 (Bakouch et al., 2025). All baseline models use the detailed annotator system prompt, which includes the complete rubric with examples, decision criteria, and language-specific guidelines. The propella-1 models use their compact prompt. We also evaluated some baseline models with the compact propella-1 prompt, which consistently led to worse performance, as the short prompt lacks the detail needed for models not specifically trained for this task.

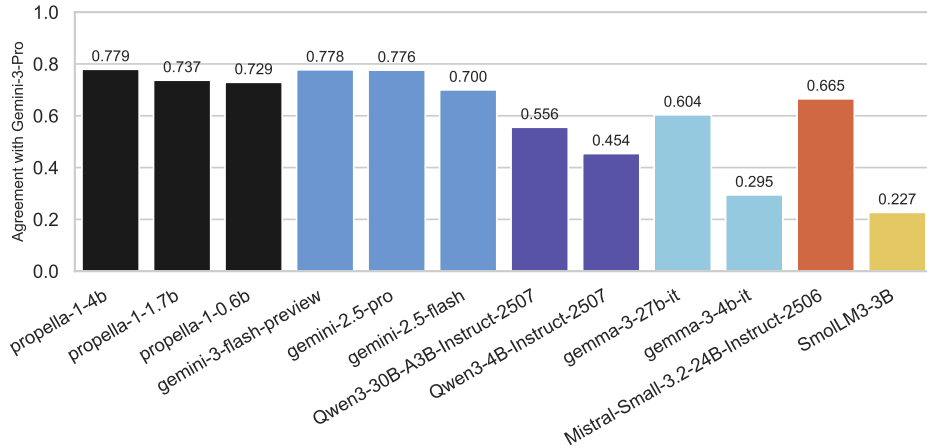


Figure 1: Overall annotation agreement scores across all evaluated models. propella-1-4b exceeds Gemini-3-Flash and significantly larger open models.

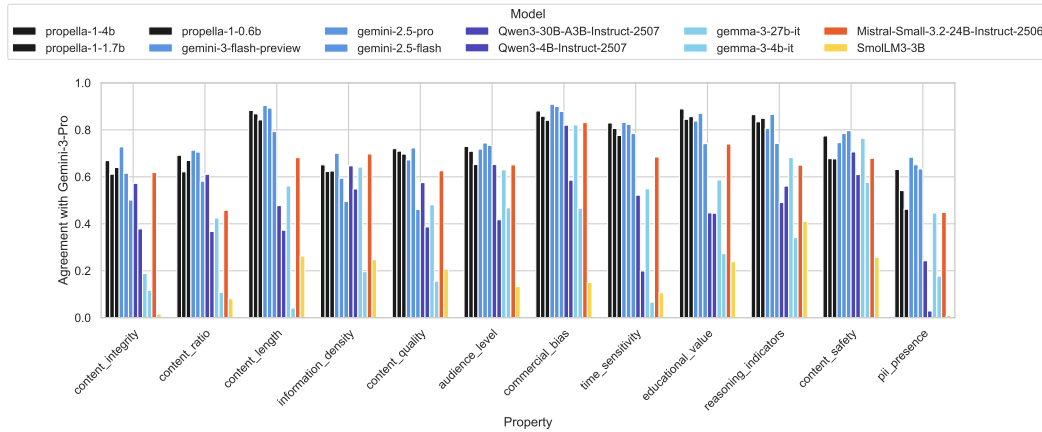


Figure 2: Per-property annotation agreement scores across all evaluated models (12 properties). See Figure 7 in Appendix D for the full breakdown of all 17 properties.

5.4 RESULTS

Figure 1 shows the overall performance comparison across all evaluated models. The 4B model achieves an overall score of 0.779, exceeding Gemini-3-Flash and all open-weight baselines despite being significantly smaller. Even the 0.6B model achieves a strong score of 0.729, demonstrating that small, specialized models can approach the annotation quality of much larger general-purpose models on this task. Figure 2 presents the per-property score breakdown, revealing that propella-1 models perform well across all properties. We also find that fp8 inference preserves annotation quality with negligible differences compared to bf16 (see Appendix E).

6 ANNOTATION AT SCALE

6.1 SERVING INFRASTRUCTURE

We serve propella-1 models using SGLang (Zheng et al., 2024) with the llguidance structured output backend (Guidance AI, 2025), enforcing strict JSON schema conformance and no whitespace during generation, guaranteeing that every output is a valid annotation object with correct enumerated values. This eliminates the need for response schema validation or retry logic.

Table 2: Inference throughput in documents per second and tokens-per-second (TPS) for propella-1 models. GPU-hours per 1M documents estimates the compute cost for large-scale annotation. Throughput is dominated by prefill, as annotation inputs are full documents while outputs are compact JSON objects.

Model	GPU	Docs/s	h / 1M docs	Prompt TPS	Output TPS
propella-1-4b	A100 80GB	10.3	27.0	19.1K	1.5K
propella-1-4b	H100 96GB	22.4	12.4	41.6K	3.2K
propella-1-4b (fp8)	H100 96GB	27.0	10.3	50.1K	3.9K
propella-1-1.7b	A100 80GB	17.8	15.6	33.0K	2.6K
propella-1-1.7b	H100 96GB	35.8	7.8	66.5K	5.2K
propella-1-1.7b (fp8)	H100 96GB	39.1	7.1	72.7K	5.7K
propella-1-0.6b	A100 80GB	21.5	12.9	40.0K	3.1K
propella-1-0.6b	H100 96GB	39.9	7.0	74.2K	5.7K

Table 2 reports throughput benchmarks. After a short warmup phase, we run inference for 5,000 documents, sending 1,000 concurrent requests to the SGLang server, to obtain a rough estimate of the GPU-hours required to annotate one million documents. On a single H100 GPU with fp8 precision, the 4B model processes 27.0 documents per second, corresponding to approximately 10.3 GPU-hours per million documents. Since annotation inputs are full documents while outputs are compact JSON objects, throughput is dominated by prefill: the 4B fp8 configuration achieves 50.1K prompt tokens per second versus 3.9K output tokens per second.

Scaling. For large-scale annotation, we scaled inference across thousands of GPUs using inference-hive,² an open-source distributed inference orchestration system for SLURM clusters. In one deployment, we ran the 4B model on 3,936 A100 GPUs and annotated ~500 million documents from FineWeb-2 in approximately 3.5 hours.

6.2 THE PROPELLA-ANNOTATIONS DATASET

We release propella-annotations, a dataset of over three billion document annotations covering major pretraining corpora. All annotations were produced with the propella-1-4b model using the serving infrastructure described above. Table 3 summarizes the dataset.

The annotated corpora span a diverse set of pretraining sources: web crawls (FineWeb-2, HPLT 3.0, Nemotron-CC, German Commons), extracted PDFs (FinePDFs), encyclopedic text (finewiki), and synthetic data (SYNTH, a collection of synthetic instruction-following and conversation data). While the propella-1 models support 57 languages, the current annotation release covers English and 14 European languages, with the largest volumes in German, Spanish, French, and Italian. The dataset will be extended with additional corpora and languages over time.

Each annotation record contains the full set of 18 property values for a single document, keyed by the document identifier from the source dataset. This enables direct joining with the original corpora for downstream filtering and selection without redistributing the source text. Because the annotations consist of categorical values and brief text descriptions rather than full documents, their small storage and memory footprint makes it practical to analyze and filter large pretraining datasets on modest hardware.

The dataset is hosted on the Hugging Face Hub,³ including a detailed dataset card with usage examples. All annotations are released under the CC-BY-4.0 license.

²<https://github.com/ellamind/inference-hive>

³<https://hf.co/datasets/openeurolm/propella-annotations>

Table 3: Overview of the propella-annotations dataset. All annotations were produced with propella-1-4b.

Source Dataset	Languages	Annotations
FineWeb-2	DE, ES, FR, IT, SV, FI	1,632,650,735
HPLT 3.0	DE, FI	694,920,477
FinePDFs	EN + 14 EU languages	365,048,869
Nemotron-CC	EN (high-quality split)	155,688,999
SYNTH	Multilingual	77,908,583
finewiki	Multilingual	43,097,138
German Commons	DE	35,716,016
Total		3,005,080,817

7 DATASET INSIGHTS

We demonstrate the analytical value of multi-property annotation through three case studies. These findings are only possible with structured multi-property annotation at scale and illustrate why a single quality score is insufficient for understanding pretraining data.

7.1 GERMAN MULTI-SOURCE PROFILING

We compare property distributions across four German-language pretraining sources: FineWeb-2 (Penedo et al., 2025a) (496M docs), FinePDFs (Penedo et al., 2025b) (36M docs), HPLT 3.0 (Oepen et al., 2025) (645M docs), and German Commons (Gienapp et al., 2025) (36M docs). By controlling for language, we isolate the effect of data source on content characteristics.

Figure 3 reveals that these pools of German text have dramatically different quality profiles. FinePDFs contains substantially more content rated as “excellent” quality (21.4% vs. 2.4% in FineWeb-2), approximately 12 \times more content with analytical reasoning indicators, and 7.4% high educational value compared to 0.6% in FineWeb-2. HPLT 3.0, as a less curated web crawl, shows higher rates of content fragments and degraded documents compared to FineWeb-2.

These differences have direct implications for data mixture design: a practitioner seeking reasoning-heavy training data should prioritize FinePDFs, while one optimizing for scale with acceptable quality should consider FineWeb-2. This information is difficult to capture with single-score quality classifiers.

7.2 NEMOTRON-CC QUALITY TIER AUDIT

Nemotron-CC (Su et al., 2025) classifies documents into quality tiers using a classifier ensemble. We analyze 10,000 documents from each of the five quality tiers using propella-1 annotations to audit their multi-dimensional quality profiles. Specifically, we examine whether documents in the “high” quality tier are uniformly high-quality across all dimensions.

Figure 4 shows the prevalence of specific quality issues across tiers. While Nemotron-CC’s classifier tiers do correlate with several propella-1 quality dimensions (higher tiers have better content quality and information density on average), the “high” tier still contains documents with issues that multi-property annotation reveals. These include documents with heavy or pure marketing commercial bias, thin information density, and incomplete content integrity, among others.

This finding illustrates a key limitation of single-score quality classification: documents can score well on one aggregate measure while having undesirable characteristics on specific dimensions that matter for particular downstream applications.




Figure 3: Property distributions across four German-language pretraining sources. Sources differ dramatically across quality dimensions despite all being German text corpora. FinePDFs shows substantially higher rates of excellent quality, analytical reasoning, and high educational value.

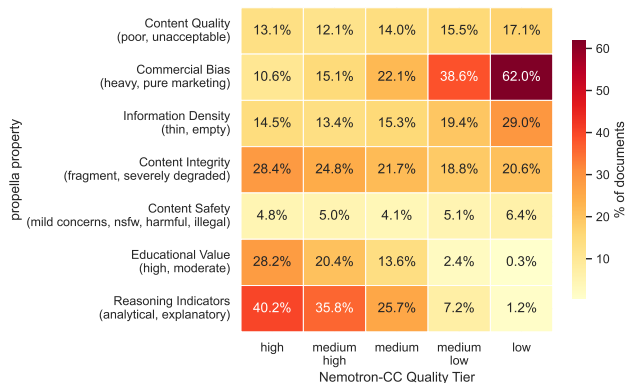


Figure 4: Prevalence of quality issues in Nemotron-CC quality tiers. Even the “high” quality tier contains documents with issues on specific dimensions (commercial bias, information density, content integrity) that the single-score classifier does not capture.

7.3 CROSS-LANGUAGE CONTENT VARIATION IN FINEWEB-2

FineWeb-2 applies the same crawl and processing pipeline across languages (Penedo et al., 2025a). We examine whether this produces comparable quality profiles by comparing property distributions across six languages: German, Spanish, French, Italian, Swedish, and Finnish.

Figure 5 shows that the same pipeline yields substantially different content profiles by language. Languages differ in their distributions of content quality, commercial bias, educational value, and content type composition. These differences likely reflect the underlying composition of each language’s web presence rather than pipeline artifacts.

The implication is that language-specific filtering thresholds are needed for effective multilingual data curation. A single quality threshold applied uniformly across languages may be too aggressive for some languages and too lenient for others.

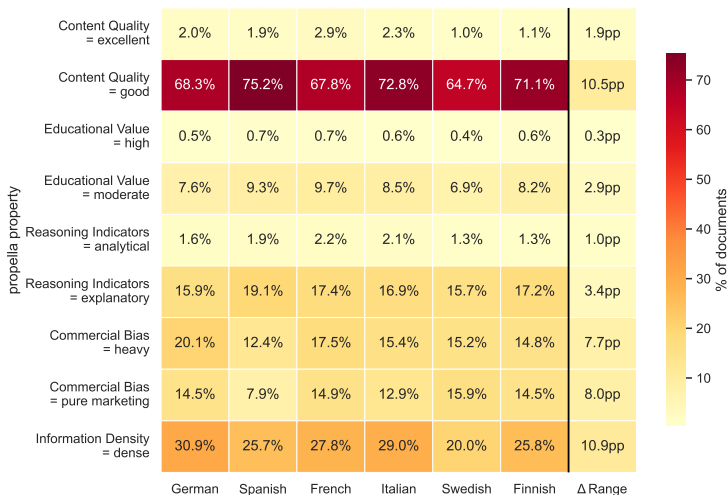


Figure 5: Property distributions across six languages in FineWeb-2. The rightmost column shows the range (max – min) across languages. Commercial bias and information density exhibit the largest cross-language variation, while educational value and reasoning indicators are more uniform. These differences motivate language-specific filtering strategies.

8 DISCUSSION AND FUTURE WORK

Value of the release. The propella-annotations dataset enables researchers to study which quality dimensions matter for downstream LLM training without incurring the cost of annotation. By providing pre-computed annotations for over three billion documents across major pretraining corpora, we lower the barrier to data-centric research.

Limitations. Our evaluation relies on agreement with Gemini-3-Pro rather than human annotations, which introduces a dependency on a specific model’s judgments. Because the distilled models and the evaluation reference may share systematic biases originating from frontier LLM outputs, high agreement could mask shared errors that human evaluation would reveal. This risk is compounded by the tendency of LLM evaluators to favor outputs resembling their own generations (Panickssery et al., 2024). Additionally, the propella-1 models inherit biases from the frontier LLMs used to generate training labels, which may affect annotation consistency for underrepresented languages or content types. LLM-based judges are also known to vary in reliability across languages (Fu & Liu, 2025), and our per-language evaluation samples may be too small to detect such variation reliably. The property rubric itself reflects one team’s design choices and may not capture all dimensions relevant to every use case. Most importantly, we have not yet evaluated whether filtering with propella-1 annotations improves downstream model training, which is the ultimate measure of their utility. However, the case studies in §7 demonstrate the analytical value of multi-property annotation, and the public release of all annotations enables any researchers to conduct downstream training experiments.

Future work. The most important next step is evaluating the downstream training impact of multi-property filtering. Does selecting documents that are high on educational value, reasoning indicators, and content quality simultaneously produce better models than selecting on any single dimension? Our annotations and their public release make this experiment accessible to any research group. We plan to perform various ablation experiments to evaluate the impact of different quality dimensions on downstream task performance. Additional directions include property-aware curriculum learning, where the training data mixture evolves based on multiple quality dimensions over the course of training, and integration with data transformation and synthetic generation pipelines, where multi-property annotations could guide which documents to rephrase, augment, or discard.

9 CONCLUSION

We presented propella-1, a family of small multilingual LLMs for multi-property document annotation, and propella-annotations, a large-scale open multi-property annotation of LLM pretraining data covering over three billion documents. The annotations are accurate (the 4B model exceeds Gemini-3-Flash in agreement with Gemini-3-Pro), efficient to produce (27 documents per second per H100 GPU), highly multilingual (57 languages), and openly available. Through three analytical case studies, we demonstrated that multi-property annotation reveals compositional structure in pretraining datasets that single-score approaches cannot capture: different data sources have dramatically different quality profiles, single-score quality tiers mask heterogeneity, and the same processing pipeline produces different quality distributions across languages. We release all model weights and annotations to enable the community to build on this resource for data mixture design, compliance-aware governance, and multilingual data curation research.

ACKNOWLEDGMENTS

- This project used compute resources made available via the EuroHPC Joint Undertaking (EuroHPC JU) AI Factories initiative (AI for Industrial Innovation – Large Scale Access - EHPC-AIF-2025LS01-028) on the EuroHPC supercomputer LEONARDO operated by CINECA and the LEONARDO consortium.
- This project used compute resources made available via the EuroHPC Joint Undertaking (EuroHPC JU) AI Factories initiative (AI for Industrial Innovation – Large Scale Access - EHPC-AIF-2025LS16-024) on the EuroHPC supercomputer MareNostrum 5 operated by the Barcelona Supercomputing Center (BSC).
- This project is supported by the OpenEuroLLM project, co-funded by the Digital Europe Programme under GA no. 101195233. For more information see <https://openeurollm.eu>.
- This project is supported by the LLMs4EU project, co-funded by the Digital Europe Programme under GA no. 101198470. For more information see the LLMs4EU website.
- This work was funded by the German Federal Ministry for Economic Affairs and Climate Action (BMWK) under grant 13IPC040J (SOOFI) based on a resolution of the German Bundestag, and by the European Union – NextGenerationEU. The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.
- ellamind thanks the AI Service Center for Sensitive and Critical Infrastructures (KISSKI), operated by GWDG, for additional compute access.

ETHICS STATEMENT

The propella-annotations dataset includes safety and PII annotations for over three billion documents. While these labels are intended to support filtering out unsafe or privacy-sensitive content, they could in principle be misused to locate such content. We note that the annotations contain only categorical labels and brief text descriptions, not the source documents themselves, which limits this risk. The safety and PII labels produced by propella-1 are not a substitute for dedicated content moderation systems and should be used as one signal among many in a responsible data curation pipeline. The annotation rubric reflects one team’s design choices and may encode biases, particularly for underrepresented languages and cultural contexts. The models inherit biases from the frontier LLMs used to generate training labels, and users should be aware of this when applying annotations to sensitive filtering decisions.

REPRODUCIBILITY STATEMENT

We release all artifacts necessary to reproduce and build upon this work.⁴ Model weights for all three propella-1 variants (0.6B, 1.7B, 4B) are released under the Apache 2.0 license. The propella-annotations dataset of over three billion annotations is hosted on the Hugging Face Hub under CC-BY-4.0. The complete annotation rubric is provided in Appendix F and released alongside the model weights. Serving configuration, prompt design, and inference throughput details are described in Sections 4 and 6, and are included in the model repository. Training data composition is detailed in Appendix A.

⁴<https://hf.co/collections/ellamind/propella-1>

REFERENCES

- Mehdi Ali, Manuel Brack, Max Lübbering, Elias Wendt, Abbas Goher Khan, Richard Rutmann, Alex Jude, Maurice Kraus, Alexander Arno Weber, Felix Stollenwerk, David Kaczér, Florian Mai, Lucie Flek, Rafet Sifa, Nicolas Flores-Herr, Joachim Koehler, Patrick Schramowski, Michael Fromm, and Kristian Kersting. Judging quality across languages: A multilingual approach to pretraining data filtering with language models. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 8859–8898, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.449. URL <https://aclanthology.org/2025.emnlp-main.449/>.
- Elie Bakouch, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, Lewis Tunstall, Carlos Miguel Patiño, Edward Beeching, Aymeric Roucher, Aksel Joonas Reedi, Quentin Gallouédec, Kashif Rasul, Nathan Habib, Clémentine Fourier, Hynek Kydlicek, Guilherme Penedo, Hugo Larcher, Mathieu Morlon, Vaibhav Srivastav, Joshua Lochner, Xuan-Son Nguyen, Colin Raffel, Leandro von Werra, and Thomas Wolf. SmolLM3: smol, multilingual, long-context reasoner. <https://huggingface.co/blog/smollm3>, 2025. Accessed: 2026-02-07.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Xiyan Fu and Wei Liu. How reliable is multilingual LLM-as-a-judge? In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 11040–11053, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-335-7. doi: 10.18653/v1/2025.findings-emnlp.587. URL <https://aclanthology.org/2025.findings-emnlp.587/>.
- Lukas Gienapp, Christopher Schröder, Stefan Schweter, Christopher Akiki, Ferdinand Schlatt, Arden Zimmermann, Phillipe Genêt, and Martin Potthast. The german commons – 154 billion tokens of openly licensed text for german language models. *CoRR*, abs/2510.13996, 2025. doi: 10.48550/ARXIV.2510.13996. URL <https://doi.org/10.48550/arXiv.2510.13996>.
- Guidance AI. Ilguidance: Fast structured output for LLMs. GitHub repository, 2025. <https://github.com/guidance-ai/llguidance>.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. Bag of tricks for efficient text classification. In Mirella Lapata, Phil Blunsom, and Alexander Koller (eds.), *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 427–431, Valencia, Spain, April 2017. Association for Computational Linguistics. URL <https://aclanthology.org/E17-2068/>.
- Jeffrey Li, Alex Fang, Georgios Smyrnis, Hongyu Kang, Suchin Gururangan, Samir Yitzhak Gadre, Ludwig Schmidt, and Tatsunori Hashimoto. DataComp-LM: In search of the next generation of training sets for language models. In *Advances in Neural Information Processing Systems (NeurIPS 2024) Datasets and Benchmarks Track*, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/19e4ea30dded58259665db375885e412-Abstract-Datasets_and_Benchmarks_Track.html.
- Pratyush Maini, Vineeth Dorna, Parth Doshi, et al. BeyondWeb: Lessons from scaling synthetic data for trillion-scale pretraining. *arXiv preprint arXiv:2508.10975*, 2025.
- Bettina Messmer, Vinko Sabolčec, and Martin Jaggi. Enhancing multilingual LLM pretraining with model-based data selection. In *NeurIPS 2025 Datasets and Benchmarks Track*, 2025. URL <https://openreview.net/forum?id=ufhH8YXxOB>. Poster.
- Mistral AI. Mistral-small-3.2-24b-instruct-2506. <https://huggingface.co/mistralai/Mistral-Small-3.2-24B-Instruct-2506>, 2025. Accessed: 2026-02-07.

- Amanda Myntti, Liina Repo, Elian Freyermuth, Antti Kanner, Veronika Laippala, and Erik Henriksson. Intersecting register and genre: Understanding the contents of web-crawled corpora. In Mika Hämmäläinen, Emily Öhman, So Miyagawa, Khalid Alnajjar, and Yuri Bizzoni (eds.), *Proceedings of the 4th International Conference on Natural Language Processing for Digital Humanities*, pp. 386–397, Miami, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.nlp4dh-1.38. URL <https://aclanthology.org/2024.nlp4dh-1.38/>.
- Amanda Myntti, Erik Henriksson, Veronika Laippala, and Sampo Pyysalo. Register always matters: Analysis of LLM pretraining data through the lens of language variation. In *Second Conference on Language Modeling*, 2025. URL <https://openreview.net/forum?id=FqXXtSZWEZ>.
- Stephan Oepen, Nikolay Arefev, Mikko Aulamo, Marta Bañón, Maja Buljan, Laurie Burchell, Lucas Charpentier, Pinzhen Chen, Mariya Fedorova, Ona de Gibert, Barry Haddow, Jan Hajič, Jindřich Helcl, Andrey Kutuzov, Veronika Laippala, Zihao Li, Risto Luukkonen, Bhavitvya Malik, Vladislav Mikhailov, Amanda Myntti, Dayyán O’Brien, Lucie Poláková, Sampo Pyysalo, Gema Ramírez Sánchez, Janine Siewert, Pavel Stepachev, Jörg Tiedemann, Teemu Vahtola, Dušan Variš, Fedor Vitiugin, Tea Vojtěchová, and Jaume Zaragoza. Hplt 3.0: Very large-scale multilingual resources for llm and mt. mono- and bi-lingual data, multilingual evaluation, and pre-trained models. 2025. URL <https://arxiv.org/abs/2511.01066>.
- Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. Asynchronous pipelines for processing huge corpora on medium to low resource infrastructures. *Proceedings of the Workshop on Challenges in the Management of Large Corpora (CMLC-7) 2019*. Cardiff, 22nd July 2019, pp. 9–16, Mannheim, 2019. Leibniz-Institut für Deutsche Sprache. doi: 10.14618/ids-pub-9021. URL <http://nbn-resolving.de/urn:nbn:de:bsz:mh39-90215>.
- Arjun Panickssery, Samuel R. Bowman, and Shi Feng. LLM evaluators recognize and favor their own generations. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-2197. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/7f1f0218e45f5414c79c0679633e47bc-Abstract-Conference.html.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben Allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. The fineweb datasets: decanting the web for the finest text data at scale. In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS ’24*, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9798331314385. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/370df50ccfd8bde18f8f9c2d9151bda-Paper-Datasets_and_Benchmarks_Track.pdf.
- Guilherme Penedo, Hynek Kydlíček, Vinko Sabolčec, Bettina Messmer, Negar Foroutan, Amir Hossein Kargar, Colin Raffel, Martin Jaggi, Leandro Von Werra, and Thomas Wolf. Fineweb2: One pipeline to scale them all – adapting pre-training data processing to every language. In *Second Conference on Language Modeling*, 2025a. URL <https://openreview.net/forum?id=jnRBe6zatP>.
- Guilherme Penedo, Hynek Kydlíček, et al. FinePDFs: High-quality scientific pdfs for pretraining. <https://huggingface.co/datasets/HuggingFaceFW/FinePDFs>, 2025b. Accessed: 2026-02-07.
- Ru Peng, Kexin Yang, Yawen Zeng, Junyang Lin, Dayiheng Liu, and Junbo Zhao. Dataman: Data manager for pre-training large language models. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 9747–9799, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/1abed6ee581b9ceb4e2ddf37822c7fcb-Paper-Conference.pdf.
- Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, H. Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, Eliza Rutherford, Tom Hennigan, Jacob Menick, Albin Cassirer, Richard Powell, George van den Driessche, Lisa Anne Hendricks, Maribeth Rauh, Po-Sen Huang, Amelia Glaese, Johannes Welbl, Sumanth Dathathri, Saffron Huang, Jonathan Uesato, John Mellor, Irina Higgins, Antonia Creswell, Nat McAleese, Amy Wu,

- Erich Elsen, Siddhant M. Jayakumar, Elena Buchatskaya, David Budden, Esme Sutherland, Karen Simonyan, Michela Paganini, Laurent Sifre, Lena Martens, Xiang Lorraine Li, Adhiguna Kuncoro, Aida Nematzadeh, Elena Gribovskaya, Domenic Donato, Angeliki Lazaridou, Arthur Mensch, Jean-Baptiste Lespiau, Maria Tsimpoukelli, Nikolai Grigorev, Doug Fritz, Thibault Sottiaux, Mantas Pajarskas, Toby Pohlen, Zhitao Gong, Daniel Toyama, Cyprien de Masson d’Autume, Yujia Li, Tayfun Terzi, Vladimir Mikulik, Igor Babuschkin, Aidan Clark, Diego de Las Casas, Aurelia Guy, Chris Jones, James Bradbury, Matthew J. Johnson, Blake A. Hechtman, Laura Weidinger, Iason Gabriel, William Isaac, Edward Lockhart, Simon Osindero, Laura Rimell, Chris Dyer, Oriol Vinyals, Kareem Ayoub, Jeff Stanway, Lorraine Bennett, Demis Hassabis, Koray Kavukcuoglu, and Geoffrey Irving. Scaling language models: Methods, analysis & insights from training gopher. *CoRR*, abs/2112.11446, 2021. URL <https://arxiv.org/abs/2112.11446>.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.
- Dan Su, Kezhi Kong, Ying Lin, Joseph Jennings, Brandon Norrick, Markus Kliegl, Mostofa Patwary, Mohammad Shoeybi, and Bryan Catanzaro. Nemotron-cc: Transforming common crawl into a refined long-horizon pretraining dataset. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025*, pp. 2459–2475. Association for Computational Linguistics, 2025. ISBN 979-8-89176-251-0. URL <https://aclanthology.org/2025.acl-long.123/>.
- Gemma Team. Gemma 3 technical report. *CoRR*, abs/2503.19786, 2025. doi: 10.48550/ARXIV.2503.19786. URL <https://doi.org/10.48550/arXiv.2503.19786>.
- Team OLMo, Allyson Ettinger, Amanda Bertsch, et al. OLMo 3. *arXiv preprint arXiv:2512.13961*, 2025.
- Maurice Weber, Dan Fu, Quentin Anthony, Yonatan Oren, Shane Adams, Anton Alexandrov, Xiaozhong Lyu, Huu Nguyen, Xiaozhe Yao, Virginia Adams, et al. Redpajama: an open dataset for training large language models. In *Advances in Neural Information Processing Systems*, volume 37, pp. 116462–116492, 2024. URL https://papers.nips.cc/paper/2024/hash/d34497330b1fd6530f7afd86d0df9f76-Abstract-Datasets_and_Benchmarks_Track.html.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. CCNet: Extracting high quality monolingual datasets from web crawl data. In Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis (eds.), *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pp. 4003–4012, Marseille, France, May 2020. European Language Resources Association. ISBN 979-10-95546-34-4. URL <https://aclanthology.org/2020.lrec-1.494/>.
- Alexander Wettig, Aatmik Gupta, Saumya Malik, and Danqi Chen. Qurating: Selecting high-quality data for training language models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 52915–52971, 2024. URL <https://proceedings.mlr.press/v235/wettig24a.html>.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jian Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. *CoRR*, abs/2505.09388, 2025. doi: 10.48550/ARXIV.2505.09388. URL <https://doi.org/10.48550/arXiv.2505.09388>.

Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Livia Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E Gonzalez, et al. Sglang: Efficient execution of structured language model programs. In *Advances in Neural Information Processing Systems*, volume 37, pp. 62557–62583, 2024.

Xinlin Zhuang, Jiahui Peng, Ren Ma, Yinfan Wang, Tianyi Bai, Xingjian Wei, Qiu Jiantao, Chi Zhang, Ying Qian, and Conghui He. Meta-rater: A multi-dimensional data selection method for pre-training language models. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 10856–10896, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.533. URL <https://aclanthology.org/2025.acl-long.533/>.

A TRAINING DATA COMPOSITION

Tables 4 and 5 show the language and source distributions of the training data, respectively.

Table 4: Language distribution of the training data. Language codes follow the FLORES-200 convention: ISO 639-3 language codes combined with ISO 15924 script codes (e.g., eng_Latn = English in Latin script).

Language	%	Language	%	Language	%
eng_Latn	35.08	por_Latn	0.90	glg_Latn	0.83
spa_Latn	3.98	kat_Geor	0.90	gle_Latn	0.83
ita_Latn	3.97	ben_Beng	0.90	nno_Latn	0.83
fra_Latn	3.95	fas_Arab	0.89	ltg_Latn	0.77
deu_Latn	3.86	nob_Latn	0.86	yue_Hant	0.49
pol_Latn	3.81	ekk_Latn	0.89	cmn_Hant	0.48
code	2.82	fin_Latn	0.89	hrv_Latn	0.43
math	2.77	tur_Latn	0.89	arb_Arab	0.39
sft	2.41	swe_Latn	0.88	bos_Latn	0.39
ukr_Cyrl	0.95	ind_Latn	0.88	mkd_Cyrl	0.39
nld_Latn	0.95	ces_Latn	0.88	srp_Latn	0.37
tha_Thai	0.95	lit_Latn	0.88	cmn_Hani	0.37
jpn_Jpan	0.94	eus_Latn	0.87	hin_Deva	0.36
heb_Hebr	0.94	vie_Latn	0.87	srp_Cyrl	0.36
ell_Grek	0.93	slv_Latn	0.87	als_Latn	0.35
kor_Hang	0.93	mlt_Latn	0.86	sqi_Latn	0.03
isl_Latn	0.93	bul_Cyrl	0.86	est_Latn	0.02
dan_Latn	0.92	lvs_Latn	0.86	lav_Latn	0.02
cat_Latn	0.92	urd_Arab	0.85	nor_Latn	0.02
slk_Latn	0.92	hun_Latn	0.85	swa_Latn	0.02
rus_Cyrl	0.91	ron_Latn	0.84		

Table 5: Source distribution of the training data.

Source	%	Source	%
HPLT 3.0 (unfiltered)	39.59	Nemotron SFT code	0.41
FineWeb	16.01	Nemotron-CC diverse QA	0.41
FineWeb-2	13.23	RedPajama arXiv sample	0.40
FinePDFs	8.09	RedPajama Wikipedia sample	0.40
FineWeb-2 (removed)	6.28	Nemotron-CC high-quality synth	0.39
FineWeb-Edu (dedup)	3.92	Nemotron-CC translated diverse	0.38
The Stack	2.04	Nemotron synthetic code	0.38
FineMath	2.00	RedPajama GitHub sample	0.37
OpenHermes	2.00	RedPajama StackExchange sample	0.36
FineWiki	1.35	Nemotron-CC math	0.35
RedPajama C4 sample	0.42	Nemotron-CC high-quality	0.33
Nemotron SFT general	0.42	RedPajama book sample	0.07
Nemotron SFT math	0.42		

B EXAMPLE ANNOTATION OUTPUT

```
{
  "content_integrity": "complete",
  "content_ratio": "mostly_content",
  "content_length": "substantial",
  "one_sentence_description": "Research paper on ...",
  "content_type": ["analytical"],
  "business_sector": ["academic_research"],
  "technical_content": ["scientific", "data_heavy"],
  "content_quality": "excellent",
  "information_density": "dense",
  "educational_value": "high",
  "reasoning_indicators": "analytical",
  "audience_level": "expert",
  "commercial_bias": "none",
  "time_sensitivity": "slowly_changing",
  "content_safety": "safe",
  "pii_presence": "no_pii",
  "regional_relevance": ["european"],
  "country_relevance": ["germany"]}

```

Figure 6: Representative propella-1 annotation output (JSON, formatted for readability). During inference, the output contains no whitespace to minimize token count.

C PROMPT TEMPLATE

The propella-1 models use a compact system prompt that lists all 18 properties with their enumerated values and brief descriptions. The user message wraps the document text in delimiter tags. During training data creation, a substantially longer annotator prompt (~14K tokens) was used, which includes the full rubric from Appendix F together with the JSON schema.

System prompt.

```
Annotate the document. Any language; assess quality
within its linguistic norms. Respond with a JSON object:
content_integrity: technical completeness
  (complete|mostly_complete|fragment|severely_degraded)
content_ratio: content vs navigation/boilerplate ratio
  (complete_content|mostly_content|mixed_content|
  mostly_navigation|minimal_content)
content_length: substantive words
```

```
(substantial 2k+|moderate 500-2k|brief 100-500|
minimal <100)
one_sentence_description: neutral ~10 word summary
in English
content_type[]: functional purpose
(analytical|instructional|reference|procedural|
qa_structured|conversational|creative|transactional|
boilerplate|news_report|opinion_editorial|
review_critique|technical_documentation|
specification_standard|legal_document|press_release|
structured_data|source_code)
business_sector[]: industry domain
(academic_research|education_sector|...|other)
technical_content[]: specialized knowledge
(code_heavy|math_heavy|scientific|data_heavy|
engineering|basic_technical|non_technical)
content_quality: writing/presentation quality
(excellent|good|adequate|poor|unacceptable)
information_density: signal vs padding
(dense|adequate|moderate|thin|empty)
educational_value: teaching potential
(high|moderate|basic|minimal|none)
reasoning_indicators: logical analysis depth
(analytical|explanatory|basic_reasoning|minimal|none)
audience_level: assumed background
(expert|advanced|general|beginner|youth|children)
commercial_bias: promotional influence
(none|minimal|moderate|heavy|pure_marketing)
time_sensitivity: temporal decay
(evergreen|slowly_changing|regularly Updating|
time_sensitive)
content_safety: harmful content
(safe|mild_concerns|nsfw|harmful|illegal)
pii_presence: private individual data
(no_pii|contains_pii)
regional_relevance[]: geographic/cultural context
(european|north_american|east_asian|...|indeterminate)
country_relevance[]: specific countries as ISO names,
or supranational|none
```

User message.

```
<start_of_document>
{document text}
<end_of_document>
```

Multi-select properties are indicated by [] after the property name. The `business_sector` property (37 values) and `regional_relevance` property (14 values) are abbreviated above; the full enumerations are listed in Appendix F.

Annotator system prompt (training data creation). During training data creation, frontier LLMs receive a substantially longer system prompt that includes the full annotation rubric and JSON response schema. The template is shown below with placeholders: `{property_descriptions}` expands to the complete rubric (Appendix F), and `{json_schema}` expands to a JSON Schema derived from the Pydantic response model.

```
You are an expert content analysis assistant specializing
in document annotations for LLM pretraining data. Your
team is curating a multilingual dataset for language model
training. Your task is to analyze documents and annotate
them with specific properties that will later on be used
to filter the dataset. The user will provide a document
inside of "<start_of_document>" and "<end_of_document>"
tags. Analyze the content of the document systematically
```

and objectively. Respond with your annotations in JSON format, following the annotation framework below.

```
# Annotation Framework
## Output Requirements
- Respond with a JSON object matching the specified schema.
- Use the exact enum values provided.
- Ensure all fields are included.
- Multi-select fields: content_type, business_sector,
  technical_content, regional_relevance,
  country_relevance. All other properties are
  single-select strings.

## Key Principles
* Objective assessment
* Completeness: Address all properties for every document.
* Consistency: Apply the same standards across documents.
* Multilinguality: Language should not influence
  annotations.

## Properties to Annotate
[Property summary with enum values -- same content as
the compact prompt above, organized by category]

{property_descriptions}

## JSON Schema for the Response
```json
{json_schema}
```

## Multilingual Annotation Guidelines

### Universal Principles
1. Evaluate content quality within language context
2. Consider linguistic norms
3. Respect script directionality
4. Account for morphological complexity

### Language-Specific Considerations
[Guidelines for Chinese/Japanese, Arabic/Hebrew/Persian,
Indian languages, and European languages]

# Annotation Workflow
- Analyze the document systematically and objectively
- Respond with a valid JSON object matching the schema
- Use the exact enum values provided
- Ensure all required fields are included
- For multi-select properties, always return arrays
- Do not include any explanatory text or formatting
```

Annotator user message.

Analyze the following document and provide annotations in JSON format according to the annotation framework. Return only the JSON object.

```
<start_of_document>
{document text}
<end_of_document>
```

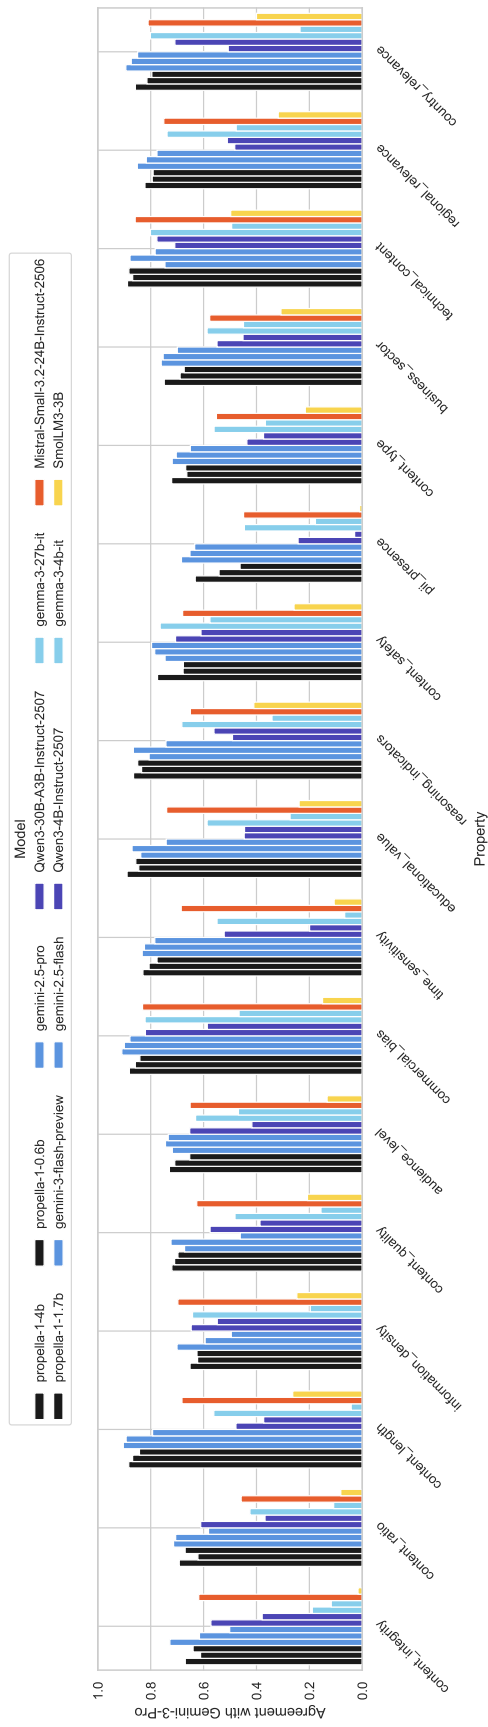


Figure 7: Per-property annotation agreement scores across all evaluated models and all 17 evaluated properties.

D FULL PER-PROPERTY EVALUATION RESULTS

E INFERENCE PRECISION: FP8 VS. BF16

propella-1 models are trained with fp8 precision and work well in both bf16 and fp8 inference modes. As Table 6 shows, fp8 inference preserves annotation quality, with negligible differences compared to bf16.

Table 6: Effect of inference precision on annotation quality (overall score).

| Model | bf16 | fp8 | Diff |
|-----------------|-------------|------------|-------------|
| propella-1-4b | 0.779 | 0.783 | +0.51% |
| propella-1-1.7b | 0.737 | 0.731 | -0.81% |

F COMPLETE PROPERTY REFERENCE

Below is the complete annotation rubric for all 18 propella-1 properties. For each property, we provide the full operational definition including enumerated values, decision criteria, borderline examples, decision trees, and language-specific guidelines. The rubric is also released alongside the model weights.

F.1 CORE CONTENT PROPERTIES

F.1.1 CONTENT INTEGRITY

What we measure: Completeness and technical quality of the content itself, regardless of navigation ratio.

complete Full, intact content as intended.

- Content appears complete with proper beginning, middle, and end.
- All essential elements present (introduction, body, conclusion where appropriate).
- No obvious truncation or missing sections.
- Example: Complete articles, full tutorials, intact documents.

mostly_complete Minor elements missing but core content intact.

- Core content is complete but some secondary elements may be missing.
- Minor truncation that does not affect main message.
- Example: Article with truncated comments, missing sidebar content, partial author bio.

fragment Incomplete content, missing significant portions.

- Missing introduction, conclusion, or substantial middle sections.
- Truncated mid-sentence or mid-paragraph.
- Content feels incomplete or cut off.
- Example: Search result snippets, article excerpts, broken crawls, partial downloads.

severely_degraded Broken, unreadable, or corrupted content.

- Encoding errors, scrambled text, missing characters.
- Severely malformed HTML rendering as gibberish.
- Technical corruption making content unreadable.
- Example: Garbled characters, completely broken formatting, corrupted files.

Key Decision Points.

- **Content completeness:** Does the content feel like a complete unit of information?
- **Technical integrity:** Is the content technically readable and properly formatted?
- **Fragment vs. complete:** Independent of navigation; is the actual content complete?
- **Degraded vs. fragment:** Degraded has technical issues; fragment is just incomplete.

Note. Documents may end with a special `<content_truncated>` tag, indicating upstream length-based truncation due to processing constraints. Content Integrity should not be penalized for this truncation signal; assess integrity based on the visible content's coherence and technical readability, ignoring the artificial cutoff.

F.1.2 CONTENT RATIO

What we measure: How much of the document is actual content vs. navigation, UI elements, and structural markup.

complete_content 90–100% meaningful content.

- Full articles, papers, tutorials with minimal navigation.

- Clean text with proper paragraphs and structure.
- Example: A Wikipedia article, academic paper, complete blog post.

mostly_content 70–89% meaningful content.

- Complete documents with some navigation elements (header, footer, sidebar).
- Minor UI elements that do not disrupt reading.
- Example: News articles with standard website navigation.

mixed_content 40–69% meaningful content.

- Significant navigation mixed throughout content.
- Multiple sidebars, ads, or UI elements interrupting text.
- Example: E-commerce product pages with reviews mixed with purchase options.

mostly_navigation 10–39% meaningful content.

- Predominantly menus, links, headers, footers.
- Content overwhelmed by structural elements.
- Example: Site maps, navigation pages, heavily UI-focused pages.

minimal_content 0–9% meaningful content.

- Almost entirely navigation, UI elements, or structural markup.
- Very little readable content present.
- Example: Empty pages, pure navigation menus, error pages with minimal text.

Key Decision Points.

- Focus on the ratio of readable text to navigation/UI elements.
- Count only substantive content; ignore boilerplate and structural elements.
- **Mixed vs. mostly_navigation:** Can you read it as coherent content despite distractions?

F.1.3 CONTENT LENGTH

What we measure: Amount of substantive content, ignoring navigation and boilerplate.

substantial 2,000+ words of meaningful content.

- Long-form, comprehensive content that provides in-depth coverage of a topic.
- Typically includes detailed analysis, multiple sections or chapters, extensive research, or thorough exploration of complex subjects.
- Examples: White papers, research reports, e-books, long-form journalism.

moderate 500–2,000 words of meaningful content.

- Standard-length content that offers meaningful coverage while remaining focused and digestible.
- Balances depth with accessibility; provides enough detail to be informative without overwhelming readers.
- Examples: Typical blog posts, news articles, product reviews, how-to guides.

brief 100–500 words of meaningful content.

- Short, focused content that delivers key information quickly and efficiently.
- Gets straight to the point while still providing value and context.
- Examples: News briefs, product descriptions, FAQs, short blog posts.

minimal Under 100 words of meaningful content.

- Very short content that provides only essential information or serves as a quick reference.
- Designed for rapid consumption or specific micro-purposes.
- Examples: Social media posts, announcements, abstracts, snippets, navigation pages.

Measurement Tips.

- **Count only readable content of value:** include article body and substantive headings/captions; exclude headers/footers, menus/sidebars, related links, share/consent UI, pagination, ads, and boilerplate.
- Focus on substantive information, not filler words.
- Complete thoughts matter more than exact word counts.
- **Contextual adjustment:** Thresholds are guidelines and can be adjusted based on specific use cases. Academic contexts may shift ranges upward, while social media contexts may shift them downward.

F.2 CONTENT CLASSIFICATION

F.2.1 ONE-SENTENCE DESCRIPTION

What we produce: A very short, neutral description of what the document contains.

one_sentence_description Ultra-short neutral description of the document.

- Exactly one sentence; target length <100 characters.
- Focus on the main topic and, if useful, the document’s function (e.g., tutorial, policy, news report, product page, navigation page).
- Neutral, descriptive tone (no hype or marketing language).

To Avoid.

- Boilerplate intros: “This document. . .”, “This article. . .”, “In this guide. . .”
- Calls to action: “Learn how to. . .”, “Discover. . .”, “Find out. . .”
- User-facing phrasing: “You will learn. . .”, “How do I. . .”
- Non-essential details (dates, numbers) unless central to the topic.

Examples.

- “Beginner tutorial on React hooks and basic state management.”
- “News report on European Central Bank interest rate decisions.”
- “Internal policy for customer data retention and deletion.”
- “API reference for payment processing endpoints and error codes.”
- “Research paper analyzing housing price trends in major US cities.”
- “FAQ answering common questions about employee parental leave.”
- “Opinion essay arguing for stricter international climate change legislation.”

Examples for low-quality or problematic documents (still annotate).

- “Fragment of article discussing proposed changes to European data privacy laws.”
- “Keyword-stuffed promotional page about cheap car insurance quotes.”
- “Website navigation page listing links to product categories and help pages.”
- “Error page explaining that the requested resource could not be found.”
- “Affiliate landing page promoting multiple online casino bonus offers.”
- “Corrupted text with no identifiable topic or meaningful content.”

F.2.2 CONTENT TYPE

What we measure: The functional structure and purpose of content.

Content can be assigned multiple type labels if it genuinely serves multiple purposes. Choose all applicable types rather than forcing a single primary choice. Always output an array for this property, even if only one type applies.

analytical In-depth analysis, research, and critical examination. Provides detailed analysis or research on a topic; develops arguments, evaluates evidence, or presents findings. Example: research analysis, investigative reports, academic articles, expert commentary.

instructional Teaching and how-to content. Explicitly teaches skills, concepts, or procedures; step-by-step guidance or educational explanations. Example: tutorials, how-to guides, educational content, training materials.

reference Lookup materials, definitions, specifications. Designed for looking up specific information rather than reading through; often organized alphabetically, categorically, or as lists. Example: dictionaries, encyclopedias, API references, product catalogs.

procedural Step-by-step processes and procedures. Sequential instructions or workflows; process documentation with clear steps. Example: recipes, installation guides, standard operating procedures, workflows.

qa_structured Structured question-answer content. Formal Q&A format with clear questions and answers; often expert responses to specific questions. Example: Stack Overflow, FAQ sections, structured Q&A sites.

conversational Multi-party or turn-based dialogues (humans, bots, or both). Casual or structured conversations between two or more participants; may include human–AI chats, forum threads, or comment chains. Example: Reddit threads, forum discussions, support chats, assistant chat logs.

creative Entertainment, artistic, fictional content. Primary purpose is entertainment or artistic expression; not primarily informational or instructional. Example: short stories, poems, movie reviews, game content, fiction.

transactional Commercial, shopping, service-oriented. Primary purpose is to facilitate a transaction or service; focuses on products, services, or business processes. Example: product listings, service descriptions, checkout pages.

boilerplate Legal, policy, standard template text. Standard legal or policy language; often repeated across multiple sites with minimal variation. Example: terms of service, privacy policies, disclaimers, cookie banners, standard notices.

news_report Straight reporting of events with minimal analysis. Describes events or facts in a neutral, descriptive tone; time-bound news, updates, or reports. Example: wire-service news articles, breaking-news updates.

opinion_editorial Persuasive/opinionated commentary or editorials. Expresses a stance or argument; aims to persuade; may cite evidence but prioritizes viewpoint. Example: op-eds, opinion columns, personal essays with clear stance.

review_critique Evaluative reviews of products, media, or services. Provides judgments, ratings, or critiques; may include pros/cons, scoring systems. Example: product reviews, film/book critiques, app store reviews (long-form).

technical_documentation Manuals, API docs, developer guides, READMEs. Primary goal is to instruct usage of software/hardware/APIs; includes reference sections, examples, parameters, version notes. Example: API reference, library README, user manual.

specification_standard Normative standards and formal specifications. Defines requirements, must/shall language, compliance criteria; maintained by standards bodies or authoritative groups. Example: RFCs, ISO standards, formal protocol specs.

legal_document Statutes, case law, contracts, regulatory texts. Binding or authoritative legal content; formal legal language and structure. Example: court opinions, legislation, contracts, regulatory rules.

press_release Organization-issued announcements and PR materials. Promotional announcements framed as information; quotes from executives, product/service announcements. Example: company press releases, launch announcements.

structured_data Tables, datasets, indices, catalogs with minimal prose. Predominantly tabular/listed data meant for lookup; minimal narrative or explanatory text. Example: product catalogs, schedules, statistical tables.

source_code Code listings as primary content. Dominant content is program source code or scripts; may include lightweight comments or snippets without narrative. Example: code files, gist-like pages, competitive programming solutions.

Multi-Type Examples.

- Tutorial that analyzes different approaches → [instructional, analytical]
- Educational reference manual → [instructional, reference]
- Research paper with step-by-step methodology → [analytical, procedural]
- Q&A site with analytical responses → [qa_structured, analytical]
- API guide with examples → [technical_documentation, reference, instructional]
- RFC with rationale → [specification_standard, analytical]
- Film review with interview snippets → [review_critique, conversational]
- Helpdesk chat with an AI → [conversational, transactional]

F.2.3 BUSINESS SECTOR

What we measure: Business sector(s) or industry domain(s) for training sector-specific LLMs.

Content can be assigned multiple sector labels if it genuinely spans multiple industries. Choose all applicable sectors rather than forcing a single primary choice or using “other”. Always output an array.

academic_research Scholarly and research content (journal articles, conference papers, dissertations).

education_sector Educational institutions and pedagogy (curricula, edtech, teaching resources).

technology_software Software and information technology (development, platforms, IT services).

hardware_electronics Hardware devices and electronics (semiconductors, embedded systems, datasheets).

healthcare_medical Healthcare and medical sector (clinical practice, medical devices, public health).

pharmaceutical_biotech Pharmaceutical and biotechnology (drug development, clinical trials, life sciences).

financial_services Banking and financial services (investment, fintech, asset management).

legal_services Legal sector and jurisprudence (law firms, compliance, litigation).

government_public Government and public administration (agencies, policy, regulatory bodies).

manufacturing_industrial Manufacturing and heavy industry (production, supply chain, industrial engineering).

mining_resources Mining and natural resources (exploration, extraction, commodity operations).

chemicals_materials Chemicals and advanced materials (petrochemicals, polymers, materials science).

energy_utilities Energy and utilities (power generation, renewables, grid management).

retail_commerce Retail and e-commerce (retail operations, merchandising, consumer experience).

wholesale_distribution Wholesale trade and distribution (B2B, procurement, fulfillment).

real_estate_construction Real estate and construction (property development, architecture).

transportation_logistics Transportation and logistics (shipping, freight, fleet management).

travel_aviation Travel industry and commercial aviation (airlines, OTA platforms, IATA regulations).

automotive_industry Automotive manufacturing and services (vehicle technology, electric vehicles).

telecommunications Telecommunications (network infrastructure, mobile services, 5G).

media_entertainment Media and entertainment (film, television, music, streaming, publishing).

gaming_industry Video games and interactive entertainment (development, esports, live ops).

gambling_betting Gambling, betting, and online casinos (sportsbooks, affiliate content).

advertising_marketing Advertising, marketing, and PR (brand strategy, campaign planning, martech).

hospitality_tourism Hospitality and tourism (hotels, destination management, event planning).

food_beverage_hospitality Food & beverage and restaurant operations (menu engineering, food safety).

agriculture_food Agriculture and food production (farming, agtech, food processing).

environmental_services Environmental and sustainability services (ESG, remediation, impact assessments).

aerospace_defense Aerospace and defense (aircraft manufacturing, space technology, military systems).

insurance_industry Insurance sector (actuarial science, underwriting, risk assessment).

nonprofit_ngo Nonprofit and NGO sector (charitable organizations, international development).

consulting_professional Professional services and consulting (management consulting, business advisory).

human_resources HR and people operations (talent acquisition, workforce planning, L&D).

security_cyber Security and cybersecurity (threat intelligence, incident response, compliance).

consumer_goods Consumer products and CPG (household products, personal care, FMCG).

general_interest General audience content (broad topics, cross-sector, lifestyle).

other Highly specialized or unclassifiable niches.

Multi-Sector Examples.

- Medical device regulations → healthcare_medical + pharmaceutical_biotech + government_public
- Fintech software documentation → financial_services + technology_software
- Agricultural biotechnology research → agriculture_food + pharmaceutical_biotech

F.2.4 TECHNICAL CONTENT

What we measure: Type and intensity of specialized technical knowledge.

Content can be assigned multiple technical content labels if it genuinely combines multiple technical domains. Always output an array.

code_heavy Significant programming content. Multiple code examples, algorithms, or implementations; software development focus. Example: programming tutorials, API documentation, software guides.

math_heavy Substantial mathematical content. Mathematical equations, proofs, or statistical analysis; quantitative analysis and mathematical reasoning. Example: mathematical papers, statistical analysis, quantitative research.

scientific Research and scientific methodology content. Scientific research findings, experimental data; peer-reviewed research content. Example: research papers, scientific studies, experimental reports.

data_heavy Substantial datasets, tables, and data analysis. Contains significant data tables, charts, or datasets; focus on data interpretation. Example: research data, statistical reports, survey results.

engineering Engineering and applied technical content. Engineering design, systems, and applied technical solutions; non-software engineering disciplines. Example: mechanical engineering, civil engineering, technical specifications, design documents.

basic_technical Some technical elements but not dominant. Light technical content mixed with general explanations; technical concepts explained for general audience. Example: technology articles for general audience, basic technical explanations.

non_technical No significant technical content. General audience content without specialized technical knowledge; no programming, mathematical, engineering, or scientific focus. Example: general articles, humanities content, basic informational content.

Multi-Technical Examples.

- Data science tutorial with code examples → [code_heavy, math_heavy, data_heavy]
- Engineering research with statistical analysis → [engineering, scientific, data_heavy]
- Computational biology paper → [code_heavy, scientific]

F.3 QUALITY AND VALUE ASSESSMENT

F.3.1 CONTENT QUALITY

What we measure: Overall quality of content considering writing excellence, substantive value, and presentation quality regardless of authorship origin.

excellent Outstanding quality across all dimensions.

- Sophisticated writing with varied sentence structures and engaging style.
- Rich, appropriate vocabulary with error-free grammar and punctuation.
- High substantive value with clear insights or information.
- Professional presentation and formatting; natural flow and logical organization.
- Example: high-quality publications, expert analyses, polished educational content.

good High quality with minor imperfections.

- Grammatically correct with proper sentence structure; appropriate vocabulary and tone.
- Solid substantive value and clear information; good organization and readable flow.
- Only occasional minor issues (1–2 typos per section).
- Example: quality journalism, professional websites, well-written blog posts.

adequate Acceptable quality for most purposes.

- Generally clear and understandable writing; some grammatical errors but meaning remains clear.
- Reasonable substantive value though may lack depth; basic organization present.
- Example: casual blogs, user reviews, basic informational content, simple guides.

poor Significant quality issues impacting utility.

- Multiple errors affecting comprehension or credibility.
- Unclear expression, confusing organization, or awkward phrasing.
- Limited substantive value or questionable information.
- Example: low-quality web content, poorly edited materials, confusing instructions.

unacceptable Quality too low for productive use.

- Severely impaired communication with major errors; incoherent, nonsensical, or corrupted content.
- No reliable substantive value; broken formatting or technical corruption.
- Example: corrupted text, severe translation errors, spam content, SEO content.

Quality Assessment Guidelines.

- **Comprehension:** Can the intended message be clearly understood?
- **Substantive value:** Does the content provide useful information or insights?
- **Technical presentation:** Is the content properly formatted and readable?
- **Error impact:** Do errors significantly impede understanding or credibility?
- **Professional standards:** Does the content meet basic standards for its intended purpose?

Language-Specific Quality Indicators.

- For non-Latin scripts (Arabic, Chinese, Japanese): check for proper character encoding.
- For agglutinative languages (Turkish, Finnish): adjust expectations for word count/density.
- For languages with different formality levels (Japanese, Korean): assess appropriate register.
- Mixed-language documents: evaluate code-switching quality and appropriateness.

F.3.2 INFORMATION DENSITY

What we measure: Ratio of valuable information to redundancy, padding, and repetition.

dense Efficient, information-packed content. Every sentence adds new information or insight; minimal redundancy or unnecessary elaboration; little to no repetition. Example: technical specifications, concise academic writing, quality reference material.

adequate Good information content with reasonable elaboration. Most content adds value with some acceptable elaboration; minimal repetition; good balance of information and explanation. Example: well-written articles, good tutorials with examples.

moderate Mixed substantive content with noticeable padding. Some valuable information mixed with unnecessary elaboration; noticeable repetition of key points. Example: blog posts with some fluff, articles with repetitive conclusions.

thin Low information content with significant problems. Much content does not add new information; high internal repetition and excessive redundancy; significant padding. Example: SEO-optimized content, poorly edited writing.

empty Dominated by repetition and meaningless content. Minimal actual information value; dominated by repetition and copy-paste artifacts; same ideas repeated without development. Example: spam content, template-filled pages, keyword-stuffed articles.

Common Repetition Patterns to Watch For.

- Same phrases repeated throughout (especially in SEO content).
- Identical paragraphs or sections (copy-paste errors).
- Circular reasoning (saying the same thing in different ways).
- Template artifacts (repeated boilerplate mixed with content).

F.3.3 EDUCATIONAL VALUE

What we measure: Potential for teaching, learning, and knowledge transfer.

high Clear instructional design and learning objectives.

- Explicitly teaches concepts or skills; progressive skill building from basic to advanced.
- Clear learning objectives and outcomes; comprehensive explanations with examples.
- Example: quality tutorials, textbooks, structured courses, educational guides.

moderate Good instructional value with some learning potential.

- Some instructional elements present; explanations help build understanding.
- Transferable knowledge to other contexts; good examples or illustrations.
- Example: how-to articles, explanatory content, informative guides.

basic Limited educational content.

- Some explanations but not systematically instructional; basic explanations of concepts.
- Limited learning potential or skill building.
- Example: basic explanations, simple informational content.

minimal Little educational value.

- Primarily informational rather than instructional; no clear learning objectives.
- Entertainment or commercial focus.
- Example: entertainment content, basic news, commercial content.

none No educational content.

- No instructional value or learning potential; purely transactional, entertainment, or administrative.
- Example: pure entertainment, transactions, legal boilerplate.

Disambiguation Tips.

- Explanatory vs. educational: explanations alone \neq educational design; require intent to teach plus scaffolding for Basic+.
- Reference docs: typically Minimal; promote to Basic/Moderate when guided “how-to” segments or curated examples exist.
- Reviews/op-eds: None/Minimal unless they include actionable how-to guidance designed for learning.

Automation Heuristics.

- Keywords: Objectives/Outcomes, Lesson, Exercise/Quiz, Homework, Assessment, Syllabus, Module, Unit, Learning Goals.
- Structure: numbered steps + prerequisites/requirements \rightarrow Basic; add practice tasks/solutions \rightarrow Moderate; syllabus/modules/assessments \rightarrow High.
- Signals of non-educational mix: heavy CTAs/ads or product pitches \rightarrow cap at Minimal unless clear instructional scaffolding.

Quick Decision Tree.

- Are there explicit learning goals or a syllabus? → High.
- Else, are there step-by-step instructions with examples/exercises? → Moderate.
- Else, are there explanatory sections intended to teach basics? → Basic.
- Else, is there any minor instructional element? → Minimal.
- Otherwise → None.

Borderline Examples.

- API reference with examples but no guidance → Minimal to Basic (depending on clarity/examples).
- Blog post explaining concept with analogies and one example → Basic.
- Tutorial with tasks, checkpoints, and solutions → High.
- Product documentation with “Getting Started” and “How-To” flows → Moderate.

Educational Indicators.

- **Learning objectives:** Clear goals for what reader should learn.
- **Skill progression:** Builds from basic to advanced concepts.
- **Examples and practice:** Provides concrete examples or exercises.
- **Knowledge transfer:** Concepts applicable beyond immediate context.

F.3.4 REASONING INDICATORS

What we measure: Presence and quality of logical reasoning, analysis, and explanatory content.

analytical Complex reasoning and systematic analysis.

- Multi-step arguments with logical progression; cause-effect analysis and systematic thinking.
- Considers multiple perspectives or variables; draws conclusions from evidence and reasoning.
- Example: research analysis, complex problem-solving, systematic evaluations.

explanatory Clear explanations with logical flow.

- Explains how or why things work; shows cause-effect relationships clearly.
- Educational reasoning that builds understanding; logical connections between concepts.
- Example: good tutorials, educational content, how-to explanations.

basic reasoning Simple logical connections.

- Some logical connections between ideas; basic explanations of concepts or processes.
- Elementary analytical thinking; simple cause-effect relationships.
- Example: basic explanations, simple arguments, elementary analysis.

minimal Limited reasoning, mostly descriptive.

- Primarily describes what rather than why or how; few logical connections between ideas.
- Mostly factual statements without analysis; little explanatory content.
- Example: basic descriptions, simple factual content, minimal analysis.

none No clear reasoning present.

- Purely descriptive content; simple factual listing without connections.
- Narrative content without analysis; no logical argumentation or explanation.
- Example: simple lists, basic narratives, pure description.

Thinking-Trace Signals.

- Stepwise structure: numbered steps in proofs/derivations/solutions; “First. . . therefore. . . hence. . . so. . .”
- Hypothesis and test: assumptions, intermediate results, counterexamples, sanity checks.
- Tool- or method-calls: named algorithms, theorems, lemmas, or procedures invoked and justified.
- Error analysis or reflection: “we tried X, failed because Y, so we. . .”, “limitations,” “edge cases.”
- Intermediate artifacts: scratch calculations, partial code reasoning, sub-problems and sub-claims.

Disambiguation Rules.

- Explanatory vs. analytical: explanations tell how; analytical shows multi-step inference with evidence and intermediate claims.
- Worked example vs. mere answer: worked examples expose steps and justification; mere answers without steps are not reasoning-rich.
- Procedural vs. reasoning: procedural lists actions; reasoning links actions via logic, evidence, or constraints.

Automation Heuristics.

- Lexical cues: because, therefore, thus, hence, suppose/assume, we conclude, by induction, lemma/theorem/proof, $O(n)$, hypothesis, counterexample.
- Structure cues: presence of proof blocks, derivations (e.g., “Proof.”, “QED”, \TeX environments), multi-step numeric calculations.
- Program reasoning: code comments like `// invariant, // complexity, pre/post-conditions, test reasoning`.
- Thresholding: count reasoning cues per 1k tokens; with ≥ 2 structural cues or ≥ 5 lexical cues \rightarrow at least explanatory; proofs/derivations \rightarrow analytical.

Quick Decision Tree.

- Is there a proof/derivation or multi-step argument with intermediate claims? \rightarrow Analytical.
- Else, does it explain why/how with cause-effect and logical links? \rightarrow Explanatory.
- Else, are there simple logical connections or one-step justifications? \rightarrow Basic reasoning.
- Else, does it mostly describe without connecting ideas? \rightarrow Minimal/None.

Borderline Examples.

- Answer-only solutions (final numeric result without steps) \rightarrow Minimal.
- Step-by-step math solution with intermediate equations \rightarrow Analytical.
- “How it works” article connecting 2–3 causal steps without data \rightarrow Explanatory.
- Troubleshooting log with attempts and justifications \rightarrow Analytical if causal chain is explicit; otherwise Explanatory.

Key Reasoning Patterns to Identify.

- **Cause-effect:** “Because X, therefore Y.”
- **Problem-solution:** Identifies problems and proposes solutions.
- **Comparison:** Analyzes similarities and differences.
- **Logical progression:** Ideas build on previous ideas.
- **Evidence-based conclusions:** Draws conclusions from presented evidence.

F.4 AUDIENCE AND PURPOSE

F.4.1 AUDIENCE LEVEL

What we measure: Intended sophistication level and background knowledge assumptions of the target audience.

expert Highly specialized professional/academic content.

- Assumes deep domain expertise and advanced training; uses technical terminology without explanation.
- Content for practitioners actively working in specialized fields.
- Example: climate modeling methodology in Nature Climate Change, research papers, technical specifications.

advanced Educated adult audience with analytical skills.

- Assumes higher education and critical thinking ability; explains specialized concepts but uses sophisticated language.
- Intellectually challenging but accessible to educated generalists.
- Example: complex climate change analysis in The Atlantic, quality journalism, policy analysis.

general General adult audience.

- Accessible to most educated adults without specialized background; explains technical concepts when introduced.
- Uses clear language while maintaining intellectual substance.
- Example: quality journalism, general interest articles, accessible explanations of complex topics.

beginner Introductory level with minimal prerequisites.

- Explains basic concepts and terminology; builds up from fundamental principles.
- Assumes minimal prior knowledge of the subject area.
- Example: introductory tutorials, beginner guides, basic explanations, getting-started content.

youth Targeted at teenagers and young adults (ages 13–19).

- Age-appropriate complexity with contemporary cultural references.
- Sophisticated enough for developing critical thinking but accessible.
- Example: high school educational content, young adult literature, college prep materials.

children Designed specifically for children.

- Simple language and concepts appropriate for young readers.
- Educational content designed for elementary/middle school levels.
- Example: children’s educational content, elementary school materials, simple explanations for young learners.

Assessment Guidelines.

- **Professional context:** Is this content designed for workplace use vs. general learning?
- **Terminology density:** How much specialized vocabulary is used without explanation?
- **Concept complexity:** How sophisticated are the ideas and their development?
- **Background assumptions:** What education level and domain knowledge does the author assume?

Cross-Linguistic Considerations.

- Expert terminology density varies by language (German allows more compound terms).
- Formality markers differ across cultures.
- Educational level assumptions vary by country’s education system.
- Age-appropriate content differs across cultures.

F.4.2 COMMERCIAL BIAS

What we measure: How much commercial interests influence the objectivity and informational value of content.

none No commercial influence detected. Objective, informational presentation; no promotional language or commercial agenda; focus purely on informing or educating. Example: academic papers, objective journalism, educational content.

minimal Slight commercial context but maintains objectivity. May mention products/services but in informational context; maintains balanced, objective tone. Example: product reviews with balanced analysis, informational articles mentioning relevant products.

moderate Some commercial influence on content. Mix of informational and promotional content; some promotional language but still provides useful information. Example: company blogs with useful information, sponsored content with actual value.

heavy Strong commercial bias throughout. Primarily promotional with some informational elements; heavy use of marketing language and persuasive techniques. Example: marketing articles disguised as information, heavily biased product comparisons.

pure marketing Entirely commercial/promotional content. No genuine informational value beyond promotion; pure marketing copy or advertising material; designed solely to drive sales or conversions. Example: sales pages, pure advertising copy, promotional brochures.

Key Indicators.

- **Language tone:** Objective vs. promotional language.
- **Primary purpose:** Inform vs. persuade/sell.
- **Balance:** Are alternatives/drawbacks mentioned?
- **Call-to-action:** Subtle information vs. obvious sales pitch.

F.4.3 TIME-SENSITIVITY

What we measure: How time-sensitive the content is, i.e., whether its value degrades over time or remains stable.

evergreen Content remains valuable indefinitely. Fundamental concepts, principles, theories; historical information and established facts; skills and techniques that do not change. Example: mathematical proofs, language grammar guides, classical literature analysis, basic cooking techniques.

slowly changing Content remains valuable for years. Best practices that evolve slowly; technical content that updates every few years; cultural and social topics with gradual change. Example: programming language tutorials, academic textbooks, industry standards, educational curricula.

regularly updating Content valuable for months to a year. Industry trends and market analysis; technology reviews and comparisons; policy discussions and current research. Example: software framework guides, business strategies, product reviews, research summaries.

time sensitive Content value degrades quickly. News and current events; time-bound information (prices, schedules, availability); temporary situations or short-term trends. Example: stock prices, weather reports, breaking news, event announcements, sales/promotions.

Key Decision Points.

- **Core question:** If someone reads this in 2 years, will it still be valuable?
- **Update frequency:** How often does this type of information typically change?
- **Temporal references:** Does the content heavily reference “now,” “recently,” “currently”?
- **Subject matter stability:** Is this about unchanging principles or evolving situations?

F.5 SAFETY AND COMPLIANCE

F.5.1 CONTENT SAFETY

What we measure: Presence of inappropriate, harmful, or legally problematic content.

safe Appropriate for all contexts. No concerning content of any type; professional, appropriate language throughout; suitable for general audiences including workplace settings.

mild concerns Minor issues that do not constitute major problems. Occasional mild profanity in context; brief mentions of sensitive topics handled appropriately. Example: historical discussions of sensitive topics, professional content with mild language.

nsfw Not safe for work or general audiences. Explicit sexual content or graphic descriptions; adult themes requiring content warnings; graphic violence or disturbing imagery descriptions. Example: adult content, graphic medical descriptions, explicit violence.

harmful Potentially harmful content requiring careful handling. Content promoting dangerous activities or self-harm; hate speech targeting individuals or groups; violent content glorifying harm to others. Example: self-harm content, hate speech, dangerous “how-to” guides.

illegal Illegal content requiring immediate rejection. Content promoting clearly illegal activities; material that violates laws in major jurisdictions. Example: terrorist content, child exploitation.

Safety Assessment Guidelines.

- **Context matters:** Medical/educational discussions of sensitive topics may be appropriate.
- **Intent matters:** Discussing harmful topics for educational purposes vs. promoting them.
- **Audience consideration:** Content appropriate for experts may not be safe for general audiences.

F.5.2 PII PRESENCE

What we measure: Whether the content contains personally identifiable information that could identify private individuals.

no_pii No personal information detected.

- No names of private individuals; no contact information (emails, phones, addresses); no identification numbers.
- Public figures and officials mentioned by name are acceptable.
- Example: news articles about politicians, technical documentation, general information.

contains_pii Contains potentially identifiable information.

- Names of private individuals (non-public figures).
- Email addresses, phone numbers, physical addresses.
- ID numbers (SSN, passport, driver’s license, employee IDs).
- Medical information about identifiable individuals; financial account information.
- Example: personal blogs with full names, leaked databases, medical case studies with identifying info.

Key Decision Points.

- **Public vs. Private figures:** Politicians, celebrities, CEOs = public (no PII flag); private citizens = PII.
- **Context matters:** Academic paper authors and their institutional emails = typically no PII; personal emails in forums = PII.
- **Aggregated vs. Individual:** Statistical data = no PII; individual records = PII.

F.6 GEOGRAPHIC RELEVANCE

F.6.1 REGIONAL RELEVANCE

What we measure: Primary regional, cultural, or geopolitical sphere(s) that the content relates to, regardless of language used.

Content can be assigned multiple regional labels if it genuinely spans multiple regions. Choose all applicable regions rather than forcing a single primary choice. Always output an array.

european European context (EU and broader Europe). Content about European countries, EU policies, or pan-European topics; European cultural perspectives, social systems, or business practices. Includes EU member states, UK, Switzerland, Norway, Balkans, etc. Example: GDPR compliance, European Parliament elections, Schengen area travel, European football leagues.

north.american North American context. Content about US, Canada, or Mexico; North American cultural perspectives, USMCA/NAFTA region topics. Example: FDA regulations, Silicon Valley tech, NHL, US constitutional law, Canadian healthcare.

east.asian East Asian context. Content about China, Japan, Korea (North/South), Taiwan, Mongolia; East Asian cultural perspectives, Confucian-influenced societies. Example: Gaokao exams, K-pop, Shenzhen tech hub, Japanese work culture, Taiwan semiconductor industry.

south.asian South Asian context. Content about India, Pakistan, Bangladesh, Sri Lanka, Nepal, Bhutan, Afghanistan, Maldives; South Asian cultural perspectives. Example: IIT entrance exams, Bollywood, cricket leagues, monsoon impacts.

southeast.asian Southeast Asian context. Content about ASEAN countries (Indonesia, Thailand, Vietnam, Philippines, Malaysia, Singapore, etc.); Southeast Asian regional perspectives. Example: ASEAN economic community, Indonesian elections, Singapore financial sector, Thai tourism.

middle.eastern Middle Eastern and North African context. Content about Arab states, Iran, Turkey, Israel, North Africa (MENA region); Islamic finance, regional conflicts. Example: Gulf Cooperation Council, OPEC decisions, Middle East peace process, Islamic banking.

sub.saharan.african Sub-Saharan African context. Content about African countries south of the Sahara; African Union topics, sub-Saharan development issues. Example: M-Pesa mobile banking, African Union policies, safari tourism, ubuntu philosophy.

latin.american Latin American context. Content about Central and South America, Caribbean; regional integration (Mercosur, etc.). Example: Mercosur trade, telenovelas, Amazon rainforest, Latin American revolutions.

oceanian Oceanian context. Content about Australia, New Zealand, Pacific Island nations; Oceanian perspectives, Pacific regional issues. Example: ANZAC relations, Pacific Island climate change, Australian mining, Māori culture.

central.asian Central Asian context. Content about Kazakhstan, Uzbekistan, Turkmenistan, Tajikistan, Kyrgyzstan; post-Soviet regional dynamics, Silk Road region. Example: Silk Road initiatives, Caspian Sea resources, post-Soviet transitions.

russian.sphere Russian/Post-Soviet context. Content about Russia, Belarus, and strong Russian influence areas; CIS (Commonwealth of Independent States) topics. Example: Russian federal politics, CIS integration, post-Soviet economic transitions.

global Genuinely international or universal. Content with truly global scope or application; international organizations, worldwide phenomena, global comparisons. Example: UN reports, climate change (global perspective), international standards, pandemic response.

culturally neutral No clear regional focus. Abstract, theoretical, or technical content without regional markers; universal scientific, mathematical, or philosophical content. Example: mathematical proofs, chemical formulas, abstract philosophy, programming concepts.

indeterminate Cannot determine regional relevance. Insufficient content to identify regional focus; mixed or contradictory regional signals; fragment or corrupted content. Example: technical specifications without context, isolated data tables.

Multi-Regional Examples.

- EU-China trade relations → [european, east_asian]
- NAFTA/USMCA impact on Mexican agriculture → [north_american, latin_american]
- Indian diaspora in the Gulf states → [south_asian, middle_eastern]
- Comparative study of healthcare systems globally → [global]

Regional Identification Guidelines. Primary indicators:

- **Geographic references:** Countries, cities, regions, landmarks mentioned.
- **Institutional references:** Governments, companies, universities, organizations specific to region.
- **Cultural markers:** Holidays, customs, cultural phenomena, sports, entertainment.
- **Political/economic systems:** References to regional political structures, economic blocs.
- **Legal/regulatory frameworks:** Region-specific laws, regulations, standards.
- **Language context:** While not determinative, language can provide regional hints.

Important distinctions:

- **Language ≠ Region:** Spanish content about Asian markets = [east_asian], not [latin_american].
- **Company origin vs. topic:** Apple (US company) operating in India = consider actual content focus.
- **Historical vs. current:** Historical content about ancient Rome = [european] if discussing modern implications.
- **Diaspora content:** Content about diaspora communities should include both origin and current regions.

Quality checks:

- If content is in a non-English language but discusses global topics → still mark as [global].
- If content compares multiple regions → mark all regions discussed substantially.
- If content is about a specific place but has universal applications → consider both regional and global tags.

F.6.2 COUNTRY RELEVANCE

What we measure: Which specific country or countries (if any) the content is relevant to, globally.

Always output an array of country names for this property (even when only a single country applies). Use standard country names from any region worldwide (e.g., “germany”, “france”, “united_states”, “united_kingdom”, “china”, “japan”, “brazil”, “india”, “south_africa”, “australia”, “canada”, etc.). The array may also contain the special values `supranational` or `none`.

{COUNTRY_NAME} Content specifically relevant to a single country.

- Content explicitly about that country’s politics, culture, institutions, or regulations.
- Content written from that country’s cultural perspective.
- Content about that country’s cities, companies, institutions, or country-specific topics.
- Only use country names listed in ISO-3166. Use “united_kingdom” instead of “england”, “wales”, etc.
- Example: For “germany” → German election coverage, Bundesliga content, German legal analysis.
- Example: For “japan” → Japanese politics, J-League content, Japanese cultural analysis.

supranational Content focused on supranational entities or regions.

- International organizations, regional blocs, global institutions.
- Content about supranational policies, international organizations, global governance.
- Example: UN resolutions, NATO discussions, EU policy analysis, ASEAN agreements, WTO trade rules.

none Content not specifically relevant to any country.

- Abstract, theoretical, or universal content without geographic specificity.
- Technical/scientific content that applies globally without country focus.
- Example: mathematical proofs, universal scientific principles, abstract philosophical discussions.

Country Identification Criteria.

- **Political content:** Elections, government policies, political parties, political figures specific to the country.
- **Cultural content:** National traditions, cultural phenomena, historical events specific to the country.
- **Institutional references:** Government bodies, national companies, universities specific to the country.
- **Geographic focus:** Cities, regions, landmarks within the country as primary subjects.
- **Legal/regulatory:** Laws, regulations, legal frameworks specific to the country.
- **Economic content:** National economic policies, country-specific market analysis.
- **Sports/media:** National sports leagues, national teams, country-specific media outlets.
- **Social issues:** Social policies, demographic topics, social movements specific to the country.