

PrionNER: A Named Entity Recognition Dataset for Prion Disease Biomedical Literature

An Dao^{1,5*} Nhan Ly^{2*} Thao Tran^{2*} Yuji Matsumoto^{3,4} Akiko Aizawa⁵

¹The University of Tokyo, Tokyo, Japan

²Medical Doctor, Independent Researcher

³ Center for Language AI Research, Tohoku University, Sendai, Japan

⁴ RIKEN Center for Advanced Intelligence Project, Tokyo, Japan

⁵ National Institute of Informatics, Tokyo, Japan

dtan@g.ecc.u-tokyo.ac.jp, trinkanly1996@gmail.com, thaotran1490@gmail.com,
yuji.matsumoto.a4@tohoku.ac.jp, aizawa@nii.ac.jp

Abstract

Prion diseases are rare, rapidly progressive, and fatal neurodegenerative disorders that remain difficult to diagnose, particularly in their early stages because of nonspecific clinical presentations. However, to our knowledge, there is no publicly available prion-disease-focused dataset designed to capture a broad range of clinically relevant entities from the biomedical literature. We introduce **PrionNER**, a manually annotated named entity recognition dataset for prion disease clinical information in PubMed abstracts. The current release comprises 317 abstracts, 2,943 sentences, and 6,955 text-bound entity annotations spanning 15 coarse-grained and 31 fine-grained clinically oriented entity types covering diseases, symptoms, diagnostics, findings, anatomy, treatments, and temporal and statistical evidence. Inter-annotator agreement reaches 81.78 exact-match F1, indicating strong annotation consistency. We benchmark supervised BERT baselines, W2NER, and zero-shot extractors on PrionNER. W2NER is the strongest supervised model, and Gemma-4-31B is the strongest zero-shot model, but the benchmark remains challenging, especially for structurally complex mentions and fine-grained clinically adjacent label distinctions. PrionNER provides a clinically grounded benchmark for prion-disease information extraction and supports research on rare-disease biomedical NLP under low-resource, fine-grained, and non-flat extraction conditions. The dataset, annotation guidelines, and evaluation scripts are available at <https://github.com/daotuanan/PrionNER/>.

1 Introduction

Prion diseases are rapidly progressive, fatal neurodegenerative disorders caused by the misfolding and accumulation of pathological prion pro-

tein (Prusiner, 1998). They remain untreatable while carrying a severe risk of iatrogenic transmission (Centers for Disease Control and Prevention, 2026). Early diagnosis is difficult because these diseases are rare and therefore unfamiliar to many general physicians, while their initial symptoms are often vague and overlap with other psychiatric and neurological conditions (Geschwind, 2015). By the time the disease reaches a clearly recognizable stage, it is often already terminal, limiting opportunities for intervention and increasing the risk of inadvertent spread during routine medical procedures (Vallabh et al., 2020). Consequently, recent research efforts have emphasized earlier diagnosis both to reduce transmission risk and to identify patient cohorts for future therapeutic trials (Inada Shimamura and Satoh, 2025). To achieve this, consolidating scattered clinical knowledge into a unified, foundational dataset is essential.

Named entity recognition (NER) has substantially advanced clinical data extraction by structuring key concepts from biomedical text, including diseases, symptoms, diagnostic tests, and biomarkers (Lee et al., 2020; Gu et al., 2021). This progress is supported by widely used biomedical and clinical corpora such as NCBI Disease (Dogan et al., 2014), BC5CDR (Li et al., 2016), CRAFT (Bada et al., 2012), MedMentions (Mohan and Li, 2019), i2b2/VA (Uzuner et al., 2011), and ShARE/CLEF (Pradhan et al., 2013). However, most existing resources target broad biomedical domains, disease mentions, or general classifications rather than the clinically rich diagnostic evidence needed for rare-disease recognition in biomedical literature. Rare-disease-oriented datasets also focus more narrowly on disease, symptom, sign, or phenotype extraction than on the broader diagnostic schema needed for prion disease (Martínez-deMiguel et al., 2022; Groza et al., 2015; Shyr et al.,

*These authors contributed equally to this work.

2024). As a result, there is still no publicly available prion-focused NER benchmark that captures the heterogeneous evidence clinicians rely on when distinguishing prion disease and its subtypes from related conditions. To address this gap, we introduce **PrionNER**, a manually annotated dataset for named entity recognition in prion disease clinical narratives derived from PubMed abstracts.

PrionNER contains 317 abstracts with 15 coarse-grained and 31 fine-grained entity types, and its pre-adjudication double-annotated test split reaches 81.78 entity-level exact agreement F1. We evaluate PrionNER using both supervised biomedical encoders and zero-shot extraction models. Among supervised models, W2NER is strongest in both coarse-grained and fine-grained flat-ner, reaching 81.86 F1 and 80.46 F1, respectively. Gemma-4-31B is the strongest zero-shot model in flat-ner, reaching 71.41 coarse-grained F1 and 68.41 fine-grained F1. However, non-flat-ner remains difficult for all models, with the best F1 scores of W2NER (supervised) reaching 13.48 in coarse evaluation and 13.70 in fine evaluation. These results show that PrionNER is learnable but still challenging: performance consistently drops from coarse-grained to fine-grained prediction, and the remaining difficulty reflects a long-tailed label distribution, clinically adjacent type distinctions, and nested or discontinuous mentions.

This work provides a foundation for prion-disease information extraction and clinically oriented biomedical NLP. More broadly, PrionNER illustrates a general challenge in rare-disease biomedical NLP: clinically useful extraction often depends on modeling heterogeneous evidence types, fine-grained diagnostic distinctions, and non-flat mention structures. In this sense, the dataset is relevant beyond prion disease itself and can serve as a compact case study for building clinically grounded resources in other specialized biomedical subdomains.

2 Related Work

In the biomedical literature, JNLPBA (Collier et al., 2004), NCBI Disease Corpus (Dogan et al., 2014), BC5CDR (Li et al., 2016), CRAFT (Bada et al., 2012), and MedMentions (Mohan and Li, 2019) provide benchmark resources for entity recognition, normalization, and concept annotation in MEDLINE or PubMed texts. In the clinical domain, i2b2/VA (Uzuner et al., 2011), ShARE/CLEF (Prad-

han et al., 2013), and MedDec (Elgaar et al., 2024) provide corpora from patient reports, clinical notes, and discharge summaries that focus on clinically meaningful entities or decision spans such as problems, tests, treatments, disorder mentions, and medical decisions.

Resources closest to our setting are mainly rare-disease or phenotype-oriented. The HPO corpora (Groza et al., 2015) target phenotype concept recognition and normalization, while RareDis (Martínez-deMiguel et al., 2022) extends this line with annotations for diseases, rare diseases, symptoms, signs, and anaphoric mentions. However, based on our preliminary inspection of the released RareDis annotations, prion-disease coverage remains limited, and these datasets do not provide the broader clinically oriented schema needed for prion-disease information extraction. More broadly, prior resources reflect different tradeoffs in domain breadth and label granularity: corpora such as NCBI Disease and BC5CDR emphasize wider biomedical coverage with a small number of target categories (Dogan et al., 2014; Li et al., 2016), whereas our goal is richer evidence coverage within a single specialized domain.

For prion disease specifically, PDDB (Gehlenborg et al., 2009) is a structured transcriptomic resource rather than a text annotation benchmark. To our knowledge, no publicly available dataset currently targets prion-disease named entity recognition in the biomedical literature with a broad clinically oriented annotation schema. PrionNER fills this gap by providing a manually annotated dataset for prion disease clinical narratives in PubMed abstracts. We next describe how the corpus was collected, filtered, and annotated.

3 Dataset Construction

3.1 Data Sources

We constructed the corpus from PubMed abstracts retrieved with a keyword-based Boolean query over the title and abstract fields; the full query is provided in Appendix Section A.1. The query combined general and subtype-specific prion disease terms, including *Prion Diseases*, *Creutzfeldt-Jakob Disease*, *CJD*, *sporadic CJD*, *familial/genetic CJD*, *variant CJD*, *iatrogenic CJD*, *Kuru*, *Gerstmann-Straussler-Scheinker*, and *Fatal Familial Insomnia (FFI)*, with clinically oriented terms such as *diagnosis*, *clinical*, *symptoms*, *case*, *progression*, and *treatment*. To bias retrieval toward human clinical

Dataset	Domain	Task	Text Source	Entity Coverage	Size
JNLPBA (Collier et al., 2004)	Biomedical	NER	MEDLINE abstracts	Protein, DNA, RNA, cell line, cell type	2,404 abstracts
i2B2/VA 2010 (Uzuner et al., 2011)	Clinical	NER + RE	Patient reports	Problems, tests, treatments, assertions	871 reports
CRAFT (Bada et al., 2012)	Biomedical	Concept Ann.	Full-text articles	Ontology-linked biomedical concepts	97 articles
SHARE/CLEF (Pradhan et al., 2013)	Clinical	NER + Norm.	Clinical reports	Disorder mentions	300 reports
NCBI DISEASE (Dogan et al., 2014)	Biomedical	NER + Norm.	PubMed abstracts	Disease mentions linked to MeSH/OMIM	793 abstracts
HPO CORPORA (Groza et al., 2015)	Rare disease	NER + Norm.	PubMed abstracts	Phenotype/HPO concept mentions	228 abstracts
BC5CDR (Li et al., 2016)	Biomedical	NER + RE	PubMed abstracts	Chemicals, diseases, chemical-disease links	1,500 abstracts
MEDMENTIONS (Mohan and Li, 2019)	Biomedical	Concept Ann.	PubMed abstracts	UMLS-linked biomedical concepts	4,392 abstracts
RAREDIS (Martínez-deMiguel et al., 2022)	Rare disease	NER + RE	Rare-disease texts	Disease, rare disease, symptom, sign, anaphor	1,041 texts
PrionNER (ours)	Prion disease	NER	PubMed abstracts	15 coarse / 31 fine clinically relevant types	317 abstracts

Table 1: Comparison of PrionNER with representative prior datasets and related resources.

narratives, we excluded terms commonly associated with animal or basic-science studies, including *mice*, *mouse*, *rat*, *animal*, *cell*, *protein*, and *in vitro*. The PubMed query returned 3,414 abstracts, and 3,138 remained after basic preprocessing, including removal of records with empty abstracts.

A pilot manual screening of approximately 500 abstracts by two annotators showed that the retrieved set still contained many off-target papers unrelated to clinical prion disease narratives. Common exclusion cases included basic science, animal or other non-human research, non-clinical analyses such as economic or purely epidemiological studies, and papers in which prion disease was not the main focus. Based on this pilot review, we defined an operational criterion for *related* abstracts and used GPT-5.4 (OpenAI, 2026), prompted as described in Appendix Section A, to screen the 3,138 preprocessed abstracts and reduce manual review workload. Under this criterion, abstracts were retained if they primarily concerned human prion disease in a clinical context, including diagnosis, symptoms, disease progression, or treatment. Across the full 3,138-abstract pool, GPT-5.4 predicted 1,304 abstracts as related and 1,834 as unrelated. We then manually reviewed abstracts from the screened pool together with the pilot-screened set to confirm relevance and remove duplicates. In total, we manually reviewed 1,383 abstracts, including 868 abstracts rated as related by GPT-5.4. Within this set, 772 abstracts were labeled related and 611 were labeled not related by human review. The GPT-5.4 model achieved 90.60 accuracy and 97.80 recall for the related class; it missed only 17 relevant abstracts but incorrectly marked 113 irrel-

evant abstracts as related, indicating a high-recall but over-inclusive screening strategy. This bias was appropriate for corpus construction because missing clinically relevant rare-disease abstracts would have been more costly than forwarding some extra candidates to manual review. Although we initially intended to annotate all 772 eligible abstracts, practical time constraints limited the current release to 317 annotated abstracts. We therefore treat the present release as a high-quality first benchmark for this specialized domain rather than an exhaustive survey of all eligible prion-disease abstracts. Appendix Section A also provides the relevance-filter audit.

3.2 Entity Schema

We designed the schema to capture clinically meaningful evidence needed for prion-disease diagnostic reasoning and future knowledge graph construction. It was developed through pilot annotation and iteratively revised after joint review of 20 shared annotated abstracts by two medical doctors with experience in neurological diseases together with one coordinator, so that the labels remained clinically meaningful while also remaining coherent from an NLP annotation perspective. Because the dataset is intended to support future clinical decision-support work for prion disease diagnosis, the schema reflects the evidence integration a clinician uses when evaluating a suspected case. Compared with earlier biomedical and rare-disease corpora, which often focus on narrower mention types such as diseases, chemicals, phenotypes, symptoms, signs, or anaphora, our schema is broader because prion diagnosis depends on combining multiple kinds of

diagnostic evidence (Dogan et al., 2014; Li et al., 2016; Groza et al., 2015; Martínez-deMiguel et al., 2022). To reflect this workflow, we organize the schema into three groups: *Case Input*, *Case Diagnosis*, and *Clinical Course and Context*. *Case Input* captures the information available at a patient’s presentation, including Age, Symptom, Test_name, Sequences, Anatomic_location, and Findings. *Case Diagnosis* captures the interpretation of those inputs, including disease names, subtype labels (Generic_Prion, Sporadic_Prion, Familial_Prion, Acquired_Prion), and alternative conditions considered during evaluation (Differential_Diagnosis). *Clinical Course and Context* captures clinically relevant supporting information, including Treatment, Complication, Time, and Stats. This broad-entity, focused-domain design differs from prior related datasets such as RareDis, which do not combine the same breadth of clinically oriented entities within a single prion-focused annotation schema (Martínez-deMiguel et al., 2022). We include these categories selectively because of their direct clinical relevance and their potential to support diagnosis in realistic settings. The schema comprises 15 coarse-grained types and 31 fine-grained entity types, and annotation was performed at the mention level using minimal span selection. Table 2 provides a short schema summary for the main text, and Appendix Table 8 provides full definitions with representative examples.

3.3 Annotation Workflow and Quality Control

Pilot annotation. Annotation began with a pilot phase of 20 PubMed abstracts. During this stage, two annotators independently labeled entity mentions in order to identify ambiguities in span boundaries, type assignment, and difficult clinical expressions. Disagreements were reviewed jointly, and these pilot annotations were used to establish the initial annotation guidelines before large-scale annotation. To preserve anonymity, we refer to the two annotators as *Annotator 1* and *Annotator 2*.

Training annotation and guideline refinement. Using the pilot guidelines, the remaining training abstracts were annotated individually, with 151 abstracts annotated by *Annotator 1* and 96 abstracts annotated by *Annotator 2*. During this stage, the guidelines were iteratively refined to address recurring ambiguities and improve consistency across the corpus.

Double annotation, agreement measurement,

and final test-set selection. After the training-stage guideline refinement, an additional 70 abstracts were independently annotated by both annotators. Inter-annotator agreement was computed on these independent annotations before any disagreement resolution. Only after agreement measurement did the annotators jointly adjudicate disagreements in this 70-abstract subset to produce a single finalized annotation layer containing 787 sentences and 1,806 text-bound entity annotations. We designated this double-annotated and adjudicated subset as the test split because it is the highest-confidence evaluation subset in the corpus: both annotators labeled these abstracts independently, agreement was measured on the pre-adjudication annotations, and disagreements were then resolved to create the final released test annotations. The clarified adjudication decisions were then propagated back to the training annotations to ensure consistency with the final guideline version, and the final dataset was validated for annotation integrity, including span correctness, label consistency, and the absence of duplicate or conflicting entity annotations.

4 Dataset Statistics

4.1 Corpus Overview

Overall, PrionNER contains 317 abstracts, 2,943 sentences, and 6,955 text-bound entity annotations across the training and test splits. The dataset covers 15 coarse-grained and 31 fine-grained entity types spanning diagnostically relevant clinical evidence in prion disease literature. The training split contains 247 abstracts, and the finalized test split contains 70 abstracts. As described in Section 3.3, the released test split reflects a single merged annotation layer derived from the adjudicated 70-abstract double-annotated subset. To support reproducibility, Appendix Section A, Appendix Section B, and Appendix Sections A.4–A.5 document the PubMed search query, the abstract-screening prompt, the annotation guidelines, and the model settings used in this work. The dataset annotations, annotation guidelines, and evaluation scripts are publicly available at <https://github.com/daotuanan/PrionNER/>.

Coarse-grained Type	Fine-grained Types	Definition Summary
Case Input (Willis, 2006)		
Age	Age	Age or age-at-onset expressions.
Symptom	Symptom	Symptoms, signs, and other explicit clinical manifestations.
Test_name	Imaging_test, Electrophysio_test, Blood_biomarker_test, Genetic_test, Molecular_assay, Autopsy	Diagnostic procedures, assays, and examination types.
Sequences	Imaging_sequence	Imaging acquisition or sequence terms.
Anatomic_location	Anatomic_location	Anatomical sites linked to findings or symptoms.
Findings	Imaging_finding, Autopsy_finding	Radiologic and pathological findings.
Case Diagnosis (Washington State Department of Health, 2022)		
Generic_Prion	Generic_Prion	Prion disease mentions without explicit subtype specification.
Sporadic_Prion	sCJD, sFI, VPSPr	Sporadic prion disease subtypes.
Familial_Prion	fCJD, GSS, FFI	Familial or inherited prion disease subtypes.
Acquired_Prion	vCJD, iCJD, Kuru	Acquired prion disease subtypes.
Differential_Diagnosis	Differential_Diagnosis	Non-prion diseases considered in differential diagnosis.
Clinical Course and Context (Thabane et al., 2013; Willis, 2006; Murad et al., 2023)		
Treatment	Treatment	Therapies, medications, and care strategies.
Complication	Complication	Secondary conditions or end-stage outcomes.
Time	Duration, Time_point	Temporal spans and specific time anchors.
Stats	Sensitivity, Specificity, Prevalence, Incidence	Diagnostic and epidemiological quantitative measures.

Table 2: Short summary of the PrionNER entity schema. Full type definitions and representative examples are provided in Appendix Table 8.

Split	Abstracts	Sentences	Entities
Train	247	2,156	5,149
Test	70	787	1,806
Total	317	2,943	6,955

Table 3: Corpus-level statistics of PrionNER.

4.2 Inter-Annotator Agreement

Inter-annotator agreement. We report agreement on the pre-adjudication independent annotations of the 70-abstract double-annotated subset described in Section 3.3. Agreement is assessed using three complementary symmetric measures: entity-level exact agreement F1, Jaccard similarity, and Cohen’s kappa. Entity-level exact agreement F1 summarizes span-level agreement under the standard NER criterion that an entity counts as correct only when both annotators select the same text span and assign the same entity type (Tjong Kim Sang and De Meulder, 2003). Jaccard similarity measures the overlap between the two annotators’ entity sets as the size of the intersection relative to the union (Jaccard, 1901). Cohen’s kappa is computed on token-level BIO labels to quantify agreement while correcting for agreement expected by chance (Cohen, 1960). On the 70 compared documents, entity-level exact agreement is 81.78 F1. Jaccard similarity over the annotated entity sets is 69.17, and token-level BIO Cohen’s kappa is 81.05. Together, these metrics provide complementary views of annotation consistency at both the span and token levels, and Appendix Table 17 provides the full per-label agreement breakdown.

Entity Type	Train n	Train %	Test n	Test %
<i>Most Frequent</i>				
Symptom	791	16.99	248	15.03
Generic_Prion	741	15.92	233	14.12
Anatomic_location	646	13.88	223	13.52
Imaging_test	261	5.61	75	4.55
Duration	217	4.66	49	2.97
<i>Least Frequent Observed</i>				
sFI	3	0.06	0	0.00
Prevalence	4	0.09	1	0.06
Incidence	6	0.13	4	0.24
Sensitivity	0	0.00	4	0.24
Specificity	0	0.00	4	0.24

Table 4: Selected most frequent and least frequent observed schema-defined fine-grained entity types in the train and test splits. The full split-wise fine-grained distribution is provided in Appendix Table 18.

4.3 Label Distribution and Structural Complexity

Tables 3 and 4 show that PrionNER covers a broad range of diagnostically relevant evidence while remaining strongly long-tailed. The most frequent fine-grained labels are *Symptom*, *Generic_Prion*, and *Anatomic_location*, whereas several clinically meaningful types remain rare. The adjudicated test set also retains structurally complex mentions, including 34 discontinuous entities, 80 nested pairs, and 2 overlapping pairs. Figure 1 gives representative examples of these discontinuous, nested, and overlapping annotations from the test set. Full split-wise label counts and structural statistics are provided in Appendix Section C. With the corpus characteristics established, we next describe the experimental setup used to benchmark the dataset.

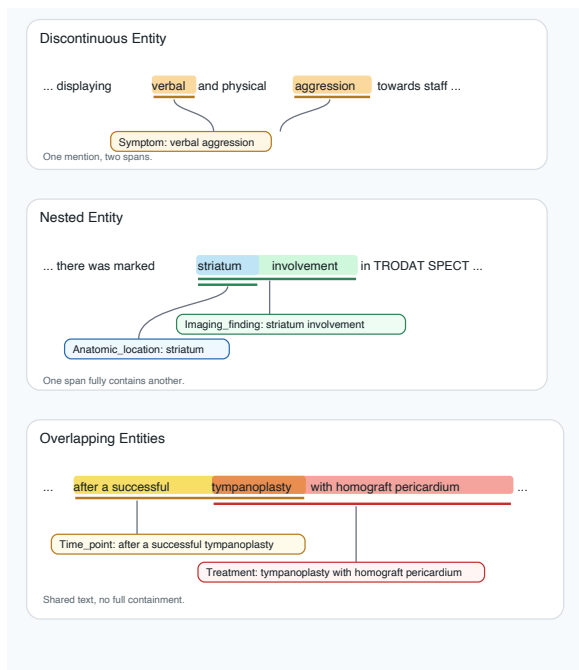


Figure 1: Representative discontinuous, nested, and overlapping entity annotations from the test set.

5 Experiments

5.1 Experimental Settings

We formulate PrionNER as a sequence-labeling task over schema-defined entity mentions. For supervised models, this sequence-labeling setup is implemented with standard BIO tagging, where each entity type is paired with B- and I- prefixes together with the outside label O (e.g., B-Symptom, I-Symptom, O). Models are trained on the training split and evaluated on the test set, with corpus statistics summarized in Table 3. Because the corpus is relatively small and highly long-tailed, we do not introduce a separate development split, as doing so would further reduce training coverage for already sparse labels. Instead, we use fixed model configurations applied uniformly within each model family, and no hyperparameters are selected based on test-set performance. We evaluate the corpus under two structural settings. In the `flat-ner` setting, both training and evaluation are restricted to the flat-compatible portion of the corpus, so only contiguous, non-overlapping gold entities are included. In the `non-flat-ner` setting, evaluation is restricted to the structurally complex portion of the corpus, so the gold annotations include only nested, discontinuous, and overlapping entities. We evaluate the dataset under two label-granularity settings. In the *fine-grained* setting, models are trained and evalu-

ated using the 31 schema-defined entity types. In the *coarse-grained* setting, models are trained and evaluated separately using the 15 coarse-grained categories. This coarse setting is implemented as a separate task rather than by collapsing fine-grained predictions at evaluation time. We report entity-level precision, recall, and F1-score using exact span-and-label matching. Our primary metric is micro-averaged F1 on the test set, and we report it for both coarse-grained and fine-grained settings. The search query and screening prompt are documented in Appendix Section A, and model configurations are documented in Appendix Sections A.4 and A.5, so that the dataset construction and evaluation pipeline can be reproduced more directly.

5.2 Models

5.2.1 Supervised Models

We evaluate supervised baselines fine-tuned on the PrionNER training data. These include BioBERT (Lee et al., 2020), ClinicalBERT (Alsentzer et al., 2019), and PubMedBERT (Gu et al., 2021), which are widely used pretrained encoders for biomedical and clinical text mining, together with W2NER (Li et al., 2022) as an additional structured supervised baseline. We include W2NER because it is a supervised model that can also handle the `non-flat-ner` setting, allowing a supervised comparison beyond standard BIO-style flat tagging. For W2NER, we use PubMedBERT embeddings as the underlying biomedical text representation. The encoder baselines are implemented as BIO-based token classification systems and trained under the same data and model-selection protocol. The BERT baselines operate only in the standard `flat-ner` setting, whereas W2NER is additionally evaluated in the `non-flat-ner` setting. Detailed checkpoints and training hyperparameters are described in Appendix Section A.4.

5.2.2 Zero-shot Models

We additionally evaluate extraction performance using model families that do not require task-specific supervised fine-tuning on PrionNER. These models include OpenAI GPT-5.4 (OpenAI, 2026), Gemma 4 models (Team et al., 2024), the GLiNER2 variants GLiNER2-short and GLiNER2-def (Zaratianna et al., 2025), and GLiNER-BioMed (Yazdani et al., 2025). Because these systems typically return entity strings rather than reliable character offsets, their outputs require a determinis-

Model	Coarse			Fine		
	P	R	F1	P	R	F1
<i>Supervised</i>						
W2NER	84.23	79.63	81.86	83.07	78.01	80.46
PubMedBERT	81.88	81.37	81.62	78.93	80.06	79.49
BioBERT	80.96	81.47	81.22	78.61	78.61	78.61
ClinicalBERT	79.45	79.10	79.28	78.29	78.73	78.51
<i>Zero-shot</i>						
GPT-5.4	72.50	58.81	64.94	72.01	58.42	64.50
Gemma-4-31B	78.33	65.61	71.41	80.14	59.67	68.41
Gemma-4-26B	77.43	55.25	64.48	79.71	46.67	58.87
GLiNER2-def	67.62	59.41	63.25	73.29	43.10	54.28
GLiNER-BioMed	29.50	42.71	34.90	28.99	45.35	35.37
GLiNER2-short	32.63	32.67	32.65	39.51	34.79	37.00

Table 5: flat-ner results for all supervised and zero-shot models on PrionNER under coarse-grained and fine-grained label spaces. Best scores are bolded separately within the supervised and zero-shot sections.

tic span-alignment step before exact-match evaluation. For GLiNER2, we consider two input modes: GLiNER2-short, which provides only entity type names, and GLiNER2-def, which provides entity type names together with their definitions. GLiNER-BioMed is evaluated only in the *short* setting. In our current GLiNER-based pipeline, non-flat behavior can be represented for nested and overlapping entities through multiple contiguous spans, but the model does not support explicit discontinuous multi-span entity objects. We also explored additional small local biomedical LLMs, including Llama3-OpenBioLLM-8B and BioMistral-7B, but found them unreliable for strict exact-span extraction in this pipeline. This setting allows us to assess how well general or biomedical zero-shot extractors transfer to rare-disease clinical literature without supervised adaptation. Detailed model checkpoints, prompting, and inference settings are described in Appendix Section A.5.

5.3 Main Results

Tables 5 and 6 report the main results under coarse-grained and fine-grained label spaces in the flat-ner and non-flat-ner settings. All models are evaluated in flat-ner, whereas non-flat-ner is reported only for W2NER and the zero-shot models because the BERT baselines operate only in the standard flat setting. This design keeps discontinuous and other structurally complex gold annotations in the benchmark rather than discarding them.

In the flat-ner setting, W2NER is the strongest supervised model in both label spaces, reaching

Model	Coarse			Fine		
	P	R	F1	P	R	F1
<i>Supervised / Structured</i>						
W2NER	12.40	14.81	13.48	12.74	14.81	13.70
<i>Zero-shot</i>						
GPT-5.4	4.19	9.63	5.84	4.19	9.63	5.84
Gemma-4-31B	4.82	12.59	6.97	4.96	9.63	6.55
Gemma-4-26B	5.76	11.85	7.75	0.56	0.74	0.63
GLiNER2-def	3.53	4.44	3.93	2.78	0.74	1.17
GLiNER-BioMed	0.00	0.00	0.00	0.00	0.00	0.00
GLiNER2-short	2.07	2.22	2.14	2.90	2.96	2.93

Table 6: non-flat-ner results on PrionNER under coarse-grained and fine-grained label spaces. This setting includes only nested, discontinuous, and overlapping entities. Best scores are bolded separately within the supervised/structured and zero-shot sections.

81.86 F1 in the coarse-grained setting and 80.46 F1 in the fine-grained setting. Among the zero-shot systems, Gemma-4-31B performs best, with 71.41 coarse-grained F1 and 68.41 fine-grained F1.

A notable pattern in flat-ner is the relatively high precision of the strongest zero-shot models under the span-recovery pipeline described in Appendix Section A.5. Gemma-4-31B reaches 78.33 precision in the coarse-grained setting and 80.14 precision in the fine-grained setting, suggesting that the schema is semantically coherent enough for strong zero-shot label assignment from definitions and task instructions alone, even though recall remains well below the supervised baselines. Appendix Sections A.5 and D.3 further suggest that the strongest zero-shot model is limited more by omission than by label precision: Gemma-4-31B remains close to PubMedBERT in entity-only precision while trailing much more clearly in recall.

In the non-flat-ner setting, performance drops sharply for all models, but W2NER remains clearly stronger than the zero-shot systems, reaching 13.48 coarse-grained F1 with 12.40 precision and 14.81 recall, and 13.70 fine-grained F1 with 12.74 precision and 14.81 recall. Among the zero-shot systems, the best coarse-grained non-flat-ner score is 7.75 for Gemma-4-26B, while the best fine-grained score is 6.55 for Gemma-4-31B. GPT-5.4 is second among these models in the fine-grained non-flat-ner setting with 5.84 F1 and remains competitive in the coarse-grained setting, but it does not lead there. Within the GLiNER family, coarse-grained non-flat-ner performance improves slightly relative to fine-grained evaluation, but the GLiNER variants remain clearly behind the

top LLMs, and GLiNER-BioMed fails to recover any non-flat entities in this setting.

Overall, fine-grained prediction is generally harder than coarse-grained prediction, confirming that finer label distinctions remain substantially harder, especially once structurally complex mentions are retained in evaluation. Appendix Figure 2 further visualizes the fine-grained precision–recall trade-offs across the evaluated models.

5.4 Per-Type Analysis

Appendix Figure 3 summarizes the per-type fine-grained F1 patterns for the main supervised and zero-shot models, and Appendix Table 16 gives the full scores. Performance is uneven across the schema. Supervised models dominate many frequent labels such as *Anatomic_location*, *Generic_Prion*, *Imaging_test*, *Symptom*, and *sCJD*, which helps explain their strong aggregate performance. Zero-shot models remain competitive on several semantically clearer or clinically distinctive labels, with GPT-5.4 or Gemma 4 achieving the best score for types such as *FFI*, *Electrophysio_test*, *fCJD*, *Genetic_test*, *Molecular_assay*, and *Differential_Diagnosis*. GLiNER2-def is a notable exception within the GLiNER family, reaching 1.00 F1 on *Sensitivity*, but sparse or boundary-sensitive labels such as *Time_point*, *iCJD*, *Imaging_finding*, and especially *Prevalence* remain difficult for nearly all systems. These extreme per-type values should be interpreted cautiously for very rare labels, where isolated perfect or zero scores may reflect only a handful of test mentions rather than stable schema-wide behavior.

6 Discussions

The main results and per-type patterns motivate a broader discussion of why the benchmark remains difficult and how well it aligns with prion-disease clinical evidence.

6.1 Remaining Challenges

First, the entity distribution is long-tailed. As shown in Table 4, a small number of frequent labels account for a large share of mentions: *Symptom*, *Generic_Prion*, and *Anatomic_location* account for 791, 741, and 646 mentions in the training split and 248, 233, and 223 in the test split, respectively. At the same time, many clinically meaningful categories remain sparse, including *Prevalence* (4 in the training split, 1 in the test split) and *Sensitivity* and

Specificity (0 in the training split and 4 each in the test split). This skew is not merely a corpus artifact but reflects the structure of real diagnostic reporting, where a few common evidence types dominate while rarer findings and subtypes remain clinically important. Difficulty also reflects lexical diversity rather than frequency alone: even common labels such as *Symptom*, *Anatomic_location*, and *Duration* cover hundreds of distinct normalized surface forms in the training split. There is also substantial train–test surface-form mismatch, so models must generalize beyond memorized mention dictionaries even for relatively common categories.

Second, the schema requires distinctions among diagnostically adjacent categories, including prion subtypes, test names, findings, temporal expressions, and differential diagnoses, so errors are not only boundary errors but often clinically meaningful type confusions. In practice, subtype mentions such as *fCJD*, *iCJD*, and “corneal transplant-related CJD” are sometimes predicted as *Generic_Prion*. Others arise when late-stage downstream consequences such as *dysphagia* or *death* appear lexically symptom-like even though they are better labeled as *Complication* in context. The imaging cluster is also challenging, with *Imaging_test*, *Imaging_sequence*, and *Imaging_finding* often requiring fine-grained distinction among modality names, acquisition terms, and radiologic abnormalities. Appendix Section D.3 provides a more detailed entity-only confusion analysis with illustrative examples.

Third, the corpus contains discontinuous, nested, and overlapping mentions that are difficult for standard extraction pipelines. This difficulty is partly a consequence of dataset design itself: PrionNER preserves clinically adjacent labels and non-flat mention structures because simplifying them away would reduce the realism of the target task. At the structural level, the training split contains 97 discontinuous entities, 235 nested pairs, and 5 overlapping pairs, while the test split contains 34 discontinuous entities, 80 nested pairs, and 2 overlapping pairs. In practice, this means that most structural difficulty comes from nesting rather than true overlap. Appendix Tables 10–13 provide the full structural breakdown, and Figure 1 illustrates representative examples from the test set.

6.2 Clinical Alignment

PrionNER captures key prion-disease characteristics, including symptoms, tests, and disease pro-

gression described in the academic literature and clinical guideline documents, rather than arbitrary lexical patterns (Washington State Department of Health, 2022; Centers for Disease Control and Prevention, 2026). CJD-related terminology is strongly represented in the labeled terms, with *CJD* appearing 576 times, which is consistent with the prominence of Creutzfeldt-Jakob disease in scientific reports. The Symptom entity (1,279 mentions) aligns closely with current medical guidance and captures hallmark clinical features such as *dementia* (79), *myoclonus* (79), and *ataxia* (47). The dataset also preserves rarer but clinically important prion-disease subtypes and variants such as FFI (85), GSS (63), and the Heidenhain variant (4). Finally, the Duration entity (290 mentions) captures meaningful disease-course information, with frequent expressions such as *within a year* (7) and *12 months* (7), reflecting the typical one-year survival duration described in the clinical literature. Taken together, these patterns show that PrionNER aligns well with expert knowledge and published guidelines. This makes the dataset useful not only for benchmarking NER systems but also for downstream knowledge extraction and future diagnostic-support settings.

7 Conclusions

In this paper, we introduce PrionNER, a manually annotated named entity recognition dataset for prion disease biomedical literature derived from 317 PubMed abstracts. We present a clinically grounded schema that captures fine-grained diagnostic evidence and non-flat entity structure, together with a benchmark spanning supervised and zero-shot extraction settings. PrionNER provides a reliable annotation resource, with 81.78 entity-level exact agreement F1 on the pre-adjudication double-annotated test split. Our experiments show that W2NER provides the strongest supervised results, PubMedBERT is the strongest BERT baseline, and Gemma-4-31B is the strongest zero-shot model, while the remaining performance gaps confirm that PrionNER is a useful but challenging benchmark for prion-disease information extraction. More broadly, the benchmark surfaces challenges that are likely to recur in other rare-disease settings, including sparse but clinically important labels, fine-grained diagnostic distinctions, and nested or discontinuous spans. We hope PrionNER will support future research on rare-disease information extrac-

tion, structured knowledge construction, and clinically oriented biomedical NLP, including relation extraction, document-level evidence consolidation, and normalization or retrieval settings that connect extracted mentions to structured rare-disease knowledge resources.

Limitations

PrionNER has several limitations. First, it is built from 317 PubMed abstracts, which limits corpus size and coverage. Second, its focus on prion diseases may reduce generalizability to other disorders. Third, the data come from biomedical abstracts rather than clinical notes, so the language is more summarized and structured than real-world records. Future work should expand the corpus, improve coverage of sparse labels, and incorporate additional clinical text sources.

Ethics Statement

This study uses publicly available PubMed abstracts and does not involve human subjects or protected health information. No identifiable personal data were collected or processed, so institutional review board approval was not required.

Acknowledgement

This work was supported by Cross-ministerial Strategic Innovation Promotion Program (SIP) on “Integrated Health Care System” Grant Number JPJ012425.

References

- Emily Alsentzer, John Murphy, William Boag, Weihung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. Publicly available clinical bert embeddings. In *Proceedings of the 2nd clinical natural language processing workshop*, pages 72–78.
- Michael Bada, Miriam Eckert, Donald Evans, Kristin Garcia, Krista Shipley, Dmitry Sitnikov, William A Baumgartner Jr, K Bretonnel Cohen, Karin Verspoor, Judith A Blake, et al. 2012. Concept annotation in the craft corpus. *BMC bioinformatics*, 13(1):161.
- Rebeca Benavente and Rodrigo Morales. 2024. Therapeutic perspectives for prion diseases in humans and animals. *PLoS pathogens*, 20(12):e1012676.
- Centers for Disease Control and Prevention. 2026. About prion diseases. <https://www.cdc.gov/prions/about/index.html>. Accessed: 2026-04-15.

- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Nigel Collier, Tomoko Ohta, Yoshimasa Tsuruoka, Yuka Tateisi, and Jin-Dong Kim. 2004. Introduction to the bio-entity recognition task at JNLPBA. In *Proceedings of the International Joint Workshop on Natural Language Processing in Biomedicine and Its Applications*, pages 70–75.
- Rezarta Islamaj Dogan, Robert Leaman, and Zhiyong Lu. 2014. [Ncbi disease corpus: a resource for disease name recognition and concept normalization](#). *Journal of Biomedical Informatics*, 47:1–10.
- Mohamed Elgaar, Jiali Cheng, Nidhi Vakil, Hadi Amiri, and Leo Anthony Celi. 2024. Meddec: A dataset for extracting medical decisions from discharge summaries. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 16442–16455.
- Nils Gehlenborg, Daehee Hwang, Inyoul Y Lee, Hyun-ae Yoo, David Baxter, Brianne Petritis, Rose Pitstick, Bruz Marzolf, Stephen J DeArmond, George A Carlson, et al. 2009. The prion disease database: a comprehensive transcriptome resource for systems biology research in prion diseases. *Database*, 2009:bap011.
- Michael D. Geschwind. 2015. Prion diseases. *CONTINUUM: Lifelong Learning in Neurology*, 21(6, Neuroinfectious Disease):1612–1638.
- Tudor Groza, Sebastian Köhler, Sandra Doelken, Nigel Collier, Anika Oellrich, Damian Smedley, Francisco M Couto, Gareth Baynam, Andreas Zankl, and Peter N Robinson. 2015. Automatic concept recognition using the human phenotype ontology reference and test suite corpora. *Database*, 2015:bav005.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. Domain-specific language model pretraining for biomedical natural language processing. *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23.
- Mika Inada Shimamura and Katsuya Satoh. 2025. [Challenges and revisions in diagnostic criteria: Advancing early detection of prion diseases](#). *International Journal of Molecular Sciences*, 26(5):2037.
- Paul Jaccard. 1901. Étude comparative de la distribution florale dans une portion des alpes et du jura. *Bulletin de la Société Vaudoise des Sciences Naturelles*, 37:547–579.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Jiao Li, Yueping Sun, Robin J Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn J Mattingly, Thomas C Wieggers, and Zhiyong Lu. 2016. Biocreative v cdr task corpus: a resource for chemical disease relation extraction. *Database*, 2016.
- Jingye Li, Hao Fei, Jiang Liu, Shengqiong Wu, Meishan Zhang, Chong Teng, Donghong Ji, and Fei Li. 2022. Unified named entity recognition as word-word relation classification. In *proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 10965–10973.
- Claudia Martínez-deMiguel, Isabel Segura-Bedmar, Esteban Chacón-Solano, and Sara Guerrero-Aspizua. 2022. The raredis corpus: a corpus annotated with rare diseases, their signs and symptoms. *Journal of biomedical informatics*, 125:103961.
- Sunil Mohan and Donghui Li. 2019. Medmentions: A large biomedical corpus annotated with umls concepts. *arXiv preprint arXiv:1902.09476*.
- M Hassan Murad, Lifeng Lin, Haitao Chu, Bashar Hasan, Reem A Alsibai, Alzhras S Abbas, Reem A Mustafa, and Zhen Wang. 2023. The association of sensitivity and specificity with disease prevalence: analysis of 6909 studies of diagnostic test accuracy. *Cmaj*, 195(27):E925–E931.
- OpenAI. 2026. Introducing GPT-5.4. <https://openai.com/index/introducing-gpt-5-4/>. Accessed: 2026-04-11.
- Sameer Pradhan, Noemie Elhadad, Brett R South, David Martinez, Lee M Christensen, Amy Vogel, Hanna Suominen, Wendy W Chapman, and Guergana K Savova. 2013. Task 1: Share/clef ehealth evaluation lab 2013. *CLEF (working notes)*, 1179.
- Stanley B Prusiner. 1998. Prions. *Proceedings of the National Academy of Sciences*, 95(23):13363–13383.
- Andrea Saltelli, Ksenia Aleksankina, William Becker, Pamela Fennell, Federico Ferretti, Niels Holst, Sushan Li, and Qiongli Wu. 2019. Why so many published sensitivity analyses are false: A systematic review of sensitivity analysis practices. *Environmental modelling & software*, 114:29–39.
- Cathy Shyr, Yan Hu, Lisa Bastarache, Alex Cheng, Rizwan Hamid, Paul Harris, and Hua Xu. 2024. [Identifying and extracting rare diseases and their phenotypes with large language models](#). *Journal of Healthcare Informatics Research*, 8(2):438–461.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- S Tenny and MR Hoffman. 2017. Prevalence.

- Lehana Thabane, Lawrence Mbuagbaw, Shiyuan Zhang, Zainab Samaan, Maura Marcucci, Chenglin Ye, Marroon Thabane, Lora Giangregorio, Brittany Dennis, Daisy Kosa, et al. 2013. A tutorial on sensitivity analyses in clinical trials: the what, why, when and how. *BMC medical research methodology*, 13(1):92.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Ozlem Uzuner, Brett R. South, Shuying Shen, and Scott L. DuVall. 2011. 2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556.
- Sonia M Vallabh, Eric V Minikel, Stuart L Schreiber, and Eric S Lander. 2020. Towards a treatment for genetic prion disease: trials and biomarkers. *The Lancet Neurology*, 19(4):361–368.
- Washington State Department of Health. 2022. Prion disease reporting and investigation guideline. <https://doh.wa.gov/sites/default/files/2025-08/420-069-Guideline-Prion.pdf>. Last revised December 2022; accessed 2026-04-15.
- Marjorie Canfield Willis. 2006. *Medical terminology: The language of health care*. Lippincott Williams & Wilkins.
- Anthony Yazdani, Ihor Stepanov, and Douglas Teodoro. 2025. Gliner-biomed: A suite of efficient models for open biomedical named entity recognition. *arXiv preprint arXiv:2504.00676*.
- Urchade Zaratiana, Gil Pasternak, Oliver Boyd, George Hurn-Maloney, and Ash Lewis. 2025. Gliner2: An efficient multi-task information extraction system with schema-driven interface. *arXiv preprint arXiv:2507.18546*.

A Data Collection and Model Setup

A.1 PubMed Search Query

The following Boolean query is used to retrieve candidate abstracts from PubMed:

```
(
  "Prion Diseases"[Title/Abstract] OR
  "Creutzfeldt-Jakob Disease"[Title/Abstract] OR
  "CJD"[Title/Abstract] OR
  "sporadic CJD"[Title/Abstract] OR
  "familial CJD"[Title/Abstract] OR
  "genetic CJD"[Title/Abstract] OR
  "variant CJD"[Title/Abstract] OR
  "iatrogenic CJD"[Title/Abstract] OR
  "Kuru"[Title/Abstract] OR
  "Gerstmann-Straussler-Scheinker"[Title/Abstract] OR
  "Fatal Familial Insomnia"[Title/Abstract] OR
  "FFI"[Title/Abstract]
)
AND
(
  diagnosis[Title/Abstract] OR
  clinical[Title/Abstract] OR
  symptoms[Title/Abstract] OR
  case[Title/Abstract] OR
  progression[Title/Abstract] OR
  treatment[Title/Abstract]
)
NOT
(
  mice[Title/Abstract] OR
  mouse[Title/Abstract] OR
  rat[Title/Abstract] OR
  animal[Title/Abstract] OR
  cell[Title/Abstract] OR
  protein[Title/Abstract] OR
  in vitro[Title/Abstract]
)
```

A.2 Prompt Used for Abstract Relevance Screening

Before applying model-based screening to the full retrieval set, two annotators manually screened approximately 500 abstracts and observed that a large fraction of the keyword-matched results were not actually suitable for corpus construction. The most common false matches were papers centered on basic science, animal or other non-human studies, non-clinical analyses such as economic or purely epidemiological reports, or papers in which prion disease was only a secondary topic. We used this pilot review to formulate an operational definition of *related* abstracts for automated screening.

The GPT-5.4 screening prompt asks the model to determine whether an abstract is relevant to *human prion diseases in a clinical context*, based only on the abstract text. Relevant abstracts are defined as those primarily focused on diagnosis, symptoms, disease progression, or treatment in human prion disease. The prompt explicitly instructed the model to reject abstracts centered on basic science, animal or non-human studies, protein mechanisms without clinical human focus, unrelated primary diseases, or non-clinical reports such as economic, policy, surveillance, or purely epidemiological analyses.

The model is required to return a strict JSON object with three fields: 'is_relevant', 'reason', and 'evidence_spans'. If relevance checking is enabled, we verify that the response is a JSON object with is_relevant as a Boolean value, reason as a non-empty string, and evidence_spans as a list. We then retain only evidence spans that can be matched to the source abstract after light text normalization. A simplified version of the prompt structure is shown below.

```
System: You are a biomedical abstract screening
assistant.
Decide whether the abstract is relevant to human
prion
diseases in a clinical context.
Return one strict JSON object only.
```

User:

- Use the provided relevance definition.
- Judge based only on the abstract text.
- Exclude basic science, animal, non-human, unrelated-topic, and non-clinical abstracts.
- Return JSON only in the form:

```
{
  "is_relevant": true,
  "reason": "<short reason>",
  "evidence_spans": [
    "<short exact span from abstract>"
  ]
}
```

A.3 Audit of the Relevance Filter

After the initial retrieval and preprocessing stage produced 3,138 abstracts, GPT-5.4 screened the full set with the relevance definition above and labeled 1,304 abstracts as related and 1,834 as unrelated. We then manually reviewed abstracts from the screened pool together with the earlier pilot screening results to confirm relevance and remove duplicates. In total, this yielded 1,383 manually reviewed abstracts for auditing the screening process, including 868 abstracts rated as related by GPT-5.4.

Within this manually reviewed set, 772 abstracts were judged truly related and 611 were judged truly unrelated. GPT-5.4 predicted 868 abstracts as related and 515 as unrelated. Among the 868 abstracts rated as related by GPT-5.4, 755 were truly related and 113 were actually unrelated, corresponding to a precision of 86.98 for the related class. The resulting confusion pattern was strongly asymmetric: only 17 truly related abstracts were missed, whereas 113 unrelated abstracts were incorrectly marked as related. This shows that GPT-5.4 was aggressive in calling abstracts relevant.

On this 1,383-abstract audit set, GPT-5.4 achieved 90.60 accuracy and 89.65 balanced accuracy. For the related class, precision was 86.98, recall was 97.80, and F1 was 92.07. For the unrelated class, precision was 96.70, recall was 81.51, and F1 was 88.45. These headline metrics indicate that about 90.60% of all abstracts were classified correctly. Balanced accuracy of 89.65% is a useful summary here because it averages recall on the related and unrelated classes, which is important given the model’s asymmetric behavior. For the related class, GPT-5.4 was correct 86.98% of the time when it predicted *related*, and it recovered 97.80% of all truly related abstracts, indicating that it was very effective at not missing relevant papers. For the unrelated class, GPT-5.4 was correct 96.70% of the time when it predicted *not related*, but its recall for that class was lower at 81.51%, meaning that it did not label unrelated abstracts as *not related* often enough. In practical terms, GPT-5.4 functioned as a strong high-recall relevance screener: it is well suited to settings where the priority is to avoid missing relevant abstracts, but it is weaker at filtering out irrelevant abstracts cleanly. In plain terms, the model tends to over-include, which may be acceptable for screening but increases downstream manual review effort and can introduce more false positives for later NER annotation.

A.4 Supervised Model Details

The supervised BERT model aliases and checkpoints are as follows: `biobert = dmis-lab/biobert-base-cased-v1.2`, `clinicalbert = emilyalsentzer/Bio_ClinicalBERT`, and `pubmedbert = microsoft/BiomedNLP-PubMedBERT-base-uncased-abstract-fulltext`. To avoid further reducing training coverage for rare labels, we do not create a separate development split; instead, we fix the training hyperparameters in advance and apply the same settings across the BERT baselines. No hyperparameter choices are made based on test-set results. All BERT models use the same hyperparameters: `max_length = 512`, `num_train_epochs = 5.0`, `learning_rate = 5e-05`, `weight_decay = 0.01`, `train_batch_size = 8`, `eval_batch_size = 16`, and `seed = 42`. BERT flat-ner results for this 5-epoch setup are as follows. In coarse-grained evaluation, BioBERT reaches 80.96 precision, 81.47 recall, and 81.22 F1; ClinicalBERT reaches 79.45 precision, 79.10

recall, and 79.28 F1; and PubMedBERT reaches 81.88 precision, 81.37 recall, and 81.62 F1. In fine-grained evaluation, BioBERT reaches 78.61 precision, 78.61 recall, and 78.61 F1; ClinicalBERT reaches 78.29 precision, 78.73 recall, and 78.51 F1; and PubMedBERT reaches 78.93 precision, 80.06 recall, and 79.49 F1. For W2NER (Li et al., 2022), we use PubMedBERT embeddings as the encoder input representation and train for 10 epochs. We include it specifically because it remains supervised while also supporting the non-flat-ner setting, unlike the BIO-based BERT baselines. We additionally report aggregate W2NER runs in both flat-ner and non-flat-ner settings. In coarse-grained flat-ner, W2NER reaches 84.23 precision, 79.63 recall, and 81.86 F1; in fine-grained flat-ner, it reaches 83.07 precision, 78.01 recall, and 80.46 F1. For non-flat-ner, W2NER reaches 12.40 precision, 14.81 recall, and 13.48 F1 in the coarse-grained setting, and 12.74 precision, 14.81 recall, and 13.70 F1 in the fine-grained setting.

A.5 Zero-shot Prompting and Inference Details

For zero-shot evaluation, we apply the GLiNER2 variants GLiNER2-short and GLiNER2-def (Zaratiána et al., 2025) together with GLiNER-BioMed (Yazdani et al., 2025), without supervised fine-tuning on PrionNER. For both GLiNER2-short and GLiNER2-def, we use the checkpoint `fastino/gliner2-large-v1`; for GLiNER-BioMed, we use `Ihor/gliner-biomed-base-v1.0`. GLiNER2-short supplies only the entity type names, whereas GLiNER2-def supplies both the entity type names and short type definitions. GLiNER-BioMed is evaluated only with entity type names. In our current GLiNER-based pipeline, non-flat behavior can be represented for nested and overlapping entities through multiple contiguous spans, but the model does not support explicit discontinuous multi-span entity objects. We also tested smaller biomedical LLMs, including aaditya/Llama3-OpenBioLLM-8B and BioMistral/BioMistral-7B, but excluded them from the main comparison because in our pipeline they frequently produced chat-style outputs, malformed or truncated JSON, corrupted labels, and non-literal spans, making them unreliable for strict exact-span NER extraction. The models are run at the abstract level.

For the LLM-based setting, we evaluate OpenAI GPT-5.4 (OpenAI, 2026), google/gemma-4-31B-it, and google/gemma-4-26B-A4B-it from the Gemma 4 family (Team et al., 2024) using a fixed prompt template that introduces the task, specifies the target entity schema, and requests structured outputs. Each input instance is processed at the abstract level, and the model outputs are required to follow a constrained format such as JSON with entity text.

After inference, we apply a verification and normalization pipeline before evaluation. For NER outputs, we first verify the output shape: the model response must be a JSON object containing an entities list, and for OpenAI/Gemma outputs the returned text field must match the input abstract after collapsing whitespace. We then validate entity labels against the schema, rejecting invalid coarse_type/fine_type combinations. In limited cases, we repair partially correct predictions: in the coarse-only schema setting, fine_type is normalized to the same value as coarse_type, and if a predicted fine_type is valid but the coarse_type is missing or incorrect, we infer the corresponding coarse label from the schema. Next, we verify that each predicted mention can be aligned back to the source text and attach start/end offsets with deduplication. This alignment step is necessary for the zero-shot systems because LLM- and GLiNER-style outputs typically return entity strings but not reliable character offsets, whereas exact-match NER evaluation requires position-specific spans. When the same predicted surface form occurs multiple times in the same abstract, we map that mention to all exact string-matching occurrences in order to recover candidate offsets deterministically. This expansion step should therefore be understood as an offset-recovery procedure for span-text outputs rather than an additional modeling component, although it introduces an evaluation asymmetry relative to supervised token-level baselines that predict positions directly. For OpenAI/Gemma outputs, mentions that cannot be aligned are skipped and logged with warnings.

The GPT-5.4 prompt template used in our experiments is shown below.

```
System:
You are an information extraction model for
  biomedical case reports and reviews about
  prion disease.
Your job is to label entity spans using only the
```

```
provided schema and output one valid JSON
object.
```

```
User: Read the schema carefully and follow it
      exactly.
```

```
Schema JSON: {entity_schema}
```

```
Rules:
1. Extract only explicit spans from the text.
2. Use exact labels from the schema.
3. Use exact text spans from the input.
4. Do not output start or end offsets.
5. Do not output comments, markdown, or extra
   keys.
6. If uncertain, omit the span instead of
   guessing.
```

```
Input text: {text}
```

```
Output exactly one JSON object with this
structure:
```

```
{
  "text": "<original input text>",
  "entities": [
    {
      "mention": "<exact span>",
      "coarse_type": "<schema coarse type>",
      "fine_type": "<schema fine type>",
      "normalized": "<normalized form or same as
                    mention>"
    }
  ]
}
```

B Annotation Guidelines and Schema

B.1 Guideline Principles

Global rules: annotate only explicit mentions; prefer the most specific label available; keep spans minimal but complete; annotate both long form and abbreviation when both are explicit; treat molecular subtype strings such as *MM1*, *MV2*, and *VV2* as sCJD mentions, with finer subtype interpretation handled in normalization; treat exposure or treatment phrases such as *corneal transplantation*, *cadaveric-derived hormone treatment* (e.g., growth hormone), and *cadaveric dura mater graft* as Treatment mentions rather than disease mentions; include generic heads such as *disease*, *syndrome*, and *symptom* when they are part of the explicit annotated mention; and for statistical expressions annotate only the value span (e.g., *100%*, *0.06%*) rather than surrounding cue phrases.

Do not assign labels by keyword alone. A span is a disease mention only if it names a disease entity in context, and a span is a symptom mention only if it expresses a clinical manifestation in context. For example, *disease duration* is not a disease label, and *symptom onset* is not a Symptom label.

B.2 Train Annotation Consistency and Guideline Drift

The training annotations were produced over multiple rounds by the two annotators and then consolidated through adjudication into a single final corpus version. Because the train split was not preserved as two separately frozen annotator-specific versions, we do not report a formal per-annotator label-distribution comparison here. Instead, Table 7 summarizes the main recurrent sources of early annotation-style drift that were resolved during guideline refinement and then applied consistently in the final schema. Most of these changes affected boundary selection or distinctions between nearby label types rather than the overall entity inventory.

No.	First Guideline / Early Practice	Final Guideline
1	<i>brain</i> was not labeled as <code>Anatomic_location</code> .	<i>brain</i> is labeled as <code>Anatomic_location</code> .
2	<i>died</i> and <i>death</i> were not labeled as <code>Complication</code> .	<i>died</i> and <i>death</i> are labeled as <code>Complication</code> .
3	<i>FLAIR</i> , <i>T1-weighted</i> , <i>T2-weighted</i> , <i>ADC</i> , and <i>DWI</i> were labeled as <code>Imaging_test</code> .	These sequence terms are labeled as <code>Imaging_sequence</code> rather than <code>Imaging_test</code> .
4	Time expressions such as <i>in year 2013</i> , <i>in 2013</i> , and <i>12 months after onset of symptoms</i> were not consistently separated into anchor vs. duration components.	<code>Time_point</code> captures anchor expressions such as <i>2013</i> and <i>onset of symptoms</i> , while <code>Duration</code> captures spans such as <i>12 months</i> .
5	Autopsy was restricted to explicit procedure mentions such as <i>autopsy</i> and <i>autopsy examination</i> .	Autopsy was expanded to include additional procedure-reporting forms such as <i>pathologically</i> , <i>autopsy</i> , and <i>neuropathological examination</i> .
6	Statistical entities were often annotated as full cue phrases, e.g., <i>sensitivity of 61%</i> or <i>prevalence of 1/1 million people per year</i> .	Only the value span is annotated: e.g., <i>61%</i> as <code>Sensitivity</code> and <i>1/1 million people per year</i> as <code>Prevalence</code> .
7	<i>Heidenhain variant</i> was not labeled as <code>sCJD</code> .	<i>Heidenhain variant</i> is labeled as <code>sCJD</code> .

Table 7: Examples of guideline drift observed during training annotation and the final adjudicated guideline used in PrionNER.

B.3 Full Entity Schema

Table 8: Full coarse-grained and fine-grained entity schema of PrionNER, including definitions and representative examples.

Coarse-grained Type	Fine-grained Type	Definitions
Case Input (Geschwind, 2015)		
Age	Age	The age of a patient at disease onset. Examples: age 62; 45-year-old; at 58 years
Symptom	Symptom	Any subjective experience reported by the patient or objective observation made by a clinician. Examples: memory loss; fatigue; hyperreflexia; Babinski sign; rigidity; startle response
Test_name	Imaging_test	MRI; CT scan; PET scan
	Electrophysio_test	nerve conduction study, Polysomnography
	Blood_biomarker_test	CSF 14-3-3 assay; tau assay; complete blood count
	Genetic_test	PRNP genetic testing; mutation analysis; gene panel testing
	Molecular_assay	RT-QuIC; PCR; Western blot
	Autopsy	brain autopsy; neuropathological examination; postmortem analysis
Sequences	Imaging_sequence	DWI; FLAIR; diffusion-weighted imaging; ADC maps
Anatomic_location	Anatomic_location	A specific anatomical structure where a clinical finding or pathological change is observed. Examples: cortex; basal ganglia; thalamus; caudate nucleus; cerebral cortex
Findings	Imaging_finding	pulvinar sign; restricted diffusion; cortical ribboning
	Autopsy_finding	spongiform change; neuronal loss; astrogliosis
Case Diagnosis (Centers for Disease Control and Prevention, 2026)		

Continued on next page

Table 8: Full coarse-grained and fine-grained entity schema of PrionNER, including definitions and representative examples (continued).

Coarse-grained Type	Fine-grained Type	Definitions
Generic_Prion	Generic_Prion	Prion disease mentions without specific subtype classification. <i>Examples:</i> prion disease; CJD, bovine spongiform encephalopathy, BSE, mad cow disease, Transmissible Spongiform Encephalopathy, TSE
Sporadic_Prion	sCJD	Sporadic Creutzfeldt-Jakob Disease
	sFI	Sporadic Fatal Insomnia
	VPSPr	Variably Protease-Sensitive Prionopathy
Familial_Prion	fCJD	Familial Creutzfeldt-Jakob Disease
	GSS	Gerstmann-Sträussler-Scheinker Syndrome
	FFI	Fatal Familial Insomnia
Acquired_Prion	vCJD	Variant Creutzfeldt-Jakob Disease
	iCJD	Iatrogenic Creutzfeldt-Jakob Disease
	Kuru	Kuru
Differential_Diagnosis	Differential_Diagnosis	Non-prion diseases used for differential diagnosis (no fine-grained subtypes in this schema). <i>Examples:</i> Alzheimer’s disease; autoimmune encephalitis; viral encephalitis; Parkinson’s disease
Clinical Course and Context		
Treatment	Treatment	A therapeutic intervention, medication used to manage a disease or its symptoms. <i>Examples:</i> supportive care; quinacrine; doxycycline; symptomatic treatment (Benavente and Morales, 2024)
Complication	Complication	A secondary medical condition that arises as a consequence of disease progression. <i>Examples:</i> pneumonia; aspiration pneumonia; respiratory failure (Willis, 2006)
Time	Duration	A span or length of time over which a clinical event, symptom, or disease progression occurs. <i>Examples:</i> within 3 months; over 2 years; rapidly progressive over weeks
	Time_point	A specific point in time associated with clinical symptoms or a specific event. <i>Examples:</i> at onset; at age 62; in 2021
Stats	Sensitivity	The true positive rate for a diagnostic test (Saltelli et al., 2019). <i>Examples:</i> sensitivity of 85%
	Specificity	The true negative rate for a diagnostic test (Saltelli et al., 2019). <i>Examples:</i> specificity of 92%
	Prevalence	Disease frequency within a population (Tenny and Hoffman, 2017). <i>Examples:</i> 1 per million; prevalence of 0.5%
	Incidence	The rate at which new cases of a disease occur in a population during a specified time period. <i>Examples:</i> annual incidence of 2 per million

B.4 Compact Annotation Guideline

Table 9: Compact PrionNER annotation guideline.

Label	Use For	Annotate	Do Not Annotate
Generic_Prion	Prion disease mentions without explicit subtype classification	<i>prion disease, prion diseases, Creutzfeldt-Jakob disease, CJD, bovine spongiform encephalopathy, BSE, mad cow disease</i>	Specific subtype mentions such as <i>sporadic CJD</i> or <i>variant CJD</i> ; protein mentions such as <i>prion protein, PrP</i> , or <i>PrP^{Sc}</i> ; pathological finding uses such as <i>spongiform encephalopathy</i> when used as an autopsy finding
sCJD	Sporadic Creutzfeldt-Jakob disease	<i>sporadic Creutzfeldt-Jakob disease, sporadic CJD, sCJD, MM1, MV1, VV2</i> when clearly used as sCJD subtypes	Generic <i>CJD</i> without a sporadic cue
sFI	Sporadic fatal insomnia	<i>sporadic fatal insomnia, sFI</i>	<i>fatal familial insomnia, FFI, FFI-1, FFI-2</i>

Label	Use For	Annotate	Do Not Annotate
VPSPr	Variably protease-sensitive prionopathy	<i>variably protease-sensitive prionopathy, VPSPr</i>	<i>PrP, PrP^{Sc}, PrP-res</i> , or other protein-level mentions
fCJD	Familial or hereditary CJD	<i>familial CJD, familial Creutzfeldt-Jakob disease, hereditary CJD, genetic CJD, fCJD</i>	Mutation or genotype mentions alone
GSS	Gerstmann-Sträussler-Scheinker syndrome	<i>Gerstmann-Sträussler-Scheinker syndrome, GSS, GSS102, GSS105</i>	Mutation names alone; pathological descriptors such as <i>kuru plaques</i>
FFI	Fatal familial insomnia	<i>fatal familial insomnia, FFI, FFI-1, FFI-2</i>	Isolated <i>insomnia</i> ; use Symptom instead
vCJD	Variant CJD	<i>variant Creutzfeldt-Jakob disease, variant CJD, vCJD, new variant CJD, nvCJD</i>	Generic uses of <i>variant</i> without clear disease reference
iCJD	Iatrogenic CJD	<i>iatrogenic Creutzfeldt-Jakob disease, iatrogenic CJD, growth hormone-associated CJD, dural graft associated CJD, iCJD, dCJD</i>	<i>iatrogenic</i> alone when the CJD referent is unclear
Kuru	Kuru disease mention	<i>Kuru, kuru</i>	<i>kuru plaques</i> ; use Autopsy_finding instead
Differential_Diagnosis	Non-prion diseases considered as alternatives or comparators	<i>Alzheimer disease, autoimmune encephalitis, Huntington disease</i>	Prion diseases; symptom phrases unless they are clearly used as alternative diagnoses
Symptom	Symptoms, signs, and clinical manifestations	<i>dementia, myoclonus, cerebellar ataxia, psychotic symptoms, hyperreflexia</i>	Disease names; imaging findings; autopsy findings
Imaging_test	Imaging modalities or procedures	<i>MRI, CT, PET, SPECT</i>	Sequences such as <i>FLAIR, DWI</i> ; findings such as <i>pulvinar sign</i>
Electrophysio_test	Electrophysiology procedures	<i>EEG, electroencephalogram, polysomnography</i>	Electrophysiologic findings or waveform targets
Blood_biomarker_test	Biomarker assays and specimen-as-test shorthand	<i>CSF, tau assay, blood tests</i>	Molecular methods such as <i>PCR, RT-QuIC</i> , or <i>Western blot</i>
Genetic_test	Genetic testing procedures	<i>genetic testing, Prion gene analysis, genetic analysis</i>	Gene names and mutation names alone
Molecular_assay	Molecular or biochemical assays	<i>RT-QuIC, PCR, Western blot, molecular analysis</i>	Biopsy or autopsy procedures; gene targets without a testing method
Autopsy	Postmortem or tissue examination procedures	<i>brain autopsy, postmortem examination, brain biopsy</i>	Pathological findings themselves; anatomy alone
Imaging_sequence	Acquisition or sequence terms	<i>DWI, FLAIR, T2 weighted, ADC maps</i>	Imaging modality names; imaging abnormalities
Imaging_finding	Radiologic abnormalities or named signs	<i>pulvinar sign, restricted diffusion, cortical ribboning</i>	Test modality alone; anatomy alone unless separately annotated
Autopsy_finding	Pathological or postmortem findings	<i>spongiform change, neuronal loss, gliosis, florid plaques, kuru plaques, spongiform encephalopathy</i> when used as a pathological finding	Procedure terms; disease names
Anatomic_location	Body locations linked to findings or symptoms	<i>thalamus, basal ganglia, caudate nucleus, cortex, deep white matter</i>	Non-anatomical descriptors; whole finding phrase when only one part is anatomy
Age	Age or age-at-onset expression	<i>58-year-old, age 62, young adults</i>	Durations; calendar dates

Label	Use For	Annotate	Do Not Annotate
Treatment	Therapies, medications, or care strategies	<i>palliative care, quinacrine, doxycycline, antipsychotics</i> ; annotate the full treatment mention when expressed as a therapy or care strategy	Diagnostic tests; non-intervention goals unless expressed as actual therapy
Complication	Secondary adverse conditions or end-stage outcomes	<i>pneumonia, respiratory failure, SIADH, death, died</i>	The primary prion disease itself
Duration	Time span or length	<i>within 3 months, over 2 years, 13 months</i>	Calendar dates; age expressions
Time_point	Specific time anchor	<i>at onset, on admission, 1996, January 2012</i>	Durations; exclude discourse connectives such as <i>in, after</i> when they are not themselves the time anchor
Sensitivity	Diagnostic sensitivity expression	<i>100%, 91%</i> when they are explicitly the sensitivity value, e.g., annotate only <i>100%</i> in <i>sensitivity of 100%</i>	The cue word alone, e.g., <i>sensitivity</i> , when no value is included; unrelated percentages
Specificity	Diagnostic specificity expression	<i>92%, 95%</i> when they are explicitly the specificity value	The cue word alone, e.g., <i>specificity</i> , when no value is included; unrelated percentages
Prevalence	Disease frequency expression	<i>1–2 people per million annually, 0.5%</i> when it is explicitly the prevalence value	The cue word alone, e.g., <i>prevalence</i> , when no value is included; sample size counts
Incidence	New-case rate expression	<i>0.06%, 2 per million annually, 0.37 cases/million, 1 in 1 000 000</i> when they are explicitly the incidence value	The cue word alone, e.g., <i>incidence</i> , when no value is included; sample size counts; surrounding time or cue phrases such as <i>in 2023, million per year, or annual incidence of</i>

C Extended Dataset Statistics

C.1 Structural Annotation Statistics

The train and test splits both contain non-trivial structural complexity at the span level (Table 10). Discontinuous entities remain relatively uncommon in absolute terms, but they appear in both splits and are distributed across clinically salient labels rather than being confined to a single category. In both train and test, *Symptom* is the most common discontinuous label, followed by *Anatomic_location*; train also shows notable discontinuous cases for *Imaging_finding*, while test includes comparatively more *Blood_biomarker_test* discontinuities (Table 11).

Nested structures are much more frequent than overlapping ones, and they are concentrated in recurring clinically meaningful label pairs (Tables 12 and 13). In the training split, the most common nested envelope patterns are *Symptom* within *Symptom*, *Duration* with *Time_point*, and *Anatomic_location* paired with *Symptom*; in the test split, the dominant patterns shift toward *Anatomic_location* with *Symptom* or *Imaging_finding*, with an additional cluster of *Treatment*–*iCJD* nesting. By contrast, overlapping envelope pairs are rare in both splits, suggesting that the benchmark’s structural difficulty is driven primarily by discontinuity and especially nesting rather than widespread partial-overlap phenomena.

Split	Docs	Entities	Disc.	Nested Env.	Overlap Env.
Train	247	5,149	97	235	5
Test	70	1,806	34	80	2

Table 10: Span-structure summary for the train and test splits. Discontinuous entities are text-bound entities with more than one atomic span. Nested and overlapping pair counts are reported at the entity-envelope level.

Label	Train	Test
Symptom	36	10
Anatomic_location	15	9
Imaging_finding	14	3
Blood_biomarker_test	4	5
Differential_Diagnosis	4	2
Treatment	4	1
Duration	3	2
Imaging_sequence	4	0
Imaging_test	4	0
Autopsy_finding	3	1
sCJD	3	0
vCJD	1	1
fCJD	1	0
iCJD	1	0

Table 11: Discontinuous entities by label in the train and test splits. The train split contains discontinuous entities in 54 documents, and the test split contains discontinuous entities in 20 documents.

Label A	Label B	Train	Test
Symptom	Symptom	42	9
Anatomic_location	Symptom	28	13
Duration	Time_point	30	2
Anatomic_location	Anatomic_location	18	10
Anatomic_location	Imaging_finding	14	11
Anatomic_location	Autopsy_finding	16	4
Imaging_finding	Imaging_sequence	9	2
Anatomic_location	Differential_Diagnosis	8	0
Blood_biomarker_test	Blood_biomarker_test	4	4
Treatment	iCJD	1	6
Imaging_sequence	Imaging_sequence	6	0
Differential_Diagnosis	Differential_Diagnosis	4	2
Generic_Prion	Time_point	4	2
Duration	Duration	2	3
Imaging_finding	Imaging_finding	4	0
Autopsy_finding	Kuru	3	1
Anatomic_location	Treatment	3	0
Imaging_sequence	Imaging_test	3	0
Differential_Diagnosis	Symptom	2	1
Blood_biomarker_test	Differential_Diagnosis	0	3

Table 12: Nested entity-envelope pairs by label pair in the train and test splits.

Label A	Label B	Train	Test
Anatomic_location	Anatomic_location	1	0
Autopsy_finding	Generic_Prion	1	0
FFI	sFI	1	0
Imaging_test	Imaging_test	1	0
Symptom	Symptom	1	0
Blood_biomarker_test	Differential_Diagnosis	0	1
Time_point	Treatment	0	1

Table 13: Overlapping entity-envelope pairs by label pair in the train and test splits.

Split	Abstracts	Entity Types	Mentions	Unique Surface Forms
Train	247	28	4,655	1,889
Test	70	29	1,650	729

Table 14: Split-level summary of schema-defined non-meta entity annotations and unique normalized surface forms.

Entity Type	Train Unique	Test Unique	Train Mentions	Test Mentions
Symptom	424	148	791	248
Anatomic_location	212	89	646	223
Duration	163	41	217	49
Time_point	127	39	152	47
Treatment	102	59	126	88
Autopsy_finding	110	32	215	42
Differential_Diagnosis	113	27	188	36
Age	108	25	137	29
Imaging_finding	92	38	153	71
Imaging_test	72	22	261	75
Generic_Prion	61	27	741	233
Autopsy	54	25	199	64
Imaging_sequence	38	27	117	77
Electrophysio_test	42	10	173	38
Molecular_assay	20	29	23	45
Blood_biomarker_test	23	14	64	39
sCJD	18	16	86	80
iCJD	18	10	32	16
Complication	18	6	78	19
GSS	14	8	32	16

Table 15: Train/test comparison of unique normalized surface forms and non-meta mention counts for the most frequent entity types.

C.2 Split-Level Entity Summary

C.3 Entity Surface-Form Dictionary Summary

Unique surface forms are counted per entity type after lowercasing, trimming edge whitespace, collapsing internal whitespace, and joining discontinuous spans with spaces.

D Extended Results

D.1 Overall Performance

Figure 2 complements the main fine-grained results by showing the precision–recall trade-offs across the evaluated models. The supervised BERT baselines cluster in the strongest region overall, while the zero-shot models show a wider spread and sharper trade-offs; among them, Gemma-4-31B occupies the strongest overall position in the fine-grained setting.

D.2 Per-Type Fine-grained Results

Figure 3 and Table 16 show that fine-grained performance is highly uneven across the schema. The heatmap makes clear that there is no single difficulty level for PrionNER: some labels are solved reliably by most strong systems, while others re-




Figure 2: Precision–recall trade-offs for fine-grained entity extraction across the evaluated models.

main unstable even for the best models. This heterogeneity tracks both frequency and semantic specificity. Common and lexically anchored categories such as *Generic_Prion*, *sCJD*, *Imaging_test*, *Symptom*, *Age*, and *Anatomic_location* are comparatively robust, whereas sparse or context-dependent labels such as *Prevalence*, *Specificity*, *Sensitivity*, *Time_point*, *Molecular_assay*, and subtype distinctions such as *iCJD* and *fCJD* remain much harder.

The supervised encoders occupy the strongest region overall and are more consistent across label families. PubMedBERT in particular is among the top performers on many of the clinically common labels, reaching 91.77 on *Generic_Prion*, 90.79 on *Imaging_test*, 89.74 on *sCJD*, 88.89 on *Age*, 84.95 on *Symptom*, and 81.48 on *Blood_biomarker_test*. BioBERT and ClinicalBERT show a similar profile and outperform PubMedBERT on a few categories, such as *Anatomic_location*, *vCJD*, *Complication*, *Autopsy*, *Autopsy_finding*, *Incidence*, and *iCJD*. Taken together, the supervised models suggest that once training data are available, the main gains come from broad coverage across the entire schema rather than isolated wins on a handful of labels.

The zero-shot models show a much sharper specialization pattern. Gemma-4-31B is the strongest zero-shot model overall, but its strengths are concentrated in a narrower subset of distinctive labels, including *FFI* (100.00), *fCJD* (77.78), *Genetic_test* (71.43), *Molecular_assay* (65.79), *Duration* (65.98), and *Differential_Diagnosis* (61.29). GPT-5.4 is competitive on several labels and is best on *Electrophysio_test* (89.19) and *Time_point* (40.78), while GLiNER2 occasionally produces

isolated best scores on very sparse categories such as *Sensitivity*, *Specificity*, and *Duration*. However, these zero-shot wins are scattered and do not translate into the same schema-wide stability seen in the supervised baselines.

Two further patterns are worth noting. First, the imaging and diagnostic-context families remain difficult even when overall F1 is moderate: *Imaging_finding*, *Imaging_sequence*, *Blood_biomarker_test*, *Autopsy_finding*, and *Time_point* all show larger cross-model variation than canonical disease-name labels. Second, the most extreme values should be interpreted with caution for very rare labels. For example, *Prevalence* remains at 0.00 for all models, and isolated perfect or near-perfect results on labels such as *Sensitivity* or *FFI* reflect very small test counts rather than uniformly solved clinical reasoning. Overall, the per-type view reinforces the main conclusion of the paper: PrionNER rewards models that combine strong lexical grounding on common biomedical entities with fine-grained contextual discrimination on rarer, semantically adjacent categories.

D.3 Entity-only Confusion Analysis

Figures 4a and 4b provide a complementary entity-only view of fine-grained errors for the two strongest models. Under this evaluation, PubMedBERT is clearly stronger overall, reaching 0.794 flat entity-level F1 compared with 0.684 for Gemma-4-31B. The main difference is recall: PubMedBERT reaches 0.793 recall versus 0.597 for Gemma-4-31B, while their precision remains very similar at 0.796 and 0.801, respectively. This pattern shows that Gemma-4-31B is comparatively conservative, preserving Gemma-level precision but missing many more entities.

In these row-normalized, entity-label-only matrices, PubMedBERT has the cleaner diagonal overall. More labels retain darker diagonal mass with less off-diagonal spill, indicating more consistent label assignment across the schema. Gemma-4-31B looks reasonable on a few canonical labels, but its diagonal is visibly weaker for many harder clinical categories, including *Autopsy*, *Autopsy_finding*, *Blood_biomarker_test*, *Time_point*, and *Treatment*.

Several confusion patterns are shared by both models. *Prevalence* and *Incidence* remain difficult to separate, the imaging cluster shows substantial overlap among *Imaging_finding*, *Imaging_sequence*, and *Imaging_test*, and the prion sub-

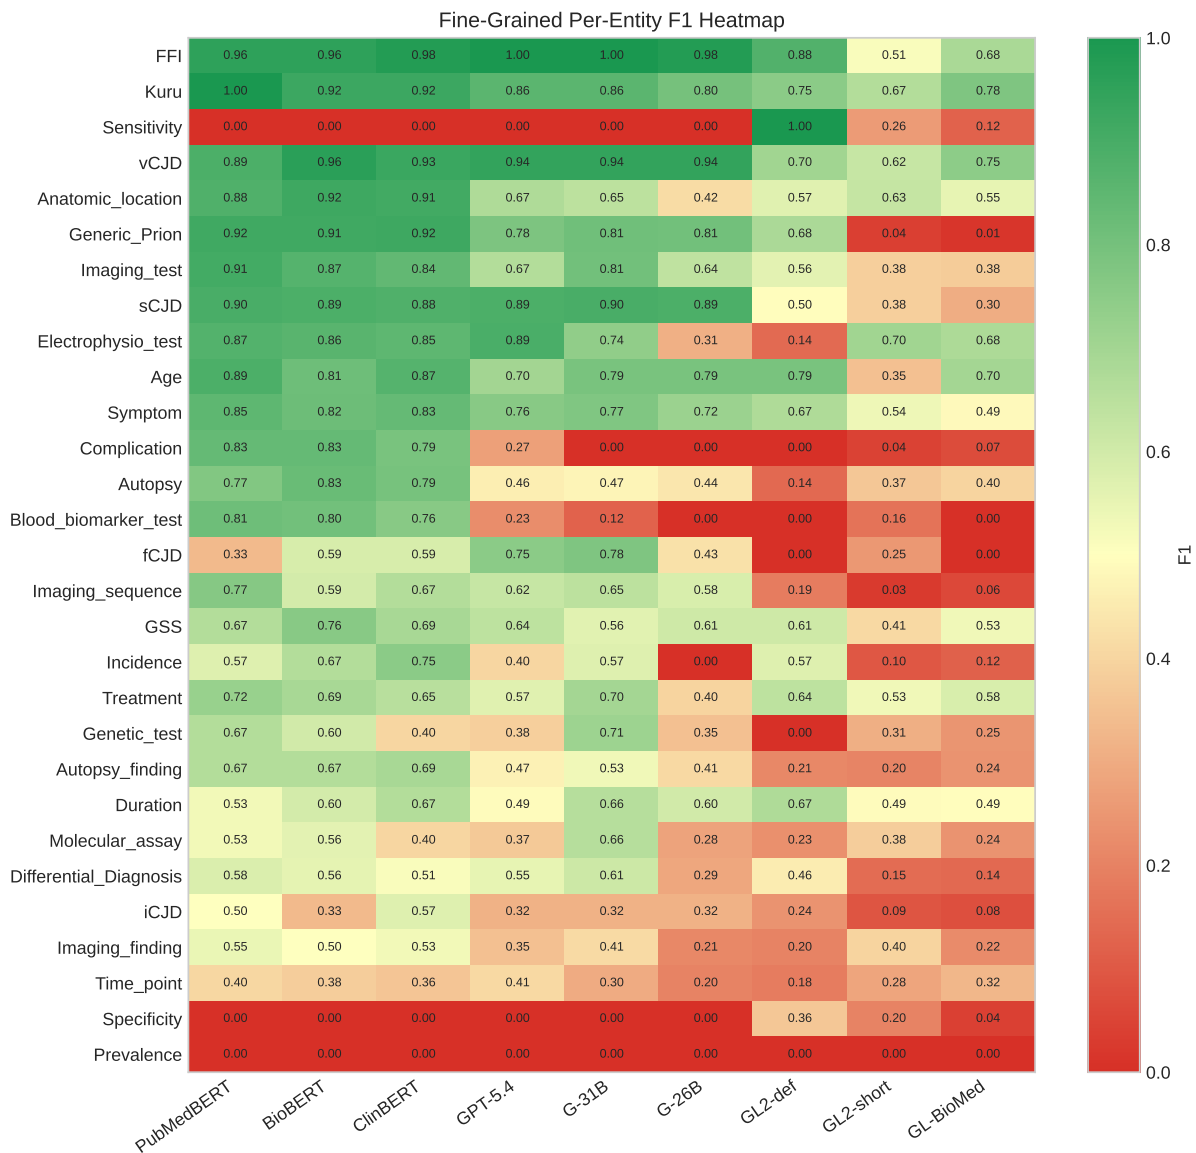


Figure 3: Per-type fine-grained F1 heatmap for the main supervised and zero-shot models on the test set.

type region also remains imperfectly separated, especially among *fCJD*, *iCJD*, *GSS*, *sFI*, and nearby subtype labels. These are largely fine-grained semantic confusions between clinically adjacent categories rather than arbitrary label noise. One recurring shared error is the prediction of *Generic_Prion* for *fCJD* and *iCJD*. For *fCJD*, a plausible explanation is that abstracts often realize the subtype as phrases such as “familial CJD” or “familial form of CJD,” where the diagnostically important modifier appears before the base disease name. In these cases, the model appears to anchor on “CJD” while failing to consistently preserve the prefix that signals the familial subtype. The same pattern likely explains part of the *iCJD* → *Generic_Prion* confusion: the model often cap-

tures “CJD” but misses the preceding cue that indicates an iatrogenic or transmitted form. This error is amplified by longer compositional mentions such as “corneal transplant-related CJD,” where the treatment-related phrase can be separated from the disease name; in such cases, the model may label “corneal transplant” as *Treatment* and the remaining “CJD” span as *Generic_Prion* rather than assigning the full mention to *iCJD*. Another plausible shared confusion is between *Complication* and *Symptom*. Some complication mentions are lexically symptom-like in isolation, but are labeled as *Complication* because they denote downstream consequences of disease progression rather than primary manifestations. For example, *dysphagia* can be a symptom in many settings, but in prion

Entity Type	Supervised			Zero-shot					
	PubMedBERT	BioBERT	ClinBERT	GPT-5.4	G-31B	G-26B	GL2-def	GL2-short	GL-BioMed
FFI	95.65	95.65	97.78	100.00	100.00	97.78	87.80	51.43	68.42
Kuru	100.00	92.31	92.31	85.71	85.71	80.00	75.00	66.67	77.78
Sensitivity	0.00	0.00	0.00	0.00	0.00	0.00	100.00	25.81	12.24
vCJD	88.70	96.23	92.73	94.34	94.34	94.34	70.37	62.26	74.77
Anatomic_location	88.28	91.88	91.18	67.19	64.79	41.61	56.81	62.83	55.35
Generic_Prion	91.77	91.45	91.81	78.38	80.95	80.69	68.33	3.86	1.35
Imaging_test	90.79	86.96	84.34	66.67	80.50	63.83	55.93	38.37	37.70
sCJD	89.74	89.03	87.90	88.89	89.66	88.89	49.54	38.37	30.19
Electrophysio_test	87.18	86.08	85.00	89.19	73.85	31.11	14.29	70.42	67.74
Age	88.89	81.48	86.79	70.18	79.25	79.25	79.25	34.92	69.84
Symptom	84.95	81.97	83.21	76.16	77.41	71.86	67.28	53.69	48.72
Complication	83.33	83.33	78.95	27.27	0.00	0.00	0.00	4.40	6.90
Autopsy	77.31	82.76	79.31	45.83	47.42	43.90	13.89	36.54	39.66
Blood_biomarker_test	81.48	80.49	75.86	22.64	12.20	0.00	0.00	16.33	0.00
fCJD	33.33	58.82	58.82	75.00	77.78	42.86	0.00	25.00	0.00
Imaging_sequence	76.51	59.42	66.67	62.30	64.62	58.33	18.60	2.50	6.20
GSS	66.67	75.86	68.97	64.29	56.25	60.61	60.87	40.91	53.33
Incidence	57.14	66.67	75.00	40.00	57.14	0.00	57.14	9.52	11.76
Treatment	72.39	69.01	65.43	56.94	69.57	39.64	64.43	53.27	58.45
Genetic_test	66.67	60.00	40.00	38.46	71.43	34.78	0.00	30.77	24.56
Autopsy_finding	66.67	66.67	68.82	46.67	53.12	40.68	21.28	20.34	24.18
Duration	53.06	59.57	66.67	49.02	65.98	59.77	67.35	49.02	49.48
Molecular_assay	52.78	56.00	40.00	36.92	65.79	28.07	23.08	38.20	23.91
Differential_Diagnosis	57.97	55.56	51.43	55.17	61.29	28.57	45.61	14.81	13.79
iCJD	50.00	33.33	57.14	31.58	31.58	31.58	24.00	9.23	7.69
Imaging_finding	54.55	50.38	53.03	34.92	41.18	21.43	20.25	39.71	21.90
Time_point	40.43	38.10	36.00	40.78	29.85	20.00	18.46	28.32	32.43
Specificity	0.00	0.00	0.00	0.00	0.00	0.00	36.36	20.00	3.92
Prevalence	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 16: Per-type fine-grained F1 scores for selected supervised and zero-shot models on the test set, sorted by the best score achieved for each entity type. Header abbreviations: ClinBERT = ClinicalBERT; GL2-def = GLiNER2-def; GL2-short = GLiNER2-short; GL-BioMed = GLiNER-BioMed; G-26B and G-31B = Gemma 4 26B and 31B.

disease abstracts it may appear as a late-stage secondary consequence of earlier neurological decline, which makes the boundary between the two labels difficult to recover from local context alone. Similarly, mentions such as “died” or “death” may be pulled toward *Symptom* or other outcome-like categories even though they are not direct symptomatic expressions of prion disease itself. Instead, they typically reflect the final downstream result of severe functional deterioration and accumulated health complications over the course of the disease.

For PubMedBERT, the strongest diagonal cells are concentrated on common, high-volume labels such as *Symptom*, *Anatomic_location*, *Imaging_test*, *Generic_Prion*, and *sCJD*. Its remaining errors mostly look like local mix-ups within semantically related label families rather than broad collapse across the label space. Gemma-4-31B, by contrast, is weaker on context-heavy diagnostic labels and on imaging-related distinctions; its diag-

onal is often present but noticeably lighter, suggesting less stable fine-grained label assignment even when an entity span is recovered.

This entity-only view should also be read together with the full confusion analysis that includes *O*. Because these matrices exclude non-entity predictions, they likely understate Gemma-4-31B’s main weakness, namely omission of the entity altogether. In other words, PubMedBERT appears stronger and more balanced, with errors concentrated within semantically adjacent labels, whereas Gemma-4-31B shows reasonable zero-shot behavior but is less reliable on subtle distinctions and is likely more omission-heavy in the full setting.


At the label level, PubMedBERT remains strong on common clinical categories, including *Generic_Prion* (0.918), *Imaging_test* (0.908), *sCJD* (0.897), *vCJD* (0.887), *Anatomic_location* (0.883), and *Symptom* (0.850). Gemma-4-31B is still strong on a narrower set of distinctive labels,

including *Generic_Prion* (0.917), *FFI* (0.977), *Kuru* (0.923), *Age* (0.792), *Imaging_test* (0.781), and *Symptom* (0.773), but its performance drops sharply on rarer or more context-dependent categories. Its weakest flat F1 values include *Complication* (0.000), *Prevalence* (0.000), *Sensitivity* (0.000), *Specificity* (0.000), *Blood_biomarker_test* (0.172), *iCJD* (0.286), *Time_point* (0.317), and *Imaging_finding* (0.321). PubMedBERT also shows weak spots, but they are milder and are concentrated in rare epidemiology labels and subtype distinctions: *Prevalence* (0.000), *Sensitivity* (0.000), *Specificity* (0.000), *fCJD* (0.333), *Time_point* (0.404), *iCJD* (0.500), and *Molecular_assay* (0.528). Overall, these figures reinforce the main result: PubMedBERT is the better model for PrionNER because it maintains Gemma-4-31B-level precision while recovering many more entities, whereas Gemma-4-31B remains the strongest zero-shot model but is still limited primarily by recall.

D.4 Per-label Annotation Agreement

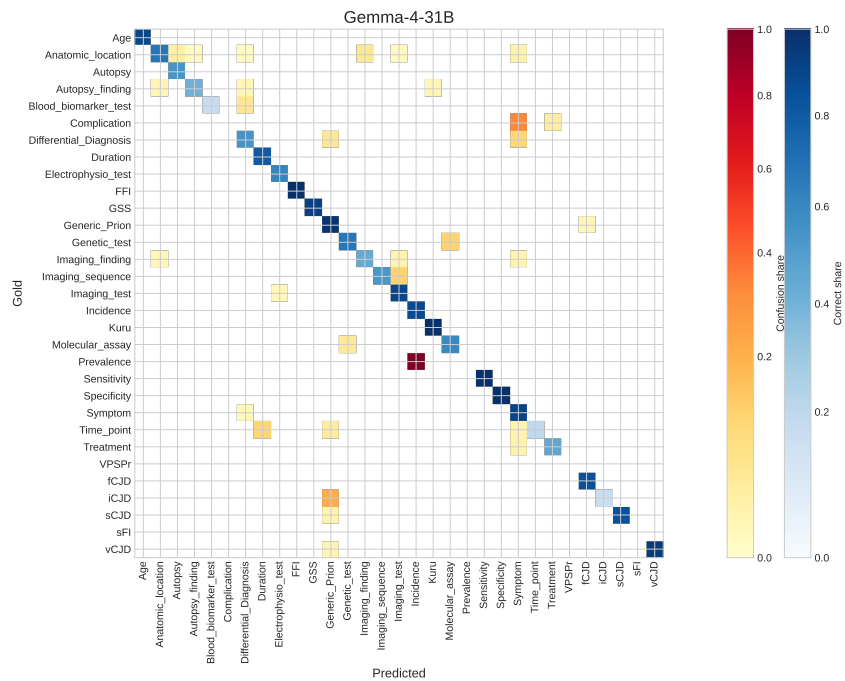
Table 17 provides the full per-label agreement breakdown for the test split. Agreement is highest for common and semantically distinctive disease labels such as *vCJD*, *GSS*, *FFI*, and *sCJD*, while lower agreement is concentrated in sparse or boundary-sensitive categories such as *Complication*, *Prevalence*, *Incidence*, and *Imaging_finding*.

Fine-Grained Token Confusion Matrix: PubMedBERT (entity labels only, row-normalized, blue diagonal + warm confusions)



(a) PubMedBERT

Fine-Grained Token Confusion Matrix: Gemma-4-31B (entity labels only, row-normalized, blue diagonal + warm confusions)



(b) Gemma-4-31B

Figure 4: Entity-only fine-grained confusion matrices for the strongest supervised model (PubMedBERT) and the strongest zero-shot model (Gemma-4-31B).

Label	F1	Matches	Ann. 1	Ann. 2	Union	Jaccard
Complication	0.00	0	12	2	14	0.00
Prevalence	0.00	0	2	4	6	0.00
Incidence	40.00	1	2	3	4	25.00
Imaging_finding	47.19	21	49	40	68	30.88
Autopsy_finding	51.61	16	37	25	46	34.78
Blood_biomarker_test	52.63	20	32	44	56	35.71
Time_point	60.61	20	34	32	46	43.48
Treatment	63.01	46	58	88	100	46.00
Differential_Diagnosis	63.64	28	36	52	60	46.67
Autopsy	76.54	31	43	38	50	62.00
Age	76.92	20	26	26	32	62.50
Duration	77.65	33	40	45	52	63.46
iCJD	78.26	9	11	12	14	64.29
Imaging_sequence	80.00	50	66	59	75	66.67
Genetic_test	80.00	4	4	6	6	66.67
Symptom	81.00	194	235	244	285	68.07
Specificity	83.33	5	5	7	7	71.43
Electrophysio_test	84.51	30	34	37	41	73.17
Sensitivity	85.71	6	6	8	8	75.00
Molecular_assay	87.64	39	47	42	50	78.00
Anatomic_location	90.30	149	159	171	181	82.32
Imaging_test	92.31	60	64	66	70	85.71
Kuru	93.33	7	7	8	8	87.50
Generic_Prion	94.83	220	234	230	244	90.16
fCJD	95.24	10	11	10	11	90.91
sCJD	95.54	75	79	78	82	91.46
FFI	95.65	22	22	24	24	91.67
vCJD	96.55	56	58	58	60	93.33
GSS	96.77	15	16	15	16	93.75

Table 17: Per-label annotation agreement under exact label-and-span matching on the 70-abstract test split, with one abstract excluded from the final agreement comparison. Ann. 1 and Ann. 2 denote the number of entities annotated by Annotator 1 and Annotator 2, respectively, Matches denotes exact label-and-span matches, and Union is the size of the union of annotated entities for that label.

Coarse-grained Type	Fine-grained Type	Train Count	Train %	Test Count	Test %
Generic_Prion	Generic_Prion	741	15.92	233	14.12
Sporadic_Prion	sCJD	86	1.85	80	4.85
Sporadic_Prion	sFI	3	0.06	0	0.00
Sporadic_Prion	VPSPr	0	0.00	0	0.00
Familial_Prion	fCJD	26	0.56	10	0.61
Familial_Prion	GSS	32	0.69	16	0.97
Familial_Prion	FFI	39	0.84	22	1.33
Acquired_Prion	vCJD	102	2.19	55	3.33
Acquired_Prion	iCJD	32	0.69	16	0.97
Acquired_Prion	Kuru	28	0.60	7	0.42
Differential_Diagnosis	Differential_Diagnosis	188	4.04	36	2.18
Symptom	Symptom	791	16.99	248	15.03
Test_name	Imaging_test	261	5.61	75	4.55
Test_name	Electrophysio_test	173	3.72	38	2.30
Test_name	Blood_biomarker_test	64	1.37	39	2.36
Test_name	Genetic_test	16	0.34	8	0.48
Test_name	Molecular_assay	23	0.49	45	2.73
Test_name	Autopsy	199	4.27	64	3.88
Sequences	Imaging_sequence	117	2.51	77	4.67
Findings	Imaging_finding	153	3.29	71	4.30
Findings	Autopsy_finding	215	4.62	42	2.55
Anatomic_location	Anatomic_location	646	13.88	223	13.52
Age	Age	137	2.94	29	1.76
Treatment	Treatment	126	2.71	88	5.33
Complication	Complication	78	1.68	19	1.15
Time	Duration	217	4.66	49	2.97
Time	Time_point	152	3.27	47	2.85
Stats	Sensitivity	0	0.00	4	0.24
Stats	Specificity	0	0.00	4	0.24
Stats	Prevalence	4	0.09	1	0.06
Stats	Incidence	6	0.13	4	0.24
Total		4655	100.00	1650	100.00

Table 18: Full fine-grained entity distribution in the PrionNER train and test splits. Percentages are computed separately within the train and test sets over schema-defined fine-grained entity mentions only.

E Additional Reference Tables

E.1 Full Fine-grained Entity Distribution

E.2 Top Surface Forms by Entity Type

Table 19 reports the 10 most frequent normalized surface forms for each entity type. Numbers in parentheses after each surface form indicate mention counts, and the left-column notation “U./M.” denotes total unique surface forms and total mentions for that entity type.

Table 19: Top 10 normalized surface forms for each entity type in the combined dataset.

Entity Type (U./M.)	Top Surface Forms (mentions)
Symptom (524 / 1279)	dementia (79); myoclonus (79); ataxia (47); rapidly progressive dementia (36); psychiatric symptoms (24); akinetic mutism (16); progressive dementia (16); myoclonic jerks (15); aggression (15); neurological symptoms (14)
Anatomic_location	

Continued on next page

Entity Type (U./M.)	Top Surface Forms (mentions)
(276 / 837)	brain (48); basal ganglia (41); white matter (29); cortical (26); cerebellum (23); cerebellar (23); thalamus (22); striatum (21); cerebral cortex (19); pulvinar (18)
Duration (201 / 290)	7 months (8); within a year (7); 12 months (7); four months (6); 13 months (6); within one year (5); 4 months (5); 3 months (5); 5 months (5); disease duration (4)
Time_point (181 / 248)	age at onset (5); at onset (5); 1996 (5); after onset (5); in 1996 (5); 1985 (4); at the onset (4); the end of the third stage (4); clinical onset (3); early in the course (3)
Imaging_finding (161 / 237)	pulvinar sign (17); brain atrophy (13); hyperintensity (11); high signal intensity (7); atrophy (6); hyperintense signal abnormalities (4); signal intensity abnormalities (4); high signal (4); hyperintensities (3); cortical atrophy (3)
Autopsy_finding (155 / 276)	neuronal loss (22); spongiform change (20); spongiosis (9); spongiform changes (8); gliosis (8); status spongiosus (8); astrocytosis (7); neuronal degeneration (5); spongiform encephalopathy (5); kuru plaques (5)
Differential_Diagnosis (136 / 278)	alzheimer's disease (17); dlb (16); dementia (14); ad (12); stroke (9); neurodegenerative disorder (7); non-prion disorders (6); insomnia (5); dementias (4); thalamic dementia (4)
Treatment (131 / 241)	palliative care (23); quinacrine (16); antipsychotics (9); quetiapine (6); haloperidol (6); treatment (4); anesthesia (4); supportive care (4); hospice care (4); dura mater graft (4)
Age (124 / 182)	61-year-old (5); elderly (5); 59-year-old (5); 49-year-old (5); 54-year-old (4); 65-year old (3); 48-year-old (3); 70-year-old (3); 75-year-old (3); 58-year-old (3)
Imaging_test (100 / 383)	mri (82); magnetic resonance imaging (37); ct (28); spect (27); mr imaging (17); mr (13); mr images (13); ct scan (9); positron emission tomography (7); brain mri (7)
Generic_Prion (85 / 1219)	cjd (576); creutzfeldt-jakob disease (327); prion disease (50); bse (33); prion diseases (30); bovine spongiform encephalopathy (24); spongiform encephalopathy (16); creutzfeldt-jakob disease (14); creutzfeldt-jacob disease (9); subacute spongiform encephalopathy (6)
Autopsy (74 / 271)	autopsy (62); brain biopsy (36); necropsy (20); biopsy (13); neuropathological (11); neuropathological examination (9); neuropathologically (7); pathological (5); pathologically (5); histologically (5)
Imaging_sequence (63 / 241)	dwi (34); flair (34); diffusion-weighted (15); diffusion-weighted imaging (12); t2-weighted (12); t2 (11); diffusion-weighted images (9); dw (8); adcs (8); dw images (8)
Electrophysio_test (49 / 247)	eeg (119); electroencephalogram (29); electroencephalographic (15); electroencephalography (10); eegs (8); electroencephalograms (6); erg (4); polysomnography (4); br (3); electroencephalographic findings (3)
Blood_biomarker_test (44 / 140)	csf (57); cerebrospinal fluid (14); blood tests (9); csf analysis (6); csf tau protein (4); csf virology (3); biomarkers (3); csf tau-pt181 (2); csf studies (2); serum workup (2)
Molecular_assay (40 / 111)	cdi (16); csf rt-quick (6); immunohistochemistry (6); ihc (6); western blot (5); immunocytochemistry (5); molecular analysis (4); densitometric analysis (4); protein assay (3); csf rt-qulc (3)
sCJD (26 / 241)	scjd (90); sporadic cjd (52); sporadic creutzfeldt-jakob disease (29); sporadic (17); vv2 (8); heidenhain variant (4); creutzfeldt-jakob disease (4); sporadic form (4); mv1 (4); mm1 (4)
Complication (22 / 52)	death (20); died (7); myocarditis (3); deaths (2); bronchopneumonia (2); fatal (2); dysphasia (1); loss of independence (1); dead (1); acute myocarditis (1)
iCJD (19 / 51)	iatrogenic cjd (8); iatrogenic creutzfeldt-jakob disease (8); dcjd (8); icjd (5); iatrogenic (4); iatrogenic forms (3); iatrogenic cases (2); dural graft associated cjd (2); iatrogenic transmission (1); iatrogenic transmission of cjd (1)
GSS	

Continued on next page

Entity Type (U./M.)	Top Surface Forms (mentions)
(18 / 63)	gss (19); gerstmann-sträussler-scheinker disease (14); gss102 (6); gerstmann-straussler-scheinker disease (4); gss105 (4); gssd (3); gerstmann-straussler-scheinker syndrome (2); gerstmann-sträussler-scheinker's disease (1); gerstmann-sträussler-scheinker's syndrome (1); gerstmann-strässler-scheinker's syndrome (1)
vCJD (18 / 218)	vcjd (104); nvcjd (37); variant creutzfeldt-jakob disease (26); variant cjd (26); nv-cjd (4); acquired (4); variant creutzfeldt-jakob (2); infectious (2); new variant cjd (2); variant (2)
fCJD (16 / 47)	familial (9); familial cjd (8); inherited prion disease (6); inherited (4); fcjd (3); genetic cjd (3); genetic (3); hereditary cjd (2); cjd178 (2); genetic forms (1)
Genetic_test (15 / 26)	molecular genetic analysis (5); genetic testing (3); genetic analysis (2); prp gene analysis (2); restriction-enzyme analysis (2); analysing dna (2); prion gene analysis (2); genetic examination (1); genotyping (1); genetic tests (1)
Incidence (11 / 12)	high incidence (2); 0.06% in 2023 (1); 0.10% in 2024 (1); 1 to 2 cases per million people per year (1); annual incidence of 0.37 cases/million (1); incidence (1); annual incidence of 0.5-1.5 cases of cjd per million (1); incidence of approximately .5-1 new cases per million population per year (1); incidence 1 in 1 000 000 (1); 1 to 2 cases per million people per year. (1)
Prevalence (9 / 10)	1–2 people per million annually (2); 1–2/million/year (1); one case per million people per year (1); one in one million (1); 1 to 2 cases per million people per year (1); 35% (1); 1–2 people per million annually. (1); 0.06% (1); 0.10% (1)
FFI (7 / 85)	ffi (44); fatal familial insomnia (33); ffi-1 (4); ffi-2 (1); fatal familial insomnia (1); met-met subtype (1); fatal insomnia (1)
Sensitivity (7 / 15)	sensitivity (4); sensitivity of 100% (2); sensitivity of 87% (2); sensitivity (96%) (2); 91% sensitive (2); sensitivity higher (2); sensitive (1)
Specificity (5 / 12)	specificity (4); specificity of 92% (2); specificity of 97% (2); specificity (97%) (2); 95% specific (2)
Kuru (3 / 43)	kuru (40); kuru plaques (2); kuru type (1)
sFI (3 / 3)	ffi-1 (1); ffi-2 (1); fatal familial insomnia (1)