

Binary Token-Level Classification with DeBERTa for All-Type MWE Identification: A Lightweight Approach with Linguistic Enhancement

Diego Rossini and Lonneke van der Plas

Università della Svizzera Italiana

diego.rossini@usi.ch and lonneke.vanderplas@usi.ch

Abstract

We present a comprehensive approach for multiword expression (MWE) identification that combines binary token-level classification, linguistic feature integration, and data augmentation.¹ Our DeBERTa-v3-large model achieves 69.8% F1 on the CoAM dataset, surpassing the best results (Qwen-72B, 57.8% F1) on this dataset by 12 points while using 165× fewer parameters. We achieve this performance by (1) reformulating detection as binary token-level START/END/INSIDE classification rather than span-based prediction, (2) incorporating NP chunking and dependency features that help discontinuous and NOUN-type MWEs identification, and (3) applying oversampling that addresses severe class imbalance in the training data. We confirm the generalization of our method on the STREUSLE dataset, achieving 78.9% F1. These results demonstrate that carefully designed smaller models can substantially outperform LLMs on structured NLP tasks, with important implications for resource-constrained deployments.

1 Introduction

Multiword expressions (MWEs) challenge NLP systems due to their syntactic flexibility and semantic non-compositionality (Constant et al., 2017; Sag et al., 2002). They are also widely acknowledged to be frequent in natural language (Baldwin and Kim, 2010). Accurate detection nevertheless remains difficult, particularly for discontinuous patterns in which components may be separated by intervening material, as in *look the information up*, where *look...up* forms the MWE while *the information* does not (Rohanian et al., 2019).

MWE identification is traditionally formulated as sequence-labeling tasks using BIO (Ramshaw and Marcus, 1995) or extended tagsets (e.g.,

¹Model available at <https://huggingface.co/DiegoRossini/mwe-detection-deberta>. Code at <https://github.com/DiegoRossini/binary-mwe-detection>.

BILOU (Ratinov and Roth, 2009)) to capture discontinuities. These multi-class labeling schemes assign each token a single label from a predefined tagset. Current state-of-the-art approaches often focus on specific MWE types, e.g., transformer-based models (Swaminathan and Cook, 2023; Tanner and Hoffman, 2023) that incorporate linguistic features like dependency parsing for verbal discontinuous patterns (Nasr et al., 2015) and chunk features for nominal detection (Cordeiro et al., 2016). Existing methods for all-type MWE rely on massive language models with billions of parameters (Ide et al., 2025) and lack systematic integration of linguistic knowledge. Moreover, while binary classification approaches have proven successful for span-based NER (Eberts and Ulges, 2020; Li et al., 2023; Tan et al., 2020; Fei et al., 2021), to the best of our knowledge, this has not been explored yet for the task of MWE identification.

Several datasets support MWE research. PARSEME shared tasks (Savary et al., 2017; Ramisch et al., 2018, 2020; Walsh et al., 2018) provide multilingual corpora for verbal MWEs. DiMSUM (Schneider et al., 2016) covers multiple types but suffers from annotation inconsistencies (Tanner and Hoffman, 2023). STREUSLE (Schneider et al., 2018) augments the English Web Treebank with various MWE types and emphasizes supersense labeling over explicit MWE type distinctions. The recently introduced CoAM (Ide et al., 2025) addresses these limitations by offering comprehensive coverage of all MWE types with explicit type annotations (NOUN, VERB, MOD/CONN, CLAUSE, OTHER), enabling fine-grained analyses previously unavailable. While limited to 1,301 sentences, CoAM’s careful curation provides crucial evaluation capabilities.

Leveraging CoAM’s comprehensive annotations for primary evaluation, we present an approach that reformulates MWE identification as binary token-level classification—departing from traditional

multi-class sequence labeling to predict three independent binary labels per token. Our DeBERTa-v3-large model (He et al., 2021) achieves 69.8% F1 on CoAM, surpassing the state-of-the-art on this dataset, Qwen-72B (57.8% F1), by 12 points while using 165× fewer parameters. We further validate generalization on STREUSLE, achieving 78.9% F1 after mapping its categories to CoAM-compatible types, confirming that our approach transfers effectively across datasets with different annotation schemes.

2 Method

We detail our approach, beginning with the dataset description and our binary token-level classification framework, followed by linguistic feature integration and data augmentation.

2.1 Datasets

CoAM (Ide et al., 2025) contains 1,301 English sentences (780 train, 521 test) with 867 annotated MWEs, distributed over the following types: NOUN (34.3% compounds and idiomatic nominals), VERB (37.5% phrasal and light verb constructions), MOD/CONN (23.9% complex modifiers), CLAUSE (1.6% sentential idioms), and OTHER (2.8%). MWEs are 88.7% continuous and 11.3% discontinuous, spanning 2–13 tokens.

STREUSLE (Schneider et al., 2018) v4.6 (UD v2.16) provides a complementary evaluation setting with 1,530 training (2,448 MWEs), 217 development (287 MWEs), and 209 test (284 MWEs) sentences. We map STREUSLE’s fine-grained categories (e.g., *verb-particle/construction*, *noun/compound*) to CoAM-compatible types: verbal constructions to VERB, nominal expressions to NOUN, adverbial/prepositional phrases to MOD/CONN, and remaining categories to OTHER. No STREUSLE annotations correspond to the CLAUSE category. MWEs span 2–7 tokens, with 10.2% being discontinuous.

2.2 Binary Token-Level Classification and MWE Reconstruction

We reformulate MWE identification as binary token-level classification, inspired by its success in NER (Eberts and Ulges, 2020; Li et al., 2023; Tan et al., 2020; Fei et al., 2021): instead of assigning each token a single label from a multi-class tagset (e.g., B-MWE, I-MWE, O), we make three independent binary decisions per token, predicting

whether it serves as a START, END, or INSIDE component of a MWE. Our approach predicts three independent probabilities for each token x_i : p_i^{start} , p_i^{end} , and p_i^{inside} . For instance, in *looked the information up*:

- *looked*: START=1, END=0, INSIDE=0
- *the*: START=0, END=0, INSIDE=0
- *information*: START=0, END=0, INSIDE=0
- *up*: START=0, END=1, INSIDE=0

Our model would ideally predict high START probability for *looked*, high END probability for *up*, and low INSIDE probabilities for the intervening words, correctly identifying the discontinuous MWE $\{\textit{looked}, \textit{up}\}$. Crucially, our binary token-level approach scales linearly with sequence length $O(n)$, as we make three independent predictions per token, whereas span-based enumeration requires quadratic complexity $O(n^2)$ to evaluate all possible token pairs.

To train this binary classification framework, we must first convert CoAM’s original span-based annotations. CoAM provides span-based annotations (token indices for each MWE), which we convert to token-level labels through a projection system generating three binary vectors that form our START/END/INSIDE encoding scheme: $\text{start}[i] = 1$ if token i is the first token of any MWE, $\text{end}[i] = 1$ if token i is the last token of any MWE, and $\text{inside}[i] = 1$ if token i appears strictly between start/end positions. Projections are stored as versioned JSON artifacts with SHA256 hashing for reproducibility.

At inference, these independent predictions must be reconstructed into MWE spans. We form candidate spans from token pairs (s, e) with $e > s$ such that $p_s^{\text{start}} \geq \tau_{\text{start}}$ and $p_e^{\text{end}} \geq \tau_{\text{end}}$. A span is accepted if its width $(e - s + 1) \leq 13$ (dataset maximum), intermediate tokens $t \in (s, e)$ are included when $p_t^{\text{inside}} \geq \tau_{\text{inside}}$, and the reconstructed expression contains 2–6 members, matching CoAM’s MWE length distribution. Candidates are stored as sorted sets of token indices. After training, we tune the three thresholds $(\tau_{\text{start}}, \tau_{\text{end}}, \tau_{\text{inside}})$ independently on the development set via grid search, finding optimal values that maximize the F1 score. Our approach thus focuses on MWE segmentation and boundary detection; type classification could be added as a second-stage task but is beyond the scope of this work (see Appendix A).

		Q-72B	BBS	BBT	BLT	DBT	DLT	DBT+l	DLT+l	DBT+o	DLT+o	DBT+a	DLT+a	DBT+lo	DLT+lo	DBT+la	DLT+la
CoAM	Prec	63.8	51.6	62.5	67.4	62.1	64.0	64.0	52.6	65.0	61.3	68.2	61.1	63.6	69.3	66.4	70.0
	Rec	52.8	55.9	55.6	53.3	65.4	61.2	68.0	69.3	61.9	63.5	58.5	69.3	66.4	70.3	63.8	66.4
	F1	57.8	53.7	58.9	59.5	63.7	62.6	65.9	59.8	63.4	62.4	63.0	64.9	65.0	69.8	65.1	68.2
STREUSLE	Prec	–	11.1	58.1	68.4	68.8	64.2	71.3	75.5	67.1	62.1	68.4	63.3	71.5	72.4	73.5	70.0
	Rec	–	90.1	78.2	72.5	78.5	84.5	80.6	82.4	82.7	87.3	81.7	83.8	82.0	82.4	85.2	84.5
	F1	–	19.8	66.8	70.4	73.3	72.9	75.7	78.8	74.1	72.6	74.5	72.1	76.4	77.1	78.9	76.5

Table 1: Performance comparison on CoAM and STREUSLE test sets. On CoAM, DLT+lo achieves 69.8% F1, outperforming Qwen-72B (57.8% F1). On STREUSLE, DBT+la achieves 78.9% F1. Best results per dataset in bold.

		Q-72B	BBS	BBT	BLT	DBT	DLT	DBT+l	DLT+l	DBT+o	DLT+o	DBT+a	DLT+a	DBT+lo	DLT+lo	DBT+la	DLT+la
CoAM	CL	28.6	42.9	28.6	42.9	42.9	42.9	42.9	42.9	42.9	42.9	28.6	42.9	28.6	85.7	42.9	71.4
	M/C	57.7	67.6	75.7	76.6	82.9	76.6	80.2	82.0	76.6	79.3	67.6	78.4	76.6	84.7	73.0	80.2
	NOUN	42.1	43.8	38.0	35.5	57.9	51.2	60.3	65.3	53.7	52.9	53.7	62.7	60.3	66.9	63.6	58.7
	VERB	59.7	59.0	56.8	51.1	59.7	59.0	66.9	64.7	59.0	62.6	58.3	66.2	66.9	66.9	59.0	63.3
STREUSLE	M/C	–	90.6	81.2	78.1	84.3	89.1	85.9	87.5	92.2	90.6	89.1	89.1	87.5	84.4	92.2	89.1
	NOUN	–	88.4	76.9	69.2	78.5	86.1	83.1	84.6	82.3	86.9	81.5	84.6	83.1	82.3	86.1	85.4
	VERB	–	90.9	74.2	72.7	74.2	75.7	75.0	75.7	75.7	81.8	77.2	75.7	75.7	80.3	78.9	80.3
	OTH	–	95.8	87.5	75.0	75.0	87.5	72.7	75.0	79.2	95.8	75.0	83.8	79.2	83.3	79.2	79.2

Table 2: Type-specific recall (%) across MWE types. CoAM: CL=CLAUSE (N=7), M/C=MOD/CONN (N=111), NOUN (N=121), VERB (N=139). STREUSLE: M/C (N=64), NOUN (N=130), VERB (N=66), OTH=OTHER (N=24). CoAM OTHER (N=3) excluded; STREUSLE has no CLAUSE annotations. Note: on STREUSLE, DLT+o achieves higher recall than DBT+la in 3/4 categories but lower overall F1 (72.6% vs 78.9%, see Table 1) due to over-prediction (399 vs 329 predictions against 284 gold MWEs).

2.3 Linguistic Feature Integration

We enrich DeBERTa-v3-large with two forms of syntactic knowledge:

NP Chunk Features. Tokens are tagged as inside or outside a noun phrase (NP) using spaCy `en_core_web_lg` (Honnibal et al., 2020). These binary tags are mapped to 16-dimensional learnable embeddings and concatenated with contextualized token states. Chunk embeddings help the model detect nominal MWEs such as *stock market*, a class noted in the CoAM study as particularly difficult to recall compared to MOD/CONN or VERB MWEs.

Dependency features. Following extensive prior work on dependency-based MWE identification (Lin, 1999; Wehrli, 2000; Green et al., 2013; Nasr et al., 2015), we compute shortest dependency path lengths between all token pairs using NetworkX (Hagberg et al., 2008) graph traversal, capped at 5. These distances serve dual purposes: (1) as learnable embeddings during encoding, where we average each token’s dependency distances to other tokens and map them to 32-dimensional embeddings, and (2) as hard constraints during reconstruction, where we reject discontinuous MWE candidates with consecutive

members having dependency distance greater than 4, filtering out implausible token combinations.

2.4 Data Augmentation

Beyond linguistic features, addressing the severe class imbalance in the training data proved crucial for model performance. We evaluated two data augmentation approaches (He and Garcia, 2009): (1) oversampling of training sentences containing MWEs, and (2) lexical substitution, replacing single tokens outside MWE boundaries with semantically similar words. Both approaches were tested by selecting 10%, 20%, 30%, or 40% of the training sentences for augmentation.

The effectiveness of each strategy depends on dataset size. On the smaller CoAM (486 training MWEs), oversampling proved more effective, with 30% selection ratio proving optimal. We attribute this to CoAM’s limited size combined with its carefully curated annotations: the dataset appears too small for the model to generalize from lexical variations, while any modifications risk disrupting the subtle patterns that define MWEs. Conversely, on the larger STREUSLE (2,448 training MWEs), lexical substitution proved more effective, as sufficient

		Q-72B	BBS	BBT	BLT	DBT	DLT	DBT+l	DLT+l	DBT+o	DLT+o	DBT+a	DLT+a	DBT+lo	DLT+lo	DBT+la	DLT+la
CoAM	C-P	–	51.6	77.4	81.1	77.5	78.9	74.1	65.7	79.1	76.5	80.2	76.5	72.5	76.9	75.8	80.4
	C-R	–	63.0	59.8	58.3	70.4	66.3	72.8	74.9	66.0	67.5	62.1	73.1	70.1	74.9	67.8	70.4
	C-F1	57.3	56.6	67.4	67.8	73.8	72.0	73.4	70.0	71.9	71.7	70.0	74.7	71.3	75.9	71.6	75.1
	D-P	–	0.0	12.8	10.3	11.7	11.3	17.8	9.4	16.0	14.4	20.0	15.6	22.5	25.9	21.9	23.1
	D-R	–	0.0	23.3	14.0	25.6	20.9	30.2	25.6	30.2	32.6	30.2	39.5	37.2	34.9	32.6	34.9
	D-F1	17.1	0.0	16.5	11.9	16.1	14.6	22.4	13.8	21.0	20.0	24.1	22.4	28.1	29.7	26.2	27.8
STREUSLE	C-P	–	29.5	78.3	84.2	85.7	84.3	84.4	89.4	87.6	84.1	89.3	83.6	86.2	85.7	88.3	84.7
	C-R	–	91.8	82.3	75.6	82.3	88.6	85.1	86.3	86.2	91.3	85.5	87.8	85.5	85.1	88.6	87.4
	C-F1	–	44.6	80.3	79.7	84.0	86.4	84.7	87.8	86.9	87.6	87.4	85.6	85.8	85.4	88.4	86.1
	D-P	–	1.4	10.5	18.0	16.4	13.2	18.7	21.9	15.1	12.3	14.7	13.0	20.5	21.9	24.3	21.2
	D-R	–	75.8	41.3	44.8	44.8	48.3	41.4	48.3	51.7	51.7	48.3	48.3	51.7	55.2	58.6	58.6
	D-F1	–	2.8	16.7	25.7	24.1	20.7	25.8	30.1	23.4	19.9	22.6	20.4	29.4	31.3	34.4	31.2

Table 3: Performance on continuous (C) and discontinuous (D) MWEs. CoAM: N=338 continuous, N=43 discontinuous. STREUSLE: N=255 continuous, N=29 discontinuous. Best model per dataset in bold.

training examples enable the model to benefit from lexical variation.

3 Experiments

We conduct a systematic ablation study with 15 model configurations to isolate the contributions of each component: our binary token-level classification approach, model architecture choice, linguistic feature integration, and data augmentation.

3.1 Experimental Design

Our experiments build incrementally across four phases. First, we compare BERT-base span-based prediction (BBS) against BERT-base/large token-level classification (BBT/BLT). Second, we compare BBT/BLT against DeBERTa-v3-base/large (DBT, DLT), whose disentangled attention should better capture discontinuous patterns (He et al., 2021). Third, we compare DBT/DLT against variants augmented with linguistic features (+l) targeting NOUN and discontinuous MWEs. Finally, we compare DBT+l/DLT+l against data-augmented variants using oversampling (+o) and lexical substitution (+a), both independently and combined with linguistic features (see Table 1).

3.2 Setup

We convert the span-based annotations from both datasets to our START/END/INSIDE format by projecting the provided token indices. For STREUSLE, we additionally apply the type mapping described in Section 2.1. This harmonization enables direct comparison across datasets while preserving type-specific analysis capabilities.

We train all models on NVIDIA GH200 hardware. Our baseline for CoAM is the dataset authors’

Qwen-72B (57.8% F1, Ide et al., 2025), and BERT-base span-based (BBS) as additional comparison across both datasets. We report exact span matching micro F1 with breakdowns by type and continuity. Hyperparameters were optimized via Optuna across learning rate, batch size, dropout, and weight decay. Our best model on CoAM (DLT+lo) uses learning rate $3e-5$, batch size 16, dropout 0.3, weight decay 0.01 with early stopping at epoch 8. Our best model on STREUSLE (DBT+la) uses learning rate $2e-5$, batch size 4, dropout 0.3, weight decay 0.05 with early stopping at epoch 10. Training requires 10–25 minutes per model, making our approach highly practical compared to large model fine-tuning.

3.3 Main Results

Our systematic ablation results in Table 1 reveal several critical insights about each component’s contribution. Our binary token-level classification yields +5.2 F1 improvement over span-based prediction (BBS: 53.7 → BBT: 58.9), validating that three independent START/END/INSIDE decisions per token provide richer gradient signals than sparse span-level feedback. DBT achieves 63.7% F1, surpassing BBT by +4.8 points and even outperforming BLT, confirming that disentangled attention mechanisms are particularly suited for MWE detection where content and position interact in an intricate way.

Adding linguistic features to DBT improves F1 by +2.2 points (63.7 → 65.9), with particularly strong recall improvements. However, DLT shows a surprising decrease with linguistic features alone (62.6 → 59.8), suggesting potential overfitting on the small dataset. A critical pattern emerges in the

interaction between augmentation and model capacity: without augmentation, base models often outperform large counterparts, but with augmentation, large models excel—DLT+lo achieves 69.8% F1 while DBT+lo reaches only 65.0%. Oversampling consistently outperforms lexical substitution when combined with linguistic features, which may be attributed to CoAM’s high annotation quality where lexical changes risk disrupting subtle MWE patterns (see Appendix A.2 and Appendix B).

Our DeBERTa-v3-large model achieves 69.8% F1 on CoAM, surpassing the state-of-the-art on this dataset, Qwen-72B (57.8% F1), by 12 points while using 165× fewer parameters. Other comparisons to previous work are complex because we focus on all types of MWEs, which differs fundamentally from recent frameworks like Swaminathan and Cook (2023) that are tailored specifically to verbal MWEs or binary idiomaticity classification rather than comprehensive span identification across all MWE types (see Framework Comparison Limitations).

On STREUSLE, the ablation pattern replicates. The baseline BBS achieves 19.8% F1 due to massive over-prediction (11.1% precision), while binary token-level classification provides substantial improvement: DBT reaches 73.3% F1, a 53.5-point gain mirroring CoAM’s pattern. Linguistic features consistently improve both backbones, with DLT+I achieving 78.8% F1. Crucially, the optimal augmentation strategy differs: lexical substitution outperforms oversampling (DBT+la: 78.9% F1 vs DBT+lo: 76.4% F1), contrasting with CoAM where oversampling proved superior—we attribute this to STREUSLE’s larger training set (2,448 MWEs) enabling the model to benefit from lexical variation.

3.4 Analysis by MWE Type and Continuity

Table 3² highlights dramatic improvements for discontinuous MWEs. The baseline span-based model completely fails on discontinuous MWEs (0% recall), while our binary token-level classification immediately enables detection (BBT: 23.3% recall). Adding linguistic features to DBT increases discontinuous recall by +4.6 points, confirming that explicit syntactic knowledge helps identify separated MWE components. Our best model on

²Type-specific precision cannot be computed as our model predicts only reconstructed span boundaries without type labels. See Limitations (Type Prediction Limitations) for a detailed discussion.

CoAM (DLT+lo) achieves 34.9% discontinuous recall while maintaining 25.9% precision (see Appendix B for detailed error analysis).

In addition to handling discontinuities, we also analyze performance across MWE types to understand where improvements are most pronounced. Table 2 demonstrates type-specific improvements across categories. CLAUSE expressions improve from Qwen-72B’s 28.6% to 85.7% recall, though with only 7 test instances, this represents detecting 6 versus 2 examples. NP chunking knowledge integration helps NOUN expressions, improving recall from 57.9% to 60.3% (DBT → DBT+I), with stronger effects in large models. MOD/CONN expressions show consistent gains, with DLT+lo achieving 84.7% recall compared to Qwen’s 57.7%. Importantly, type-specific recall must be interpreted alongside precision: models with aggressive prediction strategies can achieve high recall across categories while suffering from low overall F1 due to excessive false positives.

On STREUSLE, type-specific patterns confirm generalization: MOD/CONN achieves 92.2% recall for DBT+la, followed by NOUN (86.1%) and VERB (78.9%). Regarding contiguity patterns, DBT+la achieves 88.4% F1 on continuous MWEs and 34.4% F1 on discontinuous ones. Linguistic features consistently boost discontinuous performance: DLT+I reaches 30.1% discontinuous F1 versus DLT’s 20.7%, confirming that syntactic structure helps across datasets.

4 Conclusion

We reformulate MWE identification as binary token-level classification with three independent START/END/INSIDE predictions per token, combining DeBERTa-v3 with NP chunking and dependency features plus data augmentation. On CoAM, this achieves 69.8% F1, outperforming Qwen-72B (57.8%) by 12 points with 165× fewer parameters; for discontinuous MWEs, we achieve 29.7% F1 versus Qwen-72B’s 17.1%, doubling baseline performance though these patterns remain challenging. We confirm generalization on STREUSLE, achieving 78.9% F1 with the same methodological choices, demonstrating that binary token-level classification and linguistic feature integration capture fundamental aspects of MWE structure rather than dataset-specific artifacts.

Limitations

While our approach achieves substantial improvements on CoAM, several limitations warrant discussion:

Dataset Scale and Model Generalization. The first dataset we used, CoAM, is of high quality and diverse but it is small. This may lead to potential overfitting (see Appendix A.2 for detailed analysis), but meaningful generalization is evident: the test set remains unseen during training and is guaranteed to not contain MWE patterns present in the training data, yet our model maintains strong performance. The performance difference between training and test sets (Table 5) is typical of small datasets.

On the second dataset, STREUSLE, which is larger in its training set, we confirm cross-dataset transfer (78.9% F1). The finding that simple duplication outperforms lexical substitution on CoAM while the opposite holds for STREUSLE (Section 2.4) suggests that dataset size influences optimal augmentation strategy, though preserving exact MWE patterns remains crucial as even minor lexical variations can fundamentally alter MWEs’ idiomatic properties.

Statistical Robustness. Preliminary experiments with different random seeds on a subset of models (5 out of 15 configurations) showed no significant performance variation from the reported results using seed 42, suggesting stability in our approach. However, due to limited time for resource access, we could not systematically test all model configurations across multiple seeds.

Framework Comparison Limitations. Recent work by Swaminathan and Cook (2023) achieves strong results on MWE-related tasks but addresses fundamentally different problems than comprehensive span detection, preventing direct comparison. Their MTLB-STRUCT framework targets PARSEME 1.2, which exclusively annotates verbal MWEs (light-verb constructions, verb-particle constructions, inherently reflexive verbs, and verbal idioms), making it architecturally incompatible with CoAM where 62.5% of MWEs are non-verbal (NOUN, MOD/CONN, CLAUSE, OTHER categories). Their second approach addresses SemEval 2022 Task 2, a binary sentence-level classification task determining whether a known potentially-idiomatic expression is used idiomatically or literally—fundamentally different from identifying which token spans constitute MWEs.

Discontinuous MWE Performance. Despite doubling the baseline’s recall for discontinuous MWE identification (34.9% vs 17.1%), the overall discontinuous F1 remains modest at 29.7%. The dependency distance constraints we introduce (rejecting candidates with consecutive members having distance >4) improve precision but may be overly restrictive for certain discontinuous patterns. This fundamental challenge suggests that even with linguistic enhancements, current architectures struggle with long-range discontinuous dependencies.

Type Prediction Limitations. Our binary token-level formulation predicts only span boundaries without MWE type labels, creating fundamental constraints on type-specific evaluation. While we can compute type-specific recall by checking which gold MWEs of each type were correctly identified, type-specific precision is impossible to calculate. When the model predicts a span like *systemic risk* as an MWE (false positive), we have no type prediction to determine whether the model believes this is a NOUN, VERB, or MOD/CONN expression—we only know it identified these tokens as forming some MWE. Without type-specific precision, F1 scores per type cannot be computed. This limitation means Table 2 reports only recall metrics, preventing complete type-wise performance analysis. For instance, while we know the model correctly identifies 66.9% of VERB expressions (recall), we cannot determine what percentage of its verb-like predictions are actually correct (precision) because the predictions lack type labels. This design choice prioritizes boundary detection accuracy over categorization but limits practical applications requiring typed MWE identification and restricts our ability to analyze whether false positives tend to be specific types of spurious expressions. Note that this issue only arises in the type-specific evaluation results in Table 2.

Knowledge Source Dependencies. Our linguistic features rely on spaCy’s accuracy for NP chunking and dependency parsing. Errors in these upstream components propagate through our pipeline: incorrect chunk boundaries or erroneous dependency trees directly impact model performance. Additionally, such models and pipeline components (parsers, taggers) may not be available for all languages.

Computational Requirements. Experiments were conducted on NVIDIA GH200 hardware, which may limit reproducibility for researchers without access to high-end GPUs. While train-

ing requires only 10-25 minutes per model, the hyperparameter search with multiple trials across multiple oversampling ratios demands substantial compute. Our fastest configuration still requires ~4GB GPU memory, potentially excluding older hardware.

Language Coverage. Evaluation is restricted to English, despite MWEs being a pervasive multilingual phenomenon with language-specific characteristics. The syntactic features we exploit (NP chunks, dependency paths) may transfer poorly to languages with different syntactic structures or morphological complexity.

Acknowledgments

This research was funded by the NCCR Evolving Language, Swiss National Science Foundation Agreement No. 51NF40_180888. We thank the anonymous reviewers for their helpful feedback. We also thank the people participating in the workshop *Linguistic Representations: Through the lens of compounding*, organized by Prajit Dhar, for the fruitful discussions.

References

- Timothy Baldwin and Su Nam Kim. 2010. Multiword expressions. In *Handbook of Natural Language Processing*, pages 267–292. CRC Press.
- Mathieu Constant, Gülşen Eryiğit, Johanna Monti, Lonneke van der Plas, Carlos Ramisch, Michael Rosner, and Amalia Todirascu. 2017. [Multiword expression processing: A survey](#). *Computational Linguistics*, 43(4):837–892.
- Silvio Cordeiro, Carlos Ramisch, and Aline Villavicencio. 2016. LIF&UFRGS at SemEval-2016 Task 10: Rule-based MWE identification and predominant-supersense tagging. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 910–917. Association for Computational Linguistics.
- Markus Eberts and Adrian Ulges. 2020. [Span-based joint entity and relation extraction with transformer pre-training](#). In *24th European Conference on Artificial Intelligence (ECAI 2020)*, pages 2006–2013. IOS Press.
- Hao Fei, Yafeng Ren, and Donghong Ji. 2021. Boundary enhanced neural span classification for nested named entity recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9016–9023.
- Spence Green, Marie-Catherine de Marneffe, and Christopher D. Manning. 2013. Parsing models for identifying multiword expressions. *Computational Linguistics*, 39(1):195–227.
- Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. 2008. NetworkX: Network analysis in Python. In *Proceedings of the 7th Python in Science Conference (SciPy 2008)*, pages 11–15.
- Haibo He and Edwardo A. Garcia. 2009. [Learning from imbalanced data](#). *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [DeBERTa: Decoding-enhanced BERT with disentangled attention](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength natural language processing in Python](#).
- Yusuke Ide, Joshua Tanner, Adam Nohejl, Jacob Hoffman, Justin Vasselli, Hidetaka Kamigaito, and Taro Watanabe. 2025. CoAM: Corpus of all-type multiword expressions. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 27004–27021, Vienna, Austria. Association for Computational Linguistics.
- Yongliang Li, Xinyu Ma, Xiaoqing Zhou, and Jizhong Qi. 2023. [T2-ner: A two-stage span-based framework for unified named entity recognition with templates](#). *Transactions of the Association for Computational Linguistics*, 11:224–240.
- Dekang Lin. 1999. Automatic identification of non-compositional phrases. In *Proceedings of the 37th annual meeting of the Association for Computational Linguistics*, pages 317–324.
- Alexis Nasr, Carlos Ramisch, José Deulofeu, and André Valli. 2015. Joint dependency parsing and multiword expression tokenization. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 1116–1126.
- Carlos Ramisch, Silvio Ricardo Cordeiro, Agata Savary, Veronika Vincze, Verginica Barbu Mititelu, Archana Bhatia, Maja Buljan, Marie Candito, Polona Gantar, Voula Giouli, and 1 others. 2018. Edition 1.1 of the PARSEME shared task on automatic identification of verbal multiword expressions. In *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*, pages 222–240. Association for Computational Linguistics.
- Carlos Ramisch, Agata Savary, Bruno Guillaume, Jakub Waszczuk Abigail Walsh, Marie Candito, Ashim Gupta Mehmet Kurfalı, Verginica Barbu Mititelu, Gözde Berk, Johanna Monti, and 1 others. 2020.

- Edition 1.2 of the PARSEME shared task on semi-supervised identification of verbal multiword expressions. In *Proceedings of the Joint Workshop on Multiword Expressions and Electronic Lexicons*, pages 107–118. Association for Computational Linguistics.
- Lance A. Ramshaw and Mitchell P. Marcus. 1995. Text chunking using transformation-based learning. In *Third Workshop on Very Large Corpora*, pages 82–94. Association for Computational Linguistics.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, pages 147–155.
- Omid Rohanian, Shiva Taslimipour, Samaneh Kouchaki, Le An Ha, and Ruslan Mitkov. 2019. Bridging the gap: Attending to discontinuity in identification of multiword expressions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2692–2698, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ivan A Sag, Timothy Baldwin, Francis Bond, Ann Copestake, and Dan Flickinger. 2002. Multiword expressions: A pain in the neck for NLP. In *International conference on intelligent text processing and computational linguistics*, pages 1–15. Springer.
- Agata Savary, Carlos Ramisch, Silvio Cordeiro, Federico Sangati, Veronika Vincze, Behrang QasemiZadeh, Marie Candito, Fabienne Cap, Voula Giouli, Ivelina Stoyanova, and Antoine Doucet. 2017. The PARSEME shared task on automatic identification of verbal multiword expressions. In *Proceedings of the 13th Workshop on Multiword Expressions (MWE 2017)*, pages 31–47, Valencia, Spain. Association for Computational Linguistics.
- Nathan Schneider, Dirk Hovy, Anders Johannsen, and Marine Carpuat. 2016. SemEval-2016 Task 10: Detecting minimal semantic units and their meanings (DiMSUM). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 546–559, San Diego, California. Association for Computational Linguistics.
- Nathan Schneider, Jakob Prange, Emily Danchik, Luke Gessler, Samantha Onuffer, Nora Mordowanec, Nour Kazour, and Aviram Nguyen. 2018. Comprehensive supersense disambiguation of English prepositions and possessives. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1853–1865. Association for Computational Linguistics.
- Raghavendra Swaminathan and Paul Cook. 2023. Token-level identification of multiword expressions using pre-trained multilingual language models. In *Proceedings of the 19th Workshop on Multiword Expressions (MWE 2023)*, pages 1–6, Dubrovnik, Croatia. Association for Computational Linguistics.
- Chuanqi Tan, Wei Qiu, Mosha Chen, Rui Wang, and Fei Huang. 2020. Boundary enhanced neural span classification for nested named entity recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9016–9023.
- Joshua Tanner and Jacob Hoffman. 2023. MWE as WSD: Solving multiword expression identification with word sense disambiguation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 181–193. Association for Computational Linguistics.
- Kevin Walsh, Carlos Ramisch, Agata Savary, and Silvio Cordeiro. 2018. A dataset of verbal multiword expressions annotated with syntactic variants. In *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*, pages 101–110. Association for Computational Linguistics.
- Eric Wehrli. 2000. Incremental parsing and multiword expression. In *Proceedings of the Workshop on Multiword Expressions: Analysis, Acquisition and Treatment*, pages 33–40.

A Implementation and Performance Details

A.1 Threshold Values and Hyperparameters

Table 4 presents the optimal threshold values obtained through grid search on the development set for each model configuration. The thresholds control the sensitivity of START, END, and INSIDE predictions during MWE reconstruction. Lower thresholds generally increase recall at the cost of precision.

Model	τ_{start}	τ_{end}	τ_{inside}
<i>CoAM</i>			
BBS	–	–	–
BBT	0.45	0.20	0.20
BLT	0.60	0.20	0.20
DBT	0.45	0.60	0.20
DLT	0.50	0.20	0.20
DBT+l	0.30	0.20	0.20
DLT+l	0.50	0.30	0.20
DBT+o	0.20	0.20	0.20
DLT+o	0.50	0.60	0.40
DBT+a	0.50	0.60	0.30
DLT+a	0.40	0.40	0.20
DBT+lo	0.30	0.50	0.40
DLT+lo	0.50	0.60	0.20
DBT+la	0.60	0.40	0.20
DLT+la	0.60	0.40	0.20
<i>STREUSLE</i>			
DBT+la	0.60	0.40	0.20

Table 4: Optimal threshold values for each model configuration. BBS uses span-based prediction and does not require thresholds. Best models per dataset shown in bold.

All models were trained with AdamW optimizer and linear learning rate scheduling. The best performing model (DLT+lo) used: learning rate 3e-5, batch size 16, dropout 0.3, weight decay 0.01, and early stopping with patience 3.

A.2 Training Dynamics and Model Performance

Table 5 shows training and test performance across all configurations. The train-test gaps (17.0-30.7% F1) are expected given the limited training data (780 sentences) and the complexity of distinguishing idiomatic from compositional expressions. Notably, our best model maintains 69.8% test F1

on unseen data—the test set contains MWE patterns and combinations not present in the training data—validating our architectural choices despite the challenging dataset size.

Model	Train F1 (%)	Test F1 (%)	Gap (Δ)
<i>CoAM</i>			
BBS	70.7	53.7	17.0
BBT	89.6	58.9	30.7
BLT	90.1	59.5	30.6
DBT	83.2	63.7	19.5
DLT	88.9	62.6	26.3
DBT+l	92.3	65.9	26.4
DLT+l	86.1	59.8	26.3
DBT+o	89.3	63.4	25.9
DLT+o	88.7	62.4	26.3
DBT+a	86.5	63.0	23.5
DLT+a	89.6	64.9	24.7
DBT+lo	94.1	65.0	29.1
DLT+lo	95.0	69.8	25.2
DBT+la	93.9	65.1	28.8
DLT+la	94.6	68.2	26.4
<i>STREUSLE</i>			
DBT+la	86.8	78.9	7.9

Table 5: Train vs test F1 scores across all configurations. Best models per dataset shown in bold. STREUSLE’s smaller gap reflects its larger training set.

Figure 1 illustrates the training dynamics of our best model. The plateauing development performance after epoch 5, while training loss continues decreasing, demonstrates effective early stopping preventing overtraining.

B Error Analysis

This appendix analyzes our best model’s (DLT+lo) predictions on the test set to understand both its strong performance (69.8% F1) and areas for improvement. The analysis reveals systematic patterns in the model’s decisions, demonstrating that it has learned meaningful linguistic regularities rather than arbitrary memorization.

B.1 Overall Performance Patterns

The model generates 387 predictions for 381 gold MWEs, achieving well-balanced precision (69.3%) and recall (70.3%). Notably, 359 of 521 gold sentences (68.9%) are predicted perfectly, indicating robust pattern learning. Among the 162 imperfect sentences out of 521 gold ones, 87 contain only false negatives (missed MWEs), 51 contain only false positives (spurious detections), and 24 show both error types, typically representing partial recognition where the model correctly identifies MWE presence but misaligns boundaries.

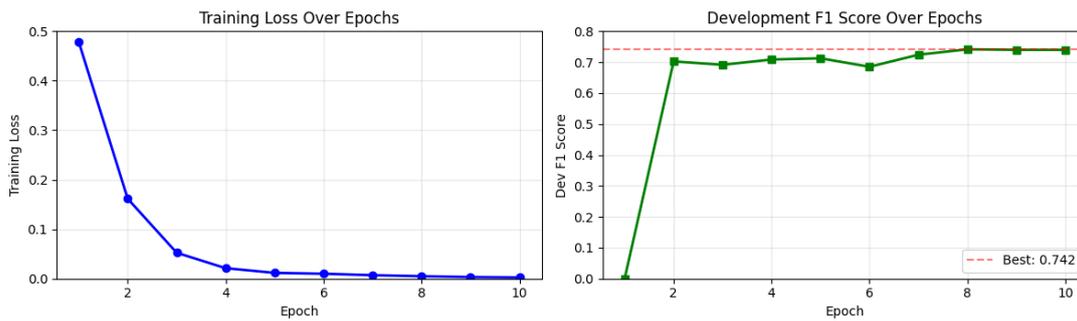


Figure 1: Learning curves for DLT+Io: training loss decreases while development F1 plateaus at epoch 8 (approximately 0.74), triggering early stopping. The initial fluctuation (epochs 2-3) reflects the model adapting to the challenging task of distinguishing genuine MWEs from compositional phrases.

The prediction distribution reveals interesting patterns: 329 continuous and 58 discontinuous predictions against 338 continuous and 43 discontinuous gold MWEs. While the continuous predictions aligns well, discontinuous predictions show higher variance, with 43 false positives (36.1% of all FPs) despite discontinuous MWEs comprising only 11.3% of the gold standard.

B.2 Continuous MWE Recognition

Continuous expressions achieve strong 74.9% recall (253/338 correctly identified). This 74.9% recall represents a significant achievement, correctly identifying three-quarters of continuous MWEs despite training on only 780 sentences. Performance varies systematically by expression type:

High-confidence categories: Discourse markers (*In fact*, *Of course*, *According to*) and phrasal verbs (*took over*, *brought down*, *turn out*) achieve near-perfect detection when components are adjacent. Fixed nominals like *Prime Minister* (correctly identified in all test instances) and *exchange rate* demonstrate reliable institutional term recognition.

False positives (76 instances): These split into interpretable categories. Technical overdetection (23 cases) flags domain bigrams like *systemic risk* as MWEs—specialized usage of this bigram with high co-occurrence, but arguably not a MWE. Incomplete spans (19 cases) systematically miss required elements, such as predicting *open the way* instead of the gold *open the way for*.

False negatives (85 instances): The model fails to detect MWEs that appear infrequently in training (<7 occurrences). These include: (1) metaphorical phrases like *watershed moment* and *rose-tinted glasses* that require semantic understanding beyond

surface patterns, and (2) short functional MWEs like *act on* and *yet again*, where the individual words (e.g., "act", "on") appear frequently throughout the corpus in non-MWE contexts, making it difficult for the model to recognize when they specifically form a MWE.

B.3 Discontinuous MWE Progress

Discontinuous detection demonstrates meaningful progress with 34.9% recall (15/43), doubling baseline performance and achieving 25.9% precision.

Successful patterns: The 15 correct predictions follow learnable structures: (1) verb-particle with single interveners like *turn [the situation] into* where dependency distance ≤ 3 , (2) possessive frames like *on [their] own*, and (3) simple modifier insertions like *free [X] markets*.

Error analysis: The 43 discontinuous false positives primarily result from independent START/END predictions creating semantically invalid spans (*not...risk* spanning 8 tokens) or partial captures missing critical elements (predicting *turning...on* instead of the gold *turning...gun...on*). The 28 false negatives often exceed our dependency distance threshold, like *called...back* with an 11-token gap, or involve multiple discontinuities our architecture cannot represent.

B.4 Type-Specific Performance

The model demonstrates strong performance across different MWE types. Formulaic expressions (MOD/CONN) achieve excellent 84.7% recall, successfully identifying phrases like *In fact* and *According to* that follow consistent syntactic patterns. CLAUSE expressions reach 85.7% recall despite having only 7 test instances, demonstrating the

Type	Rec%	N
CLAUSE	85.7	7
MOD/CONN	84.7	111
VERB	66.9	139
NOUN	66.9	121

Table 6: Type-specific recall performance across MWE categories for our best model (DLT+lo).

model’s ability to generalize from limited examples.

VERB and NOUN categories both achieve 66.9% recall, a notable accomplishment given their inherent challenges. VERB expressions require distinguishing idiomatic uses (e.g., *take off* meaning "succeed") from compositional ones (physically removing something). NOUN expressions span from technical terms to creative compounds, requiring the model to recognize both conventional phrases (*stock market*) and novel compositional combinations.

Overall, the error analysis reveals that our model’s mistakes are systematic and interpretable, often reflecting genuine linguistic ambiguities (e.g., compositional vs. idiomatic uses) or structural challenges rather than random failures.