

# Creating a Hybrid Rule and Neural Network Based Semantic Tagger using Silver Standard Data: the PyMUSAS framework for Multilingual Semantic Annotation

Andrew Moore<sup>\*1</sup>, Paul Rayson<sup>\*1,8</sup>, Dawn Archer<sup>2</sup>, Tim Czerniak<sup>3</sup>, Dawn Knight<sup>4</sup>, Daisy Lal<sup>1</sup>, Gearóid Ó Donnchadha<sup>5</sup>, Mícheál Ó Meachair<sup>6</sup>, Scott Piao<sup>1</sup>, Elaine Uí Dhonnchadha<sup>3</sup>, Johanna Vuorinen<sup>†5</sup>, Yan Yabo<sup>7</sup>, Xiaobin Yang<sup>7</sup>

<sup>1</sup>UCREL, Lancaster University, UK; <sup>2</sup>Manchester Metropolitan University, UK;

<sup>3</sup>Centre for Language and Communication Studies, Trinity College, Dublin, Ireland;

<sup>4</sup>School of English, Communication and Philosophy, Cardiff University, Wales;

<sup>5</sup>Independent Researcher;

<sup>6</sup>Fiontar & Scoil na Gaeilge, Dublin City University, Ireland;

<sup>7</sup>Hubei University, China; <sup>8</sup>VinUniversity, Vietnam

a.p.moore, p.rayson@lancaster.ac.uk

## Abstract

Word Sense Disambiguation (WSD) has been widely evaluated using the semantic frameworks of WordNet (Maru et al., 2022), BabelNet (Pasini et al., 2021), and the Oxford Dictionary of English (Gadetsky et al., 2018; Chang et al., 2018). However, for the UCREL Semantic Analysis System (USAS) framework, no open extensive evaluation has been performed beyond lexical coverage or single language evaluation. In this work, we perform the largest semantic tagging evaluation of the rule based system that uses the lexical resources in the USAS framework covering five different languages using four existing datasets and one novel Chinese dataset. We create a new silver labelled English dataset, to overcome the lack of manually tagged training data, that we train and evaluate various mono and multilingual neural models in both mono and cross-lingual evaluation setups with comparisons to their rule based counterparts, and show how a rule based system can be enhanced with a neural network model. The resulting neural network models, including the data they were trained on, the Chinese evaluation dataset, and all of the code have been released as open resources.

**Keywords:** Semantic tagging, Lexicons, Multilingual Annotation, Machine Learning

## 1. Introduction and Related Work

Word Sense Disambiguation (WSD) is the task of assigning a word with a pre-defined sense inventory according to a given context. WSD as a field has progressed significantly from early uses of feature-based Support Vector Machines (SVMs) per word-type<sup>1</sup> (Zhong and Ng, 2010) to more recent fine-tuning of Pre-trained Language Models (PLMs) (Barba et al., 2021). Current Large Language Models (LLMs) (Basile et al., 2025), with prompting of LLMs (Meconi et al., 2025; Basile et al., 2025) achieve State-Of-The-Art (SOTA) performance on various evaluation setups. In addition to the development of models, WSD has grown with respect to datasets for both training (e.g. SemCor and the WordNet Gloss Corpus) (Miller et al., 1994; Langone et al., 2004) and evaluating (Raganato et al., 2017; Maru et al., 2022) on English as well as training (Conia et al., 2024) and evaluating (Pasini et al., 2021) on a variety of other languages. However, all of this has been created within the WordNet

(Miller, 1995) or BabelNet (Navigli et al., 2021) semantic frameworks.

In this work, we focus on evaluating and extending semantic tagging tools and datasets within the USAS semantic framework (Rayson et al., 2004) which, compared to WordNet and BabelNet, provides a more coarse-grained sense inventory, as shown in Table 1<sup>23</sup>. The framework is language independent, similar to other semantic frameworks such as BabelNet and Open Multilingual WordNet (Bond and Foster, 2013).

Most of the previous work within the USAS framework has focused on developing rule-based semantic taggers, which at their core rely on language-specific semantic lexicons, including the original C version of the English semantic tagger (Rayson et al., 2004)<sup>4</sup>. More recently, the system has been re-developed and expanded to 10 languages through a Python based framework called Py-

<sup>2</sup>Statistics for WordNet taken from: <https://wordnet.princeton.edu/documentation/wnstats7wn>

<sup>3</sup>Statistics for BabelNet taken from: <https://babelnet.org/statistics>

<sup>4</sup>C is the programming language that the original English rule based tagger was written in.

<sup>\*</sup>Corresponding authors. <sup>†</sup>Johanna Vuorinen and the referenced Laura Löfberg refer to the same individual.

<sup>1</sup>Word-type here refers to the word and its Part Of Speech.

MUSAS<sup>5</sup> that utilises language specific lexicons which have been developed manually and semi-automatically through previous works and a large community effort (Piao et al., 2015, 2016).

In this research, we demonstrate the power of combining neural networks and the USAS rule-based model to create a novel hybrid model for semantic tagging. We extend the work of Ezeani et al. (2019) for the Welsh language by fine-tuning larger pre-trained models, specifically Pre-trained Language Models (PLMs). The PLMs were fine-tuned on over 5 million English-only silver labelled tokens generated by the English rule-based tagger, thus alleviating the need for a manually annotated training dataset. As the PLMs are both English and multilingual based, in addition to the Welsh neural tagger, we produce the first neural network models trained specifically for USAS tagging for English, Irish, Finnish and Chinese.

We collate existing manually annotated datasets across the four different languages, with an additional newly labelled Chinese dataset, to generate a new benchmark covering the five languages, allowing us for the first time to compare rule-based, neural network, and hybrid models. In addition, this is the first time models have been contextually evaluated for USAS tagging beyond a single language with open datasets (Piao et al., 2015; Ezeani et al., 2019; Czerniak and Uí Dhonnchadha, 2024). All data within this paper have been made publicly available or are accessible upon request<sup>6</sup>.

Sense Inventory	No. Senses
USAS	232
WordNet 3.0	117,659
BabelNet v4.0	15,780,364
BabelNet v5.3	22,892,310

Table 1: The number of unique senses in each sense inventory (No. Senses), also known as synsets, for each sense inventory. WordNet 3.0 is displayed here as it is the sense inventory used in the most popular benchmarking datasets (Raganato et al., 2017; Maru et al., 2022) and BabelNet v4.0 and v5.3 are shown as they are used in a popular multilingual benchmarking dataset (Pasini et al., 2021) and is the latest version respectively.

The main contributions of this work include:

1. The first neural network based English and multilingual semantic taggers that have been trained specifically for the USAS tagset on English silver labelled data.

<sup>5</sup><https://github.com/UCREL/pymusas>

<sup>6</sup>The Irish data that is used for evaluation in this paper is available upon request, please email your request to: Dr. Elaine Uí Dhonnchadha

2. A demonstration of how a neural model can enhance an existing rule-based system, through the creation of a hybrid rule/neural based model.
3. An evaluation of the contextual correctness of tags in the existing rule-based systems for four languages for the first time.
4. The first comparison of rule, neural, and hybrid systems in five languages.
5. The release of the first open access manually annotated corpus for USAS semantic tagging in Chinese.

All of the resources mentioned in the paper have been released open access with the publication of this paper, or will be available on request<sup>7</sup>.

## 2. USAS Tagset

The USAS tagset (Rayson et al., 2004; Löfberg and Rayson, 2019), which was originally derived from the Longman Lexicon of Contemporary English (McArthur, 1981), contains 21 major discourse fields that expand to 232 category labels. USAS has a hierarchical tagset that has up to three levels of sub-division. An example of the tagset is shown in Table 3. These 232 category labels can be enhanced with affixed symbols to indicate additional linguistic meaning, such as rarity (%@), gender (mf), antonyms (+-), and can be applied to both single words and Multi-Word Expressions (MWE). The semantic tags can be combined with double or triple membership through the use of a slash ("/"). For instance, in Table 2 the token "**Coffee**" is tagged with "*F2/O2*", indicating the token is related to both Drinks (*F2*) and Objects (*O2*), which is correct. Examples of semantically tagged tokens with the full USAS tagset can be seen in table 2, and a full explanation of the USAS tagset can be found in Archer et al. (2002).

In this work, we focus on evaluating on the 232 category labels without any of the affixed symbols nor modelling the MWE. Unlike the rule-based models, the neural models were trained to predict the 232 category labels without the dual or triple tag membership, in order to simplify the training procedure as it reduces the number of category labels<sup>8</sup> and reduces the number of training examples per category.

<sup>7</sup>Code for training the models; <https://github.com/UCREL/experimental-wsd> and the trained models, training dataset, and evaluation datasets; <https://huggingface.co/collections/ucrelnlp/usas-neural-taggers-10>

<sup>8</sup>For dual membership the number of category labels increase by the power of 2 and for triple membership by the power of 3.

As the neural model requires a description of each category label, also known in WSD as the gloss, we used the combination of the title and if one exists description of the tag from Archer et al. (2002) as the gloss.

### 3. Dataset

In this section, we detail the silver-labelled training data that we have created for the neural based models and the evaluation data.

#### 3.1. Training Data

Inspired by Conia et al. (2024) who created a large silver labelled dataset named MOSAICo that was used for WSD, Semantic Role Labelling, Semantic Parsing, and Relation Extraction, we have also created a silver labelled dataset for WSD using the USAS tagset instead of the BabelNet tagset.

In comparison to MOSAICo, where they used pre-trained models that have already been trained on a smaller manually annotated dataset (229,517 + 496,776 annotated words from SemCor and WordNet GlossTag corpora (Vial et al., 2018)), we used the original rule-based English-only C version of the semantic tagger (Rayson et al., 2004) for creating the silver standard training data. The C version of the tagger was used as it is more complex and accurate compared to the newer PyMUSAS version, however we use the PyMUSAS version later in the paper to evaluate on and compare to other methods as the C-version is not open source or widely available<sup>9</sup>.

We also used the higher quality data from the MOSAICo dataset, denoted as *MOSAICo Core*, to create our silver labelled training dataset. This data contains English Wikipedia documents labelled as either "good" or "featured". Conia et al. (2024) demonstrated that by using a higher quality but smaller silver training dataset, they could achieve comparable results to a model trained on a large Wikipedia dataset that has not been filtered by quality.

The Wikipedia data was tokenized, sentence-split, and POS-tagged by the CLAWS tagger and then lemmatized and semantically tagged by the C-version of the semantic tagger. The composition of this corpus can be seen in table 4, where we have also divided the data into training and validation splits. The semantic tagger labels each token with one or more USAS tags derived from all possible senses determined by the lexicons. If a tag contains dual or triple membership, we split the tag up into

<sup>9</sup>A version of the tagger was given to us from the original authors for this work and have given us permission to distribute the tags generated from the tagger under the same license as the dataset.

2 or 3 separate tags respectively. We removed all tags marked as punctuation or containing the unmatched USAS ("Z99"), as these tags do not have any semantic meaning.

As this dataset only contains positive USAS tags per token, whereby all positive senses for the target token  $i$  is represented by  $S_{pi} = s_1, \dots, s_c$  and all USAS tags are represented by  $S$  where  $S_{pi} \subset S$ , to train the neural model we need negative examples per token as well.

We randomly sampled three negative USAS tags per positive tag from three different weighted distributions; 1. the USAS tag distribution from the silver labelled training data split (*Original*), 2. the inverse frequency of the USAS tag distribution (*Inverse*), and 3. the log to the base 2 of the inverse frequency (*Log Inverse*). The reason for the 3 sampling distributions is to keep a balance between; sampling from the original distribution, sampling labels that are under-represented (*Inverse*), and by sampling *Log Inverse* to keep a more even distribution that is more likely to sample from tags that are neither massively over or under representative. This sampling strategy should, to a point, even out the skewed label distribution that is present in USAS labelled data<sup>10</sup>, which is common in most datasets that contain large label sets. This inverse frequency negative sampling strategy was inspired by Blevis and Zettlemoyer (2020) who weighted the word sense labels by their inverse frequency within the model's loss function so that the model performed better on less frequent word senses for WSD.

These negative samples, which can be represented as  $S_{ni} = s_1, s_2, s_3$  for target word  $i$ , are constrained so that no negative USAS tag is a member of the positive USAS tags for that target word  $i$  thus  $S_{ni} \cap S_{pi} = \{\}$ .

This negative sampling strategy is used for both the training and validation split of the silver training data.

#### 3.2. Evaluation Data

All the evaluation data has either been manually tagged or has been manually checked. For English and Finnish, we used the dataset from Löfberg et al. (2003), a tagged corpus of texts from a Finnish coffee website<sup>11</sup>, the English corpus is a machine translated version of the Finnish that was post edited by a native Finnish speaker.

For the Welsh language, we used the entire dataset released in Ezeani et al. (2019) for eval-

<sup>10</sup>The distribution of the USAS tagset within this silver labelled training dataset split can be seen in figure 2 in appendix A.

<sup>11</sup>The website URL: <http://www.kahvilasi.net/>. This website is no longer available.

**Oxygen** {O1.3} , {PUNC} **light** {O4.3} and {Z5} **moisture** {O1.2}  
**Coffee** {F2/O2[i135.2.1]} **pot** {F2/O2[i135.2.2]}  
**Erik** {Z1mf} {Z3c} **Adolf** {Z1mf} {Z3c} **von** {Z1mf} {Z3c} **Willebrand** {Z99}

Table 2: Example of **tokens** semantically annotated with USAS tags. Each token is tagged with one or more *USAS tag groups*, indicated by the curly braces (`{}`). The "PUNC" or sometimes seen as "PUNCT" indicates a punctuation token and is not part of the USAS tagset. The first two lines come from a coffee related article and the last comes from the Wikipedia article on Erik Adolf von Willebrand.

Tag	Short Description
F2	Drinks
O1.2	Substances and materials generally: Liquid
O1.3	Substances and materials generally: Gas
O2	Objects generally
O4.3	Colour and colour patterns
Z1	Personal names
Z3	Other proper names
Z5	Grammatical bin
Z99	Unmatched

Table 3: A selection of USAS tags with a short description.

uation. The dataset consists of 8 extracts from 4 diverse data sources; Kynulliad3<sup>12</sup> (Welsh Assembly proceedings), Meddalwedd<sup>13</sup> (translations of software instructions), Kwici<sup>14</sup> (Welsh Wikipedia articles), and LERBIML<sup>15</sup> (multi-domain spoken corpora). In the original work with manually corrected POS tags, 90% of this dataset was used for training a joint POS and semantic tagger, while 10% was used for evaluation.

For the Irish language, we used the data released by Czerniak and Uí Dhonnchadha (2024) which contains 3 texts;<sup>16</sup> a paragraph from an online news article, and the first 226 and 301 tokens of two Wikipedia articles about the TV show *The Wire* and the author George Orwell respectively. The Irish data also has corrected lemma and POS tags.

The Chinese dataset that we have created is a manually tagged text from the "News Report" genre of the ToRCH2019 corpus (Jialei Li, Mingchen Sun, and Jiajin Xu, 2022), specifically about the 2019 mil-

<sup>12</sup><https://kevindonnely.org.uk/kynulliad3/>

<sup>13</sup><https://techiaith.cymru/corpws/Moses/Meddalwedd/>

<sup>14</sup><https://kevindonnely.org.uk/kwici/>

<sup>15</sup><https://www.lancaster.ac.uk/fass/projects/biml/bimls3corpus.htm>

<sup>16</sup>This dataset has been extended since Czerniak and Uí Dhonnchadha (2024) to 10 texts, but we decided to use the original 3 so that readers could compare it to the results of the original work that used a different evaluation metric. In the future we will use the extended dataset for evaluation.

itary world games in Wuhan, China. This dataset was manually annotated following a three-stage procedure; 1. independent tagging, 2. independent reviews of their own tagging, and 3. reaching consensus between two trained researchers.

Similar pre-processing to the training data was performed on these evaluation datasets; we removed all tag tokens marked as punctuation or containing the unmatched USAS ("Z99"), as these tags do not have any semantic meaning. In addition, any labelled tokens that could not be matched to the USAS tagset were removed. The statistics for these datasets are shown in Table 5<sup>17</sup>. Unlike the training dataset, all labelled tokens in the checked evaluation datasets contain only one USAS tag each.

## 4. Models

In this section, we outline the rule-based semantic taggers used in this experiment, along with the architecture of the neural based model, and we discuss how the neural model is incorporated into the rule-based model that creates the hybrid model.

### 4.1. Rule-Based Models

We use the PyMUSAS framework for our rule-based methods which uses one or more lexicons to assign one or more USAS tags to a token. The framework is very flexible, but it works most effectively when additional tools are used in conjunction (e.g. lemmatisers and POS taggers), as these linguistic features are used by the framework to more accurately disambiguate tokens to their respective USAS tag(s) through the lexicons. If a token has more than one USAS tag after other disambiguation methods have been applied, then the first USAS tag is considered the most likely one and the last the least likely one.

The lexicons come in two types, **single word** and **MWE**, and both are in essence dictionary lookups. Examples of these lexicons are shown in Tables 6 and 7 respectively. The **single word**

<sup>17</sup>The Chinese and Irish dataset for 301 and 6 tokens contain more than one USAS label assigned to them, therefore we have used the first USAS label as the only label for these tokens assuming that label is the most likely.

Split	Documents	Sentences	Tokens	L. Tokens	Labels	Labels per Token
Train	1,083	444,880	6,651,919	5,378,222	11,362,935	2.11
Validation	21	20,000	258,561	204,871	453,300	2.21

Table 4: The number of documents, sentences, tokens, Labelled Tokens (L. Tokens), and labels in the silver labelled training dataset per split. Each labelled token could have more than one label associated it hence the labels per token ratio.

Language	Text Level	Texts	Tokens	L. Tokens	Multi Tag Membership (%)
Chinese	sentence	46	2,312	1,747	1 (0%)
English	sentence	73	3,899	3,468	212 (6.1%)
Finnish	sentence	72	2,439	2,068	254 (12.3%)
Welsh	sentence	611	14,876	12800	1,311 (10.2%)
Irish	paragraph	3	711	618	55 (8.9%)

Table 5: Evaluation dataset statistics. The number of texts, tokens, Labelled Tokens (L. Tokens). Multi Tag Membership is the number of labelled tokens, and the percentage (%), whereby the USAS tag either has dual, triple, or quadruple membership.

lexicon matches a word to a list of USAS tags based on its lemma and POS. The **MWE** lexicon uses a MWE template system (a simplified pattern matching code similar to regular expression) of “{token/lemma}\_{POS}” to match more than one word to a list of semantic tags whereby each word will then be assigned those semantic tags. For example, the MWE template “\*\_ Ocean\_NOUN” would match to the MWE “Pacific Ocean”, and both “Pacific” and “Ocean” would be assigned the same semantic tag of “Z2”. For some languages, the lexicons have been completely manually created (e.g. English and Finnish), whereas for others they have been automatically created by translation and then partially manually checked to varying degrees (e.g. Welsh, Spanish and Italian).

Lemma	POS	USAS tags
coffee-house	NOUN	H1/F1
programming	VERB	Y2 P1

Table 6: Example single word lexicon where the USAS tags assigned to a lemma and POS are whitespace separated.

MWE Template	USAS tags
*_ Ocean_NOUN	Z2
*_ VERB over_ADV	T2- M1 M6

Table 7: Example MWE lexicon where the USAS tags assigned to a MWE template are whitespace separated.

The PyMUSAS framework follows a set of heuristics based on those stated on page 4 column 2 of Piao et al. (2003)<sup>18</sup> which in essence ranks all

<sup>18</sup>For a detailed list of these rules see: <https://ucrel.github.io/pymusas/api/rankers/>

lexicon matches for a token. For example, **MWE** matches are ranked higher than **single word** lexicon matches, and the search for a match can be expanded through dropping the POS tag requirement or lower casing the token etc.

For each language we evaluate on, we start by using the existing resources and methods with PyMUSAS, i.e. each language has its own rule-based tagger along with the resources, as shown in Table 8<sup>19</sup>. The Irish lexicons come from Czerniak and Uí Dhonnchadha (2024)<sup>20</sup>, the English, Chinese and Finnish lexicons come from Piao et al. (2016), and the Welsh lexicons come from Piao et al. (2018)<sup>21</sup>.

## 4.2. Neural Models

We used the WSD Bi-Encoder Model (BEM) from Blevis and Zettlemoyer (2020), where the model is trained to predict the correct sense definition (gloss) for a given ambiguous word given numerous possible sense definitions, where only one of the sense definitions is correct. The architecture of this model is illustrated in Figure 1.

In detail, the task is given a text  $T = t_i, \dots, t_n$  of  $n$  words disambiguate the target word at position  $i$  within the text  $t_i$  to a sense  $s_j$  from the set of all possible senses  $s_j \in S$  within the given sense inventory  $S = s_1, \dots, s_m$ , where the sense inventory is the 232 USAS categories. Each sense in the

lexicon\_entry#contextualrulebasedranker

<sup>19</sup>All spaCy Transformer and large language models can be found at: <https://spacy.io/models>

<sup>20</sup>The Irish tagger from the original paper has been used here. A later improved version will be used in future testing.

<sup>21</sup>All the lexicons apart from Irish are available from: <https://github.com/UCREL/Multilingual-USAS>

Language	Lemmatizer and POS tagger	Single word lexicon	MWE lexicon
Irish	(Uí Dhonnchadha and Van Genabith, 2006)	30,190	6,825
English	spaCy Transformer	54,797	19,042
Chinese	spaCy Transformer	64,541	19,040
Finnish	spaCy Large	46,226	0
Welsh	CyTag (Neale et al., 2018)	143,292	240

Table 8: Overview of the resources used for each language’s rule-based tagger. When spaCy Transformer or Large is listed this is a language specific spaCy model.

sense inventory will have a corresponding gloss text  $G = g_1, \dots, g_m$  describing the sense, each gloss text is represented by a sequence of gloss text tokens  $g_n = gt_1, \dots, gt_p$ .

As shown in Figure 1, the model encodes the text that contains the target word ( $i$ )<sup>22</sup> to disambiguate using the context encoder ( $C$ ) which is a Pre-trained Language Model (PLM), creating the vector  $u = C(T)$  where  $u \in \mathbb{R}^{1 \times d}$  and  $d$  is the size of the hidden dimension of the PLM.

The text from each gloss is encoded using a PLM from the gloss encoder ( $GE$ ), denoted as  $GE_{plm}$ . Each gloss text is represented by the vector  $j_n \in \mathbb{R}^{1 \times d}$  which is the mean representation of the gloss encoder PLM  $j_n = \frac{1}{p} \sum_{i=1}^p GE_{plm}(gt_i)[i]$ , whereby the gloss encoder and context encoder use the same shared PLM<sup>23</sup>. All of the senses are then represented as the concatenation of all senses’ representation denoted by the matrix  $J \in \mathbb{R}^{m \times d}$ .

The score for each target word is the dot product between the encoded target word and the sense representations  $score_i = u \cdot J^T$  where  $score_i \in \mathbb{R}^{1 \times m}$ , the sense with the highest score is the sense we assign to the target word. The score for a single sense can be represented as  $score_{i,n} = u \cdot j_n^T$ .

When training the model, using the silver labelled data described in section 3.1, we use the cross-entropy loss function, as shown in equation 1, using the 3 negative samples per target word  $i$ . For each sample during training, the model has to predict between 4 senses, 3 are negative and 1 is positive, which is represented by the symbol  $v$  in equation 1.

$$\mathcal{L}(t_i) = -\log \frac{\exp(score_{i,v})}{\sum_{q=1}^4 \exp(score_{i,q})} \quad (1)$$

The BEM model was chosen over more accurate and recent models from the WSD literature due to the model’s computational efficiency, while

<sup>22</sup>If the word is made up of more than one sub-word tokens from the PLM sub-word tokenizer, then the average vector of the sub-word tokens is used to represent the word.

<sup>23</sup>In the original BEM model the context and gloss encoders did not share a PLM model, they used independent PLM models. We decided to share the PLM to reduce the number of parameters within the model.

the more recent and accurate models (Barba et al., 2021; Zhang et al., 2022) are all cross-encoders. To use cross-encoders at inference time to disambiguate all words in a text would require running the PLM  $n \times m$  times, whereby  $n$  is the number of words in the given text to disambiguate and  $m$  are the number of senses in the sense inventory compared to  $m + 1$  for the BEM model<sup>24</sup>. The computational efficiency difference between bi-encoders and cross-encoders is highlighted in previous works (Reimers and Gurevych, 2019; Humeau et al., 2019).

### 4.3. Hybrid Model

The major limitation of the rule-based models is lexical coverage, that is, the extent to which words in a text are recognised. This coverage is determined by the lexicon size; in general, the larger the lexicon, the greater its coverage. The benefit of the neural network model is that it can tag any word because the PLM of the context and gloss encoder is capable of embedding any word in context. We integrate the neural network model into the rule-based model, so that when the rule-based model fails to make a prediction due to a word being not defined in the lexicon(s), the neural model can step in as a back-off model.

## 5. Experimental Setup

For the neural network model, we test 4 different variants with different PLMs; 2 English PLMs from Weller et al. (2025), and 2 Multilingual PLMs from Marone et al. (2025), which are outlined in Table 10. These PLMs were chosen as they have been shown to perform best in their size category, and are completely open including the data they were trained on, and the multilingual models have been pre-trained on all the languages we are evaluating on. Only the Multilingual PLMs will be tested on the non-English datasets.

Each of the neural networks are fine-tuned on the English silver labelled training dataset from section 3.1, using the loss function detailed in section

<sup>24</sup>The BEM model can pre-compute the  $m$  sense embeddings in this case disambiguating all words in a given text would only require running the context encoder PLM once.

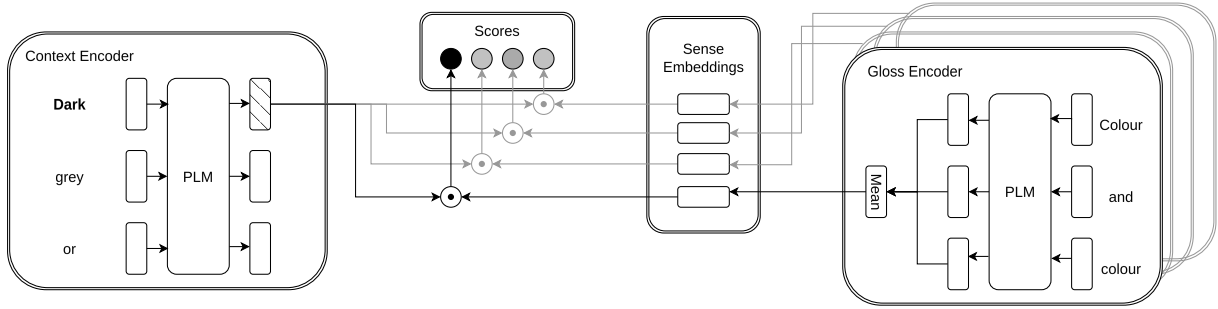


Figure 1: Architecture of the WSD Bi-Encoder Model (BEM) from Blevins and Zettlemoyer (2020).

4.2. The models are trained with early stopping on the validation split of the silver labelled training dataset, whereby we track the accuracy of the model predicting the correct sense rather than the 3 negative examples. We checkpoint the models each time it has trained on 20% of the training samples. All models were trained with a batch size of 64 and a learning rate of  $1e-5$ . Due to computing resource constraints we only managed to train the  $N_{EngS}$  model completely, as in the early stopping constraint stopped the model training. For all other models, we used the best performing checkpoint. The number of epochs the models trained for can be seen in table 9.

Model	Epochs	Validation accuracy (%)
$N_{EngS}$	3	99.39
$N_{EngB}$	3	99.67
$N_{MulS}$	2.2	99.62
$N_{MulB}$	1.36	99.63

Table 9: The neural network model, the number of epochs it trained for on the silver labelled training data, and the validation accuracy on the validation split of the silver labelled training data.

As outlined in section 4.1, we only use one rule-based model per language that will be represented as  $R$  in the results. The hybrid models using the rule-based model per language will be represented as  $H_{neural\ model}$ , e.g. the hybrid model using the small multilingual neural network will be  $H_{MulS}$ .

All models are evaluated on top-n accuracy, whereby top-1 only considers the first tag generated by the model, and top-5 considers the first 5 tags generated by the model. In this evaluation, a single tag can include Multi Tag Membership for both the true tag and the predicted. In these cases, for it to be correct, the true and predicted tags must be identical. For example, if the predicted tag is  $F2/O2$  and the correct tag is  $F2/O1$ , then the prediction is considered wrong. This way of evaluating means that, for the neural network based predictions, they will always fail to predict the Multi Tag Membership true tags, because it has not been trained to predict these types of tags.

When evaluating on the non-English datasets, this will be a cross-lingual evaluation setup, as the neural network based models have only been fine-tuned on English data for this WSD task.

Name	Parameters (M)
$N_{EngS}$	16.8
$N_{EngB}$	68.1
$N_{MulS}$	140.5
$N_{MulB}$	306.9

Table 10: The name of the neural network models where English Small (EngS) and English Base (EngB) use the base model *Ettin-Enc-17m* and *Ettin-Enc-68m* respectively from Weller et al. (2025) and Multilingual Small (MulS) and Multilingual Base (MulB) use the base model *MMBERT small* and *MMBERT base* respectively from Marone et al. (2025). The number of parameters in millions (M).

## 6. Results

The top-n accuracy results are shown in Table 12<sup>25</sup>, where we can see across all languages and values of  $n$  that the the best model is either neural network model or the hybrid model, demonstrating the benefit of training a neural network model.

There is a large difference in the evaluation data results in table 12 and the validation accuracy results in table 9 on the silver labelled data. This is due to the fact the validation accuracy is performed on an easier task in that it is predicting the most likely sense between 1 positive sense and 3 negative senses, whereas the evaluation task is

<sup>25</sup>The Irish language results are not directly comparable with the rule based semantic tagging results of Czerniak and Uí Dhonnchadha (2024), this is due to; 1. different pre-processing, in this paper, of the evaluation data whereby punctuation is not treated as a labelled token and affixed symbols such as %, @, m, f, +, and - were removed from the USAS tags, and 2. different evaluation metrics have been used, the "Correctness" for all tokens metric from Czerniak and Uí Dhonnchadha (2024) is comparable to top-1 accuracy within this paper.

identifying the correct sense from all of the USAS tag senses, therefore very high validation accuracy is expected.

For  $n = 1$ , all but the Chinese rule-based models performed better than the neural network models, but in all cases the hybrid models performed best. The same findings can be seen for  $n = 5$  except that the English rule-based model performs worse than the neural network models.

For Chinese, it can be seen that, for both  $n = 1$  and  $n = 5$ , the neural network models ( $\mathbf{N}_{\text{MultiB}}$  and  $\mathbf{N}_{\text{MultiB}}$ ) outperformed both the rule-based and hybrid models due to the poor performance of the rule-based model. This is also the case for English when  $n = 5$ . This highlights a drawback of the hybrid model whereby if the recall of the rule-based model is high and the precision is low then the hybrid model will be limited by the rule-based model’s performance. It will be less likely to fall back to the neural network model due to the rule-based model’s high recall. This, we hypothesise is the reason why there is little difference in performance between the rule-based model and the hybrid model for English for both values of  $n$ .

In general, we find that the larger the neural network model the better the performance. However, for English the two base models,  $\mathbf{N}_{\text{EngB}}$  and  $\mathbf{N}_{\text{MultiB}}$ , their performances are comparable, showing that a smaller language specific model can be as performant as a larger multilingual model.

Language	Pre-Training Tokens (B)
Welsh	2.2
Irish	3
Finnish	25.5
Chinese	139.2
English	803.3

Table 11: Number of pre-training tokens in billions (B), in which both multilingual PLMs were pre-trained. Statistics is taken from table 9 of [Marone et al. \(2025\)](#).

We can see that the performance of the multilingual neural network models is significantly higher for Chinese and English. We assume this is, in part, due to the huge amount of Chinese and English data available for pre-training the model, as shown in Table 11. The implication of this result is that it is worth investigating: a) using a more language specific PLM to enhance the performance of the tagger, b) fine-tuning the PLM to the target specific languages.

## 7. Conclusion and Future Work

In this work, we have created the first neural network English and multilingual semantic taggers that

have been trained specifically for the USAS tagset without any manually annotated data using only English silver labelled data created from a English rule-based tagger. We have also demonstrated how a hybrid neural/rule-based method can be created.

Through the most extensive evaluation so far, with respect to number of languages, we have shown that either the hybrid or neural network based models across all languages outperform the existing rule-based methods, showing the advantages of these new neural network or hybrid based models for USAS tagging. In future work, we will evaluate the hybrid model further to disentangle the performance gains and losses from the neural network model and rule-based model respectively. We will also evaluate the neural network and rule-based models on tokens that are defined in the rule-based model’s lexicon(s) so that we can better isolate the disambiguate performance of the models, rather than evaluating both coverage and disambiguate, in addition this evaluation will inform us of whether neural network models can go beyond the performance of the rule-based model that it learnt from.

Our results show that even though we have not fine-tuned on any language other than English the neural network models have performed well for Chinese, which we believe is due to the amount of Chinese in the pre-training data that PLM was pre-trained on (see table 11), which can also explain the reason why they performed worse for the low resourced languages (Finnish, Irish, and Welsh). Future work could explore the affect of using a language specific PLM across more languages, as we found in this work that a smaller English specific PLM ( $\mathbf{N}_{\text{EngB}}$ ) performed similarly to a larger multilingual PLM ( $\mathbf{N}_{\text{MultiB}}$ ).

In this work, we have released the first open manually annotated corpus for USAS semantic tagging in Chinese. Our initial qualitative comparison of the rule-based Chinese model with the manually annotated Chinese corpus revealed that Chinese measure words, dates, and time expressions are semantic categories where the model can be improved. These categories will be monitored and systematically evaluated in future work.

We have also released the first silver labelled dataset in English that has been shown to be useful for training both mono and multilingual neural networks for USAS tagging. Future work could explore extending the silver labelled datasets into other languages using the same methodology of using existing rule-based USAS tagger, which will allow exploration of fine-tuning of language specific neural networks models.

All of the resources created or used in this paper have been made available as open resources. The

	n=1					n=5				
	Chinese	English	Finnish	Irish	Welsh	Chinese	English	Finnish	Irish	Welsh
<b>R</b>	32.6	<u>72.4</u>	58.4	<u>56.6</u>	70.6	43.6	81.8	64.0	<u>62.1</u>	73.2
<b>N<sub>EngS</sub></b>	-	66.4	-	-	-	-	87.6	-	-	-
<b>N<sub>EngB</sub></b>	-	70.1	-	-	-	-	<u>90.0</u>	-	-	-
<b>H<sub>EngS</sub></b>	-	<b>72.5</b>	-	-	-	-	<u>81.9</u>	-	-	-
<b>H<sub>EngB</sub></b>	-	<b>72.5</b>	-	-	-	-	82.0	-	-	-
<b>N<sub>MuIS</sub></b>	<u>42.2</u>	66.0	15.8	28.5	21.7	<u>66.3</u>	88.9	32.8	47.6	40.8
<b>N<sub>MuIB</sub></b>	<b>47.9</b>	70.2	25.9	35.6	42.0	<b>70.4</b>	<b>90.1</b>	42.4	51.6	56.4
<b>H<sub>MuIS</sub></b>	39.8	<b>72.5</b>	<u>59.1</u>	<b>57.1</b>	<u>71.3</u>	55.6	82.0	<u>65.8</u>	<b>63.3</b>	<u>75.5</u>
<b>H<sub>MuIB</sub></b>	39.8	<b>72.5</b>	<b>60.3</b>	<b>57.1</b>	<b>72.4</b>	56.3	82.0	<b>67.3</b>	<b>63.3</b>	<b>75.9</b>

Table 12: Top  $n$  accuracy results for all models and languages. The best performing results per language per value of  $n$  is denoted in **bold** and the second best is underlined. The results are divided by a horizontal line by rule-based method (**R**), then English only neural and hybrid models, and then multilingual neural and hybrid models. The neural and hybrid models are divided within the English and multilingual by the dashed horizontal line. The dash (-) denotes a result that is not applicable, which in all cases is using a English neural model for non-English texts.

neural network and hybrid models are also available to use within the PyMUSAS framework<sup>26</sup>, which is an easy to use Python semantic tagging framework that already contains the rule based models.

## 8. Limitations

This work acknowledges that creating the silver labelled training data that we used to train the neural network based models does require an existing rule-based model. The effect of the accuracy of such rule-based model on the performance of the trained neural network model is currently unknown and could be explored in the future as this is likely to be a limitation for future work which may create silver labelled datasets for low resource languages with potentially less performant rule-based models. However, given our successful cross lingual experiments based on English, this may be unnecessary.

With respect to future work of low resource languages, such as Irish and Welsh, we understand that the amount of data to create a silver labelled corpus would be a limitation as compared to a higher resourced language where the amount of text available for these languages is much less.

## 9. Acknowledgements

This research was partially funded by the 4D Picture Project<sup>27</sup>. Andrew Moore was also funded by the University Centre for Computer Corpus Research on Language (UCREL) at Lancaster University, UK<sup>28</sup>. The 4D Picture project research leading to these results has received funding from the EU

<sup>26</sup><https://ucrel.github.io/pymusas/>

<sup>27</sup><https://4dpicture.eu/>

<sup>28</sup><https://ucrel.lancs.ac.uk/>

research and innovation programme HORIZON Europe 2021 under grant agreement 101057332 and by the Innovate UK Horizon Europe Guarantee Programme, UKRI Reference Number 10041120.

We are grateful to the UCREL research group at Lancaster University for providing us access to the Hex computer cluster (Vidler and Rayson, 2024) where the experiments were processed.

## 10. Bibliographical References

Dawn Archer, Andrew Wilson, and Paul Rayson. 2002. [Introduction to the USAS category system](#).

Edoardo Barba, Tommaso Pasini, and Roberto Navigli. 2021. [ESC: Redesigning WSD with extractive sense comprehension](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4661–4672, Online. Association for Computational Linguistics.

Pierpaolo Basile, Lucia Siciliani, Elio Musacchio, and Giovanni Semeraro. 2025. [Exploring the word sense disambiguation capabilities of large language models](#). *arXiv preprint arXiv:2503.08662*.

Terra Blevins and Luke Zettlemoyer. 2020. [Moving down the long tail of word sense disambiguation with gloss informed bi-encoders](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1006–1017, Online. Association for Computational Linguistics.

- Francis Bond and Ryan Foster. 2013. [Linking and extending an open multilingual Wordnet](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1352–1362, Sofia, Bulgaria. Association for Computational Linguistics.
- Ting-Yun Chang, Ta-Chung Chi, Shang-Chi Tsai, and Yun-Nung (Vivian) Chen. 2018. [xSense: Learning sense-separated sparse representations and textual definitions for explainable word sense networks](#). *ArXiv*, abs/1809.03348.
- Simone Conia, Edoardo Barba, Abelardo Carlos Martinez Lorenzo, Pere-Lluís Huguet Cabot, Riccardo Orlando, Luigi Procopio, and Roberto Navigli. 2024. [MOSAICo: a multilingual open-text semantically annotated interlinked corpus](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7990–8004, Mexico City, Mexico. Association for Computational Linguistics.
- Tim Czerniak and Elaine Uí Dhonnchadha. 2024. [Towards semantic tagging for Irish](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16643–16652, Torino, Italia. ELRA and ICCL.
- Ignatius Ezeani, Scott Piao, Steven Neale, Paul Rayson, and Dawn Knight. 2019. [Leveraging pre-trained embeddings for Welsh taggers](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (Repl4NLP-2019)*, pages 270–280, Florence, Italy. Association for Computational Linguistics.
- Artyom Gadetsky, Ilya Yakubovskiy, and Dmitry Vetrov. 2018. [Conditional generators of words definitions](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 266–271, Melbourne, Australia. Association for Computational Linguistics.
- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2019. [Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring](#). In *International Conference on Learning Representations*.
- Helen Langone, Benjamin R. Haskell, and George A. Miller. 2004. [Annotating WordNet](#). In *Proceedings of the Workshop Frontiers in Corpus Annotation at HLT-NAACL 2004*, pages 63–69, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Laura Löfberg, Dawn Archer, Scott Piao, Paul Rayson, Tony McEnery, Krista Varantola, and Jukka-Pekka Juntunen. 2003. [Porting an English semantic tagger to the Finnish language](#). In *Proceedings of the Corpus Linguistics 2003 conference.*, pages 457–464.
- Laura Löfberg and Paul Rayson. 2019. *Developing multilingual automatic semantic annotation systems*. Cambridge University Press.
- Marc Marone, Orion Weller, William Fleshman, Eugene Yang, Dawn Lawrie, and Benjamin Van Durme. 2025. [mmBERT: A modern multilingual encoder with annealed language learning](#). *arXiv preprint arXiv:2509.06888*.
- Marco Maru, Simone Conia, Michele Bevilacqua, and Roberto Navigli. 2022. [Nibbling at the hard core of Word Sense Disambiguation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4724–4737, Dublin, Ireland. Association for Computational Linguistics.
- Tom McArthur. 1981. *Longman Lexicon of Contemporary English*. Longman London.
- Domenico Meconi, Simone Stirpe, Federico Martelli, Leonardo Lavallo, and Roberto Navigli. 2025. [Do large language models understand word senses?](#) In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 33885–33904, Suzhou, China. Association for Computational Linguistics.
- George A. Miller. 1995. [WordNet: a lexical database for English](#). *Commun. ACM*, 38(11):39–41.
- George A. Miller, Martin Chodorow, Shari Landes, Claudia Leacock, and Robert G. Thomas. 1994. [Using a semantic concordance for sense identification](#). In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Roberto Navigli, Michele Bevilacqua, Simone Conia, Dario Montagnini, and Francesco Cecconi. 2021. [Ten Years of BabelNet: A Survey](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4559–4567. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- Steven Neale, Kevin Donnelly, Gareth Watkins, and Dawn Knight. 2018. [Leveraging lexical resources and constraint grammar for rule-based part-of-speech tagging in Welsh](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC*

- 2018), Miyazaki, Japan. European Language Resources Association (ELRA).
- Tommaso Pasini, Alessandro Raganato, and Roberto Navigli. 2021. [Xl-wsd: An extra-large and cross-lingual evaluation framework for word sense disambiguation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13648–13656.
- Scott Piao, Francesca Bianchi, Carmen Dayrell, Angela D’Egidio, and Paul Rayson. 2015. [Development of the multilingual semantic annotation system](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1268–1274, Denver, Colorado. Association for Computational Linguistics.
- Scott Piao, Paul Rayson, Dawn Archer, Francesca Bianchi, Carmen Dayrell, Mahmoud El-Haj, Ricardo-María Jiménez, Dawn Knight, Michal Křen, Laura Löfberg, Rao Muhammad Adeel Nawab, Jawad Shafi, Phoey Lee Teh, and Olga Mudraya. 2016. [Lexical coverage evaluation of large-scale multilingual semantic lexicons for twelve languages](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 2614–2619, Portorož, Slovenia. European Language Resources Association (ELRA).
- Scott Piao, Paul Rayson, Dawn Knight, and Gareth Watkins. 2018. [Towards a Welsh semantic annotation system](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Scott S. L. Piao, Paul Rayson, Dawn Archer, Andrew Wilson, and Tony McEnery. 2003. [Extracting multiword expressions with a semantic tagger](#). In *Proceedings of the ACL 2003 Workshop on Multiword Expressions: Analysis, Acquisition and Treatment*, pages 49–56, Sapporo, Japan. Association for Computational Linguistics.
- Alessandro Raganato, Jose Camacho-Collados, and Roberto Navigli. 2017. [Word sense disambiguation: A unified evaluation framework and empirical comparison](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 99–110, Valencia, Spain. Association for Computational Linguistics.
- Paul Rayson, Dawn E Archer, Scott L Piao, and Tony McEnery. 2004. [The UCREL semantic analysis system](#). In *Proceedings of the workshop on beyond named entity recognition semantic labelling for nlp tasks, in association with Irec-04*, pages 7–12. European Language Resources Association.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- E. Uí Dhonnchadha and J. Van Genabith. 2006. [A part-of-speech tagger for Irish using finite-state morphology and constraint grammar disambiguation](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC’06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Loïc Vial, Benjamin Lecouteux, and Didier Schwab. 2018. [UFSAC: Unification of sense annotated corpora and tools](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- John Vidler and Paul Rayson. 2024. UCREL - Hex; a shared, hybrid multiprocessor system. <https://github.com/UCREL/hex>. Accessed: 2024.
- Orion Weller, Kathryn Ricci, Marc Marone, Antoine Chaffin, Dawn Lawrie, and Benjamin Van Durme. 2025. [Seq vs seq: An open suite of paired encoders and decoders](#). *arXiv preprint arXiv:2507.11412*.
- Guobiao Zhang, Wenpeng Lu, Xueping Peng, Shoujin Wang, Baoshuo Kan, and Rui Yu. 2022. [Word sense disambiguation with knowledge-enhanced and local self-attention-based extractive sense comprehension](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4061–4070, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Zhi Zhong and Hwee Tou Ng. 2010. [It makes sense: A wide-coverage word sense disambiguation system for free text](#). In *Proceedings of the ACL 2010 System Demonstrations*, pages 78–83, Uppsala, Sweden. Association for Computational Linguistics.

# 11. Language Resource References

Jialei Li, Mingchen Sun, and Jiajin Xu. 2022. *ToRCH2019 Balanced Corpus of Modern Chinese*. Center for Foreign Languages and Education, Beijing Foreign Studies University.

## Appendix A. Additional Data Details




Figure 2: Probability of a USAS label within the silver labelled training data, training split, for each distribution.