

# Entity Recognition at First Sight: Improving NER with Eye Movement Information

Nora Hollenstein  
ETH Zurich  
noraho@ethz.ch

Ce Zhang  
ETH Zurich  
ce.zhang@inf.ethz.ch

## Abstract

Previous research shows that eye-tracking data contains information about the lexical and syntactic properties of text, which can be used to improve natural language processing models. In this work, we leverage eye movement features from three corpora with recorded gaze information to augment a state-of-the-art neural model for named entity recognition (NER) with gaze embeddings. These corpora were manually annotated with named entity labels. Moreover, we show how gaze features, generalized on word type level, eliminate the need for recorded eye-tracking data at test time. The gaze-augmented models for NER using token-level and type-level features outperform the baselines. We present the benefits of eye-tracking features by evaluating the NER models on both individual datasets as well as in cross-domain settings.

## 1 Introduction

The field of natural language processing includes studies of tasks of different granularity and depths of semantics: from lower level tasks such as tokenization and part-of-speech tagging up to higher level tasks of information extraction such as named entity recognition, relation extraction, and semantic role labeling (Collobert et al., 2011). As NLP systems become increasingly prevalent in society, how to take advantage of information passively collected from human readers, e.g. eye movement signals, is becoming more interesting to researchers. Previous research in this area has shown promising results: Eye-tracking data has been used to improve tasks such as part-of-speech tagging (Barrett et al., 2016), sentiment analysis (Mishra et al., 2017), prediction of multiword expressions (Rohanian et al., 2017), and word embedding evaluation (Søgaard, 2016).

However, most of these studies focus on either relatively lower-level tasks (e.g. part-of-speech tagging and multiword expressions) or relatively

global properties in the text (e.g. sentiment analysis). In this paper, we test a hypothesis on a different level: *Can eye movement signals also help improve higher-level semantic tasks such as extracting information from text?*

The answer to this question is not obvious. On one hand, the quality improvement attributed to eye movement signals on lower-level tasks implies that such signals do contain linguistic information. On the other hand, it is not clear whether these signals can also provide significant improvement for tasks dealing with higher-level semantics. Moreover, even if eye movement patterns contain signals related to higher-level tasks, as implied by a recent psycholinguistic study (Tokunaga et al., 2017), noisy as these signals are, it is not straightforward whether they would help, if not hurt, the quality of the models.

In this paper, we provide the first study of the impact of gaze features to automatic named entity recognition from text. We test the hypothesis that eye-tracking data is beneficial for entity recognition in a state-of-the-art neural named entity tagger augmented with embedding layers of gaze features. Our contributions in the current work can be summarized as follows:

1. First, we manually annotate three eye-tracking corpora with named entity labels to train a neural NER system with gaze features. This collection of corpora facilitates future research in related topics. The annotations are publicly available.
2. Beyond that, we present a neural architecture for NER, which in addition to textual information, incorporates embedding layers to encode eye movement information.
3. Finally, we show how gaze features generalized to word types eliminate the need for recorded eye-tracking data at test time. This

makes the use of eye-tracking data in NLP applications more feasible since recorded eye-tracking data for each token in context is not required anymore at prediction time. Moreover, type-aggregated features appear to be particularly useful for cross-domain systems.

Our hypotheses are evaluated not only on the available eye-tracking corpora, but also on an external benchmark dataset, for which gaze information does not exist.

## 2 Related Work

The benefits of eye movement data for machine learning have been assessed in various domains, including NLP and computer vision. Eye-trackers provide millisecond-accurate records on where humans look when they are reading, and they are becoming cheaper and more easily available by the day (San Agustin et al., 2009; Sewell and Komogortsev, 2010). Although eye-tracking data is still being recorded in controlled experiment environments, this will likely change in the near future. Recent approaches have shown substantial improvements in recording gaze data while reading by using cameras of mobile devices (Gómez-Poveda and Gaudioso, 2016; Papoutsaki et al., 2016). Hence, eye-tracking data will probably be more accessible and available in much larger volumes in due time, which will facilitate the creation of sizable datasets enormously.

Tokunaga et al. (2017) recently analyzed eye-tracking signals during the annotation of named entities to find effective features for NER. Their work proves that humans take into account a broad context to identify named entities, including predicate-argument structure. This further strengthens our intuition to use eye movement information to improve existing NER systems. And going even a step further, it opens the possibility for real-time entity annotation based on the reader’s eye movements.

The benefit of eye movement data is backed up by extensive psycholinguistic studies. For example, when humans read a text they do not focus on every single word. The number of fixations and the fixation duration on a word depends on a number of linguistic factors (Clifton et al., 2007; Demberg and Keller, 2008). First, readers are more likely to fixate on open-class words that are not

predictable from context (Rayner, 1998). Reading patterns are a reliable indicator of syntactical categories (Barrett and Søgaard, 2015a). Second, word frequency and word familiarity influence how long readers look at a word. The frequency effect was first noted by Rayner (1977) and has been reported in various studies since, e.g. Just and Carpenter (1980) and Cop et al. (2017). Moreover, although two words may have the same frequency value, they may differ in familiarity (especially for infrequent words). Effects of word familiarity on fixation time have also been demonstrated in a number of recent studies (Juhász and Rayner, 2003; Williams and Morris, 2004). Additionally, the positive effect of fixation information in various NLP tasks has recently been shown by Barrett et al. (2018), where an attention mechanism is trained on fixation duration.

**State-of-the-art NER** Non-linear neural networks with distributed word representations as input have become increasingly successful for any sequence labeling task in NLP (Huang et al., 2015; Chiu and Nichols, 2016; Ma and Hovy, 2016). The same applies to named entity recognition: State-of-the-art systems are combinations of neural networks such as LSTMs or CNNs and conditional random fields (CRFs) (Strauss et al., 2016). Lample et al. (2016) developed such a neural architecture for NER, which we employ in this work and enhance with eye movement features. Their model successfully combines word-level and character-level embeddings, which we augment with embedding layers for eye-tracking features.

## 3 Eye-tracking corpora

For our experiments, we resort to three eye-tracking data resources: the *Dundee corpus* (Kennedy et al., 2003), the *GECO corpus* (Cop et al., 2017) and the *ZuCo corpus* (Hollenstein et al., 2018). For the purpose of information extraction, it is important that the readers process longer fragments of text, i.e. complete sentences instead of single words, which is the case in all three datasets.

Table 1 shows an overview of the domain and size of these datasets. In total, they comprise 142,441 tokens with gaze information. Table 1 also shows the differences in mean fixation times between the datasets (i.e. fixation duration (the average duration of a single fixation on a word in

	<b>Dundee</b>	<b>GECO</b>	<b>ZuCo</b>	<b>Total</b>
domain(s)	news articles	literature	movie reviews, Wikipedia articles	-
number of sentences	2367	5424	700	8491
mean sentence length	24.75	12.65	22.12	19.84
number of words	58598	68606	15237	142441
unique word types	9131	5283	4408	13937
mean word length	4.29	3.76	4.44	4.16
fixation duration (ms)	202	214	226	214
gaze duration (ms)	237	232	265	244.7

Table 1: Descriptive statistics of the eye-tracking corpora, including domain, size and mean fixation and gaze duration per token.

	<b>Dundee</b>		<b>GECO</b>		<b>ZuCo</b>		<b>Total</b>	
	all	unique	all	unique	all	unique	all	unique
PERSON	732	415	1870	108	657	446	3259	955
ORGANIZATION	475	261	26	12	156	95	657	364
LOCATION	431	177	101	23	366	155	898	1646
<b>total</b>	<b>1638</b>	<b>853</b>	<b>1997</b>	<b>143</b>	<b>1179</b>	<b>696</b>	<b>4814</b>	<b>1646</b>
		52%		7%		59%		34%

Table 2: Number and distribution of named entity annotations in all three eye-tracking corpora.

milliseconds) and gaze duration (the average duration of all fixations on a word)).

**Dundee Corpus** The gaze data of the Dundee corpus (Kennedy et al., 2003) was recorded with a *Dr. Bouis Oculometer Eyetracker*. The English section of this corpus comprises 58,598 tokens in 2,367 sentences. It contains eye movement information of ten native English speakers as they read the same 20 newspaper articles from *The Independent*. The text was presented to the readers on a screen five lines at a time. This data has been widely used in psycholinguistic research to analyze the reading behavior of subjects while reading sentences in context under relatively naturalistic conditions.

**GECO Corpus** The Ghent Eye-Tracking Corpus (Cop et al., 2017) is a more recent dataset, which was created for the analysis of eye movements of monolingual and bilingual subjects during reading. The data was recorded with an *EyeLink 1000* system. The text was presented one paragraph at a time. The subjects read the entire novel *The Mysterious Affair at Styles* by Agatha Christie (1920) containing 68,606 tokens in 5,424 sentences. We use only the monolingual data recorded from the 14 native English speakers for this work to maintain consistency across corpora.

**ZuCo Corpus** The Zurich Cognitive Language Processing Corpus (Hollenstein et al., 2018) is a combined eye-tracking and EEG dataset. The gaze data was also recorded with an *EyeLink 1000* system. The full corpus contains 1,100 English sentences read by 12 adult native speakers. The sentences were presented at the same position on the screen one at a time. For the present work, we only use the eye movement data of the first two reading tasks of this corpus (700 sentences, 15,237 tokens), since these tasks encouraged natural reading. The reading material included sentences from movie reviews from the Stanford Sentiment Treebank (Socher et al., 2013) and the Wikipedia dataset by Culotta et al. (2006).

For the purposes of this work, all datasets were manually annotated with named entity labels for three categories: PERSON, ORGANIZATION and LOCATION. The annotations are available at <https://github.com/DS3Lab/ner-at-first-sight>.

The datasets were annotated by two NLP experts. The IOB tagging scheme was used for the labeling. We followed the ACE Annotation Guidelines (Linguistic Data Consortium, 2005). All conflicts in labelling were resolved by adjudication between both annotators. An inter-

<b>Basic</b>	
$n$ fixations	total number of fixations on a word $w$
fixation probability	the probability that a word $w$ will be fixated
mean fixation duration	mean of all fixation durations for a word $w$
<b>Early</b>	
first fixation duration	duration of the first fixation on a word $w$
first pass duration	sum of all fixation durations during the first pass
<b>Late</b>	
total fixation duration	sum of all fixation durations for a word $w$
$n$ re-fixations	number of times a word $w$ is fixated (after the first fixation)
re-read probability	the probability that a word $w$ will be read more than once
<b>Context</b>	
total regression-from duration	combined duration of the regressions that began at word $w$
$w-2$ fixation probability	fixation probability of the word before the previous word
$w-1$ fixation probability	fixation probability of the previous word
$w+1$ fixation probability	fixation probability of the next word
$w+2$ fixation probability	fixation probability of the word after the next word
$w-2$ fixation duration	fixation duration of the word before the previous word
$w-1$ fixation duration	fixation duration of the previous word
$w+1$ fixation duration	fixation duration of the next word
$w+2$ fixation duration	fixation duration of the word after the next word

Table 3: Gaze features extracted from the Dundee, GECO and ZuCo corpora.

annotator reliability analysis on 10,000 tokens (511 sentences) sampled from all three datasets yielded an agreement of 83.5% on the entity labels ( $\kappa = 0.68$ ).

Table 2 shows the number of annotated entities in each dataset. The distribution of entities between the corpora is highly unbalanced: Dundee and ZuCo, the datasets containing more heterogeneous texts and thus, have a higher ratio of unique entity occurrences, versus GECO, a homogeneous corpus consisting of a single novel, where the named entities are very repetitive.

#### 4 Eye-tracking features

The gaze data of all three corpora was recorded for multiple readers by conducting experiments in a controlled environment using specialized equipment. It is important to consider that, while we extract the same features for all corpora, there are certainly practical aspects that differ across the datasets. The following factors are expected to influence reading: experiment procedures; text presentation; recording hardware, software and quality; sampling rates; initial calibration and filtering, as well as *human factors* such as head movements and lack of attention. Therefore, separate normalization for each dataset should better preserve the

signal within each corpus and for the same reason the type-aggregation was computed on the normalized feature values. This is especially relevant for the type-aggregated features and the cross-corpus experiments described below.

In order to add gaze information to the neural network, we have selected as many features as available from those present in all three corpora. Previous research shows benefits in combining multiple eye-tracking features of different stages of the human reading process (Barrett et al., 2016; Tokunaga et al., 2017).

The features extracted follow closely on Barrett et al. (2016). As described above, psycholinguistic research has shown how fixation duration and probability differ between word classes and syntactic comprehension processes. Thus, the features focus on representing these nuances as broadly as possible, covering the complete reading time of a word at different stages. Table 3 shows the eye movement features incorporated into the experiments. We split the 17 features into 4 distinct groups (analogous to Barrett et al. (2016)), which define the different stages of the reading process:

1. *BASIC* eye-tracking features capture characteristics on word-level, e.g. the number of all

fixations on a word or the probability that a word will be fixated (namely, the number of subjects who fixated the word divided by the total number of subjects).

2. *EARLY* gaze measures capture lexical access and early syntactic processing and are based on the first time a word is fixated.
3. *LATE* measures reflect the late syntactic processing and general disambiguation. These features are significant for words which were fixated more than once.
4. *CONTEXT* features capture the gaze measures of the surrounding tokens. These features consider the fixation probability and duration up to two tokens to the left and right of the current token. Additionally, regressions starting at the current word are also considered to be meaningful for the syntactic processing of full sentences.

The eye movement measurements were averaged over all native-speaking readers of each dataset to obtain more robust estimates. The small size of eye-tracking datasets often limits the potential for training data-intensive algorithms and causes overfitting in benchmark evaluation (Xu et al., 2015). It also leads to sparse samples of gaze measurements. Hence, given the limited number of observations available, we normalize the data by splitting the feature values into quantiles to avoid sparsity issues. The best results were achieved with 24 bins. This normalization is conducted separately for each corpus.

Moreover, special care had to be taken regarding tokenization, since the recorded eye-tracking data considers only whitespace separation. For example, the string *John's* would constitute a single token for eye-tracking feature extraction, but would be split into *John* and *'s* for NER, with the former token holding the label PERSON and the latter no label at all. Our strategy to address this issue was to assign the same values of the gaze features of the originating token to split tokens.

#### 4.1 Type aggregation

Barrett and Søgaard (2015b) showed that type-level aggregation of gaze features results in larger improvements for part-of-speech tagging. Following their line of work, we also conducted exper-

iments with type aggregation for NER. This implies that the eye-tracking feature values were averaged for each word type over all occurrences in the training data. For instance, the sum of the features of all  $n$  occurrences of the token “island” are averaged over the number of occurrences  $n$ . As a result, for each corpus as well as for the aggregated corpora, a lexicon of lower-cased word types with their averaged eye-tracking feature values was compiled. Thus, as input for the network, either the type-level aggregates for each individual corpus can be used or the values from the combined lexicon, which increases the number of word types with known gaze feature values.

The goal of type aggregation is twofold. First, it eliminates the requirement of eye-tracking features when applying the models at test time, since the larger the lexicon, the more tokens in the unseen data receive type-aggregated eye-tracking feature values. For those tokens not in the lexicon, we assign a placeholder for unknown feature values. Second, type-aggregated features can be used on any dataset and show that improvements can be achieved with aggregated gaze data without requiring large quantities of recorded data.

## 5 Model

The experiments in this work were executed using an enhanced version of the system presented by Lample et al. (2016). This hybrid approach is based on bidirectional LSTMs and conditional random fields and relies mainly on two sources of information: character-level and word-level representations.

For the experiments, the originally proposed values for all parameters were maintained. Specifically, the bidirectional LSTMs for character-based embeddings are trained on the corpus at hand with dimensions set to 25. The lookup table for the word embeddings was initialized with the pre-trained GloVe vectors of 100 dimensions (Pennington et al., 2014). The model uses a single layer for the forward and backward LSTMs. All models were trained with a dropout rate at 0.5. Moreover, all digits were replaced with zeros.

The original model<sup>1</sup> was modified to include the gaze features as additional embedding layers to the network. The character-level representation, i.e. the output of a bidirectional LSTM, is concatenated with the word-level representation from

<sup>1</sup><https://github.com/glample/tagger>




Figure 1: Main architecture of the network. Character and word embeddings concatenated with gaze features are given to a bidirectional LSTM.  $l_i$  represents the word  $i$  and its left context,  $r_i$  represents the word  $i$  and its right context. Concatenating these two vectors yields a representation of the word  $i$  in its context,  $c_i$ .

a word lookup table. In the augmented model with eye-tracking information, the embedding for each discrete gaze feature is also concatenated to the input. The dimension of the gaze feature embeddings is equal to the number of quantiles. This architecture is shown in Figure 1. Word length and word frequency are known to correlate and interact with gaze features (Tomanek et al., 2010), which is why we selected a base model that allows us to combine the eye-tracking features with word- and character-level information.

## 6 Results

Our main finding is that our models enhanced with gaze features consistently outperform the baseline. As our baseline, we trained and evaluated the original models with the neural architecture and parameters proposed by Lample et al. (2016) on the GECO, Dundee, and ZuCo corpora and compared it to the models that were enriched with eye-tracking measures. The best improvements on  $F_1$ -score over the baseline models are significant under one-sided t-tests ( $p < 0.05$ ).

All models were trained with 10-fold cross validation (80% training set, 10% development set, 10% test set) and early stopping was performed after 20 epochs of no improvement on the development set to reduce training time.

First, the performance on the individual datasets is tested, together with the performance of one combined dataset consisting of all three corpora (consisting of 142,441 tokens). In addition, we

evaluate the effects of the type-aggregated features using individual type lexicons for each datasets, and combining the three type lexicons of each corpus. Finally, we experiment with cross-corpus scenarios to evaluate the potential of eye-tracking features in NER for domain adaptation. Both settings were also tested on an external corpus without eye-tracking features, namely the CoNLL-2003 dataset (Sang and De Meulder, 2003).

### 6.1 Individual dataset evaluation

First, we analyzed how augmenting the named entity recognition system with eye-tracking features affects the results on the individual datasets. Table 4 shows the improvements achieved by adding all 17 gaze features to the neural architecture, and training models on all three corpora, and on the combined dataset containing *all* sentences from the Dundee, GECO and ZuCo corpora. Noticeably, adding token-level gaze features improves the results on all datasets individually *and* combined, even on the GECO corpus, which yields a high baseline due to the homogeneity of the contained named entities (see Table 2).

Furthermore, Table 4 also presents the results of the NER models making use of the type-aggregated features instead of token-level gaze features. There are two different experiments for these type-level features: Using the features of the word types occurring in the corpus only, or using the aggregated features of all word types in the three corpora (as describe above). As can be seen, the performance of the different gaze fea-

	P	R	F
<b>Dundee</b>			
baseline	79.29	78.56	78.86
with gaze	79.55	79.27	79.35
type individual	<b>81.05</b>	<b>79.37</b>	<b>80.17*</b>
type combined	80.27	79.26	79.67
<b>Geco</b>			
baseline	96.68	97.24	96.95
with gaze	<b>98.08</b>	<b>97.94</b>	<b>98.01*</b>
type individual	97.72	97.42	97.57*
type combined	97.76	97.16	97.46*
<b>ZuCo</b>			
baseline	84.52	81.66	82.92
with gaze	<b>86.19</b>	<b>84.28</b>	<b>85.12*</b>
type individual	84.21	82.61	83.30
type combined	83.26	83.37	83.31
<b>All</b>			
baseline	86.92	86.58	86.72
with gaze	88.72	89.39	89.03*
type combined	<b>89.04</b>	<b>89.52</b>	<b>89.26*</b>

Table 4: Precision (P), recall (R) and  $F_1$ -score (F) for all models trained on individual datasets (best results in bold; \* indicates statistically significant improvements on  $F_1$ -score). *With gaze* are models trained on the original eye-tracking features on token-level, *type individual* are the models trained on type-aggregated gaze features of this corpus only, while *type combined* are the models trained with type-aggregated features computed on all datasets.

ture levels varies between datasets, but both the original token-level features as well as the individual and combined type-level features achieve improvements over the baselines of all datasets.

To sum up, the largest improvement with eye-tracking features is achieved when combining all corpora into one larger dataset, where an additional 4% is gained in  $F_1$ -score by using type-aggregated features. Evidently, a larger mixed-domain dataset benefits from the type aggregation, while the original token-level gaze features achieve the best results on the individual datasets. Moreover, the additional gain when training on all datasets is due to the higher signal-to-noise ratio of type-aggregated features from multiple datasets.

**Evaluation on CoNLL-2003** Going on step further, we evaluate the type-aggregated gaze features on an external corpus with no eye movement information available. The CoNLL-2003 corpus (Sang and De Meulder, 2003) has been

CoNLL-2003	P	R	F
baseline	93.89	94.16	94.03
type combined	<b>94.38</b>	<b>94.32</b>	<b>94.35*</b>

Table 5: Precision (P), recall (R) and  $F_1$ -score (F) for using type-aggregated gaze features on the CoNLL-2003 dataset (\* marks statistically significant improvement).

widely used as a benchmark dataset for NER in different shared tasks. The English part of this corpus consists of Reuters news stories and contains 302,811 tokens in 22,137 sentences. We use this dataset as an additional corpus without gaze information. Only the type-aggregated features (based on the combined eye-tracking corpora) are added to each word. Merely 76% of the tokens in the CoNLL-2003 corpus also appear in the eye-tracking corpora described above and thus receive type-aggregated feature values. The rest of the tokens without aggregated gaze information available receive a placeholder for the unknown feature values.

Note that to avoid overfitting we do not train on the official train/test split of the CoNLL-2003 dataset, but perform 10-fold cross validation. Applying the same experiment setting, we train the augmented NER model with gaze features on the CoNLL-2003 data and compare it to a baseline model without any eye-tracking features. We achieve a minor, but nonetheless significant improvement (shown in Table 5), which strongly supports the generalizability effect of the type-aggregated features on unseen data.

## 6.2 Cross-dataset evaluation

In a second evaluation scenario, we test the potential of eye-tracking features for NER across corpora. The goal is to leverage eye-tracking features for domain adaptation. To show the robustness of our approach across domains, we train the models with token-level and type-level features on 100% of corpus A and a development set of 20% of corpus B and test on the remaining 80% of the corpus B, alternating only the development and the test set for each fold.

Table 6 shows the results of this cross-corpus evaluation. The impact of the eye-tracking features varies between the different combinations of datasets. However, the inclusion of eye-tracking features improves the results for all combinations, except for the models trained on the ZuCo corpus

		Dundee			GECO			ZuCo		
		P	R	F	P	R	F	P	R	F
<b>Dundee</b>	baseline				74.20	70.71	72.40	75.36	75.62	75.44
	token				75.68	71.54	73.55*	<b>78.85</b>	74.51	77.02
	type				<b>76.44</b>	<b>77.09</b>	<b>76.75*</b>	78.33	<b>76.49</b>	<b>77.35</b>
<b>GECO</b>	baseline	58.91	34.91	43.80				68.88	42.49	52.38
	token	<b>59.61</b>	35.62	<b>44.53</b>				<b>69.18</b>	<b>44.22</b>	<b>53.81</b>
	type	58.39	<b>35.99</b>	44.44				67.69	42.36	52.01
<b>ZuCo</b>	baseline	65.85	<b>54.01</b>	59.34	83.00	<b>78.11</b>	<b>80.48</b>			
	token	<b>72.62</b>	50.76	59.70	82.92	75.35	78.91			
	type	69.21	53.05	<b>59.95</b>	<b>83.68</b>	74.57	78.85			

Table 6: Cross-corpus results: Precision (P), recall (R) and  $F_1$ -score (F) for all models trained on one dataset and tested on another (rows = training dataset; columns = test dataset; best results in bold; \* indicates statistically significant improvements). The baseline models are trained without eye-tracking features, *token* models on the original eye-tracking features, and *type* are the models trained with type-aggregated features computed on all datasets.

and tested on the GECO corpus. Presumably, this is due to the combination of the small training data size of the ZuCo corpus and the homogeneity of the named entities in the GECO corpus.

CoNLL-2003	P	R	F
baseline	72.80	56.97	63.92
type combined	<b>74.56</b>	<b>60.20</b>	<b>66.61*</b>

Table 7: Precision (P), recall (R) and  $F_1$ -score (F) for using type-aggregated gaze features trained on all three eye-tracking datasets and tested on the CoNLL-2003 dataset (\* marks statistically significant improvement).

**Evaluation on CoNLL-2003** Analogous to the individual dataset evaluation, we also test the potential of eye-tracking features in a cross-dataset scenario on an external benchmark dataset. Again, we use the CoNLL-2003 corpus for this purpose. We train a model on the Dundee, GECO and ZuCo corpora using type-aggregated eye-tracking features and test this model on the CoNLL-2003 data. Table 7 shows that compared to a baseline without gaze features, the results improve by 3%  $F_1$ -score. These results underpin our hypothesis of the possibility of generalizing eye-tracking features on word type level, such that no recorded gaze data is required at test time.

## 7 Discussion

The models evaluated in the previous section show that eye-tracking data contain valuable semantic information that can be leveraged effectively by NER systems. While the individual datasets are

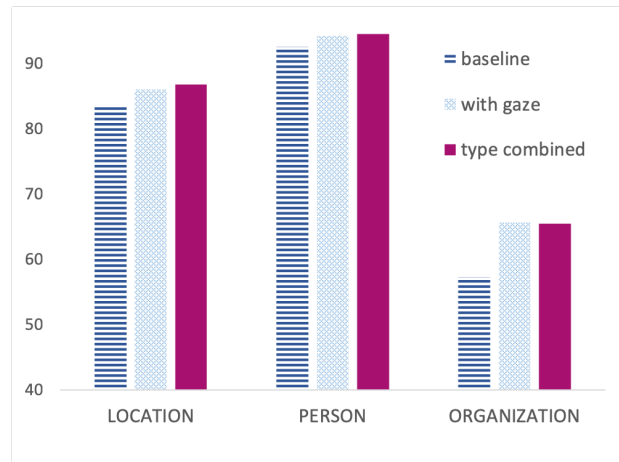


Figure 2: Results per class for the models trained on all gaze datasets combined.

still limited in size, the largest improvement is observed in the models making use of *all* the available data.

At a closer look, the model leveraging gaze data yield a considerably higher increase in recall when comparing to the baselines. In addition, a class-wise analysis shows that the entity type benefiting the most from the gaze features over all models is ORGANIZATION, which is the most difficult class to predict. Figure 2 illustrates this with the results per class of the models trained on all three gaze corpora jointly.

In the individual dataset evaluation setting, the combined type-level feature aggregation from all datasets does not yield the best results, since each sentence in these corpora already has accurate eye-tracking features on token-level. Thus, it is under-



standable that in this scenario the original gaze features and the gaze features aggregated only on the individual datasets result in better models. However, when evaluating the NER models in a cross-corpus scenario, the type-aggregated features lead to significant improvements.

Type aggregation evidently reduces the fine-grained nuances contained in eye-tracking information and eliminates the possibility of disambiguation between homographic tokens. Nevertheless, this type of disambiguation is not crucial for named entities, which mainly consist of proper nouns and the same entities tend to appear in the same context. Especially noteworthy is the gain in the models tested on the CoNLL-2003 benchmark corpus, which shows that aggregated eye-tracking features from other datasets can be applied to any unseen sentence and show improvements, even though more than 20% of the tokens have unknown gaze feature values. While the high number of unknown values is certainly a limitation of our approach, it shows at once the possibility of not requiring original gaze features at prediction time. Thus, the trained NER models can be applied robustly on unseen data.

## 8 Conclusion

We presented the first study of augmenting a NER system with eye-tracking information. Our results highlight the benefits of leveraging cognitive cues such as eye movements to improve entity recognition models. The manually annotated named entity labels for the three eye-tracking corpora are freely available. We augmented a neural NER architecture with gaze features. Experiments were performed using a wide range of features relevant to the human reading process and the results show significant improvements over the baseline for all corpora individually.

In addition, the type-aggregated gaze features are effective in cross-domain settings, even on an external benchmark corpus. The results of these type-aggregated features are a step towards leveraging eye-tracking data for information extraction at training time, without requiring real-time recorded eye-tracking data at prediction time.

## References

- Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei, and Anders Søgaard. 2018. Sequence classification with human attention. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 302–312.
- Maria Barrett, Joachim Bingel, Frank Keller, and Anders Søgaard. 2016. Weakly supervised part-of-speech tagging using eye-tracking data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, volume 2, pages 579–584.
- Maria Barrett and Anders Søgaard. 2015a. Reading behavior predicts syntactic categories. In *Proceedings of the 19th Conference on Computational Natural Language Learning*, pages 345–349.
- Maria Barrett and Anders Søgaard. 2015b. Using reading behavior to predict grammatical functions. In *Proceedings of the Sixth Workshop on Cognitive Aspects of Computational Language Learning*, pages 1–5.
- Jason PC Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional LSTM-CNNs. *Transactions of the Association for Computational Linguistics*, 4:357–370.
- Agatha Christie. 1920. *The Mysterious Affair at Styles*. Retrieved from Project Gutenberg, [www.gutenberg.org](http://www.gutenberg.org).
- Charles Clifton, Adrian Staub, and Keith Rayner. 2007. Eye movements in reading words and sentences. In *Eye Movements*, pages 341–371. Elsevier.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*, 12(Aug):2493–2537.
- Uschi Cop, Nicolas Dirix, Denis Drieghe, and Wouter Duyck. 2017. Presenting GECO: An eyetracking corpus of monolingual and bilingual sentence reading. *Behavior research methods*, 49(2):602–615.
- Aron Culotta, Andrew McCallum, and Jonathan Betz. 2006. Integrating probabilistic extraction models and data mining to discover relations and patterns in text. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, pages 296–303.
- Vera Demberg and Frank Keller. 2008. Data from eye-tracking corpora as evidence for theories of syntactic processing complexity. *Cognition*, 109(2):193–210.
- Jose Gómez-Poveda and Elena Gaudioso. 2016. Evaluation of temporal stability of eye tracking algorithms using webcams. *Expert Systems with Applications*, 64:69–83.
- Nora Hollenstein, Jonathan Rotsztein, Marius Troendle, Andreas Pedroni, Ce Zhang, and Nicolas Langer. 2018. ZuCo, a simultaneous EEG and eye-tracking resource for natural sentence reading. *Scientific Data*.

- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Barbara J Juhasz and Keith Rayner. 2003. Investigating the effects of a set of intercorrelated variables on eye fixation durations in reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6):1312.
- Marcel A Just and Patricia A Carpenter. 1980. A theory of reading: From eye fixations to comprehension. *Psychological review*, 87(4):329.
- Alan Kennedy, Robin Hill, and Joël Pynte. 2003. The Dundee corpus. In *Proceedings of the 12th European Conference on Eye Movement*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270.
- Linguistic Data Consortium. 2005. ACE (Automatic Content Extraction) English annotation guidelines for entities. *Version*, 5(6):2005–08.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 1:1064–1074.
- Abhijit Mishra, Diptesh Kanojia, Seema Nagar, Kuntal Dey, and Pushpak Bhattacharyya. 2017. Leveraging cognitive features for sentiment analysis. *Proceedings of The 20th Conference on Computational Natural Language Learning*, pages 156–166.
- Alexandra Papoutsaki, Patsorn Sangkloy, James Laskey, Nediya Daskalova, Jeff Huang, and James Hays. 2016. WebGazer: Scalable webcam eye tracking using user interactions. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence-IJCAI 2016*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. **Glove: Global vectors for word representation**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543.
- Keith Rayner. 1977. Visual attention in reading: Eye movements reflect cognitive processes. *Memory & Cognition*, 5(4):443–448.
- Keith Rayner. 1998. Eye movements in reading and information processing: 20 years of research. *Psychological bulletin*, 124(3):372.
- Omid Rohanian, Shiva Taslimipoor, Victoria Yaneva, and Le An Ha. 2017. Using gaze data to predict multiword expressions. In *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pages 601–609.
- Javier San Agustín, Henrik Skovsgaard, John Paulin Hansen, and Dan Witzner Hansen. 2009. Low-cost gaze interaction: ready to deliver the promises. In *CHI'09 Extended Abstracts on Human Factors in Computing Systems*, pages 4453–4458.
- Erik F Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the 7th Conference on Natural Language Learning*, volume 4, pages 142–147.
- Weston Sewell and Oleg Komogortsev. 2010. Real-time eye gaze tracking with an unmodified commodity webcam employing a neural network. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems*, pages 3739–3744.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Anders Søgaard. 2016. Evaluating word embeddings with fMRI and eye-tracking. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 116–121.
- Benjamin Strauss, Bethany Toma, Alan Ritter, Marie-Catherine de Marneffe, and Wei Xu. 2016. Results of the wnut16 named entity recognition shared task. In *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*, pages 138–144.
- Takenobu Tokunaga, Hitoshi Nishikawa, and Tomoya Iwakura. 2017. An eye-tracking study of named entity annotation. *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pages 758–764.
- Katrin Tomanek, Udo Hahn, Steffen Lohmann, and Jürgen Ziegler. 2010. A cognitive cost model of annotations based on eye-tracking data. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1158–1167.
- Rihana Williams and Robin Morris. 2004. Eye movements, word familiarity, and vocabulary acquisition. *European Journal of Cognitive Psychology*, 16(1-2):312–339.
- Pingmei Xu, Krista A Ehinger, Yinda Zhang, Adam Finkelstein, Sanjeev R Kulkarni, and Jianxiong Xiao. 2015. Turkergaze: Crowdsourcing saliency with webcam based eye tracking. *arXiv preprint arXiv:1504.06755*.