

NLP in Human Rights Research: Extracting Knowledge Graphs About Police and Army Units and Their Commanders

Daniel Bauer^{‡*}, Tom Longley^{†*}, Yuen Ma^{‡*}, Tony Wilson^{†*}

[‡]Computer Science Department
Columbia University
{db2711, ym2745}@columbia.edu

[†] Security Force Monitor
Human Rights Institute
Columbia Law School
{tom,tony}@securityforcemonitor.com

* All authors contributed equally.

Abstract

In this paper we explore the use of an NLP system to assist the work of Security Force Monitor (SFM). SFM creates data about the organizational structure, command personnel and operations of police, army and other security forces, which assists human rights researchers, journalists and litigators in their work to help identify and bring to account specific units and personnel alleged to have committed abuses of human rights and international criminal law. This paper presents an NLP system that extracts from English language news reports the names of security force units and the biographical details of their personnel, and infers the formal relationship between them. Published alongside this paper are the system’s code and training dataset. We find that the experimental NLP system performs the task at a fair to good level. Its performance is sufficient to justify further development into a live workflow that will give insight into whether its performance translates into savings in time and resource that would make it an effective technical intervention.

Keywords: Ethics and Legal Issues; Information Extraction, Information Retrieval; Knowledge Discovery, Representation; Named Entity Recognition; Tools, Systems, Applications

1. Introduction

Human rights organizations around the world gather large amounts of information for the purposes of promoting and protecting human rights. The promise offered by automated information extraction and processing technologies is of making these rivers of information easier to comprehend and take action on. This promise is much touted; it also feels like such capacities might be more accessible to everyone in a world where software can drive cars or defeat a 9 dan ranked Go master. What, however, does this promise mean in practice for the basic daily work of human rights researchers, rather than their counterparts in commercial, scientific and industrial domains? In this paper we try to provide some insight into this question by reporting the initial outcomes of a multi-disciplinary collaboration to explore the value of Natural Language Processing (NLP) methods as components of information extraction systems used to gather detailed data about state security and defense forces implicated in human rights abuses.

Security Force Monitor¹ (SFM) (Wilson, 2017) is a human rights research project that compiles and analyzes public information to create detailed data on the or-

ganizational structure, command personnel and operations of police, army and other security forces. They provide this data to other human rights researchers, investigative journalists and litigators to help them identify and bring to account specific units and command personnel alleged to have committed abuses of human rights and international criminal law. SFM’s research has been used in the investigation of drug-related killings by police in the Philippines (Security Force Monitor, 2019), allegations of war crimes committed by the army in Mexico (Longley, 2018), and the use of lethal force by the Nigerian military against protesters in Nigeria (Searcey and E., 2018).

SFM’s approach² is to identify salient material (“sources”) through targeted web searches, extract up to 80 specific pieces of information about people, locations and organizations from these sources, and arrange them into a graph-like data structure. These data are transformed into hierarchical organograms or other visualizations of security force structures, showing additional information about personnel (name, rank, role, title), geographic footprint (facilities, bases, camps) and areas of operation. For example, figures 1 and 2 show the command structure and areas of operation of the Western Regional Military Command of the Myan-

¹Security Force Monitor is part of the Columbia Law School Human Rights Institute (<https://securityforcemonitor.org>).

²See “Research Handbook for Security Force Monitor”, <https://help.securityforcemonitor.org/>

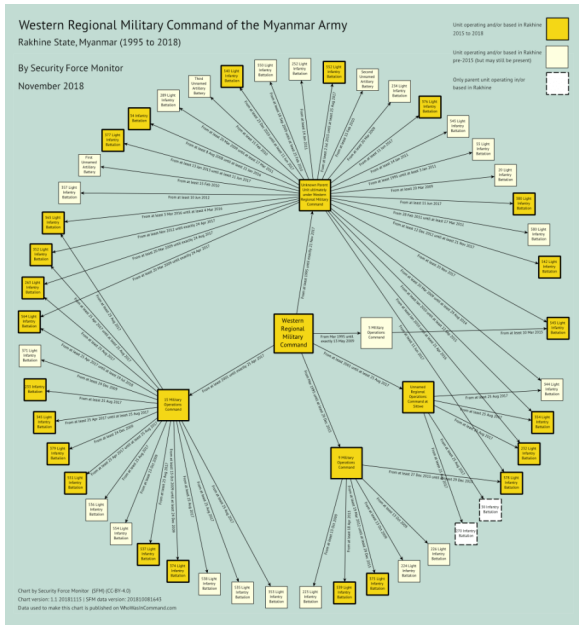


Figure 1: Western Regional Military Command of the Myanmar Army – Rakhine State (1995 to 2018) – a chart by Security Force Monitor (CC-BY-4.0) (Security Force Monitor, 2018)

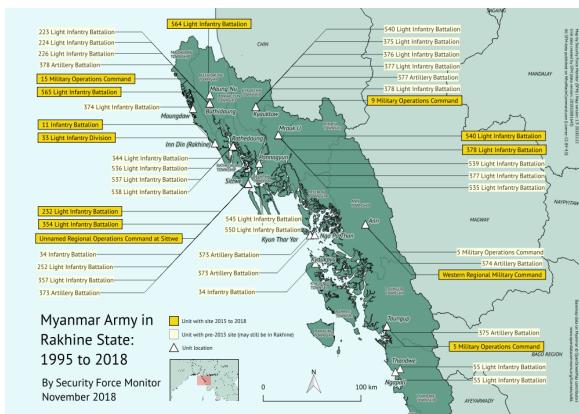


Figure 2: Western Regional Military Command of the Myanmar Army – Rakhine State (1995 to 2018) – a map by Security Force Monitor (CC-BY-4.0) (Security Force Monitor, 2018)

mar Army. SFM performs this extraction work mostly “by hand”. For example, in the course of its research SFM would review the following extract from an article published in Nigeria’s *Vanguard* newspaper in 2012 (Obateru, 2012):

General Officer Commanding 3 Armoured Division of the Nigerian Army, Major General Jack Nwaogbo, has again re-assured Nigerians that the Boko Haram insurgency would soon be contained.

From this, SFM would extract the following pieces of

information and enter them into a database:

- Name of person: “Jack Nwaogbo”
- Rank of person: “Major General”
- Title of person: “General Officer Commanding”
- Organization/unit: “3 Armoured Division”
- Role of person in unit: “Commander”

SFM also applies a number of integrity measures to every data point: they are specifically evidenced by one or more sources, are time-bound (valid from, valid until) and are rated for confidence in their accuracy (from low to high). SFM also extracts and encodes geographical information about the emplacements and operations of specific units. The resulting datasets are made public by SFM, and can be searched through a public website³. At the time of writing, SFM has manually analyzed over 8,000 documents (of which 130 have been annotated for use in this experiment), assembling data on 10,900 specific units, 2,700 command personnel and over 200 alleged human rights violations in 19 countries, going back a decade. Already faced with managing a rich and complicated dataset, SFM faces challenges of scale on numerous fronts: extending coverage to include new force branches and new countries, updating existing data as relevant new material appears online, and working in a number of different languages.

Given the centrality of this type of text analysis to SFM’s research process, NLP would seem to hold potential in automating - partly, or fully - time-consuming tasks like identifying and relating a specific person to a specific unit, and extracting contextual biographical data such as rank and official title. The value to SFM is in picking out common named entities (like Persons and Organizations), and in establishing and extracting the relationships between them in a format that can be quickly appraised for accuracy.

This paper explores this potential in the form of a pilot study and is organized as follows. Section 2 describes the research task and an annotated dataset of 130 news reports about the defense and security forces of Nigeria, which is released with this paper. Section 3 presents initial results of a pilot/baseline system implemented by the authors. We show the system’s results with respect to both named entity recognition and relation extraction tasks. Finally, in Section 4 we look at some of the limitations of the system and the experimental results.

2. Data and task description

SFM identified 130 of the most information-rich text documents from which it has extracted material to create its data on the Nigerian Army and Nigerian Police Force. This document corpus contains 4,711 lines. A single, expert annotator annotated the text of these articles to create a gold standard dataset for use in the

³<https://WhoWasInCommand.com>




Figure 3: An example from our dataset

development of an NLP system. A single expert annotator annotated the 130 documents in the corpus. In the future we plan to develop this dataset by introducing additional annotators and monitoring inter-rater agreement to assure quality. The annotations describe the relationships between persons and the units to which they are posted, which is one of the main information extraction tasks done by hand by SFM. SFM have published the corpus of annotated texts online, along with extensive documentation on the document selection, processing and annotation process.⁴

The SFM expert annotator created the annotations using the Berkeley Rapid Annotation Tool (BRAT) (bra, 2020a), which has a graphical interface that the annotator can use to highlight and connect different information contained in the text. The types of information that can be annotated are described in a BRAT configuration file:

```
[entities]
Person
Organization
Rank
Title
Role
[relations]
is_posted Arg1:Person, Arg2:Organization
has_title Arg1:Person, Arg2:Title
has_role Arg1:Person, Arg2:Role
has_rank Arg1:Person, Arg2:Rank
<OVERLAP> Arg1:Role, Arg2:Rank, <OVL-TYPE>:<ANY>
<OVERLAP> Arg1:Title, Arg2:Role, <OVL-TYPE>:<ANY>
```

The “entities” section above shows us that an annotator can decide a particular word or extract describes a “Person” or an “Organization”, as well as biographical information like the “Title”, “Rank” or “Role” the person may have within that organization. The “relations” section shows how these building blocks can be connected to each other: a “Person” can be “Posted” to an “Organization”, and a “Person” can have a “Role” and a “Rank” during the course of that posting.

Each document in the annotated corpus has two corresponding files: the first stores the raw text, the second the annotations made to that text. A third file contains the document title, date, and other metadata and is not used in the research. The text and annotation files share the same file name (a 36-character UUID) while the suffixes are ‘.txt’ and ‘.ann’, respectively.

Annotations follow the BRAT Standoff format (bra, 2020b). Named entities are identified by text-bounds, which use two numbers to locate the first character and the last character of a name entity. Relations are identified with two arguments, which are the ID’s of the two

⁴https://github.com/security-force-monitor/nlp-starter_dataset

Class	Unique	Mentions
Person	409	531
Rank	103	513
Organization	320	735
Title	151	360
Role	44	167
All Classes	1028	2307

Table 1: Annotations - Unique Named Entities and mentions

Class	Occurrence
has_rank	507
has_title	450
has_role	168
is_posted	391
All Relations	1416

Table 2: Annotations - Relationships between Named Entities

name entities involved in the relation. Relations are directed. The following is an example of the annotations for the sentence mentioned above:

```
T1 Title_Role 35 38 GOC
T2 Title_Role 52 70 Officer Commanding
T3 Organization 71 90 3 Armoured Division
T4 Organization 98 111 Nigerian Army
R1 has_rank Arg1:T1 Arg2:T2
R2 is_posted Arg1:T1 Arg2:T3
R3 has_title Arg1:T1 Arg2:T4
```

The annotations can be also be visualized in the BRAT tool, an example of which is shown in Figure 3. Summaries of the occurrence of Named Entities in the corpus, and the relationships between them, are included in Table 1 and Table 2.

In addition to this dataset of annotated documents, SFM also provided lists of known names of Nigerian military and police units.⁵ These additional datasets were extracted from SFM’s own research into the Nigerian security forces, as well as from lists of named Nigerian military units that the United States government has provided with assistance and training since 2000.

Our research task is to automate the extraction of such annotations from a raw-text input, and gain insight into way in which this capability could replace, substantially augment or otherwise assist researchers in performing this task. The experimental system is not required to reconcile entities across documents or with

⁵Available at https://github.com/security-force-monitor/nlp-starter_dataset/tree/master/other_training_data

Class	True Positives	False Positives	False Negatives	Precision	Recall	F1 Score
Person	87	13	6	0.87	0.94	0.90
Rank	80	14	11	0.85	0.88	0.86
Organization	103	33	31	0.76	0.77	0.76
Title/Role	85	20	23	0.81	0.79	0.80
All Classes	355	80	71	0.82	0.83	0.82

Table 3: NER model evaluation

an external dataset (entity linking). Even though our pilot system is basic, our aim is to quantify the system’s performance and identify the key factors that affect this performance, sufficiently to say whether or not the task can be accomplished in a way that would make it worth implementing as part of SFM’s research workflow. The following section discusses some of these factors.

3. Pilot System

Our pilot system⁶ addresses only two sub-tasks of the full knowledge graph extraction task: Named Entity recognition, and relation extraction. We aim to address entity linking in a future paper.

3.1. Named Entity Recognition (NER)

To extract name entities, we use BiLSTM-CNN-CRF model (Ma and Hovy, 2016) in the traditional inside–outside–beginning (IOB) tagging framework (Sang and Veenstra, 1999). Here is an example of the IOB tagging for the sentence in Figure 3:

```

General Officer Commanding 3 Armoured Division
B-TTL I-TTL I-TTL B-ORG I-ORG I-ORG
of the Nigerian Army , Major General Jack
I-ORG I-ORG I-ORG I-ORG B-RNK I-RNK B-PER
Nwaogbo, has again re-assured Nigerians that the
I-PER O O O O O O O O
Boko Haram insurgency would soon be contained.
O O O O O O O O

```

We first compute input representations for each token by applying a convolutional neural network (CNN) to compute character-level representations. CNNs have been shown to be effective at extracting morphological information (Dos Santos and Zadrozny, 2014; Chiu and Nichols, 2016). The output of the CNN layer is concatenated with pre-trained GloVe word embedding (Pennington et al., 2014) to represent each token.

Next, we compute context representations from the word-level representations by encoding the context using a BiLSTM (Hochreiter and Schmidhuber, 1997; Pascanu et al., 2012). Because we are using the IOB tagging format (Sang and Veenstra, 1999), the label sequence follows certain rules. For instance, I-ORG cannot follow I-PER. Therefore, label sequences are modeled jointly using a conditional random field (CRF) (Lafferty et al., 2001).

⁶Available at <https://github.com/security-force-monitor/sfm-graph-extractor>

Because the data created by SFM for this task is relatively small (Security Force Monitor, 2020), and part of it is needed for testing, we retrained the model with two additional datasets. First, we added part of the CoNLL-2003 dataset (Tjong Kim Sang and De Meulder, 2003). The CoNLL-2003 dataset has two classes which also appears in our dataset: ‘Person’ and ‘Organization’. In addition, as mentioned above, SFM provided a list of known organizations that could be added our dataset. Since it is hard to draw a clear line between the class ‘Title’ and the class ‘Role’, which were specified as distinct in SFM’s knowledge graph, we decided to collapse them into a single class.

The performance of the named entity model is shown in Table 3.

3.2. Relation Extraction

In our pilot system, we experiment with three approaches to relation extraction: nearest person, shortest dependency path, and a neural network based approach. These approaches all share the same underlying idea: starting at each non-Person named entity e , our system tries to identify an entity of type Person in the same sentence that stands in a relation with e , and then tries to predict the type of this relation. We will look at each approach in turn.

3.2.1. Nearest Person

This baseline algorithm is based on the simple idea that a non-Person named entity is often related to a Person named entity nearby. For example, one could say “General Lamidi Adeosun” where “Lamidi Adeosun” has the rank “General”, as in one of the documents in our corpus. The algorithm, therefore, merely relates a non-Person named entity to the Person named entity immediately to the right; if there is no Person named entity to its right, then the algorithm relates it to the nearest Person entity no matter which side the Person entity is on. Although we did not expect performance of this system to be competitive, an immediate advantage of this simple technique is that its decisions are transparent.

3.2.2. Shortest Dependency Path

Instead of using the distances between named entities in raw text, we can take syntactic information into account. This method relates a non-Person named entity to the person named entity to which the dependency path is shortest. This relies largely on how well the de-

Method	True Positives	False Positives	False Negatives	Precision	Recall	F1 score
Nearest Person (Baseline)	993	759	423	0.567	0.701	0.627
Shortest Dep. Path (No constraint)	1083	651	333	0.625	0.765	0.687
Shortest Dep. Path (With constraint)	1180*	559	236*	0.679	0.833*	0.748*
Neural Network (No constraint)	1086	667	330	0.620	0.767	0.685
Neural Network (With constraint)	1103	450*	313	0.710*	0.779	0.743

Table 4: RE algorithms evaluation




Figure 4: The architecture of the neural network: The weight matrix W_1 for path pattern information is shared across all path pattern one-hot vectors, but it is distinct from the weight matrix W_2 for type one-hot vector. The softmax-activated output vector represents the probability of potential Person name entities.

pendency parser performs, so we used a state-of-the-art dependency parser (Nguyen and Verspoor, 2018). Since one named entity could span multiple tokens, we only use the shortest path among the various possible paths between the tokens of two named entities.

We find that constraining the algorithm to only choose between the two Persons that appear immediately to the left and to the right increases performance, at least on our data set. We use this constraint in the final version of our system.

Assuming reasonable performance by the dependency parser, decisions of this heuristic dependency based approach are easy to trace.

3.2.3. Neural Network

In this approach, we use machine learning to predict relations based on dependency paths and named entity types. We use the phrase “path pattern” to refer to the list of edge types along a dependency path. The path patterns are encoded into one-hot vectors. Uncommon path patterns are treated as a single “unknown” category. Since there are multiple possible persons in a sentence, there will be a one-hot vector for each Person. We pass these multiple vectors as input to the network, so that it makes a joint decision over multiple candidate entities of type Person.

In addition, the path length is concatenated to the one-hot vector to compensate for the loss of information when we replace the less frequent path patterns. Multiple one-hot + length vectors of different persons are concatenated together along with a small one-hot vector which encodes the type of the non-person name en-

tity, which makes up the input of the neural network.

The network architecture is shown in Figure 4 and its performance is shown in Table 4. The first layer has a set of shared weights for those one-hot + length vectors and a separate set of weights for the name entity type vector. The second layer is a dense layer whose output is activated by a softmax layer. The softmax layer outputs a vector where the largest element corresponds to the target Person.


The number of persons that the model could process within a sentence is limited to 7. If there are more than 7 persons in a sentence, we set the target to an all-zero vector. If a prediction does not correspond to any person, such as when there are only 3 persons in a sentence while the model predicts the fourth person, then we do not build any relations for the name entity.

To include the constraint mentioned in Section 3.2.2, we made the target a 3-element vector. If the first element is the largest, then the Person on the left side is the predicted person; if the second element is the largest, then the Person on the right side is the predicted Person; if the third element is the largest, then the model predicts that the correct Person should be some Person other than the two nearest Persons. When the model predicts that the Person is not nearby, we select the Person that has shortest dependency path excluding the two nearest Persons.


4. Experimental results

For sentences with relatively few named entities (for example, a sentence has only 1 Person named entity), both Shortest Dependency Path and Neural Network perform very well since there are not many choices to make. When there are more named entities in a sentence, it becomes harder to make a correct choice and that is where the two methods make different predictions.

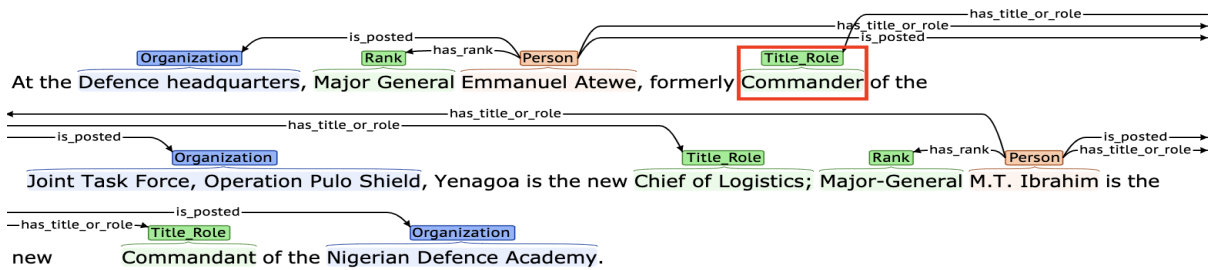
The neural network method looks at the path length, its pattern and the named entity type, so it sees more information than the shortest dependency path method, which only looks at path length. These two additional pieces of information sometimes do help make a better prediction, but they can confuse the neural network as well. In Figure 5, the shortest dependency path method falsely related the name entity “Chief of Logistics” to the person “M. T. Ibrahim”. “Chief of Logistics” is closer to “M. T. Ibrahim” than to “Emmanuel Atewe”



(a) True Annotations



(b) Shortest Dependency Path




(c) Neural Network


Figure 5: Two methods make different mistakes

General Osuji who is the GOC of the 1 Brigade, which is made up of the Nigerian Battalion (NIBATT)36 in the United Nations Mission in Liberia (UNMIL) said, they must work with the rules of engagement of the mission and respect the cultural and religious sensitivity of Liberia people.

(a) True Annotations



(b) Shortest Dependency Path



(c) Neural Network

Figure 6: Neural Network is more conservative

Component	Time (Seconds)	Model Size (Parameters)
NER	1.54	6153100
Dep. Parsing	0.70	8791858
Shortest Dep. Path	0.0039	N/A
Neural Network	0.051	294

Table 5: Average processing time per line and model sizes

in the dependency tree, but the additional information helped the neural network method make the correct choice. But when predicting “Commander”, simply using the length of the dependency path yields the right relation.

The neural network method has fewer false positives, since the shortest dependency path method is forced to build a relation for every non-Person name entity while the neural network has the option not to build a relation. Therefore, the neural network has fewer false positives and more false negatives. This indicates that the neural network is slightly too conservative when deciding whether there exists a correct relation for a given non-Person name entity. One example is shown in Figure 6.

There are some obvious limitations of our algorithms. First, all three algorithms only relate two named entities when they are in the same sentence. When there are relations that cross sentences, our algorithms will not be able to capture them. Second, sometimes, one named entity can be related to multiple other entities. Our algorithms only relate one non-Person name entity to a single Person name entity. Third, we have not implemented the functionality to collect and reconcile the information of Person named entities that exist in different documents.

We also measured the average processing time of different components in the pipeline, as shown in Table 5. The results are based on three independent but identical measurements of the processing time on our entire dataset. The numbers are obtained by averaging the measured time after discarding obvious outliers. The measurements are conducted on a MacBook Pro with a quad-core CPU @ 4.1GHz Max. It is worth noting that the processing time of NER and dependency parsing is dependent on the length of the sentences and the processing time of the shortest dependency parsing algorithm or the neural network is dependent on the amount of named entities in a line.

5. Discussion and conclusion

The results show that an NLP system can perform the research task with a *fair to good* degree of accuracy, albeit with some clear limitations that we must acknowledge.

Our experiment was conducted in an unusual domain in which attention to NLP is in its infancy and there are no pre-existing, domain-specific systems against

which we could compare our results. That said, our experiment would have benefited from the inclusion of a comparison between the system designed and trained to perform the present task and a different, off-the-shelf NER system. Though we saw some improvement to the system’s performance in identifying Organizations by including items from CoNLL-2003 and SFM unit lists, bench-marking against a distinct, untrained system would have provided us with an additional insight.

The size of the material available for training, even if it were not further restricted to the context of the Nigerian security forces, also poses a challenge. The annotated documents were drawn from English-language news sources in Nigeria, and as such are similar to the news-wire material used to train many English language NER models. There are three points worth making, however. First, the approach that we took to annotation could be strengthened in future work with the inclusion of additional annotators and monitoring inter-rater agreement. Second, we did not tune the NER model to better identify Nigerian given and patronymic names, which could have helped boost the system’s performance in detecting Persons. Finally, the names of security force units are often generic and contain numbers (“4 Motorized Regiment”, “25th Division”). It is an open question about whether there is, in the universe of military naming, sufficient material to identify and distinguish these from non-military entities.

The system skews a little towards recall, which is preferable where a subsequent human review is intended. Demonstrating this capability meets the first objective of this collaboration. However, the results alone do not tell us whether its performance is tolerable within SFM’s research workflow (and by extension that of other human rights researchers). Understanding this requires implementing the system in a way that enables SFM to accept, reject or quickly update the proposals it makes, and assessing whether this creates savings in time and resource as compared to doing the work wholly “by hand”. Subsequent work will focus on this next, implementation step.

Although the present paper is mostly technical and focused on practical application, throughout the authors’ collaboration we have ranged across the wider matters of NLP and the challenges of technology implementation within the human rights domain. The potential that NLP represents is set against the sector’s considerable financial and technical capacity constraints and a dearth in transparent examples of successful NLP use within it. Surrounding this are concerns about the human rights implications of NLP methods themselves: the discriminatory potential of the datasets used to train them; the dominance of government and corporate actors in their technical development; and, implementations that infringe human rights directly. In future papers drawn from our research, we aim to assess how these affect the desirability and feasibility of NLP use within the wider non-profit domain.

6. References

- brat, (2020a). *BRAT Rapid Annotation Tool*.
- brat, (2020b). *BRAT Standoff Format*.
- Chiu, J. P. and Nichols, E. (2016). Named entity recognition with bidirectional LSTM-CNNs. *Transactions of the Association for Computational Linguistics*, 4:357–370.
- Dos Santos, C. N. and Zadrozny, B. (2014). Learning character-level representations for part-of-speech tagging. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14*, pages II–1818–II–1826. JMLR.org.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November.
- Lafferty, J. D., McCallum, A., and Pereira, F. C. N. (2001). Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Longley, T. (2018). Prosecutor of the international criminal court receives complaint of crimes against humanity by the mexican army. <https://securityforcemonitor.org/2018/06/11/prosecutor-of-the-international-criminal-court-receives-complaint-of-crimes-against-humanity-by-the-mexican-army/>.
- Ma, X. and Hovy, E. (2016). End-to-end sequence labeling via bi-directional lstm-cnns-crf.
- Nguyen, D. Q. and Verspoor, K. (2018). An improved neural network model for joint POS tagging and dependency parsing. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 81–91, Brussels, Belgium, October. Association for Computational Linguistics.
- Obateru, T. (2012). Boko haram ‘ll soon be contained – goc. <http://www.vanguardngr.com/2012/07/boko-haram-ll-soon-be-contained-goc/>, 7.
- Pascanu, R., Mikolov, T., and Bengio, Y. (2012). On the difficulty of training recurrent neural networks.
- Pennington, J., Socher, R., and Manning, C. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Sang, E. F. T. K. and Veenstra, J. (1999). Representing text chunks. In *Proceedings of the Ninth Conference on European Chapter of the Association for Computational Linguistics, EACL '99*, pages 173–179, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Searcey, D. and E., A. (2018). Nigeria says soldiers who killed marchers were provoked. video shows otherwise. *New York Times*, 12.
- Security Force Monitor. (2018). The structure and operations of the myanmar army in rakhine state: a review of open source evidence. <https://securityforcemonitor.org/2018/11/20/myanmar-army-in-rakhine-state-structure-and-operations/>.
- Security Force Monitor. (2019). Investigating drug related killings in the philippines. <https://securityforcemonitor.org/2019/08/19/investigating-drug-related-killings-in-the-philippines/>.
- Security Force Monitor. (2020). WhoWasInCommand. <https://whowasincommand.com/>.
- Tjong Kim Sang, E. F. and De Meulder, F. (2003). Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Wilson, T. (2017). Why I Started The Security Force Monitor. <https://securityforcemonitor.org/2017/04/11/why-i-started-the-security-force-monitor/>, 4.