NEFUSI: NeuroFuzzy Similarity. Final Report[1]

Jorge Martinez-Gil
December 20, 2022

Abstract

This research work presents the final report for the NEFUSI project. In fact, we present here our research findings on building neurofuzzy models that automatically evaluate semantic textual similarity in an accurate and timely manner. We show that neural networks and fuzzy logic have different features that make them suitable for certain problems but unsuitable for others. Neural networks, on the one hand, are valuable tools for identifying patterns. However, they need to make it easier for people to comply with the decisions. On the other hand, interpretation is possible within fuzzy logic systems, but they cannot automatically derive the rules they use to make those decisions. These constraints served as the primary reason for developing a novel intelligent hybrid system, which combines two approaches to circumvent the individual effects of both limitations simultaneously. The source code for this approach can be found in the following repositories: `https://github.com/jorge-martinez-gil/nefusi` and `https://github.com/jorge-martinez-gil/nefusi2`.

Keywords: Neurofuzzy systems, Semantic Textual Similarity

1. Introduction

In this research, we present a summary of all the work done in the framework of the NEFUSI project. This project was about making new neurofuzzy architectures that let computers figure out how similar two pieces of text are semantically similar without giving up the ability to understand the solutions.

The main issue was that because human language has so many nuances, it is frequently challenging to gauge how much different texts are similar in meaning. Because deep learning (DL) models have been shown to produce the most significant outcomes across the vast majority of fields where artificial intelligence can assist in process automation, we have used them in this context. However, these models have specific areas for improvement, such as the inability to make it simpler for human operators to comprehend the insights they offer and the excessive demand for data to be trained. These two reasons explain why the usage of these technologies is less ubiquitous than its potential

[1] This report is deliverable number 4 of the NEFUSI project `https://www.ngi.eu/funded_solution/nefusi/`

benefits. On the other hand, our working premise was that if these DL models are suitably integrated with other techniques, such as fuzzy logic (FL) more suited to facilitating interpretability, it should be possible to produce an appropriate compromise [5].

Semantic similarity measurement frequently requires the interpretability of models to understand the information that has been derived. The vast majority of the solutions address some form of explainable artificial intelligence. They are primarily focused on technical solutions that are catered to users with an extensive understanding of mathematics. Nevertheless, the last few years have paved the way for the foundation of yet another family of alternative methods. As a result, this research field has seen an explosion of research to make models clearer and easier to comprehend.

In the frame of the NEFUSI project, we examine neurofuzzy models. As a result, the contribution of this work is that we propose a model whereby a fuzzy system operates concurrently with an artificial neural network (ANN) in the hope of ameliorating the shortcomings of each of these two kinds of systems [20]. This solution yields a fuzzy system that works flawlessly with an ANN. By approaching the problem in this manner, our goal is to develop a more effective, reliable, and straightforward model for assessing the semantic similarity among pieces of text.

Furthermore, in the context of this work, we have concentrated on using well-known models such as BERT [8], ELMo [30], or USE [2] with Mamdani inference [17]. Using these models has allowed us to come to some interesting results. Additionally, the various parameters of the fuzzy component have been set up under the seminal work done by Angelov and Buswell [1], which means that this research is founded on that work. It is also possible to investigate the attainable success level by employing Takagi Sugeno-type models [35].

The remaining parts of this work are organized in the following manner: The works relevant to applying neurofuzzy systems to textual processing information are presented in Section 2. Section 3 presents the technical details for using neurofuzzy systems in the NEFUSI project. Section 4 shows the experimental setup and user instructions for running the software. Finally, we remark the most important conclusions from this research.

2. Related works

Regarding the automatic semantic similarity assessment, there is a significant body of research on approaches and resources that can be used to solve the problem [7, 10, 11, 21, 23, 29]. DL in general, and word embeddings [26] in particular, are currently the most popular family of methods because it can produce the best outcomes in various applications [14]. However, its efficacy has yet to be demonstrated in more specialized fields with distinctive characteristics [36]. It is a common belief that models focusing on whole-sentence comprehension instead of the more traditional approach of breaking down each word will produce superior results. The NEFUSI project presents a novel neurofuzzy architecture to address the problem of semantic similarity at the level of words and sentences.

From the viewpoint of the FL, the analytically derived fuzzy systems have one obvious benefit over other types of solutions: the human operator does not need to initially specify the structure of the model before beginning to work with it. However, this benefit is not without its drawbacks, the most significant of which is the computation time required. Exploring the search space containing all possible configurations is usually unfeasible. Therefore, sophisticated algorithms must be utilized to explore this space to find the appropriate model to fit the data [34].

The provision of computational tools by FL can help develop inference algorithms [4]. Mamdani fuzzy inference is something we are looking into because it makes it simpler to build systems managed by a set of rules that are very similar to those of natural language [18]. Due to its application, any rule put into place in Mamdani fuzzy inference systems will always produce a fuzzy set. The rule base of these systems is very similar to natural language, making it relatively simple to understand and very similar to natural language. We will be using the IEC 61131-7 standard because it provides the foundation for a programming fuzzy systems.

3. Contribution

The structure of a fuzzy system and the number of fuzzy rules it employs are typically the defining characteristics of a fuzzy system [21]. In contrast, ANNs, which were developed to solve problems involving pattern recognition, are limited to only learning from the data initially fed into the model as training. If it is necessary to model a trade-off between the model's accuracy and its interpretability, it is also possible to use multi-objective algorithms [6].

The ANN component makes use of transformers, which are models that are suitable for the transformation of abstract representations into one another [33]. The transformer models are based on an encoder-decoder architecture because it allows the encoder to learn a representation of the input and then pass that representation to the decoder component [25].

The interpretability of FL systems can be improved by utilizing IF-THEN rules [18], but acquiring this knowledge takes time and effort. It is usually important to rely on the help of a specialist in the subject because there are so many variables that need to be adjusted. However, neurofuzzy systems combine the benefits of traditional and current methods to address these concerns. When these two approaches (ANN+FL) are combined correctly, a neurofuzzy model is generated [32]. In addition, it is possible to use multi-objective evolutionary techniques to model a trade-off between the model's accuracy and its interpretability [22].

4. User instructions

The NEFUSI project has been divided into four deliverables. The first was a state-of-the-art neurofuzzy system based on a preliminary literature survey. The two central deliverables, 2 and 3, consisted of developing a prototype that works with words and sentences. Furthermore, finally, deliverable 4 is the present report as a summary. We look at the two central deliverables and their associated prototypes in the following.

Please be aware that running the software under the current configuration takes some time to finish. We run the software in an Intel(R) Core(TM) i7-1185G7 @ 3.00 GHz and 1.80 GHz.

4.1. Words

The prototype consisted of a neurofuzzy solution to measure the semantic similarity of words. The source code for this approach can be found under the following URL: `https://github.com/jorge-martinez-gil/nefusi`

4.1.1. Compilation

```
...\nefusi\src>javac -cp combined.jar nefusi/*.java
```

4.1.2. Execution

```
...\nefusi\src>java -cp .;combined.jar nefusi.nefusi
```

4.1.3. Datasets

Using a well-known dataset designed for general purposes will give us an idea of how effectively our method operates. This dataset is known as the MC30 dataset, and it was created by Miller and Charles [27]. It comprises 30-word pairs relevant to everyday life and could be used by anyone.

4.1.4. Results

Table 1 summarizes the results obtained at the word level for the MC30 dataset and the comparison with other existing techniques.

4.2. Sentences

The second prototype consisted of a neurofuzzy solution to determine the semantic similarity of sentences and paragraphs. The source code for this approach can be found at the following URL: `https://github.com/jorge-martinez-gil/nefusi2`

4.2.1. Compilation

```
...\nefusi-sentences\src>javac -cp combined-sentences.jar nefusi/*.java
```

Approach	Score
Google distance [3]	0.470
Huang et al. [12]	0.659
Jiang & Conrath [13]	0.669
Resnik [31]	0.780
Leacock & Chodorow [15]	0.807
Lin [16]	0.810
Faruqui & Dyer [9]	0.817
Mikolov et al. [26]	0.820
CoTO [18]	0.850
FLC [21]	0.855
BERT+Mamdani (median)	0.861
BERT+Mamdani (maximum)	0.867

Table 1: Results from the different approaches over the MC30 dataset using Pearson Correlation

4.2.2. Execution

```
...\nefusi-sentences\src>java -cp .;combined-sentences.jar nefusi.nefusi_Sentences
```

4.2.3. Datasets

We work here with the lawsentence200, which is composed of 200 pairs of paragraphs taken from documents of a legal nature and having had their degree of semantic similarity manually labeled by a group of legal experts using a scale ranging from 1 (not similar at all) to 5 (totally equivalent).

4.2.4. Results

Table 2 summarizes the results obtained at the sentence level for the lawsentence200 dataset and the comparison with other existing techniques.

5. Discussion

In engineering and other industrial applications, neurofuzzy systems have been the focus of a significant amount of research [28]. However, their application to natural language comprehension, and the challenge of semantic similarity assessment in particular, has yet to be thoroughly researched. Despite this, it is a fact that text can be transformed into numerical vectors by utilizing innovative techniques founded on neural-based problem-solving methods. The positional information regarding the words associated with the text can be preserved through this transformation. Consequently, they are an excellent option for processing complex words, sentences, and paragraphs. After that, an FL component can assist in configuring a semantic similarity score according to the user's requirements.

Approach	Score
USE Angular [2]	0.398
ELMo Euclidean [30]	0.546
ELMo Manhattan [30]	0.553
ELMo Angular [30]	0.618
ELMo Cosine [30]	0.624
USE Euclidean [2]	0.649
USE Manhattan [2]	0.660
USE Cosine [2]	0.705
BERT Pairwise [8]	0.736
BERT Manhattan [8] [8]	0.740
BERT Euclidean [8]	0.743
BERT Cosine [8]	0.780
BERT+Mamdani (median)	0.788
BERT Inner Product [8]	0.803
BERT+Mamdani (maximum)	0.826

Table 2: Results over the lawSentence200 dataset using Pearson Correlation

6. Conclusions

In the NEFUSI project, we have seen that, when taken into consideration separately, ANNs and FL each have several significant advantages and disadvantages [19]. ANNs can be utilized to acquire knowledge, but the learning process takes a very long time because it requires collecting enormous amounts of data; it is also typically complicated to understand the final model [24]. Additionally, it can be difficult to extract structured knowledge from the developed model and to include domain knowledge to speed up the learning process. The use of FL is the perfect complement since it tends to excel in the aspects that architectures of a neural nature suffer from. Furthermore, vice versa, it is weaker in most aspects that ANNs excel at. As a result of this project, we have developed a prototype capable of working at word and sentence levels that have proven to work quite well. Our prototype is available to the public for experimentation and improvement.

Since neurofuzzy systems have traditionally been utilized in control applications, this novel approach represents an innovation in bringing this computational model to the field of semantic similarity measurement. This approach makes it possible to look into how these hybrid systems can be used in a wide range of language-related situations, opening up a new field of possibilities. Although in the context of the NEFUSI project, we have focused on the measurement of semantic similarity as one of the critical activities for information retrieval on the Web, there are many more possibilities for building different ensembles and solutions in various domains: question answering or ontology matching, for example.

Acknowledgments

The development of NEFUSI is funded in the project NGI Zero Discovery by the NLnet Foundation and the European Commission. Project number: 2021-04-069

References

[1] Plamen P. Angelov and Richard A. Buswell. Automatic generation of fuzzy rule-based models from data by genetic algorithms. *Inf. Sci.*, 150(1-2):17–31, 2003.

[2] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. Universal sentence encoder for english. In Eduardo Blanco and Wei Lu, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, pages 169–174. Association for Computational Linguistics, 2018.

[3] Rudi Cilibrasi and Paul M. B. Vitányi. The google similarity distance. *IEEE Trans. Knowl. Data Eng.*, 19(3):370–383, 2007.

[4] Pablo Cingolani and Jesús Alcalá-Fdez. jfuzzylogic: a java library to design fuzzy logic controllers according to the standard for fuzzy control programming. *Int. J. Comput. Intell. Syst.*, 6(sup1):61–75, 2013.

[5] Oscar Cordón. A historical review of evolutionary learning methods for mamdani-type fuzzy rule-based systems: Designing interpretable genetic fuzzy systems. *Int. J. Approx. Reason.*, 52(6):894–913, 2011.

[6] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.*, 6(2):182–197, 2002.

[7] Scott C. Deerwester, Susan T. Dumais, Thomas K. Landauer, George W. Furnas, and Richard A. Harshman. Indexing by latent semantic analysis. *J. Am. Soc. Inf. Sci.*, 41(6):391–407, 1990.

[8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *Preprint arXiv:1810.04805*, 2018.

[9] Manaal Faruqui and Chris Dyer. Improving vector space word representations using multilingual correlation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2014, April 26-30, 2014, Gothenburg, Sweden*, pages 462–471, 2014.

[10] Lushan Han, Abhay L. Kashyap, Tim Finin, James Mayfield, and Jonathan Weese. Umbc_ebiquity-core: Semantic textual similarity systems. In Mona T. Diab, Timothy Baldwin, and Marco Baroni, editors, *Proceedings of the Second Joint Conference on Lexical and Computational Semantics, *SEM 2013, June 13-14, 2013, Atlanta, Georgia, USA*, pages 44–52. Association for Computational Linguistics, 2013.

[11] Sébastien Harispe, Sylvie Ranwez, Stefan Janaqi, and Jacky Montmain. *Semantic Similarity from Natural Language and Ontology Analysis*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2015.

[12] Eric H. Huang, Richard Socher, Christopher D. Manning, and Andrew Y. Ng. Improving word representations via global context and multiple word prototypes. In *The 50th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, July 8-14, 2012, Jeju Island, Korea - Volume 1: Long Papers*, pages 873–882, 2012.

[13] Jay J. Jiang and David W. Conrath. Semantic similarity based on corpus statistics and lexical taxonomy. In *Proceedings of the 10th Research on Computational Linguistics International Conference, ROCLING 1997, Taipei, Taiwan, August 1997*, pages 19–33, 1997.

[14] Juan J. Lastra-Díaz, Josu Goikoetxea, Mohamed Ali Hadj Taieb, Ana García-Serrano, Mohamed Ben Aouicha, and Eneko Agirre. A reproducible survey on word embeddings and ontology-based methods for word similarity: Linear combinations outperform the state of the art. *Eng. Appl. Artif. Intell.*, 85:645–665, 2019.

[15] Claudia Leacock and Martin Chodorow. Combining local context and wordnet similarity for word sense identification. *WordNet: An electronic lexical database*, 49(2):265–283, 1998.

[16] Dekang Lin. An information-theoretic definition of similarity. In *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998), Madison, Wisconsin, USA, July 24-27, 1998*, pages 296–304, 1998.

[17] E. H. Mamdani and S. Assilian. An experiment in linguistic synthesis with a fuzzy logic controller. *Int. J. Hum.-Comput. Stud.*, 51(2):135–147, 1999.

[18] Jorge Martinez-Gil. Coto: A novel approach for fuzzy aggregation of semantic similarity measures. *Cogn. Syst. Res.*, 40:8–17, 2016.

[19] Jorge Martinez-Gil. Semantic similarity aggregators for very short textual expressions: a case study on landmarks and points of interest. *J. Intell. Inf. Syst.*, 53(2):361–380, 2019.

[20] Jorge Martinez-Gil. A comprehensive review of stacking methods for semantic similarity measurement. *Machine Learning with Applications*, 10:100423, 2022.

[21] Jorge Martinez-Gil and Jose M. Chaves-Gonzalez. Automatic design of semantic similarity controllers based on fuzzy logics. *Expert Syst. Appl.*, 131:45–59, 2019.

[22] Jorge Martinez-Gil and Jose M. Chaves-Gonzalez. Semantic similarity controllers: On the trade-off between accuracy and interpretability. *Knowl. Based Syst.*, 234:107609, 2021.

[23] Jorge Martinez-Gil and Jose Manuel Chaves-Gonzalez. A novel method based on symbolic regression for interpretable semantic similarity measurement. *Expert Syst. Appl.*, 160:113663, 2020.

[24] Jorge Martinez-Gil and Jose Manuel Chaves-Gonzalez. Sustainable semantic similarity assessment. *Journal of Intelligent & Fuzzy Systems*, 43(5):6163–6174, 2022.

[25] Jorge Martinez-Gil, Riad Mokadem, Josef Küng, and Abdelkader Hameurlain. A novel neuro-fuzzy approach for semantic similarity measurement. In Matteo Golfarelli, Robert Wrembel, Gabriele Kotsis, A Min Tjoa, and Ismail Khalil, editors, *Big Data Analytics and Knowledge Discovery - 23rd International Conference, DaWaK 2021, Virtual Event, September 27-30, 2021, Proceedings*, volume 12925 of *Lecture Notes in Computer Science*, pages 192–203. Springer, 2021.

[26] Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 3111–3119, 2013.

[27] G. Miller and W. Charles. Contextual correlates of semantic similarity. *Language and Cognitive Processes*, 6(1):1–28, 1991.

[28] Detlef D. Nauck and Rudolf Kruse. Neuro-fuzzy systems for function approximation. *Fuzzy Sets Syst.*, 101(2):261–271, 1999.

[29] Roberto Navigli and Federico Martelli. An overview of word and sense similarity. *Nat. Lang. Eng.*, 25(6):693–714, 2019.

[30] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237, 2018.

[31] Philip Resnik. Semantic similarity in a taxonomy: An information-based measure and its application to problems of ambiguity in natural language. *J. Artif. Intell. Res.*, 11:95–130, 1999.

[32] Leszek Rutkowski and Krzysztof Cpalka. Flexible neuro-fuzzy systems. *IEEE Trans. Neural Networks*, 14(3):554–574, 2003.

[33] Barbara Rychalska, Katarzyna Pakulska, Krystyna Chodorowska, Wojciech Walczak, and Piotr Andruszkiewicz. Samsung poland NLP team at semeval-2016 task 1: Necessity for diversity; combining recursive autoencoders, wordnet and ensemble methods to measure semantic similarity. In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval-NAACL-HLT 2016, San Diego, CA, USA, June 16-17, 2016*, pages 602–608, 2016.

[34] Himanshu Singh and Yunis Ahmad Lone. *Deep Neuro-Fuzzy Systems with Python*. Springer, 2020.

[35] Tomohiro Takagi and Michio Sugeno. Fuzzy identification of systems and its applications to modeling and control. *IEEE Trans. Syst. Man Cybern.*, 15(1):116–132, 1985.

[36] Ganggao Zhu and Carlos Angel Iglesias. Computing semantic similarity of concepts in knowledge graphs. *IEEE Trans. Knowl. Data Eng.*, 29(1):72–85, 2017.