Information Fusion in Sentiment Fuzzy Rule-Based systems: How to Improve Readability and Robustness through SMART Operators

Andrea Capotorti[1], Davide Petturiti[1], and Barbara Vantaggi[2]

[1]University of Perugia, Italy
{*andrea.capotorti, davide.petturiti*}*@unipg.it*
[2]University of Rome "la Sapienza", Italy
barbara.vantaggi@uniroma1.it

Abstract

We propose the adoption of the newly introduced SMART operators to improve the readability of a Fuzzy rule-based system developed for sentiment analysis. Specifically, the S-or operator will be utilized to combine the results of various rules generated from a designated lexicon. Conversely, the S-and operator will merge the diverse inferences derived from multiple lexicons.

1 Introduction

Fuzzy rule-based systems, such as those discussed in Alonso Moral et al. (2021)), offer outputs that are inherently explainable. This is particularly evident when these systems are utilized in sentiment analysis, as demonstrated in Liu and Cocea (2017). However, traditional Mamdani-type systems (Mamdani (1974); Mamdani and Assilian (1975)), employing standard min-max aggregation operators, often yield outputs that are challenging to interpret prior to the final defuzzification step . To address this issue, we propose the adoption of the SMART disjunctive operator (referred to as S-or hereafter), which was recently introduced in Capotorti and Figà-Talamanca (2020). By utilizing the S-or operator instead of the conventional max operator for aggregation, we aim to produce more easily interpretable outputs for the consequents of the rules. This approach enhances the overall interpretability of the fuzzy rule-based system, ultimately improving the clarity and understanding of the system's outputs.

In sentiment analysis, Fuzzy rule-based systems typically rely on scores derived from lexicons to provide crisp inputs (see, e.g., Nadali et al. (2010)). However, it is important to note that different lexicons can yield varying results, as demonstrated in previous studies, such as Chauhan et al. (2023); Vashishtha and Susan (2019). Usually, only the final defuzzified crisp values are utilized to evaluate classifier performance, overlooking the richness and diversity inherent in the various Fuzzy outputs. To address this limitation, we apply a novel approach of merging multiple Fuzzy outputs using the SMART conjunctive operator (S-and hereafter) before the defuzzification process, as suggested in Capotorti and Figà-Talamanca (2020). This method allows for the creation of an ensemble of classifications, offering a unique perspective on sentiment analysis. Our approach differs significantly from existing methods, such as Shapiro and Moritz Sudhof (2020)), and provides a fresh and innovative solution to the challenges in sentiment analysis.

The paper is organized in the following way: In Section 2 we briefly illustrate a Mamdami-type Fuzzy Rule-Base apt to deal with sentiment analysis with two different scores obtained through a lexicon and how the S-or operator can be employed to obtain a more interpretable output membership. Moreover, we propose a new dufuzzification operator that is more accurate of the usual center of area (COA). In Section 3.1 we propose the adoption of the other aggregation operator S-and to join the information stemming from different sources, in particular different lexicons, obtaining and ensemble of classifiers. Finally, a concluding section summarizes the contribution and hints about future developments.

2 FUZZY RULE-BASED SYSTEM

Our proposal is based on the Fuzzy Rule-Based system outlined in the study by Vashishtha and Susan (2019). This system was chosen for its clear and easily understandable design, as well as its strong classification performance. However, it is important to note that our method is versatile and can be applied to any Fuzzy Rule-Based system of the Mamdani type.

In the aforementioned paper by Vashishtha and Susan (2019), the authors utilized a system consisting of 9 well-defined rules, as detailed in Tab. 1; Each rule is composed of two antecedent variables, each with three Fuzzy subsets (Low, Medium, High), and one consequent variable with three Fuzzy subsets (Negative, Neutral, Positive).

Table 1: The nine Mamdani Fuzzy rules proposed in Vashishtha and Susan (2019).

RULE	Pos. score	Neg. score	Sentiment
R1	Low	Low	Neutral
R2	Medium	Low	Positive
R3	High	Low	Positive
R4	Low	Medium	Negative
R5	Medium	Medium	Neutral
R6	High	Medium	Positive
R7	Low	High	Negative
R8	Medium	High	Negative
R9	High	High	Neutral

The two input variables, $X1$ and $X2$, respectively represent the negative and the positive sentiment score of a specific text obtained by adopting one specific lexicon (e.g. SentiWordNet Baccianella et al. (2010), AFINN Nielsen (2011) or VADER Hutto and Gilbert (2015)). For the three linguistic terms Low-Medium-High, authors use a standard Fuzzy partition of the domain of each input score with three triangular memberships $\mu_i(x)$, $i \in \{\text{Low, Medium, High}\}$, as those depicted in Fig.1 Usually, but not necessarily, input

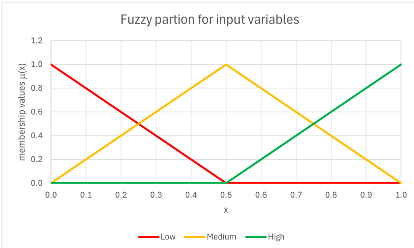

Figure 1: Fuzzy partition for input variables.

domains are normalized into the real unit interval $[0, 1]$.

For the output variable Y, representing the overall sentiment judgment of the selected text (in the aforementioned paper authors analyze Twitter posts), authors use a similar Fuzzy partition with three linguistic terms, but on the standard scale of numbers between 0 an 10, with memberships $\mu_i(y)$, $i \in \{neg, neu, pos\}$, illustrated in Fig.2

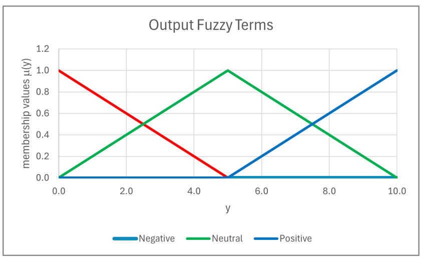

Figure 2: Fuzzy linguistic terms for the output variable.

The sentiment classification for each text is performed through the following steps.

- The two input scores $x1$ and $x2$ are fuzzified into the membership values $\mu_{j1}(x1)$ and $\mu_{j2}(x2)$, respectively, of each of the 9 rules in Tab.1.

- Through the usual minimum t-norm, the firing strength of the j-th rule is obtained as:

$$w_{Rj} = \mu_{j1}(x1) \wedge \mu_{j2}(x2). \tag{1}$$

- The usual maximum t-conorm $\vee$ is applied for combining the firing strengths derived in the previous step in order to obtain the fulfillment of each output term:

$$w_{neg} = w_{R4} \vee w_{r7} \vee w_{R8} \tag{2}$$
$$w_{neu} = w_{R1} \vee w_{R5} \vee w_{R9} \tag{3}$$
$$w_{pos} = w_{R2} \vee w_{R3} \vee w_{R6} \tag{4}$$

- The resultant output memberships for all terms are obtained through the usual minimum t-norm $\wedge$:

$$\tilde{\mu}_{neg}(y) = w_{neg} \wedge \mu_{neg}(y); \tag{5}$$
$$\tilde{\mu}_{neu}(y) = w_{neu} \wedge \mu_{neu}(y); \tag{6}$$
$$\tilde{\mu}_{pos}(y) = w_{pos} \wedge \mu_{pos}(y). \tag{7}$$

- The final output memberships $\tilde{\mu_j}$, $j \in \{neg, neu, pos\}$, are again aggregated through the usual maximum t-conorm $\vee$ before being defuzzified:

$$\mu_{\vee-\max}(y) = \tilde{\mu}_{neg}(y) \vee \tilde{\mu}_{neu}(y) \vee \tilde{\mu}_{pos}(y). \tag{8}$$

The usual aggregation produces memberships $\mu_{\vee-\max}$ that are not easy to interpret due to their potential extreme vagueness, e.g. those depicted in Fig.3.

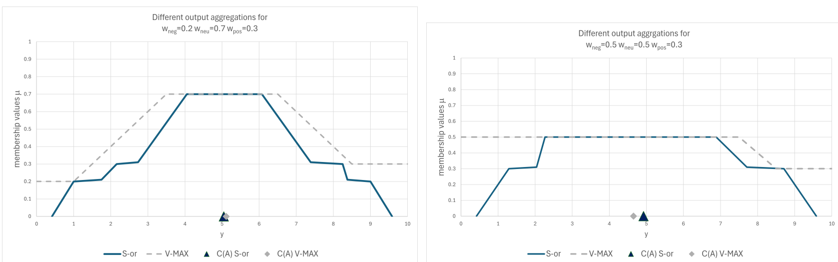

Figure 3: Comparison of aggregated output memberships obtained by $\vee-\max$ or S-or operators.

We perform here the same exercise by applying the disjunctive aggregation operator S-or in the last step; the results highlight that this alternative merge produces less vague and concave memberships. Indeed, $\mu_{S-or} \subseteq \mu_{\vee-\max}$ using usual Fuzzy inclusion relation, and μ_{S-or} displays nested α-cuts. Such membership can be easy interpreted as a truncated Fuzzy number. Fig.3 provides a visual comparison of μ_{S-or} w.r.t. $\mu_{\vee-\max}$).

2.1 The S-or aggregation operator

The S-or function operates by calculating a weighted average of the extreme values of the different α-cuts of the m memberships to be aggregated. The weights are carefully adjusted to achieve a specific behavior in the aggregation process, focusing on the outermost values. This approach aligns with the canonical max t-conorm $\vee$ if applied vertically. The weighting emphasizes the disagreement among the different α-cuts in terms of weak overlapping. To achieve this, when fixing an α-cut, the weights of the outermost values, identified by indexes in O_l for the left extrema and O_r for the right extrema, are calculated as $\frac{1}{m}(1+\epsilon_j)$, with

$$\epsilon_j = \begin{cases} \frac{\sum_{f=1}^{M} \frac{1}{f}\pi_f^j}{\Delta} & \text{if } \Delta \neq 0 \\ 0 & \text{otherwise} \end{cases} . \tag{9}$$

In (9) the π_f^j's are the lengths of the various overlappings while Δ is the range of the α-cuts, as depicted in Fig.4.

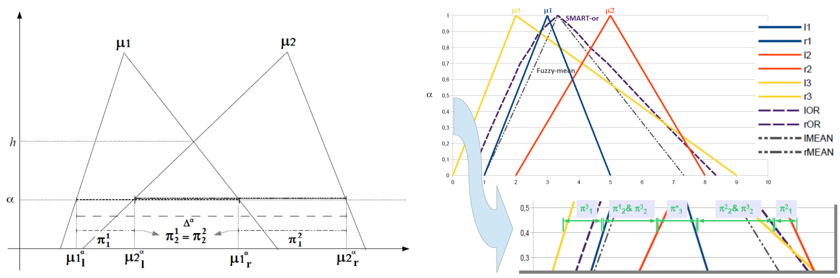

Figure 4: Overlapping lengths π_f^j and range Δ of α-cuts involved in the weights ϵ_j for the S-or aggregation: above between two memberships; below among three.

For the two inner extremes, i.e. the largest left extreme and the lowest right extreme, the weights are simply given by $\frac{1}{m}(1-\sum_{j\in O_*}\epsilon_j)$.

The memberships m that need to be aggregated are truncated by the firing strengths in equations (5, 6, 7). As a result, these memberships may not be normal, meaning there is no value $\tilde{y}$ for which $\tilde{\mu}_j(\tilde{y}) = 1$. This is different from the original formulation in Capotorti and Figà-Talamanca (2020), where the memberships were normal fuzzy numbers. Our goal is to ensure that the output of this merging process is a fuzzy number, possibly truncated. When the number of α-cuts to aggregate changes, the resulting extremes must align with those of the lower levels. This alignment can be achieved through a proper translation and deformation of the extremes.

Due to the necessary discretization of the α levels to be considered in practical applications, we can formulate the transformation by referring to two consecutive values α_m and α_{m+1}. The extremes of the transformed α_{m+1}-cut are computed recursively as follows:

$$\widehat{\mu}_l^{\alpha_{m+1}} = \widehat{\mu}_l^{\alpha_m} + \varrho^{\alpha_m} |\overline{\mu}_l^{\alpha_{m+1}} - \overline{\mu}_l^{\alpha_m}| \tag{10}$$

$$\widehat{\mu}_r^{\alpha_{m+1}} = \widehat{\mu}_r^{\alpha_m} - \varrho^{\alpha_m} |\overline{\mu}_r^{\alpha_{m+1}} - \overline{\mu}_r^{\alpha_m}| \tag{11}$$

with

$$\varrho^{\alpha_m} = \frac{\widehat{\mu}_r^{\alpha_m} - \widehat{\mu}_l^{\alpha_m}}{\overline{\mu}_r^{\alpha_m} - \overline{\mu}_l^{\alpha_m}} \tag{12}$$

and where the "overlined" extremes are those obtained by the previously described S-or and the "hatted" ones are those finally obtained by the transformation at the specified levels.

It is important to note that the typical method of discretizing the membership functions $\tilde{\mu_j}$, where $j \in \{neg, neu, pos\}$, results in step functions with a finite number of distinct α-cuts. This characteristic makes the S-or operator easily applicable and effectively implementable for our purposes. It is worth mentioning that the original formulation in Capotorti and Figà-Talamanca (2020) was intended for Fuzzy numbers, whereas we are working with more general Fuzzy quantities.

2.2 Defuzzification

The final step in text sentiment classification involves synthesizing the final output membership, denoted as $\mu_{\vee-\max}$ or μ_{S-or}, into a single real value that represents the overall sentiment of the text. This value is then used to assign a final sentiment label based on a predetermined decision rule. For example, the following decision rule adopted in Vashishtha and Susan (2019):

$$\text{Label} = \begin{cases} \text{Negative} & \text{if } COA \in [0, 3.3[\\ \text{Neutral} & \text{if } COA \in [3.3, 6.7[\\ \text{Positive} & \text{if } COA \in [6.7, 10] \end{cases} \tag{13}$$

where COA is the usual defuzzified center of the area value of their aggregated membership $\mu_{\vee-\max}$ usually approximated by

$$COA = \frac{\sum_{y_i \in \mathcal{Y}} y_i \, \mu_{\vee-\max}(y_i)}{\sum_{y_i \in \mathcal{Y}} \mu_{\vee-\max}(y_i)}, \tag{14}$$

summations being over the partition $\mathcal{Y} = \{y_0, y_1, \ldots, y_n\}$ chosen for the discrete operational representation of the memberships (in the aforementioned paper authors chose, e.g., $\mathcal{Y} = \{0, 1, \ldots, 10\}$).

Since the polygonal shape of the output membership μ_{S-or}, we suggest instead to approximate the center of the area through the more accurate value proposed in Naimi and Tahayori (2020) and defined as:

$$C(A) = \frac{\sum_{j=1}^{n}\left((y_j - y_{j-1})\left[\mu_{j-1}\frac{(2y_{j-1}+y_j)}{3} + \mu_j \frac{(y_{j-1}+2y_j)}{3}\right]\right)}{\sum_{j=1}^{n}((y_j - y_{j-1})(\mu_{j-1}+\mu_j))}. \tag{15}$$

where μ_{j-1} and μ_j stay for $\mu_{S-or}(y_{j-1})$ and $\mu_{S-or}(y_j)$, respectively, and $\mathcal{Y} = \{0, 0.1, \ldots, 9.9, 10\}$.

The center of the area is quite insensitive to the specific shape of the membership functions, consequently, we obtain almost the same values with $\mu_{\vee-\max}$ or μ_{S-or}. Hence, concerning the classification aim, our proposal reaches the same performances of Vashishtha and Susan (2019). In fact, our goal is to improve the interpretability rather than the performance and we expect that this characteristic of our approach will be valuable when applied to real and extensive datasets.

3 An Ensemble of Sentiment Classifier

As mentioned in the introduction, sentiment classifications are typically conducted using different lexicons, with each lexicon providing a classification for a given instance. However, there are instances where different lexicons may produce varying classifications for the same data point. For example, in the Sanders dataset referenced in Vashishtha and Susan (2019), tweet # 3420 is labeled as Neutral using the AFINN lexicon with a COA of 5.4, but is classified as Positive when processed through the VADER lexicon with a COA of 7.67. To address this issue, the results from different lexicons can be aggregated to provide a more comprehensive classification. One approach is to calculate the arithmetic mean of the centers of the areas $C(A)$'s in (15), but this method may not account for the unique shapes of the membership functions. A weighted average could offer a more nuanced solution, although determining the appropriate weights can be challenging.

An alternative method is to utilize the conjunctive SMART operator S-and, as introduced in Capotorti and Figà-Talamanca (2020). This operator considers the agreements among the α-cuts of the initial membership functions, resulting in a more robust aggregation of output memberships. By combining the output memberships obtained from different lexicons, a single aggregated membership can be obtained using the S-and operator. The center of the area of this aggregated membership can then be used to determine the classification label through a decision rule.

Precisely, by aggregating the different output memberships μ_{s-or}'s obtained for different lexicons, a single aggregated joined membership μ_{S-and} is obtained; by computing the center of the area (15) of μ_{S-and} we can eventually decide the classification label through some decision rule, like, e.g., (13). As a result, this ensemble approach provides a weighted mean of the different classifications, with the final center of the area of the aggregated membership representing a combination of the centers of the individual memberships, implicitly weighted based on their agreement or disagreement.

3.1 The S-and aggregation operator

In order to gain a simple understanding of the S-and operator, let us briefly outline the main steps involved. For more detailed information, please refer to the source cited as Capotorti and Figà-Talamanca (2020).

The S-and operator is a variation of the Fuzzy mean that considers the full or partial overlap among the α-cuts of the Fuzzy memberships being combined.

Let m represent the number of membership functions to be aggregated, $I_l \subset \{1, \ldots, m\}$ denote the set of indices of the last $m-1$ left extremes, and $I_r \subset \{1, \ldots, m\}$ represent the set of indices of the first $m-1$ right extremes (in ascending order) of the m α-cuts in input. The construction of the S-and operator follows a similar basic rule to the S-or operator described in Subsection 2.1, but requires a more complex formulation. The computation of the α-cuts of the result varies depending on whether α is below or above the value $h \in (0, 1)$, which represents the highest level of non-empty intersection among all the m α-cuts of the original Fuzzy numbers (see Fig.4 for the case $m = 2$).

If $\alpha \leq h$, the extremes of the α-cuts are determined as convex combinations of the original extremes using coefficients given by:

$$\frac{1}{m}(1 + \gamma_j) \qquad j \in I_l \text{ or } I_r \ , \tag{16}$$

where the quantities

$$\gamma_j = \frac{\sum_{f=1}^{m} \frac{1}{m+1-f} \pi_f^j}{\sum_{k=1}^{m} \sum_{f=1}^{m} \frac{1}{m+1-f} \pi_f^k} \tag{17}$$

represent the weighted normalized contribution of the $m-1$ most relevant extremes, specifically those in I_l for the left extremes and those in I_r for the right extremes.

It is worth to remark that the factor $\frac{1}{m+1-f}$ for each term in the numerator of γ_j is directly proportional to the number of overlaps, which represent the level of agreement. Additionally, in cases where all m α-cuts align, the values of γ_j can be consistently set to zero.

Furthermore, the coefficients associated with the m-th position, specifically those linked to the outermost left extreme in $\{1, \ldots, m\} \setminus I_l$ and the outermost right extreme in $\{1, \ldots, m\} \setminus I_r$, can be expressed as follows:

$$\frac{1}{m}(1 - \sum_{j=I_*} \gamma_j) \qquad I_* = I_l,\, I_r \text{ respectively.} \tag{18}$$

In situations where α exceeds the threshold h, the formal definition becomes more intricate as various subgroups of intersections involving two or more memberships need to be identified.

The logic underlying the method is to compute the S-and operator within each subgroup obtaining different intermediate α-levels, and merge them in a second step, by applying the S-or operator. For the technical details, we refer the reader once again to Capotorti and Figà-Talamanca (2020).

Furthermore, when utilizing the S-and operator, it is necessary that any changes in the number of memberships to be merged at a certain α-level result in the nesting of the resulting extremes within those of the lower levels. This nesting process should be achieved through translation or deformation, employing the same technique described for the S-or operator in equations (10-11).

In Figure 5, the results of the S-and and of the two S-or outputs for tweet #3420 from the Sanders dataset are plotted, as documented in Vashishtha and Susan (2019). The AFINN lexicon assigned a positive score of $x_1 = 3.0$ and a negative score of $x_2 = 0.0$ to the tweet. By applying equations (2-4), we calculated a negative firing strength of $w_{neg} = 0.0$, a neutral firing strength of $w_{neu} = 0.625$, and a positive firing strength of $w_{pos} = 0.375$.

When processed through the VADER lexicon, the tweet received a positive score of $x_1 = 0.5$ and a negative score of $x_2 = 0.0$, resulting in a negative firing strength of $w_{neg} = 0.0$, a neutral firing strength of $w_{neu} = 0.0$, and a positive firing strength of $w_{pos} = 1.0$.

The membership value μ_{S-and} is compared with the original $\mu_{\vee-\max}$ values. It is important to note that the center of the area $C(A)$ is calculated to be 6.19, which differs from the arithmetic mean of 6.99 for the two original $C(A)$ values, which were 5.65 and 8.33, respectively.

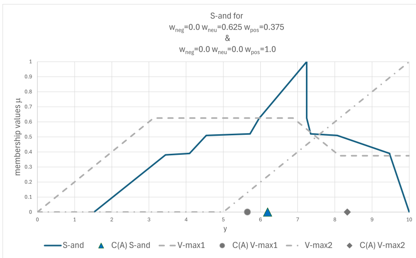

Figure 5: Comparison of the S-and membership and $C(A)$ w.r.t. two $\vee-\max$ of Vashishtha and Susan (2019) for tweet #3420 of Sanders dataset.

4 Conclusion

We have shown how the implementation of the recently introduced SMART Fuzzy aggregation operators, S-or and S-and, can improve the interpretability of a rule-based system proposed in Vashishtha and Susan (2019) for sentiment classification. Our innovative approach can be seamlessly integrated into any other Fuzzy rule-based system of Mandami type, regardless of the lexicon used.

Specifically, the S-or aggregator generates a truncated Fuzzy number as the output of the inference process, which is more precise and easier to interpret compared to the traditional V-max aggregation method. The S-or membership represents a Fuzzy number centered around its $C(A)$ value, with a weight determined by the maximum height of the membership.

On the other hand, the S-and operator enables the creation of an ensemble of sentiment classifiers utilizing various lexicons, resulting in a final crisp sentiment score that is a weighted average of scores obtained from individual lexicons. These weights implicitly capture the level of agreement or disagreement among the original Fuzzy outputs.

In conclusion, the SMART Fuzzy aggregation operators offer a versatile and effective solution for enhancing the interpretability and performance of rule-based sentiment classification systems. Further research will be devoted to apply these operators to classify the sentiment of news articles retrieved from financial US newspaper and to finally assess the impact of the fuzzy-based sentiment score in the return and volatility dynamics of major US stocks.

As a distinguished target problem, we envisage the application of the present method to find associations between the sentiment of central banks (FED and ECB) communications (via policy statements, post-meeting press conferences, economic forecasts, monetary policy reports, speeches, interviews, and testimony to parliament) and the monetary policy stance, similarly to what has been done in Hilscher et al. (2024).

Acknowledgments:

We acknowledge the support of the PRIN 2022 project "Models for dynamic reasoning under partial knowledge to make interpretable decisions" (Project number: `2022AP3B3B`, CUP Master: `J53D23004340006`, CUP: `B53D23009860006`) funded by the European Union – Next Generation EU.

References

J. M. Alonso Moral, C. Castiello, L. Magdalena, and C. Mencar. *An Overview of Fuzzy Systems*, pages 25–47. Springer International Publishing, Cham, 2021. ISBN 978-3-030-71098-9. doi: 10.1007/978-3-030-71098-9_2. URL `https://doi.org/10.1007/978-3-030-71098-9_2`.

S. Baccianella, A. Esuli, and F. Sebastiani. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the 7th International Conference on Language Resources and Evaluation, LREC 2010*, pages 2200 —- 2204, 2010. URL `https://www.scopus.com/inward/record.uri?eid=2-s2.0-85034054192&partnerID=40&md5=42db614bb847384aba6a349b4e0c98f9`. Cited by: 2502.

A. Capotorti and G. Figà-Talamanca. Smart-or and smart-and fuzzy average operators: A generalized proposal. *Fuzzy Sets and Systems*, 395:1–20, 2020. ISSN 0165-0114. doi: https://doi.org/10.1016/j.fss.2019.04.027. URL `https://www.sciencedirect.com/science/article/pii/S0165011419302349`. Aggregation Operations.

S. Chauhan, J. P. Shet, S. M. Beram, V. Jagota, M. Dighriri, M. W. Ahmad, M. S. Hossain, and A. Rizwan. Rule based fuzzy computing approach on self-supervised sentiment polarity classification with word sense disambiguation in machine translation for hindi language. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 22(5), may 2023. ISSN 2375-4699. doi: 10.1145/3574130. URL `https://doi.org/10.1145/3574130`.

J. Hilscher, K. Nabors, and A. Raviv. Information in central bank sentiment: An analysis of fed and ecb communication. *Available at SSRN*, 2024. URL `http://dx.doi.org/10.2139/ssrn.4797935`.

C. Hutto and E. Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM 2014*, 01 2015.

H. Liu and M. Cocea. Fuzzy rule based systems for interpretable sentiment analysis. In *2017 Ninth International Conference on Advanced Computational Intelligence (ICACI)*, pages 129–136, 2017. doi: 10.1109/ICACI.2017.7974497.

E. Mamdani. Application of fuzzy algorithms for control of simple dynamic plant. *Proceedings of the Institution of Electrical Engineers*, 121:1585–1588(3), December 1974. ISSN 0020-3270. URL https://digital-library.theiet.org/content/journals/10.1049/piee.1974.0328.

E. Mamdani and S. Assilian. An experiment in linguistic synthesis with a fuzzy logic controller. *International Journal of Man-Machine Studies*, 7(1):1–13, 1975. ISSN 0020-7373. doi: https://doi.org/10.1016/S0020-7373(75)80002-2. URL https://www.sciencedirect.com/science/article/pii/S0020737375800022.

S. Nadali, M. A. A. Murad, and R. A. Kadir. Sentiment classification of customer reviews based on fuzzy logic. In *2010 International Symposium on Information Technology*, volume 2, pages 1037–1044, 2010. doi: 10.1109/ITSIM.2010.5561583.

M. Naimi and H. Tahayori. Centroid of polygonal fuzzy sets. *Applied Soft Computing*, 95: 106519, 2020. ISSN 1568-4946. doi: https://doi.org/10.1016/j.asoc.2020.106519. URL https://www.sciencedirect.com/science/article/pii/S1568494620304580.

F. Å. Nielsen. A new ANEW: evaluation of a word list for sentiment analysis in microblogs. In M. Rowe, M. Stankovic, A.-S. Dadzie, and M. Hardey, editors, *Proceedings of the ESWC2011 Workshop on 'Making Sense of Microposts': Big things come in small packages*, volume 718 of *CEUR Workshop Proceedings*, pages 93–98, May 2011. URL http://ceur-ws.org/Vol-718/paper_16.pdf.

A. H. Shapiro and D. W. Moritz Sudhof. Measuring news sentiment. Working Paper 2017-01, Federal Reserve Bank of San Francisco, 2020. URL https://doi.org/10.24148/wp2017-01.

S. Vashishtha and S. Susan. Fuzzy rule based unsupervised sentiment analysis from social media posts. *Expert Systems with Applications*, 138:112834, 2019. ISSN 0957-4174. doi: https://doi.org/10.1016/j.eswa.2019.112834. URL https://www.sciencedirect.com/science/article/pii/S0957417419305366.