**University of
Zurich^{UZH}**

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2020

MuSeM: Detecting Incongruent News Headlines using Mutual Attentive Semantic Matching

Mishra, Rahul ; Yadav, Piyush ; Calizzano, Remi ; Leippold, Markus

Abstract: Measuring congruence between two texts has several useful applications, such as detecting the prevalent deceptive and misleading news headlines on the web. Many works have proposed machine learning based solutions such as text similarity between the headline and body text to detect the incongruence. Text similarity based methods fail to perform well due to different inherent challenges such as relative length mismatch between the news headline and its body content and non-overlapping vocabulary. On the other hand, more recent works that use headline guided attention to learn a headline derived contextual representation of the news body also result in convoluting overall representation due to the news body's lengthiness. This paper proposes a method that uses inter-mutual attention-based semantic matching between the original and synthetically generated headlines, which utilizes the difference between all pairs of word embeddings of words involved. The paper also investigates two more variations of our method, which use concatenation and dot-products of word embeddings of the words of original and synthetic headlines. We observe that the proposed method outperforms prior-arts significantly for two publicly available datasets.

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-190878>

Conference or Workshop Item

Published Version

Originally published at:

Mishra, Rahul; Yadav, Piyush; Calizzano, Remi; Leippold, Markus (2020). MuSeM: Detecting Incongruent News Headlines using Mutual Attentive Semantic Matching. In: International Conference on Machine Learning and Applications (ICMLA) 2020, Miami, Florida, 14 December 2020 - 17 December 2020, IEEE ICMLA.

MuSeM: Detecting Incongruent News Headlines using Mutual Attentive Semantic Matching

Rahul Mishra
University of Stavanger, Norway
rahul.mishra@uis.no

Piyush Yadav
Lero, NUI Galway, Ireland
p.yadav1@nuigalway.ie

Remi Calizzano
DFKI, Germany
remi.calizzano@dfki.de

Markus Leippold
University of Zurich, Switzerland
markus.leippold@bf.uzh.ch

Abstract—Measuring the congruence between two texts has several useful applications, such as detecting the prevalent deceptive and misleading news headlines on the web. Many works have proposed machine learning based solutions such as text similarity between the headline and body text to detect the incongruence. Text similarity based methods fail to perform well due to different inherent challenges such as relative length mismatch between the news headline and its body content and non-overlapping vocabulary. On the other hand, more recent works that use headline guided attention to learn a headline derived contextual representation of the news body also result in convoluting overall representation due to the news body’s lengthiness. This paper proposes a method that uses inter-mutual attention-based semantic matching between the original and synthetically generated headlines, which utilizes the difference between all pairs of word embeddings of words involved. The paper also investigates two more variations of our method, which use concatenation and dot-products of word embeddings of the words of original and synthetic headlines. We observe that the proposed method outperforms prior arts significantly for two publicly available datasets.

I. INTRODUCTION

In the age of the prevalence of smart handheld devices, most of the information consumption is digital. This paradigm shift in the way people consume information has also brought forth several new challenges, such as misinformation and deceptive content. News headlines that incorrectly represent the content of the news body are called incongruent or click-baits. A deceptive and misleading news headline can result in false beliefs and wrong opinions. News titles play an essential role in making first impressions to readers and thereby deciding the viral potential of news stories within social networks [1]. Most users rely only on the news title content to determine which news items are significant enough to read [2]. The curse of deceptive content gets amplified by several magnitudes when people share it without reading news body content. Consider an example of an incongruent headline in Figure 1 taken from CNN (cnn.com).¹ The headline states that “Trump says GOP working on tax plan for middle class” whereas body content mentions that “It’s unclear what tax proposal Trump was referring to on Saturday” which contradicts with the headline.

¹<https://edition.cnn.com/2018/10/20/politics/donald-trump-tax-middle-income/index.html>

Fig. 1: Examples of Incongruent Headlines related to politics and health-care.

In another example in Figure 1,² the headline reads “The Scary New Science That Shows Milk Is Bad For You” but a part of the body text clearly states that “The study was small, to be sure, and it included no women.” The headline radically generalizes the claim made in the study and exaggerates it.

Many machine learning based solutions have been investigated previously in the literature to detect click-baits and news headline incongruence. Some initial works [3][4] use sentence matching based methods to compute similarity or dissimilarity between the web claims and news headlines to detect incongruence. Researchers [5][6] have also utilized

²<https://www.motherjones.com/environment/2015/11/dairy-industry-milk-federal-dietary-guidelines/>

different classification methods and used linguistic and stylistic features to learn a classifier to identify the incongruity between news titles and news body. The authors in [7] propose neural attention-based methods to find the entailment between the news items' title and body. Some recent works [8] have identified the significance of generative methods such as generative adversarial networks for incongruence detection.

Previous methods applied to click-bait detection that use natural language processing techniques are not suitable for the detection of headline incongruence as this problem requires more facets and aspects to be covered than just stylistic features [9]. The methods that use a text similarity based approach to detect incongruence perform poorly because of the long text content of the news body. Text similarity schemes work well in case of short texts.

This paper proposes a semantic matching technique based on inter-mutual attention that uses a synthetically generated headline corresponding to the news body content and original news headline to detect the incongruence. The proposed inter-mutual attention technique is inspired by some recent works, which use the intra-mutual word to word attention within a sentence to detect sarcasm [10] and intra-mutual user to user attention within a retweet propagation path sequence to detect misinformation [11].

In sum, the major contributions of this work are as follows:

- 1) We are the first to use inter-mutual attention based semantic matching to detect incongruent news headlines. The key idea is to get the pairwise difference between word embeddings of original and synthetic headlines and compute a mutual attention matrix after applying a dense layer. Subsequently, row-wise max-pooling can be used to compute the attention scores, to be used for the classification.
- 2) We use synthetic headlines generated from various generative adversarial networks based schemes using news body content to get the effective, contextual, and low dimensional representation.
- 3) We also investigate two additional variants of the proposed model, which incorporate addition and concatenation of word embeddings of the word pairs of original and synthetic headlines.
- 4) We combine all the three variants of word embedding operations in a clubbed model, which outperforms the three variants individually.
- 5) We conduct experiments with two publicly available datasets, which show the effectiveness of the proposed models to detect incongruent news headlines.

II. RELATED WORKS

Most of the literature's initial works propose using linguistic and statistical features based classifiers to detect click-baits and incongruent headlines. The authors of [12] suggest to perform lexical and syntactic analysis for the identification of the click-baits. In [13], authors use linguistic features to learn support vector machine (SVM) based classifier to detect click-baits, and they also release a manually annotated dataset.

Authors of [14] use text features and meta-information of tweets to learn a classifier to detect click-baits.

On the other hand, some works [3][4] deal with sentence matching based stance classification which is a closely related problem to headline incongruence. These methods are not suitable directly for headline incongruence due to some challenges such as relative length mismatch between the news headline and its body content and non-overlapping vocabulary. Authors in [15] propose to use a co-training approach with myriad kinds of features such as, e.g., sentiments, textual, and informality. Some very recent works such as [7] use neural attention [16] based approach to achieve headline guided contextual representation of the news body text. This headline guided attention also results in convoluting the body content's overall representation due to its lengthiness.

We propose an improved semantic matching between the news headline and its body text via inter-mutual attention. The inter-mutual attention based matching is performed between the original and a synthetically generated headline, rather than between original news headline and its body content. The synthetic headline is generated via state-of-the-art generative adversarial methods using the body contents of the news item. The generative adversarial techniques have been very successful in generating realistic artificial images [17], videos [18][19] and text [20] contents. One of the pioneer work in artificial text generation is [21], which solves the problem of generator differentiation by updating the gradient policy directly. The authors of [20] introduce a GAN based text generation method, which uses a third module, called MAN-AGER, which helps the generator network to utilize some leaked feature information from the discriminator network. In [8], authors introduce a synthetic headline generation method based on style transfer and generative adversarial network.

The key idea behind the inter-mutual attention method is inspired by two recent contributions [10][11]. There are some significant differences between these works and the proposed model. Firstly, they use intra-mutual word to word or user to user attention, which builds on word embedding pairs or user embedding pairs occurring within a sentence or retweet path sequence correspondingly to compute the attention scores. In contrast, the proposed model uses pairs of word embeddings of words that belong to two different sentences, i.e., original and synthetic headline. Secondly, these previous methods use a concatenation of embedding vectors to form an overall pair embedding representation. In contrast, the proposed method uses the difference between the word embeddings to form an overall pair embedding representation. The intuition behind these differences is evident as we detect incongruence between two pieces of texts (headline and body), and computing the difference between word embedding pairs results in capturing the cues related to semantic similarity or dissimilarity between the two.

III. PROBLEM DEFINITION AND PROPOSED MODEL

This section introduces the problem definition and presents the overall architecture of the proposed model MuSeM in

Fig. 2: Overall Architecture of the MuSeM model. (a) A version of the MuSeM model with difference operation between word embedding pairs. (b) A depiction of concatenation operation between word embedding pairs. (c) A depiction of the dot-product operation between word embedding pairs.

detail. We describe the synthetic headline generation schemes, and then we provide details of the inter-attentive semantic matching technique. We also present two more variants of MuSeM model and devise a clubbed model, which combines all the three variants.

A. Problem Definition

Given a news item $n_i \in N$, having headline h_i and body content b_i , we need to predict whether there is incongruence between h_i and b_i by classifying the news as either “Congruent(C)” or “Incongruent(I).” The news headline h_i consists of a sequence of l words denoted as $h_i = \{w_{h1}, w_{h2}, \dots, w_{hl}\} \in W$, and the news body content b_i consists of a sequence of m words denoted as $b_i = \{w_{b1}, w_{b2}, \dots, w_{bm}\} \in W$.

B. Word Embedding Layer

For a news item n_i , the corresponding headline h_i of length l and body content b_i of length m are represented as $h_i = \{f(w_{h1}), f(w_{h2}), \dots, f(w_{hl})\}$ where $\forall j, f(w_{hj}) \in \mathbb{R}^d$ is a word embedding vector of dimension d for the j^{th} word in headline h_i , and $b_i = \{f(w_{b1}), f(w_{b2}), \dots, f(w_{bm})\}$ where $\forall k, f(w_{bk}) \in \mathbb{R}^d$ is a word embedding vector of dimension

d for the k^{th} word in body content b_i . We experiment with pre-trained GLOVE embeddings for the evaluation.

C. Synthetic Headline Generation

As discussed in the introduction, the semantic similarity between the headline and news body content is the most significant incongruence indicator. Text similarity schemes work well in case of short texts. However, the major bottleneck to compute such a similarity measure is the news body’s text length. For this reason, the methods which use a text similarity based approach to detect incongruence perform poorly. Generative adversarial networks (GANs) are recently getting much traction for various application scenarios, from the generation of realistic artificial images to compositions of meaningful poems. We utilize GANs based text generation techniques to generate a short synthetic headline for each of the news items using respective news body contents. The resulting synthetic headline is a low dimension contextual representation of the news body content.

1) *Generative Adversarial Network:* A typical generative adversarial network [22] comprises two neural nets, a discriminator and a generator. Both of these sub neural nets compete

with each other. The generator network tries to maximize the classification error, while the discriminator network tries to minimize it. Both the generator and discriminator converge together and reach an equilibrium state.

2) *Sequence Generative Adversarial Nets (SeqGAN)*: SeqGAN[21] is a sequence generation method based on generative adversarial networks. The key difference between standard GANs and SeqGAN is that SeqGAN bypasses the generator differentiation problem associated with original GANs by directly performing a gradient policy update using a reinforcement learning (RL) based approach. In the θ -parameterized generative model G_θ and start state s_0 , the expected end reward for output y_1 is defined as:

$$J(\theta) = E[R_T | s_0, \theta] = \sum_{y_1 \in Y} G_\theta(y_1 | s_0) \cdot Q_{D_\phi}^{G_\theta}(s_0, y_1), \quad (1)$$

where $Q_{D_\phi}^{G_\theta}(s_0, y_1)$ is an action-value function. A recurrent neural network (RNN) is used for the generative model for sequences, and a convolutional neural network (CNN) is used as the discriminator. We refer to the respective paper for more details.

3) *Stylized Headline Generation (SHG)*: Since incongruent headlines and click-baits usually follow a certain catchy and deceptive writing style, it can be advantageous to generate synthetic headlines that mimic these similar writing styles. Authors of very recent work SHG[8] propose a style transfer based headline generation approach, which uses a generative adversarial network with style discriminator. A gated recurrent unit (GRU) based recurrent neural network (RNN) is used as a generator that tries to minimize the following negative log-likelihood:

$$L_G(\theta_G) = E_{(x,h) \in S} [-\log_{PG}(h|y^L, z)]. \quad (2)$$

Three variants of discriminators are used, one for distinguishing the styles of the original and generated headline, a second one to maintain the aligned distributions in both the original and generated headlines, and a third for making sure that headline and body pairs are correctly classified. Again, we refer to the respective paper for more details.

D. Inter-mutual Attentive Semantic Matching

We now discuss the proposed semantic sentence matching scheme in detail. Our key idea is to model the relationships between all possible pairs of the words occurring in both the sentences, which captures the inherent semantic similarity between the sentences. We use pretrained word embeddings vectors and apply a novel inter mutual attention technique to model the similarity or dissimilarity relationship between pair of sentences.

1) *Word to Word Inter Mutual Attention*: We compute mutual attention scores between a pair of sentences, of which first $h_i^o = \{f(w_{h^o_1}), f(w_{h^o_2}), \dots, f(w_{h^o_l})\}$ where $\forall j, f(w_{h^o_j}) \in \mathbb{R}^d$ is the original headline of the news item, represented in terms of a sequence of word embeddings, and second $h_i^s = \{f(w_{h^s_1}), f(w_{h^s_2}), \dots, f(w_{h^s_p})\}$ where $\forall j, f(w_{h^s_j}) \in \mathbb{R}^d$ is the synthetically generated headline for news item n_i ,

represented in terms of a sequence of word embeddings. Here, l and p are lengths of the original headline h_i^o and the synthetic headline h_i^s , respectively.

First of all, we compute the difference between word embedding vectors of each candidate word pair W_q, W_r . Candidate pairs are formed by selecting all possible combinations of the inter sentence word pairs, such as W_q, W_r , where W_q and W_r are the q^{th} word of the original headline h_i^o and the r^{th} word of the synthetic headline h_i^s , respectively. Now we use a dense layer to project the difference of candidate embedding pairs into a scalar score:

$$C_{qr} = \theta_{diff}([f(w_{h^o_q}) - f(w_{h^s_r})]) + b_{diff}, \quad (3)$$

where $\theta_{diff} \in \mathbb{R}^{d \times 1}$ is a weight matrix and $b_{diff} \in \mathbb{R}$ a bias term. The score matrix $C = (C_{qr})$ is of dimension $l \times p$. We use row-wise avg-pooling and apply softmax to compute inter mutual attention scores for the original headline:

$$A^o = \text{Softmax}(\text{avg}_{row} C), \quad (4)$$

where A^o is the learned inter-mutual attention weight vector for original headline. We use column-wise avg-pooling and apply softmax to compute inter mutual attention scores for synthetic headline.

$$A^s = \text{Softmax}(\text{avg}_{col} C), \quad (5)$$

where A^s is the learned inter-mutual attention weight vector for synthetic headline. Subsequently, inter-mutual attended representations for original headline M_{A^o} and for synthetic headline M_{A^s} can be computed as:

$$\begin{aligned} M_{A^o} &= \sum_{i=1}^l f(w_{h^o_i}) A^o_i \\ M_{A^s} &= \sum_{i=1}^p f(w_{h^s_i}) A^s_i, \end{aligned} \quad (6)$$

where $f(w_{h^o_i})$ and $f(w_{h^s_i})$ are the word embeddings of the i^{th} word of the original and synthetic headline, respectively. Now, we compute the overall inter-mutual attended representation M_A as:

$$M_A = M_{A^o} + M_{A^s} \quad (7)$$

2) *Variants of Inter Mutual Attention*: Although the MuSeM model with difference of word embeddings of candidate word pair works well, we also investigate and experiment with two other versions of MuSeM that use the dot product of candidate word embedding pairs and concatenation of word embedding pairs, similar to [23]. We notice that the dot-product and concatenation variants do not perform better than the difference-oriented variant in our experiments. We also combine all three operations, namely difference, dot-product, and concatenation, which performs slightly better than the purely difference oriented model. Attention scores for the dot-product and concatenation oriented models can be computed

in a very similar fashion to the difference model, except for equation (3), which needs to be replaced by:

$$\begin{aligned} C_{qr} &= \theta_{dot}([f(w_{h^oq}) \cdot f(w_{h^sr})]) + b_{dot} \\ C_{qr} &= \theta_{con}([f(w_{h^oq}) \parallel f(w_{h^sr})]) + b_{con}. \end{aligned} \quad (8)$$

For the sake of brevity, we do not repeat all the other equations.

3) *Clubbed Model*: We also experiment with a combined model in which all three variants namely difference based, dot-product based and concatenation based methods are used in parallel and resulted overall embedding representation is used for computing attention scores.

$$\begin{aligned} F_{qr}^{dot} &= [f(w_{h^oq}) \cdot f(w_{h^sr})] \\ F_{qr}^{con} &= [f(w_{h^oq}) \parallel f(w_{h^sr})] \\ F_{qr}^{diff} &= [f(w_{h^oq}) - f(w_{h^sr})] \\ C_{qr} &= \theta_{dpc}([F_{qr}^{dot} \parallel F_{qr}^{con} \parallel F_{qr}^{diff}]) + b_{dpc}, \end{aligned} \quad (9)$$

where C_{qr} is overall attention score, θ_{dpc} is the weight matrix and b_{dpc} is the bias term for the clubbed model.

E. LSTM based Sequence Encoder

We concatenate the sequence of word embeddings of the words of the original headlines and words of the synthetically generated headline and use long short-term memory unit (LSTM) [24] to encode this overall sequence by using the standard LSTM equations as follows:

$$\begin{aligned} f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\ i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\ \tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \\ C_t &= f_t * C_{t-1} + i_t * \tilde{C}_t \\ o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \\ h_t &= o_t * \tanh(C_t), \end{aligned} \quad (10)$$

where h_{t-1} is previous hidden state and x_t is the current input. We use the last hidden state of LSTM as the encoded representation of the overall sequence M_E .

F. Classification

We learn a joint representation M of inter-mutual attended representation M_A and encoded representation of the overall sequence M_E using a nonlinear transformation layer with ReLU activation. Subsequently, we use a softmax layer to predict the label:

$$\begin{aligned} M &= \text{ReLU}(W_t([M_A, M_E] + b_t)) \\ \hat{y} &= \text{Softmax}(W_{cl}M + b_{cl}), \end{aligned} \quad (11)$$

where W_{cl} , b_{cl} , and $\hat{y}$ are the weight matrix, the bias term, and the predicted label, respectively. We use softmax cross-entropy with logits as loss L :

$$L = -\sum_{i=1}^m \log \frac{e^{w_{y_i}^T x_i + b_{y_i}}}{\sum_{j=1}^n e^{w_j^T x_i + b_j}}. \quad (12)$$

IV. EXPERIMENTAL SETUP

We implement the proposed models within the TensorFlow framework. For the evaluation and comparison, we use Macro F1 and AUC scores as evaluation metrics. All the parameters are tuned using a grid search. As a result, we keep learning rate as 0.001, batch size as 100, hidden states of LSTM as 100, and pretrained GloVe[25] embeddings of 300 dimensions. We use softmax cross-entropy with logits as the loss function, maximum sentence length as 50, dropout rate as 0.2, and the number of epochs as 10. For the synthetic headline generation, we use the default parameter values used in the papers mentioned above.

A. Dataset Statistics

We use two publicly available datasets, NELA17³ and Click-bait Challenge⁴ for the evaluation and comparison of the proposed model with the baseline methods. The NELA17 dataset is provided by [7]. Although they do not evaluate their model with it, they provide a script⁵ to generate the dataset from an original news collection dataset. It contains 91042 news items in total, of which 45521 news items are congruent, and 45521 news items are non-congruent. The Click-bait Challenge dataset is a collection of social media posts, which are annotated as click-bait or non-click-bait using a crowd-sourcing platform via majority voting. It contains 21033 social media posts in total, of which 16150 posts are congruent, and 4883 posts are non-congruent.

TABLE I: Dataset Statistics

Statistics	NELA17
Total	91042
Non-congruent	45521
Congruent	45521

Statistics	Click-bait Challenge
Total	21033
Non-congruent	4883
Congruent	16150

The NELA17 dataset is balanced but the Click-bait Challenge dataset has class imbalance problem. We use the `class_weight`⁶ parameter to adjust the imbalance in the dataset.

V. EVALUATION AND DISCUSSION

In this section, we compare the performance of the proposed models with state of the art methods and baseline models on two publicly available datasets. Subsequently, we also provide some explanations on the differences in the performance of the analyzed models. We also discuss a potentially useful trick from [11], which can improve the incongruence detection accuracy by capturing the newly uncovered cues.

³<https://github.com/BenjaminDHome/NELA2017-Dataset-v1>

⁴<http://www.clickbait-challenge.org/>

⁵<https://github.com/sugoiiii/detecting-incongruity-dataset-gen>

⁶https://www.tensorflow.org/tutorials/structured_data/imbalanced_data.

Fig. 3: Clubbed Model Architecture

We compare the proposed model with these baselines:

- **SVM:**[26] In this method, we use handcrafted linguistic and other statistical features to learn a classifier with support vector machine technique.
- **LSTM:**[24] In this method, we use long short term memory unit to encode both headline and body pair and apply softmax for the classification.
- **Hi-LSTM:**[27] This technique uses an LSTM based hierarchical encoder, which first encodes words and then uses another LSTM encoder to form the sentence representations.
- **Yoon:**[7] This is a state of the art, hierarchical dual encoder based model which uses headline guided attention to learn the contextual representation.

We experiment with four variants of the proposed MuSeM model.

- **MuSeM_diff_SeqGAN:** This is the MuSeM model with the difference between word embeddings, and synthetic headlines are generated using SeqGAN method.
- **MuSeM_dpc_SeqGAN:** This is the MuSeM model which uses a combination of all three operators, namely difference, concatenation, and dot-product between word embeddings and synthetic headlines are generated using SeqGAN method.
- **MuSeM_diff_SHG:** This is the MuSeM model with the difference between word embeddings, and synthetic headlines are generated using SHG method.
- **MuSeM_dpc_SHG:** This is the MuSeM model which uses a combination of all three operators, namely difference, concatenation, and dot-product between word embeddings and synthetic headlines are generated using SHG method.

TABLE II: Comparison of the proposed models with various state of the art baseline models for the NELA17 Dataset.

NELA17 Dataset		
Model	Macro F1	AUC.
SVM	0.622	0.637
LSTM	0.642	0.663
HiLSTM	0.651	0.672
Yoon	0.685	0.697
MuSeM_diff_SeqGAN	0.713	0.720
MuSeM_dpc_SeqGAN	0.719	0.727
MuSeM_diff_SHG	0.740	0.753
MuSeM_dpc_SHG	0.752	0.769

TABLE III: Comparison of the proposed models with various state of the art baseline models for the Click-bait challenge 2017 Dataset.

Click-bait challenge 2017 Dataset		
Model	Macro F1	AUC.
SVM	0.618	0.629
LSTM	0.630	0.641
HiLSTM	0.642	0.656
Yoon	0.660	0.678
MuSeM_diff_SeqGAN	0.677	0.683
MuSeM_dpc_SeqGAN	0.690	0.698
MuSeM_diff_SHG	0.729	0.734
MuSeM_dpc_SHG	0.735	0.747

A. Results for NELA17 Dataset

In the case of the NELA17 dataset, the SVM model with linguistic and statistical features achieves 0.622 and 0.637 in terms of Macro F1 and AUC, respectively. The LSTM model outperforms SVM with Macro F1 as 0.642 and AUC as 0.663. This gain can be attributed to the suitability of LSTM to learn the contextual representation of text sequences. The Hierarchical LSTM (Hi-LSTM) performs slightly better than simple LSTM with Macro F1 as 0.651 and AUC as 0.672. The probable reason for this improvement is the better representation learned by H-LSTM, in the form of the documents' hierarchical structure. Since the Yoon model uses a dual hierarchical encoder, which encodes words and paragraphs of the new body text separately using an attention mechanism guided by news headline, it outperforms the Hi-LSTM with significant gains. The Yoon⁷ model works well for long texts as it selects important paragraphs from the long body text, which reduces the effective size of the document. On the other hand, SVM, LSTM, and Hi-LSTM models do not scale well for long text sequences.

We compare the baselines and state-of-the-art models with four variants of the proposed model MuSeM. The MuSeM_diff_SeqGAN model performs better than the Yoon model with 0.713 and 0.720 in terms of Macro F1 and AUC. This performance improvement can be credited to low dimension representation of the news body content in the form of synthetic headline and inter-mutual attention based semantic matching. Although the headline guided attention to select relevant paragraphs in the Yoon model reduces the effective document length, the resultant document representation is still not of very low dimension. In contrast, the MuSeM model uses a very low dimensional representation of news body content in the form of a synthetic headline, which is more effective in semantic matching.

MuSeM captures both compositional aspects and latent cues related to similarity and dissimilarity relationships among the words using LSTM based sequence encoder and inter-mutual attention based semantic matching. The MuSeM_dpc_SeqGAN model performs slightly better than MuSeM_diff_SeqGAN due to additional patterns pertaining to similarity relationship among words, captured by the combination of all three variants, i.e., difference based, dot-product based, and concatenation based methods. The MuSeM_diff_SHG model, which uses a style transfer based headline generation method, outperforms both variants of the MuSeM models with SeqGAN by a large margin. This performance improvement can be attributed to a better headline generation by stylized headline generation (SHG) method. The MuSeM_dpc_SHG model outperforms all the other models by achieving 0.752 and 0.769 in Macro F1 and AUC, respectively.

B. Results for Click-bait Challenge Dataset

In the case of the Click-bait Challenge dataset, we observe very similar trends as with the NELA17 dataset. The deep

learning based methods perform significantly better than the handcrafted feature based SVM model. With 0.660 and 0.678 in Macro F1 and AUC, the Yoon model outperforms both the LSTM and Hi-LSTM models. All variants of the MuSeM model perform better than all the baseline methods, whereas MuSeM_dpc_SHG model achieves 0.735 and 0.747 in terms of Macro F1 and AUC, beating all the other variants.

C. Higher Order Inter-mutual Attention

Authors in [11] propose a variant of mutual attention which can be used to model higher-order relationship among the candidate words. Essentially, the proposed inter-mutual attention only models the relationship between two words at a time individually, and the presence of other words in the sentences are not taken into account. If we can also reckon the presence of other words during the computation of the inter-mutual attention, we can capture the additional contextual cues related to similarity or dissimilarity relationship. The higher-order mutual attention trick allows us to model such relationships.

There are certain challenges to be taken care of before the application of higher order mutual attention to incongruence detection problem. Firstly, the original higher order mutual attention is designed for a single sequence as it uses intra-mutual attention within a sequence, while headline incongruence detection involves two sentences or sequences. Secondly, higher order mutual attention is achieved by applying a square or cube of the attention score matrix, which is possible only if the original score matrix is a square matrix. In the headline incongruence detection task, the score matrix may or may not be a square matrix because original and synthetically generated headlines may have different lengths.

VI. CONCLUSION

This paper proposes inter-mutual attention based semantic matching (MuSeM) to detect incongruence in news headlines. We also investigate a different variant of MuSeM, which combines three operations on word embedding pairs to compute inter-mutual attention scores. The proposed models outperform all the baselines in experiments with two publicly available datasets. We notice that the performance of inter-mutual attention based semantic matching greatly depends on the accuracy of synthetic headline generation step.

In future research, we plan to use the higher order word to word attention trick used in [11] to model and capture the indirect relationships between the word pairs. We plan to investigate and devise an end to end version of MuSeM model, where synthetic headline generation and semantic matching steps are seamlessly integrated. We also plan to conduct an analysis of the efficacy of the attention mechanism by visualizing the attention weights.

REFERENCES

- [1] J. Reis, P. Olmo, R. Prates, H. Kwak, and J. An, "Breaking the News : First Impressions Matter on Online News," pp. 357–366.

⁷<https://github.com/david-yoon/detecting-incongruity>

- [2] M. Gabielkov, A. Ramachandran, A. Chaintreau, and A. Legout, "Social clicks: What and who gets read on twitter?" in *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*, ser. SIGMETRICS '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 179–192. [Online]. Available: <https://doi.org/10.1145/2896377.2901462>
- [3] W. Ferreira and A. Vlachos, "Emergent: a novel data-set for stance classification," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California: Association for Computational Linguistics, Jun. 2016, pp. 1163–1168. [Online]. Available: <https://www.aclweb.org/anthology/N16-1138>
- [4] Z. Wang, W. Hamza, and R. Florian, "Bilateral multi-perspective matching for natural language sentences," in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 2017, pp. 4144–4150. [Online]. Available: <https://doi.org/10.24963/ijcai.2017/579>
- [5] J. Blom and K. Reinecke Hansen, "Click bait: Forward-reference as lure in online news headlines," *Journal of Pragmatics*, vol. 76, pp. 87–100, Jan. 2015.
- [6] M. U. Rony, N. Hassan, and M. Yousuf, "BaitBuster : A Clickbait Identification Framework," pp. 8216–8217, 2017.
- [7] S. Yoon, K. Park, J. Shin, H. Lim, S. Won, M. Cha, and K. Jung, "Detecting incongruity between news headline and body text via a deep hierarchical encoder," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 791–800.
- [8] K. Shu, S. Wang, T. Le, D. Lee, and H. Liu, "Deep headline generation for clickbait detection," in *2018 IEEE International Conference on Data Mining (ICDM)*, 2018, pp. 467–476.
- [9] S. Chesney, M. Liakata, M. Poesio, and M. Purver, "Incongruent headlines: Yet another way to mislead your readers," in *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*. Copenhagen, Denmark: Association for Computational Linguistics, Sep. 2017, pp. 56–61. [Online]. Available: <https://www.aclweb.org/anthology/W17-4210>
- [10] Y. Tay, L. A. Tuan, S. C. Hui, and J. Su, "Reasoning with sarcasm by reading in-between," *ACL 2018 - 56th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, vol. 1, pp. 1010–1020, 2018.
- [11] R. Mishra, "Fake news detection using higher-order user to user mutual-attention progression in propagation paths," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
- [12] Y. Chen, N. J. Conroy, and V. L. Rubin, "Misleading online content: Recognizing clickbait as "false news"," in *Proceedings of the 2015 ACM on Workshop on Multimodal Deception Detection*, ser. WMDD '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 15–19. [Online]. Available: <https://doi.org/10.1145/2823465.2823467>
- [13] A. Chakraborty, B. Paranjape, S. Kakarla, and N. Ganguly, "Stop clickbait: Detecting and preventing clickbaits in online news media," in *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2016, pp. 9–16.
- [14] M. Potthast, S. Köppl, B. Stein, and M. Hagen, "Clickbait Detection," *ECIR'16*, no. 1, 2016.
- [15] W. Wei and X. Wan, "Learning to identify ambiguous and misleading news headlines," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, ser. IJCAI'17. AAAI Press, 2017, p. 4172–4178.
- [16] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *ICLR 2015 as oral presentation*.
- [17] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4396–4405.
- [18] S. Suwajanakorn, S. M. Seitz, and I. Kemelmacher-Shlizerman, "Synthesizing obama: Learning lip sync from audio," *ACM Trans. Graph.*, vol. 36, no. 4, Jul. 2017. [Online]. Available: <https://doi.org/10.1145/3072959.3073640>
- [19] K. Nagano, J. Seo, J. Xing, L. Wei, Z. Li, S. Saito, A. Agarwal, J. Fursund, and H. Li, "Pagan: Real-time avatars using dynamic textures," *ACM Trans. Graph.*, vol. 37, no. 6, Dec. 2018. [Online]. Available: <https://doi.org/10.1145/3272127.3275075>
- [20] J. Guo, S. Lu, H. Cai, W. Zhang, Y. Yu, and J. Wang, "Long Text Generation via Adversarial Training with Leaked Information," 2017.
- [21] L. Yu, W. Zhang, J. Wang, and Y. Yu, "Seqgan: Sequence generative adversarial nets with policy gradient," in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, ser. AAAI'17. AAAI Press, 2017, p. 2852–2858.
- [22] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 2672–2680. [Online]. Available: <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
- [23] Y. Tay, L. A. Tuan, and S. C. Hui, "Compare, compress and propagate: Enhancing neural architectures with alignment factorization for natural language inference," in *Proceedings of EMNLP 2018*, 2018.
- [24] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, p. 1735–1780, Nov. 1997. [Online]. Available: <https://doi.org/10.1162/neco.1997.9.8.1735>
- [25] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *In EMNLP*, 2014.
- [26] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, p. 273–297, Sep. 1995. [Online]. Available: <https://doi.org/10.1023/A:1022627411411>
- [27] D. Tang, B. Qin, and T. Liu, "Document modeling with gated recurrent neural network for sentiment classification," in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, Sep. 2015, pp. 1422–1432.