Deep Multimodal Image-Repurposing Detection

Ekraam Sabir

USC Information Sciences Institute
Marina Del Rey, CA
esabir@isi.edu

Yue Wu

USC Information Sciences Institute
Marina Del Rey, CA
yue_wu@isi.edu

Wael AbdAlmageed

USC Information Sciences Institute
Marina Del Rey, CA
wamageed@isi.edu

Prem Natarajan

USC Information Sciences Institute
Marina Del Rey, CA
pnataraj@isi.edu

ABSTRACT

Nefarious actors on social media and other platforms often spread rumors and falsehoods through images whose metadata (e.g., captions) have been modified to provide visual substantiation of the rumor/falsehood. This type of modification is referred to as image repurposing, in which often an unmanipulated image is published along with incorrect or manipulated metadata to serve the actor's ulterior motives. We present the Multimodal Entity Image Repurposing (MEIR) dataset, a substantially challenging dataset over that which has been previously available to support research into image repurposing detection. The new dataset includes location, person, and organization manipulations on real-world data sourced from Flickr. We also present a novel, end-to-end, deep multimodal learning model for assessing the integrity of an image by combining information extracted from the image with related information from a knowledge base. The proposed method is compared against state-of-the-art techniques on existing datasets as well as MEIR, where it outperforms existing methods across the board, with AUC improvement up to 0.23.

CCS CONCEPTS

• **Human-centered computing** → **Social media**; • **Computing methodologies** → **Natural language processing**; **Computer vision**; **Multi-task learning**; • **Information systems** → **Multi-media information systems**; **Multimedia and multimodal retrieval**;

KEYWORDS

Rumor detection, fake news, computer vision, deep learning, multi-task learning

ACM Reference Format:

Ekraam Sabir, Wael AbdAlmageed, Yue Wu, and Prem Natarajan. 2018. Deep Multimodal Image-Repurposing Detection. In *2018 ACM Multimedia Conference (MM '18), October 22–26, 2018, Seoul, Republic of Korea*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3240508.3240707>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.
MM '18, October 22–26, 2018, Seoul, Republic of Korea
© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-5665-7/18/10...\$15.00.
<https://doi.org/10.1145/3240508.3240707>

1 INTRODUCTION

Social media is increasingly becoming the news source of choice, replacing conventional news outlets. In [14], Marchi *et al.* present a survey, showing that teens are increasingly inclined towards evolving news articles over conventional and objective renderings of journalism. However, the unregulated use of social media makes it convenient to float rumors or present misleading information. Reasons for spreading misleading information can vary with possible malicious intent, political motivation or simply due to ignorance, as shown in Figure 1. The left-most example is relatively easy for an educated audience to identify as fake. The right three examples, however, highlight the difficulty of detecting manipulated image data when the details being manipulated are subtle, subliminal and often interleaved with the truth. The right three images each have one entity manipulated in the associated caption, while the image remains untampered. In Figure 1(b) the deceased rapper Tupac Shakur is incorrectly referenced. Figure 1(c) has the Democratic National Convention referenced instead of the Ku Klux Klan, while the image in Figure 1(d) is from a protest in Chile instead of North Dakota. It is noteworthy that the manipulated details in the right three examples are semantically coherent and leave no digital footprint in the image itself, making their detection difficult.

In the broad category of manipulation detection, the research community has invested significantly in identifying digital manipulations of images [17], which involves identifying whether an image has been tampered with at the pixel and/or object level. However, cases of image repurposing—in which the image is authentic but the accompanying metadata has been manipulated—have not been thoroughly studied. The right three images in Figure 1 are examples of image repurposing.

There are crowd-sourced attempts to vet suspicious information spreading on social media. Facebook has recently launched a reporting and fact checking feature on its platform. Snopes¹ is another fact checking website which debunks hoaxes and fake news. However the reach of these efforts is limited and requires significant human time, resource, motivation and skill. Large-scale detection of image repurposing necessitates the development of automatic systems that could vet the credibility of an image from internet sources itself.

In this paper we attempt to address these challenges by developing a deep multimodal system for image repurposing detection. To the best of our knowledge, there has been one previous attempt

¹<https://www.snopes.com/>

Figure 1: Examples of real-world manipulated information. Image (a) is an obvious hoax with digital image manipulation and is less likely to fool people. Image (b) falsely claims that Tupac Shakur is alive. In image (c), Democratic Party is incorrectly referenced. Image (d) misattributes the photograph to a protest at the North Dakota Pipeline, when it was originally taken in Chile. Images (b), (c) and (d) are not digitally manipulated, but the details accompanying them are, making detection more challenging.

at image repurposing detection in literature [9], in which the authors pose the problem as a semantic integrity assessment problem, where the semantics are potentially inconsistent between images and captions in a query *package*. The package is defined as an image and associated text caption. However, semantically inconsistent manipulations are farther from real-world examples as demonstrated in Figure 1. We use the definition of the package introduced in [9]. Meanwhile, we also extend it to include the global positioning system (GPS) information included in the image metadata. We call this a *multimodal image package*.

The contributions of this paper are:

- (1) A new and challenging dataset with relatively intelligent, labeled and semantically coherent manipulations over what is currently available for research.²
- (2) A novel deep multimodal model that uses multitask learning for image repurposing detection

The remainder of this paper is organized as follows: Section 2 discusses related works. Section 3 discusses the dataset and procedure involved in creating it. Section 4 describes our approach to image repurposing detection and the model, and Section 5 gives experimentation details and describes baselines.

2 RELATED WORK

In [25], Zampoglou *et al.* identify the importance of vetting images and accompanying information for journalistic review purposes. They rely on existing research in developing a graphical user interface (GUI) which can be used to verify the content of information found on the web. However, they are limited to discovering digital manipulations of images, which has been the focus of the research community [22][24][17][1]. Digital manipulations are mainly divided into copy-move and image-splicing problems. Copy-move involves duplicating objects within the same image, while splicing inserts objects from a foreign image. In [1], Asghar *et al.* broadly classify digital image manipulation detection techniques into active and passive (blind) methods. Active methods involve a signature or watermark which has to be extracted from the image. The blind or

passive approach looks for inconsistencies and artifacts in the image being investigated, which might result from digital manipulations.

To the best of our knowledge, there has been a single attempt at image repurposing detection—by Jaiswal *et al.* in [9]. The authors present the concept of a multimedia package, which in theory contains an untampered image and related metadata, but is limited to image caption pairs in their paper. They induce manipulations in packages by replacing captions from random packages. They also use an unlabeled knowledge base which is assumed to comprise untampered multimedia packages. The overall method involves learning joint representations from the knowledge base and performing anomaly detection on representations of query packages. The motivation for this approach is that tampering in query packages will manifest itself as an outlier compared to distribution of representations from the knowledge base.

The approach in the above formulation is useful for laying a foundation in image repurposing detection. However, it suffers from some issues. First, manipulations are semantically incorrect across images and captions, making them relatively easier to detect. Putting it in context, an image with an elephant and a manipulated caption describing a dining table are two widely different object classes and therefore relatively easy to discern. Second, as seen from examples in Figure 1, actual image repurposing involves manipulating a detail which is semantically coherent, but misrepresents the image. Manipulations involving swapped captions in [9] are farther from such real-world examples. Third, the knowledge base provided does not contain information related to the query package. It is useful in an anomaly detection framework, but is not adaptable when related information is crucial for image repurposing detection. We address some of these issues in this work.

Ruchansky *et al.* [18] provide an excellent summary of rumor detection literature. The work by Jin *et al.* in [10] tackles rumor detection and continues a similar body of work [11][12]. It performs rumor detection on image caption pairs from Weibo and Twitter. The dataset is described by Boididou *et al.* in [3]. However, it differs from our work in two key aspects. First, spliced and photoshopped images are included in their dataset, which have already been studied. They are qualitatively closer to Figure 1(a) and are easier to detect for an informed audience. Second, the attention-based model

²<https://github.com/Ekraam/MEIR>

	Reference Dataset		Manipulated Dataset
	Related Package	Manipulation Source Package	Repurposed Query Package
Image
Text	Ex Convento at Cuilapan De Guerrero completed in 1555	The Dorena Bridge has these neat windows , which kind makes you feel like you are in a house !	Ex Convento at Dorena Bridge completed in 1555
Location	Mexico, Cuilapam de Guerrero, Oaxaca, Cuilapam de Guerrero. GPS: 16.992657, -96.779133	United States, Lane, Oregon, Row River. GPS: 43.737553, -122.883646	United States, Lane, Oregon, Row River. GPS: 43.737553, -122.883646
Image
Text	The 2001 Space pod modelled in Lego by Dilip	" It was a dream , and in dreams you have no choices: ..." – Neil Gaiman , American Gods (2001)	Space pod modelled in Lego by Neil Gaiman
Location	United Kingdom, West Sussex, England, Billingshurst. GPS: 51.024538, -0.450281	United States, Riverside, California, Moreno Valley. GPS: 33.900035, -117.254384	United Kingdom, West Sussex, England, Billingshurst. GPS: 51.024538, -0.450276
Image
Text	The International Rally of Queensland Sunday competition images + podium	\$10 , 000 , 000 00 is the estimated value for the first United States gold coin Numismatic Guaranty Corporation special display ...	The Numismatic Guaranty Corporation Sunday competition images + podium
Location	Australia, Gympie Regional, Queensland, Imbil. GPS: -26.460497, -152.679491	United States, Philadelphia, Pennsylvania, Philadelphia. GPS: 39.953333, -75.156667	Australia, Gympie Regional, Queensland, Imbil. GPS: -26.461131, 152.678117

Table 1: Examples of location, person and organization manipulations from the Multimodal Entity Image Repurposing (MEIR) dataset. A manipulation source package provides a foreign entity which is used to replace corresponding entities from the manipulated package. Text in red indicates manipulation. Longer text descriptions have been truncated.

presented in their work relies solely on the information inconsistency visible in a query package to detect fake news. Similar to the setup in [9], this approach does not account for semantically coherent manipulations within query information.

3 MULTIMODAL ENTITY IMAGE REPURPOSING DATASET (MEIR)

Jaiswal *et al.* [9] presented the Multimodal Information Manipulation (MAIM) dataset of multimedia packages, in which each package contains an image-caption pair. In order to mimic image repurposing, associated captions were swapped with randomly chosen captions—potentially leading to semantically inconsistent repurposing, such as an image of a cat captioned as a car. The MAIM dataset was the first attempt to create a dataset for image

repurposing detection research. However, MAIM had several limitations. First, the resulting manipulations are often easy to detect and unlikely to fool a human. Second, the reference dataset provided in MAIM is unlabeled and does not contain entities directly related to manipulated packages. Therefore, there is no explicit use of reference dataset when detecting manipulations. MAIM dataset does not reflect difficult manipulations that are created by carefully altering caption contents to produce semantically consistent packages—for example, leaving the caption unaltered except for replacing the name of an entity.

In order to address these limitations, we introduce the Multimodal Entity Image Repurposing Dataset (MEIR), which compared to MAIM, contains semantically intelligent manipulations. Multimedia packages in MEIR are manipulated by replacing named entities

of corresponding type, instead of entire captions. We consider three named entities: person, organization and location. Unlike MAIM, the reference dataset in MEIR has directly related entities to manipulated packages. It is also labeled to identify this relationship. The process for bringing these improvements in a new dataset is described in the following.

Creating MEIR consisted of three stages—data curation, clustering packages by relatedness, and creating manipulations in a fraction of packages of each cluster. MEIR was created from Flickr images and their associated metadata, collected in the form of packages. To ensure uniformity, packages with missing modalities and non-English text were removed. The remaining packages were pre-processed to remove noise in the form of html tags. Following data collection and preprocessing, packages were clustered by relatedness. Clustering by relatedness helps allocate packages to reference, training, development and test sets, such that related packages are labeled and present in the reference dataset when detecting manipulations. Identifying relatedness before synthesizing manipulations also helps to avoid replacing entities between two related packages. Two packages are considered related when they are geographically located close to each other in terms of real-world location, and the image and text modalities between them are similar. Clustering is done in two stages: clustering by location followed by refining the clusters based on image and text similarities between packages. Initial clusters of packages with neighboring global positioning system (GPS) coordinates were created by measuring the proximity of GPS coordinates up to two decimal places of accuracy. This limited the maximum distance between two related packages to approximately 1.3 kilometers. A relaxed distance constraint between two packages was kept to ensure better recall in the first stage of clustering. In the second stage, clusters were refined based on image and text similarity between packages. Similarity was measured by scoring image and text feature representations of two packages. Feature representations used were VGG19 [21], features pretrained on imagenet and averaged word2vec [15] embeddings for image and text respectively. Cosine similarity was used to measure similarity. Selection of a suitable threshold on the similarity score was based on a sample of 200 package pairs which was manually annotated for relatedness.

In order to create manipulations, we applied named entity recognition for identifying and replacing entities between packages. We used StanfordNER's [7] three-class model for identifying person name, location and organization entities in each package. Entities of corresponding type were replaced between two packages from unrelated clusters. All instances of that entity type were manipulated to "cover-up" potential inconsistencies within a manipulated package. As an example, when making a location manipulation of *London* to *Paris*, any mention of *England* should also be manipulated along with the GPS coordinates. Approximately, a quarter of each cluster is manipulated. The reference dataset is allocated to approximately half of the packages from each cluster, such that each package is unmanipulated. The remaining half of each cluster is an equal mix of manipulated and unmanipulated packages and is used to create the train-test-validation split. The reference dataset is common to train, test and validation sets.

The reference dataset contains 82,156 untampered related packages. There are 57,940 packages in the manipulated dataset out of

Modality	Manipulated	Unmanipulated	Overall
GPS coordinates	47.0%	95.0%	71.0%
Image	65.9%	65.6%	65.8%
Text	60.0%	77.2%	68.6%
Image+GPS+Text	66.1%	93.6%	79.9%

Table 2: Top-1 retrieval accuracy of related packages when a query package is manipulated or unmanipulated. Using all modalities gives the best performance for retrieval.

Figure 2: Average feature importance of each modality with varying feature dimension. All modalities are found to contribute towards manipulation detection.

which 28,976 packages are split into 14,397 location, 7,337 person and 7,242 organization manipulations. The manipulated dataset is divided into 40,940, 7,000 and 10,000 packages for training, validation and test, respectively. Each package consists of image, location and text modalities. Location comprises country, county, region and locality names, in addition to GPS coordinates. Text is the user generated description related to the image. The dataset also contains an additional modality—time-stamp which contains the date and time associated with each package. Packages are spread across 106 countries with English-speaking countries contributing the majority.

Examples from MEIR are presented in Table 1 showing location, person and organization manipulation. It is important to note that a human is unlikely to see through these manipulations. This understanding reinforces the need for a reference dataset containing directly related information.

4 IMAGE REPURPOSING DETECTION

Manipulation detection systems often retrieve one or more related images or packages from a reference dataset first, followed by comparing the query package to the retrieved one to determine likelihood of manipulation.

Since different modalities in a given package (e.g., image, text, location, etc.) contribute differently to both retrieval and manipulation detection, we start by assessing the importance of each modality to guide the design of the manipulation detection approach.

Figure 3: Model overview. A potentially related package is retrieved from the reference dataset. Feature extraction and balancing modules create a concatenated feature vector. The package evaluation module consists of related and single package modules. All NN_i layers represent a single, dense, fully connected layer with ReLU activation.

4.1 Modality Importance

Image manipulators are assumed to leave most of the content of the package unchanged, and only change a small number of package modalities (e.g., only location) to make the manipulation subtle and hard to detect. Therefore, the query and retrieved packages are assumed to have largely overlapping information.

We use similarity scoring to retrieve a package from reference dataset $\mathbb{R}$ using every modality for every query package. Let f_{qm} and f_{rm} be the feature vectors for modality m in query q and reference r packages, respectively, and $s(.,.)$ is the similarity metric. The top retrieved package r^* is identified as shown in Equation 1.

$$r^* = \arg \max_{r \in \mathbb{R}} \sum_m s(f_{qm}, f_{rm}) \quad (1)$$

VGG19 [21] pretrained on imagenet [19] and averaged word2vec [15] features are used for images and text respectively. We use cosine distance to measure package similarities. Experimental results are presented in Table 2, which shows the top-1 retrieval results for both manipulated and unmanipulated partitions of the query subset of MEIR. The results illustrate that using all modalities for retrieving top-1 related package provides the best retrieval results.

Further, in order to estimate the importance of each modality for manipulation detection, we devise a classification experiment followed by feature importance measurement. The model comprises Gaussian random projection for dimensionality reduction followed by random forest [4] for classification. Random projections have been used previously for dimensionality reduction and in classification problems [2][6]. Each modality in query and related packages is reduced to a common feature dimension L , and features from all modalities are concatenated. A simple random forest classifier is trained on the resulting feature vector for manipulation detection. Feature importance of each dimension is measured using Gini impurity across all trees as described in [5] and implemented in

[16]. Averaged modality importance is shown in Figure 2. Each experiment configuration is repeated for 30 trials and averaged results are presented. As shown in Figure 2, all modalities contribute significantly to manipulation detection, at all feature dimensions.

4.2 Deep Manipulation Detection Model

Broadly speaking, there are two approaches for manipulation detection. The first approach depends on assessing the coherency of the content of the query package, without using any reference datasets, e.g., by matching caption against the image or vice versa [9]. The second approach assumes the existence of a relatively large reference dataset, and assesses the integrity of the query package by comparing it to one or more packages retrieved from the reference dataset. The main advantage of the second approach is when the manipulations are semantically consistent. The proposed method in this paper belongs to the second approach, since the information in the query package is potentially manipulated and requires external information for validation.

We propose a deep multimodal, multi-task learning model for image repurposing detection, as illustrated in Figure 3. The proposed model consists of four modules: (1) feature extraction, (2) feature balancing, (3) package evaluation and (4) integrity assessment. The model takes a query package and the top-1 related package, retrieved from a reference dataset, as discussed in Section 4.1.

Both query and reference dataset packages are assumed to contain image, text and GPS coordinates. Images are encoded using VGG19 [21] pretrained on imagenet [19]. GPS coordinates are a compact numerical representation of location; therefore we normalize them without any further processing. Finally, text is represented using averaged word2vec [15] features for all words in a sentence. We also explore performance improvement using an attention model over words instead of averaging word features.

Component	Combination					
Single Package	✓	✓	✓	✓	✓	✓
Related Package		✓	✓	✓	✓	✓
Multitask-loss			✓	✓	✓	✓
Feature Balancing				✓	✓	✓
Attention in Text					✓	✓
Learnable Forget Gate						✓
Metric	Scores					
F_1 tampered	0.58	0.59	0.61	0.76	0.80	0.83
F_1 clean	0.60	0.60	0.65	0.78	0.82	0.84
AUC	0.64	0.65	0.70	0.86	0.89	0.91

Table 3: This table justifies different design choices of the model. A ✓ indicates the presence of a component. Model architecture has been optimized on the development set.

Features from different modalities have widely different dimensionalities (4096D for image, 300D for text and 2D for location). As shown before in Section 4.1, varying the dimensionality of each modality does not significantly change its importance. In order to ensure features from all modalities are balanced, all features are transformed to 300D feature vectors (similar to word2vec text features) and concatenated into a single feature vector. Neural layers are used to transform feature dimensions.

The core of the proposed model is the package evaluation module, which consists of related package and single package sub-modules. As shown in Figure 3, the related package sub-module consists of two Siamese networks. The first network is a *relationship classifier* that verifies whether the query package and top-1 package are indeed related, while the second network is a *manipulation detection* network that determines whether the query package is a manipulated version of the top-1 retrieved package. Since manipulation detection is dependent on the relatedness of the two packages, the relationship classifier network controls a forget gate which scales the feature vector of the manipulation detection network down to zero if the two packages are unrelated. Two designs are considered for a forget gate. The first is a dot product between the output of *relationship classifier* and *manipulation detection feature vector*. This formulation of the forget gate is not learnable and scales all dimensions of the manipulation detection feature vector indiscriminately. An alternative is a learnable forget gate similar to LSTMs [8]. If x is the relationship classifier output, y_{inp} the manipulation detection feature vector, w the forget gate weight matrix and b the bias vector, then the output feature vector y_{out} is given by Equation 2 where $*$ is the Hadamard product and $\cdot$ is matrix multiplication. For a learnable forget gate, the feature vector of the relationship classifier is used as input. Choice of gate is justified ahead.

$$y_{out} = y_{inp} * \sigma(w \cdot x + b) \quad (2)$$

In the meantime, a single package module verifies the coherency (i.e., integrity) of the query package alone.³ The single package module is a feedforward network that takes the balanced feature vector of the query package as input and performs feature fusion to present a 100 dimensional feature vector.

³Similar to approaches that do not depend on a reference dataset for integrity assessment

Figure 4: Our baseline semantic retrieval system (SRS) retrieves similar concepts via each modality and uses an overlap between retrieved packages as an indicator for integrity score. Jacardian index is used to measure overlap.

Modalities	Method	MEIR		
		F_1 tampered	F_1 clean	AUC
Image + Text	VSM [9]	0.56	0.63	0.60
	Our Method	0.73	0.76	0.81
Image + Text + GPS	VSM [9]	0.63	0.59	0.65
	GCCA	0.75	0.71	0.81
	SRS	0.66	0.37	0.71
	Our Method	0.80	0.80	0.88

Table 4: We present results on MEIR. Experiments are performed in two settings—with image and text modalities, and image, location and text modalities. This is done for a fairer comparison with VSM which was originally designed for the image-text dataset.

To ensure the overall network behaves as intended, we use multi-task learning in the two submodules. Relationship and manipulation classifiers are trained with corresponding labels. To avoid conflict in training, the manipulation classifier in related package submodule has a third label *unknown* apart from *manipulated* and *unmanipulated*, when the retrieved package is unrelated. All neural layers in the package evaluation module are dense, fully connected 100 dimensional layers.

The integrity assessment module concatenates feature vectors from both related and single package modules for manipulation classification. A neural network classifier is trained on the concatenated feature vector with no hidden layers.

The model has multiple components that need to be validated before we evaluate it against other methods and baselines. As discussed previously, there are also choices involved with some components of the model—attention over averaged word2vec features for text feature extraction and learnable over non-learnable forget gate. These choices are tuned on the development set with a set of ablation experiments. The experiments add different components of the model sequentially and notice a performance increase justifying their addition. By default, averaged word2vec features and non-learnable forget gates are used, unless otherwise mentioned. Experimental results are shown in Table 3. A ✓ indicates the presence of a component. Improvement of 0.16 AUC from the feature

Method	MAIM			Flickr30K			MS COCO		
	F_1 tampered	F_1 clean	AUC	F_1 tampered	F_1 clean	AUC	F_1 tampered	F_1 clean	AUC
MAE [9]	0.49	0.49	-	0.49	0.50	-	0.5	0.48	-
BiDNN [9]	0.52	0.52	-	0.63	0.62	-	0.76	0.77	-
VSM [9]	0.75	0.77	-	0.89	0.88	-	0.94	0.94	-
Our Method - SPA	0.78	0.78	0.87	0.88	0.88	0.95	0.92	0.92	0.96

Table 5: We present results on Multimodal Information Manipulation (MAIM) dataset. MAIM has subjective image caption pairs. We train our method without Inter-package task module since there are no related packages to leverage in MAIM. We compare against methods presented in [9].

Missing Modality - Training	Missing Modality - Test			
	None	Image	Text	Location
None	0.88	0.77	0.66	0.63
Image	0.87	0.85	-	-
Text	0.87	-	0.75	-
Location	0.85	-	-	0.82
All	0.88	0.86	0.76	0.79

Table 6: Scores with missing modalities in retrieved packages. None and All refer to no missing and all missing modalities respectively. The training scheme improves performance for all missing modality scenarios.

balancing layer can be further broken down into 0.04 AUC from increased depth and 0.12 AUC from matching feature dimensions of different modalities. We hypothesize that dimensionality matching prevents modalities with higher dimension (e.g., image) from dominating over low dimension modalities (e.g., GPS). A model comprising all discussed components, with attention over text features and a learnable forget gate, gives the best performance. We use this configuration in all future evaluations.

The model is trained end to end. We use Adam optimizer in all examples with a learning rate of 0.001. All nonlinear transformations use the ReLU activation function. We use Keras with tensorflow backend. Undiscussed parameters are set to default values.

5 EXPERIMENTAL EVALUATION

Detection of image repurposing is a relatively new research area without established methods for evaluation. In this paper we propose using the area under receiver operating characteristic curve (AUC) and also F_1 scores for evaluating performance.

We compare the performance of our model against the state-of-the-art in [9]. The visual semantic model (VSM) is the best encoding model in the anomaly detection framework presented in [9]. Further, we also compare our performance against two baselines: the first is our baseline semantic retrieval system (SRS) shown in Figure 4. SRS retrieves similar packages corresponding to each modality using cosine distance. The output integrity score is the measured overlap between retrieved packages. Intuitively, this overlap indirectly measures whether modalities in a query package point to the same related packages in the reference dataset. Since each modality will retrieve similar packages, a rogue modality will retrieve packages pertaining to a different event from the reference dataset. If modalities in the query package are consistent with information

in the reference dataset, the overlap between retrieved packages will be significant. We use the Jaccard index for measuring overlap between retrieved packages. The second baseline approach uses generalized canonical correlation analysis (GCCA) [13] for feature embedding and random forest for classification. GCCA transforms multimodal features with varying feature dimensions into a common embedding space and has been used for classification with multimodal features [20][23].

We also explore the performance of our model without retrieving the top-1 package from a reference dataset. In this setting we remove the related package module. We call this version of our model the single package assessment (SPA) in all figures. This variation enables us to analyze the importance of related package retrieval and compare the performance on MEIR and MAIM datasets.

The results of our experiments on MEIR are presented in Table 4, which illustrates that our model performs better than the baselines and VSM [9]. It should be noted that VSM was originally designed for image and text modalities. We extend it with location modality. We also evaluate our model with image and text modality for a complete comparison with VSM. Under both circumstances our model performs better. The extended version of VSM with location modality performs better than the original version. Table 7 shows examples of location, person and organization manipulation from the test set which our model identifies correctly.

Jaiswal et al. [9] evaluate VSM and other methods on three datasets—MAIM, MSCOCO and Flickr30K. MAIM has subjective image caption pairs and is more challenging compared to Flickr30K and MSCOCO, which have objective image caption pairs. However, these datasets do not have related content in the reference dataset. This does not help the related package module in our dataset. We therefore compare the single package assessment (SPA) version of our model without retrieving any query package. They use a representation learning and outlier detection method for comparison. We compare against all encoding methods shown in the referenced paper. They do not provide AUC scores for their method. In Table 5, SPA gives superior performance on MAIM and competitive performance on Flickr30K and MS COCO.

A reference dataset may not have all modalities present, which makes missing modalities from retrieved packages an additional problem of interest. Without modifying model architecture, the problem can be tackled by improving the training scheme. During training a modality is deliberately removed from some retrieved samples and represented as a zero vector. Table 6 shows results of training with a 20% missing modality. At test time, the corresponding modality is completely removed from all samples. The

	Reference Dataset		Manipulated Dataset
	Related Package	Manipulation Source Package	Repurposed Query Package
Image
Text	Sommerpalast Suzhou Street	Grand Beach , Lake Winnipeg , Manitoba , Canada	Canada
Location	China, Beijing, Beijing, Beijing. GPS: 40.000826, 116.268825	Canada, Manitoba, Manitoba, Grand Beach. GPS: 50.562316, -96.614059	Canada, Manitoba, Manitoba, Grand Beach. GPS: 50.562316, -96.614059
Image
Text	Live Life Awards Image by Sean McGrath Photography	Film Effects - The Vault CartLike many effects seen in the films ... Ollivanders wand shop , Flourish and Blotts , the ...	Live Life Awards Image by Blotts Photography
Location	Canada, Saint John, New Brunswick, Saint John. GPS: 45.271428, -66.061981	United Kingdom, Hertfordshire, England, Watford. GPS: 51.693761, -0.422329	Canada, Saint John, New Brunswick, Saint John. GPS: 45.271428, -66.061981
Image
Text	1st Combined Convention 28th Annual National DeSoto Club Convention & 44th Annual Walter P Chrysler Club ConventionJuly 17 - 21 , 2013Lake Elmo , MinnesotaClick link below for more car	Quick-Look Hill-shaded Colour Relief Image of 2014 2m LIDAR Composite Digital Terrain Model (DTM) Data supplied by Environment Agency under the ...	1st Combined Convention 28th Environment Agency ConventionJuly 17 - 21 , 2013Lake Elmo , MinnesotaClick link below for more car pictures:
Location	United States, Chisago, Minnesota, Taylor Falls, GPS: 45.401366, -92.651065	United Kingdom, East Sussex, England, Fairlight Cove, GPS: 50.881552, 0.664108	United States, Chisago, Minnesota, Taylor Falls, GPS: 45.401366, -92.651065

Table 7: Location, person and organization manipulation examples in order. Our model identifies each of these examples correctly from test set. In location manipulation China is replaced by Canada, while unrelated names are swapped for person manipulation. The Annual National Club Convention & 44th Annual Walter P Chrysler Club has been replaced by Environment Agency for organization manipulation. Longer text descriptions have been truncated in examples.

new training scheme improves performance on missing modalities when compared to simple model training, while maintaining competitive performance on samples with no missing modalities.

6 CONCLUSION

With today's increase in fake news, image repurposing detection has become an important problem. It is also a relatively unexplored research area. The scope of image repurposing was expanded to semantically consistent manipulations which are more likely to fool people. We presented the MEIR dataset with intelligent and semantically consistent manipulations of location, person and organization entities. Our end-to-end deep multimodal learning model gives state-of-the art performance on MEIR and MAIM. The model

is also shown to be robust to missing modalities with a proper training scheme.

ACKNOWLEDGMENTS

This work is based on research sponsored by the Defense Advanced Research Projects Agency under agreement number FA8750-16-2-0204. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

REFERENCES

- [1] Khurshid Asghar, Zulfiqar Habib, and Muhammad Hussain. 2017. Copy-move and splicing image forgery detection and localization techniques: a review. *Australian Journal of Forensic Sciences* 49, 3 (2017), 281–307. <http://www.tandfonline.com/doi/abs/10.1080/00450618.2016.1153711>
- [2] Ella Bingham and Heikki Mannila. 2001. Random Projection in Dimensionality Reduction: Applications to Image and Text Data. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '01)*. ACM, New York, NY, USA, 245–250. <https://doi.org/10.1145/502512.502546>
- [3] Christina Boididou, Katerina Andreadou, Symeon Papadopoulos, Duc-Tien Dang-Nguyen, Giulia Boato, Michael Riegler, and Yiannis Kompatsiaris. 2015. Verifying Multimedia Use at MediaEval 2015. In *MediaEval*.
- [4] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (Oct. 2001), 5–32. <https://doi.org/10.1023/A:1010933404324>
- [5] Leo Breiman, Jerome H. Friedman, Richard A. Olshen, and Charles J. Stone. 1984. Classification and regression trees. <http://cds.cern.ch/record/2253780>
- [6] Sanjoy Dasgupta. 2000. Experiments with random projection. In *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 143–151.
- [7] Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd annual meeting on association for computational linguistics*. Association for Computational Linguistics, 363–370.
- [8] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [9] Ayush Jaiswal, Ekraam Sabir, Wael AbdAlmageed, and Premkumar Natarajan. 2017. Multimedia Semantic Integrity Assessment Using Joint Embedding Of Images And Text. In *Proceedings of the 2017 ACM on Multimedia Conference (MM '17)*. ACM, New York, NY, USA, 1465–1471. <https://doi.org/10.1145/3123266.3123385>
- [10] Zhiwei Jin, Juan Cao, Han Guo, Yongdong Zhang, and Jiebo Luo. 2017. Multimodal Fusion with Recurrent Neural Networks for Rumor Detection on Microblogs. In *Proceedings of the 2017 ACM on Multimedia Conference (MM '17)*. ACM, New York, NY, USA, 795–816. <https://doi.org/10.1145/3123266.3123454>
- [11] Zhiwei Jin, Juan Cao, Yazhi Zhang, and Yongdong Zhang. 2015. MCG-ICT at MediaEval 2015: Verifying Multimedia Use with a Two-Level Classification Model. In *MediaEval*.
- [12] Z. Jin, J. Cao, Y. Zhang, J. Zhou, and Q. Tian. 2017. Novel Visual and Statistical Image Features for Microblogs News Verification. *IEEE Transactions on Multimedia* 19, 3 (March 2017), 598–608. <https://doi.org/10.1109/TMM.2016.2617078>
- [13] J. R. Kettenring. 1971. Canonical analysis of several sets of variables. *Biometrika* 58, 3 (Dec. 1971), 433–451. <https://doi.org/10.1093/biomet/58.3.433>
- [14] Regina Marchi. 2012. With Facebook, Blogs, and Fake News, Teens Reject Journalistic Objectivity. *Journal of Communication Inquiry* 36, 3 (July 2012), 246–262. <https://doi.org/10.1177/1096859912458700>
- [15] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems* 26, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 3111–3119. <http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>
- [16] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (Oct. 2011), 2825–2830. <http://jmlr.csail.mit.edu/papers/v12/pedregosa11a.html>
- [17] Muhammad Ali Qureshi and Mohamed Deriche. 2015. A bibliography of pixel-based blind image forgery detection techniques. *Signal Processing: Image Communication* 39 (2015), 46–74. <http://www.sciencedirect.com/science/article/pii/S0923596515001393>
- [18] Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. CSI: A Hybrid Deep Model for Fake News Detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. ACM, 797–806.
- [19] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115, 3 (Dec. 2015), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
- [20] Cencheng Shen, Ming Sun, Minh Tang, and Carey E. Priebe. 2014. Generalized canonical correlation analysis for classification. *Journal of Multivariate Analysis* 130, C (2014), 310–322. https://econpapers.repec.org/article/eejmvana/v_3a130_3ay_3a2014_3ai_3ac_3ap_3a310-322.htm
- [21] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv:1409.1556 [cs]* (Sept. 2014). <http://arxiv.org/abs/1409.1556> arXiv: 1409.1556.
- [22] Ratnam Singh and Mandeep Kaur. 2016. Copy Move Tampering Detection Techniques: A Review. *International Journal of Applied Engineering Research* 11, 5 (2016), 3610–3615. https://www.ripublication.com/ijaerv11n5_104.pdf
- [23] Ming Sun, Carey E. Priebe, and Minh Tang. 2013. Generalized canonical correlation analysis for disparate data fusion. *Pattern Recognition Letters* 34, 2 (2013), 194–200.
- [24] Yue Wu, Wael Abd-Elmageed, and Prem Natarajan. 2017. Deep Matching and Validation Network: An End-to-End Solution to Constrained Image Splicing Localization and Detection. In *Proceedings of the 2017 ACM on Multimedia Conference (MM '17)*. ACM, New York, NY, USA, 1480–1502. <https://doi.org/10.1145/3123266.3123411>
- [25] Markos Zampoglou, Symeon Papadopoulos, Yiannis Kompatsiaris, Ruben Bouwmeester, and Jochen Spangenberg. 2016. Web and Social Media Image Forensics for News Professionals. In *SMN@ ICWSM*. <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM16/paper/download/13206/12860>