

Diachronic Cross-modal Embeddings

David Semedo
NOVALINCS

Universidade NOVA de Lisboa, Portugal
df.semedo@campus.fct.unl.pt

João Magalhães
NOVALINCS

Universidade NOVA de Lisboa, Portugal
jm.magalhaes@fct.unl.pt

ABSTRACT

Understanding the semantic shifts of multimodal information is only possible with models that capture cross-modal interactions over time. Under this paradigm, a new embedding is needed that structures visual-textual interactions according to the temporal dimension, thus, preserving data's original temporal organisation. This paper introduces a novel *diachronic cross-modal embedding* (DCM), where cross-modal correlations are represented in embedding space, throughout the temporal dimension, preserving semantic similarity at each instant t . To achieve this, we trained a neural cross-modal architecture, under a novel ranking loss strategy, that for each multimodal instance, enforces neighbour instances' temporal alignment, through subspace structuring constraints based on a temporal alignment window. Experimental results show that our DCM embedding successfully organises instances over time. Quantitative experiments, confirm that DCM is able to preserve semantic cross-modal correlations at each instant t while also providing better alignment capabilities. Qualitative experiments unveil new ways to browse multimodal content and hint that multimodal understanding tasks can benefit from this new embedding.

CCS CONCEPTS

• **Information systems** → **Multimedia and multimodal retrieval**; • **Computing methodologies** → *Neural networks*;

KEYWORDS

diachronic embeddings; neural cross-modal learning; semantic shifts; neural networks; data mining

ACM Reference Format:

David Semedo and João Magalhães. 2019. Diachronic Cross-modal Embeddings. In *Proceedings of the 27th ACM International Conference on Multimedia (MM '19), October 21–25, 2019, Nice, France*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3343031.3351036>

1 INTRODUCTION

It is known that with just a few glances, humans can quickly and accurately understand an image with minimal effort [6]. One of the reasons for this phenomenon is that images, or more generally humans' visual perception of elements of the physical world

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 '19, October 21–25, 2019, Nice, France

© 2019 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-6889-6/19/10...\$15.00

<https://doi.org/10.1145/3343031.3351036>

(e.g. objects, shapes), do not change over time. Language on the other hand, is a way for humans to express/communicate their knowledge about something. Accordingly, textual descriptions of images reflect the way humans refer to each image at a *given point in time*, thus being susceptible to change over time, either due to the occurrence of external events [16] or simply by word meaning change across time [9, 38]. Recently, word embedding models that capture language evolution over time (referred as *distributional diachronic* models) have been proposed [9, 24, 38]. In our setting we seek to obtain representations that *capture the evolution of visual and textual correlations over time*. While visual elements can be seen as anchors – i.e. pictures and videos freeze a reality and do not change afterwards – word descriptions change over time (e.g. the same image, or a semantically similar one, may be referred twice, at different points in time, but with different descriptions).

Cross-modal embeddings, structure images and texts in a common space, that is organised based on data interactions, enabling the study of such interactions towards multimedia understanding. Until now, these embeddings have been learned in a *static* manner, i.e. without preserving the time dimension, and thus ignoring the evolution of modality interactions [5, 7, 12, 20, 21, 23, 27, 30, 31, 35, 36]. Approaches have ranged from solutions that organize the space according to linear correlations [23, 33, 36] (image and texts co-occurrence), semantic [20, 30, 34, 35] (category information) and/or temporal correlations [27]. State-of-the-art methods can preserve semantic and temporal correlations but do not consider the time dimension in its original scale, thus resulting in the loss of all the information regarding the evolution of data.

In this paper we depart from previous approaches and propose a novel model for learning *Diachronic Cross-modal embeddings*, materialised by continuous projection functions, where the time dimension is preserved in order to capture data interactions over time. Figure 1 illustrates the target cross-modal diachronic embedding. The hypothesis is that cross-modal interactions evolve along the temporal dimension. Therefore, the embedding space should structure images and texts such that for each instant t , i.e., elements are organised according to semantic correlations between instances and corresponding absolute timestamp. This results in a model in which neighbours of an element (e.g. a text or an image), at time instant t_1 , may differ from the neighbours at time instant t_2 , if data correlations between the two instants change (Figure 1).

The first challenge we need to address corresponds to unveiling and quantifying, for each image and text, the evolution of the semantic correlations w.r.t. to other instances. The second, corresponds to the learning of the diachronic embedding, in which at each instant t , semantic category information is used to guide the structuring of multimodal instances' neighbourhood. This requires solving the *alignment problem*, in which embeddings of an instance

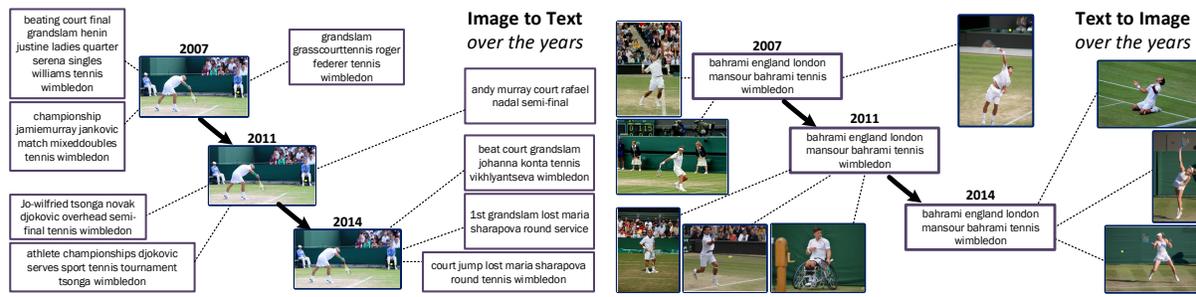


Figure 1: Diachronic Cross-modal Embeddings illustration.

at an instant t , should retain correlation with embeddings of semantically similar instances, on distinct instants. We address these challenges, by employing a two-part approach for neighbourhood structuring for an arbitrary instant time t : first, for instances of the same semantic category within a given time range, semantic correlations need to be maximal, second, instances outside the time range are placed far apart. Then, a novel ranking loss is formulated to achieve a continuous diachronic structure, by enforcing correlations from the two dimensions in the space structure organisation, by using temporal context of instances, at each instant, while aligning instances across adjacent time instants.

In summary, the key contributions of this paper are:

- (1) The first Diachronic Cross-modal embedding learning approach, where the evolution of multimodal data correlations are modelled. Time is explicitly modelled, thus allowing *conditioning on time at both training and inference time*;
- (2) A novel *temporally constrained ranking loss* formulation, aligns instances embeddings over time, and enables the learning of neural projections from timestamped multimodal data;
- (3) A principled approach that offers statistical guarantees, and allows for correct joint-inferences (image+text+time) that other methods do not, enabling it to be used for a wide number of media interpretation tasks.

2 RELATED WORK

Across the literature, several works have researched methods to model and incorporate temporal aspects to better understand data. There is enough evidence in the literature that demonstrates how information sources are correlated along a timeline of events, with different media and event types [14, 16, 25, 28, 29]. In [14] temporal clues were used to improve search relevance at query time, by modelling content streams using a multi-task regression on multivariate point processes. In Uricchio et al. [29], the value of temporal information for the tasks of image annotation and retrieval, such as tag frequency, is recognised.

In order to model the temporal behaviour of data, embeddings must retain temporal correlations [2, 9, 15, 24, 27, 38]. The challenge resides in capturing such correlations and incorporating them in cross-modal embeddings. Blei and Lafferty [4] proposed a dynamic topic model to capture temporal behaviour of data by modelling the evolution of word interactions over time. The model is an extension of static topic models (e.g. LDA), where latent topic evolution over

time was accounted. Topic model approaches treat words as *symbols* and thus lack all the properties of distributed representations [3]. Word embedding models aim at learning word representations, such that words that appear in similar contexts are structured close together in the embedding space [17]. *Diachronic Word Embeddings* consist of word embeddings that model word meaning change across time, by encoding words' usage over time [2, 9, 15, 24, 38]. Lately these models have been actively researched to aid the understanding of words' semantic evolution. A common approach to learn such embeddings has been to split text documents into bins (e.g. by year), and then train a static Skip-Gram [17] (*word2vec*) model on each bin. Embeddings of adjacent bins are then aligned by learning a linear transformation that performs the best rotational alignment, while preserving cosine similarities [9, 15, 38]. Data binning introduces several issues and limitations. Small bins are required to capture fine-grained interactions, however these may incur in bins with very few data for training. Conversely, with large bins only coarse grained representations can be obtained. To overcome this, Rosenfeld and Erk [24] recently proposed a continuous approach, in which time is taken as a continuous variable. The model learns an embedding for each word w at each time instant t . Our work goes in this direction, however two aspects invalidate the use of existing word diachronic models: a) unlike words, that are predominant across time instants, each instance is posted only once, invalidating existing alignment strategies, b) in the cross-modal scenario two modalities need to be aligned instead of only one.

Static cross-modal embedding models represent multimodal data in a common space. Early approaches [8, 13, 23, 39], learn projections based on linear correlation. State-of-the-art works adopt neural networks, to learn non-linear modality projection functions [1, 7, 20, 21, 30, 36]. The focus has been on capturing correlations between instances (using category information when available), without accounting for temporal correlations, for data organisation. Recently, a temporal cross-modal common space approach was proposed [27], where temporal correlations are used to organise instances in a static embedding space. The authors observed that for dynamic data, the incorporation of temporal insights increases retrieval performance. This embedding space has the limitation that it is not diachronic, since the temporal dimension is discarded after training, thus not allowing the study of the evolution of semantic correlations. We depart from their work by proposing the first diachronic cross-modal embedding, in which semantic evolution is preserved.

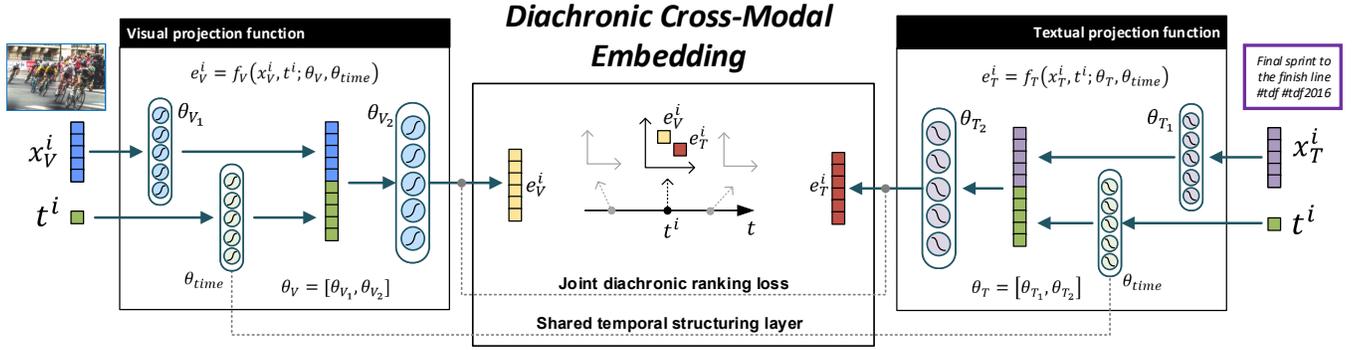


Figure 2: Diachronic cross-modal architecture overview. Visual (blue) and textual (purple) instances, at an instant t^i , are mapped to a D dimensional diachronic embedding space. A shared temporal structuring layer takes the timestamp t^i as input and learns an embedding for t^i , that is then used to independently condition modality projections on time. A diachronic ranking loss is responsible for structuring instances over time. Best viewed in color.

3 DIACHRONIC CROSS-MODAL EMBEDDING

The key element of diachronic embedding methods is the set of time-preserving projection functions, responsible for mapping the original data onto the embedding. This section will detail the proposed embedding and the corresponding project in functions.

3.1 Definitions

We start by introducing some notation and defining the task of diachronic cross-modal embedding learning. Without loss of generality, let $C = \{d_i\}_{i=1}^N$ be a set of N visual-textual instance tuples

$$d^i = (\vec{x}_V^i, \vec{x}_T^i, ts^i, c^i), \quad (1)$$

where $\vec{x}_V^i \in \mathcal{R}^{D_V}$ and $\vec{x}_T^i \in \mathcal{R}^{D_T}$ are the feature representations of the image and textual elements, respectively, ts^i the timestamp and c^i the instance (unique) semantic category. Accordingly, D_V and D_T correspond to the image and text features dimensionality, respectively. The instances timestamp have a timespan defined by $TS = [t_s, t_f]$, where t_s and t_f are the first and last instants of the dataset, respectively. Let $*$ in $\{V, T\}$ on the remainder of this paper, to avoid notation cluttering.

The goal is to obtain a common continuous (over time instants) embedding space, in which the visual and textual elements are organised according to their semantic category and timestamp. The space is formally defined as follows:

Definition 1. A diachronic cross-modal embedding space refers to a common space, that structures the visual and textual elements of the data instances over time. In this embedding, similarity between instances of the same category that are close in time, is maximised. In all other cases, similarity between instances is minimal.

In diachronic word embeddings, the temporal dimension captures *word meaning change* [9, 38]. In a cross-modal scenario, with timestamped data, the words (textual modality) that co-occur with pictures (visual modality), at different instants in time, may be different. This characteristic is quantified by the semantic alignment of the diachronic models:

Definition 2. The semantic alignment over time is the embeddings correlation, between semantically similar instances, that is retained across different time instants.

3.2 Time-preserving Projections

The diachronic cross-modal embedding embeds the modality vector \vec{x}_*^i (image or text) of an instance d^i and a time instant t , using the pair of functions:

$$\vec{e}_V = f_V(\vec{x}_V^i, t; \theta_V, \theta_{time}) \quad \vec{e}_T = f_T(\vec{x}_T^i, t; \theta_T, \theta_{time}), \quad (2)$$

where each $\vec{e}_* \in \mathbb{R}^D$ denotes the embedding of the corresponding modality, at time instant t . $\theta_V = [\theta_{V_h}; \theta_{V_o}]$ and $\theta_T = [\theta_{T_h}; \theta_{T_o}]$ are the model parameters, and θ_{time} are the key parameters responsible for controlling the temporal structuring of the \vec{x}_V^i and \vec{x}_T^i projections. As a consequence, parameters θ_{time} tie the two projection functions.

Diachronic cross-modal embedding functions are defined by the mappings:

$$f_V(\cdot) : \mathbb{R}^{D_V} \times TS \mapsto \mathbb{R}^D \quad f_T(\cdot) : \mathbb{R}^{D_T} \times TS \mapsto \mathbb{R}^D. \quad (3)$$

The output of f_V and f_T is normalised such that $\ell_2(f_*(\cdot)) = 1$. Accordingly, instances will be organised based on time and semantic similarity, over a D -dimensional space. Similarity between projected sample elements x_*^i and x_*^j , is computed through *cosine* similarity.

The diachronic projection functions $f_V(\cdot)$ and $f_T(\cdot)$ are implemented as a neural network with 2 fully connected layers. Figure 2 depicts the neural architecture. Formally, the diachronic projection functions are defined as

$$f_*(x_*^i, ts^i) = \tanh\left(\theta_{*o} \cdot [f_{*h}(x_*^i); f_{time}(ts^i)]\right), \quad (4)$$

$$f_{*h}(x_*^i) = \tanh(\theta_{*h} \cdot x_*^i), \quad f_{time}(ts^i) = \tanh(\theta_{time} \cdot ts^i), \quad (5)$$

where θ_{*h} , θ_{time} and θ_{*o} correspond to hidden (per modality), time and output layer weight matrices, respectively. $[\cdot]$ denotes the concatenation operation. An initial **encoding layer** (f_{V_h} or f_{T_h}), receives the input vector and transforms it into an internal representation that is compatible with the internal representation of data timestamps. A **shared time embedding layer** (f_{time}) maps

data timestamps to an embedding representation. The obtained time embedding is then used to condition the output projections of f_{V_h} and f_{T_h} , through a concatenation operation, making them time-dependent. A final **output layer** takes as input the result of conditioning f_{*h} and f_{time} to produce the final D -dimensional projection to a diachronic embedding space.

The resulting embeddings, produced by f_V and f_T are characterised by the properties defined in the following section.

3.3 Embedding Properties

The structure of the temporal embedding space, *i.e.* how multimodal instances will be organised, is formalised by a set of fundamental properties. These properties stem from two grounding intuitions: data is primarily associated by the time dimension, and then by their semantic categories. The model will thus capture the evolution of semantic correlations, over time instants, by maximising the similarity between instances that are within a given temporal window and share the same category. Accordingly, two embedding vectors \tilde{e}_*^i and \tilde{e}_*^j will be projected into the same neighbourhood if the two following properties are met simultaneously:

- **Property 1.** the timestamps of d^i and d^j are within the same temporal window, *i.e.* $|t^i - t^j| < w$, and the two instances d^i and d^j share the same category, *i.e.* $c^i = c^j$.

Conversely, the two embedding vectors \tilde{e}_*^i and \tilde{e}_*^j will be projected onto distant regions if:

- **Property 2.** the timestamps of d^i and d^j are outside the same temporal window, *i.e.* $|t^i - t^j| > w$, independently of the instances' semantic category;
- **Property 3.** elements do not share any semantic category.

The final and most novel property follows from the requirement that the target embedding space needs to be continuous over time. Thus, a final *semantic alignment over time* property is introduced:

- **Property 4.** For each image or text of an instance d^i , embeddings must evolve smoothly between neighbouring time instants t^1 and t^2 , with $|t^2 - t^1| \leq w$.

4 DIACHRONIC EMBEDDING LEARNING

To learn the time-dependent continuous embedding functions $f_V(\cdot)$ and $f_T(\cdot)$ it is essential to maximise correlation in the new embedding space. In this section we define the optimisation objective and show how it enforces the temporal organisation of the embedding.

4.1 From Projections to Ranking Loss

Following the definition of the diachronic projection functions, the two component correlation scheme (temporal and semantic), and respecting the properties defined in section 3.3, we start by noting that similarities in the embedding space are computed by the dot product $s(x_*^i, x_*^j) = f_*(x_*^i) \cdot f_*(x_*^j)$. Building on the most recent state-of-the-art cross-modal learning works [21, 27, 30, 33, 37], we adopt the *ranking loss function* as the model base loss. In its general formulation, triplets (x_*^a, x_*^p, x_*^n) , are composed by an anchor element x_*^a , that should be more similar to positive elements x_*^p sharing a category, than to negative elements x_*^n not sharing categories, by at least a margin m . Triplet constraints are expressed as

$s(x_*^a, x_*^p) > s(x_*^a, x_*^n) + m$, and then turned into a differentiable function, by means of a relaxation under the hinge loss function [11]:

$$\ell_\theta(x_*^a, x_*^p, x_*^n) = [m - s(x_*^a, x_*^p) + s(x_*^a, x_*^n)]_+, \quad (6)$$

where m denotes a constant margin, $[x]_+$ the function $\max(0, x)$, and $\theta = [\theta_V; \theta_T; \theta_{time}]$ is the complete set of parameters. One of such constraints would then be enforced for each sampled triplet. In the next section we detail how ranking loss is extended to cope with the temporal dimension.

4.2 Joint Diachronic Ranking Loss

The learning problem is then formulated by coupling the learning of the two individual modality and a third timestamp embedding functions, through a global loss function \mathcal{L} . The full loss function of our model, for diachronic cross-modal embedding learning, is derived by enforcing multiple constraints, for each possible anchor element, and summing all the constraint violations. To this end, we define each f_* as a neural network, capable of unveiling complex non-linear interactions. The objective function becomes

$$\arg \min_{\theta_V, \theta_T, \theta_{time}} \mathcal{L}_{\theta_V, \theta_T, \theta_{time}}(C), \quad (7)$$

with θ_V , θ_T and θ_{time} being the projection functions parameters.

State-of-the-art cross-modal retrieval interlace modalities by enforcing triplet constraints in both modality directions [19, 30, 33], *i.e.* $image \mapsto text$ and $text \mapsto image$. Thus, we formulate the final loss \mathcal{L} function for temporal cross-modal embedding model with parameters $\theta = [\theta_V; \theta_T; \theta_{time}]$ as:

$$\mathcal{L}_\theta(C) = \sum_{a,p,n} \underbrace{\mathcal{L}_\theta(x_V^a, x_V^p, x_V^n)}_{image \mapsto text} + \underbrace{\mathcal{L}_\theta(x_T^a, x_T^p, x_T^n)}_{text \mapsto image}, \quad (8)$$

where p and n denote indices of positive and negative instances, respectively, w.r.t. to an anchor element x_*^a . This function is evaluated batch-wise. Thus, at each batch, the sampled elements are used to create triplet constraints.

4.3 Continuous Diachronic Structure

Formally, for each anchor element x_*^a , of a triplet (x_*^a, x_*^p, x_*^n) , we define the following loss function:

$$\mathcal{L}_\theta(x_*^a, x_*^p, x_*^n) = \mathcal{L}_{inter}(x_*^a, x_*^n) + \mathcal{L}_{intra}(x_*^a, x_*^p), \quad (9)$$

where \mathcal{L}_{inter} and \mathcal{L}_{intra} are based on the ranking loss function, enforcing inter-category and intra-category embedding related properties, respectively. Both x_*^n and x_*^p correspond to sampled negative and positive images or texts, respectively.

Temporal triplets. To let the loss function of equation 9 enforce the required properties, special attention needs to be given to triplets' temporal and semantic correlations. Given an anchor element x_*^a , a temporal window of size w is used to enforce smoothness between temporal neighbourhoods in the embedding, for images and texts of the same category as x_*^a . This context-window formulation has been quite successful in word embeddings [18, 22]. For diachronic cross-modal embeddings, we consider temporally adjacent multimodal instances instead of a sliding window over text. Thus, triplets are obtained by sampling positive and negative instances by always using as positive image/text to the anchor x_*^a ,

its modality counterpart of the same instance a (e.g. for the anchor x_V^a we use $x_T^{p=a}$ as positive, and vice-versa). With this strategy we maximise correlation between modalities using images and texts that occur together, allowing us to better capture intra-category semantic diversity.

Intra-category and temporal smoothing. The assumption is that given an image or a text, its embedding should change smoothly between adjacent time instants. Smoothness of that change is captured by the size of the considered window (*Property 4*). The function \mathcal{L}_{intra} is responsible for embedding alignment over time, *i.e.* *Property 4*. To accomplish this, a temporal window of size w is considered. The rationale is that we want that instances of the same category to be close. But from *Property 2* definition, embeddings of instances of the same category, should be far apart, if they are temporally far apart. Thus, if we do not enforce any margin between positive instances, the optimal solution is when all instances of the same category are mapped to the same point, losing temporal evolution of semantic correlations. To overcome this, we employ a temporally decaying ranking loss formulation, for instances of the same category. Namely, given an anchor element x_*^a and a positive sampled instance x_*^p , \mathcal{L}_{intra} is defined as the following branch function:

$$\mathcal{L}_{intra}(x_*^a, x_*^p) = \begin{cases} 0 & , |t^a - t^p| \leq w \\ \rho(t^a, t^p) \cdot \ell_\theta(x_{V/T}^a, x_{T/V}^a, x_*^p) & , \text{otherwise} \end{cases} \quad (10)$$

where $\rho(t^a, t^p) = 1 - \exp(-|t^a - t^p| \cdot \lambda)$ is a temporal decaying function, and λ the decay rate. When two positive instances are less than w instants far part, no margin is enforced between the two. Otherwise, a triplet constraint is enforced, weighted by a decaying function that exponentially decreases the importance of ℓ_θ , the closer in time two instances are.

Finally, it is important to observe that the triplet sampling for batch creation makes a stochastic approximation to the optimal set of triplets (mining the optimal triplets for each batch is computationally too expensive [26]). Hence, although for a given batch it is likely that a given anchor x_*^a with timestamp t^a , has no adjacent (in time) projections, that same anchor x_*^a will occur with high-likelihood in subsequent batches with the required instances. Thus, convergence guarantees are preserved, due to the stochastic approximation made through the triplet sampling strategy.

Inter-category separation. For inter-class alignment \mathcal{L}_{inter} , we want to structure the embedding space such that instances of different categories will be far apart, independently of their timestamp. To this end, we enforce a triplet loss constraint, with a large margin (*i.e.* $m = 1$) over such triplets. Formally, \mathcal{L}_{inter} is defined as:

$$\mathcal{L}_{inter}(x_*^a, x_*^n) = \ell_t(x_*^a, x_*^{(p=a)}, x_*^n), \quad (11)$$

where if the anchor is an image (*i.e.* x_V^a), then the positive corresponds to a text (*i.e.* x_T^a), and vice-versa. This formulation achieves two goals: enforces the separation of positive from negative instances, by a margin m , and aligns the embeddings of the image and the text of the anchor instance.

4.4 Binned Diachronic Structure

The previous general formulation will behave as a binned diachronic method [9, 15] when $\mathcal{L}_{intra}(\cdot) = 0$. In this extreme case, data can be first divided into bins and a static cross-modal embedding model is trained on data from each bin. Then, embeddings of adjacent bins are aligned by solving the Orthogonal Procrustes problem [9, 15]. Namely, given two embedding matrices M_t and M_{t+1} with shape $N \times D$, containing N images and texts embeddings representative of adjacent time instants, the best rotational alignment is computed as:

$$\Omega_{t \rightarrow t+1} = \arg \min_{\Omega^T \Omega = I} \|M_t \Omega - M_{t+1}\|_F, \quad (12)$$

preserving cosine similarities within each M . For diachronic word embeddings, words present at one instant are also present in the next. Thus, one can directly align the embeddings of each word, at different instants. As in the cross-modal scenario each document occurs only once, we set M_{t+1} by projecting images and texts from instant t in bin $t + 1$, and perform alignment. This approach, while being prone to embedding alignment issues, is taken as a strong baseline. It preserves temporal locality biases by definition, as cross-modal correlations over distinct bins are never considered.

5 EVALUATION

5.1 A 20 years Flickr Images Dataset

For evaluation we construct a new large scale weakly-labeled dataset¹ with multimodal instances obtained from the Flickr² social network. Namely, we collect documents related with topics that show a dynamic behaviour over time such as spike-based and recurrent events. Figure 3 shows the temporal distributions of four sampled categories, and illustrates the diversity in terms of dynamic behaviour captured by the dataset. Data is collected over the period 1-1-1970 to 31-12-2018. The Flickr API is used³ to retrieve images and texts from a total of 21 categories: *easter-sunday, edinburgh-festival, flood, formula-one, horse-riding, independence-day, london-marathon, mountain-camping, nuclear-disaster, olympic-games, picnic, rock-climbing, scuba-diving, snowboarding, solar-eclipse, terrorism, tour-de-france, tsunami, white-house, wimbledon, world-cup*. We use the category name as keyword to query the API and collect data, and filter instances whose *date taken* is outside the considered temporal range. The models' granularity is set to *months*. To ensure that enough instances are available for each bin, we restrict the temporal range of images to the past 20 years (red line on figure 3 depicts the cut), and bins with less than 100 documents are excluded. After applying a set of SPAM filtering techniques, we obtain a total of 709033 instances. In general, images have (near) professional quality. Texts are on average 23.0 words long. We use 10% of the data for testing and split the remaining data in 90% for training and 10% for validation, resulting in 574,308, 63,804 and 70,921 instances, for training, validation and testing, respectively.

5.2 Implementation Details

Networks are jointly trained using SGD, with 0.9 momentum, and a learning rate $\eta = 0.005$. We train the model for 25 epochs and retain

¹<https://novasearch.org/multimodal-diachronic-models>

²<https://www.flickr.com/>

³Only Creative Commons licensed data was retrieved.

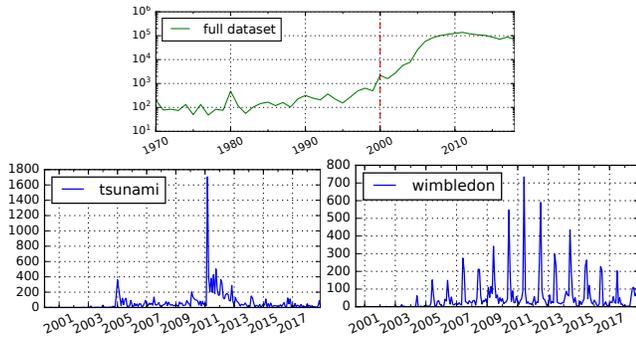


Figure 3: Temporal distribution of documents across the full dataset and four sample categories.

the best performing model based on the validation set loss. Mini-batch size is set to 64. For each neuron, we use *tanh* non-linearities. A pre-trained ResNet-50 [10], with the last fully connected layer removed (softmax), is used for image representation. We set $\lambda = 0.1$, window size $w = 4$ and pairwise-ranking loss margin $m = 1.0$. We adopt the TF-IDF bag-of-words representation for texts and CNN image representations for all methods.

The layers corresponding to the θ_{*h} and θ_{time} parameters have dimension 1024 and 200 respectively, and θ_{*o} has $D = 200$ dimensions. Thus, for an instance d^i , the visual projection network takes the CNN representation x_V^i of the image, the textual projection a bag-of-words representation of the text x_T^i , and the timestamp embedding the timestamp as input, producing the D -dimensional diachronic embedding.

5.3 Methodology

Experiments are performed in a cross-modal setting with different temporal parameters, *i.e.* when an image is considered, the evaluated neighbours are texts, and vice-versa. We refer to **DCM-Binned** and **DCM-Continuous** as the diachronic model with binned (section 4.4) and continuous (section 4.3) structure.

5.4 Diachronic Semantic Alignment

In this section we evaluate the capability of both DCM-Binned and DCM-Continuous to capture and model diachronic data behaviours. First it is important to assess the semantic alignment over time of the diachronic space. This corresponds to assessing if the obtained diachronic embedding space is capable of relating embeddings of images and texts of instant ts^i , with instances that occurred in distinct time instants $ts^j \in TS$. To accomplish this, we evaluate the semantic alignment quality of the diachronic space by designing two complementary tasks.

Coarse Semantic Alignment One operation permitted by a diachronic embedding consists of understanding, which instances are related to an image or text, from the set of all instances, spanning across all possible time instants. Hence, the *coarse* designation. An instance d^i is expected to be semantically correlated with instances not only on the d^i time instant, but also on other instants (*e.g.* recurrent events). To capture such behaviour, diachronic embedding models are required to correctly align embeddings over different time instants, such that semantic correlations are preserved. To do

Table 1: Diachronic Semantic Alignment.

Coarse Semantic Alignment			
Methods (<i>mAP</i>)	$I \mapsto T$	$T \mapsto I$	Avg.
DCM-Binned (w/ Align [9])	0.203	0.197	0.200
DCM-Continuous	0.370	0.348	0.359
Local Semantic Alignment			
Methods (<i>mAP@10</i>)	$I \mapsto T$	$T \mapsto I$	Avg.
DCM-Binned (w/ Align [9])	0.078	0.086	0.082
DCM-Continuous	0.313	0.330	0.322

this, an image or text, should be embedded on the instant ts^i corresponding to its timestamp ts^i . Its neighbourhood in embedding space, can then be analysed by comparing the similarities of the each d^i against all projected instances, on their corresponding time instant.

This is evaluated by projecting each image/text of an instance d^i from the test set, in the time instant ts^i (its timestamp) and evaluating the semantic similarity of its text/image neighbours, respectively. We use *mAP*, computed over the whole test set, to evaluate the similarity of the neighbours using semantic category information. The top part of Table 1 shows the results of this experiment. The DCM-Continuous significantly outperforms DCM-Binned, revealing superior alignment capabilities. This is justified by the binning procedure, which despite the alignment procedure, due to the stochastic nature of neural networks models (*e.g.* different initializations), leads to different organisation of data after convergence, despite its semantic correlations.

Local Semantic Alignment. In this experiment we project instances d^i onto all possible timestamps $ts^j \in TS$ and assess how each projection \vec{e}_{*}^{i,ts^j} relates to instances of that temporal neighbourhood (bin in our case). This operation is possible due to DCM’s preservation of local alignment (*w.r.t.* to time). Namely, semantically similar instances should be close in the embedding space when projected into the same time instant ts^j

We sampled 50 query instances from each category, from the test set, and project each instance d^i into each timestamp $ts^j \in TS$. Then, for each time instant ts^j , we consider only the neighbours of the embedding of d^i on that instant, *i.e.* only embeddings \vec{e}_{*}^{i,ts^j} of images and texts from ts^j instant are considered. Then we evaluate if the top-10 closest neighbours on each time instant are semantically similar (*i.e.* belong to the same category) using *mAP@10*. Table 1 shows that DCM-Continuous clearly outperforms DCM-Binned. Again, this is a result of the bad alignment of the binned approach, from which DCM-Continuous is able to overcome.

5.5 Temporally-bounded Semantics

In this section we evaluate the semantic organisation of the embedding space in temporally-bounded data. To evaluate the correctness of the embeddings neighbourhood at each instant, the test set is binned and each bin is evaluated individually. Following cross-modal learning works [7, 21, 23, 30, 32, 36], we evaluate each method on the tasks of *Image-to-Text* and *Text-to-Image* retrieval, using mean Average Precision *mAP* for *all the retrieved results*. Figure 4 shows the plots of DCM variants, with the *mAP* results per month. A static cross-modal embedding baseline, corresponding to our model without the temporal layer and with $\mathcal{L}_{intra}(\cdot) = 0$,

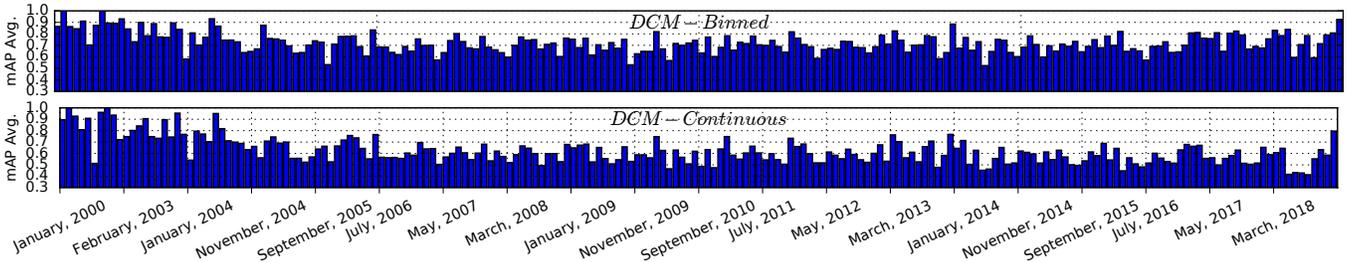


Figure 4: Temporally bounded cross-modal results (mAP) of (a) DCM-Binned (aligned) and (b) DCM-Continuous.

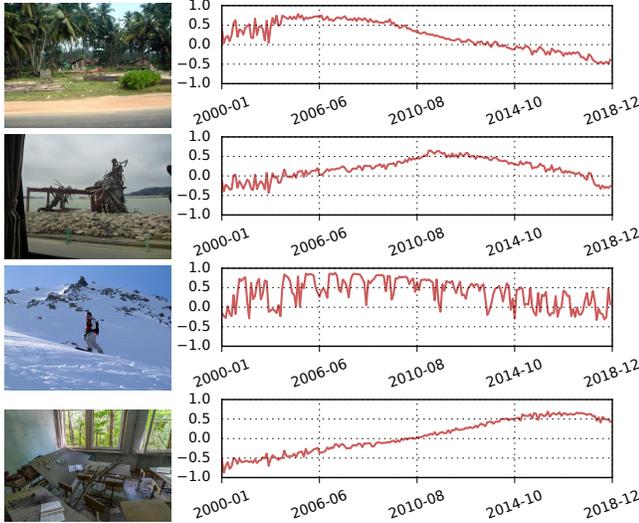


Figure 5: Semantic dispersion over time analysis for four sampled images.

was trained in all data ignoring time and obtained a mAP of 0.639. The DCM-Binned baseline and the DCM-continuous methods obtained a mAP of 0.724 and 0.623, respectively. We further evaluated against TempXNet [27], which uses temporal correlations to learn a static embedding space, obtaining a mAP of 0.645. All DCM variants have shown to be on par with state-of-the-art approaches, by obtaining scores above 0.60 mAP points. This experiment confirms our hypothesis that semantic cross-modal correlations change over time. DCM-Binned outperformed all the other methods due to its temporal locality bias: for each bin (month) a static model is trained solely on data from that bin, thus modelling the local cross-modal correlations independently, and without influences from correlations of the remaining bins. However, it lacks the advantages of a fully diachronic model.

5.6 Media Time-Period Inference

We now evaluate the enforcement of properties that define the neighbourhood of each projected instance, based on a given time window (Properties 1 and 2). Accordingly, we evaluate the neighbours of each image/text d^i , by computing a *temporally bounded mAP* ($mAP@50$): a neighbour is considered relevant if it belongs to the same category and its timestamp is within a time-window of size w , w.r.t. to the timestamp ts^i of d^i . The results are shown in

Table 2: Time-Period Inference Results.

Methods ($mAP@50$)	$I \mapsto T$	$T \mapsto I$	Avg.
Static cross-modal	0.048	0.059	0.054
TempXNet [27]	0.052	0.070	0.061
DCM-Continuous	0.126	0.144	0.135

table 2. It can be observed that DCM-Continuous significantly outperforms the two compared baselines in defining neighbourhoods that respect property 1, *i.e.* that instances from the same category that are close in time (time window w), lie close together. Section 5.5 and Figure 4 show that static cross-modal methods need to be artificially fed with the relevant time period to achieve comparable results. In contrast, this result experimentally demonstrates that DCM-Continuous does not need to receive this outside information and can indeed infer the relevant time periods on its own.

5.7 Semantic Dispersion over Time

Previously, we confirmed that diachronic embedding spaces are semantically aligned and that semantic correlations do change over time. Hence, in this section we examine the semantic dispersion of multimodal data over time. Given an image or a text, we expect that its correlations with other instances, over time, will evolve. The evolution pattern is expected to be grounded on the temporal characteristics (*e.g.* peak based, recurrent event, etc.) of the topic of each instance.

To assess this, we consider a set of target instances d^i , in which semantic evolution will be evidenced by semantic dispersion changes, over each instant $ts \in TS$. Namely, for an instance d^i with timestamp ts^i , given its embedding on instant ts^i , we define semantic dispersion as the variation of the similarity between d^i embedding and its closest neighbour instances. In practice, we sample instances d^j and project them in the time instant ts^i corresponding to its timestamp, using DCM-Continuous, obtaining the embedding \vec{e}^{i,ts^i} . Then, for each instant $ts \in TS$, we compute the semantic dispersion of the top K neighbours of \vec{e}^{i,ts^i} . Semantic dispersion on an instant ts is defined as the average of the cosine similarities between \vec{e}^{i,ts^i} and each of the K neighbours. We set $K = 5$.

Figure 5 shows the results of this experiment for four different images. The first two images belong to the *tsunami* category: the first corresponds to the Indonesia series of tsunamis in 2007, and the second to the tsunami in Japan, 2012. It can be seen that maximal similarity is achieved around the dates of the corresponding tsunami. For the first image, after the peak around 2006, similarity decreases gradually in future instants. Despite the tsunami of 2012, its similarity with content from that tsunami is low. This evidences

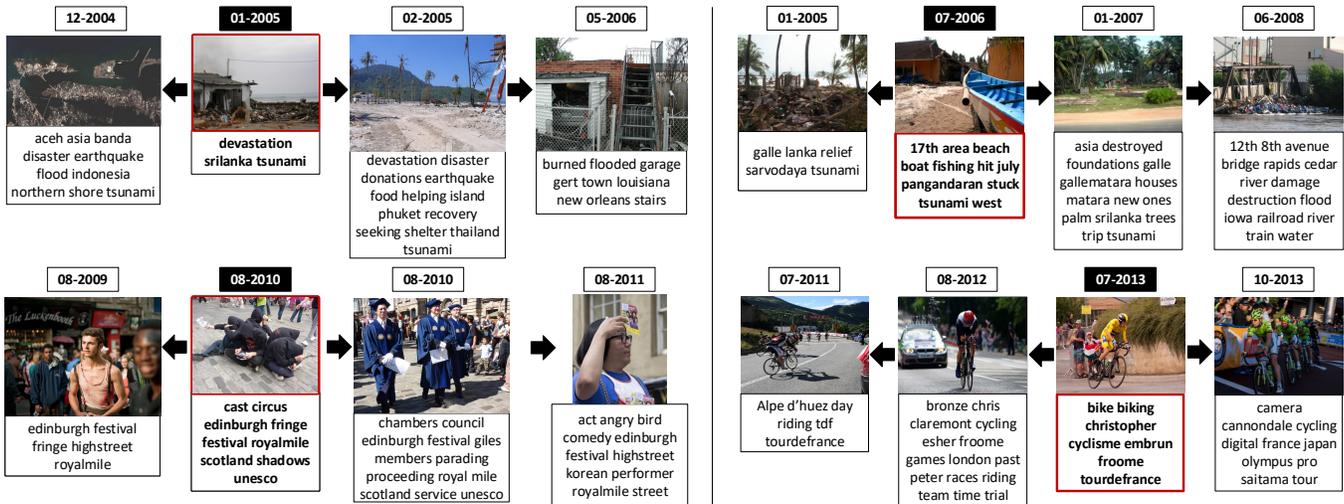


Figure 6: Evolution over time for 4 query examples (query timestamps are black-filled). Instances were retrieved from before and after the query timestamp. On the left, image queries were used to retrieve documents through their text. On the right, text queries were used to retrieve documents through their images.

that DCM-Continuous effectively delivers Property 4 (section 3.3). The third image, *snowboarding*, shows a recurrent evolution of semantic correlations over time, being more stable over the years. The last image, taken in August 2018, depicts an abandoned place due to the *nuclear disaster* of Chernobyl (pictures of contaminated areas became possible with the advent of drones and robots). Semantic similarity over time in this image gradually increases until 2015, and then stabilises until 2018. This experiment shows that diachronic embeddings obtained with DCM-Continuous encode cross-modal interactions evolution, enabling the understanding of multimodal correlations over time.

5.8 Cross-Modal Evolution

The cross-modal diachronic embedding model enables novel ways of exploring multimodal instances. One example is the analysis of the correlations evolution of an image or text, Figure 6, along the years, which is encoded in its embedding trajectory. Such operation enables one to understand the correlations shift along the years.

To perform this experiment, we sample a set of images and texts, and project them on their corresponding timestamp instant. Then, to avoid inspecting all time instants, we restrict the number of instants to the top-20 bin, *i.e.* the bins with the closest image/text, to a target text/image, respectively, based on cosine similarity. For illustration purposes, we show images/texts from 4 different instants on Figure 6. Queries are marked with the black-filled timestamps and instances were retrieved from before and after this timestamp.

The inspection of the evolution timeline (obtained with DCM-Continuous) lets us interpret the trajectory of the sampled images and texts, at particular time instants. For example, on the top-left case, the query was an image of the 2004 tsunami in Indonesia. DCM-Continuous retrieved texts of other tsunamis and floods that occurred over the years. On the right-bottom example of Figure 6 the query was a textual description of a picture with cyclist Christopher Froome. With this textual query and corresponding timestamp,

DCM-Continuous was able to illustrate the evolution of *cyclism* topic over time.

6 CONCLUSIONS

This paper introduced the first diachronic cross-modal embedding, enabling novel interpretations of cross-modal semantic shifts over time. The key novelties of the proposed diachronic model are:

- The neural architecture implements a *temporal structuring layer* that is shared across the two projection functions.
- A novel *joint diachronic ranking loss* function to control temporal structuring of data in the embedding space.

Moreover, experiments, on a 20 year span dataset, illustrated the semantic evolution and temporal flexibility of the model. The key take away messages are:

- *Cross-modal semantic evolution* is captured by the model allowing the inspection of temporal multimodal information;
- *Time is handled in a flexible manner*, *i.e.* projections are timestamped, data is organised temporally, thus supporting several diachronical operations.

As future work, we plan to investigate the geometry of the embedding space when using other temporal distributions to capture complex multimodal interactions over time.

ACKNOWLEDGMENTS

This work has been partially funded by the CMU Portugal research project GoLocal Ref. CMUP-ERI/TIC/0046/2014, by the H2020 ICT project COGNITUS with the grant agreement n° 687605 and by the FCT project NOVA LINCS Ref. UID/CEC/04516/2019. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the GPUs used for this research.

REFERENCES

- [1] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. 2013. Deep Canonical Correlation Analysis. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28 (ICML '13)*. JMLR.org, III-1247–III-1255. <http://dl.acm.org/citation.cfm?id=3042817.3043076>
- [2] Robert Bamler and Stephan Mandt. 2017. Dynamic Word Embeddings. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, International Convention Centre, Sydney, Australia, 380–389. <http://proceedings.mlr.press/v70/bamler17a.html>
- [3] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A Neural Probabilistic Language Model. *J. Mach. Learn. Res.* 3 (March 2003), 1137–1155. <http://dl.acm.org/citation.cfm?id=944919.944966>
- [4] David M. Blei and John D. Lafferty. 2006. Dynamic Topic Models. In *Proceedings of the 23rd International Conference on Machine Learning (ICML '06)*. ACM, New York, NY, USA, 113–120. <https://doi.org/10.1145/1143844.1143859>
- [5] Mengdi Fan, Wenmin Wang, Peilei Dong, Liang Han, Ronggang Wang, and Ge Li. 2017. Cross-media Retrieval by Learning Rich Semantic Embeddings of Multimedia. In *Proceedings of the 2017 ACM on Multimedia Conference (MM '17)*. ACM, New York, NY, USA, 1698–1706. <https://doi.org/10.1145/3123266.3123369>
- [6] Li Fei-Fei, Asha Iyer, Christof Koch, and Pietro Perona. 2007. What do we perceive in a glance of a real-world scene? *Journal of vision* 7 1 (2007), 10.
- [7] Fangxiang Feng, Xiaojie Wang, and Ruifan Li. 2014. Cross-modal Retrieval with Correspondence Autoencoder. In *Proceedings of the 22nd ACM International Conference on Multimedia (MM '14)*. ACM, New York, NY, USA, 7–16. <https://doi.org/10.1145/2647868.2654902>
- [8] Yunchao Gong, Qifa Ke, Michael Isard, and Svetlana Lazebnik. 2014. A Multi-View Embedding Space for Modeling Internet Images, Tags, and Their Semantics. *Int. J. Comput. Vision* 106, 2 (Jan. 2014), 210–233. <https://doi.org/10.1007/s11263-013-0658-4>
- [9] William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic word embeddings reveal statistical laws of semantic change. *arXiv preprint arXiv:1605.09096* (2016).
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [11] Ralf Herbrich, Thore Graepel, and Klaus Obermayer. 2000. Large Margin Rank Boundaries for Ordinal Regression. In *Advances in Large Margin Classifiers*, P. J. Bartlett, B. Schölkopf, D. Schuurmans, and A. J. Smola (Eds.). MIT Press, 115–132.
- [12] Xin Huang, Yuxin Peng, and Mingquan Yuan. 2017. MHTN: Modal-adversarial Hybrid Transfer Network for Cross-modal Retrieval. *CoRR abs/1708.04308* (2017). [arXiv:1708.04308](http://arxiv.org/abs/1708.04308) <http://arxiv.org/abs/1708.04308>
- [13] C. Kang, S. Xiang, S. Liao, C. Xu, and C. Pan. 2015. Learning Consistent Feature Representation for Cross-Modal Multimedia Retrieval. *IEEE Transactions on Multimedia* 17, 3 (March 2015), 370–381. <https://doi.org/10.1109/TMM.2015.2390499>
- [14] Gunhee Kim and Eric P. Xing. 2013. Time-sensitive Web Image Ranking and Retrieval via Dynamic Multi-task Regression. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining (WSDM '13)*. ACM, New York, NY, USA, 163–172. <https://doi.org/10.1145/2433396.2433417>
- [15] Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. Statistically Significant Detection of Linguistic Change. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 625–635. <https://doi.org/10.1145/2736277.2741627>
- [16] Jey Han Lau, Nigel Collier, and Timothy Baldwin. 2012. On-line Trend Analysis with Topic Models: #twitter Trends Detection Topic Model Online. In *Proceedings of COLING 2012*. The COLING 2012 Organizing Committee, 1519–1534. <http://aclweb.org/anthology/C12-1093>
- [17] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *CoRR abs/1301.3781* (2013). [arXiv:1301.3781](http://arxiv.org/abs/1301.3781) <http://arxiv.org/abs/1301.3781>
- [18] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2 (NIPS'13)*. Curran Associates Inc., USA, 3111–3119. <http://dl.acm.org/citation.cfm?id=2999792.2999959>
- [19] Niluthpol Chowdhury Mithun, Juncheng Li, Florian Metzke, and Amit K. Roy-Chowdhury. 2018. Learning Joint Embedding with Multimodal Cues for Cross-Modal Video-Text Retrieval. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval (ICMR '18)*. ACM, New York, NY, USA, 19–27. <https://doi.org/10.1145/3206025.3206064>
- [20] Yuxin Peng, Xin Huang, and Jinwei Qi. 2016. Cross-media Shared Representation by Hierarchical Learning with Multiple Deep Networks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI'16)*. AAAI Press, 3846–3853. <http://dl.acm.org/citation.cfm?id=3061053.3061157>
- [21] Y. Peng, J. Qi, X. Huang, and Y. Yuan. 2018. CCL: Cross-modal Correlation Learning With Multigrained Fusion by Hierarchical Network. *IEEE Transactions on Multimedia* 20, 2 (Feb 2018), 405–420. <https://doi.org/10.1109/TMM.2017.2742704>
- [22] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 1532–1543. <https://doi.org/10.3115/v1/D14-1162>
- [23] Nikhil Rasiwasia, Jose Costa Pereira, Emanuele Coviello, Gabriel Doyle, Gert R.G. Lanckriet, Roger Levy, and Nuno Vasconcelos. 2010. A New Approach to Cross-modal Multimedia Retrieval. In *Proceedings of the 18th ACM International Conference on Multimedia (MM '10)*. ACM, New York, NY, USA, 251–260. <https://doi.org/10.1145/1873951.1873987>
- [24] Alex Rosenfeld and Katrin Erk. 2018. Deep Neural Models of Semantic Shift. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, 474–484. <https://doi.org/10.18653/v1/N18-1044>
- [25] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors. In *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*. ACM, New York, NY, USA, 851–860. <https://doi.org/10.1145/1772690.1772777>
- [26] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. FaceNet: A Unified Embedding for Face Recognition and Clustering. *arXiv e-prints*, Article arXiv:1503.03832 (March 2015), arXiv:1503.03832 pages. [arXiv:cs.CV/1503.03832](http://arxiv.org/abs/1503.03832)
- [27] David Semedo and Joao Magalhaes. 2018. Temporal Cross-Media Retrieval with Soft-Smoothing. In *Proceedings of the 26th ACM International Conference on Multimedia (MM '18)*. ACM, New York, NY, USA, 1038–1046. <https://doi.org/10.1145/3240508.3240665>
- [28] Mikalai Tsytsarau, Themis Palpanas, and Malu Castellanos. 2014. Dynamics of News Events and Social Media Reaction. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '14)*. ACM, New York, NY, USA, 901–910. <https://doi.org/10.1145/2623330.2623670>
- [29] Tiberio Uricchio, Lamberto Ballan, Marco Bertini, and Alberto Del Bimbo. 2013. Evaluating Temporal Information for Social Image Annotation and Retrieval. In *Image Analysis and Processing - ICIAP 2013*, Alfredo Petrosino (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 722–732.
- [30] Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. 2017. Adversarial Cross-Modal Retrieval. In *Proceedings of the 2017 ACM on Multimedia Conference (MM '17)*. ACM, New York, NY, USA, 154–162. <https://doi.org/10.1145/3123266.3123326>
- [31] K. Wang, R. He, L. Wang, W. Wang, and T. Tan. 2016. Joint Feature Selection and Subspace Learning for Cross-Modal Retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 10 (Oct 2016), 2010–2023. <https://doi.org/10.1109/TPAMI.2015.2505311>
- [32] Kaiye Wang, Qiyue Yin, Wei Wang, Shu Wu, and Liang Wang. 2016. A Comprehensive Survey on Cross-modal Retrieval. *CoRR abs/1607.06215* (2016).
- [33] L. Wang, Y. Li, and S. Lazebnik. 2016. Learning Deep Structure-Preserving Image-Text Embeddings. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5005–5013. <https://doi.org/10.1109/CVPR.2016.541>
- [34] Yiling Wu, Shuhui Wang, and Qingming Huang. 2018. Learning Semantic Structure-preserved Embeddings for Cross-modal Retrieval. In *Proceedings of the 26th ACM International Conference on Multimedia (MM '18)*. ACM, New York, NY, USA, 825–833. <https://doi.org/10.1145/3240508.3240521>
- [35] Xing Xu, Jingkuan Song, Huimin Lu, Yang Yang, Fumin Shen, and Zi Huang. 2018. Modal-adversarial Semantic Learning Network for Extendable Cross-modal Retrieval. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval (ICMR '18)*. ACM, New York, NY, USA, 46–54. <https://doi.org/10.1145/3206025.3206033>
- [36] F. Yan and K. Mikolajczyk. 2015. Deep correlation for matching images and text. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3441–3450. <https://doi.org/10.1109/CVPR.2015.7298966>
- [37] T. Yao, T. Mei, and C. W. Ngo. 2015. Learning Query and Image Similarities with Ranking Canonical Correlation Analysis. In *2015 IEEE International Conference on Computer Vision (ICCV)*. 28–36. <https://doi.org/10.1109/ICCV.2015.12>
- [38] Zijun Yao, Yifan Sun, Weicong Ding, Nikhil Rao, and Hui Xiong. 2018. Dynamic Word Embeddings for Evolving Semantic Discovery. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM '18)*. ACM, New York, NY, USA, 673–681. <https://doi.org/10.1145/3159652.3159703>
- [39] X. Zhai, Y. Peng, and J. Xiao. 2014. Learning Cross-Media Joint Representation With Sparse and Semi-supervised Regularization. *IEEE Transactions on Circuits and Systems for Video Technology* 24, 6 (June 2014), 965–978. <https://doi.org/10.1109/TCSVT.2013.2276704>