

Multi-Modal Deep Analysis for Multimedia

Wenwu Zhu^{1b}, Fellow, IEEE, Xin Wang, Member, IEEE, and Hongzhi Li, Member, IEEE

Abstract—With the rapid development of Internet and multimedia services in the past decade, a huge amount of user-generated and service provider-generated multimedia data become available. These data are heterogeneous and multi-modal in nature, imposing great challenges for processing and analyzing them. Multi-modal data consist of a mixture of various types of data from different modalities such as texts, images, videos, audios etc. In this article, we present a deep and comprehensive overview for multi-modal analysis in multimedia. We introduce two scientific research problems, data-driven correlational representation and knowledge-guided fusion for multimedia analysis. To address the two scientific problems, we investigate them from the following aspects: 1) *multi-modal correlational representation*: multi-modal fusion of data across different modalities, and 2) *multi-modal data and knowledge fusion*: multi-modal fusion of data with domain knowledge. More specifically, on data-driven correlational representation, we highlight three important categories of methods, such as multi-modal deep representation, multi-modal transfer learning, and multi-modal hashing. On knowledge-guided fusion, we discuss the approaches for fusing knowledge with data and four exemplar applications that require various kinds of domain knowledge, including multi-modal visual question answering, multi-modal video summarization, multi-modal visual pattern mining and multi-modal recommendation. Finally, we bring forward our insights and future research directions.

Index Terms—Multi-modal analysis, data-driven correlational representation, knowledge-guided data fusion.

I. INTRODUCTION

WE ARE now living in the era of *Cyber, Physical and Human* (CPH) spaces. The Moore Law illustrates that the CPU speed will double every 18 months, resulting in the ubiquity of computing; the Bell Laws indicates that the chip size tends to reduce by half every 18 months, making devices including all types of sensors everywhere; the Gilder's Law shows that the network bandwidth can double every 6 months, causing communications which connect human, computers and physical identities to be ubiquitous in our daily lives. In short, data are everywhere in the era of *Cyber*,

Physical and *Human* (CPH) spaces. For example, various kinds of user-generated and service provider-generated data in social media together with the growing popularity of other information sources such as cell phones and cameras have produced a large amount of multi-modal multimedia data. Multi-modal data consists of a mixture of various types of data such as texts, images, audios, videos etc. In the past decades, most researchers, focus on analyzing data in a single modality, making uni-modal/single-medium analysis a well studied topic to date. However, we need to study multi-modal data in real life, which is particular important when we enter the “Artificial Intelligence (AI) Epoch”. Discoveries in cognitive science [1] have confirmed the fact that human are able to perceive their surrounding environment through fusing the feedback from multiple sensory organs (eyes, nose, ears, etc.) together. As such, the investigation of multi-modal analysis serves as a very promising direction in boosting the progress of research in big data and AI. Fortunately, the advent of multi-modal data brings us great opportunities for multi-modal analysis in multimedia.

Nevertheless, analyzing multi-modal multimedia data imposes great challenges. One scientific problem is how to jointly consider and fuse information from different modalities such that multi-modal approaches are able to outperform uni-modal methods which utilize information from single modality separately. Traditional approaches for multi-modal analysis can be categorized into two groups: *feature fusion* and *semantic fusion*. Feature fusion (also known as feature engineering) approaches simply conduct feature concatenations on raw features from different modalities, which is normally achieved via manual operations and has very low efficiency, as is shown in Figure 1(a). Semantic fusion first analyzes information from single modality separately in the beginning and conduct multi-modal fusion at semantic level, as is illustrated in Figure 1(b). This type of methods can maintain the explainability in semantic fusion, but fails to make full use of the rich information hidden in multi-modal.

Thanks to the success of deep neural network in computer science, a new type of approaches capable of fusing information from different modalities in hidden space at intermediate level cuts a splendid figure in multi-modal analysis, as is demonstrated in Figure 1(c). This type of methods can fully utilize the multi-modal data through learning a correlational representation for different modality in a data-driven way.

Figure 2 demonstrates a common way for multi-modal correlational representation, which is to map multi-modal data (leftmost) to a hidden representation (middle) and/or correlational representation (rightmost). Quite a few methods including deep learning can be used to learn the hidden

Manuscript received February 12, 2019; revised June 21, 2019; accepted August 28, 2019. Date of publication September 11, 2019; date of current version October 2, 2020. This work was supported in part by the National Program on Key Basic Research Project under Grant 2015CB352300, in part by the China Postdoctoral Science Foundation under Grant BX201700136, and in part by the National Natural Science Foundation of China Major Project under Grant U1611461. This article was recommended by Associate Editor F. Wu. (Corresponding author: Xin Wang.)

W. Zhu and X. Wang are with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China (e-mail: wwzhu@tsinghua.edu.cn; xin_wang@tsinghua.edu.cn).

H. Li is with Microsoft AI and Research, Redmond, WA 98052 USA (e-mail: hongzhi.li@microsoft.com).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSVT.2019.2940647

1051-8215 © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

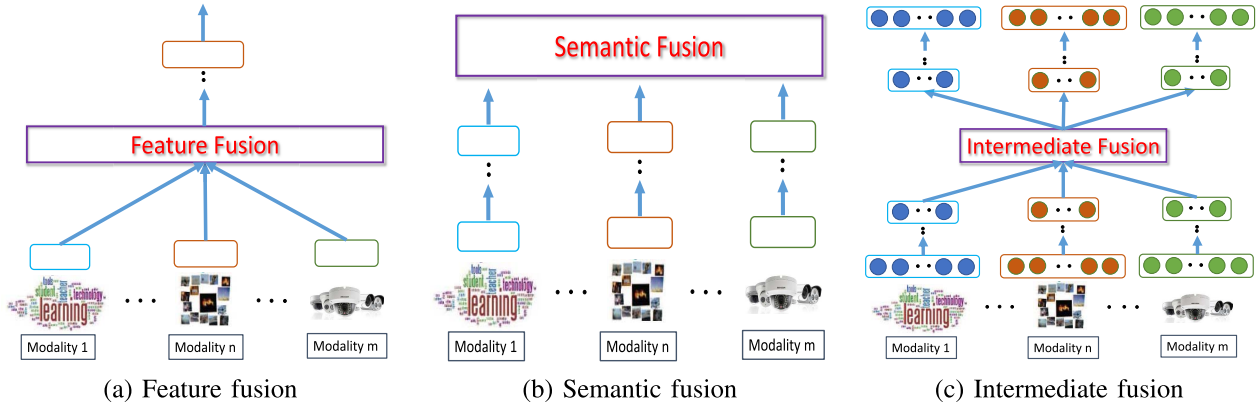


Fig. 1. Schematic diagram for semantic fusion methods and intermediate fusion methods.




Fig. 2. Correlational representation.

representation and further correlational mining techniques are necessary for the correlational representation learning.

Though being capable of handling large-scale multi-modal data, the results obtained from data-driven approaches (e.g., deep neural networks) can sometimes be unexplainable, which do not utilize too much domain knowledge, leading to a huge exploration space and low accuracy. Therefore it is very challenging to get explainable correlational representations from uncertain big data. Human, on the other hand, is capable of utilizing domain knowledge to help making decisions, resulting in high explainability and accuracy. As such, there exists a paradox between scalability and explainability, and it is desirable to figure out a balance which requires the best cooperative fusion between data and knowledge between data-driven and knowledge-driven methods.

The goal of this article is two-fold. We first give a deep and comprehensive overview for multi-modal problems in multimedia from two aspects: 1) *data-driven correlational representation*: multi-modal fusion of data from different modalities and 2) *knowledge-guided data fusion*: multi-modal fusion of data with domain knowledge. We then present our insights and thinkings on future directions for multi-modal research in the new era of artificial intelligence, and point out several promising research directions including cross-modal reasoning, cross-modal cognition and cross-modal collective intelligence, for further investigation.

One natural scientific problem is how to find a hidden representation that can best correlate information from different modalities. Several methodologies have the potential to tackle this challenge and we highlight three important categories of data-driven approaches that focus on multi-modal correlational representation: Multi-modal Deep Representation, Multi-modal Transfer Learning and Multi-modal Hashing.

Given an effective hidden correlation representation for different modality data, the next scientific research problem is how we can increase the explainability of data-driven approaches while maintaining their scalability via the guidance of domain knowledge and take advantage of their superiority. However, to the best of our knowledge, there have been no systematic or consolidated methodologies for incorporating domain knowledge into the process of cross-modal learning. We observe that there exist mainly three families of methods that may be suitable for knowledge-guided cross-modal fusion, i.e., Bayesian Inference, Teacher-student Network and Reinforcement Learning. We will elaborate our thoughts on why these three methodologies deserve further investigations for future research later in the paper.

On the other hand, many existing approaches for multi-modal oriented problems have unwittingly tried resorting to domain knowledge for the improvement of model performances. Among these methods, some utilize domain knowledge in a naive or straightforward way while some others may do it more sufficiently or elegantly. Although the existing literature is still in the preliminary stage, we believe these trials deserve attentions from researchers in the community. For a clear elaboration on the existing ideas of knowledge-guided cross-modal data fusion, we pick up four exemplar multi-modal oriented applications that require various domain knowledge, and discuss their research directions in terms of knowledge-guided multi-modal data fusion, i.e., Multi-modal Visual Question Answering, Multi-modal Video Summarization, Multi-modal Visual Pattern Mining and Multi-modal Recommendation.

In a nutshell, we present our insights on the key problems for multi-modal analysis, review some representative state-of-the-art multi-modal approaches in multimedia and summarize their characteristics in essence. Our discussions will center around the two mentioned scientific research problems in multi-modal analysis for multimedia. We discuss approaches focusing on *data-driven multi-modal correlational representation* in Section II and analyze several exemplar applications in *knowledge-guided multi-modal data fusion* in Section III. We then highlight our insights on promising research directions that may lead the next breakthrough in cross-modal intelligence, i.e., *cross-modal reasoning*, *cross-modal cognition* and

cross-modal collective intelligence. We share our opinions about why and how researchers should pay more attentions on these topics in the future in Section IV. In the end, we conclude the whole paper in Section V.

II. DATA-DRIVEN MULTI-MODAL CORRELATIONAL REPRESENTATION

In this section, we briefly introduce the concept and aim of multi-modal analysis, succinctly summarize multi-task and multi-view learning, two classic and well-documented techniques that target at learning from multiple angles, followed by our comprehensive analysis on three important categories of approaches for multi-modal correlational representation, i.e., multi-modal deep representation, multi-modal transfer learning and multi-modal hashing.

A. Multi-Task and Multi-View Learning

Multi-modal/Cross-media correlational representation seeks a way to represent different modality data in a common space such that data from every modality becomes comparable with each other and as many properties in their original spaces can be preserved in the common space as possible. As two classic methodologies, multi-task learning and multi-view learning serve as two popular ways to consider the learning process from more than one angle.

Multi-task learning aims to learn distinct tasks simultaneously by finding relationships among multiple tasks, which has been studied for roughly 20 years. One of the most important strategies in multi-task learning is to take both differences and connections among multiple tasks into account simultaneously. This strategy has been widely used in multi-label classification, face recognition, and etc. Multi-task learning can be roughly divided into two categories:

- 1) Methods forcing multiple tasks to share common parameters;
- 2) Methods mining the common latent features among multiple tasks;

Evgeniou and Pontil [2] propose *Regularized Multi-task Learning*, a representative model on common parameters which minimizes the regularization function during the learning process. Evgeniou and Pontil combine the concept of multi-task learning with single-task SVM and illustrate the connections among different single SVM tasks. They assume all tasks share a common central separation hyper plane which in turn determines the final decision boundary for the current task through an offset parameter. As for methods mining the latent features, Argyriou *et al.* [3] introduce a typical *Convex Multi-task Feature Learning* framework, laying the foundation of many later multi-task learning algorithms. Jebara, in his overview paper [4], discusses four groups of multi-task learning algorithms in terms of feature selection, kernel selection, adaptive pooling and graphical model structure. For more details, please refer to survey articles [5], [6].

Multi-view learning, as its name indicates, considers multiple views from the same input data through employing one function to model each view and jointly optimizing all the functions so that the information of multiple views can be

best exploited and the learning performance therefore can be dramatically improved. Different from multi-task learning which input data may come from multiple tasks, multi-view learning takes distinct views of the same task as input. For example, these different views can be face ID and fingerprint in recognition task, or color and words in image representation task. Multi-view learning can be categorized into three types:

- 1) Co-training: train models to achieve the maximization of the mutual consistency between two different views of the unlabeled data;
- 2) Multi-kernel learning: combine different kernels corresponding to distinct views together to achieve a performance boost;
- 3) Subspace learning: assume that there exists a common latent subspace shared by all views such that different view data can be generated from this shared latent subspace;

Besides, there are two principles widely adopted to make sure that information from multiple views can be sufficiently utilized.

- 1) Consensus principle (used by co-training): maximize the mutual consistency between two views by requiring the two hypotheses to be as consistent as possible, i.e.,

$$P(f^1 \neq f^2) \geq \max\{P_{err}(f^1), P_{err}(f^2)\}, \quad (1)$$

where $P(f^1 \neq f^2)$ is the disagreement rate between two hypotheses from the corresponding two views and $P_{err}(f^1)$, $P_{err}(f^2)$ are error rates of single hypothesis f^1 , f^2 . Thus the error rate of each single hypothesis is indirectly minimized through minimizing $P(f^1 \neq f^2)$.

- 2) Complementary principle: every distinct view has some unique information which is not possessed by others. Thus we may improve the learning performance by making full use of complementary knowledge from different views can result in an improvement for the learning performance.

Readers with interests may refer to overview papers [7], [8] on multi-view learning for more detailed information.

Besides the “pure” multi-view learning, others have investigated metric fusion [9] or similarity learning [10] based on the multi-view data as well. There are also some works combining “multi-task” and “multi-view” together whose details can be found in [11]–[14]. We note that both multi-task and multi-view learning are not customized for multi-modal correlational representation. This being the case, we highlight three promising groups of multi-modal methods designed specifically for multi-modal data and discuss them in the rest of this section.

B. Multi-Modal Deep Representation

Before deep learning is widely used in computer vision and multimedia research works, multi-modal methods can be mainly divided into two groups:

- Feature-fusion approaches [15], [16]: aggregate features extracted from each modality and feed the aggregated features to the model (similar to the process of feature engineering);

- Semantic-fusion approaches [17]–[19]: feed features from each modality into the model separately and combine the results from all the models to get the final results (similar to the methodology of ensemble learning);

A general comment on these two strategies is that feature-fusion is suitable for problems whose modalities share many correlated features while the semantic-fusion approach fits for those who have significantly uncorrelated modalities.

The prevalent success of deep learning brings us a new option for multi-modal fusion — intermediate-fusion. Thanks to deep neural network that provides variable number of layers for latent representations, it becomes flexible to choose when and which layer(s) can be used to fuse data from different modalities [20]–[24].

On the other side, it is also possible to categorize multi-modal methods based on whether they are discriminative, generative or both (hybrid). Discriminative models [19], [25]–[29] usually learn conditional distributions of labels given features. Generative models [30]–[39] tend to learn their joint distributions. And hybrid models [40]–[42] learn both conditional distributions and joint distributions by combining discriminative and generative parts. We refer readers to a survey paper on multi-modal deep learning [43] for further information. Different from the survey paper [43], in this work we focus on multi-modal scenarios in areas related to multimedia, which includes but is not limited to deep neural network based architectures.

Next, we discuss two works [44], [45] utilizing the idea of domain adaption to bridge the deep representations of different modalities, which is not covered by the referred survey paper. We start from the classic work [46] by Yosinski *et al.* on the transferability in deep neural networks. The conclusion is that for a given deep neural network, deeper representation layers are more dependent on the specific task to be solved. While the shallow layers are responsible for capturing more general features. This principle inspires us to adapt those deeper representation layers in deep learning for multi-modal tasks. As such, Tzeng *et al.* propose a representative method (called DCC) [47] to adapt deeper layers by employing *Maximum Mean Discrepancy* (MMD) [48] to reduce the disagreement between two modalities on the seventh layer (before softmax) of a eight-layer AlexNet. DCC has two drawbacks: 1) it only adapts one single layer (i.e., the seventh layer) in the deep neural network (AlexNet). It may not be enough as Yosinski *et al.* [46] point out that more than one layer is transferable and 2) it adopts a single-kernel MMD (SK-MMD) which may not serve as the optimal kernel. To tackle these weakness, Long *et al.* propose a *Deep Adaptation Network* (DAN) [44] that adapts three deep layers simultaneously through a multi-kernel MMD (MK-MMD) which is capable of constructing the final kernel by combining multiple kernels together in *Reproducing Kernel Hilbert Space* (RKHS). Figure 3 illustrates the architecture of *Deep Adaptation Network* (DAN).

The objective of DAN then consists of two components:

- 1) Deep adaptation which matches distributions of representation layers in multiple modalities,





Fig. 3. Deep adaptation network, figure from [44].



(a) Match Marginal Distributions (b) Match Joint Distributions

Fig. 4. Matching marginal distributions v.s. matching joint distributions, figure from [45].

- 2) Optimal matching which maximizes two-sample test power by MK-MMD in RKHS.

DAN, on the other side, is also not perfect because it matches the marginal distributions $P(\mathbf{x})$ and $Q(\mathbf{x})$ rather than the joint distributions $P(\mathbf{x}, \mathbf{y})$ and $Q(\mathbf{x}, \mathbf{y})$. As is shown in Figure 4, matching the joint distributions can achieve a better performance than matching the marginal distributions. Thus a model based on *Joint Adaptation Network* (JAN) [45] is proposed by Long *et al.* to match the joint distributions between deep representations from different modalities. Figure 5 illustrates the structures of Joint Adaptation Network (JAN) and its adversarial version (JAN-A). As the RKHS is normally high-dimensional or even infinite-dimensional, Gaussian kernel mapping samples to infinite spaces is usually adopted as the kernel function, and the final bandwidth parameter is selected according to empirical experiences.

We note that although these two works study the transferring representations between different (two) modalities, their learning processes are bidirectional and the proposed models can be tested on any two modalities without a fixed requirement of “source” or “target” domain in the experiments. Therefore, we group these two works in the category of multi-modal deep representation rather than multi-modal transfer learning. Besides domain adaptation, there are also works aiming at feature learning by means of deep neural networks, such as a recent work [49] by Liu *et al.*

C. Multi-Modal Transfer Learning

In the past decade, researchers have developed plenty of good models that can achieve fairly good performances on large amounts of labeled data including images, sentences etc. which is for the same task and in the same domain (e.g., predicting image class labels given images and their

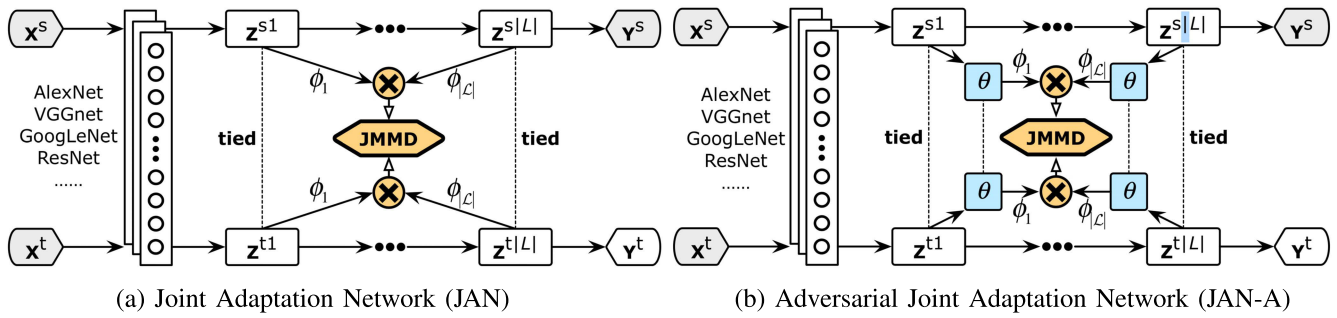


Fig. 5. Joint adaptation network and its adversarial version, figure from [45].

corresponding labels as input for training). However, these models still suffer in situations containing new scenarios that the models have never taken into account in their training phases. For instance, a model trained on detecting pedestrians during day-time may experience a deterioration in performance when being applied to detect bicyclists during night-time. Transfer learning aims to enhance the ability of a model to generalize and transfer learned knowledge (from source domain) to new scenarios (target domains), which has been an active research topic for quite a long time before the proliferation of deep neural network and we refer readers to an excellent survey [50] published in 2010 for details about early models. Multi-modal transfer learning particularly focuses on transferring knowledge from one modality (source) to a different one (target). This method enables us to handle labeled data in a new modality through leveraging the existing labeled data in the original modality. Different from Pan's work [50], we focus on multi-modal transfer learning in multi-media through adding more recent and advanced technologies including deep learning methods and so forth in this section.

One benefit brought by deep neural network is the famous deep convolutional neural network (CNN) features trained on ImageNet [51] which can be used as pre-trained features for new task(s) in the target domain. The credits should be given to CNN's capability of learning the basic components such as edges and shapes which serve as general elements in images. Thus a straightforward way to handle a new task can be simply applying some pre-trained CNN features on ImageNet to this new task, with parameters either fixed or slightly tuned under a very small learning rate [52].

However, CNN is designed specifically for image data (pixels) and what if we have data from other domains such as text or signal data? The idea of domain adaptation which tries to preserve general knowledge that does not change in different domains can serve as an appropriate candidate. Several works on natural language processing (NLP) [53], [54] and computer vision (CV) [55] have gained success through employing stacked denoising autoencoders to learn the domain-invariant deep representations.

Besides, forcing the learned representations of source domain and target domain to be similar to each other may also be an option because this procedure is able to remove domain-specific features while keeping common features shared across the two domains. It is possible to achieve this goal through either applying the strategy to the initial

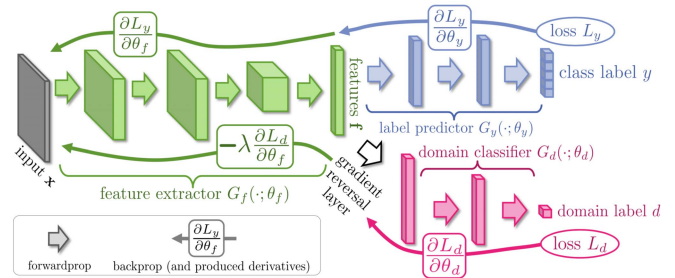


Fig. 6. Indistinguishable domains with a gradient reversal layer (in pink), figure from [59].

representations before training [56], [57] or ensuring the representations of source and target domains to be similar during the training process [47], [58].

The works [59], [60] by Ganin *et al.* propose a novel setting which makes the deep feature extraction part of the model produce features incapable of distinguishing between source and target domain. As is shown by the pink part in Figure 6, this can be done through adding a gradient reversal layer that multiplies the gradient by a certain negative constant during back propagation. In other words, the designed model in Figure 6 is able to minimize the label classification error in the source domain and fails to distinguish between different domains simultaneously, forcing the feature extractor to generate features beneficial for knowledge transfer.

Aside from the traditional images and texts data that are widely used in transfer learning, other recent works study multi-modal transfer learning based on various data including audio and video [61], head movement and co-speech [62], Alzheimer's disease (AD) and mild cognitive impairment (MCI) [63], [64] etc.

In particular, a normal characteristic of multi-modal data in transfer learning is that the labeled data in the source task \mathcal{T}_S is far more than that in target task \mathcal{T}_T and thus the source domain is more resource-rich (reliable) than the target domain, resulting a very severe unbalanced problem as demonstrated by Figure 7. To tackle this challenge, Wang *et al.* [65] in their recent work develop a deep asymmetric transfer network (DATN) that can adapt the classifier of source task to target task through learning a transfer function which maps the deep representation in the target domain to that in the source domain. The main framework of DATN is illustrated

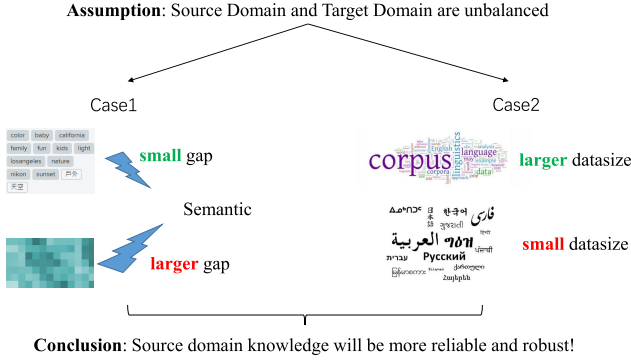


Fig. 7. The unbalanced problem between the resource-rich source domain and resource-poor target domain.

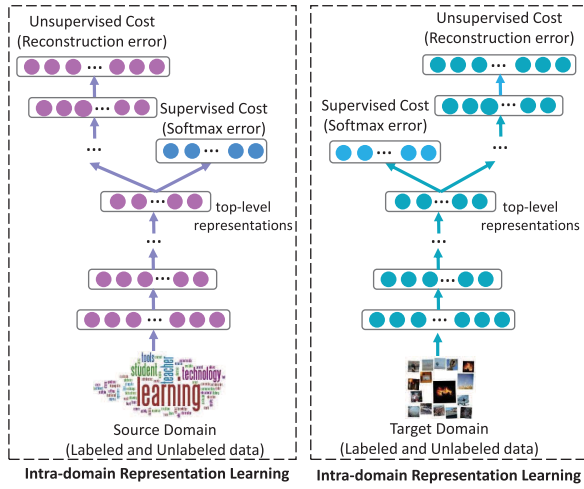


Fig. 8. Model initialization of DATN model, figure from [65].

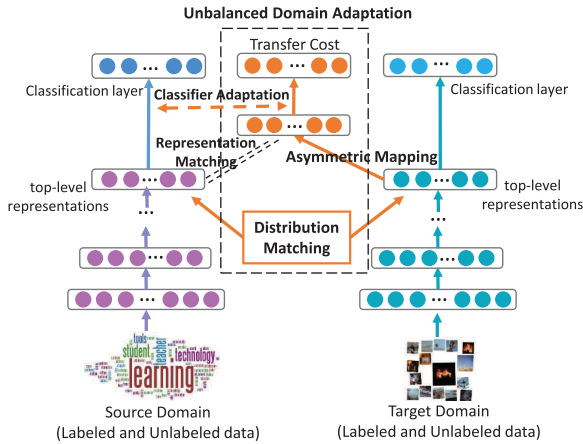


Fig. 9. Asymmetric transfer of DATN model, figure from [65].

in Figure 8 and Figure 9 where the initialization of deep representations in each modality is conducted separately through an autoencoder (shown in Figure 8) and the asymmetric transfer together with adapting the source task classifier is achieved through a transfer function (shown in Figure 9). The asymmetric transfer process consists of three parts:

- Asymmetric mapping through transfer function G :

$$\mathcal{L}_{pair} = \|Z_S^c - Z_T^c \cdot G\|_F^2 + \lambda' \|G\|_F^2 \quad (2)$$

- Source classifier adaptation:

$$\mathcal{L}_{trans} = -\frac{1}{n_T^L} \sum_{i=1}^{n_T^L} \sum_{j=1}^k 1\{y_{T_i} = j\} \log \frac{e^{\mathbf{z}_{T_i}^L \cdot G \cdot \vartheta_{S_j}}}{\sum_{l=1}^k e^{\mathbf{z}_{T_i}^L \cdot G \cdot \vartheta_{S_l}}} \quad (3)$$

- Top-level distribution matching:

$$\begin{aligned} \mathcal{L}_{unsup} &= MMD(Z_S, Z_T) \\ &= \left\| \frac{1}{n_S} \sum_{i=1}^{n_S} \mathbf{z}_{S_i} - \frac{1}{n_T} \sum_{i=1}^{n_T} \mathbf{z}_{T_i} \right\|_2^2, \end{aligned} \quad (4)$$

where Z_S, Z_T denote the top-level deep representations for the source task, target task respectively and MMD refers to *Maximum Mean Discrepancy* [48].

By putting (2), (3) and (4) together, the overall objectives can be expressed as follows:

$$\mathcal{J}^{cross} = \mathcal{L}_{pair} + \alpha \mathcal{L}_{trans} + \beta \mathcal{L}_{unsup} + \mathcal{L}_{reg}, \quad (5)$$

where \mathcal{L}_{reg} is the regularization term.

D. Multi-Modal Hashing

As the early works on multi-source hashing, multiple feature hashing [66], [67] and composite hashing [68] examine efficient hashing with multiple features or information sources taken into account. These works focus on the problem of returning the same types of items as the queries, which though have a close relation to multi-modal hashing, are not specifically designed for retrieving different sorts of items from a given query.

In the setting of multi-modal hashing, we aim at retrieving a heterogeneous type of items (e.g., images) given a corresponding input query (e.g., texts describing the images). Normally, multi-modal hashing maps data from different modalities into some common space (e.g., Hamming space) in which the hash codes obtained from multi-modality data can be directly compared.

Data from different modalities may share one unified hash code or possess separate hash codes in the new space. Good multi-modal hashing models should be capable of designing good hash functions as well as efficiently bridging the gaps between different domains for fast and accurate similarity search across multiple modalities [69]–[90]. In particular, cross view hashing [70] extends composite hashing to handle multi-view settings through summing over Hamming distance for each view:

$$d_{ij} = \sum_{k=1}^K d(y_i^{(k)}, y_j^{(k)}) + \sum_{k=1}^K \sum_{k' > k}^K d(y_i^{(k)}, y_j^{(k')}). \quad (6)$$

Multi-modal latent binary embedding [72] utilizes probability theory to learn hash function in the multi-modal setting whose graphic model is shown in Figure 10.

Co-regularized hashing, taking the regularization as an entry point, learns a multi-modal hash function with the help of a

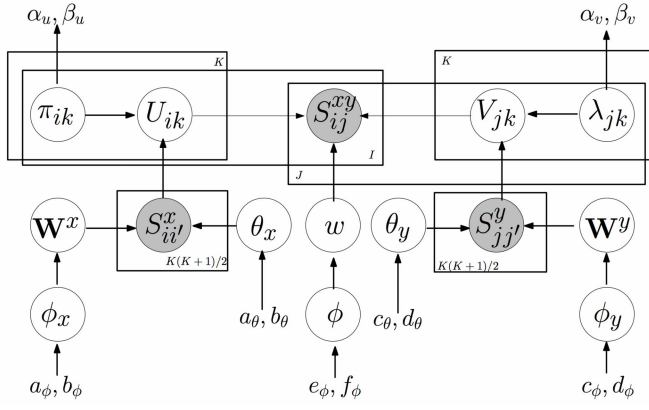


Fig. 10. Graphic model of multi-modal latent binary embedding, figure from [72].

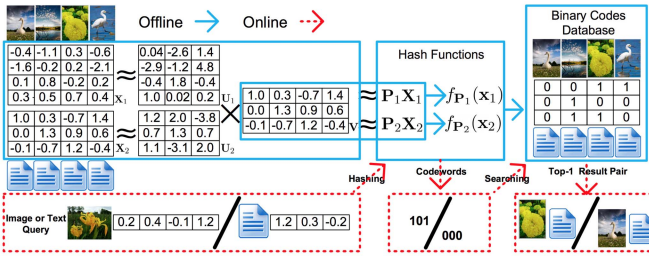


Fig. 11. Framework of collective matrix factorization hashing, figure from [77].

boosted co-regularization strategy, whose objective function is as follows:

$$\mathcal{O} = \frac{1}{I} \sum_{i=1}^I l_i^x + \frac{1}{J} \sum_{j=1}^J l_j^y + \gamma \sum_{n=1}^N \omega_n l_n^* + \frac{\lambda_x}{2} \|\omega_x\|^2 + \frac{\lambda_y}{2} \|\omega_y\|^2, \quad (7)$$

where l_i^x and l_j^y are intra-modality losses and l_n^* is inter-modality loss. Motivated by the need for scalability and training hash functions on large scale multi-modal dataset, semantic correlation maximization hashing [80] avoids explicit computation of pairwise similarity matrix through proposing a sequential hashing learning method with closed-form solution to each bit. Collective matrix factorization hashing [77] borrows the idea of collective matrix factorization to learn cross-modality hash functions by decomposing feature matrices from two different modalities (e.g., $X^{(1)}, X^{(2)}$) jointly with the constraint $V_1 = V_2 = V$:

$$\lambda \|X^{(1)} - U_1 V\|^2 + (1 - \lambda) \|X^{(2)} - U_2 V\|^2. \quad (8)$$

The whole framework of collective matrix factorization hashing is presented in Figure 11.

Quantized correlation hashing [85] is the first to integrate the process of hash function learning with quantization for multi-modal hashing by transforming multi-modal objective function into a single-modal formulation. Semantics preserving hashing [84] maps the given affinity matrix A into a

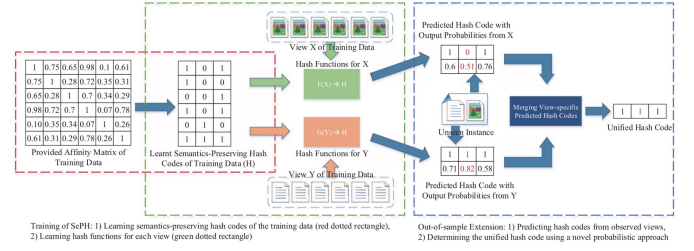


Fig. 12. Framework of semantics preserving hashing with two views, figure from [84].

probability distribution \mathcal{P} and matches it with another probability distribution \mathcal{Q} transformed from pairwise Hamming distances between hash codes in the Hamming space through minimizing the KL-divergence between \mathcal{P} and \mathcal{Q} . The objective function is:

$$\Phi = \min_{\hat{H} \in \mathbb{R}^{n \times d_c}} \sum_{i \neq j} p_{ij} \log \frac{p_{ij}}{q_{ij}} + \frac{\alpha}{C} \|\hat{H}\| - \mathbf{I}\|^2, \quad (9)$$

where \hat{H} is the relaxed hash code matrix and

$$q_{ij} = \frac{(1 + \frac{1}{4} \|\hat{H}_i - \hat{H}_j\|^2)^{-1}}{\sum_{k \neq m} (1 + \frac{1}{4} \|\hat{H}_k - \hat{H}_m\|^2)^{-1}},$$

which utilizes the Student t-distribution of degree one to transform Hamming distances into probabilities. Besides, $p_{ij} = \frac{A_{ij}}{\sum_{i \neq j} A_{ij}}$ where A_{ij} is the element of i -th row and j -th column in affinity matrix A , representing the given affinity between i and j . The overall structure of semantics preserving hashing is illustrated in Figure 12.

In general, inter-media hashing [75], cross view hashing [70], sequential spectral learning to hash [71] are unsupervised hashing models extending spectral hashing [91] to cross-modal scenario by defining the distance between documents in Hamming space and aligning the hash codes from all modalities with the given inter-document similarity. On the other hand, data fusion hashing [69], semantic correlation maximization hashing [80], collective matrix factorization hashing [77], similarity-preserving hashing [79], sparse multi-modal hashing [81], multi-modal latent binary embedding [72], semantics-preserving cross-view hashing [84] and co-regularized hashing [73] all belong to supervised hashing approaches which take the pairwise similarity information between two objects from different domains (modalities) as input and require the hash codes of these paired objects in Hamming space across different domains to be similar through the maximizing similarity-agreement criterion [69], minimizing similarity-difference criterion [80], collective matrix factorization [77] or inverted squared function [73].

Given the recent success of deep neural networks, there also have been several works [87]–[90], [92]–[96] on combining hashing with deep structures for cross-modal similarity search.

Deep cross-modal hashing [90], whose framework is presented in Figure 13, employs a convolutional neural network (CNN) that takes image data as input and a fully connected deep neural network that takes text data as input




Fig. 13. The learning architecture of deep cross-modal hashing, figure from [90].

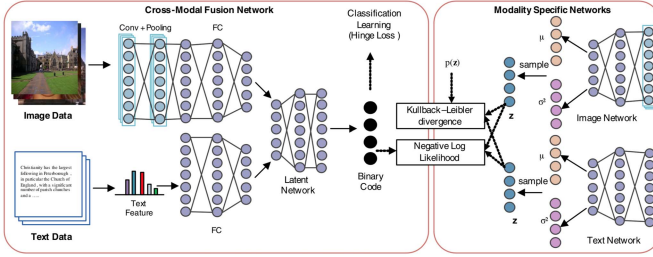


Fig. 14. The end-to-end learning framework of cross-modal deep variational hashing, figure from [92].




Fig. 15. The architecture of deep visual-semantic hashing, figure from [88].

to optimize the binary codes and parameters from two neural networks iteratively.

A variational version of deep cross-modal hashing, cross-modal deep variational hashing [92], adopts a two-step learning procedure:

- 1) Learn a fusion network and a joint binary code matrix shared by two modalities simultaneously through an alternative optimization procedure, which is similar to deep cross-modal hashing.
- 2) Learn a modality specific neural network for each modality such that the top-level representation is as similar as possible to the binary codes obtained from the fusion network and also that the approximated posterior distribution can be as close as possible to the KL-divergence regularized prior distribution.

For comparison with deep cross-modal hashing, we refer readers to Figure 14 for the end-to-end learning framework of cross-modal deep variational hashing. As is shown in Figure 15, deep visual-semantic hashing [88] proposes an end-to-end image-sentence (each image is attached with at least one sentence) cross-modal hashing algorithm which utilizes convolutional neural network (CNN) to handle image data and long short term memory (LSTM) to handle sentence data such that a joint embedding space for both modalities as well as a separate structure for each modality can be learnt under the




Fig. 16. The flowchart of deep multi-modal hashing with orthogonal regularization, figure from [87].




Fig. 17. Model pretraining of DMHOR model, figure from [87].

guidance of different losses including pairwise loss, cosine hinge loss, bit-wise margin loss, squared loss etc.

Aside from CNN, some other works [87], [89], [93] also feed data from different modalities to autoencoder (AE) which serves as an adequate tool for both model initialization and unsupervised learning. Particularly, deep multi-modal hashing with orthogonal regularization [87] whose model flowchart is displayed in Figure 16, recognizes the phenomenon of redundant information in deep multi-modal representation and proposes an orthogonal structure to reduce this redundancy with theoretical guarantee while keeping the learnt compact hash codes accurate. As is illustrated in Figure 17, a multi-modal deep belief network (DBN) consisting of several DBNs (one for each modality) and a joint Restricted Boltzmann Machine (RBM) is first developed to correlate high-level representations of data from different modalities for the purpose of pretraining. Then, to learn an adequate multi-modal representation that preserves intra-modality and inter-modality simultaneously, a multi-modal autoencoder (MAE) is developed to capture the joint correlations for different modalities and a cross-modal autoencoder (CAE) is explored to enable the reconstruction of representations in any modality from data in an arbitrary modality. The left part of Figure 18 shows the structure of MAE whose loss function is shown as follows:

$$L_{v\bar{t}}(\mathbf{x}_v, \mathbf{x}_t; \theta) = \frac{1}{2}(\|\hat{\mathbf{x}}_v - \mathbf{x}_v\|_2^2 + \|\hat{\mathbf{x}}_v - \mathbf{x}_t\|_2^2), \quad (10)$$

where $\hat{\mathbf{x}}_v$ is the reconstruction of \mathbf{x}_v and $\hat{\mathbf{x}}_t$ is the reconstruction of \mathbf{x}_t . The right part of Figure 18 demonstrates the structure of image-only CAE whose loss function can be expressed in the following:

$$L_{o\bar{t}}(x_v, x_t; \theta) = \frac{1}{2}(\|\hat{\mathbf{x}}_v^I - \mathbf{x}_v\|_2^2 + \|\hat{\mathbf{x}}_t^I - \mathbf{x}_t\|_2^2), \quad (11)$$

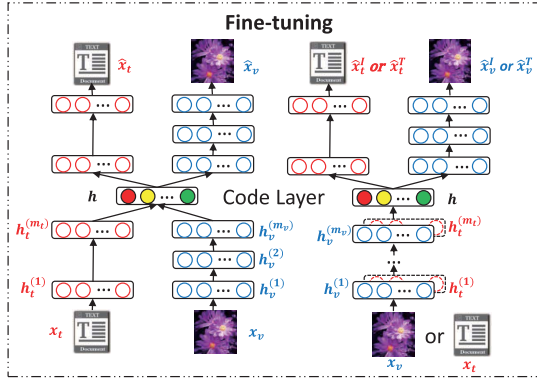


Fig. 18. Cross-modal fine tuning of DMHOR model, figure from [87].

where the subscript $v\bar{t}$ denotes the input of the provided image pathway when the corresponding text pathway is absent. \hat{x}_v^I is the reconstruction of x_v in the image pathway and \hat{x}_t^I is the reconstruction of x_t in the text pathway. The missing modality will be set to zero in the joining code layer for the calculation of \hat{x}_v^I and \hat{x}_t^I . Thus the overall objective function with only two modalities (image and text) can be formulated as follows (loss function of text-only CAE $L_{\bar{v}t}$ can be formulated in a way similar to (11)):

$$\begin{aligned} \min_{\theta} \quad & L_{MDAE}(X_v, X_t; \theta) \\ &= \frac{1}{n} \sum_{i=1}^n (L_{vt} + L_{\bar{v}t} + L_{v\bar{t}}) + L_{reg} \\ \text{s.t.} \quad & \frac{1}{n} \tilde{H}^T \cdot \tilde{H} = I, \end{aligned} \quad (12)$$

where L_{reg} is a L2-norm regularizer term of weight matrix preventing overfitting and the constraint $\frac{1}{n} \tilde{H}^T \cdot \tilde{H} = I$ ensures the orthogonality of the hash codes to reduce the redundant information.

Last but not least, the idea of adversarial training is also adopted in cross-modal hashing, such as semi-supervised cross-modal hashing [94], self-supervised adversarial hashing [95], cycle-consistent deep generative hashing [96].

III. KNOWLEDGE-GUIDED MULTI-MODAL ANALYSIS

One intelligent aspect of human being is that we are able to make decisions by resorting to domain knowledge from relevant fields or domains. This motivates the advent of knowledge-guided multi-modal approaches which adopt a more intelligent and promising multi-modal way through utilizing complementary external domain knowledge to boost the model performance in multimedia. In this section, we first present three types of methods adequate for the fusion of data and knowledge, then discuss several exemplar applications that require knowledge-guided fusion.

A. Approaches for Knowledge-Guided Fusion

There are three mainly families of methods that are suitable for knowledge-guided cross-modal fusion, i.e., Bayesian Inference, Teacher-student Network and Reinforcement Learning, which deserves further investigations for future research.

Bayesian theory [97] has been a very popular tool in statistics. Bayesian inference [98]–[100] aims to simulate the inference ability of human through encoding some “prior” knowledge into the model. Thus incorporating domain knowledge via Bayesian prior would be a good option for knowledge-guided multi-modal fusion. Since the deep neural networks usually have quite complex structures, teacher-student network [101] is originally proposed to compress the deep model (student network) via the guidance of a well-trained network (teacher network). It has also been applied for information/knowledge transfer between image sets [102], RGB images and depth images [103], as well as video sets [104]. Therefore, distilling useful domain knowledge through a teacher network and using it as guidance in cross-modal data fusion could also be an appropriate direction. Reinforcement learning [105], [106] aims at taking suitable actions to maximize rewards in certain situations. It has been a well-established machine learning research topic with wide applications, particularly in robotics [107] in the past decades. As such, utilizing domain knowledge to guide the reward/feedback in reinforcement framework seems to be another promising way to handle knowledge-guided multi-modal fusion.

B. Exemplar Applications of Data and Knowledge Fusion

Since different problems may require different domain knowledge, we discuss four exemplar research topics covering visual question answering, video summarization, visual pattern mining and recommendation from a knowledge-guided multi-modal perspective for a better illustration.

1) *Multi-Modal Visual Question Answering*: Visual Question Answering (VQA) is a challenging task, which bridges Computer Vision (CV) and Natural Language Processing (NLP) via jointly understanding visual information and natural language. Given an image and a related textual question, VQA systems are supposed to correctly answer the question based on the image, making VQA intrinsically cross-modal since it involves an image and a relevant question. In order to achieve a joint deep understanding of visual and natural language, a VQA task is designed as a practical setting to evaluate the capability of an algorithm for extracting high-level visual information and reasoning on the extracted information. VQA is very challenging not only for its requirement of bridging visual and textual modalities but also for the required versatile abilities ranging from object recognition and localization to high-level reasoning and common-sense knowledge learning. We will briefly describe the conventional cross-media architecture of VQA systems as well as several advanced techniques for connecting visual and textual modalities, followed by discussions on some issues in VQA systems and pioneering works that may lead the future research.

Conventional approaches for VQA train a neural network using (image, question, answer) triplets as supervision in an end-to-end way, establishing a mapping from the given image and question input to one of the candidate answers. Here the core idea is to learn a unified embedding of image and question. The input image will be passed through a convolutional neural network pretrained for image classification (e.g., ResNet) to obtain an image representation, i.e., a

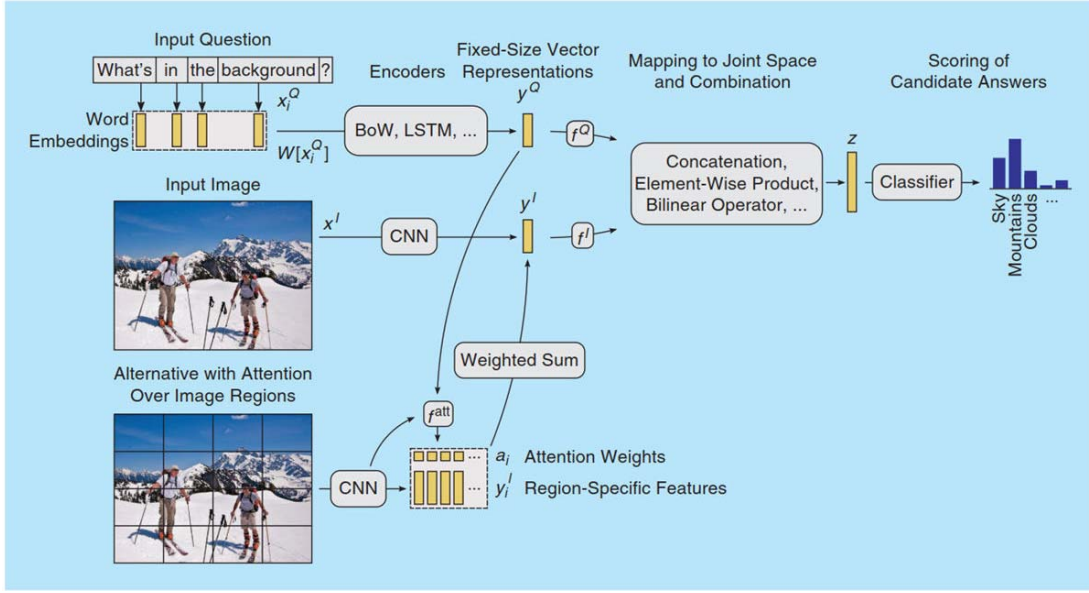


Fig. 19. An illustration of the conventional VQA approach, figure from [108].

fixed-length vector. Meanwhile, each word in the textual questions will first be embedded into a continuous space by some well-established methods (e.g., one-hot encoding, or look up in a pretrained word-embedding matrix), and then the sequence of words will be encoded into a fixed-length vector through bag-of-words or recurrent neural network to capture the sequential relationships among words. Upon obtaining the feature representations of image and question, each of them will be embedded into a common space where the combination of image and question representation will then be conducted. The embedding function is typically implemented as additional layers of neural networks, and straightforward options for combining the embedded features include concatenation and Hadamard (element-wise) multiplication in the common space. This family of works can be regarded as the simplest cross-modal fusion methods. An illustration diagram is provided in Figure 19 which is taken from [108].

Let us now turn to some advanced techniques used for modeling cross-modal interactions. Upon understanding the visual world, humans have the ability to focus on specific region(s) instead of the entire scene. Inspired by this human-possessed ability, *attention mechanism* [109] has been widely used in order to address the “where to look” problem, resulting in one of the most effective improvement for various tasks including object recognition, reading comprehension, image captioning and visual question answering etc. The core idea of attention is allowing the neural network to learn what regions to focus on, by means of modeling interactions between the content and side information in relevant regions. To adapt visual attention in VQA models, region-specific local image features are first extracted from an intermediate layer (before the last pooling operation) of a pretrained CNN. Then a scalar attention weight for each region is calculated using both textual question and local visual features, which indicates the relevance of the given region and question. Finally, the image features can be represented as a weighted sum of the local visual region

features. As an essential component for many VQA models, quite a few variations of attention mechanism have been proposed in the literature for modeling the interactions between textual and visual modalities [110]–[113]. Yang *et al.* [110] present a stacked attention network (SAN) which uses the semantic feature of the textual question as a query to search for those relevant visual regions through a multi-layer architecture. Lu *et al.* [111] propose a hierarchical co-attention (HieCoAtt) model that combines “visual attention” and “question attention” via conducting a question-guided attention on image and a image-guided attention on question, as is shown in Eq (13):

$$\begin{aligned}
 H^v &= \tanh(W_v V + (W_q Q)C), \\
 H^q &= \tanh(W_q Q + (W_v V)C^T), \\
 a^v &= \text{softmax}(w_{hv}^T H^v), \\
 a^q &= \text{softmax}(w_{hq}^T H^q), \\
 \hat{v} &= \sum_{n=1}^N a_n^v v_n, \quad \hat{q} = \sum_{t=1}^T a_t^q q_t,
 \end{aligned} \tag{13}$$

where H^v and H^q are latent deep representations of visual image features and textual question features respectively. $C \in \mathcal{R}^{T \times N}$ is an affinity matrix whose entries represent the similarities between the question features $Q \in \mathcal{R}^{d \times T}$ and image features $V \in \mathcal{R}^{d \times N}$. Actually the affinity matrix $C \in \mathcal{R}^{T \times N}$ can also be regarded as a connection between the question attention space to the image attention space. The attention weights for each image region v_n and word q_t are denoted as $a^v \in \mathcal{R}^N$ and $a^q \in \mathcal{R}^T$, respectively. Instead of performing attentions on spatial feature maps (e.g., 7×7 ResNet101 [114] *res5c* feature maps) as previous works, Anderson *et al.* [113] introduce a bottom-up visual attention mechanism that enables object-level attention based on image regions obtained through Faster R-CNN [115], as is shown in Figure 20.




Fig. 20. Spatial-based versus object-based visual features, figure from [113].

Rather than adopting the naive element-wise production or concatenation, another group of works resort to the bilinear pooling model as well as its variations [116]–[119] to achieve a great success by computing the outer product of two vectors to enable interactions among elements in both vectors. Denoting $\mathbf{v} \in \mathcal{R}^{d_v}$ and $\mathbf{q} \in \mathcal{R}^{d_q}$ as image(visual) and question feature vectors, the classification vector $\mathbf{y} \in \mathcal{R}^{|A|}$ can be calculated by Eq (14):

$$\mathbf{y} = (\mathcal{T} \times_1 \mathbf{q}) \times_2 \mathbf{v}, \quad (14)$$

where $\mathcal{T} \in \mathcal{R}^{d_q \times d_v \times |A|}$ is the parameter tensor, the operator \times_i denotes the i -mode product between a tensor and a matrix, which suffers from high dimensionality ($d_q \times d_v \times |A|$). Fukui *et al.* [116] propose a Multi-modal Compact Bilinear pooling (MCB) algorithm which adopts a sampling-based computation and projection method to reduce dimensionality while preserving the performance of full bilinear pooling. Kim *et al.* [117] present a Multi-modal Low-rank Bilinear pooling (MLB) model that forces the rank of the weight tensor to be low, as is shown in Eq (15):

$$\mathbf{y} = \mathbf{P}^\top \left(\mathbf{W}_q^\top \mathbf{q} \circ \mathbf{W}_v^\top \mathbf{v} \right) + \mathbf{b}, \quad (15)$$

where \mathbf{W} , \mathbf{P} , \mathbf{b} are model parameters and \circ denotes the Hadamard product operator. Yu *et al.* [119] propose the Multi-modal Factorized Bilinear (MFB) pooling by utilizing some tricks in matrix factorization to improve the convergence rate and reduce the number of parameters. By combining low-rank matrix constraint with Tucker decomposition, i.e., $\mathcal{T} = ((\mathcal{T}_c \times_1 \mathbf{W}_q) \times_2 \mathbf{W}_v) \times_3 \mathbf{W}_o$, Ben-Younes *et al.* [118] introduce MUTAN, and the combination is expressed in Eq (16):

$$\mathbf{y} = \left((\mathcal{T}_c \times_1 (\mathbf{q}^\top \mathbf{W}_q)) \times_2 (\mathbf{v}^\top \mathbf{W}_v) \right) \times_3 \mathbf{W}_o, \quad (16)$$

where $\mathbf{W}_q \in \mathcal{R}^{d_q \times t_q}$, $\mathbf{W}_v \in \mathcal{R}^{d_v \times t_v}$, $\mathbf{W}_o \in \mathcal{R}^{|A| \times t_o}$, and $\mathcal{T}_c \in \mathcal{R}^{t_q \times t_v \times t_o}$.

Recent studies have pointed out that current VQA models heavily rely on biases in different datasets and many existing methods overly exploit these biases to “correctly” answer

questions without considering the real visual information. For example, a model may answer “2” to any question starting with “How many” without really counting the numbers because the model learns (from biases) that answering “2” is the best guess for this dataset. As a consequence, even “blind” model can achieve satisfying results without well understanding the questions and images. Many efforts, such as building more balanced datasets [120], [121] and enforcing more transparent model designs, have been made to alleviate this issue.

a) *Multi-modal fusion*: Instead of building models merely based on visual and textual features via deep neural networks, several works seek for structural representations to handle the multi-modal nature in VQA. A series of works related to compositional models [123]–[127] have shown exciting visual reasoning abilities on synthetic datasets. Their fundamental ideas are to compose instance-specific networks based on compositional structures of questions via a collection of jointly trained neural modules. This can be regarded as a process of multi-modal information fusion where the question information is encoded inside the network architecture. An example of neural module networks is shown in Figure 22.

Another promising attempt is to exploit graph-structured representations in VQA [128], [129], where object relations and language structures are represented as graphs whose structure information can be further explored via techniques such as graph convolutional networks (GCN). As is shown in Figure 23, Norcliffe-Brown *et al.* [129] propose a graph-based approach for visual question answering. This work exploits a graph convolution-based method [130] to learn new visual representations from spatial graphs, where graph nodes are bounding boxes for object detections and graph edges are learned via an attention-based “Graph Learner” component. The graph convolution operator is defined at kernel k for node i as:

$$\mathbf{f}_k(i) = \sum_{j \in \mathcal{N}(i)} w_k(\mathbf{u}(i, j)) \mathbf{v}_j \alpha_{ij}, \quad k = 1, 2, \dots, K \quad (17)$$

where $\mathbf{u}(i, j)$ is a pseudo-coordinate function describing the relative spatial positions of vertex i and j , $w_k(\mathbf{u})$ is the k_{th} convolution kernel, $\mathcal{N}(i)$ denotes the neighbourhood of vertex i , \mathbf{v}_j is the associated feature vector of vertices j , α_{ij} is the edge weight produced by the “Graph Learner” component. In the end, the convolutional feature of vertex i is obtained through a concatenation over the K kernels.

b) *Incorporating domain knowledge*: In some situations, visual questions are not answerable by analyzing the questions and visual information themselves alone. Correctly answering visual questions may require extra information ranging from common-sense to expert domain knowledge, which is far beyond what the training dataset can provide. Thus it will be attractive to incorporate useful domain knowledge retrieved from other sources into VQA systems. Several pioneering works [131], [132] explore explicit reasoning on visual concepts and supporting facts in structural knowledge base, where raw visual signals are transformed into semantic symbols. In contrast to above symbolic-based methods, Li *et al.* [122] propose a Knowledge-incorporated Dynamic Memory Network (KDMN) framework which incorporates

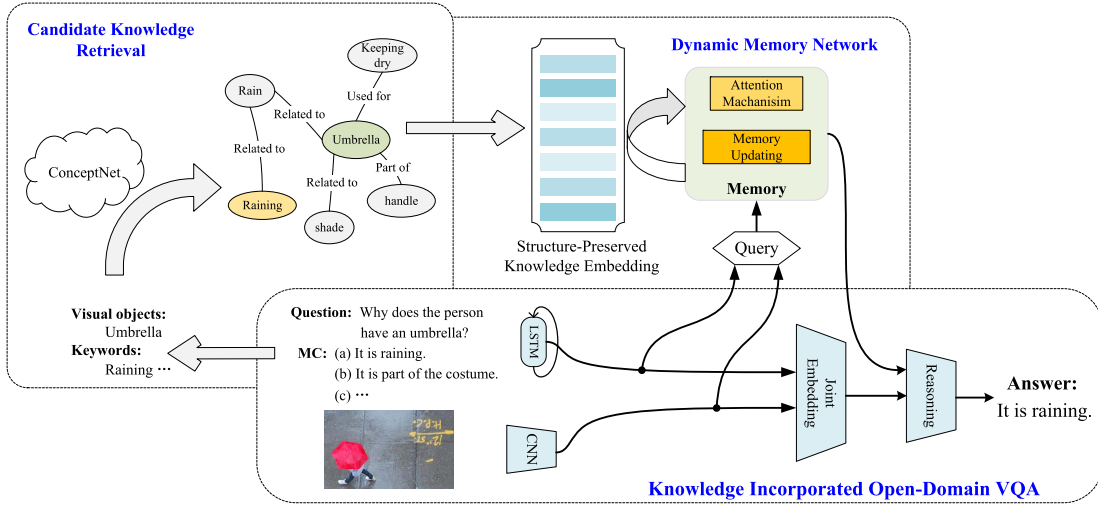


Fig. 21. The architecture of a knowledge-incorporated VQA system, figure from [122].

There is a shiny object that is right of the gray metallic cylinder; does it have the same size as the large rubber sphere?




Fig. 22. An illustration of an end-to-end module network, figure from [124].




Fig. 23. Overview of a graph-based approach for VQA, figure from [129].

massive domain knowledge into a semantic space to answer visual questions. Figure 21 provides a general picture for KDMN framework which consists of three main modules, i.e., retrieval, fusion, inference. In retrieval module, an appropriate number of candidate knowledge triplets are retrieved from the external large-scale KB through analyzing the visual content and textual question. By treating the retrieved knowledge triplets as SVO phrases in fusion module, the authors utilize an LSTM to capture the semantic meanings and embed the knowledge into memory slots, as is shown in the following Eq (19),

$$C_i^{(t)} = \text{LSTM}(\mathbf{L}[w_i^t], C_i^{(t-1)}), \quad t = \{1, 2, 3\}, \quad (18)$$

$$\mathbf{M} = [C_i^{(3)}], \quad (19)$$

where w_i^t is the t -th word of the i -th SVO phrase, \mathbf{L} is the word embedding matrix and C_i is the internal state of

LSTM cell when forwarding the i -th SVO phrase. The memory bank \mathbf{M} is designed to store a large amount of knowledge embedding. With the guidance of visual and textual features, those embedded knowledge triples are then fed into a Dynamic Memory Network [133] to obtain a distilled episodic memory vector in an iterative manner as follows:

$$\mathbf{q} = \text{Query}(\mathbf{f}^{(I)}, \mathbf{f}^{(Q)}, \mathbf{f}^{(A)}), \quad (20)$$

$$\mathbf{c}^{(t)} = \text{Attention}(\mathbf{M}; \mathbf{m}^{(t-1)}, \mathbf{q}), \quad (21)$$

$$\mathbf{m}^{(t)} = \text{Update}(\mathbf{m}^{(t-1)}, \mathbf{c}^{(t)}, \mathbf{q}), \quad (22)$$

where Query creates a context-aware query vector \mathbf{q} , Attention condenses the knowledge into a context vector $\mathbf{c}^{(t)}$ in the t -th iteration, and Update distills information into an episodic memory vector $\mathbf{m}^{(t)}$ iteratively. The final episodic memory vector $\mathbf{m}^{(T)}$ can be jointly utilized with visual features to inference the answer.

Compared with approaches based on simple explicit reasoning, methods incorporating external discrete knowledge not only maintain the superiority of deep models but also acquire the ability to exploit external knowledge for more complex reasoning.

2) *Multi-Modal Video Summarization*: Video summarization is an important and challenging research direction in computer vision (CV). It aims to produce a short video summary which contains a small portion of the video segments, so as to give users a synthetic and useful visual abstract of the video content. A great number of uni-modal approaches have been proposed to solve the problem of video summarization, among which unsupervised methods [134]–[137] normally pick frames or shots from videos with some manually designed visual criteria and supervised methods [138], [139] tend to directly leverage human-edited summary examples to learn video summarization models as well as dig the specific visual patterns for video summaries. Besides the visual features, it has also been observed that videos are often paired with abundant information from other modalities, such as audio signals, text descriptions and so on. All the modality

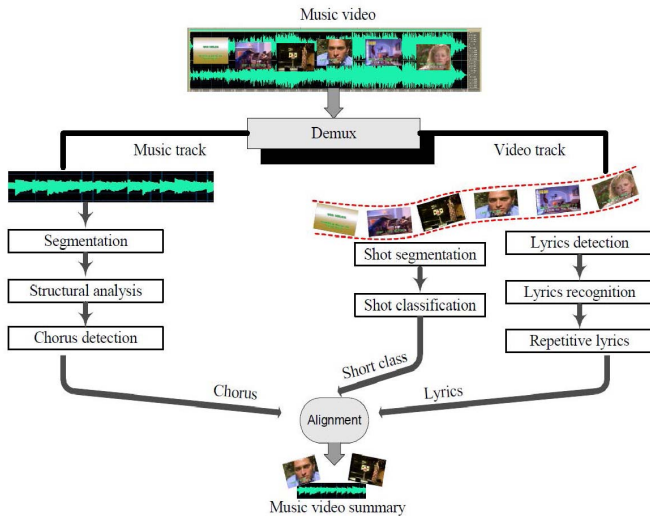


Fig. 24. Workflow of the music video summarization, figure from [140].

information is aligned or complementary with each other, and capable of reflecting video contents in different aspects. Simultaneously considering different modality information of videos can provide video summarization model with a more comprehensive view. Therefore, various multi-modal video summarization methods are proposed based on this idea and we remark that video summarization can also be treated as one application of multi-modal fusion.

a) *Conventional multi-modal video summarization*: Conventional multi-modal video summarization methods mainly focus on summarizing movies or music videos. These methods often detect and synthesize low-level visual/audio/textual cues from video itself to assess the saliency, representativeness or quality of different video parts, and then extract those informative parts to create the final video summary. Xu *et al.* [140] propose a music video summarization method based on audio-visual-text analysis and alignment. As is shown in Figure 24, they first separate the music video into a music track and a video track. For the music track, the chorus is detected based on music structure analysis. For the video track, the (video) shots are segmented and classified into close-up face shots and non-face shots, followed by extraction of the lyrics and detection of the most repeated lyrics from these shots. The music video summary is generated based on the alignment of boundaries of the detected chorus, shot class and the most repeated lyrics from the music video. Pan *et al.* [141] introduce a multi-modal story-oriented video summarization (MMSS) model through encoding both textual and scene information, as well as logos which link shots of a story as a graph. As is shown in Figure 25(a), broadcast news production commonly shows a small icon beside an anchorperson to represent the story. The same icon is usually reused later in the shots about the follow-up development of the story, as an aid for the viewers to link current coverage to past coverage. These icons are called “newslogos”. The property of logos makes them a robust feature for linking separated footages of a story. Based on the above observations, Pan *et al.* [141] build a G_{MMSS} graph as shown in Figure 25(b), which is

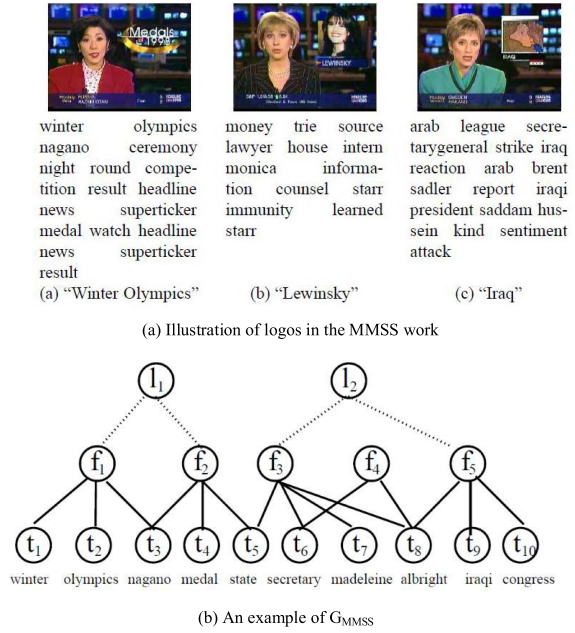


Fig. 25. Illustration of multi-modal story-oriented video summarization (MMSS), figure from [141].

a three-layer graph with three types of nodes and two types of edges. The three types of nodes are logo-node, frame-node and term-node, corresponding to the logos, keyframes (each representing a shot), and terms, respectively. The two types of edges are the term-occurrence edge and the “same-logo” edge. In logo story summarization, frames and terms forming the summary are selected based on their “relevance” to the query object, the logo (node) of the story. The strategy of random walk with restarts (RWR) is used to obtain a story-specific relevance ranking among the terms and shot key frames in the graph G_{MMSS} , then the frames (i.e., nodes) and terms (nodes) with the highest RWR scores will be selected as the story summary. Evangelopoulos *et al.* [142] formulate the detection of perceptually important video events on the basis of saliency models for the audio, visual and textual information conveyed in a video stream. Audio saliency S_a is assessed by cues that quantify multi-frequency waveform modulations. Visual saliency S_v is measured through a spatio-temporal attention model driven by intensity, color and motion. Text saliency S_t is extracted by part-of-speech tagging on the subtitle information from videos. The various modality curves are integrated into a single attention curve by a weighted linear combination of the audio, visual and text saliency,

$$S_{avt} = w_a S_a + w_v S_v + w_t S_t, \quad (23)$$

where the presence of an event may be identified in one or multiple domains. This multi-modal saliency curve is the basis of bottom-up video summarization algorithms which refine results from uni-modal or audiovisual-based skimming.

b) *Multi-modal video summarization for online videos with various side information*: With the massive growth of video websites and social networks, the problem of summarizing online web videos has attracted more and more attentions




Fig. 26. Schematic illustration of the event driven web video summarization approach, figure from [143].

from researchers. Different from traditional offline videos, online web videos are surrounded with various kinds of side information such as tags, titles, descriptions and so on, which carries rich domain knowledge. This domain knowledge often highlights crucial video contents that people focus on and therefore is quite vital for improving the performances of video summarization algorithms. Several multi-modal video summarization methods link web videos with their domain knowledge to analyze video contents and then generate the video summaries.

Wang *et al.* [143] present an approach for event-driven video summarization by tag localization and key-shot mining. As is illustrated in Figure 26, they first localize the tags associated with each video into its shots, where the conditional probability that a shot contains a tag t_k is defined as:

$$v_{ij}^k = P_t(y_{ij}|f_{ij}) = \frac{1}{1 + \exp(-(w_k f_{ij} + b_k))}, \quad (24)$$

where f_{ij} is the feature vector of the j th shot of the i th video. w_k and b_k are the parameters to be learned by the multiple instance learning. After obtaining the relevance scores of the shots with respect to all tags, the relevance score of each shot with respect to an event query can then be estimated. Denote v^k as the relevance score of a shot with respect to the k th tag, then the relevance score of this shot with respect to an event query can be defined as follows:

$$y = \frac{1}{K} \sum_k \text{sim}(q, t_k) v^k, \quad (25)$$

where q is the query and $\text{sim}(q, t_k)$ is the similarity between query q and tag t_k . Finally, a set of key-shots having high relevance scores can be identified by exploring the repeated occurrence characteristics of key sub-events.

Song *et al.* [144] observe that a video title is often carefully chosen to be maximally descriptive of its main topic, and thus images related to the title can serve as a proxy for important visual concepts of the main topic. Therefore, as is depicted in Figure 27, they leverage video titles to retrieve web images through image search engines and develop a co-archetypal analysis technique which learns canonical visual concepts shared between videos and web images. Specifically, suppose $X = [x_1, \dots, x_n] \in R^{d \times n}$ is a matrix of n video frames with each column $x_i \in R^d$ representing a frame with a certain set of image feature descriptors. $Y = [y_1, \dots, y_m] \in R^{d \times m}$ is a matrix of m retrieved images defined in a similar way. The learning of canonical visual concepts

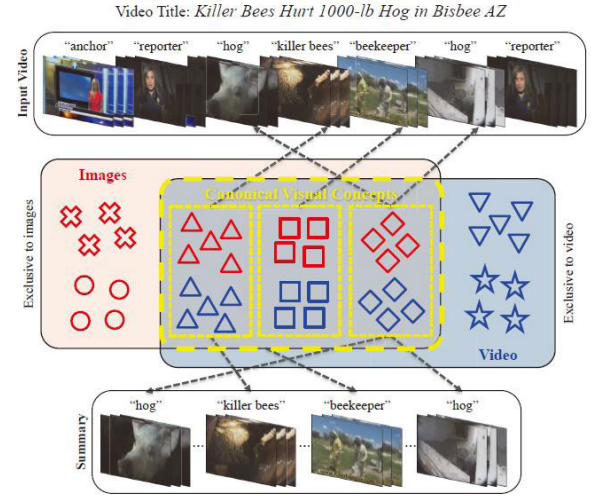


Fig. 27. An illustration of title-based video summarization, figure from [144].

$Z = [z_1, \dots, z_p] \in R^{d \times p}$ between X and Y should satisfy the following two geometrical constraints:

- 1) Each video frame x_i and image y_i should be well approximated by a convex combination of latent variables Z .
- 2) Each latent variable z_j should be well approximated jointly by a convex combination of video frames X and by a convex combination of images Y .

The co-archetypal analysis is thus formulated as an optimization problem that finds a solution set $\Omega = \{A^X, B^X, A^Y, B^Y\}$ by the following objective:

$$\min_{\Omega} \|X - ZA^X\|_F^2 + \|Y - ZA^Y\|_F^2 + \gamma \|XB^X - YB^Y\|_F^2, \quad (26)$$

where $A^X = [\alpha_1^X, \dots, \alpha_n^X] \in R^{p \times n}$, $B^X = [\beta_1^X, \dots, \beta_p^X] \in R^{n \times p}$, and similarly $A^Y \in R^{p \times m}$, $B^Y \in R^{m \times p}$. The first geometrical constraint is reflected by the first two terms in Eq (26), and the second constraint is reflected by the last term, assuming $Z = XB^X = YB^Y$. Upon learning the canonical visual concepts Z as well as the corresponding coefficient matrix A and B , video matrix X can be factorized into XBA , and the importance score of the i th video frame can then be derived as follows:

$$\text{score}(x_i) = \sum_{j=1}^n B_i \alpha_j, \quad (27)$$

which is the total contribution of the corresponding elements of BA in reconstructing the original signal X . With this frame importance measurement, video frames of higher important scores are concatenated in chronological order to form the video summaries.

Sharghi *et al.* [139] propose a query-focused extractive video summarization problem, which aims to generate video summaries based on user provided textual queries. To solve the proposed problem, they develop a probabilistic model, i.e., Sequential and Hierarchical Determinantal Point Process (SH-DPP), where the decision to include one shot in the summary jointly depends on the shot's relevance to the user

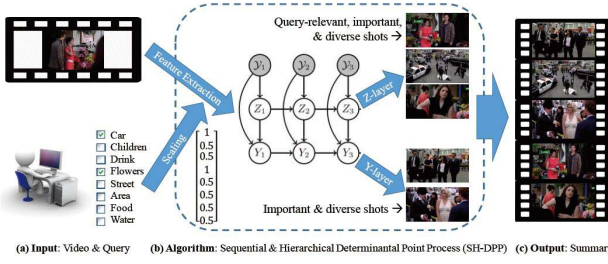


Fig. 28. The workflow of query-focused extractive video summarization, figure from [139].




Fig. 29. The graphical models of SeqDPP [145] (top) and SH-DPP (down), figure from [139].

query and its importance in the context of video. The overall workflow for SH-DPP is shown in Figure 28. Specifically, SH-DPP is established on a Sequential Determinantal Point Process (SeqDPP) method [145], which firstly partitions a video into T consecutive disjoint sets, $\cup_{t=1}^T \mathcal{Y}_t = \mathcal{V}$, where \mathcal{Y}_t represents a set consisting of only a few shots and stands as the ground set of time step t . The SeqDPP model is defined as follows (Figure 29(a) depicts its graphical model),

$$P_{SEQ}(Y|\mathcal{V}) = P(Y_1|\mathcal{Y}_1) \prod_{t=2}^T P(Y_t|Y_{t-1}, \mathcal{Y}_t), \quad \mathcal{V} = \cup_{t=1}^T \mathcal{Y}_t, \quad (28)$$

where $P(Y_t|Y_{t-1}, \mathcal{Y}_t)$ is a conditional DPP to ensure the diversities between items selected at time step t (by Y_t) and those selected in the previous time step (denoted as Y_{t-1}). In order to incorporate user queries into the video summarization procedure, the SH-DPP model (as is shown by the graphical model in Figure 29(b)) leverages the query information to guide the determinantal point process for video shot selection:

$$\begin{aligned} P_{SH}(\{Y_1, Z_1\}, \dots, \{Y_T, Z_T\}|q, \mathcal{V}) \\ = P(Z_1|q, \mathcal{Y}_1)P(Y_1|Z_1, \mathcal{Y}_1) \\ \times \prod_{t=2}^T P(Z_t|q, Z_{t-1}, \mathcal{Y}_t)P(Y_t|Z_t, Y_{t-1}, \mathcal{Y}_t). \end{aligned} \quad (29)$$

The SH-DPP first utilizes the subset selection variables Z_t to select the query-relevant video shots. Depending on the results from Z_t and Y_{t-1} , the variable Y_t in the last layer selects video

shots to further summarize the remaining content in the video segment \mathcal{Y}_t . Since annotating ground-truth for selection variables Z_t needs annotators to determine which query appears in which video shot, query-focused video summarization heavily relies on the user supervision for SH-DPP.

Besides SH-DPP, Sharghi *et al.* [146] further propose a query-focused video summarizer which employs memory network to parameterize the sequential determinantal point process. As is shown in Figure 30, unlike the hierarchical model in [139], the query-focused video summarizer does not require the costly user supervision on “which queried concept appears in which video shot” or any pre-trained concept detectors.

Yuan *et al.* [147] present a Deep Side Semantic Embedding (DSSE) model to generate video summaries by leveraging domain knowledge obtained from side information (e.g, captions, descriptions, queries) of online web videos. The basic idea of DSSE is to construct a latent subspace with the ability of directly comparing domain knowledge and video frames. In this latent subspace, the authors hope that the common information between videos and domain knowledge can be learned more completely and the semantic relevance between them can be effectively measured. As is shown in Figure 31, a latent subspace is constructed by correlating the hidden layers of two uni-modal auto-encoders which embed the video frames and domain knowledge respectively. Meanwhile, there are two components in the objective function of DSSE, i.e, L_{rel} which learns the semantic relevance and L_{rec} which learns the feature reconstruction:

$$L_{rel}(I_f, I_g; \Theta) = \|f(I_f; \Theta_f) - g(I_g; \Theta_g)\|_2^2, \quad (30)$$

$$L_{rec}(I_f, I_g; \Theta) = \|\tilde{I}_f - I_f\|_2^2 + \|\tilde{I}_g - I_g\|_2^2, \quad (31)$$

where I_f represents the visual features of the video frames and I_g represents the textual features of domain knowledge. Accordingly, $f(I_f; \Theta_f)$ is the hidden representation of I_f in the visual auto-encoder and $g(I_g; \Theta_g)$ is the hidden representation of I_g in the textual encoder. \tilde{I}_f and \tilde{I}_g denote the reconstructed features. L_{rel} requires that the matched video frames and domain knowledge be close to each other in the latent subspace and L_{rec} preserves the useful original characteristics from different modalities/media in the common latent space. By jointly minimizing L_{rel} and L_{rec} as follows:

$$\min_{\Theta} \alpha L_{rel}(I_f, I_g; \Theta) + L_{rec}(I_f, I_g; \Theta), \quad (32)$$

the semantic relevance between video frames and domain knowledge can be measured in the hidden layers of the multi-modal auto-encoders and semantically meaningful parts are selected from videos to generate video summaries by minimizing their distances to domain knowledge in the constructed latent subspace. The whole picture of DSSE model is demonstrated in Figure 32.

3) *Multi-Modal Visual Pattern Mining*: Knowledge base is a collection of entities, attributes and the relations between them. knowledge base schema is the structure of knowledge base and used to guide how the knowledge base is built. It is often constructed manually using experts with specific domain knowledge for the field of interest. Many tasks such




Fig. 30. The overview for query-focused video summarizer with memory network, figure from [146].

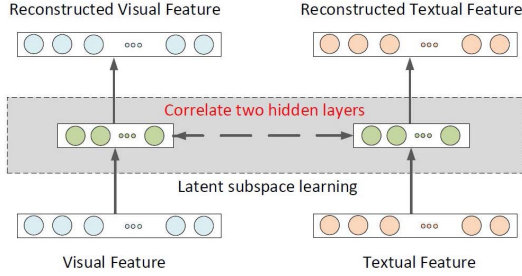


Fig. 31. The architecture of multi-modal auto-encoders, figure from [147].

as automatic content extraction highly depend on knowledge base. However, the current approaches ignore visual information that could be used to build or populate these structured ontologies. Preliminary work on visual knowledge base construction only explores limited basic objects and scene relations. A few novel multi-modal pattern mining approaches are proposed in [148]–[152], towards constructing a high-level “event” schema semi-automatically, which has the capability to extend text-only methods for schema construction. A large unconstrained corpus of weakly-supervised image-caption pairs related to high-level events is utilized to both discover visual aspects of an event, and name these visual components automatically. Li *et al.* [148] leverage the activation signal of the convolution filters to encode the visual content, and utilize the skip-gram language model to encode the textual information. The association rule mining algorithm is introduced to jointly model the visual and textual information from multi-modal data. The encoded visual and textual contents are considered as transactions in association rule mining algorithm. The visual transactions generation pipeline can be found in Figure 33.

To discover the event related multi-modal patterns for knowledge base construction, two criteria, *representative* and *discriminative*, are defined to find the high quality multi-modal visual patterns. *Discriminative* means the patterns discovered from a category should not be found in other categories. *Representative* means the discovered patterns should be commonly available in the category. In association rule mining algorithm, representative property is defined by support rate of a transaction, as is shown in (33), and the discriminative property is defined by confidence rate, as is shown in (34):

$$s(t^*) = \frac{|\{T_a | t^* \subseteq T_a, T_a \in S\}|}{m}, \quad (33)$$

$$c(t^* \rightarrow y) = \frac{s(t^* \cup y)}{s(t^*)}, \quad (34)$$

where T_a is a transaction, t^* is a set of items and y is the target category. The discovered association rules can be converted to multi-modal visual patterns by the algorithm in [148]. Mathematically, the two pattern mining requirements can be defined as:

$$\begin{aligned} c(t^* \rightarrow y) &\geq c_{min}, \\ s(t^*) &\geq s_{min}, \\ t^* \cap \mathbf{I} &\neq \emptyset, \\ t^* \cap \mathbf{C} &\neq \emptyset, \end{aligned} \quad (35)$$

where y is the event category, c_{min} is the threshold of minimum confidence rate, s_{min} is the threshold of minimum support rate, \mathbf{I} is the visual transactions, and \mathbf{C} are the text transactions. Each multi-modal pattern t^* has a set of visual items and a set of text patterns. The end-to-end multi-modal pattern discovery and naming framework can be found in Figure 34.

Multi-modal pattern mining approach can be used as a bridge to fill the gap between text analysis and visual analysis. Zhang *et al.* [150], [152] use the multi-modal visual pattern mining framework proposed in [148], [153] to improve the knowledge and event extraction problem in Natural Language Processing community. Compare to the traditional text only event extraction approach, multi-modal approach introduces the discovered domain knowledge from visual domain and achieve significantly better performance.

4) *Multi-Modal Recommendation*: With the explosive growth of various online social networks and multimedia sites, people are now getting used to engaging on different medias simultaneously to satisfy their diverse information need [154]. It is reported that each user on average has 5.54 social media accounts and is actively using 2.82 social platforms/media.¹ The cross-modal information jointly reflects each individual’s interest and preference. Therefore, organically transferring or associating cross-modal information is of significant importance in serving people intelligently [155].

Existing multi-modal recommendation works can be grouped from two angles, e.g., categorization according to association knowledge and categorization according to the entire model structure.

a) *Grouping by what knowledge to associate*: When we look through existing multi-modal models in terms of the association knowledge, one group of methods follow a user-centric way, which focuses on cross-modal information of overlapped users. A straight forward solution is to treat cross-modal

¹GWII social report: <http://www.globalwebindex.net/blog/internet-users-have-average-of-5-social-media-accounts>

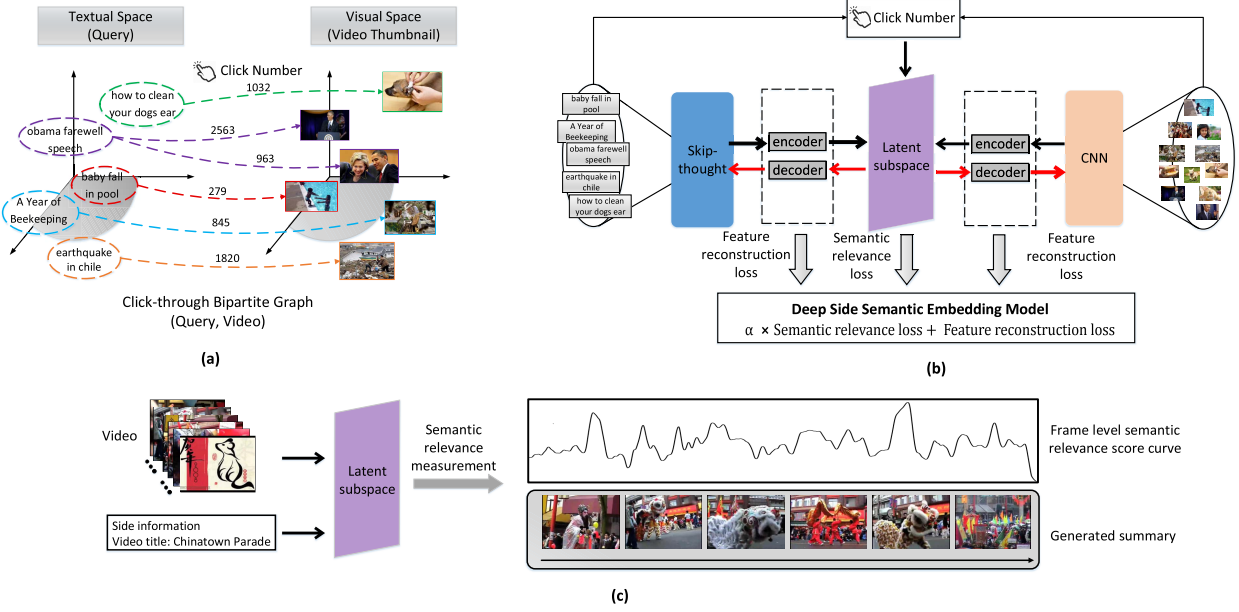


Fig. 32. The overall framework of DSSE model, figure from [147].

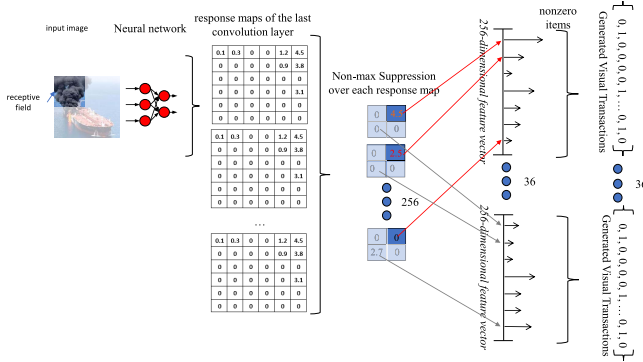


Fig. 33. The visual transaction generation pipeline utilizing the last convolutional layer of a convolutional neural network. This pipeline is used to obtain representations of each image that can localize the presence of a pattern within the image. figure from [148].

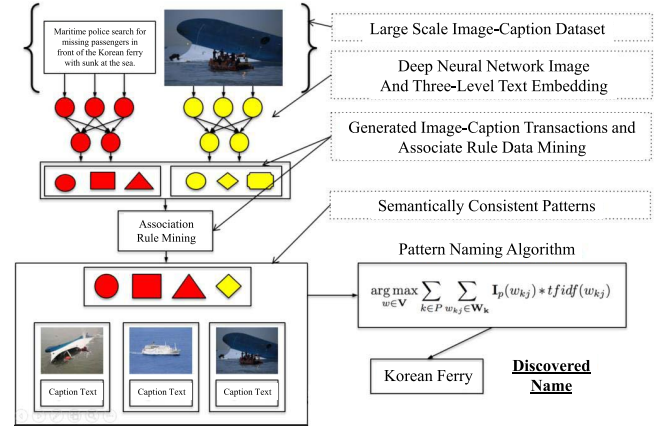


Fig. 34. Multimodal pattern discovery and naming pipeline, figure from [148].

association as a linear transfer problem, and pursue an explicit transfer matrix based on regression [156]–[158]. The objective function for this type of models can be expressed as follows:

$$\min_{\mathbf{W}} \left\| \mathbf{W}\mathbf{U}^1 - \mathbf{U}^2 \right\|_F^2 + \lambda \|\mathbf{W}\|_2, \quad (36)$$

where $\mathbf{U}^i = [\mathbf{u}_1^i, \mathbf{u}_2^i, \dots, \mathbf{u}_U^i]$. The corresponding columns are the same user's representations on two platforms/media. λ is the weighting parameter and the above ridge regression problem has an analytical solution. Instead of pursuing hard transfer, Yan *et al.* [156] propose a topic association framework based on latent attribute sparse coding. They also show that bridging information across different media in common latent space outperforms explicit matrix-oriented transfer. The objective function of the above association framework is

shown in (37):

$$\min_{\mathbf{D}^1, \mathbf{D}^2, \mathbf{S}} \left\| \mathbf{U}^1 - \mathbf{D}^1 \mathbf{S} \right\|_F^2 + \left\| \mathbf{U}^2 - \mathbf{D}^2 \mathbf{S} \right\|_F^2 + \lambda \|\mathbf{S}\|_1$$

$$s.t. \left\| \mathbf{d}_i^Y \right\|_2^2 \leq 1, \left\| \mathbf{d}_j^T \right\|_2^2 \leq 1, \quad \forall i, j, \quad (37)$$

where \mathbf{D}^i includes user factors, and \mathbf{S} includes user attribute representations. The constraint $\|\mathbf{d}\|_2^2 \leq 1$ aims to prevent \mathbf{D} from being arbitrarily large. \mathcal{L}_1 -norm penalty is adopted to encourage a compact and sparse attribute distribution space for users. This problem can be efficiently solved by the sparse coding algorithm proposed in [159] after a few transformations. As is shown in Figure 35, Maaten and Hinton [160] propose an embedding and mapping framework EMCDR in which user representations on different platforms are first

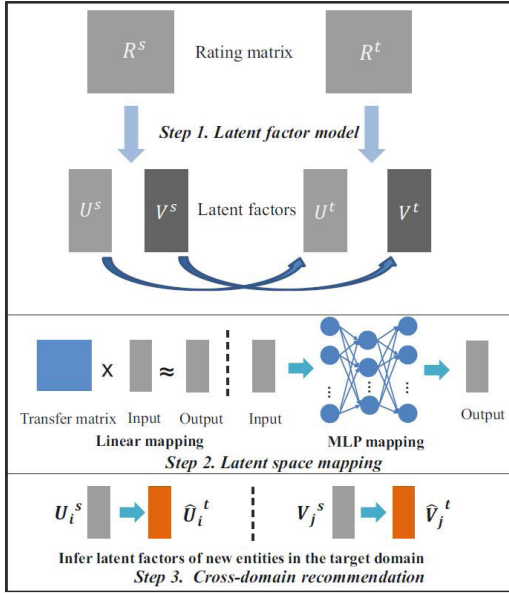


Fig. 35. Illustrative diagram of the EMCDR framework in which linear transfer and MLP are adopted as mapping functions (MLP mapping is proved to perform better according to the experiment results), figure from [160].

obtained through matrix factorization and then mapped via linear transfer or multi-layer perceptron (MLP).

The optimization problem can be formalized as:

$$\min_{\theta} \sum_{\mathbf{u} \in \mathbf{U}} \left\| f_{mlp}(\mathbf{u}^1; \theta) - \mathbf{u}^2 \right\|_2^2, \quad (38)$$

where $f_{mlp}(\cdot; \theta)$ is the MLP mapping function, and θ is its parameter set. Abel *et al.* [161] aggregate user profiles on Flickr, Twitter, Delicious, and propose a solution for the cold-start problem in recommendation. By utilizing the overlapped users and items as bridges across different media, TLRec [162] introduces a smoothness constraint and regularization for latent vectors. Later, Jiang *et al.* introduce an aligned cross-modal user behavior similarity constraint via proposing the XPTrans model [155] which exploits a small number of overlapped crowds to bridge different media optimally.

The objective function of XPTrans model is as follows:

$$\begin{aligned} J = & \left\| \mathbf{W}^1 \odot (\mathbf{R}^1 - \mathbf{U}^1 \mathbf{V}^1) \right\|_F^2 \\ & + \lambda \left\| \mathbf{W}^2 \odot (\mathbf{R}^2 - \mathbf{U}^2 \mathbf{V}^2) \right\|_F^2 \\ & + \mu \left(\left\| \mathbf{W}^{1,2} \mathbf{1}^1 \mathbf{W}^{1,2T} \odot \mathbf{U}^1 \mathbf{U}^{1T} \odot \mathbf{U}^1 \mathbf{U}^{1T} \right\| \right. \\ & + \left\| \mathbf{W}^{1,2T} \mathbf{1}^1 \mathbf{W}^{1,2} \odot \mathbf{U}^2 \mathbf{U}^{2T} \odot \mathbf{U}^2 \mathbf{U}^{2T} \right\| \\ & \left. - 2 \left\| \mathbf{U}^1 \mathbf{U}^{1T} \mathbf{W}^{1,2} \mathbf{U}^2 \mathbf{U}^{2T} \mathbf{W}^{1,2T} \right\| \right) \\ \text{s.t. } & \mathbf{U}^1 > 0, \quad \mathbf{V}^1 > 0, \quad \mathbf{U}^2 > 0, \quad \mathbf{V}^2 > 0, \end{aligned} \quad (39)$$

where the first two lines are traditional loss of matrix factorization on two platforms, and the following three lines are the derived similarity constraint.

The other group of methods are devoted to taking advantage of different media characteristics towards collaborative




Fig. 36. The graphical representation of the cross-domain collaborative learning (CDCL) algorithm. The red circles represent the shared priors to associate with the relevant information and collaboratively learn the shared feature space in different domains, figure from [166].

applications. CODEBOOK [163] investigates behavior prediction across Netflix and MovieLens without considering the overlapped users under the assumption that they share the same user-item rating patterns. Roy *et al.* [164] exploit real-time and socialized characteristics of tweets from Twitter to facilitate video recommendations on YouTube. TPCF [165] integrates three types of data, i.e., aligned users, aligned items and user-item ratings, in transfer learning for collaborative filtering. Qian *et al.* [166] propose a generic cross-domain collaborative learning (CDCL) framework based on nonparametric Bayesian dictionary learning for cross-modal data analysis as is shown in Figure 36. Min *et al.* [167] develop a multi-modal topic model capable of differentiating topics across modalities.

b) *Grouping by entire structure*: When looking through existing literature with respect to the entire structure, one group of methods are designed to build a unified framework [155], [162], [165]–[167] in which the first two works utilize matrix factorization based techniques and the latter three employ probabilistic model based strategies. Another group of works adopt a two-step procedure [156], [158], [168] by first representing users from different media in their own latent spaces and then jointly associating those representations.

The above mentioned methods hold the same core idea that all cross-modal information is consistent and should be aligned. However, a few works [168], [169] discover extra domain knowledge confirming the existence of data inconsistency phenomenon in the procedure of associating representations across different media, and attempt to solve




Fig. 37. The proposed cross-network collaboration solution framework for unified YouTube video recommendation, figure from [168].




Fig. 38. Disparity-preserved Deep Cross-platform Association model. u^T and u^Y are representations of an overlapped user on Twitter and YouTube. In latent representations, h^T and h^Y are media-specific parts preserving disparities, while h^C is the media-shared part associating representations in different media. The estimated representations \hat{u}^T and \hat{u}^Y are derived from both media-shared and media-specific parts, figure from [170].

this problem through data selection. Lu *et al.* [169] find that selecting media-consistent auxiliary data is important for cross-modal collaborative filtering. They propose a novel criterion based on empirical prediction error and variance to assess the consistency, and incorporate the criterion into a boosting framework to selectively transfer knowledge. As is shown in Figure 37, Yan *et al.* [168] divide users into three groups and propose a predefined micro-level user-specific metric to adaptively weight data while integrating heterogeneous data across different media.

In particular, Yu *et al.* [170] analyze the inconsistent behavior patterns of users in Twitter and YouTube by utilizing the domain knowledge of data inconsistency, and discover that the inconsistency is mainly caused by media-specific disparity — each individual's inherent personal preference consists of a media-shared part and a media-specific part due to users' different focuses in different media. To tackle the problem of media-specific disparity and granularity difference, they propose a disparity-preserved deep cross-platform association model whose core idea is shown in Figure 38. Their proposed model contains a partially-connected multi-modal autoencoder which explicitly captures and preserves media-specific disparities in latent representations. They divide the hidden layer into $h = [h^T, h^C, h^Y]$, where h^T , h^Y are Twitter, YouTube media-specific parts respectively, and h^C is the media-shared part. Moreover, they also introduce nonlinear mapping functions to associate cross-modal information, which is advantageous in handling the granularity difference. The detailed structure of multi-modal autoencoder can be written as follows:

$$h = g \left[\left(\sum_i w_1^i x^i \right) + b_1 \right] \\ \hat{x}^i = g \left(w_2^i h + b_2^i \right), \quad (40)$$

where $i \in \{T, Y\}$ denotes Twitter or YouTube, and the weights on the unnecessary links are all set to zero. Weight matrices W and bias units b are denoted by θ as parameters of multi-modal autoencoder. $g(\cdot)$ is the Sigmoid activation function. The total loss consists of reconstruction error, parameter regularizer (regularization penalty) and sparsity constraint, as is shown in (41):

$$L(x^i; \theta) = \sum_i \left\| \hat{x}^i - x^i \right\|_2^2 + \lambda \sum_{W \in \theta} \|W\|_F^2 + \mu \|h\|_1. \quad (41)$$

The whole framework of disparity-aware cross-modal video recommendation is presented in Figure 39.

IV. FUTURE RESEARCH DIRECTIONS

We have presented approaches on multi-modal analysis for multimedia and discussed literature on data-driven cross-modal correlational representation and knowledge-guided multi-modal fusion. With current approaches, the fusion of *continuous data* and *discrete knowledge* has been successfully handled. However, there are still great challenges in obtaining the ability of reasoning for multimedia intelligence. In this section, we share our insights on future directions for multi-modal research.

A. Cross-Modal Reasoning

If we take another look at the above two aspects from the perspective that how close the corresponding approaches/models are to the real intelligence like human beings, the results would probably be “both still have a long long way to go” — the later one may be closer to the real intelligent agent because human can always utilize knowledge from relevant domains to help make decisions. Moreover, if we think deeper about what makes current algorithms take a further step towards human intelligence, the answer will be “reasoning”. The ability of reasoning distinguishes human being from animals. One representative embodiment of reasoning lies in the process of communications among humans — the ability of reasoning meanings of spoken languages during a conversation or major ideas of written articles when reading becomes a vital necessity in understanding each other. This being the case, cross-modal intelligence reasoning over the evolution of knowledge serves as a key solution in bridging the gap between current machine learning algorithms and human intelligence. It will result in a more human-like cognition in cross-modal intelligence. Therefore, being capable of performing human-like reasoning over various kinds of knowledge in cross-modal analysis may be an great opportunity for the next breakthrough in artificial intelligence.

B. Cross-Modal Cognition

Let us consider another question: how do human learn and how can infants learn so well? The ability of cognition through “real” understanding the world could be one main answer to the question. We continuously learn different skills (tasks) since we are babies and obtaining new skills (learning new tasks) seldom deteriorates our possessions of old skills

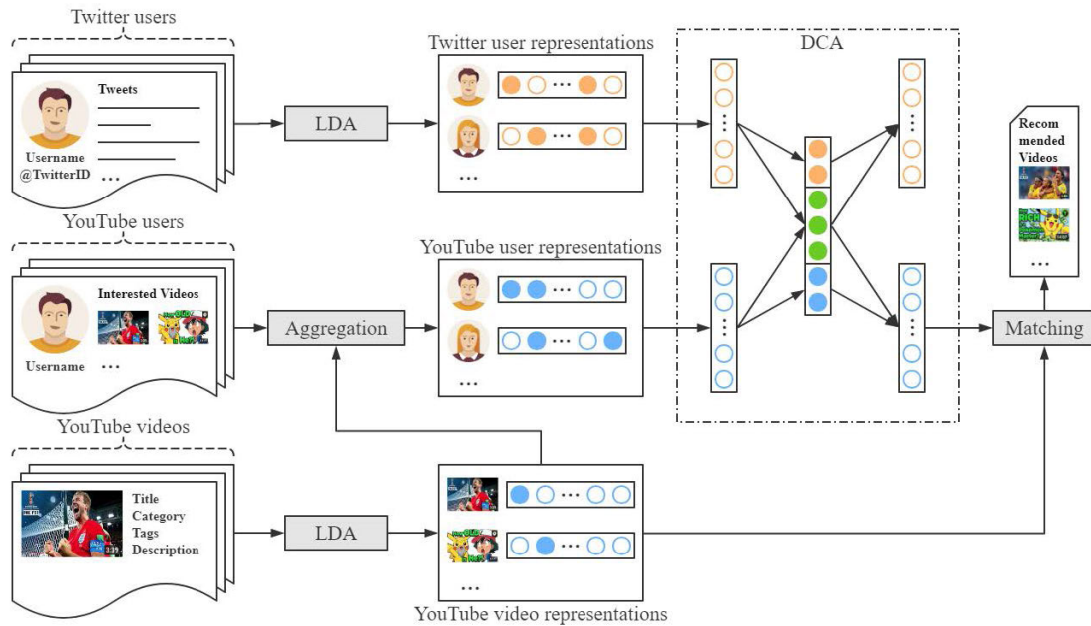


Fig. 39. Framwork for disparity-awared deep cross-modal video recommendation, figure from [170].

(learned old tasks). Most existing machine learning algorithms are capable of tackling only one single type of task. For instance, an image classification algorithm can hardly solve (or achieve a very poor performance on) the trajectory prediction problem, although both image classification and trajectory prediction can be handled by human easily. This indicates that the ability of learning to solve new tasks while maintaining the capacity to tackle previous tasks (a reflection of cognitive process) plays a crucial role in generating human-like algorithms for cross-modal analysis.

We remark that as another reflection of cognition, *commonsense learning* will be an effective path to the goal of touching real human intelligence. Just imagine what kind of scene will appear in your mind when seeing the following sentence “Tom picks up his bag and goes out”: Tom is probably a man who is at work, he stretches out his arm and holds the grip of his bag, stands up and walks to the door, opens it and goes out — Tom does not fly or crawl to the door, nor does he go out by walking through the wall. It is obvious that none of the existing models are able to obtain the above knowledge which is easily understood by human given the quoted sentence as input. We call the process of learning such commonsense knowledge *commonsense learning*, which may lead to another breakthrough in research on cross-modal intelligence.

C. Cross-Modal Collective Intelligence

The concept of collective intelligence (also refers as wisdom of crowd) was originally derived from the observations of entomologist William Morton Wheeler. On the surface, independent individuals can work very closely so that they look like a single organism. In 1911, Wheeler observed such a collaborative process indeed works on ants. An ant behaves like an animal’s cell and processed the ability of collective thinking. He called these collective ants a larger creature, namely the cluster of ant colonies seems to form a

“superorganism.” In human society, given that decisions made by a single individual tend to be inaccurate compared to decisions made by the majority, collective intelligence becomes a shared intelligence as well as the process of assembling opinions and turning them into decision-making procedure. Wikipedia, as a type of media that fully demonstrates collective intelligence, serves as an encyclopedia that can be changed by anyone at almost any time, which connects people on the web to create a huge intelligent brain. All these phenomena or instances confirm one thing, i.e., collective intelligence can produce a more power “superorganism” or brain that possesses more intelligence. With abundant cross-modal information, we believe that collective intelligence can be employed for human planning which is another unique and complex characteristic shared by human being.

In addition, it is also desirable that the advances in cross-media intelligence can indeed make some contributions to human society. Current approaches have done a good job on modality adaption, but they seldom can achieve good performances on *cross-modal generation*. Let’s take visual impaired people as an example, people with visual handicap usually wear a special-tailored helmet with a distance sensor on it. This helmet will produce some noises when there exist some obstacles within a certain distance from the people wearing it. We believe it will be a significant help towards visual impaired people if the helmet can act as an “artificial eye” by describing how far and what obstacle(s) are in which direction from him. This could be accomplished by generating logical verbal languages from understanding sensory data. In general, there is still a large potential space of improvement for cross-media intelligence in both methodologies and applications.

V. CONCLUSION

In this article, we give a comprehensive and deep investigation on multi-modal analysis. We present two

scientific problems on multi-modal analysis for multimedia. In order to address these two scientific problems, we discuss multi-modal fusion methods in two aspects: 1) data-driven multi-modal correlational representation and 2) knowledge-guided multi-modal fusion. We first give a brief summary on multi-task and multi-view learning, and target works on deep representation, transfer learning as well as hashing for data-driven correlational representation. We then present our ideas on potential methods suitable for handling the fusion of multi-modal data and domain knowledge, and discuss approaches for four promising applications, i.e., visual question answering, video summarization, visual pattern mining and recommendation, which need diverse domain knowledge for multi-modal fusion of data with knowledge. Last but not least, we highlight some insights on future research directions in the new era of artificial intelligence, and point out a few promising future directions, including: cross-modal reasoning, cross-modal cognition and cross-modal collective intelligence, for further investigation. We believe these directions have a great potential to lead the next breakthrough in cross-media intelligence.

ACKNOWLEDGMENT

The authors also thank Guohao Li, Shengze Yu, and Yitian Yuan for providing relevant materials and valuable opinions. This work will never be accomplished without their useful suggestions.

REFERENCES

- [1] H. McGurk and J. MacDonald, "Hearing lips and seeing voices," *Nature*, vol. 264, no. 5588, p. 746, 1976.
- [2] T. Evgeniou and M. Pontil, "Regularized multi-task learning," in *Proc. 10th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2004, pp. 109–117.
- [3] A. Argyriou, T. Evgeniou, and M. Pontil, "Convex multi-task feature learning," *Mach. Learn.*, vol. 73, no. 3, pp. 243–272, 2008.
- [4] T. Jebara, "Multitask sparsity via maximum entropy discrimination," *J. Mach. Learn. Res.*, vol. 12, no. 1, pp. 75–110, 2011.
- [5] Y. Zhang and Q. Yang, "A survey on multi-task learning," 2017, *arXiv:1707.08114*. [Online]. Available: <https://arxiv.org/abs/1707.08114>
- [6] S. Ruder, "An overview of multi-task learning in deep neural networks," 2017, *arXiv:1706.05098*. [Online]. Available: <https://arxiv.org/abs/1706.05098>
- [7] C. Xu, D. Tao, and C. Xu, "A survey on multi-view learning," 2013, *arXiv:1304.5634*. [Online]. Available: <https://arxiv.org/abs/1304.5634>
- [8] S. Sun, "A survey of multi-view machine learning," *Neural Comput. Appl.*, vol. 23, nos. 7–8, pp. 2031–2038, 2013.
- [9] Y. Wang, W. Zhang, L. Wu, X. Lin, and X. Zhao, "Unsupervised metric fusion over multiview data by graph random walk-based cross-view diffusion," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 28, no. 1, pp. 57–70, Jan. 2017.
- [10] X. Gao, T. Mu, J. Y. Goulermas, and M. Wang, "Topic driven multimodal similarity learning with multi-view voted convolutional features," *Pattern Recognit.*, vol. 75, pp. 223–234, Mar. 2018.
- [11] J. He and R. Lawrence, "A graphbased framework for multi-task multi-view learning," in *Proc. ICML*, 2011, pp. 25–32.
- [12] Y. Yan, E. Ricci, R. Subramanian, O. Lanz, and N. Sebe, "No matter where you are: Flexible graph-guided multi-task learning for multi-view head pose classification under target motion," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 1177–1184.
- [13] Z. Hong, X. Mei, D. Prokhorov, and D. Tao, "Tracking via robust multi-task multi-view joint sparse representation," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 649–656.
- [14] Y. Liu, Y. Zheng, Y. Liang, S. Liu, and D. S. Rosenblum, "Urban water quality prediction based on multi-task multi-view learning," in *Proc. 25th Int. Joint Conf. Artif. Intell.*, 2016, pp. 1–7.
- [15] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [16] H. P. Martínez and G. N. Yannakakis, "Deep multimodal fusion," in *Proc. Int. Conf. Multimodal Interact.*, 2014, pp. 34–41.
- [17] S. E. Kahou *et al.*, "Combining modality specific deep neural networks for emotion recognition in video," in *Proc. ACM Int. Conf. Multimodal Interact.*, 2013, pp. 543–550.
- [18] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 568–576.
- [19] D. Wu *et al.*, "Deep dynamic neural networks for multimodal gesture segmentation and recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 8, pp. 1583–1597, Aug. 2016.
- [20] D. Yi, Z. Lei, and S. Z. Li, "Shared representation learning for heterogeneous face recognition," in *Proc. 11th IEEE Int. Conf. Workshops Autom. Face Gesture Recognit. (FG)*, vol. 1, May 2015, pp. 1–7.
- [21] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and F. F. Li, "Large-scale video classification with convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 1725–1732.
- [22] C. Ding and D. Tao, "Robust face recognition via multimodal deep face representation," *IEEE Trans. Multimedia*, vol. 17, no. 11, pp. 2049–2058, Nov. 2015.
- [23] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko, "Simultaneous deep transfer across domains and tasks," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2015, pp. 4068–4076.
- [24] N. Neverova, C. Wolf, G. Taylor, and F. Nebout, "ModDrop: Adaptive multi-modal gesture recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 8, pp. 1692–1706, Aug. 2016.
- [25] A. Karpathy, A. Joulin, and F. F. Li, "Deep fragment embeddings for bidirectional image sentence mapping," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 3, 2014, pp. 1889–1897.
- [26] J. Donahue *et al.*, "Long-term recurrent convolutional networks for visual recognition and description," in *Proc. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 2625–2634.
- [27] M. Ren, R. Kiros, and R. S. Zemel, "Exploring models and data for image question answering," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 2953–2961.
- [28] F. J. Ordóñez and D. Roggen, "Deep convolutional and LSTM recurrent neural networks for multimodal wearable activity recognition," *Sensors*, vol. 16, no. 1, p. 115, 2016.
- [29] J.-H. Kim *et al.*, "Multimodal residual learning for visual QA," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 361–369.
- [30] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, "Multimodal deep learning," in *Proc. 28th Int. Conf. Mach. Learn. (ICML)*, 2011, pp. 689–696.
- [31] N. Srivastava and R. Salakhutdinov, "Learning representations for multimodal data with deep belief nets," in *Proc. Int. Conf. Mach. Learn. Workshop*, vol. 79, 2012, pp. 1–8.
- [32] N. Srivastava and R. R. Salakhutdinov, "Multimodal learning with deep boltzmann machines," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 2222–2230.
- [33] C. Yu *et al.*, "Medical image retrieval: A multimodal approach," *Cancer Inform.*, vol. 13, no. Suppl 3, pp. 125–136, 2014.
- [34] I. J. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [35] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. Int. Conf. Learn. Represent.*, 2014, pp. 1–14.
- [36] Y. Huang, W. Wang, and L. Wang, "Unconstrained multimodal multi-label learning," *IEEE Trans. Multimedia*, vol. 17, no. 11, pp. 1923–1935, Nov. 2015.
- [37] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee, "Generative adversarial text to image synthesis," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1060–1069.
- [38] M. Suzuki, K. Nakayama, and Y. Matsuo, "Joint multimodal learning with deep generative models," 2016, *arXiv:1611.01891*. [Online]. Available: <https://arxiv.org/abs/1611.01891>
- [39] G. Pandey and A. Dukkipati, "Variational methods for conditional multimodal deep learning," in *Proc. Int. Joint Conf. Neural Netw.*, 2017, pp. 308–315.
- [40] M. R. Amer, B. Siddiquie, S. Khan, and A. Divakaran, "Multimodal fusion using dynamic hybrid models," in *Proc. Appl. Comput. Vis.*, Mar. 2014, pp. 556–563.
- [41] Y. Liu, X. Feng, and Z. Zhou, "Multimodal video classification with stacked contractive autoencoders," *Signal Process.*, vol. 120, pp. 761–766, Mar. 2016.

- [42] M. R. Amer, T. Shields, B. Siddiquie, A. Tamrakar, A. Divakaran, and S. Chai, "Deep multimodal fusion: A hybrid approach," *Int. J. Comput. Vis.*, vol. 126, pp. 440–456, Apr. 2018.
- [43] D. Ramachandram and G. W. Taylor, "Deep multimodal learning: A survey on recent advances and trends," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 96–108, Nov. 2017.
- [44] M. Long, Y. Cao, J. Wang, and M. I. Jordan, "Learning transferable features with deep adaptation networks," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 97–105.
- [45] M. Long, H. Zhu, J. Wang, and M. I. Jordan, "Deep transfer learning with joint adaptation networks," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 2208–2217.
- [46] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2014, pp. 3320–3328.
- [47] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, "Deep domain confusion: Maximizing for domain invariance," 2014, *arXiv:1412.3474*. [Online]. Available: <https://arxiv.org/abs/1412.3474>
- [48] D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu, "Equivalence of distance-based and RKHS-based statistics in hypothesis testing," *Ann. Statist.*, vol. 41, no. 5, pp. 2263–2291, 2013.
- [49] X. Liu, M. Wang, Z.-J. Zha, and R. Hong, "Cross-modality feature learning via convolutional autoencoder," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 15, no. 1s, p. 7, 2019.
- [50] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
- [51] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [52] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "CNN features off-the-shelf: An astounding baseline for recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2014, pp. 806–813.
- [53] X. Glorot, A. Bordes, and Y. Bengio, "Domain adaptation for large-scale sentiment classification: A deep learning approach," in *Proc. Int. Conf. Int. Conf. Mach. Learn.*, 2011, pp. 513–520.
- [54] M. Chen, Z. Xu, K. Weinberger, and S. Fei, "Marginalized denoising autoencoders for domain adaptation," 2012, *arXiv:1206.4683*. [Online]. Available: <https://arxiv.org/abs/1206.4683>
- [55] F. Zhuang, X. Cheng, P. Luo, S. J. Pan, and Q. He, "Supervised representation learning: Transfer learning with deep autoencoders," in *Proc. Int. Conf. Artif. Intell.*, 2015, pp. 4119–4125.
- [56] H. Daumé, III, "Frustratingly easy domain adaptation," in *Proc. ACL*, 2016, pp. 1–9.
- [57] B. Sun, J. Feng, and K. Saenko, "Return of frustratingly easy domain adaptation," in *Proc. 13th AAAI Conf. Artif. Intell.*, 2016, pp. 2058–2065.
- [58] K. Bousmalis, G. Trigeorgis, N. Silberman, D. Krishnan, and D. Erhan, "Domain separation networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 343–351.
- [59] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 1180–1189.
- [60] Y. Ganin *et al.*, "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2030–2096, May 2015.
- [61] S. Moon, S. Kim, and H. Wang, "Multimodal transfer deep learning with applications in audio-visual recognition," 2016, *arXiv:1412.3121*. [Online]. Available: <https://arxiv.org/abs/1412.3121>
- [62] C. Navarretta, "Transfer learning in multimodal corpora," in *Proc. IEEE Int. Conf. Cogn. Infocommun.*, Dec. 2014, pp. 195–200.
- [63] B. Cheng, M. Liu, H.-I. Suk, D. Shen, and D. Zhang, "Multimodal manifold-regularized transfer learning for MCI conversion prediction," *Brain Imag. Behav.*, vol. 9, no. 4, pp. 913–926, 2015.
- [64] B. Cheng, M. Liu, and D. Zhang, "Multimodal multi-label transfer learning for early diagnosis of Alzheimer's disease," in *Proc. 6th Int. Workshop Mach. Learn. Med. Imag.*, vol. 9352. New York, NY, USA: Springer-Verlag, 2015.
- [65] D. Wang, P. Cui, and W. Zhu, "Deep asymmetric transfer network for unbalanced domain adaptation," in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018, pp. 444–450.
- [66] J. Song, Y. Yang, Z. Huang, H. T. Shen, and R. Hong, "Multiple feature hashing for real-time large scale near-duplicate video retrieval," in *Proc. Int. Conf. Multimedia*, Scottsdale, Az, USA, Nov./Dec. 2011, pp. 423–432.
- [67] J. Song, Y. Yang, Z. Huang, H. T. Shen, and J. Luo, "Effective multiple feature hashing for large-scale near-duplicate video retrieval," *IEEE Trans. Multimedia*, vol. 15, no. 8, pp. 1997–2008, Dec. 2013.
- [68] D. Zhang, F. Wang, and L. Si, "Composite hashing with multiple information sources," in *Proc. 34th Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, 2011, pp. 225–234.
- [69] M. M. Bronstein, A. M. Bronstein, F. Michel, and N. Paragios, "Data fusion through cross-modality metric learning using similarity-sensitive hashing," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2010, pp. 3594–3601.
- [70] S. Kumar and R. Udupa, "Learning hash functions for cross-view similarity search," in *Proc. Int. Joint Conf. Artif. Intell. (IJCAI)*, vol. 22, no. 1, 2011, p. 1360.
- [71] S. Kim, Y. Kang, and S. Choi, "Sequential spectral learning to hash with multiple representations," in *Proc. Eur. Conf. Comput. Vis.* Berlin, Germany: Springer, 2012, pp. 538–551.
- [72] Y. Zhen and D.-Y. Yeung, "A probabilistic model for multimodal hash function learning," in *Proc. 18th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2012, pp. 940–948.
- [73] Y. Zhen and D.-Y. Yeung, "Co-regularized hashing for multimodal data," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1376–1384.
- [74] X. Zhu, Z. Huang, H. T. Shen, and X. Zhao, "Linear cross-modal hashing for efficient multimedia search," in *Proc. 21st ACM Int. Conf. Multimedia*, 2013, pp. 143–152.
- [75] J. Song, Y. Yang, Y. Yang, Z. Huang, and H. T. Shen, "Inter-media hashing for large-scale retrieval from heterogeneous data sources," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2013, pp. 785–796.
- [76] M. Ou, P. Cui, F. Wang, J. Wang, W. Zhu, and S. Yang, "Comparing apples to oranges: A scalable solution with heterogeneous hashing," in *Proc. 19th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2013, pp. 230–238.
- [77] G. Ding, Y. Guo, and J. Zhou, "Collective matrix factorization hashing for multimodal data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 2075–2082.
- [78] J. Zhou, G. Ding, and Y. Guo, "Latent semantic sparse hashing for cross-modal similarity search," in *Proc. 37th Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, 2014, pp. 415–424.
- [79] J. Masci, M. M. Bronstein, A. M. Bronstein, and J. Schmidhuber, "Multimodal similarity-preserving hashing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 4, pp. 824–830, Apr. 2014.
- [80] D. Zhang and W.-J. Li, "Large-scale supervised multimodal hashing with semantic correlation maximization," in *Proc. AAAI*, 2014, vol. 1, no. 2, p. 7.
- [81] F. Wu, Z. Yu, Y. Yang, S. Tang, Y. Zhang, and Y. Zhuang, "Sparse multi-modal hashing," *IEEE Trans. Multimedia*, vol. 16, no. 2, pp. 427–439, Feb. 2014.
- [82] Z. Yu, F. Wu, Y. Yang, Q. Tian, J. Luo, and Y. Zhuang, "Discriminative coupled dictionary hashing for fast cross-media retrieval," in *Proc. 37th Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, 2014, pp. 395–404.
- [83] Y. Hu, Z. Jin, H. Ren, D. Cai, and X. He, "Iterative multi-view hashing for cross media indexing," in *Proc. 22nd ACM Int. Conf. Multimedia*, 2014, pp. 527–536.
- [84] Z. Lin, G. Ding, M. Hu, and J. Wang, "Semantics-preserving hashing for cross-view retrieval," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 3864–3872.
- [85] B. Wu, Q. Yang, W.-S. Zheng, Y. Wang, and J. Wang, "Quantized correlation hashing for fast cross-modal search," in *Proc. IJCAI*, 2015, pp. 3946–3952.
- [86] S. Moran and V. Lavrenko, "Regularised cross-modal hashing," in *Proc. 38th Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, 2015, pp. 907–910.
- [87] D. Wang, P. Cui, M. Ou, and W. Zhu, "Learning compact hash codes for multimodal representations using orthogonal deep structure," *IEEE Trans. Multimedia*, vol. 17, no. 9, pp. 1404–1416, Sep. 2015.
- [88] Y. Cao, M. Long, J. Wang, Q. Yang, and P. S. Yu, "Deep visual-semantic hashing for cross-modal retrieval," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 1445–1454.
- [89] Y. Cao, M. Long, J. Wang, and H. Zhu, "Correlation autoencoder hashing for supervised cross-modal search," in *Proc. ACM Int. Conf. Multimedia Retr.*, 2016, pp. 197–204.
- [90] Q.-Y. Jiang and W.-J. Li, "Deep cross-modal hashing," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2017, pp. 3232–3240.
- [91] Y. Weiss, A. Torralba, and R. Fergus, "Spectral hashing," in *Proc. Adv. Neural Inf. Process. Syst.*, 2009, pp. 1753–1760.
- [92] V. E. Liong, J. Lu, Y. P. Tan, and J. Zhou, "Cross-modal deep variational hashing," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 4097–4105.

- [93] X. Li, D. Hu, and F. Nie, "Deep binary reconstruction for cross-modal hashing," *IEEE Trans. Multimedia*, vol. 21, no. 4, pp. 973–985, Aug. 2018.
- [94] C. Li, C. Deng, N. Li, W. Liu, X. Gao, and D. Tao, "Self-supervised adversarial hashing networks for cross-modal retrieval," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4242–4251.
- [95] J. Zhang, Y. Peng, and M. Yuan, "SCH-GAN: Semi-supervised cross-modal hashing by generative adversarial network," *IEEE transactions on Cybern.*, 2018.
- [96] L. Wu, Y. Wang, and L. Shao, "Cycle-consistent deep generative hashing for cross-modal retrieval," *IEEE Trans. Image Process.*, vol. 28, no. 4, pp. 1602–1612, Apr. 2019.
- [97] J. M. Bernardo and A. F. Smith, *Bayesian Theory*, vol. 405. Hoboken, NJ, USA: Wiley, 2009.
- [98] A. P. Dempster, "A generalization of Bayesian inference," *J. Roy. Stat. Soc., B (Methodol.)*, vol. 30, no. 2, pp. 205–232, 1968.
- [99] F. V. Jensen, *An Introduction to Bayesian Networks*, vol. 210. London, U.K.: UCL Press, 1996.
- [100] G. E. Box and G. C. Tiao, *Bayesian inference in Statistical Analysis*, vol. 40. Hoboken, NJ, USA: Wiley, 2011.
- [101] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, *arXiv:1503.02531*. [Online]. Available: <https://arxiv.org/abs/1503.02531>
- [102] Z. Luo, J.-T. Hsieh, L. Jiang, J. C. Niebles, and L. Fei-Fei, "Graph distillation for action detection with privileged modalities," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 166–183.
- [103] S. Gupta, J. Hoffman, and J. Malik, "Cross modal distillation for supervision transfer," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 2827–2836.
- [104] C. Zhang and Y. Peng, "Better and faster: Knowledge transfer from multiple self-supervised learning tasks via graph distillation for video classification," 2018, *arXiv:1804.10069*. [Online]. Available: <https://arxiv.org/abs/1804.10069>
- [105] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, vol. 135. Cambridge, MA, USA: MIT Press, 1998.
- [106] L. P. Kaelbling, M. L. Littman, and A. W. Moore, "Reinforcement learning: A survey," *J. Artif. Intell. Res.*, vol. 4, no. 1, pp. 237–285, Jan. 1996.
- [107] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [108] D. Teney, Q. Wu, and A. van den Hengel, "Visual question answering: A tutorial," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 63–75, Nov. 2017.
- [109] K. Xu *et al.*, "Show, attend and tell: Neural image caption generation with visual attention," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2048–2057.
- [110] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola, "Stacked attention networks for image question answering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 21–29.
- [111] J. Lu, J. Yang, D. Batra, and D. Parikh, "Hierarchical question-image co-attention for visual question answering," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 289–297.
- [112] D. Yu, J. Fu, T. Mei, and Y. Rui, "Multi-level attention networks for visual question answering," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, Jun. 2017, pp. 4709–4717.
- [113] P. Anderson *et al.*, "Bottom-up and top-down attention for image captioning and visual question answering," in *Proc. CVPR*, Jul. 2018, vol. 3, no. 5, pp. 1–6.
- [114] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2016, pp. 770–778.
- [115] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [116] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach, "Multimodal compact bilinear pooling for visual question answering and visual grounding," 2016, *arXiv:1606.01847*. [Online]. Available: <https://arxiv.org/abs/1606.01847>
- [117] J.-H. Kim, K.-W. On, W. Lim, J. Kim, J.-W. Ha, and B.-T. Zhang, "Hadamard product for low-rank bilinear pooling," 2016, *arXiv:1610.04325*. [Online]. Available: <https://arxiv.org/abs/1610.04325>
- [118] H. Ben-Younes, R. Cadene, M. Cord, and N. Thome, "Mutan: Multi-modal tucker fusion for visual question answering," in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 3, Oct. 2017, pp. 2612–2620.
- [119] Z. Yu, J. Yu, J. Fan, and D. Tao, "Multi-modal factorized bilinear pooling with co-attention learning for visual question answering," in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 3, Oct. 2017, pp. 1821–1830.
- [120] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh, "Making the V in VQA matter: Elevating the role of image understanding in visual question answering," in *Proc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 6904–6913.
- [121] A. Agrawal, D. Batra, D. Parikh, and A. Kembhavi, "Don't just assume; look and answer: Overcoming priors for visual question answering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4971–4980.
- [122] G. Li, H. Su, and W. Zhu, "Incorporating external knowledge to answer open-domain visual questions with dynamic memory networks," 2017, *arXiv:1712.00733*. [Online]. Available: <https://arxiv.org/abs/1712.00733>
- [123] J. Andreas, M. Rohrbach, T. Darrell, and D. Klein, "Neural module networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 39–48.
- [124] R. Hu, J. Andreas, M. Rohrbach, T. Darrell, and K. Saenko, "Learning to reason: End-to-end module networks for visual question answering," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 804–813.
- [125] J. Johnson, B. Hariharan, L. van der Maaten, J. Hoffman, L. Fei-Fei, C. L. Zitnick, and R. Girshick, "Inferring and executing programs for visual reasoning," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2989–2998.
- [126] D. Mascharka, P. Tran, R. Soklaski, and A. Majumdar, "Transparency by design: Closing the gap between performance and interpretability in visual reasoning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2018, pp. 4942–4950.
- [127] Q. Cao, X. Liang, B. Li, G. Li, and L. Lin, "Visual question reasoning on general dependency tree," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7249–7257.
- [128] D. Teney, L. Liu, and A. van der Hengel, "Graph-structured representations for visual question answering," 2016, *arXiv:1609.05600*. [Online]. Available: <https://arxiv.org/abs/1609.05600>
- [129] W. Norcliffe-Brown, S. Vafeias, and S. Parisot, "Learning conditioned graph structures for interpretable visual question answering," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 8334–8343.
- [130] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. M. Bronstein, "Geometric deep learning on graphs and manifolds using mixture model CNNs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 5115–5124.
- [131] P. Wang, Q. Wu, C. Shen, A. van der Hengel, and A. Dick, "Explicit knowledge-based reasoning for visual question answering," 2015, *arXiv:1511.02570*. [Online]. Available: <https://arxiv.org/abs/1511.02570>
- [132] P. Wang, Q. Wu, C. Shen, A. Dick, and A. van den Hengel, "FVQA: Fact-based visual question answering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 10, pp. 2413–2427, Oct. 2017.
- [133] C. Xiong, S. Merity, and R. Socher, "Dynamic memory networks for visual and textual question answering," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 2397–2406.
- [134] Y. Cong, J. Yuan, and J. Luo, "Towards scalable summarization of consumer videos via sparse dictionary selection," *IEEE Trans. Multimedia*, vol. 14, no. 1, pp. 66–75, Feb. 2012.
- [135] S. Lu, Z. Wang, T. Mei, G. Guan, and D. D. Feng, "A bag-of-importance model with locality-constrained coding based feature learning for video summarization," *IEEE Trans. Multimedia*, vol. 16, no. 6, pp. 1497–1509, Oct. 2014.
- [136] Y. J. Lee, J. Ghosh, and K. Grauman, "Discovering important people and objects for egocentric video summarization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 1346–1353.
- [137] T. Yao, T. Mei, and Y. Rui, "Highlight detection with pairwise deep ranking for first-person video summarization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 982–990.
- [138] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman, "Video summarization with long short-term memory," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 766–782.
- [139] A. Sharghi, B. Gong, and M. Shah, "Query-focused extractive video summarization," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 3–19.
- [140] C. Xu, X. Shao, N. C. Maddage, and M. S. Kankanhalli, "Automatic music video summarization based on audio-visual-text analysis and alignment," in *Proc. 28th Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, 2005, pp. 361–368.


- [141] J. Y. Pan, H. Yang, and C. Faloutsos, "MMSS: Multi-modal story-oriented video summarization," in *Proc. IEEE Int. Conf. Data Mining*, Nov. 2004, pp. 491–494.
- [142] G. Evangelopoulos *et al.*, "Video event detection and summarization using audio, visual and text saliency," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Apr. 2009, pp. 3553–3556.
- [143] M. Wang, R. Hong, G. Li, Z.-J. Zha, S. Yan, and T.-S. Chua, "Event driven Web video summarization by tag localization and key-shot identification," *IEEE Trans. Multimedia*, vol. 14, no. 4, pp. 975–985, Aug. 2012.
- [144] Y. Song, J. Vallmitjana, A. Stent, and A. Jaimes, "TVSum: Summarizing Web videos using titles," in *Proc. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 5179–5187.
- [145] B. Gong, W.-L. Chao, K. L. Grauman, and F. Sha, "Diverse sequential subset selection for supervised video summarization," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, no. 3, 2014, pp. 2069–2077.
- [146] A. Sharghi, J. S. Laurel, and B. Gong, "Query-focused video summarization: Dataset, evaluation, and a memory network based approach," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 4788–4797.
- [147] Y. Yuan, T. Mei, P. Cui, and W. Zhu, "Video summarization by learning deep side semantic embedding," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 29, no. 1, pp. 226–237, Nov. 2017.
- [148] H. Li, J. G. Ellis, H. Ji, and S.-F. Chang, "Event specific multimodal pattern mining for knowledge base construction," in *Proc. ACM Multimedia Conf.*, 2016, pp. 821–830.
- [149] H. Li, J. G. Ellis, L. Zhang, and S.-F. Chang, "PatternNet: Visual pattern mining with deep neural network," in *Proc. ACM Int. Conf. Multimedia Retr.*, 2018, pp. 291–299.
- [150] T. Zhang, H. Li, H. Ji, and S.-F. Chang, "Cross-document event coreference resolution based on cross-media features," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2015, pp. 201–206.
- [151] D. Lu *et al.*, "Cross-media event extraction and recommendation," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Demonstrations*, 2016, pp. 72–76.
- [152] T. Zhang *et al.*, "Improving event extraction via multimodal integration," in *Proc. ACM Multimedia Conf.*, 2017, pp. 270–278.
- [153] W. Zhang, H. Li, C.-W. Ngo, and S.-F. Chang, "Scalable visual instance mining with threads of features," in *Proc. 22nd ACM Int. Conf. Multimedia*, 2014, pp. 297–306.
- [154] T. Chen, M. A. Kaafar, A. Friedman, and R. Boreli, "Is more always merrier?: A deep dive into online social footprints," in *Proc. ACM Workshop Online Social Netw.*, 2012, pp. 67–72.
- [155] M. Jiang, P. Cui, N. J. Yuan, X. Xie, and S. Yang, "Little is much: Bridging cross-platform behaviors through overlapped crowds," in *Proc. AAAI*, 2016, pp. 13–19.
- [156] M. Yan, J. Sang, and C. Xu, "Mining cross-network association for youtube video promotion," in *Proc. 22nd ACM Int. Conf. Multimedia*, 2014, pp. 557–566.
- [157] M. Jiang *et al.*, "Social recommendation with cross-domain transferable knowledge," *IEEE Trans. Knowl. Data Eng.*, vol. 27, no. 11, pp. 3084–3097, Nov. 2015.
- [158] T. Man, H. Shen, X. Jin, and X. Cheng, "Cross-domain recommendation: An embedding and mapping approach," in *Proc. 26th Int. Joint Conf. Artif. Intell.*, 2017, pp. 2464–2470.
- [159] H. Lee, A. Battle, R. Raina, and A. Y. Ng, "Efficient sparse coding algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, 2007, pp. 801–808.
- [160] J. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, Nov. 2008.
- [161] F. Abel, S. Araújo, Q. Gao, and G.-J. Houben, "Analyzing cross-system user modeling on the social Web," in *Proc. Int. Conf. Web Eng.*, Berlin, Germany: Springer, 2011, pp. 28–43.
- [162] L. Chen, J. Zheng, M. Gao, A. Zhou, W. Zeng, and H. Chen, "TLRec: Transfer learning for cross-domain recommendation," in *Proc. Int. Conf. Big Knowl.*, Aug. 2012, pp. 167–172.
- [163] B. Li, Q. Yang, and X. Xue, "Can movies and books collaborate? Cross-domain collaborative filtering for sparsity reduction," in *Proc. IJCAI*, vol. 9, 2009, pp. 2052–2057.
- [164] S. D. Roy, T. Mei, W. Zeng, and S. Li, "SocialTransfer: Cross-domain transfer learning from social streams for media applications," in *Proc. 20th ACM Int. Conf. Multimedia*, 2012, pp. 649–658.
- [165] H. Jing, A.-C. Liang, S.-D. Lin, and Y. Tsao, "A transfer probabilistic collective factorization model to handle sparse data in collaborative filtering," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, Dec. 2014, pp. 250–259.
- [166] S. Qian, T. Zhang, R. Hong, and C. Xu, "Cross-domain collaborative learning in social multimedia," in *Proc. 23rd ACM Int. Conf. Multimedia*, 2015, pp. 99–108.
- [167] W. Min, B.-K. Bao, C. Xu, and M. S. Hossain, "Cross-platform multi-modal topic modeling for personalized inter-platform recommendation," *IEEE Trans. Multimedia*, vol. 17, no. 10, pp. 1787–1801, Oct. 2015.
- [168] M. Yan, J. Sang, and C. Xu, "Unified youtube video recommendation via cross-network collaboration," in *Proc. 5th ACM Int. Conf. Multimedia Retr.*, 2015, pp. 19–26.
- [169] Z. Lu, E. Zhong, L. Zhao, E. W. Xiang, W. Pan, and Q. Yang, "Selective transfer learning for cross domain recommendation," in *Proc. SIAM Int. Conf. Data Mining*, 2013, pp. 641–649.
- [170] S. Yu, X. Wang, W. Zhu, P. Cui, and J. Wang, "Disparity-preserved deep cross-platform association for cross-platform video recommendation," 2018, *arXiv:1712.00733*. [Online]. Available: <https://arxiv.org/abs/1712.00733>



Wenwu Zhu (Fellow, IEEE) received the Ph.D. degree in electrical and computer engineering from New York University in 1996.

He was with Bell Labs, NJ, USA, as a Member of Technical Staff, from 1996 to 1999. He was the Chief Scientist and the Director of Intel Research China from 2004 to 2008. He was a Senior Researcher and a Research Manager with Microsoft Research Asia. He is currently a Professor and the Vice Chair of the Department of Computer Science and Technology, Tsinghua University, the Vice Dean of the National Research Center for Information Science and Technology, and also the Vice Director of the Tsinghua Center for Big Data. He has published over 300 referred articles in the areas of multimedia computing, communications and networking, and big data. He is an inventor or a co-inventor of over 50 patents. His current research interests include cyber-physical-human big data computing, and cross-media big data and intelligence.

Dr. Zhu is an AAAS Fellow, an SPIE Fellow, and a member of The Academy of Europe (Academia Europaea). He was a recipient of seven Best Paper Awards, including the ACM Multimedia 2012 and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY in 2001. He served as the TPC Co-Chair for the ACM Multimedia 2014 and the IEEE ISCAS 2013. He serves as the General Co-Chair for ACM Multimedia 2018 and ACM CIKM 2019. He served on the Steering Committee for the IEEE TRANSACTIONS ON MULTIMEDIA from 2015 to 2016 and the IEEE TRANSACTIONS ON MOBILE COMPUTING from 2007 to 2010. He served as a Guest Editor for the PROCEEDINGS OF THE IEEE, the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, and *ACM Transactions on Intelligent Systems and Technology* and an Associate Editor for IEEE TRANSACTIONS ON MOBILE COMPUTING, *ACM Transactions on Multimedia*, the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, and the IEEE TRANSACTIONS ON BIG DATA. He currently serves as the EiC for the IEEE TRANSACTIONS ON MULTIMEDIA.



Xin Wang (Member, IEEE) received the B.E. and Ph.D. degrees in computer science and technology from Zhejiang University, China, and the Ph.D. degree in computing science from Simon Fraser University, Canada. He is currently an Assistant Researcher with the Department of Computer Science and Technology, Tsinghua University. He has published several high-quality research articles in top conferences, including ICML, MM, KDD, WWW, and SIGIR. His research interests include cross-modal multimedia intelligence and inferable recommendation in social media. He was a recipient of the 2017 China Postdoctoral Innovative Talents Supporting Program.



Hongzhi Li (Member, IEEE) received the bachelor's degree from Zhejiang University, China, in 2010, and the master's and Ph.D. degrees from Columbia University, USA, in 2012 and 2016, respectively, all in computer science. He is currently a Principal Researcher and a Research Manager at Microsoft AI and Research. He has published in ACM Multimedia, TMM, ICMR, EMNLP, NAACL, SPIE, and other venues. His research interests include machine learning, multimodal content analysis, and cloud-based computing, with a focus on deep learning for visual intelligence and its applications on cloud computing platform. He was a recipient of the Best Poster Award in ACM ICMR 2018 and the grand challenge (First Place) in ACM Multimedia 2012. He served on the program committee of major international conferences, including ACM MM, ICME, and IJCAI. He also served as a Reviewer for journals, including the IEEE TMM, the IEEE TCSVT, TPAMI, JVCi, and JVIS.