

Multimodality in Meta-Learning: A Comprehensive Survey

Yao Ma^{a,b}, Shilin Zhao^a, Weixiao Wang^a, Yaoman Li^{a,c,*}, Irwin King^c

^a*Lenovo Machine Intelligence Center, Hong Kong Science Park, Hong Kong*

^b*Delft University of Technology, Delft, The Netherlands*

^c*The Chinese University of Hong Kong, Shatin, NT, Hong Kong*

Abstract

Meta-learning has gained wide popularity as a training framework that is more data-efficient than traditional machine learning methods. However, its generalization ability in complex task distributions, such as multimodal tasks, has not been thoroughly studied. Recently, some studies on multimodality-based meta-learning have emerged. This survey provides a comprehensive overview of the multimodality-based meta-learning landscape in terms of the methodologies and applications. We first formalize the definition of meta-learning in multimodality, along with the research challenges in this growing field, such as how to enrich the input in few-shot learning (FSL) or zero-shot learning (ZSL) in multimodal scenarios and how to generalize the models to new tasks. We then propose a new taxonomy to discuss typical meta-learning algorithms in multimodal tasks systematically. We investigate the contributions of related papers and summarize them by our taxonomy. Finally, we propose potential research directions for this promising field.

Keywords: Meta-Learning, Multimodal, Deep Learning, Few-shot Learning, Zero-shot Learning

1. Introduction

Deep learning methods have made significant progress in the fields of speech, language and vision [1, 2, 3, 4]. However, the performance of these methods heavily relies on the availability of a large amount of labeled data which may be impractical or costly to acquire in most applications. To solve this problem, many researchers have actively explored two promising directions. One is applying the “learning to learn” mechanism to gain or transfer knowledge from prior tasks to improve the learning efficiency in the new task. The other is obtaining heterogeneous modalities to enrich the model’s input, e.g., instead of only looking at the image, feed related text description and the image itself to the model simultaneously. Recently, new state-of-the-art deep learning models have been proposed by modifying the meta-learning algorithm in multimodal scenarios.

The “learning to learn” mechanism [5] used in the human learning process enables us to quickly learn new concepts from very few samples [6]. Existing evidence has shown that humans can gain experience on multiple prior tasks over bounded episodes by combining prior knowledge and context. The learned abstract experience is generalized to improve future learning performance on new concepts. Inspired by this, a computational paradigm called meta-learning [7, 8] is proposed to simulate the ability of humans to learn generalized task expe-

rience. Meta-learning allows machines to acquire prior knowledge from similar tasks and quickly adapt to new tasks. In addition, the process of extracting cross-domain task goals in a dynamic selection [9, 10] makes the meta-learning process more data-efficient than traditional machine learning (ML) models.

Due to the ability of meta-learning to generalize to new tasks, we aim to understand how meta-learning plays a role when the tasks are more complicated, for example, the data source of the task is no longer unimodal. Recent studies have focused on applying the meta-learning framework to the distribution of complex tasks [11, 12, 13], but are limited to a single modality. In particular, meta-learning has proven to be successful in both multi-task and single-task scenarios spanning various applications [7], learning prior knowledge on the optimization steps [14], data embedding [15, 16] or the model structure [17]. However, given the heterogeneous task modalities, the smart usage of meta-learning presents researchers with unique challenges. Each instance in the multimodal scenarios consists of two or multiple items of the same concept but different modalities.

Related research on meta-learning has explored limited approaches to manipulate multiple modalities. The heterogeneous features of different modalities are first explored in ZSL/generalized zero-shot learning (GSZL) for image classification. Semantic modalities provide prior solid knowledge and assist visual modalities in model training, while meta-based algorithms are widely introduced to transfer knowledge more effectively from seen classes to unseen classes to capture the attribute relationships between paired modalities. However, the training process mainly regards one modality as the primary modality and exploits additional information by adding another

*Corresponding author

Email addresses: Y.MA-11@student.tudelft.nl (Yao Ma), szhao4@lenovo.com (Shilin Zhao), wwang29@lenovo.com (Weixiao Wang), ymli@cse.cuhk.edu.hk (Yaoman Li), king@cse.cuhk.edu.hk (Irwin King)

modality. It does not involve the analysis of multiple modalities in real complex scenarios such as unpaired modalities, missing modalities, and associations between modalities.

Therefore, some studies have gone further to apply the meta-learning methods to tasks composed of other modality patterns. The data in some tasks might be missed or unbalanced. Since the modalities of different tasks usually come from different data distributions, the strengths of different modalities could be integrated into the problem to boost performance by taking full advantage of the multimodal data in the context of meta-learning. In the meantime, the training framework of meta-learning can also help improve the generalization ability of the original multimodality-based learner in new tasks.

Existing surveys only discuss meta-learning or multimodal learning separately, resulting in a research gap of the combination of multimodality with meta-learning. This paper provides the first systematic and comprehensive investigation of multimodality-based meta-learning algorithms. Our contributions can be summarized as:

- proposing a new taxonomy for multimodality-based meta-learning algorithms and providing insightful analysis for each category;
- identifying the challenges in the new promising multimodality-based meta-learning area;
- summarizing specific works to solve different challenges, including their methodologies and differences;
- highlighting current research trends and possible future directions.

The rest of the survey is organized as follows. We first briefly introduce the historical development of meta-learning and multimodality, formalizing their definitions. Then, we give the overall paradigm of challenges of multimodality-based meta-learning in Section 2. We propose a new taxonomy based on prior knowledge that meta-learning algorithms can learn in Section 3. We respectively investigate the relevant researches of how to adapt the original meta-learning methods to multimodal data in Section 4, Section 5 and Section 6, followed by a summary of these works in Section 7. Finally, we conclude the paper with current research trends in Section 8 and possible directions for future work in Section 9.

2. Foundations

2.1. A Historical Perspective

The concept of meta-learning first appeared in an article describing how to “learn to learn” [18]. It was further extended by the discovery of using evolutionary algorithms [19] to help learn the rules. More detailed methods only started to flourish after gradient descent and backpropagation were introduced to train meta-learning in 2001 [20, 21]. Table 1 briefly shows the contributions of recent surveys to meta-learning. As an alternative paradigm to further promote the performance of ML with low-shot data, meta-learning has been shown to improve the generalization ability of computing models [22, 23]. Even

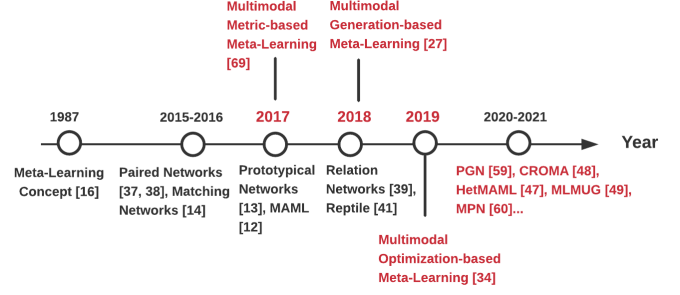


Figure 1: Chronological view of partial multimodality-based meta-learning methods. The year with **red** indicates the first introduction of multimodality for the corresponding model. Methods with **black** indicate traditional meta-learning models. Methods with **red** all belong to the category of multimodality-based meta-learning.

though ZSL/GZSL algorithms for image classification have explored auxiliary modalities in the multimodal embedding space, the most common improvement is still applied on the unimodal task sets [14, 15, 24, 25]. Multimodal research first appeared in the exploration of audiovisual interaction in the process of speech perception [26]. Later, with the massive increase of digital multimedia content, content analysis for multimodal sources became increasingly popular, along with combining deep neural networks [27] in applications such as visual question-answering [28] and image captioning [29]. The promising performance of multimodal learning models for visual semantic interaction is largely reflected in the ability to infer the potential alignment of images and semantics between modalities. However, this methodology cannot be directly applied to the case of FSL/ZSL. The pursuit of multimodality shows that training data with one modality can benefit from the alignment of different modalities [30], while maintaining robust knowledge transfer between unseen and seen classes. Figure 1 displays the historical development of meta-learning methods and highlights several time frames for the first introduction of multimodality.

Table 1: Recent meta-learning survey contributions between 2010 and the present.

Year	Reference	Taxonomy	Architecture	Trends	Applications
2012	[31]	✗	✗	✓	✗
2015	[32]	✗	✓	✓	✗
2018	[33]	✓	✓	✓	✗
	[7]	✓	✗	✓	✓
	[34]	✓	✗	✗	✓
2020	[8]	✓	✓	✓	✓
	[35]	✓	✓	✓	✓
	[36]	✗	✗	✗	✓
2021	[37]	✗	✓	✓	✓
	[38]	✗	✗	✗	✓

2.2. Formalizing Meta-Learning

A common perspective of meta-learning is to learn algorithms that can be generalized across tasks. Following the definition in [7], we denote the set of M source tasks used in the meta-training stage as $\mathcal{D}_{source} = \{(\mathcal{D}_{source}^{train}, \mathcal{D}_{source}^{val})^{(i)}\}_{i=1}^M$ where each task is trained on $\mathcal{D}_{source}^{train}$ to minimize the internal

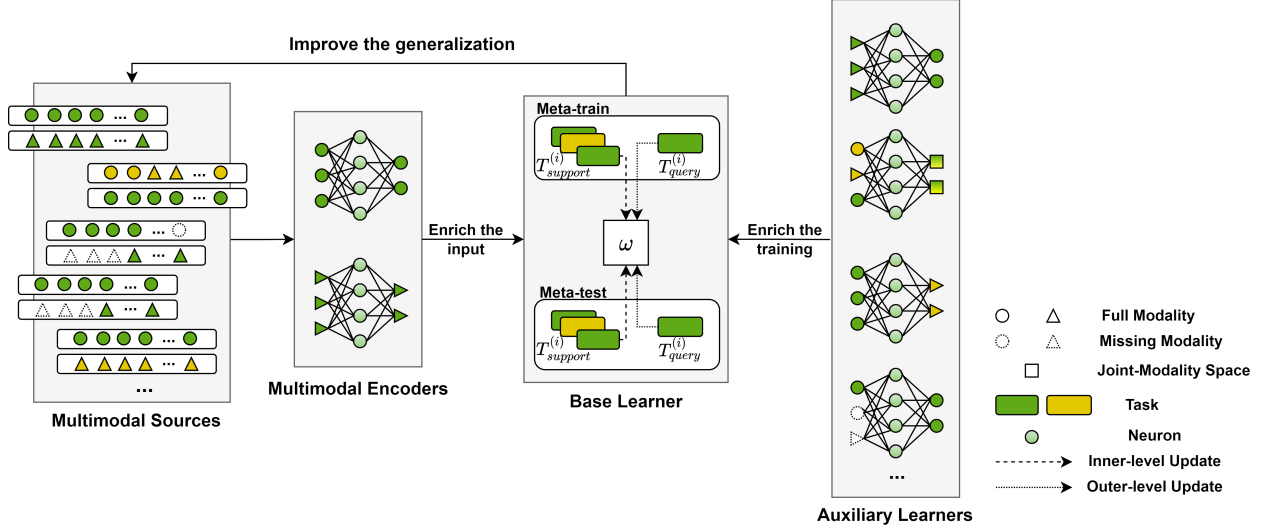


Figure 2: Paradigm of multimodality-based meta-learning. (1) Given the task segmentation of a multimodal dataset, we first list some possible task combinations with different modalities, including same tasks with same modalities, different tasks with mixed modalities, same tasks with missing modalities, different tasks with different modalities, etc. (2) Each modality needs to be encoded with the corresponding multimodal encoder as an enrichment to the input of subsequent models. (3) The base learner represents the main network trained by any meta-training algorithm that learns the meta-parameter ω^* . The dataset is split for the meta-train and meta-test stage, in the form of support set $T_{support}^{(i)}$ and query set $T_{query}^{(i)}$. (4) Auxiliary learners are provided to tackle the multimodality-related problems, such as generating conditional modalities, mapping heterogeneous task distributions to the joint-modality space, aligning cross-modal tasks, reconstructing the missing modalities, etc. These learners are integrated with the base learner to enrich the original meta-training process. (5) As a result, the multimodality-based meta-learning paradigm finally improves the generalization across tasks from multimodal sources.

loss function \mathcal{L}^{task} in Eq. (2). Then the meta-knowledge ω^* is learned when evaluating on $\mathcal{D}_{source}^{val}$ over all the source tasks:

$$\omega^* = \arg \min_{\omega} \sum_{i=1}^M \mathcal{L}^{meta}(\theta^{*(i)}(\omega), \omega, \mathcal{D}_{source}^{val(i)}), \quad (1)$$

$$\text{s.t. } \theta^{*(i)}(\omega) = \arg \min_{\theta} \mathcal{L}^{task}(\theta, \omega, \mathcal{D}_{source}^{train(i)}). \quad (2)$$

Usually, $\mathcal{D}_{source}^{train}$ and $\mathcal{D}_{source}^{val}$ are called support sets and query sets respectively. For the set of Q target tasks in meta-test stage $\mathcal{D}_{target} = \{(\mathcal{D}_{target}^{train}, \mathcal{D}_{target}^{test})^{(i)}\}_{i=1}^Q$, the learned prior knowledge ω^* in the meta-training stage will be used to train the model on unseen tasks.

The meta-learning paradigm discussed in this survey has two key properties:

- The existence of a task segmentation mechanism. The meta-training and meta-test sets [7] contain task units, where each task is divided into the support set and query set. Prior knowledge accumulated in the meta-training stage is used in the meta-test to evaluate the accuracy of the meta-model.
- The process of constructing an adaptive learner. This refers to the accumulation of meta-knowledge experience by dynamically improving the deviation between the inner base learner and the outer generalized learner. The type of meta-knowledge could be the estimation of initial parameters [14], or an embedding strategy [15]. Ultimately, it guarantees that the adapted knowledge can directly generalize to predicting new tasks instead of adding extra data.

2.3. Formalizing Multimodality

We consider data modalities to be multiple when data is collected through multiple sensors, measurement equipment or acquisition technologies [39]. The output of each sensory channel is represented as a modality, which is usually a single form of dataset associated with a medium of expression, such as vision data from seeing objects and audio data from hearing sounds. The key characteristics constituting multimodality that cover data coming from multiple modalities can be summarized as follows:

- **Complementarity.** Each modality brings a certain type of added value to the whole, and these added values cannot be inferred or obtained from any other modality in the setting [40].
- **Diversity.** The meaning is composed of different modalities from different semiotic resources.
- **Integrity.** The making of meaning involves the overall attention to the potential and limitations of each modality [41].

Our key assumption is to combine multimodal features, such as images, fine-grained descriptions, auxiliary audio and video information to force the model to recognize cross-modal discriminative features to promote usage in meta-learning applications. Under this assumption, we limit the scope of our research to models that use multiple sensors as inputs, thereby excluding studies that build their models on multimodal task distributions with disjoint modes [11, 12] or shifted domains [42].

On the other hand, there are various combinations of multimodal sources in different tasks (See Figure 2). Modalities

could be provided in pairs or across tasks. The diversity and integrity of modalities are therefore reflected in the dependence on the semantics of different modalities. New classes with insufficient training data in unimodality can benefit from previously learned features. When one of the modalities is absent, the complementarity between modalities provides a way to conditionally consider auxiliary semantics and support the global meaning.

2.4. Challenges of Multimodality-Based Meta-Learning

As a computing paradigm, meta-learning mainly learns shared meta-knowledge among different tasks and expresses it by transferring knowledge from seen classes to unseen classes. Integrating this framework with multimodality involves (i) using multimodal properties to enrich the task input of the meta-based model, and (ii) using auxiliary multimodality-based learners to enrich the task training of the meta-based model. Figure 2 illustrates the overall paradigm.

For problem (i), the research challenge is to make the learned meta-knowledge incorporate values from multimodal sources rather than a single modality. At the same time, the generalization ability of meta-knowledge in new tasks should also exceed the performance in unimodal situations. Different modalities have been encoded into relevant embeddings before being fed into the meta-learning network. Therefore, the base learner does not need to care about the specific representation of the modalities but can focus on how to use existing algorithms to adapt to multimodal inputs, including whether to use multiple modalities separately or to use fused modalities, how to combine modalities wisely and so on. In addition, there are scenarios where the meta-learning framework is only employed as an episode-based training method which could facilitate training efficiency. When multimodal inputs are added to such applications, the challenge becomes how to use the meta-learning framework to divide the multimodal task sets to achieve the ability of generalization.

For problem (ii), auxiliary tasks originating from multimodal problems have brought extra challenges to train the auxiliary learners and the base learner in a meta-learning manner. Common practices in multimodal ML [43] address how to represent, summarize, map, align and fuse multimodal data in a way that exploits the characteristics of multiple modalities. When we integrate these methods with the base learner, auxiliary learners must participate in the co-training process to deal with the multimodal challenges. In other words, the essence of meta-learning then becomes how to provide a parameterized training framework, which means that the addition of these auxiliary networks needs to establish a relationship with the training parameters of the base network before they can be jointly trained. Hence, it is non-trivial to enrich the original meta-training process by adapting the inner-level update and the outer-level update to new networks that incorporate multimodal training, thus providing extra meta-knowledge about the prior modalities.

Furthermore, although we discuss the influence of modalities on the meta-learning framework from a multimodality-based perspective, the influence between these two concepts is mutual. Since the meta-learning training framework can help tra-

ditional multimodal learning to better generalize to new tasks, additional challenges may also include elements that affect the effectiveness of generalization, such as the definition of multimodal tasks, the choice of meta-learning algorithms, and the number of modalities.

3. Proposed Taxonomy

Recently, many categorizations of meta-learning algorithms have been proposed [7, 9, 37] to explore what meta-knowledge can be learned. However, they mostly refer to application scenarios of the algorithms and ignore the influence of input data sources on the algorithms. To bridge the gap in the interaction between the input of multimodal sources and the form of algorithms in the current research, we are the first to introduce a new multimodality-based meta-learning taxonomy. We still establish the taxonomy on the type of meta-knowledge that can be learned, due to the distinct training methods employed by meta-learning algorithms.

In contrast to previously proposed categorizations, we discuss the algorithm details after combining with multimodal characteristics and the specific adjustments made for different multimodal alignments. Our taxonomy is based on three types of meta-knowledge: learn the optimization, learn the embedding, and learn the generation. For each meta-knowledge, we further provide a specific subcategory that reflects the influence of different multimodal meta-representations on the algorithms. When the relevant algorithms involve different variants operating on multimodal data, we will further discuss patterns presented in these variants as meta-attributes. Figure 3 shows the taxonomy.

Learn the Optimization Optimization-based meta-learning methods focus on learning generalized outer-level parameters across tasks [7, 9]. In a unimodal application scenario, the optimization usually only involves the main task such as the refinement of the base network structure [14], without fine-tuning possible parameters in auxiliary tasks. When applying the learning paradigm on multimodal data, such multimodality needs to be parameterized before being trained along with the base network. Mathematically, the inner-level optimization of Eq. (2) can be modified as:

$$\theta_m^{*(i)}(\omega) = \arg \min_{\theta_m} \mathcal{L}^{task}(\theta_m, \omega, \mathcal{D}_{source}^{train(i)}), \quad (3)$$

where θ_m denotes the collection of multimodal parameters, which are supposed to be diverse, such as the relationship of different modalities, the joint distribution of multiple modalities, and the weights of missing modalities.

As far as we know, current optimization-based methods are dedicated to exploring how to use the meta-learning framework to help modify the training paradigm of initial parameters that define the modalities [44] or relevant auxiliary networks [45, 46]. We, therefore, summarize the learnable meta-representations as parameterized-modal initialization and unified-modal optimizer. To capture how modality parameters

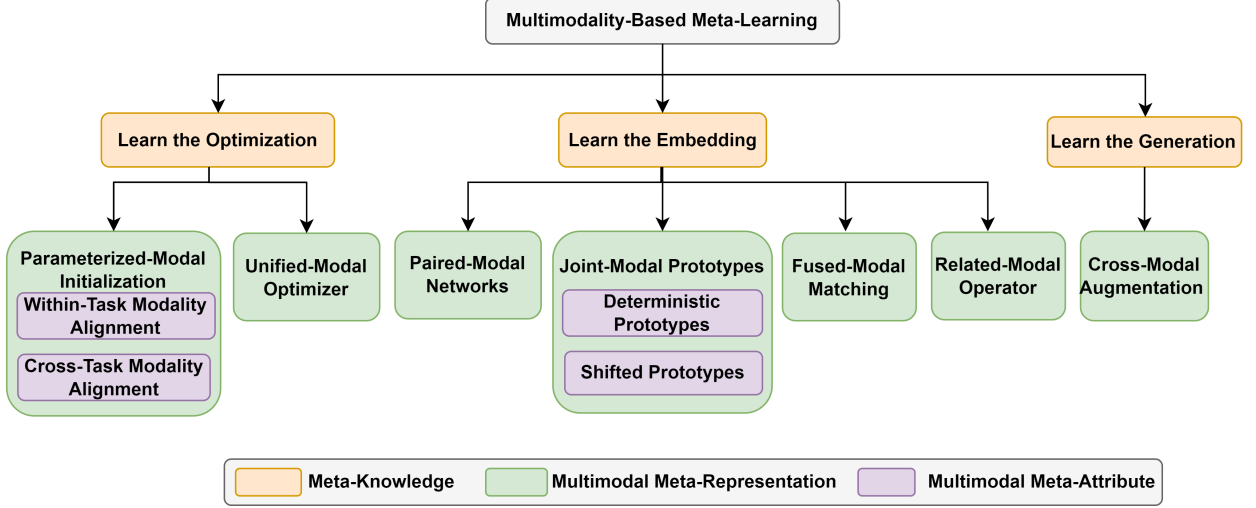


Figure 3: Taxonomy of multimodality-based meta-learning.

are optimized in different task distributions, within-task modality alignment and cross-task modality alignment will be discussed in Section 4.1.

Learn the Embedding Metric-based meta-learning methods are widely used as non-parametric algorithms for few-shot problems. The idea is to learn an embedding network that helps project the training and testing points onto the same space to implement similarity comparison. The embedding network is trained on unimodal support sets where the quality of feature embeddings is often restricted. The introduction of multimodality expands the number of original embedding spaces. We emphasize that the embedding network aligning multiple spaces or adapting to the selection of multiple spaces can be applied to match any source of multimodal information, with the ability to capture the semantic relationship between input pairs more effectively.

Here, we summarize four widely used approaches originated from Pairwise Networks [47, 48], Prototypical Networks [15], Matching Networks [16] and Relation Networks [49]. For each one in the multimodality-based scenario, modalities are respectively provided with pairs for the pairwise networks, combined to train a new joint prototype, fused to formulate the attention kernel, and related together to form the concatenation for the non-linear operators. Especially for variants of Prototypical Networks, prototypes play different roles in the model due to the different stages of multimodal fusion. Details about the difference between deterministic prototypes and shifted prototypes will be discussed in Section 5.2.

Learn the Generation Another meta-knowledge is data generation, which is widely employed for most unimodal few-shot and zero-shot applications. Since the source of unimodal data is relatively simple and scarce, we exploit prior knowledge of the data to modify the model or execute data augmentation. There are often implicit relationships between multiple modalities, which can help better describe the scarce label characteristics across tasks. Although not many, we notice that some researchers have worked on the knowledge of conditional prob-

ability to simulate distribution of the primary modality conditioned on the auxiliary modality [30, 50]. We summarize the related methods as cross-modal augmentation with generating data across modalities.

4. Learn the Optimization

Optimization-based meta-learning methods aim to solve the inner-level task as an optimization problem [7] and extract the meta-parameter ω^* through Eq. (1) to obtain the best performance for generalizing to new tasks within only a few gradient-descent updates. Most of these methods are built on top of the Model-Agnostic Meta-Learning (MAML) [14] framework, which is a popular approach to meta-learn an initial set of neural network weights adapted for fine-tuning on few-shot problems. Reptile [51] generalizes the first-order MAML by repeatedly sampling a task and training to minimize the loss on the expectation over tasks. More alternatives that fit the meta-optimization process illustrated by Eq. (1) and Eq. (2) also learn specific inner optimizer by producing the trainable ω^* as its own hyperparameters [52]. The exact optimization algorithm [17] will then generalize to unseen tasks to optimize the inner learner directly.

Multimodal information can usually be parameterized and wrapped into the inner-level optimization task, and then trained iteratively with the meta-parameter, either for the initialization or the optimization algorithm. We summarize the relevant studies in terms of where to use the meta-parameter and how it is co-trained with multimodal parameters.

4.1. Parameterized-Modal Initialization

4.1.1. Overview

Meta-learning algorithms such as MAML and Reptile are known to fine-tune parameters of unseen tasks quickly by using gradient descent methods. Since seen and unseen tasks are involved in the learning process across multiple domains, the

diversity of the task space distributions puts forward higher requirements for applying the initialized meta-parameter. One common task distribution comprises tasks with the task instances of the same manifold from the same domain, where all tasks are either unimodal or multimodal. In addition, tasks may come from the same domain, but the task instances have different structures or subspaces [46], such as using different modality features to describe the same semantic concept. We have reasons to believe that although the outer learner is still learning initial values that can be shared across different tasks, the generalization ability of these initialized parameters can be applied to more complex and multi-space task distributions. On the other hand, the heterogeneity gap [53] makes it more challenging to align multiple modalities into the same feature space across tasks. Therefore, according to the distribution pattern of parameterized multimodal information in the task space and how they are optimized to be aligned in the inner learner, we divide the current research into two branches, within-task modality alignment and cross-task modality alignment.

For the within-task modality alignment, the model’s input comes from the homogeneous task space with different modality characteristics. Each task adopts the same type and number of modality feature spaces, such as paired modalities. The inner learner does not need to pay attention to the domain adaptation problems, but only needs to process the alignment of different modality subspaces [46], such as reducing the dimensionality of some modalities, training a hybrid neural network for the mixed modality spaces, and constructing missing modality subspaces aligned with the known modalities [45]. Finally, the meta-learning framework provides an iterative method of optimizing multiple parameters from multiple networks (i.e., inner network, outer network) together. Such a training framework still keeps the original structure of MAML or Reptile, in which the inner optimization is enriched by multiple modalities or other auxiliary training networks related to multimodal scenarios.

Multimodal tasks are usually sampled from a heterogeneous and more complex task distribution for the cross-task modality alignment. Each training task may contain a specific combination of modality subspaces. Most existing model-agnostic meta-learners assume that the tasks are evenly distributed, in which all tasks belong to the same domain and have the same manifold of task instances. However, such heterogeneous task distributions are often more challenging because instances with different task subspaces cannot share the same model structure completely. Therefore, related research examines the common and unique meta-parameters of different task spaces, exploring how to meta-align the knowledge of different subspaces to generalize to new tasks. The learned meta-knowledge ω^* is supposed to incorporate task-aware parameters as part of the initialization, where the task-specific information for each type of modality subspace may share the same manifold to adapt to the model-agnostic meta-learners.

4.1.2. Methods

Within-Task Modality Alignment The parametric nature of the MAML framework enables multimodal problems to be

characterized as the internal optimization involved in meta-training, where the inner network can be extended to multiple. Ma et al. [45] focus on the implementation details by extending MAML to learn three networks in an integrated way for the main training process, missing modality reconstruction network and feature regularization by a Bayesian meta-learning based model, SMIL. There is no modality across different tasks involved obviously in training, but the use of a feature reconstruction network has leveraged the auxiliary modality to generate an approximation of the missing-modality feature efficiently. It has been highlighted in the paper that such a method is more efficient than the traditional generative methods (e.g., GAN, AE, or VAE) for not requiring full-modality data. The usage of single modality embeddings also guarantees that the approximation of full-modality is more flexible than the typical methods where feature spaces are regularized by the perturbation. Despite the successful collaboration training between modalities, an open challenge for deploying such Bayesian training framework in the real-world scenario is how to learn a set of proper modality priors for the unknown modalities.

Moreover, the inner networks can be multimodal encoders for different features. Yao et al. [44] try to reduce the risk of unstable results and knowledge transferred from a single source city. They propose a MetaST network in a meta-learning manner to learn to represent the initialization of spatio-temporal correlation based on tasks from multiple source cities. The estimated spatio-temporal values are computed by the hybrid model ST-net that combines a CNN and an LSTM to model each region’s spatial dependency and temporal evolution. The parameters of the ST-net are updated in the inner-optimization to give representations of different modalities. Meta-parameters such as time dependence, spatial proximity, and regional function will be easily adapted to the target city by fine-tuning, and finally, serve for the relevant task predictions. MetaST has overcome the risk of transferring the knowledge from a single source and adapted the spatio-temporal sequences in various scenarios rather than only discrete features. However, the type of shared knowledge is still in the state of a black box, which poses a challenge to broader scalability. Figure 4 compares the above parameterized inner learners that could enrich the training of the meta-learning framework.

In addition to the application of conventional multimodal task sets, the meta-learning framework can also help improve the training performance of neural network models in other real-world settings, such as helping agents in reinforcement learning (RL) to better adapt to different scenarios and helping solve the catastrophic forgetting of previous tasks in continual learning.

Specifically, Yan et al. [54] adopt the MAML framework in the experiment of indoor navigation from the perspective of RL. Tasks are divided based on various room scenes to help train and test the model using the aggregated features of visual and audio modalities. Different features of visual observations are encoded to aggregate the word embedding from audio features and then inputted to the memory network to carry out a sequence of action policies. The good initialization used for navigating different scenes is trained along with the inner pa-

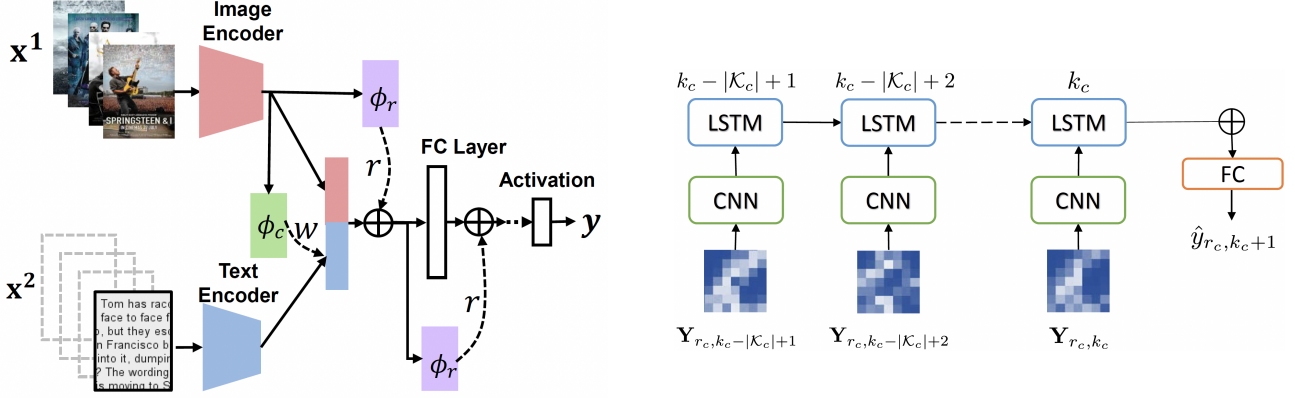


Figure 4: Network structures for parameterized inner learners. **Left [45]:** SMIL extends MAML by training the main network, the reconstruction network ϕ_c , and the regularization network ϕ_r together. x^i denotes different modalities. ϕ_c and ϕ_r together denote the inner optimization parameters while the main network is responsible for learning the meta-parameter ω^* . **Right [44]:** ST-net is parameterized by θ which encrypts the spatio-temporal correlations. θ denotes the overall parameters from CNN and LSTM where multimodal information is encoded. The initialization of ω from multiple source cities should be meta-learned.

parameters that control the optimization of the agent’s interaction loss. It is evidenced that the meta-learning framework makes the model more robust to unseen scenes, but further validation on a real programmable robot is still missing.

Verma et al. [55] introduce a new meta-continual scenario to handle unseen classes that are collected in a sequential stream dynamically. While the main architecture is based on pairing self-gating of attributes and scaled class normalization, the need for balancing even generalization across all tasks during the reservoir sampling requires the model to be trained with Reptile in a meta-learning manner. Such an idea of avoiding expensive generative models has implications for applying multimodality-based meta-learning models in a real-world setting, where the one-time adaptation paradigm of the streaming data possibly fails. Similar work can be inspired to focus on catastrophic forgetting in the sequential data.

We can conclude that for the within-task modality alignment, the methods mainly focus on transforming modality issues into the design of auxiliary learners to enrich the training process. Flexible applications can be proposed as long as the involved multiple modalities can be parameterized and added to the gradient-descent updates.

Cross-Task Modality Alignment The most common cross-task alignment is used to learn shared knowledge from different domain tasks. Modalities often appear in similar patterns in each task, but the semantic difference of different domains leads to the difficulty of transferring meta-knowledge across tasks. TGMZ [56] addresses the limitations of meta-ZSL models that are mostly optimized on the same data distribution without explicitly alleviating representation bias and prediction mismatch posed by diverse task distributions. It uses an attribute-conditioned auto-encoder to align multiple task domains in a unified distribution with auxiliary textual modalities. Different from this study [50] which searches optimal parameters for each task, the meta-training relies on the task-specific loss function to compute gradients on support and query sets for obtaining the overall optimal parameters of different modules, task encoder, task decoder, task discriminator, and task classifier. The

performance of TGMZ over ZSL/GZSL has demonstrated the necessity of task alignment in alleviating discrimination in class distributions. This suggests a potentially robust way of training meta-ZSL models by speculating more about the disjoint task distribution.

In addition to the existence of multi-domain labels in image classification applications, other multimodal scenes also include task sources from different domains. MLMUG [59] focuses on the user domain generalization issue that exists in cross-modal retrieval. To allow the model to generalize from source user domains to unknown user domains without any update or fine-tuning, the model first constructs a cross-modal embedding network to learn a shared modality feature space for cross-modal matching, which generates a shared mask used to encode transferable knowledge between different user domains. Then a meta-learning framework is implemented to learn the meta covariant attention module. Unlike the original MAML, the gradients from the meta-training set and meta-test set finally get weighted aggregation. MLMUG has provided a potential training paradigm for applications that cannot obtain unlabelled data instances from unknown fields in the real world, which possibly are more efficient than unsupervised domain adaptation. For applications with large domain spaces, this method gives access to generalized parameter learning instead of expensive joint training.

Despite the success of task-aware models in multiple task domains, different tasks are still limited to the same modality patterns. Another more generalized task distribution will be with different modality subspaces for different tasks. To align multiple subspaces, HetMAML [57] tries to extend the application of gradient-based meta-models from homogeneous task distributions to heterogeneous tasks, where tasks in the same query set have different numbers of modalities or input structures. Compared with training different types of tasks separately, HetMAML aims to train a unified meta-learner that can simultaneously capture the global meta-parameters about the concept domain shared by all tasks, as well as the customized parameters that characterize each task. A three-module shared

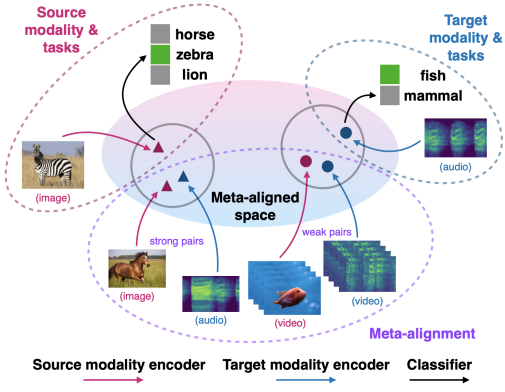
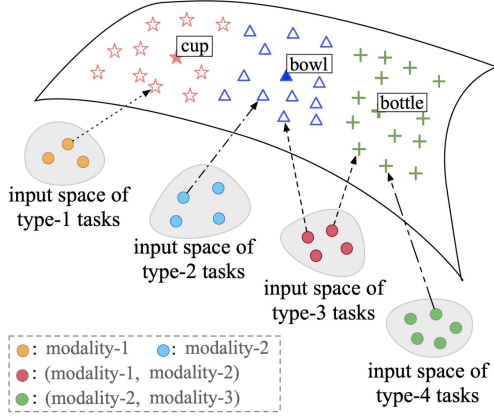


Figure 5: Meta-alignment representations of more generalized task distributions. **Left [57]: Heterogeneous task distributions.** Each task has its unique subspace describing the commonly shared concept domain. **Right [58]: Cross-modal task distributions.** Source tasks and target tasks have different modalities to be aligned.

network architecture is proposed to achieve this goal, including multi-channel feature extractors, feature aggregations, and the head module for decision making. By adopting the BRNN structure, different modality-specific feature extractors are iteratively aggregated to enable all types of tasks to be mapped from heterogeneous feature space to a unified space. A better generalization of a few examples is not only reflected in the common meta-parameters shared between different types of tasks, but also in the type-specific meta-parameters for each task. HetMAML is a potentially powerful extension of the MAML framework, which not only retains the model-agnostic property of MAML, but embeds task-aware knowledge into the outer-level update.

Furthermore, CROMA [58] goes further to establish a more generalized meta-learning paradigm based on cross-modal alignment that is trained on different source modalities, and then quickly generalize to the target modality to perform new tasks. Specifically, they jointly train the cross-modal meta-alignment space and source modality classifier in the meta-training process to perform generalization on alignment and classification tasks by Reptile. The meta-parameters learned in the training process will be used as initial parameters for the meta-test process to classify tasks in the target modality. When the labeled data of the target modality is scarce, the joint space learned by the meta-alignment of the source modality and target modality will help transfer knowledge among tasks. Figure 5 compares the meta-alignment representations of heterogeneous task distributions and cross-modal task distributions.

When generalizing gradient-based meta-models to the cross-task modality alignment, the extent of what and how much knowledge across modalities should be shared has been heavily emphasized throughout the existing research. Both tasks from different domains or the same domain concept are parameterized and updated in the outer learning process to adapt to unknown tasks quickly. The major challenges of transferring knowledge under this framework remain open, such as implementing task-specific loss, the choice between modular and hierarchical network architectures, and the construction of the

shared modality space.

4.2. Unified-Modal Optimizer

4.2.1. Overview

The meta-learning algorithm could learn the meta-optimizer by updating the inner learner at each iteration. The most typical one is the LSTM-based meta-learner [17], aiming to use the sequential features of the LSTM model to simulate the cell states during the update process. The updating process approximates the gradient descent method, where the information stored in forget and input gates during training is improved. The parametrization of the optimizer allows new tasks to be learned from a series of previous tasks. When the multimodal training network is introduced, the meta-learner still keeps the idea of optimizing the optimizer to help update a better multimodal network, depending on the specific structure of the network.

When the initial value is meta-learned for multimodal tasks, the subspace of each modality needs to be learned together through an aligned common space. However, the modality input in the real world is often dynamically changing, and the adaptation of learning frameworks to different modalities under the sequential distribution is challenging. Recent work has employed the idea of human cognitive learning to address the adaptive training of multimodal models, where various modalities do not need to be aligned during training at a fine-grained level. Therefore, the LSTM-based meta-learner becomes an appropriate alternative to take advantage of the previous experience in the sequence to avoid catastrophic forgetting.

4.2.2. Methods

While most articles apply all the modalities to each task together, another way of thinking is to add the modalities sequentially, dynamically adjusting the performance of each task on the unified multimodal model. Ge and Xiaoyang [53] specifically propose the sequential cross-modal learning (SCML) from a novel perspective in terms of learning the optimization of the unified multimodal model. Each modality is learned sequentially by the unique feature extractor and then projected

onto the same embedding space by the unified multimodal model. Then they extend the LSTM-based meta-learner [60] to effectively optimize the new unified multimodal model based on the old experience well-trained on previous tasks. It is evidenced that due to the ability of the meta-learner, the updated unified model obtains a slight increase over the ordinary gradient descent method.

From the number of articles that use LSTM-based meta-learner, we can infer that the application of the LSTM model to deal with dynamic sequential multimodal data is still limited. While mapping multiple modalities to the same embedding space simultaneously has been adopted as the common practice, sequential learning of multiple modalities by the LSTM-based meta-learner is undoubtedly promising.

5. Learn the Embedding

Metric learning approaches learn the non-linear embedding space, where intrinsic class memberships are decided by measuring the distances between points. The learned prior knowledge is often implemented as an embedding network or a projection function that transforms raw inputs into the representation suitable for similarity comparison [7] in a feed-forward manner [15]. Popular embedding networks include Siamese Network [47], Triplet Network [48], Matching Networks [16], Prototypical Networks [15] and Relation Networks [49]. Other advanced approaches such as graph-based networks [61, 62] are also widely explored to model the relationship between samples. In addition, metrics such as Euclidean distance [63], cosine distance, contrastive loss, and triplet loss [64] are typically used to measure the similarities within pair or triplet samples. The flexible embedding network structures and metric computations provide convenience for introducing multimodality, which promotes the feature extraction and interpretation to be realized in various ways.

5.1. Paired-Modal Networks

5.1.1. Overview

Pairwise networks are introduced to take pair examples and learn their sharing feature space to discriminate between two classes. For example, Siamese Network [47] employs two identical neural networks to extract embeddings from a pair of samples and computes a weighted metric to determine the similarity. Triplet Network [48] extends the networks to three with shared parameters to output the comparison probability. Although the pairwise comparators limit the possibility of training end-to-end networks directly for the few-shot problems, the learning framework is still valid to be applied to multimodal datasets. A common technique is to change the identical architectures that share parameters in the pairwise networks to different embedding networks for different modalities. New loss metrics are often proposed along with the networks to perform the matching task between inputs from different modalities.

5.1.2. Methods

As the output of paired networks, a shared feature space allows the use of different modalities under the same distance metric. Liu and Zhang [65] propose STUM to take advantage of the siamese feature space formation process [66] to learn a shared feature space by different networks for different modalities. Instead of using triplet loss during training, it also employs a simple summation of the contrastive loss function to handle the random number of streaming data representations coming from the time relation of inputs. Positive and negative samples group the inputs in one or more modalities that are adapted to non-identical networks for processing. STUM has effectively captured the formation process of the feature space within and across the modalities in time-cued data. The experiment on visual modalities has shown that the model can potentially be used for the feature organization of objects represented by multiple modalities. The high performance achieved without supervision also implies the potential extension to fast online learning.

Triplet loss is also often used as a loss variant in combination with paired networks to enhance the robustness of the model. Eloff et al. [69] focus on learning a multimodal matching space from paired images and speech in the one-shot domain. They investigate a Siamese neural network built with the semi-hard triplet mining trick to alleviate the memory issue. Each spoken test query is matched to a visual set according to the learned embedding metric from training the support set. However, they do not approach the matching problem in a true multimodal space where the mapping happens mutually between two modalities, therefore transforming it into a two-step indirect comparison in the visual embedding space.

To learn a multimodal space in one step, Nortje and Kamper [70] improved upon the above idea to learn a shared embedding space of spoken words and images from only a few paired examples by optimizing the combination of two triplet losses. The model they came up with is MTriplet, which overcomes the weakness of learning representations that rely on unimodal comparisons despite having the support set of speech-image pairs. After preprocessing the two modalities by the corresponding autoencoders, a multimodal triplet network is learned to map inputs of the same class to similar representations measured by a direct cross-modal distance metric.

Even though paired-modal networks have demonstrated the potential for feature space formation, modality patterns extracted from real-world datasets remain an open challenge. When the number of modalities increases from two to multiple, it is unclear whether the network structure could achieve high performance under conditions such as missing modalities and scarce paired samples.

5.2. Joint-Modal Prototypes

5.2.1. Overview

The term “prototype” refers to the centroid of each class within the dataset. In Prototypical Networks [15], a support set is used to calculate the prototype for each class, and the query samples are classified according to the distance from each prototype. The calculation of the prototype p_c needs to rely on the

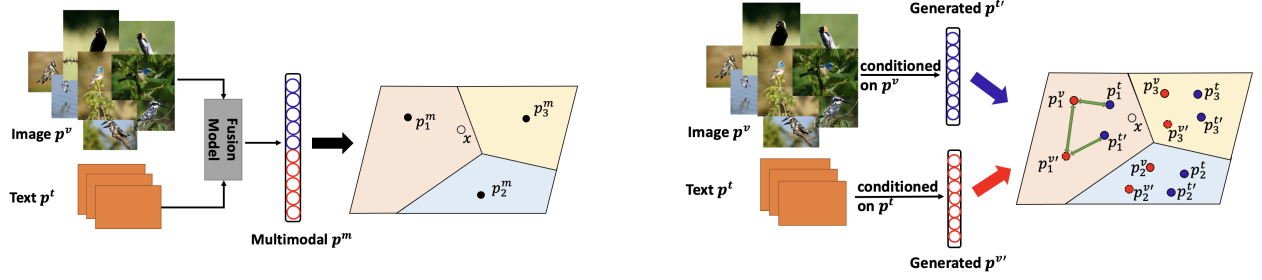


Figure 6: Representation of joint-modal prototypes. **Left: Deterministic prototypes.** Modalities are fused through the fusion model and transformed as the joint representation. Each multimodal prototype p^m in each class is projected onto the multimodal embedding space. The query sample x is then compared to different multimodal prototypes to obtain the prediction. **Right: Shifted prototypes.** New prototypes are generated by conditioning on the original prototypes in one modality. Flexible choices on original averaging prototypes and generated prototypes are introduced. Take one class p_1 as an example, AM3 [6] adaptively combines p_1^v and p_1^t . Episode-based PGN [67] modifies the objective function by integrating both of the generated prototypes $p_1^{v'}$ and $p_1^{t'}$. MPN [68] averages the conditioned $p_1^{v'}$ with the original p_1^v . The new prototype is supposed to be anywhere on the green line based on the value of the weighting factor.

averaged embeddings of all the support samples in class c for each episode of training [6]:

$$p_c = \frac{1}{|S_e^c|} \sum_{(s_i, y_i) \in S_e^c} f_\theta(s_i), \quad (4)$$

where $S_e^c \in S_e$ is the subset of support set belonging to class c , $S_e = \{(s_i, y_i)_{i=1}^{N \times K}\}$ and f_θ is the embedding network that needs to be learned. S_e contains K samples for each of the N classes. In general, after acquiring the distances of the query embedding to the embedded prototype p_c in each class, we can obtain the distribution of the query sample over all of the classes of the episode. The meta-objective of the model then becomes to minimize the expectation of the negative loglikelihood of the true class of each query sample [6]:

$$L(\theta) = E_{(S_e, Q_e)} - \sum_{i=1}^J \log p_\theta(y_i | q_i, S_e), \quad (5)$$

where the query set $Q_e = \{(q_j, y_j)_{j=1}^J\}$ contains J samples. $(y_i, q_i) \in Q_e$ and S_e are the sampled query set and support set in each episode of training.

When the prototype embeddings are computed from more than one modality, we observe that the representation of prototypes could be modified according to the way of introducing different modalities. By exploiting whether modalities are fused before passing to the prototypical networks, we divide the representation of joint-modal prototypes into two categories, **deterministic prototypes** and **shifted prototypes**.

Deterministic prototypes use a unimodal vector extracted from the fusion of other models, which means usually different modalities are trained by different encoders, followed by the connection layer to concatenate the outputs after training. The prototypical networks are expected to use the concatenated unimodal vector to perform the following similarity calculation, so there is no obvious difference from the traditional prototypical networks. This method cares more about multimodal learning rather than the meta-learning framework, such as how to remove correlations between modalities and fuse them in a lower-dimensional common subspace.

By contrast, the early fusion of modalities may add noise to the embeddings since the information retained by different modalities is homogenized in this process. The idea of relying on learned embeddings and their distances to discriminate unseen classes in the prototypical networks has a similar problem structure to alignment, a key challenge in the field of multimodal learning. Alignment focuses on how to identify the direct relationship between elements in two or more different modalities [43]. Existing research usually maps different modalities to the same semantic representation space, and then computes the similarity as a direct relationship measurement. The phenomenon of multimodal representations in the same semantic space allows extending the prototypical networks.

Relevant research often applies unimodal task training resulting in unimodal feature vectors of the prototypes. Accordingly, some studies have introduced multimodal feature vectors to shift the prototype representations. In that case, we define the method of shifted prototypes as the one that uses the various multimodal vectors directly without prior fusion and builds models upon the ensemble. Figure 6 shows the two types of joint-modal prototypes with two commonly used modalities, image and text, as an illustration.

5.2.2. Methods

Deterministic Prototypes. The computation of distances between deterministic prototypes in the joint space of different modality features follows the same approach as for traditional unimodal prototypes. The metric of the meta-learning algorithm does not require special modification. Wan et al. [71] propose to learn the prototype for each class of the social relation extracting from the combination of facial image features and text-based features. The unbalanced distribution of social relations from real-life multimodal datasets motivates them to propose prototypical networks training on a random support set of tuples on different classes of social relations by using FSL techniques. A cross-modal encoder (Illustrated as the Fusion Model in Figure 6 Left) then concatenates the normalized feature vectors of two modalities learned respectively from a pre-trained language model and FaceNet. Eventually, the prototyp-

ical networks are applied to predict the social relations in the query set by directly averaging concatenated cross-modal embeddings as the prototype for each class.

Deterministic prototypes are easy to obtain, which allows for accurately encoding different modalities before implementing the meta-learning framework. The early diffusion is also accessible to more modalities if the data pattern changes in the future. However, the effectiveness has not been justified theoretically, as most research does not favor deterministic prototypes but tends to use shifted prototypes.

Shifted Prototypes. It is expected that leveraging auxiliary modality provides the means to inject diversity into the generated sample space for novel classes during training [72, 73, 74]. Pahde et al. [68] aim to obtain more reliable prototypes and enrich the intrinsic feature sparsity in the few-shot training space. By training a generative model that maps text data into the pre-trained visual feature space, a new joint prototype p'_c for each novel class is redefined as the weighted average of both the original visual prototype p_c and the prototype computed from the generated visual feature vectors p_c^G conditioned on the text modality:

$$p'_c = \frac{p_c + \lambda * p_c^G}{1 + \lambda}, \quad (6)$$

$$p_c^G = \frac{1}{|S_e^c|} \sum_{(s_i, y_i) \in S_e^c} G_i(\Phi_T(t_i)), \quad (7)$$

where λ is a weighting factor, G_i is a generator and $\Phi_T(t_i)$ is the textual representation. Compared to Eq. (4), the embedding network f_θ is learned as the generator that generates new visual feature vectors. The method has verified that the combination of modalities allows the few-shot feature sparsity to be compensated and the class membership of novel samples to be better inferred.

Although the generative method has gained great success in the few-shot classification, it often assumes different modalities have the same information abundance and distinctiveness. In fact, for certain concepts, some modalities may be more discriminative than others due to the heterogeneous structure of their respective feature spaces. When the richness of different modalities does not match, the alignment is less useful.

Xing et al. [6] propose an adaptive modality mixture mechanism (AM3) that can be applied to the context of image and text, retaining the independent structure of the two modalities. The model is built on top of the prototype network to augment metric-based FSL methods by incorporating the label embeddings of all categories in both the train set and test set. The two modalities can be adaptively used according to different scenarios based on a convex combination of the two modalities in the prototype representation of each category.

Schwartz et al. [75] make a further improvement upon AM3 by adding multiple semantic explanations rather than just the category names to the convex combination of the visual and semantic prototypes. The introduction of multimodal prototypes and the pre-trained backbone models demonstrate the effectiveness of generating an adaptive embedding space for each task to incorporate multiple and richer semantics. Similarly, Zhang et al. [76] also adopt the adaptive mixture mechanism to combine

the generated features which are paired by two conditional variational autoencoders with different modality conditions. The adaptive learning approach is originated from:

$$p_c' = \lambda * p_c + (1 - \lambda) * g(t_c), \quad (8)$$

where λ is the adaptive mixture coefficient and $g(t_c)$ is the normalized version of the semantic embeddings lying on the space of the same dimension with the visual prototypes. The above papers have demonstrated that the adaptive method is flexible enough to project the possible modalities onto the same embedding space and provide a way to calculate their corresponding relationships. The adaptive method actually blurs the boundary between them by dictating the modality semantics through the weighted interaction to scaffold the importance of auxiliary modalities on the main modalities.

In addition to modeling the shifted prototypes explicitly in two modalities and then synthesizing a new joint multimodal prototype, Mu et al. [77] also attempt to train an embedded network to implicitly add the influence of shifted prototypes to the objective function. They propose language-shaped learning that encourages the original embedding function learned by the visual prototypes to decode the class language descriptions. The objective function $L_M(\theta)$ of the prototypical training is optimized by jointly minimizing the classification loss and the language model loss conditioned on the averaged prototypes:

$$L_M(\theta) = E_{S_e, Q_e} - \sum_{i=1}^J \log p_\theta(y_i | q_i, S_e) - \lambda \sum_{n=1}^N \sum_{j=1}^{J_n} \log g_\Phi(w_j | \tilde{c}_n). \quad (9)$$

Different from Eq. (5), the second item of Eq. (9) is added as the natural language loss where λ controls the weight and J_n is a set of language descriptions for each class. g_Φ is defined as a language model conditioned on \tilde{c}_n which is the averaged prototype from all the support and query samples of class n . There are no apparent representations of the prototypes of the semantic modality, but the shifted effect is reflected in the new embedding function which is learned after introducing the language model generating class semantics conditioned on the visual prototypes.

While most of the prototypical variants are explored for FSL, we cannot ignore the natural use of multimodal information in ZSL where auxiliary modality is essential to solving the problem of no access to visual information in the test phase. To take advantage of textual modalities in the task of zero-shot image classification, Hu et al. [78] do not implement a generative model to adjust the new prototype but use a correction module to update the initial prediction that maps text features to the class centers in the image feature space. The correction module implemented by prototypical networks serves as a meta-learner to adjust the initial prediction to be closer to the target value than the original prediction. After applying the correction, a shifted prototype for each unseen class in the image feature space is obtained and used to classify individual test samples to their nearest cluster centers.

The discriminative modality problem could also be exploited to retain projected prototypes in their opposing modality space.

Due to the challenge of mode collapse issues in training ZSL tasks, Yu et al. [67] propose an episode-based training framework to learn the embedding functions of visual prototype and class semantic prototype, contributing to the accumulated experience on predicting and adapting to unseen image classes. The support set in each episode is used to align semantic interactions between the visual and the class semantic modalities which are processed by two prototype generating networks. The query set is finally evaluated by minimizing multimodal loss $L_M(\theta)$ in both modality spaces V and T :

$$L_M(\theta) = E_{S_e, Q_e} - \sum_{t=1}^J \log p_\theta^V(y_t|q_t, S_e) - \sum_{t=1}^J \log p_\theta^T(y_t|q_t, S_e). \quad (10)$$

Arguably, shifted prototypes flexibly represent, select, and combine multiple modalities (usually two). It is particularly fundamental for the imbalanced semantic distribution of modalities in the low-data context. Although it is not complicated to apply Prototypical Networks, cautions are indispensable for encoding different modalities into the objective function. The choice of the weighting factor and the mapping network structure for the alignment will affect the modality knowledge that the meta-learning framework can eventually transfer. Another concern is that good prototypes are built on the assumption of the existence of a single prototypical representation for each class, which may not apply to some skewed data.

5.3. Fused-Modal Matching

5.3.1. Overview

Vinyals et al. [16] propose Matching Networks to predict the one-hot encoded labels of the query set by adding attention and memory mechanisms into one-shot training. It explicitly implements the set-to-set paradigm to learn from the given support set to minimize a loss over a batch. The predictions for each new query sample is defined by the parametric neural network $P(\hat{y}|\hat{x}, S)$, which can be formulated as a linear combination of the labels in the support set S with the help of the attentional kernel $a(\hat{x}, x_i)$:

$$P(\hat{y}|\hat{x}, S) = \sum_{i \in S} a(\hat{x}, x_i) y_i. \quad (11)$$

With results mapping from unimodal representation to the multimodal space, computed attention is modified through the normalized similarity comparison in which the query sample and each sample in the support set are supposed to incorporate multiple modalities into their embeddings.

5.3.2. Methods

The idea of implicitly representing modalities in Matching Networks inspires researchers to wrap multimodal information into the kernel function. Yao-Hung and Ruslan [79] introduce the statistical dependence measurement of fusing auxiliary modalities with the original image data, and then combine the attention mechanism of Matching Networks to learn the probability distribution of “lots-of-examples” on the labels

of “one-example” classes. They express auxiliary modalities as prior knowledge representing relationships between categories, including supervised human-crafted features and unsupervised word embeddings from language models. During the modality fusion stage, they adopt HSIC [80] to learn the kernels that can integrate different types of modalities and transfer knowledge across classes. The learned label affinity kernels $a_r(y, y'_i)$ will then be used for the prediction of label probability distribution $P(y'|y; R, R')$ on the label space for “one-example” classes:

$$P(y'|y; R, R') = \sum_{i \in S'} a_r(y, y'_i) y'_i, \quad (12)$$

where $a_r(\cdot, \cdot)$ is the attentional kernel mapping from the modalities R of “lots-of-examples” classes to the modalities R' of “one-example” classes. S stands for the support set with data-label pairs $\{x_i, y_i\}$ and S' denotes a different support set with “one-example” data $\{x'_i, y'_i\}$. The computation is similar to the label prediction for each query sample in Matching Networks (See Eq. (11)), with the difference of whether the kernels contain multimodal information. Finally, they combine the regression model and the non-parametric attention model in the loss function, and construct a meta-learning strategy [16] to achieve the adaptation from “lots-of-examples” classes to “one-example” classes.

The end-to-end architecture of Matching Networks is convenient to produce predictions for the unseen classes while new changes to the networks are bypassed. Besides, the attention kernel is inclusively adapted to multimodal information in tandem with specifying the final classifier. However, due to the weak comprehensibility of the kernel function, it is difficult to understand how different modalities work together in conveying semantics.

5.4. Related-Modal Operator

5.4.1. Overview

While the above networks focus on the embeddings that can be compared by a pre-defined metric or a linear classifier, Sung et al. [49] propose Relation Networks that learn the similarity comparison by a learnable non-linear operator and produces a similarity score between 0 and 1 within a relation module. The objective function is to minimize the sum of mean square losses L of the relation scores $r_{i,j}$ to their ground truths contained in the support set and query set:

$$L = \sum_{i \in S_e} \sum_{j \in Q_e} (r_{i,j} - \mathbb{I}(y_i == y_j))^2. \quad (13)$$

When the features are multimodal, instead of only feeding the unimodal query and the support samples to the relation module after being preprocessed by the same embedding module, the samples are first processed by heterogeneous embedding modules based on their modality pattern and then concatenated together as a whole feature map for the input of the relation module.

5.4.2. Methods

Sung et al. [49] expand the application of Relation Networks from FSL to ZSL/GZSL, and train two auxiliary parameterized networks for visual and semantic modalities on the embedding modules. They perform episode-based learning to exploit the support set and the query set in the network optimization process of meta-parameters. The relation score for each query sample is obtained by applying the relation net on the concatenated multimodal feature map. However, such Relation Networks can be considered simple binary decision-making regression models, which directly measure the difference between the ground truth and the predicted label, resulting in the limited ability to predict labels in an end-to-end framework.

To further exploit the discriminating power of the metric loss for GZSL problems, CPDN [82] formulates the relation model as a cross-modal metric learning problem. Similarly, the embedding module is still retained as the feature encoder to extract visual features. By conditioning on the textual attributes of unseen classes, a generative network of generating the visual prototype p_j for each class is constructed. Then, a new discriminative relation network is presented for learning a non-linear metric space to measure similarities between the image samples and the generated class-prototypes. The learning process adopts the standard episode-based meta-learning framework [14]. One major modification is that two novel metric losses, prototype-sample metric loss L_{PS} and class-prototype scatter loss L_{CS} are added to the original label prediction loss L_{LP} . The new loss function L_{new} [82] is shown in:

$$L_{new} = L_{LP} + \lambda L_{PS} + \beta L_{CS}, \quad (14)$$

where

$$L_{LP} = \sum_{i=1}^n \sum_{j=1}^c (r_{i,j} - \mathbb{I}(x_i, p_j))^2, \quad (15)$$

$$L_{PS} = 2 \sum_{i=1}^n \sum_{j=1}^c \left(\frac{1}{2} - \mathbb{I}(x_i, p_j) \right) r_{i,j}, \quad (16)$$

$$L_{CS} = \sum_{i=1}^n \sum_{j=1}^c (1 - \mathbb{I}(p_i, p_j)) \tilde{r}_{i,j}. \quad (17)$$

$r_{i,j}$ is the relation score of the i -th sample and the j -th class-prototype. $\tilde{r}_{i,j}$ is the relation score of the i -th and j -th prototypes.

Due to the scalability of the embedding module, another modality can be freely merged with the main modality, thereby achieving the realization of the multimodality-based Relation Networks. This also explains the fluency of migrating the network structures from FSL to ZSL/GZSL. Nonetheless, the risk is to assume that one modality has full predictive power for another. The above two methods have not realized the simultaneous modeling of two modalities that describe the same object, let alone extended to more modalities. Research gaps in the multimodality-based FSL still exist. In addition, we may infer that since the Relation Networks can be applied as a simple binary decision-maker, it is desirable to implement them as the final step to assist prediction in other models. For example, the

implemented solution [82] includes the metric loss calculated by various prototypes to the original MSE loss.

6. Learn the Generation

The distribution of the data itself can be modeled as learnable information, which is common in traditional statistical learning and ML. With the continuous emergence of neural networks, powerful networks provide more complex modeling methods to learn data distribution, especially for FSL/ZSL. The meta-learning framework also tries to enhance the generalization ability by learning more meta-knowledge about the data. The most common practice is introducing a generative neural network to learn its ability to generate data and use it in new tasks to enrich a small number of samples. We thus review relevant studies utilizing multimodal data to model flexible generation processes combined with the meta-learning framework.

6.1. Cross-Modal Augmentation

6.1.1. Overview

Data augmentation is another intuitive method for training small amounts of labeled data. The common techniques are often applied to generate extra samples in the image domain simply by a set of transformations, such as flipping, rotating, adding Gaussian perturbations and deformation [83, 84, 85]. However, these transformations only modify the geometric structure of the image and preserve some original properties like parallelism of lines or ratios of distances. Due to the lack of prior knowledge about data distributions, it is difficult to generate more complex new samples.

In order to prevent overfitting in supervised learning, it is important to increase the number of data samples. As an analogy, a meta-learning algorithm's performance and generalization abilities mainly depend on the number and diversity of tasks. However, it is not always feasible to obtain sufficient meta-training tasks. Therefore, similar to data augmentation, it is intuitive to use task augmentation to generate more meta-training tasks. Liu et al. [86] simply rotate the images to increase the number of classes. Yao et al. [87] linearly combine both features and labels of samples within the same meta-training task. Yao et al. [88] generate new tasks by interpolating features and labels of a pair of samples that are randomly sampled from different meta-training tasks. Task augmentation transforms and combines features and labels to generate more tasks, but it is not guaranteed that the generated meta-training tasks are complex and diversified due to the lack of prior task distribution.

Recently, more advanced data augmentation methods have been studied to learn new concepts from few labeled samples. Hariharan and Girshick [89] consider the intra-class variations and learn the analogical relations between paired samples. VI-Net [64] constructs a sampling space where the distribution is generated from the sharing class-specific pool of support sets and query sets. Such a variational Bayesian network can then be combined with any classifier-free mechanism to predict the novel classes. DAGAN [90] is built upon WGAN [91] to augment the support sets in the one-shot learning and sample from

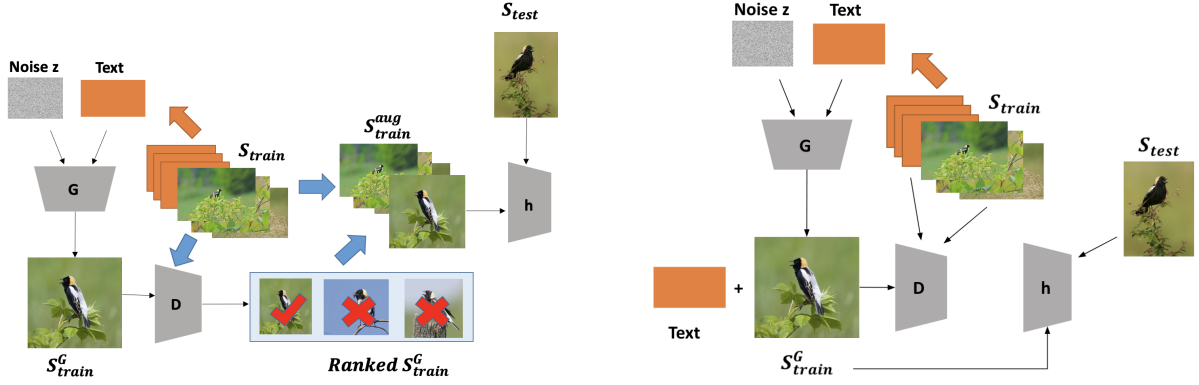


Figure 7: Network structures for cross-modal augmentation. **Left:** Adapted from [81, 74]. Meta-learning with hallucination. Given an initial training set S_{train} with two modalities, we combine a ranked S_{train}^G generated by the self-paced tcGAN with S_{train} to create an augmented training set S_{train}^{aug} . S_{train}^G is obtained by conditioning on the textual modality and noise vectors z . The generative network is trained jointly with the classifier h . **Right:** Adapted from [50]. For each task, the support set is used by the meta-learners G and D in the inner loop. The classifier h takes the generated sample from G and classifies it into the original class.

the space that is related to the best manifolds, facilitating the classification performance of matching networks. However, generic generative methods mostly transform the FSL into standard supervised learning after generating data distributions or enough samples. They usually start with base classes and then migrate the learned data pattern to generate new samples of the novel classes. One major weakness is that the generative model is trained separately from the classifier so that no explicit prior knowledge can be quickly adapted to novel classes.

The recently emerging meta-learning framework has spawned relevant work on the meta-based generative methods to learn an optimal augmentation strategy that can be trained within a few steps given unseen classes. The main idea is to wrap up the data augmentation strategy in optimization steps of the inner meta-learning process [7], with the augmentation parameterized and learned by the outer optimization in the episode training of meta-representation. MetaGAN [92] explores the impact of generating fake data on the decision boundary and provides ways to integrate the model with MAML or Relation Networks. Wang et al. [81] rely on a simple multi-layer perceptron as a hallucinator learned on the base classes to synthesize new data for unseen classes. However, the generalization ability could still be limited due to the lack of additional modalities that are helpful for the model to identify highly discriminative features across modalities [30] in the low data regime. Cross-modal data is considered an input that can effectively assist in data augmentation. By mining the relationship between the two modalities, the modality with insufficient information can be enriched.

6.1.2. Methods

Pahde et al. [30] extend the model from Wang et al. [81] in a multimodal and progressive manner by hallucinating images conditioned on fine-grained textual modality within the few-shot scenario. They propose a text-conditional GAN (tcGAN) to learn the mapping from the textual space to visual space by a self-paced based generative model [74], which ensures that only the top-ranked generated image samples are selected and

aggregated with the original novel samples. As shown in Figure 7 (Left), the hallucinator is parameterized as $G(x_i^{text}, z_i; w_G)$ and the discriminator is parameterized as $D(x_i^{image}; w_D)$, both connected to the inner classifier $h(x, S_{train}^{aug}; w_C)$ through the generated S_{train}^{aug} . The parameters w_G and w_D from the generative network and w_C from the few-shot classifier are updated directly by a joint back-propagation in the meta-training stage to augment data that are useful for discriminating classes.

The generative model requires optimized generator and discriminator parameters, which are naturally suitable for the inner-optimization of the MAML training framework. Verma et al. [50] identify the key challenge of using generative models to train ZSL/GZSL by leveraging auxiliary modalities. They propose a meta-based generative framework that integrates MAML with a generative model conditioned on class attributes for unseen samples generation. The key difference with the standard meta-learning is to mimic the ZSL behavior by setting disjoint meta-train and meta-validation partitions for each task. The meta-learning protocol modifies the standard adversarial generation process to provide an efficient discriminator and generator by enhancing their learning capability.

After the data augmentation strategy is parameterized as the inner optimizer to participate in the training, it still relies on the optimization-based meta-learning framework, except that the inner learner no longer pays attention to the learning of multimodal relations. More diversified data generators beyond conditional GANs are expected to be introduced in the future and combined with metric-based meta-learners.

7. Summary of Methods

This section provides a summary of recent works about multimodality-based meta-learning methods. The attributes of each method are concluded in Table 2. Specifically, we highlight the following attributes and discuss some common applications.

Table 2: Summary of multimodality-based meta-learning methods. “V”, “T”, “S”, “Te” and “A” indicate visual modalities, textual modalities, spatio modalities, temporal modalities and audio modalities, respectively. “M” and “F” indicate missing and full modalities, respectively.

Reference	Type	Modality	Dataset	Application
Ma et al. [45]	Optimization	V(F), T(M)/A(M)	MM-IMDb, CMU-MOSI, avMNIST	Image/multi-genre/binary sentiment classification
Yao et al. [44]	Optimization	S(F), Te(F)	City mobility/water quality datasets	Traffic volume/water quality prediction
Yan et al. [54]	Optimization	V(F), A(F)	AI2-THOR	Indoor navigation
Verma et al. [55]	Optimization	V(F), T(F)	CUB-200, aPY, AWA1, AWA2, SUN	Image classification
Liu et al. [56]	Optimization	V(F), T(F)	CUB-200, aPY, AWA1, AWA2	Image classification
Ma et al. [59]	Optimization	V(F), T(F)	USer-InstaPIC/-YFCC	Cross-modal retrieval
Chen and Zhang [46]	Optimization	V(F), T(F)	CUB-200, ModelNet40	Image classification
Liang et al. [58]	Optimization	V(F), T(F); V(F), A(F); T(F), A(F)	Yummy-28K, CIFAR-10/-100, ESC-50, Wilderness	Image/speech classification
Ge and Xiaoyang [53]	Embedding	V(F), T(F)	Wiki, MIRFLICKR, NUS-WIDE	Cross-modal retrieval
Liu and Zhang [65]	Embedding	V(F), A(F)	COIL-100	Visual learning
Eloff et al. [69]	Embedding	V(F), A(F)	MNIST, TIDigits	Cross-modal matching
Nortje and Kamper [70]	Embedding	V(F), A(F)	MNIST, TIDigits, Buckeye	Cross-modal matching
Wan et al. [71]	Embedding	V(F), T(F)	TV series, e-book pieces	Social relation extraction
Pahde et al. [68]	Embedding	V(F), T(F)	CUB-200, Oxford-102	Image classification
Xing et al. [6]	Embedding	V(F), T(F)	CUB-200, miniImageNet, tieredImageNet	Image classification
Schwartz et al. [75]	Embedding	V(F), T(F)	CUB-200, miniImageNet	Image classification
Zhang et al. [76]	Embedding	V(F), T(M)	CUB-200, miniImageNet, CIFAR-FS	Image classification
Mu et al. [77]	Embedding	V(F), T(F)	CUB-200, ShapeWorld	Image classification
Hu et al. [78]	Embedding	V(F), T(F)	CUB-200, NAB	Image classification
Yu et al. [67]	Embedding	V(F), T(F)	CUB-200, AWA1, AWA2, Oxford-102	Image classification
Yao-Hung and Ruslan [79]	Embedding	V(F), T(F)	CUB-200, AWA1	Image classification
Sung et al. [49]	Embedding	V(F), T(F)	CUB-200, miniImageNet, AWA1, AWA2, Omniglot	Image classification
Huang et al. [82]	Embedding	V(F), T(F)	CUB-200, AWA1, AWA2	Image classification
Pahde et al. [30]	Generation	V(F), T(F)	CUB-200	Image classification
Verma et al. [50]	Generation	V(F), T(F)	CUB-200, aPY, AWA1, AWA2, SUN	Image classification

- **Reference:** It denotes the individual papers that propose the method. Table 3 lists the open-source libraries for each reference if available.
- **Type:** It indicates the type of meta-knowledge learned by the method, including learning the optimization, learning the embedding, and learning the generation.
- **Modality:** It summarizes the modalities and their patterns employed in the method. Missing modalities apply to those with partial data missing, and full modalities mean that the involved modalities are not missing in the datasets. Most studies address the two most common ones, visual and textual modalities, to analyze image-related applications.
- **Dataset:** It introduces datasets that are used to evaluate the method. Many works employ CUB-200, AWA1, and AWA2 as the major datasets to derive multimodal information in the application of image classification. Table 4 lists open-source datasets used in papers if available.
- **Application:** It gives the application fields of the method, most of which are summarized based on the evaluation experiments of the paper. A large number of studies address image classification by using auxiliary textual modality. Other important areas are classification-based (e.g., sentiment, speech), prediction-based (e.g., traffic volume, water quality), and cross-modal retrieval/matching.

Furthermore, we summarize the advantages and disadvantages for each method in Table 5. The advantage that commonly appears in papers involves the flexible architecture designed for auxiliary modalities or specific tasks. The integration of feature spaces has been widely explored, though requirements for transferring models to multiple modalities or more generalized task distributions are still needed.

Table 3: List of open-source libraries for multimodality-based meta-learning methods.

Reference	Link
Ma et al. [45]	https://github.com/mengmenm/SMIL
Yao et al. [44]	https://github.com/huaxiuyao/MetaST
Liang et al. [58]	https://github.com/peter-yh-wu/xmodal
Eloff et al. [69]	https://github.com/rpeloff/multimodal_one_shot_learning
Nortje and Kamper [70]	https://github.com/LeanneNortje/DirectMultimodalFew-ShotLearning
Wan et al. [71]	https://github.com/sysulic/FL-MSRE
Xing et al. [6]	https://github.com/ElementAI/am3
Mu et al. [77]	https://github.com/jayelm/ls1
Yu et al. [67]	https://github.com/yunlongyu/EPGN
Sung et al. [49]	https://github.com/floodsung/LearningToCompare_FSL
Verma et al. [50]	https://github.com/vkverma01/meta-gzsl

7.1. Common Applications

Image classification: Image classification is one of the most common applications in the field of supervised meta-learning, where predictions are made on unknown images based on a model trained on a set of tasks. With the addition of multi-modalities in FSL/ZSL, the model can better utilize textual or audio information to understand the visual modalities, such as alignment, matching, and fusion.

Multi-genre classification: Multi-genre classification is a specific type of multi-label classification, which assumes that each sample can be assigned to more than one category. Genres are often accompanied by multiple modalities such as textual attributes, meta-attributes, movie images, etc. The introduction of meta-learning helps to better identify the effects of different modalities on different categories.

Cross-modal retrieval: Cross-modal retrieval is concerned with mutual retrieval across modalities and focuses on getting semantically similar instances in another modality (e.g., text) by searching in one modality (e.g., image). It is often necessary to find a common representation space to compare samples

Table 4: List of open-source datasets used for experimental evaluations of multimodality-based meta-learning methods.

Dataset	Link
MM-IMDb	http://lisci1.unal.edu.co/mmimdb/
CMU-MOSI	http://multicomp.cs.cmu.edu/resources/cmu-mosei-dataset/
avMNIST	https://arxiv.org/pdf/1808.07275.pdf
A12-THOR	https://github.com/allenai/ai2thor
CUB-200	http://www.vision.caltech.edu/visipedia/CUB-200.html
aPY	https://vision.cs.uiuc.edu/attributes/
AWA1	https://cvml.ist.ac.at/AwA/
AWA2	https://cvml.ist.ac.at/AwA2/
SUN	https://groups.csail.mit.edu/vision/SUN/hierarchy.html
USer-InstaPIC/-YFCC	https://ieeexplore.ieee.org/document/9465686/
ModelNet40	https://modelnet.cs.princeton.edu/
Yummly-28K	http://www.lherranz.org/datasets/
CIFAR-10/-100	https://www.cs.toronto.edu/~kriz/cifar.html
ESC-50	https://github.com/karolpiczak/ESC-50
Wilderness	https://github.com/festvox/datasets-CMU-Wilderness
Wiki	https://dl.acm.org/doi/10.1145/1873951.1873987
MIRFLICKR	https://press.liacs.nl/mirflickr/
NUS-WIDE	https://lms.comp.nus.edu.sg/wp-content/uploads/2019/research/nuswide/NUS-WIDE.html
COIL-100	https://www.cs.columbia.edu/CAVE/software/softlib/coil-100.php
MNIST	http://yann.lecun.com/exdb/mnist/
TIDigits	https://catalog.ldc.upenn.edu/LDC93S10
Buckeye	https://buckeyecorpus.osu.edu/
Oxford-102	https://www.robots.ox.ac.uk/~vgg/data/flowers/102/
miniImageNet	https://github.com/yaoyao-liu/mini-imagenet-tools
tieredImageNet	https://github.com/yaoyao-liu/tiered-imagenet-tools
ShapeWorld	https://github.com/AlexKuhnl/ShapeWorld
NAB	https://dl.allaboutbirds.org/nabirds
Omniglot	https://www.tensorflow.org/datasets/catalog/omniglot

from different modalities directly.

Speech classification: Speech classification is targeted at automatically classifying speech or audio clips. In addition to leveraging only acoustic features, cross-modal speech classification also introduces visual or textual information.

Visual learning: Visual learning is a type of learning that takes advantage of graphic information, where learners learn the material by seeing graphics, colors, and other visual information. Traditional visual learning tends to learn relevant concepts from a single pattern of data, neglecting the cross-modal learning abilities of humans. Multimodal speech-vision learning enables phonetic name features and visual object features to be combined with unsupervised meta-learning strategies for parsing visual information acquired at different times and places.

Cross-modal matching: Cross-modal matching focuses on recognizing the same objects presented in two different modalities, usually by matching a set in one modality with a matching set in another modality. For example, matching speech segments with image information, matching image information with textual information, etc.

8. Discussions and Conclusions

We have reviewed the multimodality-based meta-learning methods and classified the existing algorithms into learning the optimization, learning the embedding, and learning the generation based on the learned meta-knowledge. From the publication year of the papers reviewed in this article, it can be seen that in the past two years, meta-learning methods based on multimodal datasets have gradually attracted the community’s attention. Their main entry point is to use FSL as the research objective, and then use the characteristics of multimodality to make up for the training inefficiencies caused by the primary modalities in the few samples. At the same time, a specific meta-learning algorithm is used as the pivot and then the multimodal patterns are adjusted to adapt to the input or training process of the algorithm. As a result, the ability of meta-learning algorithms to quickly adapt to new tasks helps improve the overall performance.

The first meta-learning strategy we inspect is the optimization of model parameters (e.g., initialization, hyperparameters of the optimizer), which can help train new tasks within a few updates. Since the variant networks based on MAML or extended from MAML have achieved great success on the unimodal datasets, most researchers focus on modifying the MAML-like structure accompanied by the multimodal input and training process. We find that modifying within-task alignment mainly incorporates the parameterized multimodal model as the inner learner and still learns the initialization used for unseen tasks. A more popular way for cross-task alignment, in contrast, prefers to integrate the task-dependent parameters into the meta-learning structure and output the partially shared meta-knowledge including these parameters. Notably, the advantage of optimization-based meta-learning methods is that the variants are for the training framework without affecting how the multimodal models are encoded. True multimodalities can be guaranteed instead of being limited to common or paired modalities.

The lack of multimodal training datasets with few samples is often regarded as an obstacle to applying multimodality-based meta-learning algorithms. However, this notion is only partially correct. It can be seen from Table 2 that the application fields are extensive, and the datasets are not limited to image classification. As the applications of meta-learning in real environments are introduced with more public datasets available, the flexibility of the optimization-based framework is expected to play a more important role in adapting to new changes of modality patterns.

From the number of article distributions (See Table 2), we have seen that the metric-based meta-learning algorithms have become the preferred approach for interpreting multimodal data. In fact, metric-based algorithms provide exactly what can help multiple modalities map to the same embedding space, thus implicitly solving the problem originated from traditional multimodal ML and providing convenience to compare distances using the metric. Specifically, after reviewing relevant papers using Prototypical Networks to capture changes brought by multimodal prototypes, one would expect to be able to sum-

Table 5: Comparison of multimodality-based meta-learning methods.

Method	Reference	Advantages	Disadvantages
Parameterized-modal initialization	[45]	Efficient and flexible with the integration of feature reconstruction and training process	Sensitive to prior selection in the Bayesian framework
	[44]	Transfers sequential knowledge from multiple sources	Strong dependence on the network structure
	[54]	Adapts multimodal signals and loss functions to the RL setting	Requires the RL objectives to be identifiable and learnable
	[55]	Handles streaming data for the ZSL	Difficult to scale to multiple modalities
	[56]	Robust to class-aware task distributions	Complex learning process for synthesized features
	[59]	Generalizes well to unknown user domains without updating, scalability for larger domains	Custom weighted gradient aggregation strategy needs to be considered
	[57]	Extends MAML to heterogeneous task distributions	Dependence on the feature aggregation strategies
Unified-modal optimizer	[58]	Ensures alignment between source and target modalities, fits strongly and weakly paired data	Lacks discussions on choices of the alignment loss
	[53]	Sequentially combines multiple modalities to a unified space	Less effective without a proper updated rule
Paired-modal networks	[66], [69]	Add loss variants to help with the data deficiencies	Cost-intensive training with paired samples grows as the number of modalities increases
	[70]	Realizes cross-modal shared space	Pair mining process is time consuming
	[71]	Easy to extend to multiple modalities	Lacks theoretical assumptions to support the empirical results
Joint-modal prototypes	[68]	Leverages the auxiliary modalities to enrich the main modalities	One-side conditioned relationship between modalities makes it hard to expand to multiple modalities
	[6], [75], [76]	Adaptively combine feature across modalities rather than alignment	Limited to two modalities
	[77]	No explicit joint-modeling of prototypes is needed	Distinct properties about different modalities are ignored
	[78], [67]	Avoid the instability and expensive training of parameters of generative methods for ZSL/GZSL	Rely on the design of mapping networks
Fused-modal matching	[79]	Fast training as an end-to-end architecture	Weak comprehensibility of the kernel-based representation
Related-modal operator	[49], [82]	Compatible with multiple low-data applications by learnable non-linear operators	Rigid in combining modalities
Cross-modal augmentation	[30], [50]	Set data augmentation strategies as the optimizer	Heavily depends on the conditional generative model

marize the perfect prototype fusion method. Although different methods have proved that their performance exceeds that of unimodal prototype applications, a striking conclusion we can draw is that there is no exact solution to uniformly solve the problem of modifying multimodal prototypes in different applications. For example, in the adaptive modality mixture mechanism (AM3), many researchers [6, 75, 76] use the same prototypical networks and modality types, but the results are very different. A key aspect that is often overlooked is that the discriminative features that different modalities can provide are very different in various scenarios. Moreover, it also depends on the dataset itself. Therefore, the determination of which modality is more important is not a simple procedure. The benefits that multimodal prototypes can provide are not only to expand the possibility of more modality combinations in the embedding space, but may also be affected by the previous encoding models of multimodal data. However, these factors hardly appear in the current works.

Researchers who have achieved good performance usually distinguish aspects other than the structure of meta-learning algorithms, such as new data distributions. This philosophy is embodied in learning data generation. Although Pahde et al. [74] and Verma et al. [50] have verified the improved performance of the results brought by data augmentation strategies, there are no further discussions about whether the design of a specific meta-learning architecture containing multimodal tasks can be combined with data augmentation methods to obtain better results. Furthermore, we have noticed the potential of incorporating the meta-based ZSL/GZSL methods into our consideration. These methods divide training classes into support sets and query sets, enabling the augmentation process to learn a robust mapping when there are only auxiliary modalities or few samples. The prejudice against the seen classes during

training is thus alleviated.

In conclusion, the academic community has paid much attention to specific applications of multimodality-based meta-learning models. New frameworks are applied to multimodal datasets and compared with the baselines to illustrate their advantages. However, theoretical analyses about mathematical assumptions, generalized performance, complexity, and convergence of models remain blank.

9. Future Directions

This section provides some insights about future directions based on the current challenges and trends. We believe that the summary provided in this survey can be used as a pivot to help future research progress in specific directions, whether it is from a modality perspective or a meta-learning algorithm perspective.

Granularity of modalities. The granularity of modalities has not been studied systematically for the FSL/ZSL. Previous studies have already differentiated on the mining of semantic granularity. In some studies, only the simple word embedding or a set of several prescribed attributes [6, 68, 78] were used to encode the semantics. Nevertheless, in other studies, richer semantics could be employed as well, such as the category labels, richer ‘description level’ semantic sentences or attributes [75]. The image descriptions with fine-grained text have become a breakthrough point for those studies, which also implies that future research can explore the impact of different modalities, including the granularity’s size on the performance of the meta-learning methods.

Missing modalities. As a common data pattern, missing data is more likely to appear in multimodal scenarios in a way that is paired or separated. Considering the inherent task segmentation

in the meta-learning framework, more severe missing modalities could also appear in support sets and query sets at the same time. Existing work focuses on the direct application of the meta-learning framework, paying more attention to parameter training or modality fusion, but ignoring the fact that the lack of modalities in multiple scenes could cause unexpected effects. Although several studies [45, 76] have tried to integrate auxiliary networks to deal with the missing modalities in the whole training, a standardized framework is still to be proposed.

Multiple modalities. Current research, whether in terms of methodology or experimental objectives, mostly limits the modalities to visual and textual ones, or dual-modality that appears in pairs. Therefore it is still difficult to apply these methods to multiple modalities other than two. On the one hand, well-aligned multimodal datasets often need to be manually collected and artificially processed to be aligned with multiple modalities. On the other hand, the burden of input parameter fusion or joint training will also increase as the modalities increase. Future research needs to discover ways to construct good benchmark multimodal datasets, instead of only relying on the datasets that are originally designed for image classification. In addition, true multimodal research is supposed to be accompanied by more diverse modality distributions, especially in scenarios such as single-type tasks, cross-tasks, and heterogeneous tasks. It presents challenges to the combination of modalities under different tasks and the meta-learning framework.

More than data augmentation. The literature on learning the data is currently limited, focusing on data augmentation. Future work needs to include more data objectives that can be learned, not only use augmentation methods to learn existing data patterns, but also learn generalized data patterns from similar datasets. Apart from the data objectives, whether some data learning algorithms such as the GAN-based method [90] can be used for inner learning and optimization in outer learning is still an open challenge for the meta-learning framework.

More than classification. Most studies emphasize the use of multimodal data to improve image classification performance, which implements the visual modality as the main modality, accompanied by commonly used auxiliary modalities including text and audio. From a limited scope, this branch of method usually belongs to the category of semantic-based meta-learning, which uses the semantic features of text to enhance the performance of the few-shot classifier by using a convex combination of semantic and visual modalities. The possible future direction is to use the visual modality as an auxiliary modality to perform classification tasks such as text, video, and speech classifications. Such applications put forward high demands on the model for refining image semantics, and need to obtain unique semantics from the visual modality to match the semantic area of other modalities.

Extend to model-based methods. Previous taxonomies of meta-learning methods often include the model-based methods as synthesizing models in a feed-forward manner and learning the meta-knowledge of the single model directly. To the best of our knowledge, currently, there is no explicit work to learn such a single model that can be generalized between multi-

modal tasks. We infer that this may be limited by factors such as the complex process of parameterizing multimodalities and network structures. Future work can try to use the meta-learning models to embed the multimodal dataset into the activation state and predict the test data based on this state.

References

- [1] Y. Guo, Y. Liu, A. Oerlemans, S. Lao, S. Wu, M. S. Lew, Deep learning for visual understanding: A review, *Neurocomputing* 187 (2016) 27–48. URL: <https://doi.org/10.1016/j.neucom.2015.09.116>. doi:10.1016/j.neucom.2015.09.116.
- [2] A. Shrestha, A. Mahmood, Review of deep learning algorithms and architectures, *IEEE Access* 7 (2019) 53040–53065. URL: <https://doi.org/10.1109/ACCESS.2019.2912200>. doi:10.1109/ACCESS.2019.2912200.
- [3] T. Young, D. Hazarika, S. Poria, E. Cambria, Recent trends in deep learning based natural language processing [review article], *IEEE Comput. Intell. Mag.* 13 (2018) 55–75. URL: <https://doi.org/10.1109/MCI.2018.2840738>. doi:10.1109/MCI.2018.2840738.
- [4] Z. Song, X. Yang, Z. Xu, I. King, Graph-based semi-supervised learning: A comprehensive review, *IEEE Transactions on Neural Networks and Learning Systems* (2022) 1–21. doi:10.1109/TNNLS.2022.3155478.
- [5] S. Thrun, L. Y. Pratt, Learning to learn: Introduction and overview, in: S. Thrun, L. Y. Pratt (Eds.), *Learning to Learn*, Springer, 1998, pp. 3–17. URL: https://doi.org/10.1007/978-1-4615-5529-2_1. doi:10.1007/978-1-4615-5529-2_1.
- [6] C. Xing, N. Rostamzadeh, B. N. Oreshkin, P. O. Pinheiro, Adaptive cross-modal few-shot learning, in: H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada, 2019*, pp. 4848–4858. URL: <https://proceedings.neurips.cc/paper/2019/hash/d790c9e6c0b5e02c87b375e782ac01bc-Abstract.html>.
- [7] T. M. Hospedales, A. Antoniou, P. Micaelli, A. J. Storkey, Meta-learning in neural networks: A survey, *CoRR abs/2004.05439* (2020). URL: <https://arxiv.org/abs/2004.05439>. arXiv:2004.05439.
- [8] I. Khan, X. Zhang, M. Rehman, R. Ali, A literature survey and empirical study of meta-learning for classifier selection, *IEEE Access* 8 (2020) 10262–10281. URL: <https://doi.org/10.1109/ACCESS.2020.2964726>. doi:10.1109/ACCESS.2020.2964726.
- [9] R. Vilalta, Y. Drissi, A perspective view and survey of meta-learning, *Artif. Intell. Rev.* 18 (2002) 77–95. URL: <https://doi.org/10.1023/A:1019956318069>. doi:10.1023/A:1019956318069.
- [10] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, R. Hadsell, Meta-learning with latent embedding optimization, in: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019, OpenReview.net, 2019. URL: <https://openreview.net/forum?id=BJgklhAcK7>.
- [11] H. Sikka, Multimodal modular meta-learning (2020).
- [12] R. Vuorio, S. Sun, H. Hu, J. J. Lim, Multimodal model-agnostic meta-learning via task-aware modulation, in: H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada, 2019*, pp. 1–12. URL: <https://proceedings.neurips.cc/paper/2019/hash/e4da3b7fbbce2345d7772b0674a318d5-Abstract.html>.
- [13] Y. Li, I. King, Autograph: Automated graph neural network, in: H. Yang, K. Pasupa, A. C. Leung, J. T. Kwok, J. H. Chan, I. King (Eds.), *Neural Information Processing - 27th International Conference, ICONIP 2020, Bangkok, Thailand, November 23–27, 2020, Proceedings, Part II, volume 12533 of Lecture Notes in Computer Science*, Springer, 2020, pp. 189–201. URL: https://doi.org/10.1007/978-3-030-63833-7_16. doi:10.1007/978-3-030-63833-7_16.
- [14] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, in: D. Precup, Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*,

- Sydney, NSW, Australia, 6-11 August 2017, volume 70 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 1126–1135. URL: <http://proceedings.mlr.press/v70/finn17a.html>.
- [15] J. Snell, K. Swersky, R. S. Zemel, Prototypical networks for few-shot learning, in: I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, December 4-9, 2017, Long Beach, CA, USA, 2017, pp. 4077–4087. URL: <https://proceedings.neurips.cc/paper/2017/hash/cb8da6767461f2812ae4290eac7cbc42-Abstract.html>.
 - [16] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, D. Wierstra, Matching networks for one shot learning, in: D. D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016*, December 5-10, 2016, Barcelona, Spain, 2016, pp. 3630–3638. URL: <https://proceedings.neurips.cc/paper/2016/hash/90e1357833654983612fb05e3ec9148c-Abstract.html>.
 - [17] S. Ravi, H. Larochelle, Optimization as a model for few-shot learning, in: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, OpenReview.net, 2017. URL: <https://openreview.net/forum?id=rJY0-Kc11>.
 - [18] J. Schmidhuber, Evolutionary Principles in Self-Referential Learning. On Learning now to Learn: The Meta-Meta-Meta...Hook, Diploma thesis, Technische Universität München, Germany, 1987. URL: <http://www.idsia.ch/~juergen/diploma.html>.
 - [19] Y. Bengio, S. Bengio, J. Cloutier, Learning a synaptic learning rule, *Cite-seer*, 1990.
 - [20] S. Hochreiter, A. S. Younger, P. R. Conwell, Learning to learn using gradient descent, in: G. Dorffner, H. Bischof, K. Hornik (Eds.), *Artificial Neural Networks - ICANN 2001, International Conference Vienna, Austria, August 21-25, 2001 Proceedings*, volume 2130 of *Lecture Notes in Computer Science*, Springer, 2001, pp. 87–94. URL: https://doi.org/10.1007/3-540-44668-0_13. doi:10.1007/3-540-44668-0_13.
 - [21] A. S. Younger, S. Hochreiter, P. R. Conwell, Meta-learning with back-propagation, in: *IJCNN'01. International Joint Conference on Neural Networks. Proceedings (Cat. No. 01CH37222)*, volume 3, IEEE, 2001.
 - [22] S. Guirroy, V. Verma, C. J. Pal, Towards understanding generalization in gradient-based meta-learning, *CoRR abs/1907.07287* (2019). URL: <http://arxiv.org/abs/1907.07287>. arXiv:1907.07287.
 - [23] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, T. P. Lillicrap, Meta-learning with memory-augmented neural networks, in: M. Balcan, K. Q. Weinberger (Eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, JMLR.org, 2016, pp. 1842–1850. URL: <http://proceedings.mlr.press/v48/santoro16.html>.
 - [24] S. Khodadadeh, L. Bölöni, M. Shah, Unsupervised meta-learning for few-shot image classification, in: H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. B. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019*, December 8-14, 2019, Vancouver, BC, Canada, 2019, pp. 10132–10142. URL: <https://proceedings.neurips.cc/paper/2019/hash/fd0a5a5e367a0955d81278062ef37429-Abstract.html>.
 - [25] Y. Li, I. King, Architecture search for image inpainting, in: H. Lu, H. Tang, Z. Wang (Eds.), *Advances in Neural Networks - ISNN 2019 - 16th International Symposium on Neural Networks*, ISNN 2019, Moscow, Russia, July 10-12, 2019, Proceedings, Part I, volume 11554 of *Lecture Notes in Computer Science*, Springer, 2019, pp. 106–115. URL: https://doi.org/10.1007/978-3-030-22796-8_12. doi:10.1007/978-3-030-22796-8_12.
 - [26] B. P. Yuhas, M. H. G. Jr., T. J. Sejnowski, Integration of acoustic and visual speech signals using neural networks, *IEEE Commun. Mag.* 27 (1989) 65–71. URL: <https://doi.org/10.1109/35.41402>. doi:10.1109/35.41402.
 - [27] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, A. Y. Ng, Multimodal deep learning, in: L. Getoor, T. Scheffer (Eds.), *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, Omnipress, 2011, pp. 689–696. URL: https://icml.cc/2011/papers/399_icmlpaper.pdf.
 - [28] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, M. Rohrbach, Multimodal compact bilinear pooling for visual question answering and visual grounding, in: J. Su, X. Carreras, K. Duh (Eds.), *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, The Association for Computational Linguistics, 2016, pp. 457–468. URL: <https://doi.org/10.18653/v1/d16-1044>. doi:10.18653/v1/d16-1044.
 - [29] J. Yu, J. Li, Z. Yu, Q. Huang, Multimodal transformer with multi-view visual representation for image captioning, *IEEE Trans. Circuits Syst. Video Technol.* 30 (2020) 4467–4480. URL: <https://doi.org/10.1109/TCSVT.2019.2947482>. doi:10.1109/TCSVT.2019.2947482.
 - [30] F. Pahde, P. Jähnichen, T. Klein, M. Nabi, Cross-modal hallucination for few-shot fine-grained recognition, *CoRR abs/1806.05147* (2018). URL: <http://arxiv.org/abs/1806.05147>. arXiv:1806.05147.
 - [31] N. Bhatt, A. Thakkar, A. Ganatra, A survey and current research challenges in meta learning approaches based on dataset characteristics, *International Journal of soft computing and Engineering* 2 (2012) 234–247.
 - [32] C. Lemke, M. Budka, B. Gabrys, Metalearning: a survey of trends and technologies, *Artif. Intell. Rev.* 44 (2015) 117–130. URL: <https://doi.org/10.1007/s10462-013-9406-y>. doi:10.1007/s10462-013-9406-y.
 - [33] J. Vanschoren, Meta-learning: A survey, *CoRR abs/1810.03548* (2018). URL: <http://arxiv.org/abs/1810.03548>. arXiv:1810.03548.
 - [34] Y. Wang, Q. Yao, J. T. Kwok, L. M. Ni, Generalizing from a few examples: A survey on few-shot learning, *ACM Comput. Surv.* 53 (2020) 63:1–63:34. URL: <https://doi.org/10.1145/3386252>. doi:10.1145/3386252.
 - [35] H. Peng, A comprehensive overview and survey of recent advances in meta-learning, *CoRR abs/2004.11149* (2020). URL: <https://arxiv.org/abs/2004.11149>. arXiv:2004.11149.
 - [36] W. Yin, Meta-learning for few-shot natural language processing: A survey, *CoRR abs/2007.09604* (2020). URL: <https://arxiv.org/abs/2007.09604>. arXiv:2007.09604.
 - [37] M. Huisman, J. N. van Rijn, A. Plaat, A survey of deep meta-learning, *Artif. Intell. Rev.* 54 (2021) 4483–4541. URL: <https://doi.org/10.1007/s10462-021-10004-4>. doi:10.1007/s10462-021-10004-4.
 - [38] A. Doke, M. Gaikwad, Survey on automated machine learning (automl) and meta learning, in: 12th International Conference on Computing Communication and Networking Technologies, ICCCNT 2021, Kharagpur, India, July 6-8, 2021, IEEE, 2021, pp. 1–5. URL: <https://doi.org/10.1109/ICCCNT51525.2021.9579526>. doi:10.1109/ICCCNT51525.2021.9579526.
 - [39] D. Ramachandram, G. W. Taylor, Deep multimodal learning: A survey on recent advances and trends, *IEEE signal processing magazine* 34 (2017) 96–108.
 - [40] D. Lahat, T. Adali, C. Jutten, Multimodal data fusion: an overview of methods, challenges, and prospects, *Proceedings of the IEEE* 103 (2015) 1449–1477.
 - [41] C. Jewitt, J. Bezemer, K. O'Halloran, *Introducing multimodality*, Routledge, 2016.
 - [42] P. W. Koh, S. Sagawa, H. Marklund, S. M. Xie, M. Zhang, A. Balsubramani, W. Hu, M. Yasunaga, R. L. Phillips, I. Gao, T. Lee, E. David, I. Stavness, W. Guo, B. Earnshaw, I. Haque, S. M. Beery, J. Leskovec, A. Kundaje, E. Pierson, S. Levine, C. Finn, P. Liang, WILDS: A benchmark of in-the-wild distribution shifts, in: M. Meila, T. Zhang (Eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, PMLR, 2021, pp. 5637–5664. URL: <http://proceedings.mlr.press/v139/koh21a.html>.
 - [43] T. Baltrušaitis, C. Ahuja, L. Morency, Multimodal machine learning: A survey and taxonomy, *IEEE Trans. Pattern Anal. Mach. Intell.* 41 (2019) 423–443. URL: <https://doi.org/10.1109/TPAMI.2018.2798607>. doi:10.1109/TPAMI.2018.2798607.
 - [44] H. Yao, Y. Liu, Y. Wei, X. Tang, Z. Li, Learning from multiple cities: A meta-learning approach for spatial-temporal prediction, in: L. Liu, R. W. White, A. Mantrach, F. Silvestri, J. J. McAuley, R. Baeza-Yates, L. Zia (Eds.), *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, ACM, 2019, pp.

- 2181–2191. URL: <https://doi.org/10.1145/3308558.3313577>. doi:10.1145/3308558.3313577.
- [45] M. Ma, J. Ren, L. Zhao, S. Tulyakov, C. Wu, X. Peng, SMIL: multi-modal learning with severely missing modality, in: Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021, AAAI Press, 2021, pp. 2302–2310. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/16330>.
- [46] J. Chen, A. Zhang, Hetmaml: Task-heterogeneous model-agnostic meta-learning for few-shot learning across modalities, CoRR abs/2105.07889 (2021). URL: <https://arxiv.org/abs/2105.07889>. arXiv:2105.07889.
- [47] G. Koch, R. Zemel, R. Salakhutdinov, et al., Siamese neural networks for one-shot image recognition, in: ICML deep learning workshop, volume 2, Lille, 2015.
- [48] E. Hoffer, N. Ailon, Deep metric learning using triplet network, in: Y. Bengio, Y. LeCun (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Workshop Track Proceedings, 2015. URL: <http://arxiv.org/abs/1412.6622>.
- [49] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. S. Torr, T. M. Hospedales, Learning to compare: Relation network for few-shot learning, in: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, IEEE Computer Society, 2018, pp. 1199–1208. URL: http://openaccess.thecvf.com/content_cvpr_2018/html/Sung_Learning_to_Compare_CVPR_2018_paper.html. doi:10.1109/CVPR.2018.00131.
- [50] V. K. Verma, D. Brahma, P. Rai, Meta-learning for generalized zero-shot learning, in: The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020, AAAI Press, 2020, pp. 6062–6069. URL: <https://aaai.org/ojs/index.php/AAAI/article/view/6069>.
- [51] A. Nichol, J. Achiam, J. Schulman, On first-order meta-learning algorithms, CoRR abs/1803.02999 (2018). URL: <http://arxiv.org/abs/1803.02999>. arXiv:1803.02999.
- [52] E. Grefenstette, B. Amos, D. Yarats, P. M. Htut, A. Molchanov, F. Meier, D. Kiela, K. Cho, S. Chintala, Generalized inner loop meta-learning, arXiv preprint arXiv:1910.01727 (2019).
- [53] G. Song, X. Tan, Sequential learning for cross-modal retrieval, in: 2019 IEEE/CVF International Conference on Computer Vision Workshops, ICCV Workshops 2019, Seoul, Korea (South), October 27-28, 2019, IEEE, 2019, pp. 4531–4539. URL: <https://doi.org/10.1109/ICCVW.2019.00554>. doi:10.1109/ICCVW.2019.00554.
- [54] L. Yan, D. Liu, Y. Song, C. Yu, Multimodal aggregation approach for memory vision-voice indoor navigation with meta-learning, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020, Las Vegas, NV, USA, October 24, 2020 - January 24, 2021, IEEE, 2020, pp. 5847–5854. URL: <https://doi.org/10.1109/IROS45743.2020.9341398>. doi:10.1109/IROS45743.2020.9341398.
- [55] V. K. Verma, K. J. Liang, N. Mehta, L. Carin, Meta-learned attribute self-gating for continual generalized zero-shot learning, CoRR abs/2102.11856 (2021). URL: <https://arxiv.org/abs/2102.11856>. arXiv:2102.11856.
- [56] Z. Liu, Y. Li, L. Yao, X. Wang, G. Long, Task aligned generative meta-learning for zero-shot learning, in: Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021, AAAI Press, 2021, pp. 8723–8731. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/17057>.
- [57] J. Chen, A. Zhang, Hetmaml: Task-heterogeneous model-agnostic meta-learning for few-shot learning across modalities, in: G. Demartini, G. Zuccon, J. S. Culpepper, Z. Huang, H. Tong (Eds.), CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021, ACM, 2021, pp. 191–200. URL: <https://doi.org/10.1145/3459637.3482262>. doi:10.1145/3459637.3482262.
- [58] P. P. Liang, P. Wu, Z. Liu, L. Morency, R. Salakhutdinov, Cross-modal generalization: Learning in low resource modalities via meta-alignment, CoRR abs/2012.02813 (2020). URL: <https://arxiv.org/abs/2012.02813>. arXiv:2012.02813.
- [59] X. Ma, X. Yang, J. Gao, C. Xu, The model may fit you: User-generalized cross-modal retrieval, IEEE Transactions on Multimedia (2021).
- [60] M. Andrychowicz, M. Denil, S. G. Colmenarejo, M. W. Hoffman, D. Pfau, T. Schaul, N. de Freitas, Learning to learn by gradient descent by gradient descent, in: D. D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, R. Garnett (Eds.), Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain, 2016, pp. 3981–3989. URL: <https://proceedings.neurips.cc/paper/2016/hash/fb87582825f9d28a8d42c5e5e8b23d-Abstract.html>.
- [61] J. Kim, T. Kim, S. Kim, C. D. Yoo, Edge-labeling graph neural network for few-shot learning, in: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019, Computer Vision Foundation / IEEE, 2019, pp. 11–20. URL: http://openaccess.thecvf.com/content_CVPR_2019/html/Kim_Edge-Labeling_Graph_Neural_Network_for_Few-Shot_Learning_CVPR_2019_paper.html. doi:10.1109/CVPR.2019.00010.
- [62] V. G. Satorras, J. B. Estrach, Few-shot learning with graph neural networks, in: 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings, OpenReview.net, 2018. URL: <https://openreview.net/forum?id=BJj6qGbRW>.
- [63] Q. Luo, L. Wang, J. Lv, S. Xiang, C. Pan, Few-shot learning via feature hallucination with variational inference, in: IEEE Winter Conference on Applications of Computer Vision, WACV 2021, Waikoloa, HI, USA, January 3-8, 2021, IEEE, 2021, pp. 3962–3971. URL: <https://doi.org/10.1109/WACV48630.2021.00401>. doi:10.1109/WACV48630.2021.00401.
- [64] J. Zhang, C. Zhao, B. Ni, M. Xu, X. Yang, Variational few-shot learning, in: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019, IEEE, 2019, pp. 1685–1694. URL: <https://doi.org/10.1109/ICCV.2019.00177>. doi:10.1109/ICCV.2019.00177.
- [65] Q. Liu, Y. Zhang, Using sensory time-cue to enable unsupervised multimodal meta-learning, CoRR abs/2009.07879 (2020). URL: <https://arxiv.org/abs/2009.07879>. arXiv:2009.07879.
- [66] R. Hadsell, S. Chopra, Y. LeCun, Dimensionality reduction by learning an invariant mapping, in: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006), 17-22 June 2006, New York, NY, USA, IEEE Computer Society, 2006, pp. 1735–1742. URL: <https://doi.org/10.1109/CVPR.2006.100>. doi:10.1109/CVPR.2006.100.
- [67] Y. Yu, Z. Ji, J. Han, Z. Zhang, Episode-based prototype generating network for zero-shot learning, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, IEEE, 2020, pp. 14032–14041. URL: <https://doi.org/10.1109/CVPR42600.2020.01405>. doi:10.1109/CVPR42600.2020.01405.
- [68] F. Pahde, M. M. Puscas, T. Klein, M. Nabi, Multimodal prototypical networks for few-shot learning, in: IEEE Winter Conference on Applications of Computer Vision, WACV 2021, Waikoloa, HI, USA, January 3-8, 2021, IEEE, 2021, pp. 2643–2652. URL: <https://doi.org/10.1109/WACV48630.2021.00269>. doi:10.1109/WACV48630.2021.00269.
- [69] R. Eloff, H. A. Engelbrecht, H. Kamper, Multimodal one-shot learning of speech and images, in: IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2019, Brighton, United Kingdom, May 12-17, 2019, IEEE, 2019, pp. 8623–8627. URL: <https://doi.org/10.1109/ICASSP.2019.8683587>. doi:10.1109/ICASSP.2019.8683587.
- [70] L. Nortje, H. Kamper, Direct multimodal few-shot learning of speech and images, CoRR abs/2012.05680 (2020). URL: <https://arxiv.org/abs/2012.05680>. arXiv:2012.05680.
- [71] H. Wan, M. Zhang, J. Du, Z. Huang, Y. Yang, J. Z. Pan, FL-MSRE:

- A few-shot learning based approach to multimodal social relation extraction, in: Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021, AAAI Press, 2021, pp. 13916–13923. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/17639>.
- [72] Y. Xian, T. Lorenz, B. Schiele, Z. Akata, Feature generating networks for zero-shot learning, in: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, IEEE Computer Society, 2018, pp. 5542–5551. URL: http://openaccess.thecvf.com/content_cvpr_2018/html/Xian_Feature_Generating_Networks_CVPR_2018_paper.html. doi:10.1109/CVPR.2018.00581.
- [73] Y. Zhu, M. Elhoseiny, B. Liu, X. Peng, A. Elgammal, A generative adversarial approach for zero-shot learning from noisy texts, in: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, IEEE Computer Society, 2018, pp. 1004–1013. URL: http://openaccess.thecvf.com/content_cvpr_2018/html/Zhu_A_Generative_Adversarial_CVPR_2018_paper.html. doi:10.1109/CVPR.2018.00111.
- [74] F. Pahde, O. Ostapenko, P. Jähnichen, T. Klein, M. Nabi, Self-paced adversarial training for multimodal few-shot learning, in: IEEE Winter Conference on Applications of Computer Vision, WACV 2019, Waikoloa Village, HI, USA, January 7-11, 2019, IEEE, 2019, pp. 218–226. URL: <https://doi.org/10.1109/WACV.2019.00029>. doi:10.1109/WACV.2019.00029.
- [75] E. Schwartz, L. Karlinsky, R. S. Feris, R. Giryes, A. M. Bronstein, Baby steps towards few-shot learning with multiple semantics, CoRR abs/1906.01905 (2019). URL: <http://arxiv.org/abs/1906.01905>. arXiv:1906.01905.
- [76] Y. Zhang, S. Huang, X. Peng, D. Yang, Dizygotic conditional variational autoencoder for multi-modal and partial modality absent few-shot learning, CoRR abs/2106.14467 (2021). URL: <https://arxiv.org/abs/2106.14467>. arXiv:2106.14467.
- [77] J. Mu, P. Liang, N. D. Goodman, Shaping visual representations with language for few-shot classification, in: D. Jurafsky, J. Chai, N. Schluter, J. R. Tetreault (Eds.), Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020, Association for Computational Linguistics, 2020, pp. 4823–4830. URL: <https://doi.org/10.18653/v1/2020.acl-main.436>. doi:10.18653/v1/2020.acl-main.436.
- [78] R. L. Hu, C. Xiong, R. Socher, Correction networks: Meta-learning for zero-shot learning (2018).
- [79] Y. H. Tsai, R. Salakhutdinov, Improving one-shot learning through fusing side information, CoRR abs/1710.08347 (2017). URL: <http://arxiv.org/abs/1710.08347>. arXiv:1710.08347.
- [80] A. Gretton, O. Bousquet, A. Smola, B. Schölkopf, Measuring statistical dependence with hilbert-schmidt norms, in: International conference on algorithmic learning theory, Springer, 2005, pp. 63–77.
- [81] Y. Wang, R. B. Girshick, M. Hebert, B. Hariharan, Low-shot learning from imaginary data, in: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, IEEE Computer Society, 2018, pp. 7278–7286. URL: http://openaccess.thecvf.com/content_cvpr_2018/html/Wang_Low-Shot_Learning_From_CVPR_2018_paper.html. doi:10.1109/CVPR.2018.00760.
- [82] S. Huang, J. Lin, L. Huangfu, Class-prototype discriminative network for generalized zero-shot learning, IEEE Signal Process. Lett. 27 (2020) 301–305. URL: <https://doi.org/10.1109/LSP.2020.2968213>. doi:10.1109/LSP.2020.2968213.
- [83] Z. Chen, Y. Fu, Y. Wang, L. Ma, W. Liu, M. Hebert, Image deformation meta-networks for one-shot learning, in: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019, Computer Vision Foundation / IEEE, 2019, pp. 8680–8689. URL: http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Image_Deformation_Meta-Networks_for_One-Shot_Learning_CVPR_2019_paper.html. doi:10.1109/CVPR.2019.00888.
- [84] Z. Chen, Y. Fu, K. Chen, Y. Jiang, Image block augmentation for one-shot learning, in: The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019, AAAI Press, 2019, pp. 3379–3386. URL: <https://doi.org/10.1609/aaai.v33i01.33013379>. doi:10.1609/aaai.v33i01.33013379.
- [85] Z. Chen, Y. Fu, Y. Zhang, Y. Jiang, X. Xue, L. Sigal, Multi-level semantic feature augmentation for one-shot learning, IEEE Trans. Image Process. 28 (2019) 4594–4605. URL: <https://doi.org/10.1109/TIP.2019.2910052>. doi:10.1109/TIP.2019.2910052.
- [86] J. Liu, F. Chao, C.-M. Lin, Task augmentation by rotating for meta-learning, arXiv preprint arXiv:2003.00804 (2020).
- [87] H. Yao, L.-K. Huang, L. Zhang, Y. Wei, L. Tian, J. Zou, J. Huang, et al., Improving generalization in meta-learning via task augmentation, in: International Conference on Machine Learning, PMLR, 2021, pp. 11887–11897.
- [88] H. Yao, L. Zhang, C. Finn, Meta-learning with fewer tasks through task interpolation, arXiv preprint arXiv:2106.02695 (2021).
- [89] B. Hariharan, R. B. Girshick, Low-shot visual recognition by shrinking and hallucinating features, in: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017, IEEE Computer Society, 2017, pp. 3037–3046. URL: <https://doi.org/10.1109/ICCV.2017.328>. doi:10.1109/ICCV.2017.328.
- [90] A. Antoniou, A. J. Storkey, H. Edwards, Data augmentation generative adversarial networks, CoRR abs/1711.04340 (2017). URL: <http://arxiv.org/abs/1711.04340>. arXiv:1711.04340.
- [91] M. Arjovsky, S. Chintala, L. Bottou, Wasserstein generative adversarial networks, in: D. Precup, Y. W. Teh (Eds.), Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, volume 70 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 214–223. URL: <http://proceedings.mlr.press/v70/arjovsky17a.html>.
- [92] R. Zhang, T. Che, Z. Ghahramani, Y. Bengio, Y. Song, Metagan: An adversarial approach to few-shot learning, in: S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, 2018, pp. 2371–2380. URL: <https://proceedings.neurips.cc/paper/2018/hash/4e4e53aa080247bc31d0eb4e7aeb07a0-Abstract.html>.