

Review

On the Element-Wise Representation and Reasoning in Zero-Shot Image Recognition: A Systematic Survey

Jingcai Guo ^{1,2,*}, Zhijie Rao ¹, Zhi Chen ³, Jingren Zhou ⁴, Dacheng Tao ⁵ and Song Guo ⁶

¹ Department of Computing, Hong Kong Polytechnic University, Hong Kong SAR, China

² Department of Land Surveying and Geo-Informatics, Hong Kong Polytechnic University, Hong Kong SAR, China

³ School of Electrical Engineering and Computer Science, The University of Queensland, Brisbane QLD 4072, Australia

⁴ Alibaba Group, Hangzhou 311121, China

⁵ College of Computing and Data Science, Nanyang Technological University, Singapore 639798, Singapore

⁶ Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong SAR, China

* Correspondence: jc-jingcai.guo@polyu.edu.hk (J.G.)

How To Cite: Guo, J.; Rao, Z.; Chen, Z.; et al. On the Element-Wise Representation and Reasoning in Zero-Shot Image Recognition: A Systematic Survey. *Edge Intelligence and Systems* **2026**, *1*(1), 4.

Received: 4 April 2026

Revised: 14 June 2026

Accepted: 17 June 2026

Published: 20 July 2026

Abstract: Zero-shot image recognition (ZSIR) aims to recognize and reason in unseen domains by learning generalized knowledge from limited data in the seen domain. In resource-constrained scenarios such as edge intelligence, zero-shot learning can significantly reduce data requirements, thereby alleviating the burden of data annotation, storage and communication. The gist of ZSIR is constructing a well-aligned mapping between the input visual space and the target semantic space, which is a bottom-up paradigm inspired by the process by which humans observe the world. In recent years, ZSIR has witnessed significant progress on a broad spectrum, from theory to algorithm design, as well as widespread applications. However, to the best of our knowledge, there remains a lack of a systematic review of ZSIR from an element-wise perspective, i.e., learning fine-grained elements of data and their inferential associations. As the element-wise paradigm shows great potential for adapting to dynamic edge environments with limited memory and low latency, this paper thoroughly investigates recent advances in element-wise ZSIR and provides a sound basis for its future development. Concretely, we first integrate three basic ZSIR tasks, i.e., object recognition, compositional recognition, and foundation model-based open-world recognition, into a unified element-wise paradigm and provide a detailed taxonomy and analysis of the main approaches. Next, we summarize the benchmarks, covering technical implementations, standardized datasets, and some more details as a library. Last, we sketch out related applications, discuss vital challenges, and suggest potential future directions.

Keywords: zero-shot learning; low-resources learning; element-wise learning; edge-native AI; computer vision

1. Introduction

With the continuous advancement of deep learning techniques, automated image recognition models are gradually becoming comparable to or even better than humans. The success relies on ingenious network architectures and, more importantly, massive data along with accurate annotations. However, the deficiencies are also obvious. On the one hand, the cost of data collection can be huge. For example, despite being a representative large-scale dataset, ImageNet [1] contains 14 million images with 21,814 categories, whose coverage is still far fewer than real-world scenarios, i.e., a conservative estimate suggests that there are at least 3 million species only for the broad insect category. This enormous data burden severely hinders the application of AI in resource-



constrained environments and the development of edge-native AI. On the other hand, image samples of some categories, such as rare/extinct species, rare medical cases, and privacy-sensitive data, are usually limited or even unavailable. Moreover, the requirement for modern deep models has also become more complex, not only to know what an object is but also to describe its accurate states [2,3], e.g., the Sliced apple and an Old dog, etc. The superposition of states and objects incurs a geometrical increase in the difficulty of data collection and annotation. In recent years, more and more interests have moved from coarse- to fine-grained scenarios to endow models with the ability to discriminate microscopic characteristics that mimic the human recognition process [4,5]. For example, one can identify Zebra by recognizing an animal with black-and-white (color) and stripes (physical trait), implying that human cognition of upper concepts can be converted into the dismantling and reconstruction of base elements and further motivate the development of element-wise learning (EWL) in image recognition tasks. Concretely, the objective of EWL is to construct a bottom-up paradigm by executing element-wise representation and reasoning, i.e., a pyramidal knowledge architecture from micro to macro, attribute to object, and local to global. In particular, elements are carriers of properties [6], which describe the meta concepts of samples. The concepts or descriptors conveyed by an element vary depending on the specific tasks, like concrete components, e.g., Wings in recognizing bird species [7], or abstract descriptions, e.g., Studying in recognizing scenes [8].

Employing element-wise representation to capture discriminative elements is already customary for fine-grained classification task [4,5,9] and has also been proven effective in many other areas, such as face recognition [10], image retrieval [11], and visual question answering [12]. On top of that, element-wise reasoning is employed to assemble a complex set of elements to achieve the recognition. The very early stage began with systematic and structured combinational logic such as predicate logic, boolean logic, and so on [13], and then researchers explored the relationship between combinational elements and objects in visual recognition tasks [14,15]. Afterward, modern technologies are also designed to follow the idea of combinational reasoning, such as deformable network [16] and graph [17]. Notably, the element-wise paradigm should be distinguished from two related but narrower concepts, attribute-based learning and fine-grained learning. Attribute-based learning specifies what the descriptive primitives are, while fine-grained learning specifies which type of subtle distinctions the task targets. The element-wise perspective, in contrast, operates at the level of reasoning mechanism, which asks how visual concepts are decomposed into reusable primitives and recomposed to recognize novel cases, regardless of whether those primitives are explicit attributes, fine-grained annotations, or implicit vision-language concepts.

Zero-shot image recognition (ZSIR), as a well-defined visual-to-semantics mapping problem adapted for unseen domains, has attracted increasing attention in most recent years and made profound progress [18]. ZSIR employs texts/words as the medium to bridge the image visual disparity between seen and unseen categories, thus generalizing the models to recognize unseen domains. Simply put, the visual features are compressed into the shared semantics, and the model unifies the shared semantics and visual features of the seen categories to reason images from unseen categories. Compared to other low-shot schemes such as few-shot learning [19] and one-shot learning [20], ZSIR reduces the data requirement to zero, thus completely releasing the data collection and annotation pressure. Moreover, compared to another similar open-world problem, i.e., out-of-distribution (OOD) detection [21], ZSIR is designed to assign precise labels for unseen data instead of simply categorizing them as OOD samples, which is usually more favorable in real-world applications. Particularly in complex, dynamic and resource-constrained edge environments, ZSIR significantly reduces the costs of data collection and annotation, thereby lowering labour, storage and communication requirements. More importantly, as tasks handled by edge-native AI are often dynamic and unknown, ZSIR provides the technical support required for rapid re-adaptation.

In practice, ZSIR and EWL are highly compatible in the recognition process that well-mimics the way humans observe the world, i.e., learning the basics through fine-grained visual-to-semantics interaction to identify concepts. Moreover, ZSIR can also be inherently considered as a natural extension of EWL towards the open-environment scenarios, i.e., infer unseen categories with novel combinations of learned base elements. Therefore, their integration can provide a powerful and promising solution for real open-environment deployment. Particularly in edge computing scenarios, the element-wise paradigm not only reduces data costs and storage requirements but also offers the potential to minimise memory usage and computational load. Specifically, the element-wise paradigm breaks down visual concepts into reusable basic units. By activating only certain units or pruning redundant ones, it is possible to achieve improvements in both computational and memory efficiency. However, existing studies on ZSIR and EWL are usually fragmented and independent of each other, lacking a systematic investigation that summarizes their integration with a unified paradigm.

Although existing reviews on zero-shot learning have laid an important foundation, their scope and perspectives differ from those of this study. Early reviews [22–24] provided foundational overviews of classical zero-shot object recognition (ZSOR), covering class-level semantic space design, generative methods, and regularisation techniques. However, these reviews treated ZSOR in isolation and did not delve into element-wise

techniques for learning and reasoning regarding fine-grained visual semantic primitives. Although recent research has addressed compositional recognition or visual-language models separately, there has been no systematic integration of mainstream tasks under a unified paradigm, including zero-shot object recognition (ZSOR), zero-shot compositional recognition (ZSCR), and foundation model-based open-world recognition (FBOR). In this survey, we innovatively integrate these three mainstream tasks, offering a new perspective through element-wise learning. Specifically, we analyse the commonalities among these three tasks in terms of element-wise learning, including conceptual paradigms and shared technical primitives (such as cross-attention, prototype learning, and graph modeling). Meanwhile, we demonstrate how the principles of element-wise learning can be adapted to meet different task objectives. Overall, this survey aims to present a comprehensive overview of the latest developments in zero-shot image recognition within a unified element-wise framework.

In summary, our contributions are four-fold:

- We, for the first time, integrate three mainstream ZSIR tasks, i.e., object recognition, compositional recognition, and foundation model-based open-world recognition, to provide a comprehensive review from an element-wise perspective.
- We discuss the differences and synergies of different ZSIR tasks and analyze the key challenges. We summarize related cutting-edge techniques and provide a new taxonomy of element-wise ZSIR.
- We collect and summarize some benchmarks, including implementations, datasets, and some more details, as a library for follow-up research.
- We present the widespread applications, discuss vital challenges and suggest potential future directions.

2. Overview

2.1. Mainstream Tasks

Zero-Shot Object Recognition: Classic zero-shot learning (ZSL) is the most common setting for ZSIR [18], addressing the challenge of classifying images that come from novel unseen categories without observing any prior visual samples before [18]. In ZSL, the annotations such as attributes, words/sentences, or semantic embeddings (i.e., we use semantics uniformly later) are shared to describe categories at a high level. During training, only image samples of seen categories are available, while the model is trained to recognize unseen categories. Hence, the shared semantics serve as the crucial link between seen and unseen categories. To fully adapt to real-world scenarios, generalized ZSL (GZSL) [22] further extends the above setting to a more realistic one by requiring both seen and unseen categories to be jointly recognized during inference. Most recent studies encompass both ZSL and GZSL settings and are collectively referred to as zero-shot object recognition (ZSOR) in this survey (i.e., top-left of Figure 1).

Zero-Shot Compositional Recognition: Another critical capability in the recent ZSL research is to describe the relevant states of sample, such as recognizing whether an apple is sliced or raw instead of simply classifying it as an object. Such capability is essential for developing models with a deeper understanding of the world and is also a key component in building human-like intelligence [25]. Inspired by it, compositional learning has emerged in ZSL to explore combinational reasoning capabilities [2,3,26]. Concretely, the model is trained from image samples and their corresponding compositional states and objects, and is expected to recognize both seen and novel unseen compositions during inference. In particular, each state and object appears at least once during training. The classic setting assumes that the labels' prior is available, i.e., some state-object compositions are filtered out, speeding up the inference. In contrast, the open-world setting [27] extends compositional ZSL by expanding the inference scope to the full set, i.e., every state-object composition counts. Therefore, many irrational combinations can occur, and the model is expected to cull out on its own. We collectively refer to them as zero-shot compositional recognition (ZSCR) in this survey (i.e., top-right of Figure 1).

Foundation Model-Based Open-World Recognition: Both ZSOR and ZSCR fall short of achieving complete openness when faced with the open-world scenario, as their objectives and predictions remain fixed and well-defined. A critical question arises: can models learn broad and foundational knowledge and adapt seamlessly to unknown tasks? The advent of large vision-language pre-trained models (VLMs) [28,29] offers a promising solution. These models focus on broad vision-language interactions, learning vast amounts of elemental knowledge without a specific goal, i.e., much like humans learn about the world through observation. Remarkably, they exhibit excellent zero-shot generalization ability, performing well on various unseen domain image recognition tasks without further training. This capability has found wide applications, particularly in low-shot scenarios [30,31]. However, recent studies [32,33] suggest that the zero-shot potential of VLMs is not fully realized. In response, foundation model-based open-world recognition (FBOR) has been developed to push such a limit in foundation models (i.e., bottom of Figure 1).




(a) Zero-Shot Object Recognition**(b) Zero-Shot Compositional Recognition****(c) Foundation model-Based Open-world Recognition**

Figure 1. Three tasks of ZSIR: **(a)** ZSOR utilizes shared attributes/texts to identify unseen categories; **(b)** ZSCR infers unseen compositions by learning from seen objects and states; and **(c)** FBOR exploits the broad fundamentals learned by pre-trained VLMs to implement zero-shot recognition directly in downstream tasks.

2.2. Comparison

While ZSOR, ZSCR, and FBOR target distinct application scenarios, they are fundamentally unified by the principle of element-wise learning. We provide a systematic comparison of their commonalities and differences, revealing the common underlying paradigm and task-specific implement ways.

Commonalities. All three tasks share the common principle of disentanglement and recombination of basic elements. The deep theoretical connection is adopting a bottom-up cognitive paradigm that captures the underlying causal factors of visual concepts. Instead of treating a category (e.g., zebra) or a composition (e.g., old car) as a monolithic, one-hot entity, element-wise methods force the model to decompose the visual signal into a set of fundamental, reusable elements (e.g., stripes, hooves, old, metallic surface). Recognition of novel concepts is then achieved by recomposing these learned primitives. In terms of the specific definition of task objective, the process of deriving higher-order semantics from images is broken down into distinct stages: from images to basic concepts, and finally to higher-order semantics. This decomposition better aligns with the way humans perceive and understand the world, and also enhances the interpretability of the model's reasoning in the unknown environments. Consequently, many techniques are transferable and mutually inspiring across the three tasks.

Differences. Driven by the shared element-wise paradigm, the three tasks explore pathways of model cognition for new concepts from different perspectives. Their differences lie primarily in the definition of elements

and the objectives of the tasks. To provide a clearer analysis of the distinctions between the three tasks, we have listed more details in Table 1, including supervision granularity, reasoning approach, and so on.

Element Definition & Supervision Granularity & Element-Wise Learning. Basic elements serve as a vital bridge between link samples and higher-order semantics. How these elements are defined lies at the heart of the differences between tasks. In ZSOR, elements are typically attributes signals or text descriptions which are semantically rich (e.g., has stripes, has black eyes, has green wings). This often necessitates explicit, fine-grained visual-to-attribute alignment to ground each concept in a visual region. Since element labels are class-level, visual grounding presents a significant challenge. In ZSCR, elements are atomic labels of states and objects. The focus is on robustly decoupling their visual features and modeling their conditional dependency to handle context-variant appearances. Their labels are instance-level, but the same labels show significant variation across different instances. In FBOR, elements are broad, implicit visual-language concepts learned from large-scale data without task-specific annotations. The element-wise reasoning is thus implicitly handled by the powerful cross-modal attention within VLMs. Its element-level learning approach is more geared towards aligning with downstream tasks and general knowledge bases.

Reasoning Approach & Core Objective. Due to differences in element definitions and application scenarios, their reasoning methods and task objectives also vary. ZSOR typically uses a linear, additive composition (e.g., a predicted class score is the sum of confidence scores for its constituent attributes). The objective is encouraging the model to recognize a new class with sharing attributes. ZSCR often demands more complex non-linear or conditional reasoning to model interactions where a state's appearance (e.g., rusty) critically depends on the object (e.g., car). Its goal is generalizing to new combinations of known states and objects. FBOR represents the most flexible reasoning, operating through the non-linear processes of large-scale attention layers that can handle open-world associations from pure text-visual interactions, without a predefined set of rules or combinations. Techniques here diverge from training explicit alignments to the more flexible manipulation of knowledge retrieval through prompt engineering or test-time adaptation.

Table 1. The differences between the three tasks.

	ZSOR	ZSCR	FBOR
Element Definition	Attributes signals. Text descriptions.	Atomic states and objects.	Visual-language concepts.
Supervision Granularity	Coarse-grained (Class-level). An attribute set for all images of a class.	Medium-grained (Instance-level). A state label and an object label per image.	No supervision. Transferred knowledge from large-scale pre-training. Zero-shot is used for direct inference.
Element-Wise Learning	Visual grounding of elements, e.g., align fine-grained visual regions with semantic attributes.	Disentangle state and object visual features or model their inter-dependencies.	Align downstream visual concepts with the model's pre-learned knowledge base (e.g., through prompts).
Reasoning Approach	Linear/Additive Composition. Class score is often the sum of its present attributes.	Conditional/Non-linear Composition. Depends on the modeling way between states and objects.	Implicit Reasoning. It is essentially an interaction between images and text.
Core Objective	Generalize to new classes sharing known attributes.	Generalize to new combinations of known states and objects.	Generalize to unseen tasks and domains with large-scale knowledge base.

2.3. Challenge

Fine-Grained Visual Analysis: Element-wise ZSIR requires detailed visual anatomy to echo the semantics and capture underlying meta concepts. ZSIR has fine-grained semantic information to describe the element-wise characteristics. However, due to huge costs, images in hand typically lack enough element-wise annotations. Therefore, semantics and visual local clues are usually learned in a weakly supervised manner, and their matching accuracy cannot be fully guaranteed. Worse still, irrelevant clues, such as the backgrounds, can inevitably interfere with the learning quality. To address these issues, recent studies incorporated fine-grained visual analysis into ZSIR and have made certain progress, particularly in accuracy improvement. Studies of such kind mainly involve visual localization [34,35], visual representation [36,37], and visual-semantics correlation modeling [38,39]. However, based on our observation, three challenges still hinder the fine-grained visual analysis in ZSIR, including (i) Elements Complexity, i.e., the number of element-wise annotations tends to be larger compared to category-wise annotations, and the corresponding visual features are much more subtle [7,40]. In addition, the processing

of abstract attributes is also a challenge due to the difficulty of finding the corresponding explicit visual clues; (ii) Visual Occlusion and Overlap, i.e., visual occlusion occurs between objects and backgrounds or between elements and elements, which may result in the absence of critical visual clues. Besides, visual overlap refers to multiple semantic expressiveness of the same visual clues [41]. For example, the model can hardly completely decouple the visual clues of red and car when recognizing a red car; and (iii) Visual Interaction, i.e., the independent semantics of visual clues can be vague or even ambiguous. For example, if a piece of skin is cut off from a cow's trunk, it is hard to recognize which part or even which species the skin comes from. Thus, the semantic meaning of visual clues can also depend on the whole to varying degrees.

Domain Shift: Domain shift is a long-standing challenge in ZSIR, which refers to the distribution bias (i.e., towards seen categories) of visual features when mapped into semantic space [42]. The inducement of domain shift is the absence of visual prior in unseen categories, which makes the models favor too much on seen categories and thus learn biased mapping functions [43]. It further misleads the model in recognizing the embeddings of samples of unseen categories as somewhat seen categories that hinder the model's generalization ability. Meanwhile, some other factors can also aggravate the domain shift problem. First, different imaging mechanisms cause variational views, lightings, backgrounds, etc. [44], and thus, the visual characteristics can be diverse even for samples of the same category. Domain adaptation and generalization [44–46] partially mitigates this issue in general recognition tasks, while ZSIR is much more challenging since domain shift and unknown prediction need to be accommodated at the same time. Second and more troublesomely, context dependency, i.e., objects or states expressing the same semantic information may present very different visual characterizations in different contexts [47,48]. For example, the object tail appears hairy and straight on a horse but smooth and curly on a pig, and the state old presents itself as the metallic luster of rust on a car, while as the wrinkled skin on a person.

2.4. Organization

To sum up, Figure 2 outlines the overall organization as well as the taxonomy of this survey. Concretely, Section 3 first covers ZSOR techniques, including problem formulation and key algorithms. Similarly, Section 4 reviews ZSCR, and Section 5 introduces FBOR. Then, Section 6 details technical implementations, standardized datasets, and shared techniques. Section 7 explores widespread applications, and Section 8 suggests potential future directions. As many papers employ multiple methods, we have adopted the following primary classification rules: for each method, we first identify the technical component that most directly embodies the element-wise philosophy. Where a method integrates multiple such components (for example, incorporating both attention mechanisms and graph modeling), we classify it according to the contribution claimed in the paper itself, placing it in the category that best reflects its most distinctive innovation relative to similar methods.




Figure 2. An overview of the organization and taxonomy of this survey.

3. Zero-Shot Object Recognition

3.1. Problem Formulation

Given the seen domain $D^s = \{(x^s, y^s, a^s) | x^s \in X^s, y^s \in Y^s, a^s \in A^s\}$, where X^s, Y^s and A^s denote visual samples, category labels, and semantics (e.g., a set of attributes), and similarly, let $D^u = \{(x^u, y^u, a^u) | x^u \in X^u, y^u \in Y^u, a^u \in A^u\}$ denote the unseen domain. Notably, $a^s \cup a^u \in \mathbf{R}^{m \times k}$ are class-level semantics annotations, where m denotes the number of categories and k is the number of attributes. Without loss of generality, the task of ZSOR can be modeled as learning a mapping/relational function $\Psi: X \rightarrow A$, wherein X^u is strictly

inaccessible for training. During inference, the learned function Ψ is applied to recognize samples from the unseen domain only, i.e., ZSL, or from both seen and unseen domains, i.e., Generalized ZSL (GZSL).

The success of ZSOR relies on the shared attribute between A^s and A^u , which act as the bridge between seen and unseen domains. The intuitive idea is to learn the visual representations corresponding to attributes, which is the notion of element-wise learning. However, due to the single target and relatively closed environment of object recognition, learning for Ψ can be transformed into a plain class-wise distribution matching problem. Class-wise distribution matching concerns neither the discrepancies of the bottom attributes nor the variations of fine-grained vision. It treats both visual and attribute features as a whole, and the objective is to reduce the distributional distance between the two to learn an unbiased mapping function. Mainstream techniques include semantic space design, generative approach, and regularization. Semantic space design is to search for a suitable low-dimensional space that perfectly matches visual and attribute embeddings [43,49]. The generative approach is to simulate the visual features of unseen classes [50,51] by training a generator that produces visual features based on attributes with the help of networks such as GANs [52–54], VAEs [55,56] and generative flows [57,58]. Regularization is imposing constraints to improve the mapping performance or robustness of the model. For example, Liu et al. [59] use an information entropy loss to reduce uncertainty. Despite the impressive results achieved by these methods, they provide limited insights for micro-operations and bottom-up analyses and are, therefore, out of the scope of this survey. The readers can refer to other related surveys [23,24] for a deeper understanding.

Depending on the concrete elements the method operates on, we categorize the element-wise techniques into three broad categories, i.e., Visual Component Analysis, Visual-Attribute Matching, and Attribute Component Analysis. Then, the methods are further categorized into different sub-categories based on the specific techniques they use.

3.2. Visual Component Analysis

Visual component analysis is the refined dissection of an image or its features, involving localization, modeling, and so on. Images often contain a large amount of redundant information, only a few of which are decisive for discrimination. At the same time, images lack fine-grained annotations. Therefore, the goal of visual component analysis is to reduce the interference of irrelevant information as well as noise, and prompt the model to focus on those critical visual parts. Commonly used methods can be classified into the following three sub-categories: Local Visual Attention, Data Manipulation, and Graph Modeling.

Local Visual Attention: Attention mechanisms are effective to improve efficiency and accuracy in many visual tasks by mimicking human focuses. Motivated by this, local visual attention is to utilize the attention mechanism to identify the most valuable visual regions and reduce the influence of background or irrelevant regions. Specifically, local visual attention generates multiple masks for visual features. Each mask is a normalized weight matrix that reveals a significant region. Enhanced features can be obtained by weighting the original features using mask, thus improving the discriminability (Figure 3). Suppose that $x \in \mathbf{R}^{C \times H \times W}$ denotes the visual feature, where C , H , and W are the channel, height, and weight, respectively. Let $M = F(x)$ represents the attention mask, where F denotes a small learnable network, e.g., 1×1 convolutional network, and $M \in \mathbf{R}^{N \times H \times W}$ denotes N attention masks. Afterward, multiplying the masks with the original features yields N regional features, which can be expressed as:

$$x_{region} = \{xm_1, xm_2, \dots, xm_N\}, \quad (1)$$

where $x_{region} \in \mathbf{R}^{N \times C \times H \times W}$ and $[m_1, m_2, \dots, m_N] = M, m_i \in \mathbf{R}^{H \times W}$. How to use these regional features to improve model's performance. AREN [35] considers certain mask learned regions to be unimportant and therefore adopts an adaptive threshold to further filter out unimportant information, i.e., region features smaller than the threshold are set to 0. RSAN [60] expects the regions learned by each mask to correspond to the attributes one by one. It uses max-pooling to get the representative value in each region feature and concatenates them into a one-dimensional vector, after which it uses Mean Squared Error Loss to minimize the distance between this vector and the attribute embedding. Considering the existence of topological relationships between regions, RGEN [36] introduces graph modeling to maintain their potential connections.

Data Manipulation: Data manipulation has a same objective as local visual attention, i.e., to search and localize the most critical regions. The difference is that the latter works at the feature space, whereas the former works at the image level. An intuitive approach is employing detection techniques to recognize parts of an object, such as a bird's head, body, and claws. The detection network can provide coordinates for multiple parts simultaneously, allowing image patches to be cropped accordingly or deep features to be localized. This method introduces fine-grained semantic information that enhances the model's classification ability [61,62]. Additionally, it boosts the performance of generative methods, which often struggle with capturing fine-grained visual differences.

For example, AGAA [63] utilizes a detector to extract region information that is fed into a discriminator, along with class embeddings generated by a generator. However, this technique is limited in practice since the detector's performance heavily depends on part-level visual annotations, which are often impractical to obtain.




Figure 3. A schematic diagram of local visual attention. Visual features are passed through a sub-network to generate multiple masks, which are then multiplied to obtain enhanced features.

To alleviate this problem, a series of unsupervised detection methods have been proposed. LDF [64] employs a coordinate generating network that generates coordinates directly from image features. Specifically, suppose $x \in \mathbf{R}^{C \times H \times W}$ is the visual feature, and the coordinates can be expressed as:

$$Z = [z_h, z_w, z_l] = F(x), \quad (2)$$

where F is a learnable network. Z is the window, z_h, z_w denote the coordinates, and z_l denotes the length of the region. After that, the regions located at the coordinates are zoomed in and re-fed to the classification network. SGMA [34] incorporates local visual attention techniques to coordinates generation. Furthermore, SR2E [65] and ERPCNet [66] introduce the reinforcement learning paradigm, which induce the cropping behavior of the model by setting up special rewards and penalties to obtain the most discriminative sub-image. POPRNet [67] takes an unsupervised detection technique. It generates multiple proposal regions for an image and then sorts and selects Top-K region features based on category confidence. ZSLViT [68] utilizes the patch split mechanism of ViT [69] to get multiple region features and later filters the key visual regions by calculating the attention weights between the regional and class tokens.

Unlike the above approaches, to capture the key visual semantic embeddings, VGSE-SMO [70] takes a patch clustering approach. It utilizes an unsupervised segmentation algorithm to decompose images into multiple patches, after which individual patches are linked to class labels via a classification network. By doing so, the model can implicitly learn the relationship between classes and fine-grained visual patches that align better with each other. To alleviate semantic redundancy, Chen et al. [71] propose identifying and processing semantically irrelevant visual blocks during the input stage, thereby significantly enhancing the semantic richness of visual representations.

Graph Modeling: Visual region localization helps models focus on important local features but often overlooks the potential relationships between regions. For example, a bird's body parts are biologically connected, and treating them independently can cause semantic ambiguity due to a narrowed field of view. Sharing information between regions can mitigate this by preserving essential semantic context. Graph modeling techniques are well suited for handling such irregular non-geometric data with a structured message propagation mechanism [17]. In recent years, with the birth and development of graph convolutional networks (GCNs) [72], graph modeling techniques have been widely used in various computer vision tasks.

Regional visual features are natural graph nodes and thus can be well embedded in graph modeling techniques. For example, GKU [37] constructs a graph for each sample to capture fine-grained local structural relationships. They redefine the visual-semantic mapping problem as a graph-semantic mapping problem, thus suppressing the influence of partial visual domain shift. However, GKU requires the dataset to be supported by fine-grained regional

annotations. Without the supervision of manual annotations, RIAE [73] incorporates all regions into the graph node set. To mine the relationships between regions, it introduces cosine similarity to measure the weight of an edge between any two region nodes. GNDAN [74], on the other hand, first locates the visual regions corresponding to the attributes with cross attention techniques, and then later uses the graph for modeling.

3.3. Visual-Attribute Matching

The goal of visual-attribute matching is to align visual features with their corresponding attributes by learning fine-grained, element-wise representations through the pairing of visual-attribute embeddings. Unlike visual component analysis, this approach explicitly focuses on interacting with attributes, allowing for a better perception of subtle attribute variations and improved generalization to unseen domains. This text-driven visual learning aligns more closely with human learning patterns and presents a greater challenge. Current methods in this area can be categorized into two sub-categories: Cross Attention and Prototype Learning.

Cross Attention: Cross attention allows for the simultaneous processing of different sources of information and efficient modeling of their dependencies. Among them, dot product attention, due to its high computational efficiency and low space cost, is widely applied in many advanced models such as Transformer [75]. The core of dot product attention is to compute the similarity between query and key-value pairs to assign attention weights, and the similarity measure is the dot product of vectors. Given a query series Q , a key series K and a value series V , the attention score can be formulated as:

$$\text{Attention}(Q, K, V) = \text{soft max}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (3)$$

where d_k denotes the dimension of the key vector used to scale the dot product result. When Q , K , and V all come from the same input series, the attention score learns the dependencies between the parts of the series, i.e., self-attention. When there are two different input series involved in the computation, the attention score learns the dependencies between the different series, i.e., cross attention.

To uncover the interactions between attributes and visual regions, the cross attention technique is naturally introduced into ZSOR (Figure 4). For example, DAZLE [38] proposes a dense attention technique that assigns weights to each region based on attributes. Specifically, it utilizes the features output from the last layer of the convolutional network and treats each point on the feature as a region. The similarity between attributes and regions is then computed with cross attention. Suppose $x \in \mathbf{R}^{C \times r}$ denotes the visual feature with $r = H \times W$ regions. Let $a \in \mathbf{R}^{d \times m}$ denote the attribute vector, where m is the number of attributes, and d is the vector dimension. The cross attention score is calculated as:

$$\text{Attention}(a, x) = \frac{\exp(a^T w x)}{\sum_r \exp(a^T w x)} x, \quad (4)$$

where $w \in \mathbf{R}^{d \times C}$ is a learnable matrix that ensures the visual and attribute vectors are in the same dimension space. The weighted visual features reflect attribute preferences. After that, the attention scores of all the attributes are used for classification. GEM [76] employs the same approach, however it retains the maximum attention scores with max-pooling during classification. Composer [77] synthesizes the features of unseen classes with critical regions identified by attention mechanism, while AREES [78] exploits the regional features to improve the performance.

Some studies resort to ViT [69] to model dependencies between vision and attributes [79–81]. The essential cross attention mechanism is identical. The difference is that ViT divides the image into patches, each of which is considered a visual region. In addition, bidirectional attention is explored in depth. It is important to know not only which region the attribute is attending to, but also which attribute the region is attending to. Based on this idea, several methods are developed. For example, MSDN [82] proposes a mutual distillation network to align attribute-region attention. MSDN++ [83] further introduces causal attention to learn the true causal relationships between visual regions and features. Ge et al. [47] also integrate bidirectional attention techniques to improve the visual-attribute matching ability of the model. To promote deep integration of visual and attribute features, Jiang et al. [84] introduce strong and weak fusion techniques across different layers of the network to enhance the discriminability of visual features. Hou et al. [85] design a lightweight visual representation network with the long-range modeling mechanism of Mamba. Rao et al. [86] analyze the imbalance issue in attribute attention and propose a progressive attention scheme to mitigate attention divergence for minor attributes.

Prototype Learning: The gist of prototype learning is to perform tasks like classification and regression using a set of representative samples, known as prototypes [87] to regularize the learning process. A prototype is a standardized example that captures the essence of a category or type, such as the mean vector of a set of vectors or the center of mass of a distribution. Prototype learning mimics human analogical reasoning by drawing

conclusions based on typical examples. Its main advantage is its robustness in handling complex data distributions, particularly when there is overlap or imbalance between data categories. As a result, prototype learning is considered to be more effective for tasks like out-of-distribution detection [88] and low-shot learning [89].




Figure 4. A schematic diagram of cross attention. Attribute embeddings and region features compute similarity to obtain multiple attention maps.

For visual-attribute matching, there is semantic overlap in the visual region. The same region may contribute to multiple attributes. Although cross attention is able to measure the contribution value by similarity, the lack of an explicit means of visual decoupling may lead to pseudo-correlation of visual- attributes. Therefore, learning a separate visual prototype for each attribute is a promising solution. When the prototype is acquired, the visual region can decide which attribute it belongs to by calculating the distance between the region vector and the prototype. It can be observed that cross attention discovers pairing relations through the similarity between the attribute and visual embeddings, while prototype learning discovers visual-attribute pairing relations indirectly through the similarity between the embeddings of samples and prototypes.

APN [90] first introduces prototype learning to visual-attribute matching. It initializes a learnable prototype vector for each attribute and then calculates the similarity between the visual region and the prototype. After that, the largest region vector is taken as the attribute vector, and then regresses with the attribute embedding. In addition, it incorporates a decorrelation loss [91] to ensure the independence between individual attributes. DPPN [92] considers that the prototype possesses a seen-domain preference, which cannot capture the delicate instance-level attribute shift. Therefore, it is necessary to equip each instance with exclusive attribute prototypes. It proposes a progressive attribute localization method, which dynamically adjusts the prototype by locating attributes through the prototype while updating the prototype in turn. Yue et al. [93] propose a graph cut method to model similarities between regions, grouping similar regions into the same cluster to learn compact attribute prototype representations. In addition, some other methods [47,94–97] also achieve favorable results by using prototype learning.

3.4. Attribute Component Analysis

Attributes serve as the crucial link between seen/unseen classes and provide essential supervisory signals for learning fine-grained visual representations. Consequently, the quality of the attribute space largely determines the upper limit of a recognition system's performance. However, most attributes are manually created or pre-trained without proper adjustment, resulting in redundancy, noise, and unstructured relationships. To address this, attribute component analysis aims to optimize the attribute space through methods such as Attribute Enhancement and Refinement.

Attribute Enhancement: Attribute enhancement aims to optimize attribute embeddings to improve their discriminability, using techniques like key attribute localization and attribute space optimization. Key attribute localization functions similarly to visual attention, identifying and focusing on the most critical parts. The cross attention mechanism mentioned in Equation (4) is one of the attribute localization techniques, just changing the

input sequences. In addition, LFGAA [98] designs a latent guided attention module. It assumes that each image has different attention for attributes and, therefore, introduces multiple sub-networks to generate exclusive attribute weights for each image. Attribute space optimization adjusts pre-trained attribute embeddings to reduce semantic overlap that affects discrimination, especially in fine-grained datasets. For instance, many bird species share similar attributes, which makes distinguishing them challenging. For this reason, AREES [78] delineates the attribute embeddings with similar semantic distances into a cluster by K-means clustering algorithm, and the attributes in each cluster are subtracted from the cluster's center of mass, i.e., the mean vector, to obtain the enhanced attribute vector. APNet [99] introduces graph modeling to discover the logical relationships between classes and updates the attributes by propagating the attribute features of neighboring nodes.

Attribute Refinement: Since attributes are typically annotated by humans, the value they add from a human perspective to deep learning is uncertain. Attribute refinement seeks to address this by examining the design of attribute annotations, though few studies exist on the topic. MCZSL [100] points out that the quality of manual annotations is uncontrollable and time-consuming. Therefore, it attempts to design an automated attribute annotation scheme. Specifically, it obtains attribute descriptions from online media such as Wikipedia and designs a series of filtering strategies to filter out inappropriate attributes. DEDN [101] proposes a divide-and-conquer attribute learning scheme. It argues that attribute annotations are confusing and lead to semantic asymmetry problems. For example, for the attributes describing scenes, {studying, swimming} and {stone, tree} are not in the same dimension, with the former denoting function and the latter denoting object. Therefore, it designs a dual expert network in which one expert network focuses on processing attribute information in different dimensions.

4. Zero-Shot Compositional Recognition

4.1. Problem Formulation

Suppose $S = \{s_1, s_2, \dots, s_n\}$ is a set containing n state elements and $O = \{o_1, o_2, \dots, o_m\}$ is a set containing object elements. Then a compositional set can be obtained as $Y = \{(s_1, o_1), (s_1, o_2), \dots, (s_n, o_m) | s \in S, o \in O\}$, which containing $n \times m$ compositions. Now we have a seen domain $D^s = \{(x^s, y^s) | x^s \in X^s, y^s \in Y^s\}$, where X^s denotes the training image set and $Y^s \subseteq Y$ is the label set. Similarly, we have an unseen domain denoted by $D^u = \{(x^u, y^u) | x^u \in X^u, y^u \in Y^u\}$, where X^u denotes the test image set and $Y^u \subseteq Y$ is the label set. Data in the seen and unseen domains are disjoint, i.e., $X^s \cap X^u = \varnothing$ and $Y^s \cap Y^u = \varnothing$. The objective of ZSCR is to generalize to the unseen domain by learning the element-wise knowledge in the seen domain. Specifically, for closed-world compositional recognition, the goal is to learn a mapping function $\Psi: X \rightarrow Y^s \cup Y^u$, wherein X^u is strictly inaccessible for training. Typically, the size of $Y^s \cup Y^u$ is smaller than that of Y since some compositions are not feasible. For open-world compositional recognition [27], the goal is to learn a mapping function $\Psi: X \rightarrow Y$, assuming that which compositions are infeasible is unknown.

We divide the cutting-edge techniques of ZSCR into three categories based on their core components, including Decoupling Representation, Dependency Modeling, and Relationship Modeling. We then further split the sub-categories based on specific technical details or targets of action.

4.2. Decoupling Representation

A common baseline approach involves mapping images to deep semantic space, computing similarity to text embeddings, and optimizing with cross-entropy loss. The goal is to learn the posterior probability $p(y|\phi(x))$, where $y = (s, o)$ and ϕ is a function mapping the image into the semantic space. This approach leads to model simplifying optimization by learning direct correspondences between visual features and labels, often missing independent effects of states and objects. For example, if the model has only seen sliced apple, it may tightly associate sliced with apple and fail to recognize sliced orange. An alternative is to separately predict the labels for states and objects to avoid semantic interference, known as decoupling representation. This approach aims to disentangle states and objects, encouraging the model to learn independent visual representations. By separately recognizing states and objects, the model can better generalize to new compositions. Decoupling representation focuses on learning posterior probabilities $p(s|\phi_s(x))$ and $p(o|\phi_o(x))$, where ϕ_s and ϕ_o are mapping functions. Mainstream approaches can be categorized as: Network Design and Attention Separation.

Network Design: Many studies resort to modifying the network structure to decompose visual features. Nan et al. [102] add two sub-networks at the bottom of the backbone network to extract state and object features, respectively, and then each predicts the corresponding labels and is trained with cross-entropy loss. CASUAL [41] employs a similar network structure. The difference is that it introduces a text-to-image mapping network that reverse maps text embeddings into visual space. Meanwhile, it uses the Hilbert-Schmidt Information Criterion, which is a non-parametric method for estimating the statistical dependence between samples of two random

variables, to optimize training for ensuring state and object independence. Yuan et al. [103] propose the I2CD framework, which achieves causal decoupling between states and objects in the feature space through a reversible causal mechanism of *decoupling–combining–recoupling*. ProtoProp [104] uses a prototypical network to learn independent representations. Specifically, regional features of images and prototypes of objects and states are computed for similarity and later optimized with cross-entropy loss. Hu et al. [105] propose a two-layer decoupling network. First, three sub-networks are used to extract state, object, and composition features, respectively. In the second layer, the state and object features are fed into a reorganization network to synthesize into a composition, while the composition feature is fed into two other sub-networks for further decoupling. CAILA [106] uses CLIP as a visual coder, which inserts three adapters in the middle of the network to capture the state, object, and composition semantic information, respectively. It is worth mentioning that the three-branch network has become a common infrastructure for most approaches, i.e., three sub-networks extract state, object, and composition visual features, respectively [107]. Rao et al. [108] decouple objects and states into distinct groups using multi-expert networks, thereby enhancing the discriminability and transferability of visual features.

Attention Separation: The motivation for attention separation is to dissect semantic features by capturing the parts of objects and states attended to separately. OADis [109] introduces a cross-sample attention mechanism to refine features. It equips each sample with two companion samples, one with the same object and different state, and the other with the same state and different object. Afterwards, cross attention is performed with the corresponding samples separately to refine the semantic features. ADE [110] adopts a similar training paradigm. Unlike pre-matched pairs, Jing et al. [111] utilize image retrieval techniques to find suitable matching samples for object and state branches, after which they are used to enhance the representations. Zhang et al. [112] propose a gradient-based attention decomposition method. They view compositional recognition as a domain generalization problem, and one of the solutions is to learn domain-invariant features. To this end, they measure the variation of features by evaluating the difference in the back-propagation gradients of sample pairs to recognize the domain-invariant part. DFSP [113] firstly decouples text embeddings by prototype computation and, after that, uses the text embeddings and image features for cross-attention to decompose the visual features. Similarly, Troika [114] fuses multi-path prompt and cross attention to optimize text embeddings.

4.3. Dependency Modeling

Decoupling techniques pursue the pure semantic representations while mitigating the bias toward seen compositional concepts, introducing the bias toward the single concept, i.e., state or object. This bias comes from the fact that the same concept is represented differently in diverse contexts. Due to the existence of complex acting mechanisms in the real world, such as physical systems and biochemical reactions, the visual presentation of the same state on various objects varies greatly. For example, old is represented as metal rust spots on car, while it is shown as skin wrinkles on person. Obviously, concise text annotations cannot capture such subtle differences, and the independent visual representations learned by the model from the seen domain are biased, hindering the recognition accuracy of the unseen domain. To address this problem, many studies are devoted to modeling the dependency relationship between states and objects. By introducing a priori information about states or objects to correct each other's semantic representations or classification confidence, they can be better generalized to the unseen domain. In short, the goal of dependency modeling is to learn the conditional probabilities $p(s|\phi_s(x), o)$ and $p(o|\phi_o(x), s)$. The relevant technologies, according to the implementation objectives, can be categorized into Feature Optimization and Aid Classification.

Feature Optimization: The goal of feature optimization is to adjust each other's feature representations by leveraging the semantic relevance of states and objects. Decoupled state/object features are expected to capture dependent semantic details while maintaining their discriminative semantic information. Visual feature optimization corrects the visual embedding distribution to avoid overfitting the seen domain caused by the model focusing on a single visual feature. For example, Panda et al. [115] set up multiple encoders between the state and object branching networks to transmit object features to the state network to model the dependency of the state on the object. The same idea is taken by CoT [116], which uses object information as cuing signals while introducing an attention-based method to capture state-related visual features (Figure 5). Another approach is to optimize text features. Since text annotations are usually pre-extracted, they can hardly perceive dependency relationships. To this end, Nagarajan et al. [6] model states as operators, i.e., encode the state as a mapping function rather than an embedding, and thus explicitly learn the semantic embedding of an object *after it has been operated on*. SymNet [117] introduces two sub-networks, the decoupling network and the coupling network, which are used to learn independent and compositional representations of objects and states. CANet [118] recognizes the object labels of an image first and then corrects the text embedding of the state by using the object text embedding as a priori

condition. Li et al. [119] propose an progressive observation strategy that first assesses the similarity between object and state text and image features, subsequently generating refined conditional features.




Figure 5. A schematic diagram of dependency modeling. Object information is used as a signal to guide the extraction process of state features.

Aid Classification: Compared to adjusting the feature distribution, aid classification aims to directly introduce semantic information of objects/states to make the classifier more discriminative. In earlier work, Chen et al. [2] utilize support vector machines (SVMs) for compositional recognition. They learn the weights of SVMs by capturing state dependencies on objects and generalize to unseen compositions. SAD-SP [120] takes an intuitive approach by introducing an attention module in addition to the conventional decoupling network. State features are used to predict object soft labels and vice versa. The predicted soft labels are then used to weigh the predicted output of the conventional network. ProCC [121] performs information exchange at each layer of the state/object classifier. Specifically, it passes the output of one layer of the state classifier through a one-dimensional convolution and softmax, and then the output of the corresponding layer of the object classifier is multiplied with it and used as the input of the next layer, and vice versa.

4.4. Relationship Modeling

Dependency modeling promotes the generalizability of models by exploring the intrinsic one-to-one state-object connections. Naturally, we wonder whether many-to-many complex mesh relationships, i.e., interactions between multiple objects or states, contribute to model performance. Such relationships describe the latent logical relations of elements in the real world due to various factors, such as subject affiliation, visual similarity, co-occurrence, etc. To name a few, we would spontaneously categorize broccoli and spinach, separating beef since they belong to different subjects; a fork and a spoon are more visually similar than an umbrella, and also they have a higher probability of co-occurrence. Intuitively, an understanding of these relations facilitates logical inference. It is because humans have a remarkable ability to draw analogies, make connections, and imagine, thus reasoning and recognizing by tapping into the relationships between different elements. When we see a furry black bear and know the similarities between polar bears and black bears, it is easy to migrate the state of furry to polar bears. Driven by this idea, relationship modeling aims to enhance the inference and representation capabilities of models by uncovering and maintaining potential relationships between states and objects. The dominant techniques can be summarized as Cross Attention, Graph modeling, and Contrastive Embedding.

Cross Attention: To model the relationship, cross attention is used to capture similarities between text embeddings and thus correct text representations. BMPNet [122] introduces the cross attention mechanism to improve independent object/state representations. Specifically, for a state, the attention weights of the state and other states are first computed, and then the attention weights of an object are computed, resulting in an enhanced state representation. The object ditto. Unlike object/state relational modeling, CAPE [123] aims to reveal co-occurrence relations between compositions. It adapts the compositional features by calculating the attention weights between the text embeddings of the compositions.

Graph Modeling: As mentioned in the previous section, graphs are excellent tools for structured data and thus can be well adapted to the relationship modeling task. In earlier work, Cruz et al. [13] develop a logic tree to model relations between states and objects. They want the model to learn the notion of connectives such as {and, or, not} to understand the potential connections between nouns. CGE [124] models the entire set of objects and the set of states. It treats each object and state as a graph node and then identifies edges between nodes by the composition labels. Graph convolutional network is later introduced to propagate the information. CVGAE [125] expects the model to be able to determine on its own whether there is a connection between objects and states and

thus be adaptable to the open-world compositional recognition task. It designs a reconstruction task where the reconstruction network recovers the adjacency matrix based on latent variables. GIPCOL [126], on the other hand, fuses prompt learning and graph modeling. It employs CLIP as the backbone and takes the graph-encoded text embeddings as prompts to train. The graph neural network is involved in the learning process, thereby optimizing the text embeddings to fit the task better.

Contrastive Embedding: Contrastive embedding utilizes contrastive learning to optimize the semantic space and implicitly learn the relationships between samples. Contrastive learning [127], as a representative unsupervised learning method, demonstrates powerful semantic representational ability. It aims at clustering similar samples together and separating different samples. The learning mode is to constrain the semantic distance by comparing positive and negative examples, and the learned semantic distance reflects the affinity between the samples, i.e., the semantic relationship. Therefore, many areas exploit contrast learning to optimize semantic embeddings. Unlike the attention mechanism and graph modeling described above, contrast learning does not pre-define the relationship between samples explicitly but rather iteratively adjusts the semantic distances between samples to place them in appropriate positions. The performance of contrast learning largely depends on the number and quality of negative examples, and therefore, negative sampling is another critical aspect of contrastive learning.

Wei et al. [48] propose a quintuplet loss to optimize the text embedding space. Specifically, they use a sub-network to synthesize state and object texts into compositions. After that, the label of an image is used as a positive example and supplemented with three negative examples for contrastive learning. The negative examples are text compositions with the same state but different objects, the same object but different states, and different objects and states with the target image. SCEN [128] performs contrastive learning in the visual semantic space. It defines two databases to sample the negative examples, one with the same state as the target image but with different objects, and the other with the same object but with different states. After that, contrastive learning is implemented in the decoupled state/object embedding space. Yang et al. [129] adopt a similar approach to explore the relations among states or objects. Meanwhile, they introduce instance-level contrast learning to optimize the compositional embedding space.

4.5. Other Categories

In addition to the above mainstream techniques, many other studies provide unique perspectives to address the problem of compositional recognition. For example, TMN [26] presents a gated network that recognizes different compositions by injecting state and object information into the network to determine the switching of pathways. Similarly, TAFE-Net [130] proposes a task-aware learning paradigm that guides the studying process by using task information as cues. Some methods are devoted to analyzing the imbalance issue. MUST [131] argues that the difficulty of recognizing elements varies, e.g., mashed banana is hard to be identified as banana while relatively easy to be identified as mashed. For this reason, it introduces a weighted loss to balance the effects of different elements. ProLT [132], on the other hand, points out that the data imbalance problem caused by different sample sizes of elements affects the recognition accuracy. To this end, it similarly proposes a novel weighted loss to solve this problem. CDS-CZSL [133] argues that for an object, the importance of different states is different. For example, sliced is more important than red for strawberry, cause red is a more general state that can be reflected in a wide range of categories. In this regard, it first highlights those important states by clustering and contrasting, and then a specificity-refined primitive loss is implemented. Li et al. [134] employ optimal transport to capture fine-grained relationships between visual regions and text. To model the semantic relationships between objects and states, Wu et al. [135] employ a large language model to generate finer-grained semantic relationship trees for objects and states. Moreover, several studies develop new tasks to address larger challenges [136,137].

5. Foundation Model-Based Open-World Recognition

5.1. Problem Description

Large Vision-Language Models (VLMs) are advanced deep learning models that seamlessly integrate visual and language processing capabilities [138]. By analyzing extensive datasets of image-text pairs, these models learn the intricate relationships between vision and language, allowing them to understand and generate detailed language descriptions of image content [28]. VLMs are renowned for their robust cross-modal understanding and nuanced element perception, which make them particularly effective for zero-shot adaptation tasks. It is only necessary to associate the text description with the class name of the downstream recognition task, e.g., a photo of a [class], to achieve the purpose of classification and yield excellent results. However, recent studies suggest the capabilities of VLMs have not been fully exploited. Thus the foundation model-based open-world recognition receives more attention.

To provide a deep understanding of this topic, we first compare several prominent VLM-related research tasks. First, Better VLMs explores to improve the pre-training process of VLMs. The motivation for most studies stems from two aspects, i.e., data efficiency and learning efficiency. Data-efficient pre-training aims to reduce the data required while preserving strong representations [29,139], since large datasets, like CLIP's 400 million image-text pairs, are resource-heavy. Learning-efficient pre-training is to explore better alternative training methods like reconstruction [140], generation [141]. Next, Few-Shot Adaptation and OOD Detection, in which the former involves fine-tuning pre-trained VLMs with limited labeled data to create specialized models for specific tasks [30,31], which the latter aims to identify samples not belonging to the target categories [142,143]. Note that the OOD does not need to recognize the specific category of the foreign samples. Notably, both few-shot adaptation and OOD detection require labeled data to support fine-tuning VLMs. In contrast, Zero-Shot Adaptation releases this limitation. It consists of three basic tasks that apply VLMs to open environment recognition tasks under different constraints, including (i) Zero-shot Prediction: with only the category names of the task known and without any image data applying VLMs to the downstream task; (ii) Test-time Adaptation: know the category names of the task and allow adjustments at the online testing phase; and (iii) Unsupervised Adaptation: Know the category names of the task and have part or entire of the unlabeled test data in advance. In this survey, we refer to the three tasks of zero-shot adaptation as the category of foundation model-based open recognition.

5.2. Zero-Shot Prediction

VLMs have powerful zero-shot prediction capabilities. For a specific task, simply editing the category names into the prompt, e.g., a photo of [class], enables direct classification without fine-tuning. However, there is a performance gap between such a deployment and a model that is fully supervised and trained on the downstream task. To address this issue, zero-shot prediction presents promising solutions. A prevalent solution is to design better prompts. The commonly used a photo of [class], while intuitive, is missing a lot of detailed information. For this reason, Menon et al. [32] propose to use a large language model such as GPT-3 [144] to generate multiple word descriptions for classes. In the testing time, the similarity between the image sample and each word description is calculated, and the mean value is taken as the discriminative basis. CuPL [145] has a similar idea, which uses this description directly as the prompt for testing. CHiLS [146] also adopts a large language model, which generates multiple subclasses for each category and uses their combined similarity as the discriminative basis. APPLe [147] uses the GPT-3 descriptions to initialize prompt prototypes as the category classifiers. Unlike the design of prompt, SuS-X [148] uses Diffusion [149] to pre-generate multiple pseudo-visual samples for each class, which serves as a support set to assist the judgment at the testing time. CALIP [33] takes the perspective of embedding space, and it introduces the attention mechanism to capture the correlation between text embeddings and visual embeddings, which effectively improves the zero-shot generalization capability. PromptMix [150] proposes a unified prompt learning framework, the core of which lies in bridging data distribution gaps by constructing a modality-agnostic shared representation space, whilst utilising a large language model-assisted prompt evolution mechanism to automatically enrich and refine semantic prompts.

5.3. Test-Time Adaptation

The purpose of Test-Time Adaptation (TTA) is to dynamically adjust the model to progressively adapt to the downstream task during the testing process. The mainstream schemes are categorized into prompt adaptation and cache bank. Prompt adaptation is a technique that combines prompt learning and TTA. During testing, the prompt is set as learnable parameters and the parameters of the model are fixed. The model continuously optimizes the prompt representation and thus adapts to the target task. TPT [151] first presents this concept by augmenting the test image into multiple views and then filtering out those with low confidence. The confidence outputs of the remaining augmented samples are averaged as the final prediction, and at the same time, the information entropy is minimized to train the prompt (Figure 6). This paradigm provides a basic template for subsequent research on prompt adaptation. DiffTPT [152] argues that the augmentation approach of TPT lacks diversity and, therefore, introduces the diffusion model to generate richer samples for the test image. PromptAlign [153] argues that adjusting the prompt should be done while aligning the distribution of the target and pre-training data. Therefore it introduces additional loss to align the statistics of target and pre-training data, i.e., mean and variance. SwapPrompt [154] and DART [155] use the exponential moving average algorithm to update the prompt.

The goal of the Cache bank is to store the tested samples to improve the performance of subsequent tests. An intuitive approach is to cache the previous high-confidence samples as the support set to be queried by subsequent samples. DMN-ZS [156] introduces an attention mechanism to compute the similarity between the new test samples and the support set to obtain the weighted features. TDA [157] sets up two cache banks to store the one-

hot labels and the negative soft labels of the reliable samples, respectively. The new test samples interact with them to generate predictions. Zhang et al. [158] employ a cache bank to save high-confidence samples from learning high-quality prompts. MTA [159] proposes a new loss function to optimize visual features in the semantic space. Liang et al. [160] enhance cache quality through consensus-based perceptual entropy reweighting and propose distribution calibration to construct flexible decision boundaries, enabling reliable adaptation during testing. Zhu et al. [161] model category prototypes as dynamic distributions of textual and visual features, leveraging optimal transport fusion to enhance category representations by integrating visual information from test streams.




Figure 6. A schematic diagram of test-time prompt learning. Adjust the prompt by minimizing the entropy of the output.

5.4. Unsupervised Adaptation

Unsupervised adaptation is suitable for offline testing environments where some or all test data is available. Although there are no labels available, image samples can provide a priori information to help better prediction. Existing research combines different algorithms to address this problem. UPL [162] fuses prompt learning and self-supervised learning to learn prompt representations through a small number of high-quality pseudo-labels. MUST [163] incorporates image reconstruction techniques to facilitate the model's understanding of visual information. InMaP [164] introduces prototype learning, which mines a set of visual prototypes from the visual space to replace the text embeddings for prediction. LaFTer [165] trains a text classifier and then uses it to classify images and trains the image encoder with pseudo-labels. ZLaP [166] presents a label propagation algorithm to assign labels to images. In addition, CLIPPR [167] introduces a label distribution prior to predicting the data of the target domain.

6. Technical Details

6.1. Standardized Datasets

In Table 2, we list the commonly used standardized datasets for three tasks ZSOR, ZSCR, and FBOR, respectively. Specifically, there are 8 commonly used datasets used for ZSOR, including fine-grained datasets such as CUB [7], which is a bird dataset; SUN [8], which is a scene dataset, and coarse-grained datasets such as AWA [168], which is a dataset of various types of animals. Most of the datasets contain detailed attribute annotations. Note that Flowers102 [169] and NABirds [170] do not have attribute annotations. Flowers102 contains textual descriptions, i.e., each image is described by several sentences. NABirds has region annotations, which mark the parts of the bird and are typically used in visual component analysis methods. For ZSCR, there are 6 commonly used datasets. Different from ZSOR, the column [Class/Attr.] indicates the number of objects and the number of states they contain. It can be seen that their compositional sizes cover different orders of magnitude, ranging from 3/8 (AO-CLEVR [41]), which means up to 24 compositions, to 674/413 (C-GQA [124]), which means 278,362 compositions. Last, we show 15 commonly used datasets for FBOR that cover different categories and different data sizes. ImageNet [1] includes data with up

to 1000 categories, while EuroSAT [171] and CIFAR-10 [172] have only 10 categories. The data types are also diverse, including food [173], flowers [169], pets [174], transportation [175,176], etc. They are also fine-grained datasets. In addition, DTD [177] is a texture dataset, and EuroSAT [171] is a remote sensing dataset.

Table 2. A list of commonly used datasets for zero-shot image recognition. [Class/Attr.] means the number of classes and attributes for ZSOR, objects and states for ZSCR, classes for FBOR. Click the [link] (all accessed on 25 August 2024) to check more information in the corresponding repository.

Dataset	Year	Images	Class/Attr.	Download
Zero-Shot Object Recognition				
CUB [7]	2011	11,788	200/312	[Link]
Flowers102 [169]	2008	8189	102/-	[Link]
SUN [8]	2012	14,340	717/102	[Link]
NABirds [170]	2015	48,562	1011/-	[Link]
DeepFashion [40]	2016	289,222	46/1000	[Link]
AWA [168]	2013	30,475	50/85	[Link]
AWA2 [22]	2018	37,322	50/85	[Link]
APY [178]	2009	15,339	32/64	[Link]
Zero-Shot Compositional Recognition				
UT-Zappos [179]	2014	50,025	12/16	[Link]
MIT-States [180]	2015	53,753	245/115	[Link]
C-GQA [124]	2021	39,298	674/413	[Link]
AO-CLEVR [41]	2020	180,000	3/8	[Link]
VAW-CZSL [109]	2022	92,583	541/440	[Link]
Clothing16K	2020	16,170	8/9	[Link]
Foundation model-Based Open-world Recognition				
ImageNet [1]	2009	1,431,167	1000	[Link]
Caltech101 [181]	2004	9145	101	[Link]
Caltech256 [182]	2007	20,607	256	[Link]
SUN397 [183]	2010	108,753	397	[Link]
Food101 [173]	2014	101,000	101	[Link]
Flowers102 [169]	2008	8189	102	[Link]
StanfordCars [175]	2013	16,185	197	[Link]
FGVCAircraft [176]	2013	10,000	100	[Link]
OxfordPets [174]	2012	7349	37	[Link]
DTD [177]	2014	5640	47	[Link]
EuroSAT [171]	2019	27,000	10	[Link]
Country211 [28]	2021	64,300	211	[Link]
CIFAR-10 [172]	2009	60,000	10	[Link]
CIFAR-100 [172]	2009	60,000	100	[Link]
Birdsnap [184]	2014	49,829	500	[Link]

6.2. Implementation

To facilitate the deep understanding and fast reusing of ZSIR approaches, this survey further provides the implementation details of most representative methods in Tables 3–5. Specifically, Tables 3 and 4 show the implementation details of the related studies of ZSOR and ZSCR, including the backbone, auxiliary, dataset, and code addresses. It can be seen that commonly used backbones include ResNet [185], VGG [186], GoogleNet [187], AlexNet [188], ViT [69], and CLIP [28]. [Auxiliary] refers to additional information such as region-level annotations. We provide the dataset and code address to facilitate access to the performance of the method as well as its implementation. Finally, we label the study with up to two keywords based on the specific technique used.

In Table 5, we show more details of the related work in foundation model-based open-world recognition. We list their visual and text encoders, as well as auxiliary information. GPT-3 [144] and Diffusion [149] are the exotic models. [Statistic] denotes the mean and variance of the pre-training data. [Distribution] denotes the prior distribution of labels. We also provide their code repositories for quick reference.

Table 3. More details on selected advances for zero-shot object recognition. [Aux.] means auxiliary. [GNet] means GoogleNet, [ANet] means AlexNet, [R101] means ResNet101, [V19] means VGG19. [Region] means region-level annotations. [FLO] means Flowers102, [NAB] means NABirds, [DEP] means DeepFashion. Click the [link] (all accessed on 25 August 2024) to check more information in the corresponding repository.

	Name	Venue	Year	Backbone	Aux.	KeyWords	Dataset	Code
Visual	LH2B [62]	CVPR	2017	VGG16	Region	Detection	CUB, NAB	[Link]
	S2GA [61]	NeurIPS	2018	VGG16	Region	Detection	CUB, NAB	[Link]
	AGAA [63]	CVPR	2018	VGG16	Region	Detection	CUB, NAB	[Link]
	LDF [64]	CVPR	2018	GNet, VGG19	-	Crop	CUB, AWA	[Link]
	AREN [35]	CVPR	2019	ResNet101	-	Attention	CUB, SUN, AWA2, APY	[Link]
	SGMA [34]	NeurIPS	2019	VGG19	-	Crop	CUB, FLO, AWA	[Link]
	CANZSL [53]	WACV	2020	VGG16	Region	Consistency	CUB, NAB	[Link]
	RGEN [36]	ECCV	2020	ResNet101	-	Attention, Graph	CUB, SUN, Awa2, APY	-
	SDGZSL [56]	ICCV	2021	ResNet101	-	Disentangling	CUB, AWA2, APY, FLO	[Link]
	RSAN [60]	CIKM	2021	ResNet101	-	Attention	CUB, SUN, AWA2	-
	SR2E [65]	AAAI	2021	ResNet101	-	Crop	CUB, SUN, AWA2, APY	-
	VGSE-SMO [70]	CVPR	2022	ResNet50	-	Crop, Clustering	CUB, SUN, AWA2	[Link]
	GSMFlow [57]	TMM	2022	ResNet101	-	Generation Shifts	CUB, AWA2, APY, FLO	[Link]
	ERPCNet [66]	TCSVT	2022	ResNet101	-	Crop	CUB, SUN, AWA2, APY	-
	GKU [37]	AAAI	2023	ResNet34	Region	Graph	CUB, NAB	-
	EOPA [189]	TPAMI	2023	ANet, ResNet50	-	Graph	CUB, SUN, FLO, AWA2	-
	POPRNet [67]	TIP	2024	ResNet101	-	Crop	CUB, SUN, AWA2	[Link]
	ZSLViT [68]	CVPR	2024	ViT-B	-	Attention	CUB, SUN, AWA2	-
	SVIP [71]	ICCV	2025	ViT-B	-	Crop	CUB, SUN, AWA2	[Link]
Visual-Attribute	DAZLE [38]	CVPR	2020	ResNet101	-	Attention	CUB, SUN, AWA2, DEP	[Link]
	Composer [77]	NeurIPS	2020	ResNet101	-	Attention, Generation	CUB, SUN, AWA2, DEP	[Link]
	APN [90]	NeurIPS	2020	ResNet101	-	Prototype	CUB, SUN, AWA2	[Link]
	GEM [76]	CVPR	2021	ResNet101	Region	Attention	CUB, SUN, AWA2	[Link]
	DPPN [92]	NeurIPS	2021	ResNet101	-	Prototype	CUB, SUN, AWA2, APY	[Link]
	LAPE [190]	TCSVT	2022	ResNet101	-	Attention	CUB, SUN, AWA2, APY	-
	RIAE [73]	IS	2022	ResNet101	-	Attention, Graph	CUB, SUN, AWA2	-
	DPDN [47]	MM	2022	ResNet101	-	Attention, Prototype	CUB, SUN, AWA2	-
	GNDAN [74]	TNNLS	2022	ResNet101	-	Attention, Graph	CUB, SUN, AWA2	[Link]
	AREES [78]	TNNLS	2022	ResNet101	-	Attention, Generation	CUB, SUN, AWA, AWA2, APY	-
	MSDN [82]	CVPR	2022	ResNet101	-	Attention	CUB, SUN, AWA2	[Link]
	TransZero [39]	AAAI	2022	ResNet101	-	Attention	CUB, SUN, AWA2	[Link]
	I2DFormer [79]	NeurIPS	2022	ViT-B	-	Attention	CUB, FLO, AWA2	[Link]
	DUET [80]	AAAI	2023	ViT-B	-	Attention, Contrast	CUB, SUN, AWA2	[Link]
	PSVMA [81]	CVPR	2023	ViT-B	-	Attention	CUB, SUN, AWA2	[Link]
	HRT [191]	PR	2023	ResNet101	-	Attention	CUB, SUN, AWA2	-

Table 3. Cont.

	Name	Venue	Year	Backbone	Aux.	KeyWords	Dataset	Code
Visual-Attribute	CC-ZSL [95]	TCSVT	2023	ResNet101	-	Prototype, Contrast	CUB, SUN, AWA2	[Link]
	CoAR-ZSL [96]	TNNLS	2023	R101, ViT-L	-	Attention, Prototype	CUB, SUN, AWA2	[Link]
	GIRL [94]	IS	2023	ResNet101	-	Prototype	CUB, SUN, AWA2, APY	[Link]
	VSPCN [84]	CVPR	2025	ViT-B	-	Attention	CUB, SUN, AWA2	-
	ZeroMamba [85]	AAAI	2025	VMamba-Small	-	Attention	CUB, SUN, AWA2	[Link]
	CSAP-ZSL [93]	AAAI	2025	ViT-B	-	Prototype	CUB, SUN, AWA2	-
	HPAN [86]	MM	2025	ViT-B	-	Attention	CUB, SUN, AWA2	[Link]
Attribute	MSDN++ [83]	IJCV	2026	ResNet101	-	Attention	CUB, SUN, AWA2, FLO	-
	MCZSL [100]	CVPR	2016	VGG16	Region Crop		CUB	-
	LFCAA [98]	ICCV	2019	GNet, R101, V19	-	Attention	CUB, SUN, AWA2	[Link]
	APNet [99]	AAAI	2020	ResNet101	-	Graph	CUB, SUN, AWA, AWA2, APY	-
	DEDN [101]	IJCAI	2024	ResNet101	-	Attention, Distillation	CUB, SUN, AWA2	[Link]

Table 4. More details on selected advances for zero-shot compositional recognition. [UTZ] means UT-Zappos, [MIT] means MIT-States, [CGQ] means C-GQA, [AOC] means AO-CLEVR, [VAW] means VAW-CZSL, [CLO] means Clothing16K. Click the [link] (all accessed on 25 August 2024) to check more information in the corresponding repository.

	Name	Venue	Year	Backbone	KeyWords	Dataset	Code
Decoupling	Nan et al. [102]	AAAI	2019	ResNet18	Contrast	UTZ, MIT	-
	CAUSAL [41]	NeurIPS	2020	ResNet18	Mapping Function	UTZ, AOC	[Link]
	ProtoProp [104]	NeurIPS	2021	ResNet18	Prototype, Graph	UTZ, AOC, CGQ	[Link]
	OADis [109]	CVPR	2022	ResNet18	Attention	UTZ, MIT, VAW	[Link]
	Zhang et al. [112]	ECCV	2022	ResNet18	Attention	UTZ, AOC, CLO	[Link]
	DECA [107]	TMM	2022	ResNet18	Mapping	UTZ, MIT	[Link]
	DRANet [192]	ICCV	2023	ResNet18	Attention	UTZ, MIT, CGQ	-
	ADE [110]	CVPR	2023	ViT-B	Attention	UTZ, CGQ, CLO	[Link]
	DFSP [113]	CVPR	2023	CLIP	Attention, Prototype	UTZ, MIT, CGQ	[Link]
	Hu et al. [105]	AAAI	2023	ResNet18	Contrast	UTZ, CGQ	[Link]
	CAILA [106]	WACV	2024	CLIP	Mapping	UTZ, MIT, CGQ, VAW	[Link]
	Jing et al. [111]	AAAI	2024	CLIP	Attention	UTZ, MIT, CGQ	-
	Troika [114]	CVPR	2024	CLIP	Attention	UTZ, MIT, CGQ	[Link]
	HGRL [108]	IJCAI	2025	CLIP	Multi-Expert Network	UTZ, MIT, CGQ	-
Dependency	I2CD [103]	AAAI	2026	CLIP	Mapping Function	UTZ, MIT, CGQ	-
	Nagarajan et al. [6]	ECCV	2018	ResNet18	Contrast	UTZ, MIT	[Link]
	SymNet [117]	CVPR	2020	ResNet18	Mapping Function	UTZ, MIT	[Link]
	SAD-SP [120]	TPAMI	2023	ResNet18	Mapping Function	UTZ, MIT, CGQ	-
	CoT [116]	ICCV	2023	ResNet18, ViT-B	Attention	MIT, CGQ, VAW	[Link]
	CANet [118]	CVPR	2023	ResNet18	Mapping Function	UTZ, MIT, CGQ	[Link]
	Panda et al. [115]	PR	2024	ResNet18	Graph	UTZ, MIT, CGQ	-
	ProCC [121]	AAAI	2024	ResNet18	Mapping Function	UTZ, MIT, CGQ	[Link]
	PLO-VLM [119]	MM	2025	CLIP	Attention	UTZ, MIT, CGQ	-

Table 4. Cont.

	Name	Venue	Year	Backbone	KeyWords	Dataset	Code
Relationship	Wei et al. [48]	ICCV	2019	ResNet18	Contrast	UTZ, MIT	-
	BMPNet [122]	TMM	2021	ResNet18	Attention	UTZ, MIT	[Link]
	CGE [124]	CVPR	2021	ResNet18	Graph	UTZ, MIT, CGQ	[Link]
	SCEN [128]	CVPR	2022	ResNet18	Contrast	UTZ, MIT, CGQ	[Link]
	CVGAE [125]	MM	2022	ResNet18	Graph	UTZ, MIT, CGQ	[Link]
	CAPE [123]	WACV	2023	ResNet18	Attention	UTZ, MIT, CGQ	-
	Yang et al. [129]	TMM	2023	ResNet18	contrast	UTZ, MIT	
	GIPCOL [126]	WACV	2024	CLIP	Graph	UTZ, MIT, CGQ	[Link]
Others	TMN [26]	ICCV	2019	ResNet18	Gated Network	UTZ, MIT	[Link]
	MUST [131]	PR	2024	ResNet18	Weighted Loss	UTZ, MIT, CGQ	[Link]
	ProLT [132]	AAAI	2024	ResNet18, CLIP	Weighted Loss	UTZ, MIT, CGQ	-
	CDS-CZSL [133]	CVPR	2024	CLIP	Weighted Loss	UTZ, MIT, CGQ	-
	Tsca [134]	IJCAI	2025	CLIP	Optimal Transport	UTZ, MIT, CGQ	[Link]
	LogiCZSL [135]	CVPR	2025	CLIP	Logic Rules	UTZ, MIT, CGQ	[Link]

Table 5. More details on selected advances for foundation model-based open-world recognition. [ZSP] means zero-shot prediction. [TTA] means test-time adaptation. [UA] means unsupervised adaptation. [VEncoder] means visual encoder. [TEncoder] means text encoder. [ViT-B†]: ViT-B/32, ViT-B/16. [ViT-L†]: ViT-L/14, ViT-L/14@336px. [Statistic] means the mean and variance of pretrained dataset. [Distribution] means the label distribution. Click the [link] (all accessed on 26 August 2024) to check more information in the corresponding repository.

	Name	Venue	Year	VEncoder	TEncoder	Auxiliary	Contribution	Code
ZSP	Menon et al. [32]	arXiv	2022	ViT-B†, ViT-L†	Transformer	GPT-3	Generate description words according to category names.	-
	APPLe [147]	arXiv	2023	ViT-L/14	Transformer	GPT-3	Propose prompt prototype learning by initializing from description sentences.	-
	CuPL [145]	ICCV	2023	ViT-L/14	Transformer	GPT-3	Generate description sentences according to category names.	[Link]
	SuS-X [148]	ICCV	2023	ResNet50	Transformer	Diffusion	Generate visual support set with Diffusion.	[Link]
	CHiLS [146]	ICML	2023	ViT-L/14@336px	Transformer	GPT-3	Generate sub-class names according to category names.	[Link]
	CALIP [33]	AAAI	2023	ResNet50	Transformer	-	Propose attention interaction of image and text embeddings.	[Link]
	PromptMix [150]	INF	2026	ViT-B/16	Transformer	Qwen	A prompt evolution mechanism assisted by large language models	-
TTA	TPT [151]	NeurIPS	2022	ResNet50, ViT-B/16	Transformer	-	Propose test-time prompt adaptation.	[Link]
	DiffTPT [152]	ICCV	2023	ResNet50, ViT-B/16	Transformer	Diffusion	Propose to augment image with Diffusion.	[Link]
	PromptAlign [153]	NeurIPS	2023	ViT-B/16, ViT-B/32	Transformer	Statistic	Align the distributions of test and pre-training set.	[Link]
	SwapPrompt [154]	NeurIPS	2023	ResNet50	Transformer	-	Update the prompt using exponential moving average.	-

Table 5. Cont.

	Name	Venue	Year	VEncoder	TEncoder	Auxiliary	Contribution	Code
TTA	Zhang et al. [158]	AAAI	2024	ViT-B/16	Transformer	-	Introduce multiple manual prompts.	-
	DART [155]	AAAI	2024	ViT-B/16	Transformer	-	Learn a separate prompt for each class.	-
	TDA [157]	CVPR	2024	ResNet50, ViT-B/16	Transformer	-	Two cache banks are set up to store hard and soft labels.	[Link]
	MTA [159]	CVPR	2024	ViT-B/16	Transformer	-	Improve visual representation in the embedding space.	[Link]
	DMN-ZS [156]	CVPR	2024	ResNet50, ViT-B/16	Transformer	-	Cache high-confidence samples for subsequent predictions.	[Link]
	ReTA [160]	MM	2025	ResNet50, ViT-B/16	Transformer	-	Using entropy to update the cache bank	[Link]
	ProtoMM [161]	ICCV	2025	ResNet50, ViT-B/16	Transformer	-	Using optimal transport to update the cache bank	-
UA	CLIPPR [167]	arXiv	2022	ViT-B/32	Transformer	Distribution	Predicted label aligns a priori label distribution.	[Link]
	UPL [162]	arXiv	2022	ResNet50	Transformer	-	Select high-confidence samples for fine-tuning.	[Link]
	MUST [163]	ICLR	2023	ViT-B/16, ViT-L/14	Transformer	-	Introduce visual reconstruction and teacher-student model.	[Link]
	InMaP [164]	NeurIPS	2023	ResNet50, ViT-B/16	Transformer	-	Learn visual prototypes in the embedding space.	-
	LaFTer [165]	NeurIPS	2023	ViT-B/32	Transformer	GPT-3	Train a text classifier to categorize images.	-
	ZLaP [166]	CVPR	2024	ResNet50, ViT-B/16	Transformer	-	Introduce a label propagation algorithm.	[Link]

6.3. Evaluation Protocols

In this section, we introduce three commonly used evaluation protocols for the three tasks. For ZSOR, the test set is typically divided into seen and unseen classes, and the top-1 accuracy is calculated for each. Additionally, the harmonic mean (HM) is introduced to provide a more comprehensive assessment of model performance. Suppose S and U denote the accuracy of seen and unseen classes, respectively. $HM = \frac{2 \times S \times U}{S + U}$. For ZSCR, the test set is similarly divided into seen and unseen classes, and their top-1 accuracy is measured. Additionally, HM and AUC are introduced to comprehensively evaluate model performance. AUC is the area under the receiver operating characteristic curve. For FBOR, top-1 accuracy is typically used to evaluate model performance.

6.4. Shared Technologies

It can be seen from Tables 3–5 that some techniques are common across the three tasks. We further discuss these shared techniques to reveal more commonalities between these tasks. In particular, attention is the most commonly used technique, distributed across almost every sub-direction in the three tasks. This suggests that attention techniques are extremely generalizable and can be applied to different task goals and with different research motivations. In ZSOR, attention is used to localize key visual regions [35,36], model visual-text correlations [38,39,82], and localize key attributes [98]. In ZSCR, attention is used to model object and state

correlations [109,112,116], and enhance representations [122,123]. In FBOR, attention is used to capture correlations between images and text [33]. Prototype learning is another more generalized technique. In ZSOR, prototypes are used to represent attributes [90–92]. In ZSCR, prototypes are used to represent objects and states [104,113]. In FBOR, prototypes are used to replace the embedding of prompts in deep space [164]. Graph and contrast learning are the more commonly used techniques in ZSOR and ZSCR. The graph is used to model the topological relationships between visual regions in ZSOR [37,73], while in ZSCR, it is used to capture the logical relationships between individual objects and states [124,125]. Contrast learning is used in ZSOR to optimize attribute embeddings [80], while in ZSCR, it is used to untangle and model semantic relations [48,128]. In FBOR, thanks to the robust zero-shot performance of VLMs, fine-tuning prompts become a popular approach [151,154]. Similarly, in ZSCR, some methods with CLIP as the backbone introduce prompt learning [113,126], which effectively improves the performance of the model.

7. Application

In recent years, besides image recognition, element-wise zero-shot techniques have been applied in a wide range of areas and developed synergistically. We sketch out some following related applications.

7.1. Language and Text

Language, as one of the main media used by humans, has always received attention from researchers. Due to label pressures, zero-shot generalization is also a long-standing direction in language studies [193]. In recent years, large language models have been developed rapidly [194], and they can respond to human conversations fluently, showing their powerful zero-shot reasoning ability. Meanwhile, some studies [195] proposed fine-tuning schemes for language models to improve the zero-shot performance. In some specific tasks, zero-shot techniques have been widely applied to solve the problem of insufficient data samples, such as machine translation [196], question answering [197], etc. In addition, zero-shot techniques also play an important role in text analysis, such as sentiment classification [198], topic classification [199], etc., which is extremely helpful for research in sociology, psychology, and other disciplines. Meanwhile, the collection and labeling of text samples is a huge obstacle, and zero-shot generalization offers a promising solution [200]. Moreover, recent studies, such as named entity recognition [200] and information retrieval [201], also extensively rely on zero-shot ability.

7.2. Scene Understanding

Scene understanding is a topic based on image recognition but is more challenging. Instead of assigning a label to an entire image, it usually identifies and localizes multiple elements in an image. For example, multi-label recognition task [202] requires simultaneous recognition of multiple objects in an image. In contrast, object detection [203,204] and semantic segmentation [205,206] are two more complex and practical tasks. Object detection requires identifying and localizing multiple objects in an image, while semantic segmentation is a pixel-level classification task that requires assigning labels to each pixel according to semantic information. Scene understanding is an important research topic, which is not only the basis for AI to understand the world but also a key part of deployment to industrial scenarios. Related technologies are needed in a variety of scenarios, such as autonomous driving, medical diagnostic assistance, security monitoring, retail, and so on. However, they also require more detailed annotation and larger data for training than image recognition. Therefore, incorporating zero-shot learning into scene understanding has become an attractive direction in recent years. For example, zero-shot object detection [207–209] and zero-shot semantic segmentation [210–212] are dedicated to the deployment of models in the open world.

7.3. 3D Vision

With the increasing maturity of 3D sensors, the cost of acquiring 3D vision data has dropped significantly [213]. Compared with traditional 2D images, 3D visual data contains important depth information and, therefore, can provide richer semantic information. 3D vision data has a wide range of applications, covering automatic driving, medical diagnosis, remote sensing, etc. Correspondingly, research on 3D vision data has also received more and more attention. However, compared with 2D data, the acquisition and labeling cost of 3D data is more expensive. Therefore, many researches integrate the concept of zero-shot into 3D tasks, such as point cloud classification [214,215], object detection [216], and semantic segmentation [217]. In addition to semantic recognition tasks, some studies are also exploring the application of zero-shot learning to tasks such as

generation [218], style transfer [219], and synthesis [220,221]. There are also some studies that have explored in-depth on special tasks, such as human pose estimation [222], shape correspondence [223], and so on.

7.4. Others

Zero-shot learning offers promising solutions for many specific application scenarios. For example, medical and remote sensing images are much more difficult to collect than natural images due to sample scarcity and data protection. Therefore, the research on zero-shot medical image [224] recognition and remote sensing image [225] recognition has been developed rapidly in recent years. Anomaly detection [226] plays a key role in securing many industrial fields. However, the number and type of anomaly samples are often insufficient. Some studies resort to zero-shot learning to solve this problem. In addition, zero-shot techniques have been applied in many special tasks to alleviate the low-shot data problem as well as to enhance the generalization ability of the model, such as action recognition [227], image retrieval [228,229], video moment retrieval [230], Re-id [231] and video classification [232].

8. Future Directions

Zero-shot learning contributes powerful technical solutions for automated recognition in the open environment. Meanwhile, research based on element-wise representation and reasoning has driven the rapid development and wide application of zero-shot technology. Despite the remarkable achievements, there are still some potential limitations that have not been deeply explored. In this section, we share our perspectives on future research directions.

8.1. Continual Zero-Shot Learning

Existing studies are based on closed systems to design solutions. However, in reality, there are a large number of growing systems with dynamically changing environments and goals. For example, for zero-shot object recognition, attributes and categories may be progressively inflated. Early in the observation phase, fewer attributes are accessible, followed by a steady stream of new attribute annotations added. Similarly, the detected categories may increase over time forward. It is difficult for existing schemes to handle such dynamic recognition tasks, and continual learning has come a long way in this regard [233]. However, simply transplanting techniques from continual learning to zero-shot recognition tasks suffers from deviation. Therefore, fusing continual learning and zero-shot learning will further increase the value of applications to real-world scenarios.

8.2. Multi-Label Compositional Recognition

Zero-shot compositional recognition allows the model to not only learn to recognize objects but also to describe their state. However, current research is limited to a single object and single state. A robust compositional recognition system should have the ability to recognize multiple objects as well as multiple states. For a single object, the recognition system should describe its state from multiple perspectives to provide more detailed information [3]. For multiple objects, their individual states, as well as their interactions, relationships, etc., should be described [234]. Expanding from single labels to multiple labels is a tough choice, but it is an inescapable step toward building a smarter recognition system.

8.3. VLMs with Downstream Fine-Grained Annotations

The power of VLMs in the open world has been demonstrated by the fact that it requires only the names of the categories in the downstream task to accomplish zero-shot inference. Meanwhile, fine-grained annotations can improve the zero-shot inference performance of VLMs [32,145]. What happens if VLMs are combined with downstream exclusive attribute annotations, such as attribute texts or state descriptions? There are a number of issues that deserve deeper investigation, including domain shift between pre-trained text embeddings of VLMs and downstream text, fine-grained visual-text matching, catastrophic forgetting, and so on.

8.4. Annotation Cost and Quality

Element-wise schemes in zero-shot object recognition typically require extensive refined attribute annotations, especially for visual-attribute matching schemes. However, the attribute-level annotations are time- and labor-intensive compared to class-level labeling. Worse still, once they settle into concrete real-world scenarios, such as industrial inspection or medical pathology, the expert knowledge can be a bottleneck, which further raises the labor cost. In addition, attribute engineering is a complex crossover field. Even attributes annotated by experienced experts do not guarantee benefits for deep learning, which implies that high-quality

attribute annotations require professionals with dual knowledge of both specific domains and deep learning. Despite some studies attempting to make breakthroughs in the field of automated annotation [100], it is clear that there is still a long way to go.

8.5. Deployment Cost and Edge Feasibility

Compared to class-wise semantic modeling, many element-wise schemes have to process a higher density of information, which introduces a more luxurious deployment cost. Such cost is reflected in bloated network structures and high computational complexity in the deployment phase. As a result, they are unfriendly to edge tasks and mini-endpoints, which have to trade off performance and memory [235]. However, the minimal data requirements significantly alleviate the storage and communication pressures on edge devices, whilst also reducing the need for human intervention. The element-wise paradigm also demonstrates particular promise for edge deployment. Specifically, element-wise representations naturally decompose visual concepts into lightweight, reusable basic units. This decomposition offers tremendous potential for edge deployment. In rapidly changing edge scenarios, selective operations on a small number of basic elements are sufficient to perform category updates, decision-making calculations and resource replenishment. This means that the modular nature of element-wise reasoning makes it possible to activate only parts of the network, thereby improving computational efficiency. Furthermore, these basic conceptual units provide a heuristic approach to model compression. By identifying and pruning those element modules that are unimportant or irrelevant to the task, the model's memory requirements can be reduced.

8.6. Unified Framework

Zero-shot studies share many common techniques, such as attention, graph modeling, and so on. However, barriers between them still exist due to the different goals and contexts of the tasks. A technique and framework that is versatile across all fields is lacking. However, the significance of such a unified framework for both academia and application is immense [236]. The quest for a unified framework means exploring more commonalities across zero-shot tasks and thinking about the operational mechanisms of the entire realm from the ground-up principles. From a realistic point of view, a mature intelligence should be able to handle multiple tasks at the same time.

8.7. Robust Representation Learning

Despite some solutions have been proposed [97] to tackle with domain shift, there is a lack of quantitative assessments regarding their resistance to interference. Consequently, the robustness of existing representation methods in real-world environments remains uncertain. Recently, Shafiee et al. [237] introduce adversarial attack technique into zero-shot models. The interplay between model attacks and defenses has been a long-standing topic. Attackers aim to weaken models at minimal cost, while defenders strive to develop more robust representations to resist these perturbations. This dynamic process of learning from each other leads to continuous progress, offering a new perspective on combating domain shift.

8.8. Poor Theoretical Foundation

The development of zero-shot learning is established on the beautiful hypothesis that deep neural networks can reason logically like humans [18,22]. As mentioned, we can learn to identify basic elements and recognize new concepts by combining and reasoning about them. Nevertheless, there are not many solid theories on the compatibility between human reasoning and machine inductive ability, leading to a lack of explainability. Meanwhile, some flaws also challenge the plausibility of the hypothesis, such as the correspondence between abstract attributes and vision. Rigorous theoretical guidance is at the helm of a field moving forward, and it is of great prospective to dive into the mysterious black box in the future.

8.9. Ethical Considerations and Social Impact

The core idea behind element-wise zero-shot image recognition technology is to mimic the human bottom-up cognitive system. However, this may introduce potential ethical and social risks. Firstly, statistical bias may lead to bias in datasets and annotations. For example, specific cultural contexts may fail to reflect the global, diverse understanding of visual concepts. In zero-shot scenarios, this risk is amplified due to the lack of calibrated annotated data from affected domains. Furthermore, the harm caused by misclassification in high-risk domains has significant societal implications. Zero-shot models are increasingly being applied in safety-critical fields such

as medical diagnosis, autonomous driving and security surveillance. The inherent uncertainty in identifying unseen classes makes misclassification inevitable. In these scenarios, false negatives (failure to detect rare conditions) or false positives (labelling benign anomalies as threats) can both lead to serious consequences. Fortunately, the nature of element-wise reasoning offers potential advantages for traceability. When a model misclassifies a sample, it may be possible to trace the error back to the specific attribute or state detector that triggered the misclassification. This interpretability holds promise for enhancing the transparency and accountability of AI.

9. Conclusions

We provide a comprehensive exposition of three major tasks in the field of zero-shot image recognition that share the notion of element-wise representation and reasoning. Specifically, we show the background, commonalities, and challenges faced by several tasks. Then related recent techniques are meticulously categorized and explained. At the same time, we provide specific technical details, including datasets, training details, code repository addresses, and more. Finally, we show the wide range of applications of element-wise zero-shot learning and look at future directions.

Funding

This research was funded by Hong Kong RGC General Research Fund (Grant Nos. 15221123, 15216424, and 15211525) and the Hong Kong PolyU Internal Research Fund (Grant Nos. P0058468 and P0056171).

Data Availability Statement

The associated library, including benchmarks, technical implementations, standardized datasets, and some more details, is available at <https://github.com/eigenailab/Element-Wise-Zero-Shot-Image-Recognition>.

Conflicts of Interest

Given the role as an Executive Editor-in-Chief of *Edge Intelligence and Systems*, Jingcai Guo had no involvement in the peer review of this paper and had no access to information regarding its peer-review process. Full responsibility for the editorial process of this paper was delegated to another editor of the journal. The Hong Kong RGC General Research Fund and the Hong Kong PolyU Internal Research Fund had involvement in all stages of the research process. The authors take full responsibility for the content of the published article.

Use of AI and AI-Assisted Technologies

No AI tools were utilized for this paper.

References

1. Deng, J.; Dong, W.; Socher, R.; et al. ImageNet: A large-scale hierarchical image database. In Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL, USA, 20–25 June 2009; pp. 248–255.
2. Chen, C.-Y.; Grauman, K. Inferring analogous attributes. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 200–207.
3. Misra, I.; Gupta, A.; Hebert, M. From red wine to red tomato: Composition with context. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 1792–1801.
4. Yang, Z.; Luo, T.; Wang, D.; et al. Learning to navigate for fine-grained classification. In Proceedings of the 15th European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 420–435.
5. Akata, Z.; Reed, S.; Walter, D.; et al. Evaluation of output embeddings for fine-grained image classification. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 2927–2936.
6. Nagarajan, T.; Grauman, K. Attributes as operators: factorizing unseen attribute-object compositions. In Proceedings of the 15th European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 169–185.
7. Wah, C.; Branson, S.; Welinder, P.; et al. *The Caltech-Ucsd Birds-200-2011 Dataset*; California Institute of Technology: Pasadena, CA, USA, 2011.
8. Patterson, G.; Hays, J. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, 16–21 June 2012; pp. 2751–2758.
9. Wei, X.-S.; Song, Y.-Z.; Aodha, O.M.; et al. Fine-grained image analysis with deep learning: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *44*, 8927–8948.

10. Wang, J.; Cheng, Y.; Feris, R.S. Walk and learn: Facial attribute representation learning from egocentric video and contextual data. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 2295–2304.
11. Kovashka, A.; Parikh, D.; Grauman, K. Whittlesearch: Image search with relative attribute feedback. In Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, 16–21 June 2012; pp. 2973–2980.
12. Wu, Q.; Shen, C.; Wang, P.; et al. Image captioning and visual question answering based on attributes and external knowledge. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *40*, 1367–1381.
13. Cruz, R.S.; Fernando, B.; Cherian, A.; et al. Neural algebra of classifiers. In Proceedings of the 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), Lake Tahoe, NV, USA, 12–15 March 2018; pp. 729–737.
14. Biederman, I. Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* **1987**, *94*, 115.
15. Hoffman, D.D.; Richards, W.A. Parts of recognition. In *Readings in Computer Vision*; Elsevier: Amsterdam, The Netherlands, 1987; pp. 227–242.
16. Felzenszwalb, P.; McAllester, D.; Ramanan, D. A discriminatively trained, multiscale, deformable part model. In Proceedings of the 2008 IEEE Conference on Computer Vision and Pattern Recognition, Anchorage, AK, USA, 23–28 June 2008; pp. 1–8.
17. Battaglia, P.W.; Hamrick, J.B.; Bapst, V.; et al. Relational inductive biases, deep learning, and graph networks. *arXiv* **2018**, arXiv:1806.01261.
18. Lampert, C.H.; Nickisch, H.; Harmeling, S. Learning to detect unseen object classes by between-class attribute transfer. In Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL, USA, 20–25 June 2009; pp. 951–958.
19. Wang, Y.; Yao, Q.; Kwok, J.T.; et al. Generalizing from a few examples: A survey on few-shot learning. *ACM Comput. Surv.* **2020**, *53*, 63.
20. O'Mahony, N.; Campbell, S.; Carvalho, A.; et al. One-shot learning for custom identification tasks; a review. *Procedia Manuf.* **2019**, *38*, 186–193.
21. Yang, J.; Zhou, K.; Li, Y.; et al. Generalized out-of-distribution detection: A survey. *Int. J. Comput. Vis.* **2024**, *132*, 5635–5662.
22. Xian, Y.; Lampert, C.H.; Schiele, B.; et al. Zero-shot learning—A comprehensive evaluation of the good, the bad and the ugly. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *41*, 2251–2265.
23. Wang, W.; Zheng, V.W.; Yu, H.; et al. A survey of zero-shot learning: Settings, methods, and applications. *ACM Trans. Intell. Syst. Technol.* **2019**, *10*, 13.
24. Pourpanah, F.; Abdar, M.; Luo, Y.; et al. A review of generalized zero-shot learning methods. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**, *45*, 4051–4070.
25. Lake, B.M.; Ullman, T.D.; Tenenbaum, J.B.; et al. Building machines that learn and think like people. *Behav. Brain Sci.* **2017**, *40*, e253.
26. Purushwalkam, S.; Nickel, M.; Gupta, A.; et al. Task-driven modular networks for zero-shot compositional learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Republic of Korea, 27 October–2 November 2019; pp. 3593–3602.
27. Mancini, M.; Naeem, M.F.; Xian, Y.; et al. Open world compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 5222–5230.
28. Radford, A.; Kim, J.W.; Hallacy, C.; et al. Learning transferable visual models from natural language supervision. In Proceedings of the International Conference on Machine Learning, Virtual, 18–24 July 2021; pp. 8748–8763.
29. Jia, C.; Yang, Y.; Xia, Y.; et al. Scaling up visual and vision-language representation learning with noisy text supervision. In Proceedings of the International Conference on Machine Learning, Virtual, 18–24 July 2021; pp. 4904–4916.
30. Zhou, K.; Yang, J.; Loy, C.C.; et al. Learning to prompt for vision-language models. *Int. J. Comput. Vis.* **2022**, *130*, 2337–2348.
31. Zhou, K.; Yang, J.; Loy, C.C.; et al. Conditional prompt learning for vision-language models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 16816–16825.
32. Menon, S.; Vondrick, C. Visual classification via description from large language models. In Proceedings of the International Conference on Learning Representations (ICLR), Virtual, 25–29 April 2022.
33. Guo, Z.; Zhang, R.; Qiu, L.; et al. CALIP: Zero-shot enhancement of clip with parameter-free attention. In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, 7–14 February 2023; Volume 37, pp. 746–754.
34. Zhu, Y.; Xie, J.; Tang, Z.; et al. Semantic-guided multi-attention localization for zero-shot learning. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019; Volume 32.
35. Xie, G.-S.; Liu, L.; Jin, X.; et al. Attentive region embedding network for zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 9384–9393.

36. Xie, G.-S.; Liu, L.; Zhu, F.; et al. Region graph embedding network for zero-shot learning. In Proceedings of the European Conference on Computer Vision (ECCV), Glasgow, UK, 23–28 August 2020; pp. 562–580.
37. Guo, J.; Guo, S.; Zhou, Q.; et al. Graph knows unknowns: Reformulate zero-shot learning as sample-level graph recognition. In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, 7–14 February 2023; Volume 37, pp. 7775–7783.
38. Huynh, D.; Elhamifar, E. Fine-grained generalized zero-shot learning via dense attribute-based attention. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020; pp. 4483–4493.
39. Chen, S.; Hong, Z.; Liu, Y.; et al. TRANSZERO: Attribute-guided transformer for zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtual, 22 February–1 March 2022; Volume 36, pp. 330–338.
40. Liu, Z.; Luo, P.; Qiu, S.; et al. DeepFashion: Powering robust clothes recognition and retrieval with rich annotations. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 1096–1104.
41. Atzmon, Y.; Kreuk, F.; Shalit, U.; et al. A causal view of compositional zero-shot recognition. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–12 December 2020; Volume 33, pp. 1462–1473.
42. Fu, Y.; Hospedales, T.M.; Xiang, T.; et al. Transductive multi-view zero-shot learning. *IEEE Trans. Pattern Anal. Mach. Intell.* **2015**, *37*, 2332–2345.
43. Jiang, H.; Wang, R.; Shan, S.; et al. Transferable contrastive network for generalized zero-shot learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Republic of Korea, 27 October–2 November 2019; pp. 9765–9774.
44. Zhou, K.; Liu, Z.; Qiao, Y.; et al. Domain generalization: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**, *45*, 4396–4415.
45. Ganin, Y.; Lempitsky, V. Unsupervised domain adaptation by backpropagation. In Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015; pp. 1180–1189.
46. Rao, Z.; Guo, J.; Tang, L.; et al. SRCD: Semantic reasoning with compound domains for single-domain generalized object detection. *IEEE Trans. Neural Netw. Learn. Syst.* **2025**, *36*, 12497–12506.
47. Ge, J.; Xie, H.; Min, S.; et al. Dual part discovery network for zero-shot learning. In Proceedings of the 30th ACM International Conference on Multimedia, Lisboa, Portugal, 10–14 October 2022; pp. 3244–3252.
48. Wei, K.; Yang, M.; Wang, H.; et al. Adversarial fine-grained composition learning for unseen attribute-object recognition. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Republic of Korea, 27 October–2 November 2019; pp. 3741–3749.
49. Romera-Paredes, B.; Torr, P. An embarrassingly simple approach to zero-shot learning. In Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015; pp. 2152–2161.
50. Xian, Y.; Lorenz, T.; Schiele, B.; et al. Feature generating networks for zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 5542–5551.
51. Verma, V.K.; Arora, G.; Mishra, A.; et al. Generalized zero-shot learning via synthesized examples. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 4281–4289.
52. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; et al. Generative adversarial networks. *Commun. ACM* **2020**, *63*, 139–144.
53. Chen, Z.; Li, J.; Luo, Y.; et al. CANZSL: Cycle-consistent adversarial networks for zero-shot learning from natural language. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Snowmass Village, CO, USA, 1–5 March 2020; pp. 874–883.
54. Chen, Z.; Wang, S.; Li, J.; et al. Rethinking generative zero-shot learning: An ensemble learning perspective for recognising visual patches. In Proceedings of the 28th ACM International Conference on Multimedia, Seattle, WA, USA, 12–16 October 2020; pp. 3413–3421.
55. Kingma, D.P.; Welling, M. Auto-encoding variational bayes. In Proceedings of the International Conference on Learning Representations (ICLR), Banff, AB, Canada, 14–16 April 2014.
56. Chen, Z.; Luo, Y.; Qiu, R.; et al. Semantics disentangling for generalized zero-shot learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, QC, Canada, 11–17 October 2021; pp. 8712–8720.
57. Chen, Z.; Luo, Y.; Wang, S.; et al. GSMFlow: Generation shifts mitigating flow for generalized zero-shot learning. *IEEE Trans. Multimed.* **2022**, *25*, 5374–5385.
58. Chen, Z.; Luo, Y.; Wang, S.; et al. Mitigating generation shifts for generalized zero-shot learning. In Proceedings of the 29th ACM International Conference on Multimedia, Virtual, 20–24 October 2021; pp. 844–852.
59. Liu, S.; Long, M.; Wang, J.; et al. Generalized zero-shot learning with deep calibration network. In Proceedings of the Advances in Neural Information Processing Systems, Montréal, QC, Canada, 3–8 December 2018; Volume 31.
60. Wang, Z.; Gou, Y.; Li, J.; et al. Region semantically aligned network for zero-shot learning. In Proceedings of the 30th ACM International Conference on Information & Knowledge Management, Virtual, 1–5 November 2021; pp. 2080–2090.

61. Ji, Z.; Fu, Y.; Guo, J.; et al. Stacked semantics-guided attention model for fine-grained zero-shot learning. In Proceedings of the Advances in Neural Information Processing Systems, Montréal, QC, Canada, 3–8 December 2018; Volume 31.
62. Elhoseiny, M.; Zhu, Y.; Zhang, H.; et al. Link the head to the “beak”: Zero shot learning from noisy text description at part precision. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 5640–5649.
63. Zhu, Y.; Elhoseiny, M.; Liu, B.; et al. A generative adversarial approach for zero-shot learning from noisy texts. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 1004–1013.
64. Li, Y.; Zhang, J.; Zhang, J.; et al. Discriminative learning of latent features for zero-shot recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 7463–7471.
65. Ge, J.; Xie, H.; Min, S.; et al. Semantic-guided reinforced region embedding for generalized zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtual, 2–9 February 2021; Volume 35, pp. 1406–1414.
66. Li, Y.; Liu, Z.; Yao, L.; et al. An entropy-guided reinforced partial convolutional network for zero-shot learning. *IEEE Trans. Circuits Syst. Video Technol.* **2022**, *32*, 5175–5186.
67. Liu, M.; Zhang, C.; Bai, H.; et al. Part-object progressive refinement network for zero-shot learning. *IEEE Trans. Image Process.* **2024**, *33*, 2032–2043.
68. Chen, S.; Hou, W.; Khan, S.; et al. Progressive semantic-guided vision transformer for zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 23964–23974.
69. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; et al. An image is worth 16x16 words: Transformers for image recognition at scale. In Proceedings of the International Conference on Learning Representations (ICLR), Vienna, Austria, 3–7 May 2021.
70. Xu, W.; Xian, Y.; Wang, J.; et al. VGSE: Visually-grounded semantic embeddings for zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 9316–9325.
71. Chen, Z.; Zhao, Z.; Guo, J.; et al. SVIP: Semantically contextualized visual patches for zero-shot learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Honolulu, HI, USA, 19–20 October 2025.
72. Kipf, T.N.; Welling, M. Semi-supervised classification with graph convolutional networks. In Proceedings of the International Conference on Learning Representations (ICLR), Toulon, France, 26–24 April 2017.
73. Hu, Z.; Zhao, H.; Peng, J.; et al. Region interaction and attribute embedding for zero-shot learning. *Inf. Sci.* **2022**, *609*, 984–995.
74. Chen, S.; Hong, Z.; Xie, G.; et al. GNDAN: Graph navigated dual attention network for zero-shot learning. *IEEE Trans. Neural Netw. Learn. Syst.* **2022**, *35*, 4516–4529.
75. Vaswani, A.; Shazeer, N.; Parmar, N.; et al. Attention is all you need. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; Volume 30.
76. Liu, Y.; Zhou, L.; Bai, X.; et al. Goal-oriented gaze estimation for zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 3794–3803.
77. Huynh, D.; Elhamifar, E. Compositional zero-shot learning via fine-grained dense feature composition. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–12 December 2020; Volume 33, pp. 19849–19860.
78. Liu, Y.; Dang, Y.; Gao, X.; et al. Zero-shot learning with attentive region embedding and enhanced semantics. *IEEE Trans. Neural Netw. Learn. Syst.* **2022**, *35*, 4220–4231.
79. Naeem, M.F.; Xian, Y.; Gool, L.V.; et al. I2DFORMER: Learning image to document attention for zero-shot image classification. In Proceedings of the Advances in Neural Information Processing Systems, New Orleans, LA, USA, 28 November–9 December 2022; Volume 35, pp. 12283–12294.
80. Chen, Z.; Huang, Y.; Chen, J.; et al. DUET: Cross-modal semantic grounding for contrastive zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, 7–14 February 2023; Volume 37, pp. 405–413.
81. Liu, M.; Li, F.; Zhang, C.; et al. Progressive semantic-visual mutual adaption for generalized zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 15337–15346.
82. Chen, S.; Hong, Z.; Xie, G.-S.; et al. MSDN: Mutually semantic distillation network for zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 7612–7621.
83. Chen, S.; Chen, S.; Xie, G.-S.; et al. Mutually causal semantic distillation network for zero-shot learning. *Int. J. Comput. Vis.* **2026**, *134*, 246.

84. Jiang, H.; Li, Z.; Yu, X.; et al. Visual and semantic prompt collaboration for generalized zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 11–15 June 2025; pp. 20275–20285.
85. Hou, W.; Fu, D.; Li, K.; et al. Zeromamba: Exploring visual state space model for zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Philadelphia, PA, USA, 25 February–4 March 2025; Volume 39, pp. 3527–3535.
86. Rao, Z.; Guo, J. Balancing cross-modal attention for generalized zero-shot learning. In Proceedings of the 33rd ACM International Conference on Multimedia, Dublin, Ireland, 27–31 October 2025; pp. 3360–3369.
87. Yang, H.-M.; Zhang, X.-Y.; Yin, F.; et al. Robust classification with convolutional prototype learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 3474–3482.
88. Liu, J.; Ren, Y.; Li, W.; et al. Improving out-of-distribution detection with margin-based prototype learning. In Proceedings of the International Conference on Neural Information Processing, Changsha, China, 20–23 November 2023; pp. 149–160.
89. Wang, K.; Liew, J.H.; Zou, Y.; et al. PANET: Few-shot image semantic segmentation with prototype alignment. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Republic of Korea, 27 October–2 November 2019; pp. 9197–9206.
90. Xu, W.; Xian, Y.; Wang, J.; et al. Attribute prototype network for zero-shot learning. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–12 December 2020; Volume 33, pp. 21969–21980.
91. Jayaraman, D.; Sha, F.; Grauman, K. Decorrelating semantic visual attributes by resisting the urge to share. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 1629–1636.
92. Wang, C.; Min, S.; Chen, X.; et al. Dual progressive prototype network for generalized zero-shot learning. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021; Volume 34, pp. 2936–2948.
93. Yue, Q.; Cui, J.; Liang, J.; et al. Class semantic attribute perception guided zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Philadelphia, PA, USA, 25 February–4 March 2025; Volume 39, pp. 22281–22289.
94. Guo, T.; Liang, J.; Xie, G.-S. Group-wise interactive region learning for zero-shot recognition. *Inf. Sci.* **2023**, *642*, 119135.
95. Cheng, D.; Wang, G.; Wang, N.; et al. Discriminative and robust attribute alignment for zero-shot learning. *IEEE Trans. Circuits Syst. Video Technol.* **2023**, *33*, 4244–4256.
96. Du, Y.; Shi, M.; Wei, F.; et al. Boosting zero-shot learning via contrastive optimization of attribute representations. *IEEE Trans. Neural Netw. Learn. Syst.* **2023**, *35*, 16706–16719.
97. Chen, Z.; Zhang, P.; Li, J.; et al. Zero-shot learning by harnessing adversarial samples. In Proceedings of the 31st ACM International Conference on Multimedia, Ottawa, ON, Canada, 29 October–3 November 2023; pp. 4138–4146.
98. Liu, Y.; Guo, J.; Cai, D.; et al. Attribute attention for semantic disambiguation in zero-shot learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Seoul, Republic of Korea, 27 October–2 November 2019; pp. 6698–6707.
99. Liu, L.; Zhou, T.; Long, G.; et al. Attribute propagation network for graph zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; Volume 34, pp. 4868–4875.
100. Akata, Z.; Malinowski, M.; Fritz, M.; et al. Multi-cue zero-shot learning with strong supervision. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 59–68.
101. Rao, Z.; Guo, J.; Lu, X.; et al. Dual expert distillation network for generalized zero-shot learning. In Proceedings of the International Joint Conference on Artificial Intelligence, Jeju, Republic of Korea, 3–9 August 2024.
102. Nan, Z.; Liu, Y.; Zheng, N.; et al. Recognizing unseen attribute-object pair with generative model. In Proceedings of the AAAI Conference on Artificial Intelligence, Honolulu, HI, USA, 27 January–1 February 2019; Volume 33, pp. 8811–8818.
103. Yuan, Z.; Wang, Z.; Pan, Y.; et al. I2CD: An invertible causal framework for compositional zero-shot learning via disentangle-compose-disentangle. In Proceedings of the AAAI Conference on Artificial Intelligence, Singapore, 20–27 January 2026; Volume 40, pp. 12295–12303.
104. Ruis, F.; Burghouts, G.; Bucur, D. Independent prototype propagation for zero-shot compositionality. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021; Volume 34, pp. 10641–10653.
105. Hu, X.; Wang, Z. Leveraging sub-class discrimination for compositional zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, 7–14 February 2023, Volume 37, pp. 890–898.
106. Zheng, Z.; Zhu, H.; Nevatia, R. CAILA: Concept-aware intra-layer adapters for compositional zero-shot learning. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 3–8 January 2024; pp. 1721–1731.
107. Yang, M.; Xu, C.; Wu, A.; et al. A decomposable causal view of compositional zero-shot learning. *IEEE Trans. Multimed.* **2022**, *25*, 5892–5902.

108. Rao, Z.; Guo, J.; Li, M.; et al. Exploring transferable homogeneous groups for compositional zero-shot learning. In Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-25), Montreal, QC, Canada, 16–22 August, 2025.
109. Saini, N.; Pham, K.; Shrivastava, A. Disentangling visual embeddings for attributes and objects. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 13658–13667.
110. Hao, S.; Han, K.; Wong, K.-Y.K. Learning attention as disentangler for compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 15 315–15 324.
111. Jing, C.; Li, Y.; Chen, H.; et al. Retrieval-augmented primitive representations for compositional zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 26–27 February 2024; Volume 38, pp. 2652–2660.
112. Zhang, T.; Liang, K.; Du, R.; et al. Learning invariant visual representations for compositional zero-shot learning. In Proceedings of the European Conference on Computer Vision (ECCV), Tel Aviv, Israel, 23–27 October 2022; pp. 339–355.
113. Lu, X.; Guo, S.; Liu, Z.; et al. Decomposed soft prompt guided fusion enhancing for compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 23 560–23 569.
114. Huang, S.; Gong, B.; Feng, Y.; et al. TROIKA: Multi-path cross-modal traction for compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 24 005–24 014.
115. Panda, A.; Mukherjee, D.P. Compositional zero-shot learning using multi-branch graph convolution and cross-layer knowledge sharing. *Pattern Recognit.* **2024**, *145*, 109916.
116. Kim, H.; Lee, J.; Park, S.; et al. Hierarchical visual primitive experts for compositional zero-shot learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 5675–5685.
117. Li, Y.-L.; Xu, Y.; Mao, X.; et al. Symmetry and group in attribute- object compositions. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020; pp. 11316–11325.
118. Wang, Q.; Liu, L.; Jing, C.; et al. Learning conditional attributes for compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 11197–11206.
119. Li, L.; Chen, G.; Wang, Z.; et al. Compositional zero-shot learning via progressive language-based observations. In Proceedings of the 33rd ACM International Conference on Multimedia, Dublin, Ireland, 27–31 October 2025; pp. 3827–3836.
120. Liu, Z.; Li, Y.; Yao, L.; et al. Simple primitives with feasibility-and contextuality-dependence for open-world compositional zero-shot learning. *IEEE Trans. Pattern Anal. Mach. Intell.* **2023**, *46*, 543–560.
121. Huo, F.; Xu, W.; Guo, S.; et al. PROCC: Progressive cross-primitive compatibility for open-world compositional zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 26–27 February 2024; Volume 38, pp. 12689–12697.
122. Xu, Z.; Wang, G.; Wong, Y.; et al. Relation-aware compositional zero-shot learning for attribute-object pair recognition. *IEEE Trans. Multimed.* **2021**, *24*, 3652–3664.
123. Khan, M.G.Z.A.; Naeem, M.F.; Van Gool, L.; et al. Learning attention propagation for compositional zero-shot learning. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 2–7 January 2023; pp. 3828–3837.
124. Naeem, M.F.; Xian, Y.; Tombari, F.; et al. Learning graph embeddings for compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 953–962.
125. Anwaar, M.U.; Pan, Z.; Kleinsteuber, M. On leveraging variational graph embeddings for open world compositional zero-shot learning. In Proceedings of the 30th ACM International Conference on Multimedia, Lisboa, Portugal, 10–14 October 2022; pp. 4645–4654.
126. Xu, G.; Chai, J.; Kordjamshidi, P. GIPCOL: Graph-injected soft prompting for compositional zero-shot learning. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 3–8 January 2024; pp. 5774–5783.
127. Jaiswal, A.; Babu, A.R.; Zadeh, M.Z.; et al. A survey on contrastive self-supervised learning. *Technologies* **2020**, *9*, 2.
128. Li, X.; Yang, X.; Wei, K.; et al. Siamese contrastive embedding network for compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 9326–9335.
129. Yang, Y.; Pan, R.; Li, X.; et al. Dual-stream contrastive learning for compositional zero-shot recognition. *IEEE Trans. Multimed.* **2023**, *26*, 1909–1919.

130. Wang, X.; Yu, F.; Wang, R.; et al. Tafe-Net: Task-aware feature embeddings for low shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 1831–1840.
131. Jiang, C.; Ye, Q.; Wang, S.; et al. Mutual balancing in state-object components for compositional zero-shot learning. *Pattern Recognit.* **2024**, *152*, 110451.
132. Jiang, C.; Zhang, H. Revealing the proximate long-tail distribution in compositional zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 26–27 February 2024; Volume 38, pp. 2498–2506.
133. Li, Y.; Liu, Z.; Chen, H.; et al. Context-based and diversity-driven specificity in compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 17037–17046.
134. Li, M.; Guo, J.; Da Xu, R.Y.; et al. TSCA: On the semantic consistency alignment via conditional transport for compositional zero-shot learning. In Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-25), Montreal, QC, Canada, 16–22 August 2025.
135. Wu, P.; Lu, X.; Hu, H.; et al. Logiczsl: Exploring logic-induced representation for compositional zero-shot learning. In Proceedings of the Computer Vision and Pattern Recognition Conference, Nashville, TN, USA, 11–15 June 2025; pp. 30301–30311.
136. Karthik, S.; Mancini, M.; Akata, Z. KG-SP: Knowledge guided simple primitives for open world compositional zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 9336–9345.
137. Hu, X.; Wang, Z. A dynamic learning method towards realistic compositional zero-shot learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 26–27 February 2024; Volume 38, pp. 2265–2273.
138. Zhang, J.; Huang, J.; Jin, S.; et al. Vision-language models for vision tasks: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2024**, *46*, 5625–5644.
139. Cui, Q.; Zhou, B.; Guo, Y.; et al. Contrastive vision-language pre-training with limited resources. In Proceedings of the European Conference on Computer Vision (ECCV), Tel Aviv, Israel, 23–27 October 2022; pp. 236–253.
140. Singh, A.; Hu, R.; Goswami, V.; et al. FLAVA: A foundational language and vision alignment model. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 15638–15650.
141. Huang, R.; Long, Y.; Han, J.; et al. NLIP: Noise-robust language-image pre-training. In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, 7–14 February 2023; Volume 37, pp. 926–934.
142. Esmailpour, S.; Liu, B.; Robertson, E.; et al. Zero-shot out-of-distribution detection based on the pre-trained model clip. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtual, 22 February–1 March 2022; Volume 36, pp. 6568–6576.
143. Wang, H.; Li, Y.; Yao, H.; et al. CLIPN for zero-shot OOD detection: Teaching clip to say no. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 1802–1812.
144. Brown, T.; Mann, B.; Ryder, N.; et al. Language models are few-shot learners. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–12 December 2020; Volume 33, pp. 1877–1901.
145. Pratt, S.; Covert, I.; Liu, R.; et al. What does a platypus look like? generating customized prompts for zero-shot image classification. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 15691–15701.
146. Novack, Z.; McAuley, J.; Lipton, Z.C.; et al. CHILS: Zero-shot image classification with hierarchical label sets. In Proceedings of the International Conference on Machine Learning, Honolulu, HI, USA, 23–29 July 2023; pp. 26342–26362.
147. Chen, Z. Don't paint everyone with the same brush: Adaptive prompt prototype learning for vision-language models. In Proceedings of the Twelfth International Conference on Learning Representations, Vienna, Austria, 7–11 May 2024.
148. Udandarao, V.; Gupta, A.; Albanie, S. SUS-X: Training-free name-only transfer of vision-language models. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 2725–2736.
149. Rombach, R.; Blattmann, A.; Lorenz, D.; et al. High-resolution image synthesis with latent diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 10 684–10 695.
150. Chen, Y.; Zhang, Q.; Shi, X.; et al. PromptMix: Llm-aided prompt learning for generalizing vision-language models. *Inf. Fusion* **2026**, *131*, 104186.
151. Shu, M.; Nie, W.; Huang, D.-A.; et al. Test-time prompt tuning for zero-shot generalization in vision-language models. In Proceedings of the Advances in Neural Information Processing Systems, New Orleans, LA, USA, 28 November–9 December 2022; Volume 35, pp. 14274–14289.
152. Feng, C.-M.; Yu, K.; Liu, Y.; et al. Diverse data augmentation with diffusions for effective test-time prompt tuning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 2704–2714.

153. Samadh, J.A.; Gani, M.H.; Hussein, N.; et al. Align your prompts: Test-time prompting with distribution alignment for zero-shot generalization. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; Volume 36.
154. Ma, X.; Zhang, J.; Guo, S.; et al. Swapprompt: Test-time prompt adaptation for vision-language models. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; Volume 36.
155. Liu, Z.; Sun, H.; Peng, Y.; et al. DART: Dual-modal adaptive online prompting and knowledge retention for test-time adaptation. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 26–27 February 2024; Volume 38, pp. 14106–14114.
156. Zhang, Y.; Zhu, W.; Tang, H.; et al. Dual memory networks: A versatile adaptation approach for vision-language models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 28718–28728.
157. Karmanov, A.; Guan, D.; Lu, S.; et al. Efficient test-time adaptation of vision-language models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 14162–14171.
158. Zhang, D.-C.; Zhou, Z.; Li, Y.-F. Robust test-time adaptation for zero-shot prompt tuning. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 26–27 February 2024; Volume 38, pp. 16714–16722.
159. Zanella, M.; Ayed, I.B. On the test-time zero-shot generalization of vision-language models: Do we really need prompt learning? In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 23783–23793.
160. Liang, Y.; Chen, H.; Xiong, Y.; et al. Advancing reliable test-time adaptation of vision-language models under visual variations. In Proceedings of the 33rd ACM International Conference on Multimedia, Dublin, Ireland, 27–31 October 2025; pp. 4788–4797.
161. Zhu, X.; Wang, S.; Zhu, B.; et al. Dynamic multimodal prototype learning in vision-language models. In Proceedings of the IEEE/CVF international conference on computer vision, Honolulu, HI, USA, 19–20 October 2025; pp. 2501–2511.
162. Huang, T.; Chu, J.; Wei, F. Unsupervised prompt learning for vision-language models. *arXiv* **2022**, arXiv:2204.03649.
163. Li, J.; Savarese, S.; Hoi, S.C. Masked unsupervised self-training for label-free image classification. In Proceedings of the International Conference on Learning Representations (ICLR), Kigali, Rwanda, 1–5 May 2023.
164. Qian, Q.; Xu, Y.; Hu, J. Intra-modal proxy learning for zero-shot visual categorization with clip. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; Volume 36.
165. Mirza, M.J.; Karlinsky, L.; Lin, W.; et al. LAFTER: Label-free tuning of zero-shot classifier using language and unlabeled image collections. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; Volume 36.
166. Kalantidis, Y.; Tolias, G. Label propagation for zero-shot classification with vision-language models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 23209–23218.
167. Kahana, J.; Cohen, N.; Hoshen, Y. Improving zero-shot models with label distribution priors. *arXiv* **2022**, arXiv:2212.00784.
168. Lampert, C.H.; Nickisch, H.; Harmeling, S. Attribute-based classification for zero-shot visual object categorization. *IEEE Trans. Pattern Anal. Mach. Intell.* **2013**, *36*, 453–465.
169. Nilsback, M.-E.; Zisserman, A. Automated flower classification over a large number of classes. In Proceedings of the 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing, Bhubaneswar, India, 16–19 December 2008; pp. 722–729.
170. Van Horn, G.; Branson, S.; Farrell, R.; et al. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 595–604.
171. Helber, P.; Bischke, B.; Dengel, A.; et al. EUROSAT: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2019**, *12*, 2217–2226.
172. Krizhevsky, A.; Hinton, G. Learning multiple layers of features from tiny images. 2009. Available online: <https://cave.cs.toronto.edu/kriz/learning-features-2009-TR.pdf> (accessed on 24 July 2024).
173. Bossard, L.; Guillaumin, M.; Van Gool, L. Food-101—mining discriminative components with random forests. In Proceedings of the European Conference on Computer Vision (ECCV), Zurich, Switzerland, 6–12 September 2014; pp. 446–461.
174. Parkhi, O.M.; Vedaldi, A.; Zisserman, A.; et al. Cats and dogs. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Providence, RI, USA, 16–21 June 2012; pp. 3498–3505.
175. Krause, J.; Deng, J.; Stark, M.; et al. Collecting a Large-Scale Dataset of Fine-Grained Cars. 2013. Available online: <https://ai.stanford.edu/~jkrause/papers/fgvc13.pdf> (accessed on 23 July 2024).
176. Maji, S.; Rahtu, E.; Kannala, J.; et al. Fine-grained visual classification of aircraft. *arXiv* **2013**, arXiv:1306.5151.
177. Cimpoi, M.; Maji, S.; Kokkinos, I.; et al. Describing textures in the wild. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 3606–3613.

178. Farhadi, A.; Endres, I.; Hoiem, D.; et al. Describing objects by their attributes. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Miami, FL, USA, 20–25 June 2009; pp. 1778–1785.
179. Yu, A.; Grauman, K. Fine-grained visual comparisons with local learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 192–199.
180. Isola, P.; Lim, J.J.; Adelson, E.H. Discovering states and transformations in image collections. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 1383–1391.
181. Fei-Fei, L.; Fergus, R.; Perona, P. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In Proceedings of the 2004 Conference on Computer Vision and Pattern Recognition Workshop, Washington, DC, USA, 27 June–2 July 2004; p. 178.
182. Griffin, G.; Holub, A.; Perona, P.; et al. *Caltech-256 Object Category Dataset*; Technical Report 7694; California Institute of Technology Pasadena: Pasadena CA, USA, 2007.
183. Xiao, J.; Hays, J.; Ehinger, K.A.; et al. Sun database: Large-scale scene recognition from abbey to zoo. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), San Francisco, CA, USA, 13–18 June 2010; pp. 3485–3492.
184. Berg, T.; Liu, J.; Lee, S.W.; et al. BIRDSNAP: Large-scale fine-grained visual categorization of birds. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 2011–2018.
185. He, K.; Zhang, X.; Ren, S.; et al. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
186. Simonyan, K.; Zisserman, A. Very deep convolutional networks for large-scale image recognition. In Proceedings of the International Conference on Learning Representations (ICLR), San Diego, CA, USA, 7–9 May 2015.
187. Szegedy, C.; Liu, W.; Jia, Y.; et al. Going deeper with convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 1–9.
188. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. In Proceedings of the Advances in Neural Information Processing Systems, Lake Tahoe, NV, USA, 3–6 December 2012; Volume 25.
189. Chen, X.; Deng, X.; Lan, Y.; et al. Explanatory object part aggregation for zero-shot learning. *IEEE Trans. Pattern Anal. Mach. Intell.* **2023**, *46*, 851–868.
190. Wang, Z.; Gou, Y.; Li, J.; et al. Language-augmented pixel embedding for generalized zero-shot learning. *IEEE Trans. Circuits Syst. Video Technol.* **2022**, *33*, 1019–1030.
191. Cheng, D.; Wang, G.; Wang, B.; et al. Hybrid routing transformer for zero-shot learning. *Pattern Recognit.* **2023**, *137*, 109270.
192. Li, Y.; Liu, Z.; Jha, S.; et al. Distilled reverse attention network for open-world compositional zero-shot learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 1782–1791.
193. Hou, Y.; Zhang, J.; Lin, Z.; et al. Large language models are zero-shot rankers for recommender systems. In Proceedings of the European Conference on Information Retrieval, Glasgow, UK, 24–28 March 2024; pp. 364–381.
194. Kojima, T.; Gu, S.S.; Reid, M.; et al. Large language models are zero-shot reasoners. In Proceedings of the Advances in Neural Information Processing Systems, New Orleans, LA, USA, 28 November–9 December 2022; Volume 35, pp. 22199–22213.
195. Wei, J.; Bosma, M.; Zhao, V.Y.; et al. Finetuned language models are zero-shot learners. In Proceedings of the International Conference on Learning Representations (ICLR), Virtual, 25–29 April 2022.
196. Johnson, M.; Schuster, M.; Le, Q.V.; et al. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Trans. Assoc. Comput. Linguist.* **2017**, *5*, 339–351.
197. Baek, J.; Aji, A.F.; Saffari, A. Knowledge-augmented language model prompting for zero-shot knowledge graph question answering. In Proceedings of the 1st Workshop on Natural Language Reasoning and Structured Explanations (NLRSE), Toronto, ON, Canada, 13 July 2023.
198. Kumar, P.; Pathania, K.; Raman, B. Zero-shot learning based cross-lingual sentiment analysis for sanskrit text with insufficient labeled data. *Appl. Intell.* **2023**, *53*, 10096–10113.
199. Chen, Q.; Wang, W.; Huang, K.; et al. Zero-shot text classification via knowledge graph embedding for social media data. *IEEE Internet Things J.* **2021**, *9*, 9205–9213.
200. Bari, M.S.; Joty, S.; Jwalapuram, P. Zero-resource cross-lingual named entity recognition. In Proceedings of the AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; Volume 34, pp. 7415–7423.
201. Thakur, N.; Reimers, N.; Rücklé, A.; et al. BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. In Proceedings of the Thirty-Fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track, Virtual, 6–14 December 2021.
202. Huynh, D.; Elhamifar, E. A shared multi-attention framework for multi-label zero-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020; pp. 8776–8786.

203. Ren, S.; He, K.; Girshick, R.; et al. Faster r-CNN: Towards real- time object detection with region proposal networks. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 7–12 December 2015; Volume 28.
204. Liu, W.; Anguelov, D.; Erhan, D.; et al. SSD: Single shot multibox detector. In Proceedings of the European Conference on Computer Vision (ECCV), Amsterdam, The Netherlands, 11–14 October 2016; pp. 21–37.
205. Long, J.; Shelhamer, E.; Darrell, T. Fully convolutional networks for semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 3431–3440.
206. Kirillov, A.; Mintun, E.; Ravi, N.; et al. Segment anything. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 4015–4026.
207. Bansal, A.; Sikka, K.; Sharma, G.; et al. Zero-shot object detection. In Proceedings of the 15th European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 384–400.
208. Rahman, S.; Khan, S.; Porikli, F. Zero-shot object detection: Learning to simultaneously recognize and localize novel concepts. in ACCV. Springer, 2018; pp. 547–563.
209. Li, Z.; Yao, L.; Zhang, X.; et al. Zero-shot object detection with textual descriptions. In Proceedings of the AAAI Conference on Artificial Intelligence, Honolulu, HI, USA, 27 January–1 February 2019; Volume 33, pp. 8690–8697.
210. Gu, Z.; Zhou, S.; Niu, L.; et al. Context-aware feature generation for zero-shot semantic segmentation. In Proceedings of the 28th ACM International Conference on Multimedia, Seattle, WA, USA, 12–16 October 2020; pp. 1921–1929.
211. Huynh, D.; Kuen, J.; Lin, Z.; et al. Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 7020–7031.
212. He, S.; Ding, H.; Jiang, W. Primitive generation and semantic-related alignment for universal zero-shot segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 11 238–11 247.
213. Guo, Y.; Wang, H.; Hu, Q.; et al. Deep learning for 3D point clouds: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2020**, *43*, 4338–4364.
214. Naeem, M.F.; Ornek, E.P.; Xian, Y.; et al. 3D compositional zero-shot learning with compositional consensus. In Proceedings of the European Conference on Computer Vision (ECCV), Tel Aviv, Israel, 23–27 October 2022; pp. 713–730.
215. Cheraghian, A.; Rahman, S.; Chowdhury, T.F.; et al. Zero-shot learning on 3D point cloud objects and beyond. *Int. J. Comput. Vis.* **2022**, *130*, 2364–2384.
216. Zhang, D.; Liang, D.; Yang, H.; et al. SAM3D: Zero-shot 3D object detection via segment anything model. *Sci. China Inf. Sci.* **2024**, *67*, 149101.
217. Abdelreheem, A.; Skorokhodov, I.; Ovsjanikov, M.; et al. SATR: Zero-shot semantic segmentation of 3D shapes. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 15166–15179.
218. Liu, R.; Wu, R.; Van Hoorick, B.; et al. Zero-1-to-3: Zero-shot one image to 3D object. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023.
219. Liu, K.; Zhan, F.; Chen, Y.; et al. Style3D: Zero-shot 3D style transfer of neural radiance fields. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 8338–8348.
220. Xu, J.; Wang, X.; Cheng, W.; et al. DREAM3D: Zero-shot text-to-3D synthesis using 3D shape prior and text-to-image diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 20 908–20 918.
221. Li, L.; Dai, A. Genzi: Zero-shot 3D human-scene interaction generation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 20465–20474.
222. Jiang, Z.; Zhou, Z.; Li, L.; et al. Back to optimization: Diffusion-based zero-shot 3D human pose estimation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, 3–8 January 2024; pp. 6142–6152.
223. Abdelreheem, A.; Eldesokey, A.; Ovsjanikov, M.; et al. Zero-shot 3D shape correspondence. In Proceedings of the SIGGRAPH Asia 2023 Conference Papers, Sydney, NSW, Australia, 12–15 December 2023; pp. 1–11.
224. Javed, S.; Mahmood, A.; Ganapathi, I.I.; et al. CPLIP: Zero-shot learning for histopathology with comprehensive vision-language alignment. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 16–22 June 2024; pp. 11450–11459.
225. Li, X.; Wen, C.; Hu, Y.; et al. RS-Clip: Zero shot remote sensing scene classification via contrastive vision-language supervision. *Int. J. Appl. Earth Obs. Geoinf.* **2023**, *124*, 103497.
226. Li, A.; Qiu, C.; Kloft, M.; et al. Zero-shot anomaly detection via batch normalization. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 10–15 December 2024; Volume 36.
227. Chen, S.; Huang, D. Elaborative rehearsal for zero-shot action recognition. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, QC, Canada, 11–17 October 2021; pp. 13638–13647.

- 228. Song, L.; Shang, X.; Yang, C.; et al. Attribute-guided multiple instance hashing network for cross-modal zero-shot hashing. *IEEE Trans. Multimed.* **2022**, *25*, 5305–5318.
- 229. Wang, S.; Chang, J.; Wang, Z.; et al. Content-aware rectified activation for zero-shot fine-grained image retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.* **2024**, *46*, 4366–4380.
- 230. Jiang, X.; Xu, X.; Zhou, Z.; et al. Zero-shot video moment retrieval with angular reconstructive text embeddings. *IEEE Trans. Multimed.* **2024**, *26*, 9657–9670.
- 231. Wang, Z.; Hu, R.; Liang, C.; et al. Zero-shot person re-identification via cross-view consistency. *IEEE Trans. Multimed.* **2015**, *18*, 260–272.
- 232. Hong, M.; Zhang, X.; Li, G.; et al. Fine-grained feature generation for generalized zero-shot video classification. *IEEE Trans. Image Process.* **2023**, *32*, 1599–1612.
- 233. Wang, L.; Zhang, X.; Su, H.; et al. A comprehensive survey of continual learning: theory, method and application. *IEEE Trans. Pattern Anal. Mach. Intell.* **2024**, *46*, 5362–5383.
- 234. Zhou, F.; Huang, S.; Liu, B.; et al. Multi-label image classification via category prototype compositional learning. *IEEE Trans. Circuits Syst. Video Technol.* **2021**, *32*, 4513–4525.
- 235. Guo, J.; Zhou, Q.; Li, R.; et al. Parsnets: A parsimonious orthogonal and low-rank linear networks for zero-shot learning. In Proceedings of the International Joint Conference on Artificial Intelligence, Macao, China, 19–25 August 2023.
- 236. Rahman, S.; Khan, S.; Porikli, F. A unified approach for conventional zero-shot, generalized zero-shot, and few-shot learning. *IEEE Trans. Image Process.* **2018**, *27*, 5652–5667.
- 237. Shafiee, N.; Elhamifar, E. Zero-shot attribute attacks on fine-grained recognition models. In Proceedings of the European Conference on Computer Vision (ECCV), Tel Aviv, Israel, 23–27 October 2022; pp. 262–282.