

Seeing Unseen: Discover Novel Biomedical Concepts via Geometry-Constrained Probabilistic Modeling

Jianan Fan¹, Dongnan Liu¹, Hang Chang², Heng Huang³, Mei Chen⁴, and Weidong Cai¹

¹University of Sydney ²Lawrence Berkeley National Laboratory ³University of Maryland at College Park ⁴Microsoft
 jfan6480@uni.sydney.edu.au dongnan.liu@sydney.edu.au hchang@lbl.gov
 henghuanghh@gmail.com Mei.Chen@microsoft.com tom.cai@sydney.edu.au

Abstract

Machine learning holds tremendous promise for transforming the fundamental practice of scientific discovery by virtue of its data-driven nature. With the ever-increasing stream of research data collection, it would be appealing to autonomously explore patterns and insights from observational data for discovering novel classes of phenotypes and concepts. However, in the biomedical domain, there are several challenges inherently presented in the cumulated data which hamper the progress of novel class discovery. The non-i.i.d. data distribution accompanied by the severe imbalance among different groups of classes essentially leads to ambiguous and biased semantic representations. In this work, we present a geometry-constrained probabilistic modeling to resolve the identified issues. First, we propose to parameterize the approximated posterior of instance embedding as a marginal von Mises-Fisher distribution to account for the interference of distributional latent bias. Then, we incorporate a suite of critical geometric properties to impose proper constraints on the layout of constructed embedding space, which in turn minimizes the uncontrollable risk for unknown class learning and structuring. Furthermore, a spectral graph-theoretic method is devised to estimate the number of potential novel classes. It inherits two intriguing merits compared to existent approaches, namely high computational efficiency and flexibility for taxonomy-adaptive estimation. Extensive experiments across various biomedical scenarios substantiate the effectiveness and general applicability of our method.

1. Introduction

The rapid progress in deep learning (DL) has invigorated enormous interests for automating essential stages in biomedical image analysis through data-driven methodologies [7, 42]. However, the transformative potential of DL to facilitate scientific discovery at the fundamental level remains largely untapped. Specifically, the success of existing works hinges on expert-curated data annotations in

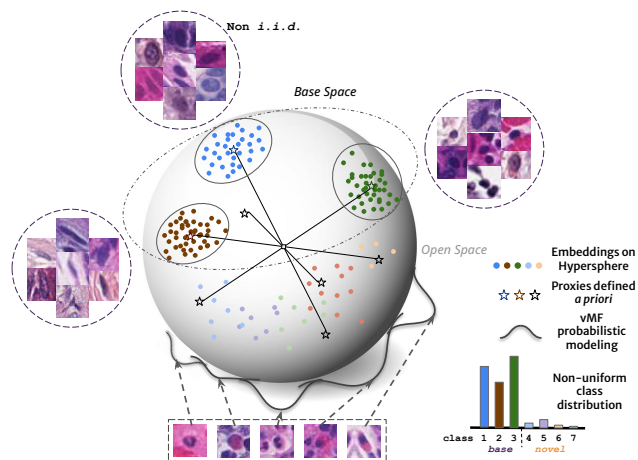


Figure 1. Conceptual illustration of the main insight. In the biomedical domain, violation of the i.i.d. assumption incurred by inconsistent imaging protocols across cohorts and non-uniform class distributions due to scarcity of rare classes could deteriorate the generalizability of learned representations for novel class discovery. We propose to address those issues via probabilistic modeling on a hyperspherical manifold and incorporation of geometrical inductive biases for countering semantic ambiguity and open space risk.

which associated biomedical structures (e.g., cells, tissues) and disease categories are defined a priori [30, 13]. They have little to offer for exploring unknown concepts and discovering expressive profiles for novel/undefined biomedical phenotypes from observational data alone. For instance, cell, the basic unit of human biological system, exhibits vast diversity with hundreds of different types and states [33]. Compared with cumbersome manual pattern identification and grouping of cell populations, autonomously mining new classes of cells with their morphological and functional characteristics could greatly speed up the scientific discovery progress of biomedical research.

To this end, we study the problem of discovering novel concepts from biomedical imaging data, under the *Generalized Novel Class Discovery* setting [44]. Given a reference dataset within a set of well-defined base categories,

it aims at identifying and clustering unseen classes in an external collection of image instances without supervision [22, 20]. The technical challenge of this task lies in how to leverage the structured prior knowledge secured from previously known biomedical object categories to guide the unsupervised discovery of new concepts. Recently, a stream of methods based on representation learning and clustering [16, 21, 44] have been proposed to detect unseen objects in generic visual perception scenes. However, the basic assumption those works built upon may not hold in the biomedical domain from two aspects: **i)** It is generally assumed that the labeled and unlabeled data are independent and identically distributed (i.i.d.) [53, 23], whereas for biomedical images, data distribution shifts inevitably arise from the variations in targeted biology and image acquisition among multiple cohorts [32, 14]. The statistical discrepancy between the reference set of data and the external collection for discovery would consequently result in ambiguous semantic representations given the interference of task-irrelevant attributes [31]. **ii)** Existing works hypothesize a balanced class distribution a priori which expects that the quantity and occurrence frequency of unseen classes should be on par with the base classes [8, 52]. They dismiss the long-tail nature of class distribution in scientific discovery that instances of classes never seen before are essentially rare and hard to collect. For instance, rare diseases are inherently atypical and much fewer in caseload than their common counterpart [9, 34]. The formulated embedding space is consequently prone to be biased towards the base classes with dominatedly copious training samples. Despite the latest efforts delivering explicit modeling of the distribution biases [50, 48], their self-training founded solutions intrinsically fall short due to erratic representation space structuring. Without consideration of the dark information [6] regarding unknown class representations, there is no guarantee that a structured open space can be preserved for learning new classes. It deteriorates the generalizability of those methods to uncover and group the implicit patterns of unseen phenotypes.

In this work, we present a geometry-constrained probabilistic treatment for discovering novel biomedical classes which resolves the above-mentioned challenges from a new perspective, as illustrated in Fig. 1. First, to counteract the semantic ambiguities induced by data distribution shifts, we suggest to circumvent deterministic image encoding and instead model latent representations as directional distributions. Specifically, for each instance, we parameterize the approximated posterior of its semantic embedding with a marginal von Mises-Fisher (vMF) distribution on a unit hyperspherical manifold. The probabilistic modeling presents an elegant solution to decouple latent bias incurred by inconsistent imaging protocols from informative semantic context as an implicit uncertainty measure. Sec-

ond, we pinpoint the necessity to incorporate *inductive biases* for imposing proper constraints on the geometric layout of learned embedding space and thereby minimizing the uncontrollable risk for unknown classes. Aside from the *hyperspherical manifold* prior, we identify two critical geometric properties for embedding space structuring, i.e., *boundness* and *uniformity*. Here, *boundness* implies that, in the latent space, both base and novel classes should be limited to their certain bounded ranges with a large margin separation between the two groups. *Uniformity* refers to the degree of distributional diversity for the set of semantic proxies on the hypersphere. These principles cooperatively shape the layout of embedding space towards risk-bounded novel class discovery and perspicuous geometric interpretation. With explicit constraints to integrate the underlined inductive biases into representation learning, our proposed method inherits guarantees to attain maximal inter-class separability and intra-class compactness, which in turn addresses the inordinate dominance of base classes and favors the generalization of learned embeddings to recognize unseen concepts.

Additionally, with respect to the practice of existing works assuming the number of novel classes are known a priori [54, 29], we argue that such supposition is over-optimistic and unrealistic in the open world of scientific discovery. We therefore propose a versatile and plug-and-play solution to estimate the number of categories in unlabeled data. The spectral graph theory-based approach is computationally efficient and does not introduce recursive clustering as in [44]. It also holds an appealing merit that it supports taxonomy-adaptive estimation, which means the users can decide the granularity level (e.g., coarse or fine-grained) in class hierarchy with regards to their specific needs and get the corresponding outcome.

Contributions. Our contributions are four-fold: **(i)** We introduce a novel paradigm to unleash the tremendous promise of DL to facilitate data-driven biomedical discovery and account for the distribution biases inherently presented in the domain. **(ii)** We propose geometry-constrained probabilistic modeling to facilitate risk-bounded novel class discovery by exploiting the *boundness* and *uniformity* geometrical inductive biases for representation space structuring. We perform theoretical justifications from two perspectives to illustrate the merits of our method. **(iii)** We devise a spectral graph-theoretic method to estimate the potential number of novel classes in unlabeled observational data. **(iv)** The proposed method is extensively validated on a diverse suite of challenging scenarios and attains superior performance compared to existing solutions.

2. Related Works

Novel Class Discovery. Inspired by the observation that human visual system can effortlessly recognize an unseen

class of objects based on previously learnt category concepts, a stream of methods have been proposed for novel class discovery (NCD) in an unsupervised manner by transferring prior knowledge gathered from a set of seen categories with annotations [22, 20, 21]. Follow-up works further identify the limitations of NCD in that it disregards the compositional nature of real-world upcoming data and supposes encountered classes will not appear again in the unlabeled set. To this end, they formalize the *Generalized Novel Class Discovery* (GNCD) setting which assumes both base and novel classes could simultaneously appear in the unlabeled data [44]. ComEx [49] and DPN [1] propose to decouple the two sets of categories for capturing discriminative representations. In a similar spirit, PromptCAL [51] and DCCL [36] resort to synergistic learning and seek to improve the constructed representations via auxiliary training objectives. Nevertheless, those works are rooted in over-optimistic assumptions of i.i.d. data and balanced class distribution [44] and thereby tend to suffer from severe learning bias in real-world scientific discovery with intricate circumstances. Despite the recent efforts to account for non-i.i.d. and non-uniform data distribution in consonance with NCD [50, 48], their solutions are limited to the commonly adopted self-training scheme. Without explicit regularizations imposed on the geometric layout structure of the embedding space, their performance could fluctuate vastly given the volatile quality of pseudo-labels [46].

Open-World Semi-Supervised Learning. The technical challenge we intend to unravel is in line with an arising topic, i.e., open-world semi-supervised learning [3]. It takes a step further than the closed-world setting held in standard semi-supervised learning and presumes the emergence of new classes in unlabeled training and test dataset. Following the insight, several recent works suggest the usage of pseudo-labels and pair-wise similarity to group and recognize unseen categories [37, 4, 38, 19]. However, their promise in biomedical scientific discovery is unfavorably undermined by the challenges posed by distribution bias.

von Mises-Fisher Distribution. Probabilistic representation learning provides an elegant viewpoint to measure predictive uncertainty and mitigate feature ambiguity arising from deterministic mappings [41]. The multivariate Gaussian distribution is commonly employed as a prominent choice for probabilistic modeling in Euclidean space, whilst its applicability on spherical embeddings is limited due to inherent conflicts with the underlying manifold geometry [11]. On the contrary, the von Mises-Fisher (vMF) distribution is introduced as a directional statistical model that is defined on the hypersphere. It has been applied in tasks including facial recognition [27, 47], out-of-distribution detection [35], and long-tailed learning [45], leading to remarkable advancements in these areas. Nevertheless, the significance of incorporating crucial inductive biases per-

taining to geometric properties within the embedding space, such as uniformity, has been overlooked in the literature. Distinguished from existing works, to the best of our knowledge, our study is the first to explicitly structure the geometric layout of learned representations and account for the statistical discrepancy across data with distributional concentration modeling.

3. Methodology

3.1. Problem Statement

The objective of this study is to achieve automated discovery of clusters that correspond to novel concepts using observational data, which is accomplished by transferring the knowledge acquired from a set of base classes with labeled image samples. Following the practical setting in [44], we assume that there exists a labeled base set $\mathcal{D}_B = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{N_B} \in \mathcal{X}_B \times \mathcal{Y}_B$, and an unlabeled set $\mathcal{D}_U = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{N_U} \in \mathcal{X}_U \times \mathcal{Y}_U$ with potential novel classes, where \mathcal{Y}_U are not available. All instances in \mathcal{D}_B are of class set \mathcal{S}_B . As base classes also appear in the unlabeled set, i.e., $\mathcal{Y}_B \subset \mathcal{Y}_U$, the set of novel classes to be discovered can be formulated as $\mathcal{S}_N = \mathcal{Y}_U \setminus \mathcal{Y}_B$, with the total number $|\mathcal{S}_N|$ unknown a priori. Notably, distinguished from [44], the two sets of data are drawn from different distributions due to inconsistent biology processing and imaging protocols, i.e., $\mathcal{X}_B \sim \mathcal{P}_B, \mathcal{X}_U \sim \mathcal{P}_U, \mathcal{P}_B \neq \mathcal{P}_U$. The class distribution of \mathcal{D}_U is innately non-uniform as well. Given a new instance from \mathcal{P}_U , the goal is to either recognize it as one of the base classes or find it is closest to which identified novel class amongst \mathcal{S}_N .

3.2. Geometry-Constrained Probabilistic Modeling

Overview. The technical challenge to adapt base knowledge for novel biomedical concept discovery lies in two aspects: semantic ambiguity caused by data distribution shifts and uncontrollable open space risk as a result of major class dominance. To this end, we propose a novel geometry-constrained probabilistic treatment. The framework is illustrated in Fig. 2. It explicitly decouples the underlying distributional bias from informative semantic representations and structures the geometric layout of the formulated embedding space based upon inductive biases on latent manifold. First, we incorporate the *uniformity* inductive bias by positioning semantic proxies a priori with uniform distribution guarantees. It helps to mitigate the representation collapse incurred by the inherent imbalance between base and novel classes. Subsequently, we leverage vMF probabilistic modeling to promote latent space *boundness* by firstly establishing a compact manifold for each base class and then regularizing the open space to be separated from the base manifold by a large angular margin. At last, to demystify the dark information regarding unknown class representations,




Figure 2. Overview of the proposed method. We propose to incorporate the *uniformity* and *boundness* geometrical inductive biases by establishing preorganized proxies as anchors and then structuring the geometric layout of learned embedding space successively with hyperspherical probabilistic modeling.

we propose to structure the embedding space and minimize the open space risk by imposing statistical discrepancy constraints on unlabeled instances with semantic consensus.

Uniform Proxies Founded A Priori with Energy Minimization. Prior to representation learning, we propose to at first organize the embedding space with pre-defined proxies. By disentangling proxy foundation as an individual step, we seamlessly integrate the *uniformity* inductive bias to shape model output space without disturbing discriminative class learning. Here, uniformity indicates the degree of distributional diversity for the set of proxies. With the proxies distributed uniformly throughout the entire manifold, we minimize the open space risk by explicitly structuring the feature space for unknown classes. The equidistributional organization also facilitates inter-class separation and intra-class compactness. Specifically, given a d -dimensional embedding space, we define a set of prior proxies $\hat{\mathbf{Y}} = \{\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_n\} \in \mathbb{S}^{d-1}$ on a unit hyperspherical manifold. The number of proxies n should be generally larger than the total number of classes. We then propose to characterize the distributional uniformity of those proxies with their hyperspherical potential energy [39]. Energy behaves as a measure of distribution redundancy where lower energy implies the proxies are evenly positioned with large margin separation. We thereby suggest to minimize the potential energy with the following objective:

$$\arg \min_{\{\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_n\} \in \mathbb{S}^{d-1}} E_{s,d}(\hat{\mathbf{v}}_i)_{i=1}^n = \sum_{i=1}^n \sum_{j=1, j \neq i}^n \xi_{s,d}(\hat{\mathbf{v}}_i, \hat{\mathbf{v}}_j), \quad (1)$$

where the unit hypersphere $\mathbb{S}^{d-1} = \{\hat{\mathbf{v}} \in \mathbb{R}^d \mid \|\hat{\mathbf{v}}\| = 1\}$. $\xi_{s,d}(\cdot, \cdot)$ represents the measurement of the correlation between proxies, which is typically implemented with the fol-

lowing Riesz s -kernel function:

$$\xi_{s,d}(\hat{\mathbf{v}}_i, \hat{\mathbf{v}}_j) := \begin{cases} \Gamma(\hat{\mathbf{v}}_i, \hat{\mathbf{v}}_j)^{-s}, & s > 0, \\ \log(\Gamma(\hat{\mathbf{v}}_i, \hat{\mathbf{v}}_j)^{-1}), & s = 0, \end{cases} \quad (2)$$

where $\Gamma(\cdot, \cdot)$ is defined as the hyperspherical geodesic distance.

Distributional Characterization on Hypersphere. To account for the semantic ambiguities induced by data distribution shifts, we propose probabilistic modeling to parameterize latent embeddings as directional distributions on a hyperspherical manifold. Specifically, we first let f_θ represent a projection network that maps input images to their hyperspherical embedding, i.e., $f_\theta : \mathbf{x} \in \mathcal{X} \mapsto \mathbf{z} \in \mathbb{S}^{d-1}$. Then, we characterize the approximated posterior of each instance embedding as a marginal von Mises-Fisher (vMF) distribution, i.e., $q_\theta(\mathbf{z}|\mathbf{x}) \sim \mathbf{vMF}(\tilde{\boldsymbol{\mu}}_{\mathbf{x}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}})$. Formally, a vMF distribution is defined as:

$$q(\mathbf{z}|\tilde{\boldsymbol{\mu}}_{\mathbf{x}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) = C_d(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) \exp(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}}^T \mathbf{z}), \quad (3)$$

$$C_d(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) = \frac{\tilde{\boldsymbol{\kappa}}_{\mathbf{x}}^{d/2-1}}{(2\pi)^{d/2} \mathcal{I}_{d/2-1}(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}})}, \quad (4)$$

where $\mathcal{I}_{d/2-1}(\cdot)$ denotes the modified Bessel function of the first kind at order $(d/2 - 1)$. $\tilde{\boldsymbol{\mu}}_{\mathbf{x}}$ and $\tilde{\boldsymbol{\kappa}}_{\mathbf{x}}$ are statistical parameters representing the direction and concentration of the distribution. With this formulation, the discriminative semantic context can be derived from the directional placement on the hypersphere and hence get decoupled from task-irrelevant attributes. Besides, the concentration parameter $\tilde{\boldsymbol{\kappa}}_{\mathbf{x}}$ provides an intuitive measurement of semantic embedding ambiguity [40] and can therefore act as an instance-wise temperature factor. It adaptively scales the contributions of different image samples based on their representation uncertainty to deliver dynamic rectifications on the latent space structuring procedure, which is theoretically analyzed afterwards.

Base Space Bounding and Open Space Dispersion. With the modeled probabilistic representations, the subsequent objective is to organically structure the geometric layout of embedding space for discriminative novel class discovery. Ideally, the space covered by well-defined base categories should be bounded within a certain range while the unknown classes corresponding to the open space should be modeled apart from the periphery of base space and dispersed over the entire manifold [6]. We therefore propose to leverage the pre-defined proxies to bound the base space and keep the open space isolated. For each base class, we select one proxy as the class anchor and reformulate it with Dirac delta δ to derive a statistical modeling over its deter-

ministic representation:

$$\Delta(\mathbf{z}) := \begin{cases} \int_{\mathbb{S}^{d-1}} \delta(\mathbf{z} - \hat{\mathbf{v}}_{\mathcal{B}}) d\mathbf{z} = 1; \\ \int_{\mathbb{S}^{d-1}} \delta(\mathbf{z} - \hat{\mathbf{v}}_{\mathcal{B}}) \phi(\mathbf{z}) d\mathbf{z} = \phi(\hat{\mathbf{v}}_{\mathcal{B}}). \end{cases} \quad (5)$$

Here $\hat{\mathbf{v}}_{\mathcal{B}}$ denotes the fixed positional values of the proxies for base classes. It is noted that the selection of proxies is random since they have been forced to be uniformly distributed on a symmetric manifold and all proxy pairs have the same distance. Then, the proxies can behave as desired latent priors to regularize the distribution of base classes. Specifically, the optimization objective is to enforce the instance-wise vMF distribution $q(\mathbf{z}|\tilde{\boldsymbol{\mu}}_{\mathbf{x}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}})$ approximating the latent prior $\Delta(\mathbf{z})$ in terms of KL divergence:

$$\min_q \mathbb{E}_{\mathcal{D}_{\mathcal{B}}} [- \int_{\mathbb{S}^{d-1}} \Delta(\mathbf{z}) \log q(\mathbf{z}|\tilde{\boldsymbol{\mu}}_{\mathbf{x}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) d\mathbf{z} - \mathcal{H}_{\Delta}(\mathbf{z})], \quad (6)$$

where $\mathcal{H}_{\Delta}(\mathbf{z})$ denotes the differential entropy over $\Delta(\mathbf{z})$, which is a constant as $\Delta(\mathbf{z})$ is fixed a priori. Combined with Eq. (3)-(5), the objective is equivalent to minimizing the following loss function:

$$\begin{aligned} \mathcal{L}_{\mathcal{B}}(\tilde{\boldsymbol{\mu}}_{\mathbf{x}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) &= - \int_{\mathbb{S}^{d-1}} [\tilde{\boldsymbol{\kappa}}_{\mathbf{x}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}}^T \mathbf{z} + \log C_d(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}})] \Delta(\mathbf{z}) d\mathbf{z} \\ &= -\tilde{\boldsymbol{\kappa}}_{\mathbf{x}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}}^T \hat{\mathbf{v}}_{\mathcal{B}} - (d/2 - 1) \log \tilde{\boldsymbol{\kappa}}_{\mathbf{x}} + \log \mathcal{I}_{d/2-1}(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) + \check{\gamma}, \end{aligned} \quad (7)$$

where $\check{\gamma} = (d/2) \log 2\pi$ is a constant constituent. The regularization term constrains the intra-class variability and shrinks the base space within a bounded radius. Thereafter, we propose open space dispersion aimed at separating the space reserved for novel classes away from the base one. The challenging aspect is how to recognize the unseen concepts from the hybrid unlabeled set of data where base and novel classes concurrently exist. In this regard, we first derive the magnitude of probabilistic deviation as an indication of open-set instances and hereby enforce large angular margins between those embeddings and the base space. Specifically, we measure the probability density of each instance-wise vMF distribution with respect to base class proxies and estimate the ranking over $\mathcal{X}_{\mathcal{U}}$:

$$\underset{\{\mathbf{x}_1, \dots, \mathbf{x}_{N_{\mathcal{U}}}\} \in \mathcal{X}_{\mathcal{U}}}{\text{argsort}} \left\{ \max_{\{\hat{\mathbf{v}}_{\mathcal{B}_i}, \forall i \in \mathcal{S}_{\mathcal{B}}\}} C_d(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) \exp(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}}^T \hat{\mathbf{v}}_{\mathcal{B}_i}) \right\}. \quad (8)$$

The batch-wise instances with maximal divergence from the base class proxies are hereafter pushed towards the open space:

$$\mathcal{L}_{dis}(\tilde{\boldsymbol{\mu}}_{\mathbf{x}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) = -\log \frac{\exp(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}}^T \hat{\mathbf{v}}_{\mathcal{R}})}{\exp(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}}^T \hat{\mathbf{v}}_{\mathcal{R}}) + \exp(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}}^T \hat{\mathbf{v}}_{\mathcal{B}})}, \quad (9)$$

where $\hat{\mathbf{v}}_{\mathcal{R}} = \hat{\mathbf{Y}} \setminus \hat{\mathbf{v}}_{\mathcal{B}}$ denotes the remaining proxies located outside of the base space. The optimization objectives of

Eq. (7) and (9) cooperatively promote *boundness* on the underlying manifold. We then propose to explicitly introduce the *uniformity* inductive bias and accordingly structure the open space to control its associated risk, which in turn appreciates the generalizability of learned representations for recognizing unseen concepts.

Open Space Structuring. The dark information in the unknown space which has not been properly characterized raises uncontrollable risks for novel classes learning and discrimination [6]. To this end, we propose to structure the open space organically with proxy-driven semantic consensus, which encourages the constructed embeddings to be uniformly distributed over the manifold with semantically-meaningful clusters such that the representations of different novel classes can be spaced apart with large angular distance. Our intuition is that instances with adjacent geometric placements on the underlying manifold should possess analogous semantics and exhibit marginal distributional discrepancy. Specifically, given each pair of unlabeled instances $(\mathbf{x}_{\alpha}, \mathbf{x}_{\beta})$ and corresponding probabilistic modeling $q_{\theta}(\mathbf{z}_{\alpha}|\mathbf{x}_{\alpha}) \sim \mathbf{vMF}(\tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\alpha}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\alpha}})$, $q_{\theta}(\mathbf{z}_{\beta}|\mathbf{x}_{\beta}) \sim \mathbf{vMF}(\tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\beta}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\beta}})$, we measure their distributional overlaps with all pre-set proxies as $\mathcal{O}_{(\alpha, \beta)}^{\hat{\mathbf{Y}}} = \{C_d(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}}) \exp(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}}^T \hat{\mathbf{v}}_i), \forall \hat{\mathbf{v}}_i \in \hat{\mathbf{Y}}, \mathbf{x} \in (\mathbf{x}_{\alpha}, \mathbf{x}_{\beta})\}$, which indicate the statistical distances between instance embeddings and proxies. Since the proxies have been priorly positioned covering the entire manifold with uniformity guarantees, the derived $\mathcal{O}_{\alpha}^{\hat{\mathbf{Y}}}$ and $\mathcal{O}_{\beta}^{\hat{\mathbf{Y}}}$ can comprehensively characterize the geometric location of instance embeddings. As such, we suggest \mathbf{x}_{α} is semantically correlated with \mathbf{x}_{β} and their distributional discrepancy should be minor if $\mathcal{O}_{\alpha}^{\hat{\mathbf{Y}}} \doteq \mathcal{O}_{\beta}^{\hat{\mathbf{Y}}}$. Formally, the condition holds true if $\arg \max_{\text{topk}}(\mathcal{O}_{\alpha}^{\hat{\mathbf{Y}}}) = \arg \max_{\text{topk}}(\mathcal{O}_{\beta}^{\hat{\mathbf{Y}}})$. Let \mathcal{G} denote the collection of instance pairs satisfying the above condition. In these cases, we aim at minimizing the discrepancy between the two vMF distributions to escalate their semantic affinities. We resort to modeling the distributional discrepancy with KL divergence and accordingly derive the following optimization objective:

$$\min_q \mathbb{E}_{\mathcal{G}} [- \int_{\mathbb{S}^{d-1}} q(\mathbf{z}|\tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\alpha}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\alpha}}) \log \left(\frac{q(\mathbf{z}|\tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\beta}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\beta}})}{q(\mathbf{z}|\tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\alpha}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\alpha}})} \right) d\mathbf{z}]. \quad (10)$$

Combined with Eq. (3) and (4), we can derive the corresponding regularization term as follows:

$$\begin{aligned} \mathcal{L}_{str}(\tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\alpha}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\alpha}}, \tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\beta}}, \tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\beta}}) &= - \left[\frac{\mathcal{I}_{d/2}(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\alpha}})}{\mathcal{I}_{d/2-1}(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\alpha}})} \right. \\ &\quad \left. (\tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\alpha}} - \tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\beta}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\alpha}} \tilde{\boldsymbol{\mu}}_{\mathbf{x}_{\beta}}^T) + \log \left(\frac{C_d(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\alpha}})}{C_d(\tilde{\boldsymbol{\kappa}}_{\mathbf{x}_{\beta}})} + 1 \right) \right]^{-1}. \end{aligned} \quad (11)$$

The constraints advocate explicit associations between the geometric layout of embedding manifold and the semantic

representations of instances, which in turn systematically structure the open space. By integrating the optimization objective along with the ones in Eq. (7) and (9), we jointly incorporate the *boundness* and *uniformity* inductive biases to shape the learned embedding space and thereby favor its generalization to unseen biomedical concepts.

3.3. Spectral-Based Class Number Estimation

With the well-structured embedding space, the subsequent target is to estimate the number of classes in unlabeled data and accordingly perform clustering to assign class label for each instance in \mathcal{D}_U . Different from previous methods counting on the results of empirical evaluations and introducing costly recursive clustering [44], we propose a spectral graph-theoretic approach based upon the finding that the eigenvalues of graph Laplacians inherently characterize the underlying connectivity of the latent manifold [15]. Specifically, we construct a weighted graph over all instances of \mathcal{D}_U , in which each node corresponds to the embedding \mathbf{z}_i of a data point, and the edge weights \mathbf{W} model the correlations between pairs of instances as $\{\mathbf{z}_i \mathbf{z}_j^T\}$. Let $\mathbf{\Lambda}$ denote the diagonal matrix with the row-wise sum of \mathbf{W} as its entries, the optimization objective for graph partitioning with minimized energy is:

$$\min_{\mathbf{h}} \frac{\mathbf{h}^T \mathbf{\Lambda}^{-1/2} (\mathbf{\Lambda} - \mathbf{W}) \mathbf{\Lambda}^{-1/2} \mathbf{h}}{\mathbf{h}^T \mathbf{h}}, \quad (12)$$

which is equivalent to solve $\mathbf{\Lambda}^{-1/2} (\mathbf{\Lambda} - \mathbf{W}) \mathbf{\Lambda}^{-1/2} \mathbf{h} = \lambda \mathbf{h}$. The results can be intrinsically depicted by the eigenvectors of the normalized graph Laplacian $\mathbf{\Lambda}^{-1/2} (\mathbf{\Lambda} - \mathbf{W}) \mathbf{\Lambda}^{-1/2}$ [2]. In this regard, the eigenvectors can behave as soft segments to partition the overall graph into a set of semantically-meaningful splits, while the corresponding eigenvalues can be considered as indicators for the inner-connectivity of potential clusters. In other words, instances in splits with higher eigenvalues exhibit stronger internal semantic affinities and are more plausible to be assigned into the same class. We therefore propose to estimate the number of latent classes according to the spectral property of constructed graph that we rank all possible graph splits with their associated eigenvalues and identify the split index which corresponds to the largest gap within the ordering of eigenvalue differences $\{argsort \text{diff}(eigenvalues)\}$. The identified index suggests the start of graph splits which contain limited inner-connectivity and should not be considered as individual classes. The index is thereby modeled as the estimated number of classes. We then perform k-means clustering over \mathcal{D}_U in the hyperspherical space with the estimated class number for label assignment. Notably, taxonomy-adaptive estimation could be in demand that it is required to gauge the number of potential fine-grained subclasses. Our method is well suited for this scenario as it can be adapted to estimate the inherent subclasses by loosening

the above condition and instead indexing the subordinately largest eigenvalue gap according to the specific need in class hierarchy.

3.4. Theoretical Analysis

In this section, we motivate our method theoretically and provide analytical insights to validate its merits.

Monotonicity between Distributional Concentration and Semantic Ambiguity. With the intrinsic semantic ambiguity present in biomedical visual data due to disparate imaging protocols across cohorts, our probabilistic treatment, which advocates for vMF statistical modeling, adeptly characterizes the ambiguous attributes with distributional concentration $\tilde{\kappa}$:

Proposition 1 *Let $\zeta_{\mathbf{x}}$ be the continuous entropy of the posterior vMF distribution parametrized by $\tilde{\boldsymbol{\mu}}_{\mathbf{x}} \in \mathbb{S}^{d-1}$ and $\tilde{\kappa}_{\mathbf{x}} \in \mathbb{R}_{>0}^d$. We have $\zeta_{\mathbf{x}}(\tilde{\kappa}_{\mathbf{x}})$ behave as a monotonically decreasing function in the interval $(0, +\infty)$.*

It suggests the statistical parameter $\tilde{\kappa}_{\mathbf{x}}$ essentially measures the aleatoric uncertainty of the semantic context in data point \mathbf{x} . Combined with the analysis of $\tilde{\kappa}_{\mathbf{x}}$ on distributional regularization terms, as presented in the supplementary material, it is demonstrated the probabilistic modeling dynamically calibrates the biased learning procedure towards discriminative contextual information modeling.

Towards Bounded Open Space Risk with Uniform Proxies. The most challenging aspect in the GNCD problem is how to demystify the dark information regarding unknown class representations in order to restrain the associated open space risk [36]. We hereby prove the effectiveness of our proposed method towards a tighter error bound for open space risk characterization and regularization. Details can be found in the supplemental material.

4. Experiments

4.1. Experimental Setup

Datasets. To evaluate the effectiveness and general applicability of our method under miscellaneous biomedical discovery scenarios, we conduct experiments on a diverse suite of challenging tasks encompassing various biomedical concepts (cell, lesion, disease), imaging modalities (X-ray, microscopy, dermatoscopy, fundus photography), and pathologies (lung infection, skin cancer, eye complication). For each task, we adopt two different datasets with evident data distribution shifts to serve as the labeled base set and the unlabeled set comprising novel classes, respectively. In specific, we consider the discovery of novel types or degrees of pneumonia infectious organisms [25, 10], cell nuclei [17, 18], skin lesions [43, 9], and diabetic retinopathy severity [28, 24]. The quantity of instances belonging to novel classes is inherently much smaller compared

Table 1. Comparison results of our proposed method against other state-of-the-art methods for discovering novel types of pneumonia infectious organisms and cell nuclei. The subscripts denote the standard deviations. The best and second-best results are highlighted in bold and brown, respectively.

Methods	Pneumonia Infectious Organisms						Cell Nuclei					
	Acc-all	Acc-known	Acc-novel	F ₁ -all	F ₁ -known	F ₁ -novel	Acc-all	Acc-known	Acc-novel	F ₁ -all	F ₁ -known	F ₁ -novel
RankStats+	41.32 _{0.90}	40.57 _{0.60}	46.67 _{2.37}	28.96 _{1.56}	33.45 _{1.13}	24.48 _{3.13}	36.25 _{0.80}	36.17 _{0.82}	39.84 _{0.65}	27.17 _{0.68}	42.45 _{0.74}	4.26 _{0.41}
UNO+	49.59 _{1.48}	50.06 _{0.92}	46.31 _{1.49}	31.24 _{1.26}	36.35 _{2.11}	26.14 _{2.70}	41.83 _{0.92}	41.59 _{0.74}	47.83 _{1.22}	33.04 _{0.58}	45.79 _{0.69}	13.54 _{0.52}
GCD	67.77 _{0.83}	69.81 _{1.45}	53.34 _{0.67}	38.06 _{1.04}	41.26 _{1.30}	32.87 _{1.95}	45.43 _{1.78}	45.57 _{1.26}	43.73 _{2.43}	38.19 _{1.47}	44.52 _{1.05}	29.64 _{2.84}
DCCL	72.73 _{2.14}	74.48 _{2.28}	60.32 _{2.82}	44.82 _{1.63}	49.85 _{1.95}	39.79 _{4.08}	43.93 _{1.34}	43.59 _{1.86}	54.49 _{0.53}	37.93 _{1.14}	40.16 _{2.04}	36.57 _{0.77}
SLF	43.81 _{0.67}	42.31 _{0.44}	54.26 _{1.12}	30.72 _{1.39}	32.86 _{0.97}	28.57 _{2.86}	39.15 _{0.86}	39.07 _{1.04}	43.06 _{0.69}	28.95 _{0.95}	44.06 _{1.02}	18.42 _{0.81}
BYOP	65.52 _{0.95}	68.87 _{1.14}	58.95 _{2.54}	40.01 _{1.02}	43.12 _{1.30}	38.52 _{3.72}	46.28 _{0.68}	46.02 _{0.57}	55.82 _{1.44}	39.06 _{0.59}	41.01 _{0.30}	35.63 _{1.86}
Ours	78.51 _{1.05}	80.19 _{1.02}	66.67 _{2.70}	46.73 _{1.47}	44.27 _{0.31}	49.18 _{3.87}	47.78 _{1.28}	47.41 _{1.53}	57.85 _{0.42}	42.64 _{0.36}	41.48 _{0.25}	43.27 _{2.29}

Table 2. Comparison results for discovering novel types of skin cancer lesions and degrees of diabetic retinopathy severity.

Methods	Skin Cancer Lesions						Diabetic Retinopathy Severity					
	Acc-all	Acc-known	Acc-novel	F ₁ -all	F ₁ -known	F ₁ -novel	Acc-all	Acc-known	Acc-novel	F ₁ -all	F ₁ -known	F ₁ -novel
RankStats+	29.19 _{1.03}	28.65 _{0.79}	33.39 _{2.04}	21.06 _{1.60}	36.43 _{0.95}	9.53 _{2.84}	47.21 _{0.94}	47.90 _{0.58}	41.64 _{1.50}	40.82 _{0.69}	54.78 _{0.46}	19.97 _{0.82}
UNO+	37.80 _{2.01}	37.43 _{1.54}	40.74 _{2.38}	29.16 _{1.49}	40.68 _{1.41}	20.37 _{1.82}	50.43 _{0.85}	51.07 _{0.97}	39.26 _{0.68}	36.57 _{0.84}	52.45 _{0.57}	16.09 _{1.23}
GCD	31.58 _{1.44}	49.46 _{1.29}	39.58 _{1.06}	27.18 _{0.82}	39.17 _{0.53}	18.20 _{2.25}	48.43 _{2.36}	49.05 _{2.04}	44.33 _{1.22}	41.69 _{1.06}	53.39 _{0.85}	21.13 _{0.93}
DCCL	56.45 _{0.78}	58.92 _{0.43}	37.84 _{1.37}	33.46 _{1.14}	53.46 _{0.59}	18.48 _{1.50}	44.35 _{1.58}	44.93 _{1.44}	46.07 _{2.27}	43.23 _{1.55}	52.16 _{1.40}	27.51 _{1.89}
SLF	44.56 _{1.29}	45.14 _{0.73}	39.02 _{0.67}	30.78 _{1.77}	49.75 _{0.84}	16.55 _{1.96}	46.29 _{0.99}	46.93 _{0.65}	44.40 _{1.40}	40.41 _{1.13}	50.92 _{0.54}	20.10 _{1.72}
BYOP	65.79 _{0.62}	69.46 _{0.44}	36.65 _{1.51}	35.94 _{1.14}	54.19 _{0.71}	23.04 _{1.70}	51.67 _{1.24}	52.31 _{1.27}	43.63 _{0.93}	44.85 _{0.82}	55.68 _{1.23}	25.56 _{0.90}
Ours	69.62 _{0.96}	73.24 _{0.32}	41.67 _{0.55}	41.39 _{0.29}	60.41 _{0.40}	27.12 _{0.48}	58.41 _{1.35}	59.62 _{1.60}	50.52 _{0.51}	47.25 _{0.77}	55.94 _{0.56}	34.22 _{0.64}

with their base counterpart and naturally formalize the long-tailed class distribution. More details can be found in the supplemental material, where extended studies on general concept discovery are further presented and analyzed.

Implementation Details and Evaluation Protocol. We adopt the ViT-B/16 [12] with DINO pre-trained weights [5] as the backbone network for fair comparisons with the previous methods [44, 36]. Its output of [CLS] token with a dimension of 768 is used for latent embedding. Please refer to the attached code for details.

For evaluation, we measure the clustering accuracy between the predicted class assignment and the ground-truth label. Following the protocol in [44], we match the two class sets by searching for the optimal permutation using the Hungarian assignment algorithm [26]. We adopt the accuracy and F1 scores as metrics and report the averaged score across the entire class set. Particularly, we also present the scores specific to the base known \mathcal{S}_B and novel \mathcal{S}_N subset of classes in a separated manner.

4.2. Results and Discussions

Main Results. We compare our method with state-of-the-art approaches for novel class discovery (NCD), including RankStats+ [21], UNO+ [16], GCD [44], and DCCL [36] which are specifically devised for the challenging generalized NCD setting, as well as SLF [50] and BYOP [48] which take account of the non-i.i.d. data and non-uniform




Figure 3. Results of class number estimation in the unlabeled set. For pneumonia and cell nuclei, we further estimate the number of fine-grained subclasses.

label distribution issues, respectively. The overall quantitative results over four diverse biomedical concept discovery tasks are presented in Table 1 and 2. All methods are implemented based upon the same ViT-B/16 backbone with DINO pre-trained weights [44]. It can be observed that our proposed method consistently and significantly outperforms the state-of-the-art novel class discovery approaches over all four tasks. Notably, as highlighted in the gray columns which correspond to the clustering metrics specific to novel classes, our method surpasses the comparison approaches by a substantial margin especially on the identification and grouping of unseen phenotypes. The results demonstrate the advantages and effectiveness of the proposed geometry-

Table 3. Key component analysis of our method. Uni, Bnd, Disp, and Str correspond to uniform proxies, base space bounding, open space dispersion and structuring, respectively.

Components				Pneumonia			Cell Nuclei		
Uni	Bnd	Disp	Str	F ₁ -all	F ₁ -known	F ₁ -novel	F ₁ -all	F ₁ -known	F ₁ -novel
				22.82	27.37	18.28	16.21	20.76	8.64
✓	✓			35.84	48.14	23.53	34.30	38.76	27.72
✓	✓	✓		36.92	45.89	27.94	33.02	31.83	33.47
✓	✓		✓	28.71	36.51	20.91	38.23	40.49	35.59
	✓	✓	✓	40.23	37.69	31.42	31.97	34.85	29.74
✓	✓	✓	✓	46.73	44.27	49.18	42.64	41.48	43.27

constrained probabilistic modeling for discovering novel concepts despite the data and label distribution biases intrinsically present in the biomedical domain.

We further assess the efficacy of the spectral-based class number estimation method, as presented in Fig. 3. The proposed approach attains promising results that the estimated number is identical or very close to the ground-truth value. The versatile method also holds an appealing merit that it supports taxonomy-adaptive estimation, which means it is capable of estimating the number of potential fine-grained subclasses, a practical aspect not previously considered in the literature at all. Additionally, it is noteworthy that our method does not rely on recursive clustering as in [44] and is therefore computationally efficient. On the pneumonia benchmark, our method only takes **0.121s**, which is more than **10×** faster than the comparison approach [44] consuming **2.854s**.

Ablation of Key Components. Table 3 presents the ablation study results to demonstrate the effectiveness of key components in our proposed method. It can be remarked that each component is indispensable and holds a positive impact on novel class discovery. Notably, the combination of pre-defined uniform proxies and base space bounding reaches the peak accuracy in terms of F₁-known on the pneumonia infectious organism benchmark. However, it suffers from inferior efficacy in identifying and clustering potential novel concepts. The observation is in line with our argument that the consideration and explicit modeling of the dark information regarding unknown classes are crucial to ensure the generalizability of learned representations for discovering unseen concepts. Solely imposing constraints on the base space is insufficient to calibrate the learning process biased towards the dominant known classes. In contrast, our method proposes two supplementary geometric constraints which cooperatively shape the layout of embedding space and compensate for the open space risk, which in turn leads to unbiased and discriminative modeling of novel class representations.

Impacts on Geometric Layout of Embedding Space. To conspicuously illustrate the effectiveness of our method to

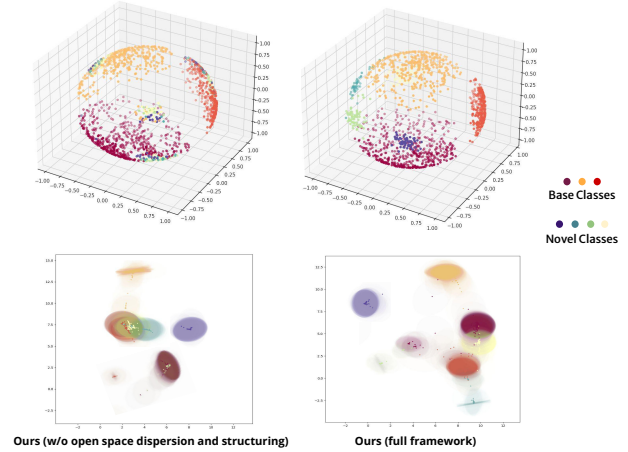


Figure 4. Hyperspherical and distributional embeddings of the skin lesion test set for two variants of our proposed method. For distributional visualization, we overlay data points over the super-level-set ellipses of the associated probabilistic distributions.

wards shaping the geometric structure of embedding space, we visualize the learned representations on the 3D spherical manifold and its associated probabilistic characterizations, as shown in Fig. 4. The covariance structure of data points is modeled as a group of super-level-set ellipses to indicate the uncertainty of constructed clusters. We observe that the removal of open space dispersion and structuring inevitably deteriorates the degree of inter-class angular separation and results in ambiguous semantic representations. The comparative analysis between the two variants of our proposed method substantiates the significance of incorporating geometric inductive biases to regulate the underlying organization of the embedding space.

5. Conclusion

In this work, we propose a geometry-constrained probabilistic modeling framework to address challenging aspects of novel class discovery in the biomedical domain. Specifically, we suggest to model the latent representation of each instance as a directional distribution to exclude the task-irrelevant data variations arising from inconsistent imaging protocols across cohorts. We then identify several geometric properties critical for organic embedding space structuring. By explicitly integrating those principles to regularize representation layout, the open space risk associated with the discovery of novel classes can be effectively bounded. Additionally, a spectral graph-theoretic algorithm is developed to estimate the number of categories in unlabeled data, which inherits taxonomy adaptivity and high efficiency. Our method exhibits appealing performance on a diverse suite of novel concept discovery applications in biomedical domains.

References

- [1] Wenbin An, Feng Tian, Qinghua Zheng, Wei Ding, Qianying Wang, and Ping Chen. Generalized category discovery with decoupled prototypical network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023.
- [2] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on pattern analysis and machine intelligence*, 23(11):1222–1239, 2001.
- [3] Kaidi Cao, Maria Brbic, and Jure Leskovec. Open-world semi-supervised learning. In *International Conference on Learning Representations*, 2022.
- [4] Tianyue Cao, Yongxin Wang, Yifan Xing, Tianjun Xiao, Tong He, Zheng Zhang, Hao Zhou, and Joseph Tighe. Pss: Progressive sample selection for open-world visual representation learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pages 278–294. Springer, 2022.
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- [6] Guangyao Chen, Limeng Qiao, Yemin Shi, Peixi Peng, Jia Li, Tiejun Huang, Shiliang Pu, and Yonghong Tian. Learning open set network with discriminative reciprocal points. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 507–522. Springer, 2020.
- [7] Mei Chen. *Computer vision for microscopy image analysis*. Academic Press, 2020.
- [8] Haoang Chi, Feng Liu, Bo Han, Wenjing Yang, Long Lan, Tongliang Liu, Gang Niu, Mingyuan Zhou, and Masashi Sugiyama. Meta discovery: Learning to discover novel classes given very limited data. In *International Conference on Learning Representations*, 2022.
- [9] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kallou, Konstantinos Liopyris, Michael Marchetti, et al. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368*, 2019.
- [10] Joseph Paul Cohen, Paul Morrison, Lan Dao, Karsten Roth, Tim Duong, and Marzyeh Ghassem. Covid-19 image data collection: Prospective predictions are the future. *Machine Learning for Biomedical Imaging*, 1:1–38, 2020.
- [11] Tim R Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M Tomczak. Hyperspherical variational auto-encoders. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2018.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [13] Jianan Fan, Dongnan Liu, Hang Chang, and Weidong Cai. Learning to generalize over subpartitions for heterogeneity-aware domain adaptive nuclei segmentation. *International Journal of Computer Vision*, pages 1–24, 2024.
- [14] Jianan Fan, Dongnan Liu, Hang Chang, Heng Huang, Mei Chen, and Weidong Cai. Taxonomy adaptive cross-domain adaptation in medical imaging via optimization trajectory distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21174–21184, 2023.
- [15] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak mathematical journal*, 23(2):298–305, 1973.
- [16] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun Zhong, Moin Nabi, and Elisa Ricci. A unified objective for novel class discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9284–9292, 2021.
- [17] Jevgenij Gamper, Navid Alemi Koohbanani, Ksenija Benet, Ali Khuram, and Nasir Rajpoot. Pannuke: an open pancreatic histology dataset for nuclei instance segmentation and classification. In *Digital Pathology: 15th European Congress, ECDP 2019, Warwick, UK, April 10–13, 2019, Proceedings 15*, pages 11–19. Springer, 2019.
- [18] Simon Graham, Mostafa Jahanifar, Ayesha Azam, Mohammed Nimir, Yee-Wah Tsang, Katherine Dodd, Emily Hero, Harvir Sahota, Atisha Tank, Ksenija Benes, et al. Lizard: A large-scale dataset for colonic nuclear instance segmentation and classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 684–693, 2021.
- [19] Lan-Zhe Guo, Yi-Ge Zhang, Zhi-Fan Wu, Jie-Jing Shao, and Yu-Feng Li. Robust semi-supervised learning when not all classes have labels. *Advances in Neural Information Processing Systems*, 35:3305–3317, 2022.
- [20] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *International Conference on Learning Representations*, 2020.
- [21] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Autonovel: Automatically discovering and learning novel visual categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6767–6781, 2021.
- [22] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8401–8409, 2019.
- [23] KJ Joseph, Sujoy Paul, Gaurav Aggarwal, Soma Biswas, Piyush Rai, Kai Han, and Vineeth N Balasubramanian. Novel class discovery without forgetting. In *European Conference on Computer Vision*, pages 570–586. Springer, 2022.
- [24] Sohier Dane Karthik, Maggie. Aptos 2019 blindness detection. *Kaggle*, 2019.
- [25] Daniel S Kermany, Michael Goldbaum, Wenjia Cai, Carolina CS Valentim, Huiying Liang, Sally L Baxter, Alex

- McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5):1122–1131, 2018.
- [26] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [27] Shen Li, Jianqing Xu, Xiqing Xu, Pengcheng Shen, Shaoxin Li, and Bryan Hooi. Spherical confidence learning for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15629–15637, 2021.
- [28] Tao Li, Yingqi Gao, Kai Wang, Song Guo, Hanruo Liu, and Hong Kang. Diagnostic assessment of deep learning algorithms for diabetic retinopathy screening. *Information Sciences*, 501:511–522, 2019.
- [29] Wenbin Li, Zhichen Fan, Jing Huo, and Yang Gao. Modeling inter-class and intra-class constraints in novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3449–3458, 2023.
- [30] Dongnan Liu, Donghao Zhang, Yang Song, Heng Huang, and Weidong Cai. Panoptic feature fusion net: a novel instance segmentation paradigm for biomedical and biological images. *IEEE Transactions on Image Processing*, 30:2045–2059, 2021.
- [31] Dongnan Liu, Donghao Zhang, Yang Song, Fan Zhang, Lauren O’Donnell, Heng Huang, Mei Chen, and Weidong Cai. Unsupervised instance segmentation in microscopy images via panoptic domain adaptation and task re-weighting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4243–4252, 2020.
- [32] Quande Liu, Cheng Chen, Jing Qin, Qi Dou, and Pheng-Ann Heng. Feddg: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1013–1023, 2021.
- [33] Evan Z Macosko, Anindita Basu, Rahul Satija, James Nemesh, Karthik Shekhar, Melissa Goldman, Itay Tirosh, Allison R Bialas, Nolan Kamitaki, Emily M Martersteck, et al. Highly parallel genome-wide expression profiling of individual cells using nanoliter droplets. *Cell*, 161(5):1202–1214, 2015.
- [34] Yassine Marrakchi, Osama Makansi, and Thomas Brox. Fighting class imbalance with contrastive learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 466–476. Springer, 2021.
- [35] Yifei Ming, Yiyu Sun, Ousmane Dia, and Yixuan Li. How to exploit hyperspherical embeddings for out-of-distribution detection? In *The Eleventh International Conference on Learning Representations*, 2023.
- [36] Nan Pu, Zhun Zhong, and Nicu Sebe. Dynamic conceptional contrastive learning for generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7579–7588, 2023.
- [37] Mamshad Nayeem Rizve, Navid Kardan, Salman Khan, Fahad Shahbaz Khan, and Mubarak Shah. Openldn: Learning to discover novel classes for open-world semi-supervised learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pages 382–401. Springer, 2022.
- [38] Mamshad Nayeem Rizve, Navid Kardan, and Mubarak Shah. Towards realistic semi-supervised learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pages 437–455. Springer, 2022.
- [39] Edward B Saff and Amo BJ Kuijlaars. Distributing many points on a sphere. *The mathematical intelligencer*, 19:5–11, 1997.
- [40] Tyler R Scott, Andrew C Gallagher, and Michael C Mozer. von mises-fisher loss: An exploration of embedding geometries for supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10612–10622, 2021.
- [41] Yichun Shi and Anil K Jain. Probabilistic face embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6902–6911, 2019.
- [42] Yang Song and Weidong Cai. Visual feature representation in microscopy image classification. In *Computer Vision for Microscopy Image Analysis*, pages 73–100. Elsevier, 2021.
- [43] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018.
- [44] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7492–7501, 2022.
- [45] Hualiang Wang, Siming Fu, Xiaoxuan He, Hangxiang Fang, Zuozhu Liu, and Haoji Hu. Towards calibrated hyper-sphere representation via distribution overlap coefficient for long-tailed learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV*, pages 179–196. Springer, 2022.
- [46] Qian Wang and Toby Breckon. Unsupervised domain adaptation via structured prediction based selective pseudo-labeling. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 6243–6250, 2020.
- [47] Jianqing Xu, Shen Li, Ailin Deng, Miao Xiong, Jiaying Wu, Jiaxiang Wu, Shouhong Ding, and Bryan Hooi. Probabilistic knowledge distillation of face ensembles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3489–3498, 2023.
- [48] Muli Yang, Liancheng Wang, Cheng Deng, and Hanwang Zhang. Bootstrap your own prior: Towards distribution-agnostic novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3459–3468, 2023.
- [49] Muli Yang, Yuehua Zhu, Jiaping Yu, Aming Wu, and Cheng Deng. Divide and conquer: Compositional experts for generalized novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14268–14277, 2022.

- [50] Qing Yu, Daiki Ikami, Go Irie, and Kiyoharu Aizawa. Self-labeling framework for novel category discovery over domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3161–3169, 2022.
- [51] Sheng Zhang, Salman Khan, Zhiqiang Shen, Muzammal Naseer, Guangyi Chen, and Fahad Shahbaz Khan. Promptcal: Contrastive affinity learning via auxiliary prompts for generalized novel category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3479–3488, 2023.
- [52] Yuyang Zhao, Zhun Zhong, Nicu Sebe, and Gim Hee Lee. Novel class discovery in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2022.
- [53] Jiyang Zheng, Weihao Li, Jie Hong, Lars Petersson, and Nick Barnes. Towards open-set object detection and discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3961–3970, 2022.
- [54] Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10867–10875, 2021.