

Deep Low-Shot Learning for Biological Image Classification and Visualization From Limited Training Samples

Lei Cai, Zhengyang Wang^{ID}, Rob Kulathinal, Sudhir Kumar, and Shuiwang Ji^{ID}, *Senior Member, IEEE*

Abstract—Predictive modeling is useful but very challenging in biological image analysis due to the high cost of obtaining and labeling training data. For example, in the study of gene interaction and regulation in *Drosophila* embryogenesis, the analysis is most biologically meaningful when *in situ* hybridization (ISH) gene expression pattern images from the same developmental stage are compared. However, labeling training data with precise stages is very time-consuming even for developmental biologists. Thus, a critical challenge is how to build accurate computational models for precise developmental stage classification from limited training samples. In addition, identification and visualization of developmental landmarks are required to enable biologists to interpret prediction results and calibrate models. To address these challenges, we propose a deep two-step low-shot learning framework to accurately classify ISH images using limited training images. Specifically, to enable accurate model training on limited training samples, we formulate the task as a deep low-shot learning problem and develop a novel two-step learning approach, including data-level learning and feature-level learning. We use a deep residual network as our base model and achieve improved performance in the precise stage prediction task of ISH images. Furthermore, the deep model can be interpreted by computing saliency maps, which consists of pixel-wise contributions of an image to its prediction result. In our task, saliency maps are used to assist the identification and visualization of developmental landmarks. Our experimental results show that the proposed model can not only make accurate predictions but also yield biologically meaningful interpretations. We anticipate our methods to be easily generalizable to other biological image classification tasks with small training datasets. Our open-source code is available at <https://github.com/divelab/lsl-fly>.

Index Terms—Biological image classification, deep two-step low-shot learning, *Drosophila in situ* hybridization (ISH) images, limited training samples, model interpretation and visualization.

Manuscript received May 11, 2020; revised December 6, 2020 and July 17, 2021; accepted August 12, 2021. This work was supported in part by the National Science Foundation under Grant IIS-1908220 and Grant DBI-2028361. (Corresponding author: Shuiwang Ji.)

Lei Cai is with the School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164 USA (e-mail: lei.cai@wsu.edu).

Zhengyang Wang and Shuiwang Ji are with the Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77843 USA (e-mail: zhengyang.wang@tamu.edu; sji@tamu.edu).

Rob Kulathinal and Sudhir Kumar are with the Department of Biology, Temple University, Philadelphia, PA 19122 USA (e-mail: robkulathinal@temple.edu; s.kumar@temple.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TNNLS.2021.3106831>.

Digital Object Identifier 10.1109/TNNLS.2021.3106831

I. INTRODUCTION

IN BIOLOGICAL image analysis tasks, annotation of training images usually requires specific domain knowledge and thus is commonly performed by expert biologists. This manual practice and the incurred cost limit the number of labeled training samples. Therefore, a recurring theme in biological image analysis is how to enable efficient and accurate model training from limited labeled training samples [1]–[4]. For example, comparative analysis of *in situ* hybridization (ISH) gene expression pattern images is a key step in studying gene interactions and regulations in *Drosophila* embryogenesis. Advances in imaging technologies have enabled biologists to collect an increasing number of *Drosophila* embryonic ISH images [5]–[8]. Currently, more than 100 000 *Drosophila* ISH images are available for investigating the functions and interconnections of genes [9]–[12]. The analysis of gene expression patterns is most biologically meaningful when images from the same developmental stage are compared [13]–[15]. Moreover, identification and visualization of developmental landmarks in each stage are required to enable biologists to interpret prediction results.

However, ISH images obtained from high-throughput experiments are commonly annotated with stage range labels, as shown in Fig. 1, since determining the precise stage for each ISH image is very hard and time-consuming even for expert biologists. Currently, most of the comparative analysis is limited to stage range levels due to the lack of accurate and cost-effective methods for assigning precise stage labels to images. In [22], a small training set is manually labeled with precise stages by developmental biologists, and a computational pipeline is proposed to train classifiers from this small training set. Specifically, predefined Gabor filters are used to compute features from ISH images [11], [17] and linear classifiers are employed to predict the precise stages of ISH images. In that work, no attempt is made to explicitly consider and account for the key bottleneck of limited training samples. Due to the small size of training set and the methods used, the predictive performance is relatively poor in that work. In order for the results to be usable by routine biological studies, an accurate computational model that can be trained effectively from limited training samples is highly desirable [18]–[21].

We notice that the Gabor filter features used in [22] are computed by convolving a set of predefined filters to




Fig. 1. Examples of *Drosophila* ISH images in different stage ranges. The images are annotated with stage range labels 1–3, 4–6, 7 and 8, 9 and 10, 11 and 12, and 13–17 [10], [16].

input images. Although the predefined Gabor filters have been widely used, the filter weights are precomputed and extract features that do not adapt to each task and data specifically. In order to enable automated data-driven and task-related feature extraction, we propose to employ deep convolutional neural networks (DCNNs) for this task [23]. However, DCNNs require a large number of training samples to estimate the filter weights, while most biological applications, including the precise stage prediction task of ISH images that we are considering, have only very limited training samples. Thus, straightforward use of CNNs inevitably results in poor performance.

In this work, we propose a deep two-step low-shot learning framework that consists of a set of novel computational techniques to enable effective deep model training on a small number of annotated training samples. Although we focus on ISH image classification and visualization in this work, the proposed approach is generic and can be easily generalizable to other biological image analysis tasks. Specifically, we formulate the original classification task with limited training samples as a deep low-shot learning problem [24]–[26] and propose a deep two-step low-shot learning approach that consists of data-level learning and feature-level learning. The first data-level learning step leverages extra data to facilitate the training. In the precise stage prediction task of ISH images, we can access a large number of extra ISH images with stage range labels, which cannot be directly used to train a precise stage prediction model. However, these extra ISH images are obtained from the same experimental pipeline and follow the same distribution as images in the original small training set. Therefore, they can be exploited to guide the training. The second step of our low-shot learning method is feature-level learning. We propose a regularization method that forces a high similarity of features extracted from different samples in the same class. It results in a high-quality feature space, enabling the extracted features to generalize well on unseen data. In order to achieve the regularization for classification, we generate reference sets that contain representative samples of each class. The proposed feature-level learning step is to fine-tune the model by adding a similarity loss based on cosine similarities between features extracted from the training sample and reference samples. The proposed deep two-step low-shot learning framework enables us to effectively train deep learning models with high classification accuracy. In addition, we explore the interpretation of deep learning models through saliency maps [27], which indicates pixel-wise

contributions of input images to the prediction results. In the precise stage prediction task of ISH images, we generate masked genomewide expression maps (GEMs) [28], [29] based on saliency maps. The masked GEMs provide biologically meaningful visualizations that help identifying and visualizing developmental landmarks.

II. DEEP TWO-STEP LOW-SHOT LEARNING FRAMEWORK

In this section, we introduce our deep two-step low-shot learning framework. We start with the motivation, advantages, and challenges of applying deep learning on biological image classification in Section II-A. To overcome the challenges, we formulate the task as a deep low-shot learning problem and propose the deep two-step low-shot learning approach to address it. The two steps, namely data-level learning and feature-level learning, are described in Sections II-B and II-C, respectively. In Section II-D, we apply the proposed method with a deep residual network (ResNet) on the precise stage prediction task of ISH images. We also introduce the interpretation method for deep learning models and how to use it to generate biologically meaningful visualizations for our task in Section II-E.

A. Deep Learning for Biological Image Classification

Many traditional biological image analyses involve a handcrafted feature extraction step, which requires expertise and only focuses on specific types of features. An example is using the Gabor filter to obtain texture features for the stage prediction task of ISH images [22]. While handcrafted features achieve acceptable performance in practice, they have two main disadvantages. First, it is expensive and time-consuming to design these feature extractors. There are a variety of biological images with different tasks, each of which needs a redesign of feature extraction methods accordingly. Second, the performance is limited by human’s knowledge. However, more accurate analysis is needed to improve our understanding.

Deep learning has provided a promising way to perform automatic feature extraction in a data-driven and task-related manner [30]. Instead of relying on handcrafted feature extractors, deep neural networks can learn to extract features through training on annotated data. Specifically, experts are only required to analyze and annotate a limited amount of data manually. Then, deep neural networks are able to determine

features that are important to the task and design corresponding feature extractors. Based on them, the trained networks can perform the task on unlabeled data automatically in an efficient and accurate way.

A key challenge of applying deep learning for biological image analysis is that the number of available annotated data is very limited, while deep learning models typically require a large amount of training data. It is usually cost-effective to perform coarse annotation on large-scale data. Fine annotation with high accuracy can only be done on a very small portion of data due to practical limitations of resources. It hinders the wide use of deep learning.

To address the challenge, we formulate it as a deep low-shot learning problem. In this problem, a large dataset with coarse labels and a small dataset with fine labels are available. To be specific, we denote the large dataset as $\mathcal{D}_t = \{(x_i^t, y_i^t)\}_{i=1}^{n_t}$ and the small dataset as $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$, where $y_i^t \in \{1, 2, \dots, l_t\}$ and $y_i^s \in \{1, 2, \dots, l_s\}$, respectively. Here, n_t and n_s represent the number of annotated data in each dataset and we have $n_t \gg n_s$. Correspondingly, l_t and l_s represent the number of classes in coarse and fine annotations, respectively. In this setting, it is reasonable to assume $l_t < l_s$. Note that $y_i^t = 1$ does not have the same meaning as $y_i^s = 1$, since coarse and fine annotations have different label spaces. The data samples in both datasets, x_i^t and x_i^s , are from the same distribution. The task is to develop accurate deep learning models to perform fine annotation. However, the small n_s prevents us to directly train a deep neural network on \mathcal{D}_s , making it necessary to develop a method that is able to effectively leverage \mathcal{D}_t . To solve this problem, we propose a method combining data-level learning and feature-level learning, leading to our deep two-step low-shot learning approach.

B. Data-Level Deep Low-Shot Learning

The first step of our proposed approach is data-level learning. Namely, we propose a way to leverage the large dataset \mathcal{D}_t and learn from more data, even when these data have coarse labels only. To introduce our method, we first review the deep classification networks.

A deep classification network is typically composed of a feature extraction network and a classification network. The feature extraction network corresponds to the handcrafted feature extraction step in traditional methods and the classification network performs the task based on extracted features. We focus on the feature extraction network as extracting features of high quality is the key to achieving satisfactory classification performance. Based on good features, the task can be well solved by using simple machine learning methods, such as nearest neighbor algorithm and linear classification methods [31]. Actually, the classification networks in modern deep classification networks [23], [32] are usually equivalent to multiclass logistic regression. The classification network simply serves as a guide for feature extraction network to extract task-related features from given data.

Note that, in the deep low-shot learning problem defined in Section II-A, x_i^t and x_i^s are from the same distribution. In addition, y_t and y_s are labels of two different accuracy

levels for the same classification task. With the insights above, we point out that training a deep classification network for the coarse annotation task on \mathcal{D}_t is able to provide an informative and useful feature extraction network for the ultimate fine annotation task.

Therefore, our data-level learning step proposes to first train a deep classification network on \mathcal{D}_t , where a large amount of annotated data are available. Specifically, the deep classification network consisting of a feature extraction network $F(\cdot; \theta)$ and a classification network $G_t(\cdot; \beta_t)$ is denoted as $G_t(F(\cdot; \theta); \beta_t)$, where θ and β_t represent the parameters in $F(\cdot)$ and $G_t(\cdot)$, respectively. Note that $G_t(\cdot; \beta_t)$ performs classification based on coarse labels, whose output layer has l_t nodes. After training $G_t(F(\cdot; \theta); \beta_t)$ on \mathcal{D}_t , we build another deep classification network $G_s(F(\cdot; \theta); \beta_s)$ to perform fine annotation. In particular, we employ the same extraction network $F(\cdot; \theta)$ and simply change the classification network to $G_s(\cdot; \beta_s)$ whose output layer has l_s nodes. Finally, we train $G_s(F(\cdot; \theta); \beta_s)$ on \mathcal{D}_s . In terms of training, we let the outputs of the classification network to go through a softmax function and compute the cross-entropy loss for backpropagation [30]. Taking $G_s(F(\cdot; \theta); \beta_s)$ as an example, for one training sample (x^s, y^s) in \mathcal{D}_s , the softmax function is defined as

$$p_i = \frac{\exp(o_i)}{\sum_{j=1}^{l_s} \exp(o_j)}, \quad i = 1, 2, \dots, l_s \quad (1)$$

where o_i denote the output of the i th node, i.e.,

$$o_i = (G_s(F(x^s; \theta); \beta_s))_i. \quad (2)$$

Here, p_i can be interpreted as the probability that the input image x^s belongs to the i th class. The cross-entropy loss can be computed by

$$\mathcal{L}_{\text{CE}} = -\log(p_{y^s}) \quad (3)$$

which simply depends on the predicted probability that x^s belongs to the y^s th class.

The key step of data-level learning is to apply the same feature extraction network in the process. Intuitively, such an approach enables the task-related knowledge obtained from \mathcal{D}_t to be transferred to \mathcal{D}_s . Technically, training $F(\cdot; \theta)$ on \mathcal{D}_t provides a good initialization for training $G_s(F(\cdot; \theta); \beta_s)$ on \mathcal{D}_s , alleviating the problem of lacking enough training data for the fine annotation task.

C. Feature-Level Deep Low-Shot Learning

The second step of our proposed approach is feature-level learning, where we propose a regularization method to force a high similarity of features extracted from different samples in the same class. To achieve it, we select several samples from each class as reference samples and regularize the feature extraction network to extract similar features from other samples in the same class. We first describe how to select reference samples and then introduce the regularization method.

After the data-level learning step described in Section II-B, we obtain a deep classification network $G_s(F(\cdot; \theta); \beta_s)$ that has been trained on \mathcal{D}_s . Using the softmax function in (1),

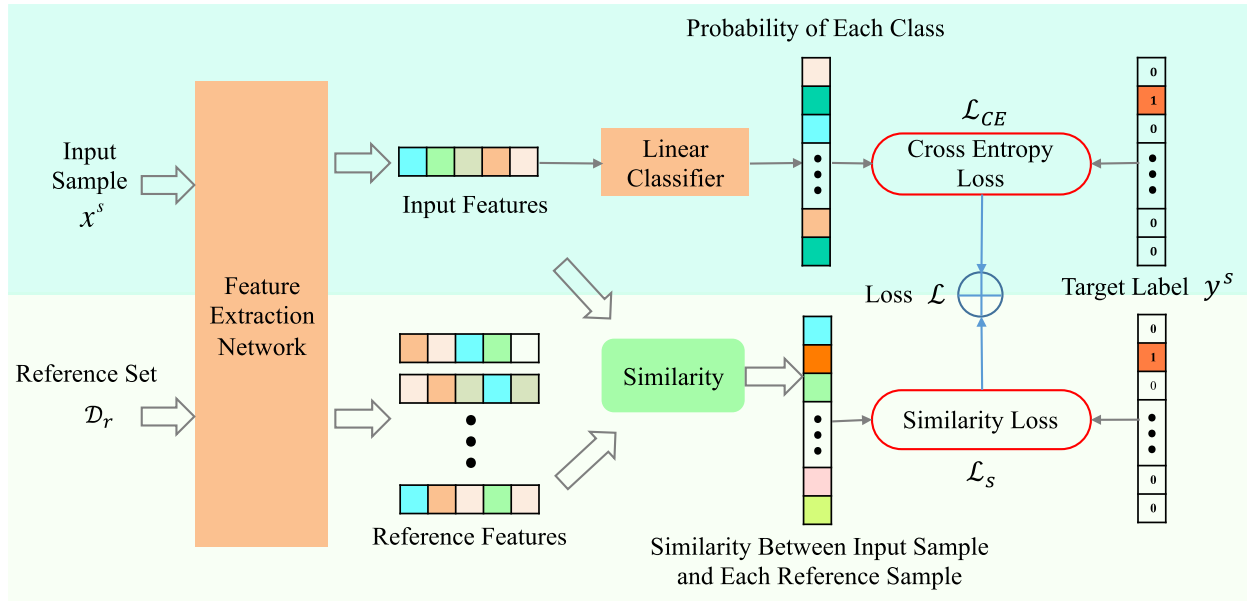


Fig. 2. Illustration of feature-level deep low-shot learning. A training sample and a reference set are fed into the feature extraction network to generate features. We compute the cosine similarity of features between the input sample and each reference sample and propose the similarity loss \mathcal{L}_S . Then, the similarity loss \mathcal{L}_S and the cross-entropy loss \mathcal{L}_{CE} are combined to train the network.

for each training sample (x^s, y^s) in \mathcal{D}_s , we can compute the probability p_{y^s} that x^s belongs to y^s . We set a threshold τ and select those samples with $p_{y^s} > \tau$. The setting of τ should make sure that $p_{y^s} > p_i, \forall i \in \{1, 2, \dots, l_s\} \setminus \{y^s\}$. Specifically, for each class $c = 1, 2, \dots, l_s$, we obtain a subset of \mathcal{D}_s defined as $\mathcal{D}_{s,c} = \{(x_c^s, c) \mid p_c > p_i, \forall i \in \{1, 2, \dots, l_s\} \setminus \{c\} \text{ and } p_c > \tau\}$. Note that these subsets $\mathcal{D}_{s,c}, c = 1, 2, \dots, l_s$ are mutually exclusive. Then, we randomly select one sample from each $\mathcal{D}_{s,c}$ to form a reference set $\mathcal{D}_r = \{(x_c^r, c)\}_{c=1}^{l_s}$, which contains exactly l_s reference samples. The last step is iterated for k times to form k reference sets $\mathcal{D}_{r_i}, i = 1, 2, \dots, k$.

The feature-level learning step follows the data-level learning by fine-tuning $G_s(F(\cdot; \theta); \beta_s)$ on \mathcal{D}_s with the help of reference sets $\mathcal{D}_{r_i}, i = 1, 2, \dots, k$. When fine-tuning with one training sample (x^s, y^s) in \mathcal{D}_s , we randomly pick a reference set $\mathcal{D}_r = \{(x_c^r, c)\}_{c=1}^{l_s}$ from $\mathcal{D}_{r_i}, i = 1, 2, \dots, k$. The feature-level learning aims to force the features extracted from input images in the same class to be similar, i.e., $F(x^s; \theta)$ and $F(x_{y^s}^r; \theta)$ should be similar. In opposite, for $c \in \{1, 2, \dots, l_s\} \setminus \{y^s\}$, $F(x^s; \theta)$ and $F(x_c^r; \theta)$ should be less similar. It requires a quantitative way to evaluate the similarity between features. In this work, we quantify the similarity by calculating

$$\begin{aligned} \text{sim}(x^s, x_c^r) &= \text{cosine}(F(x^s), F(x_c^r)) \\ &= \frac{F(x^s) \cdot F(x_c^r)}{\|F(x^s)\| \times \|F(x_c^r)\|}. \end{aligned} \quad (4)$$

For each $(x_c^r, c) \in \mathcal{D}_r$, we can compute the corresponding $\text{sim}(x^s, x_c^r)$. Then, we normalize the similarities through a softmax function

$$\text{SIM}(x^s, x_c^r) = \frac{\exp(\text{sim}(x^s, x_c^r))}{\sum_{i=1}^{l_s} \exp(\text{sim}(x^s, x_i^r))}, \quad c = 1, 2, \dots, l_s \quad (5)$$

where a large $\text{SIM}(x^s, x_c^r)$ means that $F(x^s; \theta)$ and $F(x_c^r; \theta)$ are similar. Finally, we propose the similarity loss

$$\mathcal{L}_S = -\text{SIM}(x^s, x_{y^s}^r) + \frac{1}{l_s - 1} \sum_{c \in \{1, 2, \dots, l_s\} \setminus \{y^s\}} \text{SIM}(x^s, x_c^r). \quad (6)$$

Minimizing the similarity loss \mathcal{L}_S will result in maximizing $\text{SIM}(x^s, x_{y^s}^r)$ while minimizing $\text{SIM}(x^s, x_c^r)$ for $i \in \{1, 2, \dots, l_s\} \setminus \{y^s\}$. This corresponds to our purpose of forcing features extracted from input images in the same class to be similar as well as making those in different classes distinguishable. Therefore, the similarity loss \mathcal{L}_S serves as a feature-level regularization method. When fine-tuning the deep classification network in the feature-level learning step, we apply a combined loss function defined as

$$\mathcal{L} = \mathcal{L}_{CE} + \mathcal{L}_S \quad (7)$$

where \mathcal{L}_{CE} is defined by (3). In the feature-level deep low-shot learning, all parameters feature extraction networks θ and classification networks β_s are updated by optimizing the loss function (7). In the feature-level learning, cross-entropy loss is still employed in the loss function, which can guarantee that the knowledge learned in the data level learning cannot be forgotten.

An overview of the proposed feature-level learning step is shown in Fig. 2. The feature-level learning provides better features for the fine annotation task and improves $G_s(F(\cdot; \theta); \beta_s)$.

D. Deep Two-Step Low-Shot Learning for the Precise Stage Prediction of ISH Images

Combining data-level learning and feature-level learning leads to our proposed deep two-step low-shot learning approach, as shown in Algorithm 1. We apply this approach to solve the precise stage prediction task of *Drosophila* ISH

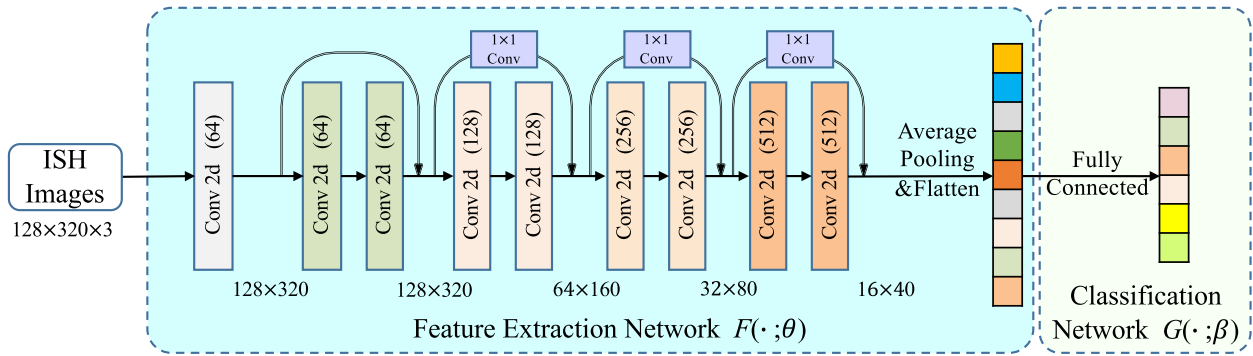


Fig. 3. Architecture of our deep ResNet in the precise stage prediction task of ISH images. The feature extraction network $F(\cdot; \theta)$ contains four residual blocks. An average pooling with a kernel size of 4×4 is employed after the last residual block. As shown in Algorithm 1, the same $F(\cdot; \theta)$ is employed through the whole two-step process. The classification network $G(\cdot; \beta)$ refers to $G_t(\cdot; \beta_t)$ and $G_s(\cdot; \beta_s)$ in the corresponding steps.

Algorithm 1 Deep Two-Step Low-Shot Learning for Biological Image Classification

Input:

Dataset with Coarse Labels $\mathcal{D}_t = \{(x_i^t, y_i^t)\}_{i=1}^{n_t}$ where $y_i^t \in \{1, \dots, l_t\}$
 Dataset with Fine Labels $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$ where $y_i^s \in \{1, \dots, l_s\}$
 Threshold τ

Step 1 Data-Level Low-shot Learning:

1. Train $G_t(F(\cdot; \theta); \beta_t)$ on \mathcal{D}_t .
2. Train $G_s(F(\cdot; \theta); \beta_s)$ with the same feature extraction network $F(\cdot; \theta)$ on \mathcal{D}_s .

Step 2: Feature-Level Low-shot Learning:

3. Construct reference sets \mathcal{D}_{r_i} , $i = 1, 2, \dots, k$ from \mathcal{D}_s using $G_s(F(\cdot; \theta); \beta_s)$ and τ .
 4. Fine-tune $G_s(F(\cdot; \theta); \beta_s)$ and update parameters θ and β_s using the loss function defined in Equation (7).
-

images [22]. In this task, the large dataset with coarse labels \mathcal{D}_t refers to the ISH images with stage range labels, as shown in Fig. 1. These labels only indicate that the developmental stage of an ISH image falls into one of the six ranges. Meanwhile, a small dataset with fine labels \mathcal{D}_s is available, where the data are annotated with precise stages. Here, $l_t = 6$ and $l_s = 14$. The precise stage prediction task can be well formulated as a deep low-shot learning problem and addressed by our proposed method.

DCNNs [30] have achieved great success in various image tasks. Convolutional layers are proved to be capable of performing effective feature extraction. Therefore, we use DCNNs as our deep learning model for the precise stage prediction task. In particular, we apply a deep ResNet by adopting the technique of adding residual connections proposed in [23]. Residual connections are known to facilitate the training of deep neural networks.

Fig. 3 shows our deep ResNet for the stage prediction task. Our feature extraction network $F(\cdot; \theta)$ contains four residual blocks, each of which is composed of two convolutional layers and a residual connection. A convolutional layer contains a convolution followed by a batch normalization [33] and a

rectified linear unit (ReLU) activation function [34]. A residual connection is employed to add the inputs to the outputs for two consecutive convolutional layers and forms a residual block. The residual connection is simply an identity connection for the first residual block. For the latter three residual blocks, we set the stride to 2 for the first convolutional layer, resulting in outputs of reduced spatial sizes. In this case, we use a 1×1 convolutional layer in the residual connection, which changes the spatial sizes of inputs accordingly to accommodate the addition. We apply an average pooling with a kernel size of 4×4 to obtain the extracted features after the last residual block. The classification networks $G_t(\cdot; \beta_t)$ and $G_s(\cdot; \beta_s)$ are designed as a single fully connected layer with 6 and 14 output nodes, respectively. Our deep learning model is trained by Algorithm 1 and outperforms the previous state-of-the-art model [22] significantly, as shown in Section IV.

E. Interpretation and Visualization Using Saliency Maps

Our proposed deep two-step low-shot learning approach yields deep learning models with high accuracy. We further introduce an interpretation method for deep learning models based on saliency maps. Saliency maps quantify the pixel-wise contributions of an input image to its prediction result by computing the gradient of the prediction outputs with respect to the input image [27]. To be specific, to interpret our deep classification network $G_s(F(\cdot; \theta); \beta_s)$ with input image x^s , we compute the saliency maps for each output node

$$S(x^s, i) = \frac{\partial o_i}{\partial x^s}, \quad i = 1, 2, \dots, l_s \quad (8)$$

where o_i is defined in (2). Suppose that the prediction result of x^s is y^s . We pay attention to $S(x^s, y^s)$ as it tells which pixels in x^s support the prediction result most. In this work, we use the guided backpropagation positive saliency map for visualization [35], where only positive gradients are preserved.

In the precise stage prediction task of ISH images, we apply saliency maps to generate masked GEMs to help identifying and visualizing developmental landmarks. GEMs provide a visualization method for precise stage prediction models and a visualization method for developmental landmark analysis.

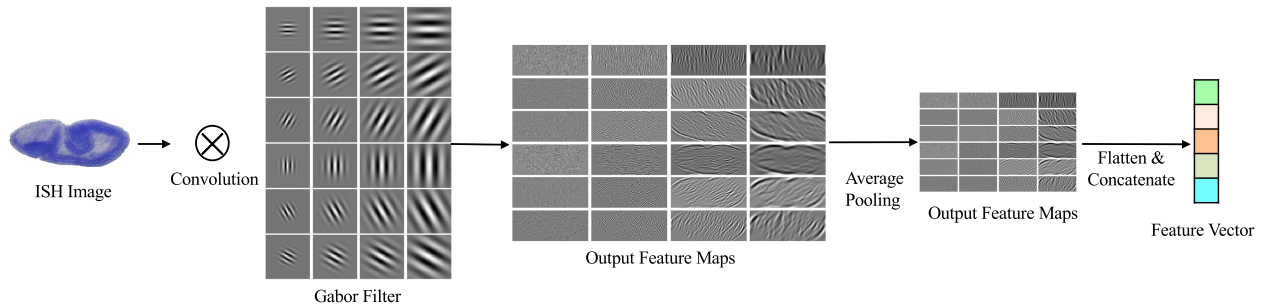


Fig. 4. Illustration of precise stage prediction using Gabor filters [22]. In this approach, 24 filters are generated with four different wavelet scales and six different filter orientations. The average pooling layer with kernel size 8×8 is employed to reduce the spatial size of the feature maps and extract the invariant features. Finally, the extracted features are flattened and concatenated together.

Originally, GEMs are generated by aggregating and normalizing the ISH images from the same stage [28], [29]. We propose to use saliency maps as masks in generating GEMs, selecting characteristic pixels from ISH images that contribute to the prediction results. We show that the resulted masked GEMs can provide more biologically meaningful visualizations in Section IV.

III. RELATED WORK

A. Precise Stage Prediction

Currently, the state-of-the-art method [22] for the precise stage prediction task of ISH images is to extract features using predefined convolutional filters and train a linear classifier based on the extracted features. Texture features of different scales and orientations can be computed with different filters. The filter weights are constructed by using log Gabor filters [11], [17]. Given a set of predefined wavelet scales and orientations, we can generate corresponding Gabor filters to extract the specific features by convolving on input images. The precise stage prediction model in [22] based on Gabor filter features is shown in Fig. 4.

The Gabor filter features focus on subtle textures, which enables the classification model to distinguish ISH images from different stages. However, there are still some limitations for Gabor filter features as Gabor filters are handcrafted and cannot learn features automatically for different tasks and datasets. Moreover, Gabor filters only extract low-level texture features. Extracting features hierarchically from low level to high level is appealing for improving the precise stage prediction accuracy.

B. Few-/Low-Shot Learning

Few-shot learning is proposed to solve classification or regression problems with very limited training samples. Specifically, we consider a K -way- N -shot classification problem, where K is the number of classes and N is the number of training samples in each class. Therefore, in the K -way- N -shot classification problem, the training dataset contains $K \times N$ samples. In few-shot learning problems, N is commonly very small like 1 or 5. Recently, few-/low-shot learning methods were proposed to overcome the challenge of limited training samples using meta-learning [36], [37],

multitask learning, learn-to-learn [38], embedding [31], [39], and other technologies [40]–[42]. For meta-learning-based methods, the model randomly selects $K \times N$ samples from training sets as support set for meta-task in each epoch. The model can learn the shared knowledge from different meta-tasks. Therefore, the model can perform well when dealing with a new meta-task. For multitask-based methods, we have several highly related tasks and use other tasks to improve the performance of the model on the task with few training samples. Based on the strategy of learning parameters, multitask-based methods can be categorized into two groups: parameter sharing and parameter tying. For parameter sharing, the model shares the first few layers from other tasks and learns the final layers for the target task. For parameter tying, parameters are constrained to be similar for different tasks. In embedding-based learning, the model learns general embedding from the training dataset and computes the similarity between training and test samples. In this work, we consider a different task from the traditional K -way- N -shot classification problem. In our problem, a large dataset with coarse labels and a small dataset with fine labels are provided. The training size is larger than that used in traditional few-shot learning but still limited for deep learning models to achieve satisfying performance. We aim at developing a model to predict the fine label while taking advantage of samples with coarse labels in the learning procedure.

IV. EXPERIMENTAL STUDIES

In this section, we evaluate our proposed deep two-step low-shot learning framework for the precise stage prediction task of ISH images on the Berkeley Drosophila Genome Project (BDGP) dataset. We perform the qualitative evaluation in terms of the prediction accuracy as well as qualitative evaluation based on feature visualization and GEMs [28], [29]. Finally, we show the saliency maps and masked GEMs for interpretation and visualization.

A. Experimental Setup

The BDGP dataset contains 36908 *Drosophila* ISH images in the lateral view. Among them, only 3474 ISH images are annotated with precise stage labels. In addition, all images are labeled into six stage ranges: 1–3, 4–6, 7 and 8, 9 and

TABLE I
OVERVIEW OF THE DATASET USED IN THIS STUDY

Stage	# images with precise stage labels	# images with stage range labels	Stage	# images with precise stage labels	# images with stage range labels
1-3	0	4203	11	246	9248
4	251	7344	12	255	
5	274		13	251	
6	224		14	252	8680
7	236	15	232		
8	260	16	243		
9	248	17	254		
10	248	3768	Total	3474	36908

TABLE II

COMPARISON OF PERFORMANCE BETWEEN DIFFERENT METHODS IN TERMS OF PRECISE STAGE PREDICTION ACCURACY. THE FIRST MODEL IS OUR BASELINE [22]. THE LAST THREE ROWS CORRESPOND TO RESNET MODELS TRAINED USING THREE DIFFERENT APPROACHES: TRAIN IT DIRECTLY ON \mathcal{D}_s (PLAIN RESNET), TRAIN IT USING DATA-LEVEL LOW-SHOT LEARNING ONLY (DATA-LEVEL), AND TRAIN IT USING BOTH DATA-LEVEL AND FEATURE-LEVEL LOW-SHOT LEARNING (DATA&FEATURE-LEVEL)

Model	Stage 4	Stage 5	Stage 6	Stage 7	Stage 8	Stage 9	Stage 10	Stage 11
Gabor Filter [17]	84%	83%	91%	95%	82%	62%	64%	85%
Plain ResNet	81%	77%	81%	90%	80%	57%	53%	83%
Data-level	76%	91%	85%	99%	84%	66%	66%	86%
Data&Feature-level	81%	89%	87%	98%	87%	75%	64%	84%

Model	Stage 12	Stage 13	Stage 14	Stage 15	Stage 16	Stage 17	Average
Gabor Filter [17]	80%	94%	74%	86%	73%	80%	80.97%
Plain ResNet	80%	89%	77%	77%	72%	80%	77.28%
Data-level	88%	92%	84%	74%	71%	84%	82.40%
Data&Feature-level	89%	94%	79%	81%	74%	87%	83.92%

10, 11 and 12, and 13–17, as shown in Fig. 1. According to Section II-D, we have \mathcal{D}_t with $n_t = 36908$ and \mathcal{D}_s with $n_s = 3474$. The details are listed in Table I. Note that, since the *Drosophila* embryogenesis is not very meaningful from stage 1 to stage 3, the precise stage prediction task does not involving this range. Therefore, there is no precise stage annotation for this range. All ISH images are standardized by the pipeline proposed in [28], and the size of processed images is 128×320 . To perform evaluation, we randomly select 80% of images \mathcal{D}_s for training and test our proposed model on the rest images.

We use the previous state-of-the-art model proposed in [22] that uses Gabor filters as our baseline. For deep learning, we use the ResNet introduced in Section II-D. To show the effectiveness of our proposed deep two-step low-shot learning approach, we compare the performance of our ResNet trained using three different approaches: train it directly on \mathcal{D}_s (Plain ResNet), train it using data-level low-shot learning only (Data-level), and train it using both data-level and feature-level low-shot learning (Data&Feature-level). For all training processes, we employ the stochastic gradient descent algorithm [30]. The learning rate starts with 0.001 and is multiplied by 0.1 at the 20th and 30th epochs. The batch size is set to 16.

B. Quantitative Analysis of Predictive Performance

There exists randomness for neural network-based models due to different weight initialization. To guarantee that the comparison is fair between different models, we run each model five times and take the average accuracy as the evaluation metric. The precise stage prediction accuracy of each method is reported in Table II. Note that the accuracy

of Plain ResNet is lower than that of our baseline model, proving that deep learning models can hardly achieve satisfactory performance with limited training samples. Training the ResNet with our data-level deep low-shot learning method addresses this problem effectively and improves the performance significantly. By adding the feature-level learning step, our proposed deep two-step low-shot learning model achieves better predictive performance and becomes the new state-of-the-art model for the precise stage prediction task of ISH images.

In order to analyze the prediction errors of our model, we compare the confusion matrices of the prediction results. The confusion matrix is obtained based on the prediction results. We randomly choose one model from five experiments for prediction and generate the confusion matrix, as shown in Fig. 5. We observe that most wrongly predicted samples are classified into neighboring stages, within the correct stage ranges. In particular, for deep learning models trained using our proposed low-shot learning methods, the number of samples that are categorized into wrong stage ranges is fewer than the baseline. It indicates that the data-level deep low-shot learning step effectively incorporates information from samples with stage range labels and improves the performance of precise stage prediction.

C. Qualitative Analysis of Predictive Performance

We perform two qualitative analyses on the baseline and the proposed deep two-step low-shot learning model. First, we compare the features extracted by Gabor filters with those generated by the first convolutional layer in the proposed model. Fig. 6 shows the visualizations of the features.

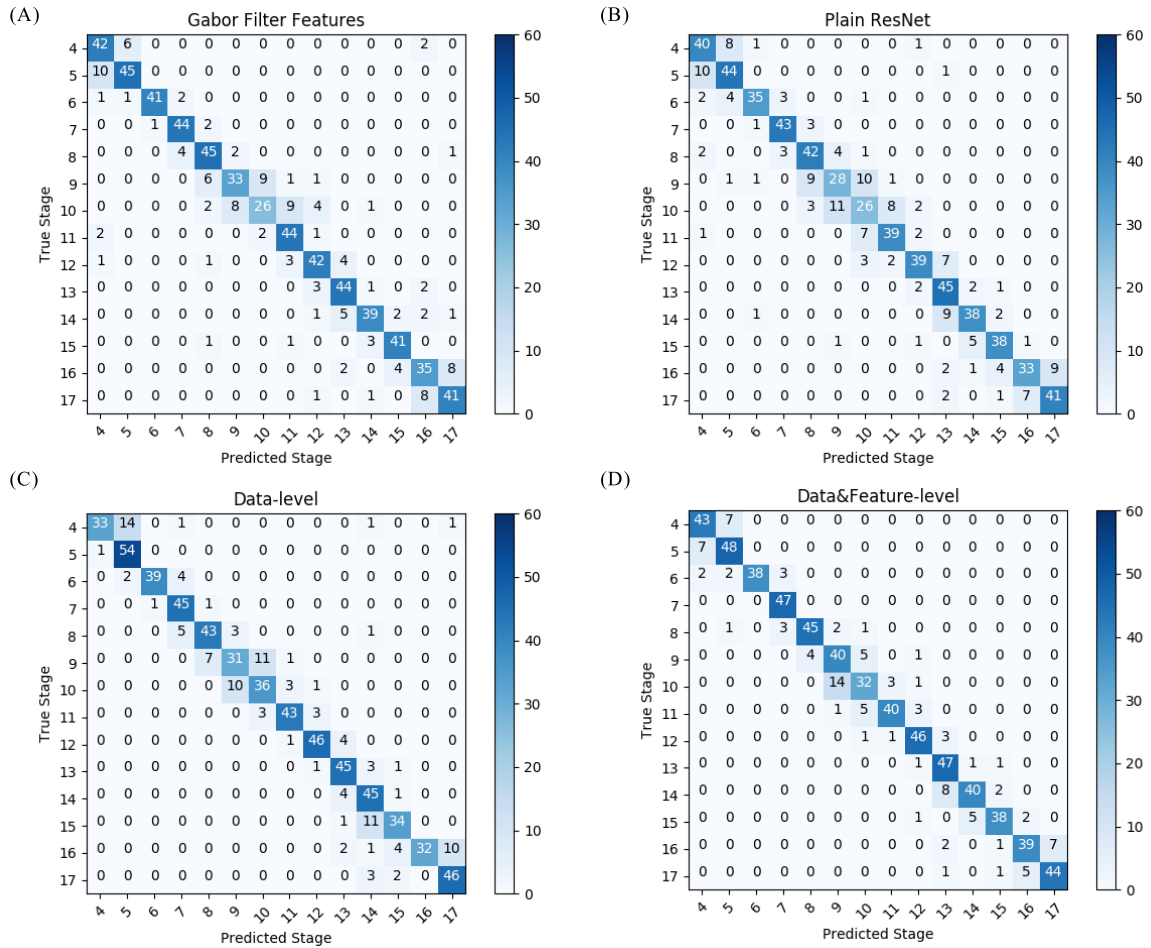


Fig. 5. Comparison of performance between different methods in terms of confusion matrices. (a) Our baseline [22]. (b)–(d) ResNet models trained using three different approaches: train it directly on \mathcal{D}_s (Plain ResNet), train it using data-level low-shot learning only (Data-level), and train it using both data- and feature-level low-shot learning (Data&Feature-level), respectively.




Fig. 6. Comparison between feature maps generated by Gabor filters in [22] and those created by our proposed deep two-step low-shot learning model.

Clearly, Gabor filters can only extract predefined low-level texture features, while our deep learning model is able to extract data-driven and task-related features automatically. The Gabor filters contain hard-wired weights to compute texture features at different scales and orientations. Thus, it is clear from the visualizations that the resulting features are embryonic patterns with different scales and orientations. These features are not adapted to the image dataset. In comparison,

the features computed by our deep learning methods are adapted to the stage classification task in which shapes and boundaries of embryos are important. Thus, we can clearly see the boundaries of embryos in the features. In addition, we can see that the germ band has been highlighted in some of the deep learning computed features, as germ-band movements are important features to distinguish different developmental stages.

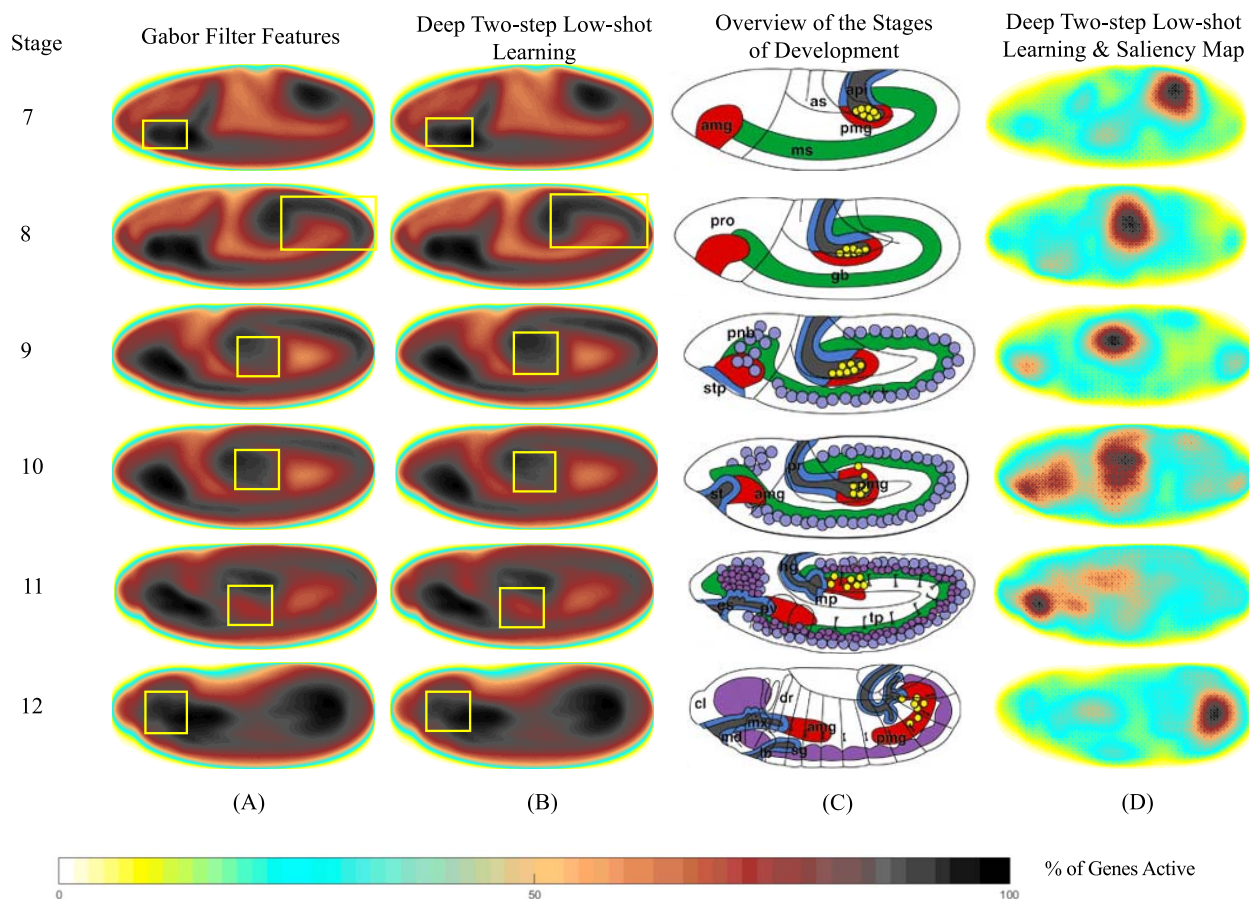


Fig. 7. Visualizations of GEMs and masked GEMs generated using different methods for stages 7–12. (a) GEMs generated using prediction results of the baseline. (b) GEMs generated using prediction results of our proposed two-step low-shot learning model. (c) Overview of the stages of development [43]. (d) Masked GEMs generated using prediction results of our proposed two-step low-shot learning model and corresponding saliency maps. The differences can be better observed when zooming in.




Fig. 8. Examples of saliency maps generated by our proposed two-step low-shot learning model. The first row is the original *Drosophila* ISH images. The second row is the corresponding saliency maps. White color indicates that the pixel contributes significantly to the prediction result of the image.

Second, we visualize GEMs [28], [29] generated by the baseline and our deep two-step low-shot learning model. Computing GEMs relies on the predictive performance and accurate precise stage prediction results produce more biologically meaningful GEMs. The first two columns in Fig. 7 compare GEMs generated using prediction results of the baseline and our model. According to the third column, we can see that the developing trend shown in both GEMs is consistent with that shown in [43]. In addition, since our proposed method achieves better predictive performance, the generated GEMs contain more biologically meaningful details as

demonstrated by yellow bounding boxes. The differences can be better observed when zooming in Fig. 7.

D. Visualization Through Saliency Maps and Masked GEMs

As introduced in Section II-E, we can compute saliency maps to measure the pixel-wise contributions of an input image to its prediction result. We visualize some saliency maps generated by our proposed two-step low-shot learning model in Fig. 8. From the saliency maps, we can infer which pixels contribute most to the prediction results. We propose to apply

saliency maps as masks and use these pixels to generate GEMs instead of using the whole input images. The resulted masked GEMs are shown in the last column in Fig. 7. The masked GEMs also show a consistent developing trend as shown in the third column. Moreover, they provide better visualizations to help identifying and visualizing developmental landmarks.

V. CONCLUSION

In this work, we propose a deep two-step low-shot learning framework and apply it on the precise stage prediction task of ISH images. We first propose to employ deep learning models to automatically extract data-driven and task-related features instead of using handcrafted feature extraction methods. In order to build accurate deep learning models from limited training samples, we formulate the task as a deep low-shot learning problem. To solve it, we propose a deep two-step low-shot learning approach that is composed of data-level learning and feature-level learning. The first data-level learning step leverages extra data to facilitate the training. The second step, feature-level learning, introduces a regularization method, which forces a high similarity of features extracted from different samples in the same class. The proposed deep two-step low-shot learning framework enables us to effectively train deep learning models with high classification accuracy. We conduct thorough experiments on the BDGP dataset using a deep ResNet as our base model. Both quantitative and qualitative experimental results demonstrate that our proposed deep two-step low-shot model outperforms the previous state-of-the-art model significantly. In addition, we explore visualization based on saliency maps [27], which computes pixel-wise contributions of input images to the prediction results. In the precise stage prediction task of ISH images, we apply saliency maps to generate masked GEMs, which provides biologically meaningful visualizations that help identifying and visualizing developmental landmarks [44].

REFERENCES

- [1] T. Zeng, R. Li, R. Mukkamala, J. Ye, and S. Ji, "Deep convolutional neural networks for annotating gene expression patterns in the mouse brain," *BMC Bioinf.*, vol. 16, no. 1, p. 147, 2015.
- [2] Q. Sun *et al.*, "Image-level and group-level models for Drosophila gene expression pattern annotation," *BMC Bioinf.*, vol. 14, no. 1, p. 350, 2013.
- [3] S. Ji, Y.-X. Li, Z.-H. Zhou, S. Kumar, and J. Ye, "A bag-of-words approach for Drosophila gene expression pattern annotation," *BMC Bioinf.*, vol. 10, no. 1, p. 119, 2009.
- [4] S. Ji, L. Yuan, Y.-X. Li, Z.-H. Zhou, S. Kumar, and J. Ye, "Drosophila gene expression pattern annotation using sparse features and term-term interactions," in *Proc. 15th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2009, pp. 407–416.
- [5] A. Cardona and P. Tomancak, "Current challenges in open-source bioimage informatics," *Nature Methods*, vol. 9, no. 7, p. 661, 2012.
- [6] T. Walter *et al.*, "Visualization of image data from cells to organisms," *Nature Methods*, vol. 7, no. 3, pp. S26–S41, Mar. 2010.
- [7] H. Peng, F. Long, J. Zhou, G. Leung, M. B. Eisen, and E. W. Myers, "Automatic image analysis for gene expression patterns of fly embryos," *BMC Cell Biol.*, vol. 8, no. 1, p. S1, Jul. 2007.
- [8] W. Zhang *et al.*, "A mesh generation and machine learning framework for Drosophila gene expression pattern image analysis," *BMC Bioinf.*, vol. 14, no. 1, p. 372, 2013.
- [9] E. Lécuyer *et al.*, "Global analysis of mRNA localization reveals a prominent role in organizing cellular architecture and function," *Cell*, vol. 131, no. 1, pp. 174–187, Oct. 2007.
- [10] P. Tomancak *et al.*, "Systematic determination of patterns of gene expression during Drosophila embryogenesis," *Genome Biol.*, vol. 3, no. 12, 2002, Art. no. research0088.
- [11] J. G. Daugman, "Two-dimensional spectral analysis of cortical receptive field profiles," *Vis. Res.*, vol. 20, no. 10, pp. 847–856, Jan. 1980.
- [12] M. Levine and E. H. Davidson, "Gene regulatory networks for development," *Proc. Nat. Acad. Sci. USA*, vol. 102, no. 14, pp. 4936–4942, 2005.
- [13] S. Kumar, K. Jayaraman, S. Panchanathan, R. Gurunathan, A. Marti-Subirana, and S. J. Newfeld, "BEST: A novel computational approach for comparing gene expression patterns from early stages of Drosophila melanogaster development," *Genetics*, vol. 162, no. 4, pp. 2037–2047, Dec. 2002.
- [14] E. Frise, A. S. Hammonds, and S. E. Celniker, "Systematic image-driven analysis of the spatial Drosophila embryonic expression landscape," *Mol. Syst. Biol.*, vol. 6, no. 1, p. 345, Jan. 2010.
- [15] H. Peng and E. W. Myers, "Comparing *in situ* mRNA expression patterns of Drosophila embryos," in *Proc. 8th Annu. Int. Conf. Comput. Mol. Biol.*, New York, NY, USA, 2004, pp. 157–166.
- [16] P. Tomancak *et al.*, "Global analysis of patterns of gene expression during Drosophila embryogenesis," *Genome Biol.*, vol. 8, no. 7, p. R145, 2007.
- [17] D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," *J. Opt. Soc. Amer. A, Opt. Image Sci.*, vol. 4, no. 12, pp. 2379–2394, 1987.
- [18] S. Ji, L. Sun, R. Jin, S. Kumar, and J. Ye, "Automated annotation of Drosophila gene expression patterns using a controlled vocabulary," *Bioinformatics*, vol. 24, no. 17, pp. 1881–1888, Sep. 2008.
- [19] J. Ye, J. Chen, Q. Li, and S. Kumar, "Classification of Drosophila embryonic developmental stage range based on gene expression pattern images," in *Computational Systems Bioinformatics*. Singapore: World Scientific, 2006, pp. 293–298.
- [20] X. Cai, H. Wang, H. Huang, and C. Ding, "Joint stage recognition and anatomical annotation of Drosophila gene expression patterns," *Bioinformatics*, vol. 28, no. 12, pp. 16–24, Jun. 2012.
- [21] V. Hartenstein, *Atlas Drosophila Development*. Cold Spring Harbor, NY, USA: Cold Spring Harbor Laboratory Press, 1995.
- [22] L. Yuan *et al.*, "Automated annotation of developmental stages of Drosophila embryos in images containing spatial patterns of expression," *Bioinformatics*, vol. 30, no. 2, pp. 266–273, Jan. 2014.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 770–778.
- [24] L. Fei-Fei, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 4, pp. 594–611, Apr. 2006.
- [25] B. Lake, R. Salakhutdinov, J. Gross, and J. Tenenbaum, "One shot learning of simple visual concepts," in *Proc. Annu. Meeting Cognit. Sci. Soc.*, 2011, vol. 33, no. 33, pp. 1–7.
- [26] Y.-X. Wang, R. Girshick, M. Hebert, and B. Hariharan, "Low-shot learning from imaginary data," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7278–7286.
- [27] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *Proc. Workshop Int. Conf. Learn. Represent.*, 2014, pp. 1–8.
- [28] C. E. Konikoff, T. L. Karr, M. McCutchan, S. J. Newfeld, and S. Kumar, "Comparison of embryonic expression within multigene families using the flyexpress discovery platform reveals more spatial than temporal divergence," *Develop. Dyn.*, vol. 241, no. 1, pp. 150–160, Jan. 2012.
- [29] S. Kumar *et al.*, "FlyExpress: Visual mining of spatiotemporal patterns for genes and publications in Drosophila embryogenesis," *Bioinformatics*, vol. 27, no. 23, pp. 3319–3320, Dec. 2011.
- [30] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [31] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 4077–4087.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *Computer Vision—ECCV 2016 (Lecture Notes in Computer Science)*, vol. 9908, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham, Switzerland: Springer, 2016, doi: [10.1007/978-3-319-46493-0_38](https://doi.org/10.1007/978-3-319-46493-0_38).

- [33] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [34] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [35] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.
- [36] J. Gordon, J. Bronskill, M. Bauer, S. Nowozin, and R. Turner, "Meta-learning probabilistic inference for prediction," in *Proc. Int. Conf. Learn. Represent.*, 2018, pp. 1–22.
- [37] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. 34th Int. Conf. Mach. Learn.*, vol. 70, Aug. 2017, pp. 1126–1135.
- [38] R. Salakhutdinov, J. B. Tenenbaum, and A. Torralba, "Learning to learn with compound hd models," in *Proc. 24th Int. Conf. Neural Inf. Process. Syst.*, 2011, pp. 2061–2069.
- [39] O. Vinyals *et al.*, "Matching networks for one shot learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 3630–3638.
- [40] L. Zhu and Y. Yang, "Label independent memory for semi-supervised few-shot video classification," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Jul. 7, 2020, doi: [10.1109/TPAMI.2020.3007511](https://doi.org/10.1109/TPAMI.2020.3007511).
- [41] A. A. Rusu *et al.*, "Progressive neural networks," 2016, *arXiv:1606.04671*. [Online]. Available: <http://arxiv.org/abs/1606.04671>
- [42] Y. Liu *et al.*, "Learning to propagate labels: Transductive propagation network for few-shot learning," in *Proc. Int. Conf. Learn. Represent.*, 2018, pp. 1–14.
- [43] V. Hartenstein, *Atlas of Drosophila Development*. Plainview, NY, USA: Cold Spring Harbor Laboratory Press, 1993.
- [44] H. Wang, C. Deng, H. Zhang, X. Gao, and H. Huang, "Drosophila gene expression pattern annotations via multi-instance biological relevance learning," in *Proc. 13th AAAI Conf. Artif. Intell.*, 2016, pp. 1324–1330.



Lei Cai received the Ph.D. degree in computer science from Washington State University, Pullman, WA, USA, in 2020.

He is currently a Data and Applied Scientist at Microsoft, Redmond, WA, USA. His research interests include data mining, machine learning, and deep learning.



Zhengyang Wang received the M.Sc. degree in mathematics and computer science from New York University, New York, NY, USA, in 2015, and the Ph.D. degree in computer science from Texas A&M University, College Station, TX, USA, in 2020.

He is currently an Applied Scientist at Amazon.com LLC, Seattle, WA, USA. His research interests include machine learning, deep learning, and data mining.



Rob Kulathinal was born in Toronto, ON, Canada. He received the B.Sc. degree (Hons.) in biology from McMaster University, Hamilton, ON, Canada, and the Ph.D. degree in evolutionary genetics from McMaster University in 2001.

He was trained as a Post-Doctoral Fellow at Harvard University, Cambridge, MA, USA, and taught a year-long course on the genomics of speciation at Harvard College. He is currently an Associate Professor at the Center for Computational Genetics and Genomics, Department of Biology, Temple University, Philadelphia, PA, USA, where he teaches several courses at both the undergraduate and graduate levels, including evolutionary genomics and topics in bioinformatics. He is also the Co-Director of the Ph.D.-Bioinformatics Program and a Steering Committee Member of the Professional Science Master's (PSM) Bioinformatics Program at Temple University. His current research focuses on examining the genetic mechanisms that underlie rapidly evolving systems and drive functional divergence with his lab employing a combination of experimental, theoretical, and computational approaches to investigate rapidly evolving genes, traits, and genetic networks and the role they play in phenotypic and species divergence.



Sudhir Kumar received the B.Eng. degree (Hons.) in electrical and electronics engineering and the M.Sc. degree (Hons.) in biological sciences from the Birla Institute of Technology and Sciences, Pilani, India, in 1990, and the Ph.D. degree in genetics from Pennsylvania State University, University Park, PA, USA, in 1996.

In 1998, he joined the Department of Biology, Arizona State University (ASU), Tempe, AZ, USA, as an Assistant Professor, where he became a Full Professor in 2012. He was appointed as a Regents' Professor at the School of Life Sciences, ASU, in 2014. In 2014, he became a Laura H. Carnell Professor at the Department of Biology, Temple University, Philadelphia, PA, USA, where he is also the Founding Director of the Institute for Genomics and Evolutionary Medicine. He is also a Full Professor at the Department of Computer and Information Sciences, Temple University. His group has made many pioneering contributions by developing new computational methods and algorithms for scalable and efficient analysis of big data in molecular evolution, functional genomics, and biomedicine. His work has been cited more than 200 000 times, with one of his scientific articles included in the Thomson Reuters Web of Science top-100 most-cited papers of all time. He has published numerous citation classics and hot papers.

Dr. Kumar is a fellow of the American Association for the Advancement of Science. He received the Innovation Award in Functional Genomics from the Burroughs-Wellcome Fund in 2000.



Shuiwang Ji (Senior Member, IEEE) received the Ph.D. degree in computer science from Arizona State University, Tempe, AZ, USA, in 2010.

He is currently a Professor at the Department of Computer Science and Engineering, Texas A&M University, College Station, TX, USA. His research interests include machine learning, deep learning, data mining, and computational biology.

Dr. Ji received the National Science Foundation CAREER Award in 2014. He is also an Associate Editor of *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, *ACM Transactions on Knowledge Discovery from Data*, and *ACM Computing Surveys*. He regularly serves as the Area Chair or equivalent roles for data mining and machine learning conferences, including the Association for the Advancement of Artificial Intelligence (AAAI), the International Conference on Learning Representations (ICLR), the International Conference on Machine Learning (ICML), the International Joint Conference on Artificial Intelligence (IJCAI), the Knowledge Discovery and Data Mining (KDD), and the Neural Information Processing Systems (NeurIPS).