# SelectionConv: Convolutional Neural Networks for Non-rectilinear Image Data

David Hart, Michael Whitney, and Bryan Morse

Brigham Young University, Provo, Utah, USA {davidmhart,mikeswhitney,morse}@byu.edu

Abstract. Convolutional Neural Networks have revolutionized vision applications. There are image domains and representations, however, that cannot be handled by standard CNNs (e.g., spherical images, superpixels). Such data are usually processed using networks and algorithms specialized for each type. In this work, we show that it may not always be necessary to use specialized neural networks to operate on such spaces. Instead, we introduce a new structured graph convolution operator that can copy 2D convolution weights, transferring the capabilities of already trained traditional CNNs to our new graph network. This network can then operate on any data that can be represented as a positional graph. By converting non-rectilinear data to a graph, we can apply these convolutions on these irregular image domains without requiring training on large domain-specific datasets. Results of transferring pre-trained image networks for segmentation, stylization, and depth prediction are demonstrated for a variety of such data forms.

Keywords: Graph Convolution, Transfer Learning, Irregular Images, Superpixels, Spherical Images, Texture Maps

# 1 Introduction

Convolution has been an important operator in image processing practically since its inception, long before the age of deep learning. It is the backbone of most modern deep neural networks, and learnable weights in convolution layers lead to incredible capabilities such as classification, object detection, segmentation, stylization, and many others.

Convolution is powerful, but the discrete form used for raster images requires dense rectilinear grids, typically Cartesian grids. For sparse, discontinuous, or irregular data, discrete raster convolution may not be applicable. Methods such as rasterization, interpolation, or padding are commonly used to convert the data into a form suitable for discrete convolution.

Graph convolution is more adaptable to less regularly structured data and is designed to mimic the process of 2D convolution. Instead of requiring spatial adjacency, it performs convolution based on an adjacency matrix that describes the edges that connect nodes to each other in the graph. One key difference, however, between traditional and graph convolution is that graph convolution is



Fig. 1. Our method allows pre-trained 2D CNNs to operate on non-rectilinear image domains such as superpixels, spherical images, masked images, and texture maps.

assumed to be non-orientable, meaning that it cannot treat incoming edges differently based on spatial direction or the order they are fed into the convolution. This is called the *permutation-invariance* constraint of graph convolution [4].

In image convolution, neighboring pixels are commonly given different weights to help detect shapes and other patterns. In graph convolution, all neighbors are aggregated in the same way, removing any location-based structure in the process. Thus, the weights learned in a 2D convolutional neural network are incommensurate with the weights learned in a graph convolution neural network.

Graphs are commonly used to model abstract data that do not have positional information (social networks, individual media ratings, etc.), but when the graph data is image-based, we still wish to leverage positions, shapes, and patterns in the same manner as traditional image convolution.

This paper presents a framework for working with non-rectilinear imagebased data that both traditional convolutional networks and graph networks are ill-equipped to handle natively, including the types shown in Fig. 1. We do this using a new type of selection-based graph convolution, which we name *SelectionConv*, that can retain the same shapes and patterns that a traditional convolution learns. In so doing, traditional convolution weights can be made commensurate with SelectionConv weights, allowing the transfer of weights directly from networks previously trained on standard image datasets. Thus, no special training or fine-tuning is necessary to run the network on less conventional image types. This is particularly significant because less common image types usually have far less available training data than typical image datasets.

Through this method, any network that operates on images can operate on any form of data that can be represented as a positional graph. This allows one framework to perform multiple tasks, such as depth prediction on superpixels, segmentation of spherical images, and stylization of texture maps for 3D meshes, as described in Sec. 5 and demonstrated in Sec. 6. This technique opens up a realm of possibilities for previously underused data sources.

In summary, our contributions are as follows:

- We present a selection-based graph convolution operator that assigns different weights to incoming edges without violating permutation invariance.
- We demonstrate how to transfer pre-trained 2D convolution weights to our new graph operator, thus removing the need to retrain the graph network.
- We apply this new method to various non-rectilinear image applications to demonstrate its effectiveness.

## 2 Related Work

#### 2.1 Graph Convolution Networks

The explosion of deep learning in recent years has led to many developments in both Convolutional Neural Networks (CNNs) and Graph Neural Networks (GNNs). Graph Convolution Networks (GCNs) started with the work of Kipf *et al.*, which extended the ideas of CNNs to a general graph structure [18]. Improvements on the original method have been proposed including higher-order aggregation structures [6, 27] and incorporating MLPs in the aggregation step [39]. For a more complete overview of this line of work, we direct the interested reader to recent surveys of both CNNs [21] and GNNs [41].

Interpolated Convolution Networks [25] and Spline Convolution Networks [10] are designed, like this work, to mimic the process of a traditional CNN on a point cloud or graph. However, these two approaches do not provide an explicit method for transferring weights to the new network. Additionally, both of these approaches require traditional position-based point clouds, whereas we show that our method is adaptable to many different data types and can be modified flexibly according to the users' specifications.

The works by Xu *et al.* [42] and Zhou *et al.* [48] both aim to use the structural and positional information inherent in graphs to improve graph learning. These are learnable components that can improve the training and performance of graph networks given sufficient suitable training data. Our approach aims to explicitly define the graph structure so that additional training is not needed.

#### 2.2 Transfer Learning

The goal of transfer learning is to take the information or weights learned from one network and utilize them in another network in some way. One common example of transfer learning is to use a CNN backbone trained for a classification network such as VGG-19 [32] for another task such as segmentation. Many researchers have explored the effectiveness of networks trained for one task when performing another task [24, 45], and a recent survey of transfer learning techniques can be found in [50]. Our work differs from the goal of previous transfer learning literature since our focus is not to transfer weights to a different task, but to a network that operates on a different domain.

It is worth noting that there has been an effort to theoretically unify the various types of neural networks and their various domains by focusing on their invariance and equivariance properties [4]. To our knowledge, though, no attempt has been made to state these operations in terms of each other and thus make them transferable.

#### 2.3 Spherical Images, Superpixels, and Texture Maps

This paper demonstrates the effectiveness of selection-based graph convolution on various forms of non-raster data. Here we include work relevant to the tasks we perform and the types of data used.

#### 4 Hart et al.

We demonstrate working with spherical images by performing both semantic segmentation and stylization. Several groups have worked on performing semantic segmentation on spherical images [17, 36, 46]. Notable is the work of Tateno *et al.* [36] who developed distortion aware convolutions that operate in spherical space and also have the ability to transfer weights from a standard 2D CNN. However, their method is specific to spherical images where ours extends to other image domains.

Ruder *et al.* [30] present a method for performing style transfer on  $360^{\circ}$  images by taking the six cube-projected views and stylizing each one in turn while enforcing consistency between each previously stylized view. We also perform spherical image stylization in Sec. 6.1, but we do so in a single feed-forward step without the need of fine-tuning a specialized network.

The aim of superpixels is to group similar pixels in an image into regions, simplifying the representation of the whole image. In this work, we use SLIC, a standard baseline algorithm for generating superpixels [1]. Many other classical approaches exist for generating superpixels [34]. Recent deep learning techniques have been proposed to improve superpixels [22, 38]. Superpixels have also been used to improve modern detection and segmentation methods [16, 47]. Yang *et al.* [43] use superpixels to take low-resolution results into higher resolutions, similar to the task we perform in Sec. 6.3, but they do so by pretraining a separate network that can predict superpixel associations on a grid-based structure.

Most work on meshes has focused on learning from the geometry rather than from the color information that is provided in the texture map. Some have explored generating new texture from a smaller example texture through classical texture synthesis methods [3, 8, 40] and neural approaches [12, 13, 31, 33, 49]. Textures have also been manipulated through lighting-based style transfer [11, 35] rather than through a purely image-based approach like the one we present in Sec. 6.4. Yin *et al.* [44] recently proposed a geometry and texture stylization approach that is optimization based and uses differentiable rendering, a fundamentally different approach to operating on texture map data than the one we explore in this work.

# 3 Selection-Based Convolution

Our method requires designing a graph convolution operator that treats incoming edges differently from one another during the aggregation step. Traditionally in graph convolution, all connecting edges are specified in a single adjacency matrix, and this matrix is used to described which nodes can influence each other after some set of transformation operations. For example, the original Graph Convolution Layer defined in [18] can be described as

$$\mathbf{X}^{(k+1)} = \tilde{\mathbf{A}} \mathbf{X}^{(k)} \mathbf{W}$$
(1)

where  $\mathbf{X}^{(k)}$  is the current node activations,  $\tilde{\mathbf{A}}$  is a normalized adjacency matrix, and  $\mathbf{W}$  is the learned weights. Note that the weight matrix is applied equally to all nodes, making the result invariant to the order that nodes are enumerated in  $\mathbf{X}^{(k)}$  as long as  $\tilde{\mathbf{A}}$  changes correspondingly. This is an example of the permutation invariance constraint to which all graph convolution operations must adhere.

In comparison, while standard 2D convolution is shift invariant, it is not permutation invariant, relying heavily on orientation when assigning the weight of each connecting pixel. For example, the pixel directly above the current one might be given different weight than the pixel to the bottom right, and so on. Graph convolutions are able to assign edge weights, but they are generally static or on a node-by-node basis using a mechanism such as attention [37]. Thus, the weights learned during a 2D convolution are incommensurate with the weights learned during a graph convolution.

In order to leverage the benefits of pretrained 2D convolutional networks while having the structural flexibility of graphs, we introduce a new graph convolution layer that can preserve location information. We do so by preprocessing the graph into multiple adjacency matrices, selecting edges to be assigned to different matrices based on the spatial relationship between their two nodes. This is similar to the way we can think of different adjacency relationships between pixels and their directional neighbors. There is also a unique weight matrix for each corresponding adjacency matrix so only those edges are affected. The results for the set of adjacency and weight matrices are summed together to make the final activation. This selection-based convolution is what we call *SelectionConv*.

For each graph, a given edge  $e_{ij}$  needs to be assigned to its specific adjacency matrix. We do so using a selection function  $s(v_i, v_j)$  for vertices  $v_i$  and  $v_j$  with spatial positions  $\mathbf{x}_i$  and  $\mathbf{x}_j$  respectively, which indicates which adjacency matrix includes the edge  $e_{ij}$  between these vertices. For m possible selections this gives m adjacency matrices respectively defined as

$$\mathbf{S}_{m_{ij}} = \begin{cases} 1 & \text{if } s(v_i, v_j) = m \\ 0 & \text{otherwise} \end{cases}$$
(2)

Each selection has a corresponding weight matrix. Thus our convolution becomes

$$\mathbf{X}^{(k+1)} = \sum_{m} \tilde{\mathbf{S}}_{m} \mathbf{X}^{(k)} \mathbf{W}_{m}$$
(3)

where  $\tilde{\mathbf{S}}_m$  is the normalized version of  $\mathbf{S}_m$  to account for nodes having multiple edges with the same selection m. With this structure, nodes can be treated differently based on location and other features relative to the current node without breaking the permutation invariance constraint. An example of this process is illustrated in Fig. 2.

We use PyTorch Geometric [9] to implement this process and use slices from 3D tensors rather than separate weight and adjacency matrices for efficiency, but the result is mathematically equivalent.

# 4 Selection-based Convolution for Images

Though the process described in Sec. 3 is general enough to work with any number of selections based on any number of node attributes, we now move to



Fig. 2. A graph with a selection function that weights upwards edges differently from downwards edges. Such a selection function would give two adjacency matrices,  $S_1$  and  $S_2$ , that would be applied to two different weight matrices,  $W_1$  and  $W_2$ .

the specific case of working in the image domain. To start, we will establish a baseline for our method by looking at a regular image and showing that we get results identical to a 2D convolution with our selection-based convolution.

#### 4.1 Setting Up Image Graphs

An image with pixels in a Cartesian grid can be thought of as a set of nodes at equally spaced distances. Many neural networks use  $3 \times 3$  convolutions as their primary image feature extractor, which look at the pixel and its 8-connected neighbors. Thus, when we construct our graph, we will add an edge from each pixel to its 8-connected neighbors and one to itself. This means we need m = 9 possible selections.

For pixels in a Cartesian grid arrangement, value assignment for the respective selection functions is straightforward. For the general spatial case, we project the vector defined by the position of the two nodes onto the set  $\mathbf{D}$  of unit vectors in each of the respective neighbor directions. Specifically,

$$\mathbf{D} := \frac{\left\langle -\sqrt{2}/2, -\sqrt{2}/2 \right\rangle \left\langle 0, -1 \right\rangle \left\langle \sqrt{2}/2, -\sqrt{2}/2 \right\rangle}{\left\langle -1, 0 \right\rangle \qquad \left\langle 1, 0 \right\rangle} \tag{4}$$

Whichever directional unit vector results in the largest projection (resulting dot product) corresponds to the assigned selection. Additionally, if the positions are the same (or within some small radius), the central selection is made. Thus, our selection function becomes

$$s(v_i, v_j) = \begin{cases} 0 \text{ if } \|\mathbf{x}_j - \mathbf{x}_i\| < \epsilon \\ \underset{k}{\operatorname{argmax}} \mathbf{D}_k \cdot (\mathbf{x}_j - \mathbf{x}_i) \text{ otherwise} \end{cases}$$
(5)

For simplicity, when assigning a selection number or index to each direction, we follow the mathematical convention of angles by making the direction to the right be the first selection and moving in the counterclockwise direction for assigning each subsequent direction. This is visualized on the left and right sides of Fig. 3.

#### 4.2 Weight Transfer from 2D Convolutions

Once the appropriate selections have been made, transferring the weights is simply copying the appropriate slice of the convolution kernel weights to its



Fig. 3. Elements of a  $3 \times 3$  convolution kernel are enumerated with zero for the center weight and neighboring weights from one to eight in counter-clockwise direction. The different weights from the kernel are transferred to associated weight matrices. Those weight matrices are then applied to the selected edges on the graph. Note that the points in the graph do not need to be equally spaced as in regular images. There can also be more than one node per selection.

assigned selection. For example, if selection 5 represents an edge going to the left, the left kernel convolution weights would be copied to  $\mathbf{W}_5$ . This process is illustrated in Fig. 3. When applied to raster data, this process leads to results that are identical to those using an image-based convolutional network.

#### 4.3 Handling Larger Kernels

If a network uses a kernel that is larger than  $3 \times 3$ , more weights need to be copied over, but the same graph structure can still be used. Rather than using a larger graph that is memory inefficient, we use the edges of the simpler  $3 \times 3$  graph. Through successive multiplications of the adjacency matrices, multiple edge hops (traversals) can be performed until specific kernel locations are reached. For example, if a  $5 \times 5$  kernel is used, the weight associated with the bottom middle pixel would be assigned to the action of taking the bottom selection's bottom selection.

## 4.4 Pooling Operators and Upsampling

Many CNNs use pooling layers throughout the network to combine nearby features. While GCNs have similar pooling features, it is important that our SelectionConv network's pooling layers mimic the downsampling nature of those used in traditional CNNs.

If the network contains fully connected layers, we impose a spatial grid onto the set of nodes when pooling. This grid matches the spatial size that the original image data reduces to after pooling. If the network is fully convolutional, pooling does not require a regular grid (especially if one cannot be imposed, such as on a texture map), and any pixel-clustering algorithm can be used. In both cases, any node within a cell or cluster is made into a single node during the pooling step. The average of the positions of nodes in that cell becomes the position of the pooled node. This is similar to the meta-node approach used in [29].

Pooled nodes also need to reestablish edge connections and selections. To do so, we implement a post-pooling function that makes new edges between the aggregated cluster nodes by using the previous layer's graph edges. If any edge exists between two nodes in different clusters, the aggregated nodes of those clusters will have a corresponding edge between them. That edge is assigned the average of the previous selection values of all the edges between the original nodes in the two clusters (while properly accounting for the selection value wraparound between 1 and 8).

If the network requires upsampling steps, we save each version of the graph before it is downsampled. When upsampling, we revert back to a previous version of the graph and copy appropriate values to each node using the defined clusters. Bilinear upsampling can also be approximated by using the average value of the connections to new nodes or through other point cloud interpolation methods.

## 4.5 Strides, Dilation, Padding

Traditional 2D convolution layers often have additional parameters such as the stride of the kernel, the dilation of the kernel, and the padding to be used on the border of the image. Selection-based convolution as described so far is equivalent to a stride of 1, a dilation of 1, and zero padding, but we have the ability to mimic these additional features in our SelectionConv network when needed.

Strides larger than 1 in a regular convolution layer are equivalent to a stride equal to 1 followed by downsampling by the stride amount (with no antialiasing). We use this same idea to implement large strides in the SelectionConv network. The convolution is performed as usual (equivalent to a stride of 1), then the graph is pooled using the method described in Sec. 4.4, but rather than using a max or average pooling operator, a predetermined central node in each cluster is always used as the pooled value.

Dilation is handled in a way similar to the larger kernels described in Sec. 4.3. The dilation amount defines how many times a selection is edge hopped. For example, a dilation of two would indicate that instead of taking the left selection, the left selection's left selection should be used instead. This process is repeated for each selection for the same number of times as the dilation's value.

Padding in traditional 2D CNNs helps control the size of output layers. Graph convolution layers do not change the size of the output since the number of nodes will stay the same, so padding is not usually needed. Some 2D CNN padding schemes, however, are used to help information propagate correctly along the borders of an image (such as reflective padding in stylization networks). We handle these situations not only by looking at nodes along a border, but by determining what to do with any missing selections. The following padding methods can be implemented effectively and approximate padding schemes used for images:

- **Zero**: Missing selections are not considered (default).
- **Constant**: Missing selections are assigned a predetermined value.
- **Replicate**: Missing selections are assigned the value of the current node.
- Reflect: Missing selections are assigned the value of the selection in the opposite direction.

Some of these steps we mimic from CNNs currently have nondifferentiable implementations. Though not needed for the scope of this work, which focuses on transferring weights from pre-trained networks, all parts of our SelectionConv network would need to be differentiable if any form of training or fine-tuning was desired for the network after weights have been transferred.

## 4.6 Verification

To verify that the SelectionConv network can truly be equivalent to a traditional CNN, we used a pre-trained VGG-11 network [28, 32] on CIFAR-10 [19] and transferred the original weights to our network. We compared this to the original image-based network using the 10,000-image validation set. As expected, the two methods resulted in identical predictions and identical accuracies of 84.5%. This remains true even when small random spatial perturbations are introduced to the points in the graph structure.

## 5 Example Non-rectilinear Configurations

Sec. 4 describes how to configure a SelectionConv network to work on images in exactly the same way as a regular CNN. This works as a baseline but does not provide any additional power over a regular CNN. In this section, we give examples of the flexibility of our method to work with data that cannot be processed with a traditional CNN due to its irregular structure. *Importantly,* our method can use the weights from a CNN pre-trained using standard image datasets without the need to retrain for specific data types.

#### 5.1 Panoramic and Spherical Images

Many smartphones and cameras allow users to take single or multiple pictures of their surroundings to acquire panoramic or even up to the full  $360^{\circ} \times 180^{\circ}$  view of their environment. These panoramic and spherical images have non-simple topologies but are typically stored as simple planar images. This requires projection of the content to a surface of some kind. While there are many ways to project, including spherical, equirectangular (cylindrical), and cubic, each of these have distortion or seams of some nature that would be difficult for a traditional CNN to handle due to irregular spatial sampling or topological considerations. With our method, such seams and distortion can be handled by proper construction of the graph.

As an example, we show how to construct the graph for a cubic projection often used for environment maps in computer graphics, which we have found to be effective to work with since it has low levels of distortion compared to other projections. Seams are naturally present along each edge of the cube, and a 2D image can only represent a few of those connections. In our graph, we simply need to make the connections between the rest of the seams, as illustrated in Fig. 4.a. Additionally, we orient our selection function in such a way so that the upwards selection is always pointed towards the top pole of the map.

10 Hart et al.



**Fig. 4.** a) Illustration of the graph connections for a cube map. Red arrows represent the upwards selection for nodes in that part of the graph. Green and blue lines represent connections made in the graph between faces. b) Illustration of the graph connections between the centroids of superpixels in an image.

## 5.2 Superpixel Images

Using superpixels is a common approach for simplifying a high-resolution image by representing it as a smaller set of similar regions. Because of the irregular structure of the regions, superpixels cannot be used for standard CNNs. They can, however, be easily represented as a graph and used with our approach.

To construct the graph from the set of superpixels, the centroid of each region is treated as a node, and edges are selected using a K-nearest neighbors method. Selections are then made using the dot product method described in Sec. 4.1 and the graph is pruned so only the closest neighbor to each node for a given selection is used. This process is illustrated in Fig. 4.b.

## 5.3 Masked Images and Texture Maps

Many applications require operating only on the foreground or a specified region of the image. Our graph construction can naturally handle these cases by simply dropping nodes and edges that are not part of the desired region, and any desired padding mode described in Sec. 4.5 can be used to handle the irregular border.

Texture maps for 3D meshes can also be thought of as masked regions. Not all pixels present on a texture map image will be used to determine colors on the actual mesh, so we can mask out the regions that contain pixels that represent some part of a face on the mesh. If we further connect these faces in the graph, we can operate on the texture map in the same fashion as any image. To do this in general, we start by determining which edges are only referenced once in the UV map, since these represent the boundaries of groups of faces on the UV map. We then find all closed loops of edges to separate each boundary. Finally, we do a polygon-contains-point test for each pixel to see if the pixel is inside any of our boundaries. This becomes the mask of relevant pixels on the texture map. An overview of this process is shown in Fig. 5.



**Fig. 5.** A 3D model (a) and its texture map. The model's texture seams can be determined from UV coordinates and represent boundaries on the texture map (b). From this, we construct a mask of relevant pixels (c) and connect discontiguous regions.

If it is known which geometric faces are connected to each other, this can be built into the graph construction. Otherwise, each edge boundary is paired with a corresponding edge on another face that is used to make the connections. All regions are connected regardless of where they are located on the image.

# 6 Results

To demonstrate transfer from image networks to SelectionConv networks, we present the results of applying this and the graph-construction methods from the previous section to various applications and image types. Additional results can be found in the accompanying supplemental materials.

## 6.1 Spherical Style Transfer

A simple but effective illustration of the seamless nature of a selection-based graph convolution is to perform style transfer on a spherical image. For this we use the feed-forward style transfer approach recently proposed by Li *et al.* [20]. Usually for this task, a special piece-wise optimization approach or fine-tuned network would be needed, such as that proposed by Ruder *et al.* [30]. However, by generating our spherical graph using the method shown in Sec. 5.1, we naturally avoid distortion, minimize seams, and can stylize in a single feed-forward pass. An example is shown in Fig. 6. The whole process of generating the graph, transferring the weights, and running the graph convolution can run in 15-20 seconds on a consumer GPU, while [30] would take 8-10 minutes per 360° image. Even the faster approach suggested in [30] requires fine-tuning a network for 120,000 iterations per style image. Our approach uses a state-of-the-art feed forward style transfer method, can be used for any style image, and still enforce consistency across the seams of the cube map.

## 6.2 Spherical Segmentation

We apply SelectionConv to the task of semantic segmentation on images in the Stanford 2D-3D-S [2] dataset. To do so, we first trained a standard 2D FCN [23] using a ResNet-50 [15] backbone on 2D projected views. We then

12 Hart et al.



**Fig. 6.** A  $360^{\circ}$  image (a) and its stylization using our feed-forward method (b). An example view looking downward at the lower pole of the image (c) has seams and distortion when naively stylizing the rectangular image (d), but those seams and distortion are minimized with our method (e). Image taken from [5].

transfer the weights to a SelectionConv-based version of FCN [23] and apply the network, with no additional training, to the standard test set, converting the data to spherical graphs using the method described in Sec. 5.1. This gives an improvement over naively applying the network to the equirectangular images, as shown in Fig. 7. When operating on the validation set, the naive approach gives an average IOU score of 32.57%. We compare our results with that of another transfer-based method, distortion-aware convolutions [36], who reported an average IOU score of 34.56% on the same dataset. Our results show an average IOU score of 36.29%. Although this is only a small improvement and is still below state-of-the-art performance for RGB spherical segmentation (45.6% [7]), we again note that other methods are designed and fine-tuned specifically for



**Fig. 7.** A visual comparison of semantic segmentation of images from the Stanford 2D-3D-S [2] dataset (a,b) between an FCN [23] with a ResNet-50 [15] backbone using standard convolutions (c) versus our SelectionConv operations (d). Note that the use of SelectionConv gives cleaner segmentation results along the poles of and seam of the image (located in the center of this representation).



Fig. 8. A high-resolution image (a) requires  $\sim 25.9$  seconds on a CPU to create a predicted depth map (b). A lower-resolution  $256 \times 256$  version can be processed by a network in  $\sim 0.8$  seconds on a GPU, but with low-fidelity results when upscaled to the same resolution (c). Generating approximately the same number of superpixels as the low-resolution image then using our graph-based network requires only  $\sim 5.1$  seconds on a GPU with higher-fidelity results (d).

spherical tasks, whereas we can achieve a performance boost through a simple design of our graph structure.

#### 6.3 Superpixel Depth Prediction

We now illustrate possible applications using our method with superpixels for efficient processing of high-resolution images. We use a Pytorch implementation of a monocular depth estimator [14, 26]. When operating on a 4K image, the amount of data is too large for a consumer-grade GPU, necessitating downsampling the input before processing. By comparison, if the image is first converted into a graph of neighboring SLIC superpixels [1], SelectionConv can process it on a GPU and output results that are of much higher quality than using a downsampled image. An example of such cases is shown in Fig. 8. We again note that Yang *et al.* [43] complete a similar task with state-of-the-art performance, but their method requires using superpixels generated by their trained network. The SelectionConv network, in comparison, can utilize any superpixel method.

## 6.4 Masked Image and 3D Mesh Style Transfer

Lastly, we demonstrate the ability of our network to work on data with many discontinuities by performing style transfer on masked images and texture maps.

To achieve style transfer on a masked region with a regular CNN would require stylizing the entire image or a zero-padded masked image and then reapplying the unmasked region. This means that background features can affect stylization of the foreground. In comparison, our method can handle these scenarios natively, leading to a stylization that only depends on the foreground statistics. Comparisons of these approaches are shown in Fig. 9.

As another illustration, treating a texture map as a 2D image and naively performing style transfer leads to noticeable seams in the mapped texture. In comparison, using the graph structure proposed in Sec. 5.3 leads to more continuous patterns. Visualizations of these two methods for various stylizations and meshes are shown in Fig.10.



**Fig. 9.** A content image, a masked region of interest, and a given style image (a). To stylize the masked region with a traditional CNN, the entire image can be stylized (b) or the image can be masked before stylization (c) and then the masked result can be applied back to the original. In both cases, outside statistics influence the stylization inside the region of interest (making (b) darker than expected and (c) brighter than expected). In comparison, our method (d) can generate a graph just for the masked region, which more closely matches the style image statistics in the region of interest.



(a) 3D Mesh (b) Naive (c) Close-up (d) SelectionConv (e) Close-up

**Fig. 10.** 3D mesh (a), the result of naively stylizing the texture map (b) and a magnification (c), and the result of using our method (d) and a magnification (e). Note the visible seams shown in the magnifications of the naive method (c), whereas our method in (e) minimizes the visibility of those seams.

Others have attempted style transfer between two different 3D objects [11, 35, 44], but we are not aware of other work attempting direct style transfer between the texture map of the 3D mesh and a 2D image.

# 7 Conclusion

We have presented a method that allows for information from pre-trained traditional convolutional neural networks to be transferred directly to a new kind of graph convolutional network. This makes it possible for these previously trained networks to operate on data that they could not before, such as superpixels, spherical images, and texture maps. We have demonstrated various use cases and given the general framework so that others can continue to extend this method for their needs. In theory, any set of adjacency matrices could be designed to work with the particular data of a graph. Future research could also use selection-based convolution to improve applications outside of the image domain.

# References

- Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: SLIC superpixels compared to state-of-the-art superpixel methods. IEEE Transactions on Pattern Analysis and Machine Intelligence 34(11), 2274–2282 (2012). https://doi.org/10.1109/TPAMI.2012.120
- Armeni, I., Sax, A., Zamir, A.R., Savarese, S.: Joint 2D-3D-Semantic Data for Indoor Scene Understanding. arXiv:1702.01105 (2017)
- Ashikhmin, M.: Synthesizing natural textures. In: Symposium on Interactive 3D graphics. pp. 217–226 (2001)
- Bronstein, M.M., Bruna, J., Cohen, T., Velickovic, P.: Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. arXiv:2104.13478 (2021)
- Chou, S.H., Sun, C., Wen-Yen, C., Hsu, W.T., Sun, M., Fu, J.: 360-Indoor: Towards learning real-world objects in 360° indoor equirectangular images. pp. 834–842 (2020). https://doi.org/10.1109/WACV45572.2020.9093262
- Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional neural networks on graphs with fast localized spectral filtering. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 29. Curran Associates, Inc. (2016)
- Eder, M., Shvets, M., Lim, J., Frahm, J.M.: Tangent images for mitigating spherical distortion. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12426–12434. IEEE (2020)
- Efros, A.A., Freeman, W.T.: Image quilting for texture synthesis and transfer. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 341–346 (2001)
- Fey, M., Lenssen, J.E.: Fast graph representation learning with PyTorch Geometric. In: ICLR Workshop on Representation Learning on Graphs and Manifolds (2019)
- Fey, M., Lenssen, J.E., Weichert, F., Müller, H.: SplineCNN: Fast geometric deep learning with continuous B-spline kernels. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 869–877. IEEE (2018)
- Fišer, J., Jamriška, O., Lukáč, M., Shechtman, E., Asente, P., Lu, J., Sýkora, D.: StyLit: Illumination-guided example-based stylization of 3D renderings. ACM Trans. Graph. 35(4), 92:1–92:11 (2016). https://doi.org/10.1145/2897824.2925948
- Frühstück, A., Alhashim, I., Wonka, P.: Tilegan: synthesis of large-scale nonhomogeneous textures. ACM Transactions on Graphics (TOG) 38(4), 1–11 (2019)
- Gatys, L., Ecker, A.S., Bethge, M.: Texture synthesis using convolutional neural networks. Advances in neural information processing systems 28, 262–270 (2015)
- Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. Conference on Computer Vision and Pattern Recognition (CVPR) pp. 770–778 (2016)
- Jampani, V., Sun, D., Liu, M.Y., Yang, M.H., Kautz, J.: Superpixel sampling networks. In: European Conference on Computer Vision (ECCV) (2018)
- Jiang, C.M., Huang, J., Kashinath, K., Prabhat, Marcus, P., Niessner, M.: Spherical CNNs on unstructured grids. In: International Conference on Learning Representations (ICLR) (2019)
- Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings (2017)

- 16 Hart et al.
- 19. Krizhevsky, A., Nair, V., Hinton, G.: Cifar-10 (canadian institute for advanced research) http://www.cs.toronto.edu/ kriz/cifar.html
- Li, X., Liu, S., Kautz, J., Yang, M.H.: Learning linear transformations for fast arbitrary style transfer. In: Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2019)
- 21. Li, Z., Yang, W., Peng, S., Liu, F.: A survey of convolutional neural networks: Analysis, applications, and prospects (2020)
- Lin, Q., Zhong, W., Lu, J.: Deep superpixel cut for unsupervised image segmentation. 2020 25th International Conference on Pattern Recognition (ICPR) pp. 8870–8876 (2021)
- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2015)
- Lu, Y., Pirk, S., Dlabal, J., Brohan, A., Pasad, A., Chen, Z., Casser, V., Angelova, A., Gordon, A.: Taskology: Utilizing task relations at scale. In: Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2021)
- Mao, J., Wang, X., Li, H.: Interpolated convolutional networks for 3D point cloud understanding. In: International Conference on Computer Vision (ICCV). pp. 1578–1587. IEEE (2019)
- 26. Monodepth. https://github.com/OniroAI/MonoDepth-PyTorch
- Morris, C., Ritzert, M., Fey, M., Hamilton, W.L., Lenssen, J.E., Rattan, G., Grohe, M.: Weisfeiler and leman go neural: Higher-Order graph neural networks. AAAI 33(01), 4602–4609 (2019)
- 28. PyTorch, Torchvision Models. https://pytorch.org/vision/stable/models.html
- Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. p. 5105–5114. Curran Associates Inc., Red Hook, NY, USA (2017)
- Ruder, M., Dosovitskiy, A., Brox, T.: Artistic style transfer for videos and spherical images. International Journal of Computer Vision 126(11), 1199–1219 (2018)
- Shi, W., Qiao, Y.: Fast texture synthesis via pseudo optimizer. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5498–5507. IEEE (2020)
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (ICLR) (2015)
- Snelgrove, X.: High-resolution multi-scale neural texture synthesis. In: SIGGRAPH Asia 2017 Technical Briefs, pp. 1–4 (2017)
- 34. Stutz, D., Hermans, A., Leibe, B.: Superpixels: An evaluation of the stateof-the-art. Computer Vision and Image Understanding 166, 1–27 (2018). https://doi.org/https://doi.org/10.1016/j.cviu.2017.03.007
- Sýkora, D., Jamriška, O., Texler, O., Fišer, J., Lukáč, M., Lu, J., Shechtman, E.: StyleBlit: Fast example-based stylization with local guidance. Computer Graphics Forum 38(2), 83–91 (2019)
- Tateno, K., Navab, N., Tombari, F.: Distortion-aware convolutional filters for dense prediction in panoramic images. In: European Conference on Computer Vision (ECCV) (2018)
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph Attention Networks. International Conference on Learning Representations (ICLR) (2018)
- Verelst, T., Blaschko, M.B., Berman, M.: Generating superpixels using deep image representations. arXiv:1903.04586 (2019)

17

- Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph CNN for learning on point clouds. ACM Trans. Graph. 38(5), 1–12 (2019)
- 40. Wei, L.Y., Levoy, M.: Fast texture synthesis using tree-structured vector quantization. In: Proceedings of the 27th annual conference on Computer graphics and interactive techniques. pp. 479–488 (2000)
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., Yu, P.S.: A comprehensive survey on graph neural networks. IEEE Transactions on Neural Networks and Learning Systems p. 1–21 (2020). https://doi.org/10.1109/tnnls.2020.2978386
- 42. Xu, M., Ding, R., Zhao, H., Qi, X.: Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds. In: Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2021)
- Yang, F., Sun, Q., Jin, H., Zhou, Z.: Superpixel segmentation with fully convolutional networks. Conference on Computer Vision and Pattern Recognition (CVPR) pp. 13961–13970 (2020)
- 44. Yin, K., Gao, J., Shugrina, M., Khamis, S., Fidler, S.: 3DStyleNet: Creating 3D shapes with geometric and texture style variations. In: International Conference on Computer Vision (ICCV). IEEE (2021)
- 45. Zamir, A.R., Sax, A., Shen, W.B., Guibas, L., Malik, J., Savarese, S.: Taskonomy: Disentangling task transfer learning. In: Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2018)
- Zhang, C., Liwicki, S., Smith, W., Cipolla, R.: Orientation-aware semantic segmentation on icosahedron spheres. In: International Conference on Computer Vision (ICCV). IEEE (2019)
- Zhao, G., Ge, W., Yu, Y.: GraphFPN: Graph feature pyramid network for object detection. In: International Conference on Computer Vision (ICCV). pp. 2763– 2772. IEEE (2021)
- Zhou, H., Feng, Y., Fang, M., Wei, M., Qin, J., Lu, T.: Adaptive graph convolution for point cloud analysis. International Conference on Computer Vision (ICCV) pp. 4945–4954 (2021)
- Zhou, Y., Zhu, Z., Bai, X., Lischinski, D., Cohen-Or, D., Huang, H.: Non-stationary texture synthesis by adversarial expansion. ACM Trans. Graph. 37(4) (2018). https://doi.org/10.1145/3197517.3201285
- 50. Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., He, Q.: A comprehensive survey on transfer learning. Proceedings of the IEEE 109, 43–76 (2021)