

Learning Inclusion Matching for Animation Paint Bucket Colorization

Yuekun Dai Shangchen Zhou Qinyue Li Chongyi Li Chen Change Loy
S-Lab, Nanyang Technological University

<https://ykdai.github.io/projects/InclusionMatching>




Figure 1. In the animation industry, digital painters use paint bucket tool to colorize drawn line arts frame by frame. Our proposed pipeline streamlines this process by requiring the painters to colorize just one frame, after which the algorithm autonomously propagates the color to subsequent frames, enabling automatic colorization. Compared with optical-flow-based method RAFT [35] and segment-matching-based method AnT [8], our method can achieve more robust results on challenging cases such as one-to-many matching, large deformation, and tiny region colorization. In this figure, RAFT is trained on Sintel dataset [5] and finetuned on AnimeRun [30]. We use the most frequent color in each segment to colorize each line-enclosed region. ©drawn by Nicca (Sriprachum Kongwisawamit), used with artist permission.

Abstract

Colorizing line art is a pivotal task in the production of hand-drawn cel animation. This typically involves digital painters using a paint bucket tool to manually color each segment enclosed by lines, based on RGB values predetermined by a color designer. This frame-by-frame process is both arduous and time-intensive. Current automated methods mainly focus on segment matching. This technique migrates colors from a reference to the target frame by aligning features within line-enclosed segments across frames. However, issues like occlusion and wrinkles in animations often disrupt these direct correspondences, leading to mismatches. In this work, we introduce a new learning-based inclusion matching pipeline, which directs the network to comprehend the inclusion relationships between segments rather than relying solely on direct visual correspondences. Our method features a two-stage pipeline that integrates a coarse color warping module with an inclusion matching module, enabling more nuanced and accurate colorization. To facilitate the training of our network, we also develop a unique dataset, referred to as PaintBucket-Character. This

dataset includes rendered line arts alongside their colorized counterparts, featuring various 3D characters. Extensive experiments demonstrate the effectiveness and superiority of our method over existing techniques.

1. Introduction

Colorizing line art is a critical step in hand-drawn animation production. Initially, key animators create the characters' keyframes, followed by in-between animators who generate intermediate frames and refine keyframes for seamless connection between each stroke, as described in [31]. These line arts are then scanned into binary images, featuring a primary color palette of red, blue, green, and black, which represents highlights, shadows, special instructions (e.g., eye irises), and normal lines respectively. These scanned images, with regions of transparent background as shown in Fig. 1, are forwarded to digital painters. The painters use a paint bucket tool and a designated color palette to fill each line-enclosed segment, including the color lines, which is known as paint bucket colorization. Finally, special effect artists add visual enhancements, such as anti-aliasing, to

blend character animations with backgrounds, culminating in the final video composition, detailed in [20].

Line art colorization is a notably laborious process. For instance, even the relatively simple scenario in Fig. 1, involving 93 segments, still requires hundreds of manual clicks. To mitigate this demanding task, several commercial software applications like Retas Studio Paintman, OpenToonz, and CLIP Studio Paint have been developed. These tools provide helpful features that allow users to easily adjust assigned colors, fill small areas, and aid in line colorization. Despite these advancements, fully automated colorization has yet to be realized.

To address the challenges of line art colorization, numerous segment-matching methods have been developed for identifying corresponding segments in reference and target frames. Traditional graph-based approaches [17–19, 39, 44] consider each segment as a graph node and establish edges based on adjacency relationships. These methods employ optimization techniques to minimize topological similarity loss among corresponding nodes. However, they are time-intensive, often taking several minutes to process a single image with around 100 segments. Casey *et al.* [8] introduced the Animation Transformer (AnT) and its application Cadmium, which significantly improved segment matching. AnT standardizes the size of each segment and employs convolutional neural networks (CNNs) to extract features, transforming the image into a segment sequence. This converts segment matching into a sequence-to-sequence problem. Although AnT excels in aligning segments between adjacent frames featuring minimal motion, it struggles in more complex scenarios, particularly those with occlusion or pronounced motion. During occlusions, interlaced lines tend to fragment segments into multiple pieces, as depicted in Fig. 2. This fragmentation disrupts the strict correspondence between segments, rendering the task of segment matching inherently problematic and poorly defined.

In this study, we introduce the novel notion of ‘**inclusion matching**’, which obviates the necessity for exact segment-to-segment correspondence. Instead of direct segment matching, our approach computes the likelihood of each segment in a target frame being included within a specific region of a reference frame. Our methodology is unique in that it follows a **coarse-to-fine colorization** pipeline, inspired by the approach of skilled digital painters. This approach typically begins with coloring larger, more noticeable segments across all frames before filling in details of smaller segments in each frame. Motivated by this workflow, our pipeline first computes optical flow between line sketches, warping colors from a reference to a target frame to achieve a preliminary, coarse colorization. This is followed by the application of an inclusion matching module, which refines the colorization result of tiny segments.

The sparse nature of line art requires special formula-




Figure 2. Methods relying on segment matching typically seek the most similar segment across frames, yet challenges arise in scenarios involving occlusion and wrinkles. This is particularly evident in tiny segments, as highlighted in the red box, where disruptions in correspondence lead to mismatches. Our innovative approach, based on inclusion matching, addresses this issue by estimating the inclusion relationship rather than pursuing direct correspondence, as illustrated by the comparison of red and blue matching lines.

tion in our approach. Specifically, recognizing the necessity for a large receptive field in extracting line art features, we employ a U-Net structure [25] to process concatenated color and line features. To address the difficulty in matching segments under large deformations, we extract line semantic features using CLIP [23] and concatenate them with features from the U-Net’s bottleneck. This allows the network to match the region with the same semantic information across frames. Finally, we introduce super-pixel pooling at each segment to transfer the image into tokens for sequence-to-sequence matching.

Dataset plays a key role in the success of our method. To train and evaluate our method, we select 22 character models from Mixamo [3] and Aplaybox [1]. We mimic real intermediate line art data on the animation coloring production line. We disable the anti-aliasing and provide the color lines which is quite common in animation production. Then we use 3D software Cinema 4D to synthesize a specially-designed cartoon character colorization dataset named *PaintBucket-Character*.

The contributions of our work include the introduction of the *PaintBucket-Character* dataset tailored for paint bucket colorization tasks, the novel inclusion matching pipeline addressing the challenge of lacking strict segment correspondence, and a new coarse-to-fine baseline that outperforms existing solutions, particularly in scenes with occlusion. Our experiments demonstrate that our method surpasses the performance of recent AnT application Cadmium, and outperforms the results presented by Li *et al.* [30]. Notably, the advantages of our approach extend beyond dataset improvements, highlighting the effectiveness of our newly designed inclusion matching pipeline and the coarse-to-fine baseline in the context of animation colorization.



Figure 3. Several examples of our rendered characters in *PaintBucket-Character* dataset.

2. Related Work

Line Art Colorization. Line art colorization aims to create reasonable colors in the sketch’s blank regions. To achieve precise control, many user guidance types are designed including text [15, 45], scribble [6, 33, 40, 41], colored reference [4, 7, 9, 16, 32, 34, 37, 42, 43]. These methods mainly focus on pixel-wise estimation and cannot fill flat accurate RGB values assigned by the color designer in line arts’ segments, which always leads to flickering and color bleeding. Besides, in the post-processing, the special effect artists often need to select specific colors to create various effects such as ramp or glow for the character’s specific region. Thus, current line art colorization methods are still difficult to be merged into animation industrial processes.

Animation-related Dataset. To facilitate hand-drawn animation production, many real and rendered datasets are proposed. To boost video interpolation performance, Li *et al.* [29] propose the ATD-12K dataset which consists of triplet frames from 30 real animation movies. Due to the absence of correspondence labels for real animation, addressing this challenge has led to the adoption of 3D rendering techniques. MPI-Sintel [5] and CreativeFlow+ [28] provide different annotations such as optical flow, segmentation labels, and depth maps for multiple applications using Blender. To overcome the domain gap between rendered images and real 2D cartoons, Li *et al.* [30] developed the AnimeRun dataset using open-source 3D animations and their rendered optical flow. It improves segment matching accuracy and handles complex scenes.

Segment-based Animation Colorization. In the cleaning-up process of the animation industry, animators enclose each line manually, making segments serve as basic units for colorization. Thus, how to calculate segment-level semantic information and correspondence becomes the key in animation colorization. Traditional methods [17–19, 39, 44] regards segment as node and adjacency relationship as edge which transfer the segment matching to the graph optimization problem. Recently, Dang *et al.* [12] use Hu mo-

ments [14] to extract each segment’s feature to replace the RGB value and apply a UNet structure [25] to obtain the feature map. Subsequently, the feature is averaged in each segment to calculate the distance across frames for matching. Referring to the success of Transformer [36] in image matching [27], Casey *et al.* [8] apply multiplex transformer to aggregate information across frames to obtain more accurate results. Besides, Casey *et al.* replace the Hu moments algorithm with a segment feature extraction CNN which crop and resize each segment to fixed resolution first. Since both Dang *et al.* and Casey *et al.*’s methods and dataset are not publicly available, AnimeRun [30] remains the sole benchmark for segment matching available to researchers.

3. Paint Bucket Colorization Dataset

AnimeRun [30] employs rendered optical flow to calculate segment correspondence, leading to frequent one-to-many and many-to-one matching in the dataset. This results in inaccurate benchmark evaluations, as each segment may have multiple potential ground-truth correspondences. Furthermore, the limited dataset provided by AnimeRun, consisting of only 1,760 training images, increases the risk of network overfitting. Meanwhile, due to the anti-aliasing processing of AnimeRun’s rendered images, accurate color labels cannot be obtained in each segment. Consequently, AnimeRun faces difficulties in serving as a paint bucket colorization dataset directly.

To address these challenges, we introduce a new dataset *PaintBucket-Character*, comprising 11,345 training images and 3,200 test images. The test set includes 3,000 3D rendered frames and 200 hand-drawn frames. Notably, our dataset focuses solely on character animations, with consideration for the common practice among animators of just drawing foreground character animations while employing 3D models or separate paintings for the background. Our dataset focuses on character animations, considering the industry practice where animators typically sketch only the foreground character movements, while backgrounds are created using separate 3D models or paintings. Different from AnimeRun, our evaluation focuses on

Table 1. A comparison between our *PaintBucket-Character* dataset and previous segment matching datasets. Since AnT [8]’s dataset is not publicly available, we use ‘-’ to mark the unknown information. AnT does possess a private hand-drawn dataset with 3578 frames; however, it does not provide train-test-split information, indicated by ‘+’. Compared with previous open-source datasets, our dataset is more diverse and provides color lines that are closer to real hand-drawn animation. No anti-aliasing provides extra semantic information for each segment, which enables the network to learn better semantic information rather than segment shape similarity.

Dataset	Statistics				Annotations				
	Train	Test	Clips	Seg per Fr.	No anti-aliasing	Optical flow	Index label	Color label	Color line
AnimeRun [30]	1,760	1,059	20+10	237	×	✓	×	×	×
AnT [8] (private)	9,900+	1,100+	-	< 50	-	×	✓	✓	×
Ours	11,345	3,200	180+170	169	✓	×	✓	✓	✓




Figure 4. Overview of our synthetic data generation pipeline. We extract the UV map (a) and the used UV region from the character. Then, we use a paint bucket tool to fill the used UV region, creating color label (b) and index label (c). Label images (b) and (c) are then pasted back to the 3D meshes to create flat color characters (d) and (e), respectively. Finally, we post-process the index image (e) to obtain the rendered line art image (f).

segments’ color accuracy rather than traditional matching accuracy. To build our dataset, we sample 22 characters from Mixamo [3] and Aplaybox [1] and set 12 characters for training and 10 for testing. Then, we replot the UV map for these character models in flat color style and use Cinema 4D’s *Sketch and Toon Effect* to produce the line art. To ensure consistent color in each segment, we disable the anti-aliasing and change the texture sampling method to the nearest. As depicted in Fig. 3, our dataset showcases a diverse array of characters encompassing both Japanese and Western cartoon styles.

3.1. Synthetic Dataset Construction

To include temporal and occlusion relationships in our dataset, we apply Mixamo’s actions to our characters. For each character, we add five types of actions and animate the camera motion to capture different viewpoints of the character in each sequence. We classify the camera positions into three types: face shot, long shot, and close-up, in which we separately focus on the face, whole body, and other parts of the body. Following this pipeline, we obtain 15 clips for each character and form 180 clips for training and 150 clips for test. In the animation industry, keyframes always have a




Figure 5. Our simulation method can generate color lines that closely mimic hand-drawn line arts. Red, blue, and green lines represent highlights, shadows, and other instructions respectively.

low frame rate [29]. Thus, in each clip, we set the frame rate to 10 to mimic real animation and bring larger motion and deformation. Meanwhile, to address the fine scale of character details such as pupils and buttons, we use 1920×1920 as our training size and 1024×1024 as our test set size to provide clear segments in these areas.

Flat Color Style Rendering. To render characters in a flat color style, we use Cinema 4D to extract a UV mesh map and set unused regions as black. Then, we draw contour lines of character details such as pockets and shoelace on the UV mesh. Based on this original UV map, we use a paint bucket tool to colorize the mesh as the color label. To obtain a more precise matching relationship for training, we connect the label with the same semantic information (such as palm and back of the hand) and use the connected components algorithm for each color to extract the index label as shown in Fig. 4. Since we aim to achieve accurate line art to separate each segment, we calculate the difference between each color’s segment and its erosion as the black lines. This operation ensures that each line-enclosed segment contains clear semantic information and comprises only one color.

Color Line Extraction. To simulate color lines which are common in real hand-drawn animation’s shadow and highlight regions, we mark these highlight and shadow colors and erode these regions to create color lines, as shown in Fig. 5. Compared with AnimeRun [30], our dataset features more realistic line art and provides clear semantic information for each segment. Since our characters are rendered with discrete colors, we can easily obtain specific meanings for each segment.

3.2. Hand-drawn Dataset Collection

Obtaining intermediate unprocessed data poses a challenge, as common commercial animations undergo post-

processing techniques including anti-aliasing. To address this challenge, we invite some professional animators to draw line art animation clips with different characters. Then, we use the paint bucket tool to colorize these line arts as our ground truth to synthesize paired line art and colorized data. Besides, animation creation software such as Retas Studio Paintman and CLIP Studio Paint all provide animation production tutorials. We collect animations in these tutorials and use the color line extraction method to synthesize the paired data. Finally, we collect a hand-drawn test dataset with 200 frames and 20 clips for evaluation.

4. Method

4.1. Inclusion Matching

Given two line art images $L_i, L_j \in \mathbb{R}^{H \times W \times 3}$, lines separate the blank area into N_i and N_j 's line-enclosed segments. In the reference frame L_i , each segment is colorized with M colors whose color labels can be represented as one-hot representation $c_i \in \mathbb{R}^{N_i \times M}$. To estimate the L_j 's color label $\hat{c}_j \in \mathbb{R}^{N_j \times M}$, previous segment-matching-based methods [8, 30] mainly focus on calculating segments' visual correspondence with the matching $\mathbf{m}^{j \rightarrow i} \in \mathbb{R}^{N_j \times N_i}$. Nevertheless, the absence of strict correspondence among segments imposes limitations on the model's learning capacity.

To tackle this challenge, we introduce a novel method called inclusion matching. In the reference frame, segments with identical semantic information are merged, creating K_i large regions with different labels. As each line-enclosed segment in the target frame can only possess one color with a specific meaning, it can be identified as a subset of one of the K_i regions. Specifically, in the training process, the index label stated in Fig. 4 separates the reference frame into K_i region. We train a network to estimate each segment's index label $\mathbb{R}^{N_j \times K_i}$. To enrich the diversity, we randomly merge the adjacent labels in the reference frame to create larger regions to guide the network to learn a more robust inclusion relationship, as shown in Fig. 6.

4.2. PaintBucket Colorization Pipeline

Color Warping Module. Colorized line art carries substantial information about how segments are grouped, making it crucial for inclusion matching. However, previous segment matching methods [8, 30] exclusively leverage the line art's features for matching and do not incorporate the colorized reference. To encode the color feature, we recolorize the index label image. Specifically, we partition the RGB space into K_i cubes, each with a side length of $255 \times K_i^{-\frac{1}{3}}$. Subsequently, we select the RGB value at the center of each cube and randomly assign them to index labels as the reference color image C_i . We employ the optical flow estimation model RAFT [35], fine-tuned on AnimeRun [30], to esti-




Figure 6. Our inclusion matching training pipeline employs index labels to partition the reference frame into K_i regions, which are then randomly merged with adjacent index labels with a probability of 20%. Subsequently, the model estimates the affiliation of each line-enclosed segment in the target frame with a specific region in the reference, thereby facilitating the learning of the inclusion relationship.

mate the optical flow between the reference line art L_i and the target line art L_j . Utilizing this optical flow, we warp the reference color image C_i to generate the coarse color estimation C_j for the target frame.

Feature Extraction Module. Lacking textual information and large deformation in animation make warped color image C_j hard to precisely align with the line art in the target frame. To address this, we employ deformable convolution [11] on the warped color C_j to extract aligned features. For a given deformable convolution kernel of K sampling positions, we use ω_k and \mathbf{p}_k to represent the weight and pre-specified offset. The aligned features F_a at each position \mathbf{p}_0 can be expressed as:

$$F_a(\mathbf{p}_0) = \sum_{k=1}^K \omega_k \cdot F_a(\mathbf{p}_0 + \mathbf{p}_k + \Delta \mathbf{p}_k), \quad (1)$$

where learnable offset $\Delta \mathbf{p}_k$ can be predicted from the offset estimation network as depicted in Fig. 8. This network, designed as a lightweight U-Net [25] denoted as Φ , is employed to expand the receptive field:

$$\Delta \mathbf{p}_k = \Phi([L_j, C_j]). \quad (2)$$

For the reference frame, we deactivate the offset estimation module and set the offset to zero. Subsequently, the aligned color feature is concatenated with the line feature and passed through a feature encoder network. To guide the network in learning semantic information matching, we resize the line art image to 320×320 and utilize CLIP [23]'s image encoder to extract features. To preserve the feature's positional information, we choose to use the ConvNext-Large model from OpenCLIP [10]. We only selected the initial layers of CLIP, which downsample the feature to a resolution of 40×40 . The CLIP feature is interpolated to concatenate with the encoded features. Finally, the feature decoder decodes these concatenated features to the input's

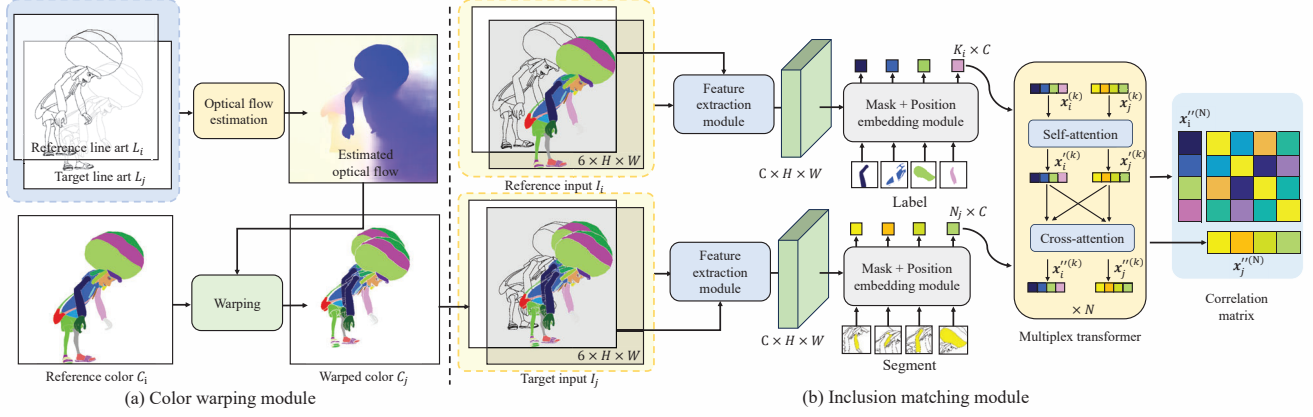


Figure 7. Our paint bucket colorization architecture. Initially, we estimate the optical flow based on the line arts and warp the reference color as the coarse colorization result. The mask and position embedding module tokenizes the segments and features extracted by the feature extraction module into a sequence, which is then fed into the multiplex transformer for information aggregation. Finally, we compute a similarity matrix between reference and target tokens to calculate matching loss and predict the final color.




Figure 8. Overall pipeline of our feature extraction module. Deformable convolution is applied to align the color features of the target frame to the line image. Then, the line features and aligned color features are concatenated and fed to the feature extraction network in a U-Net structure. Line art’s CLIP features are also concatenated to the bottleneck of the feature extraction network.

resolution with C channels. To preserve features in high resolution, shortcuts are added between the network’s encoder and decoder, following the structure of the U-Net.

Mask + Position Embedding Module. To transform different segments into tokens, we employ super-pixel pooling to average the features within each segment’s mask, resulting in a C -dimensional visual token, denoted as \mathbf{d}_i . Subsequently, a Multilayer Perceptron (MLP) layer is utilized to encode the segment’s bounding box’s four-dimensional coordinates, denoted as p_i , into a C -dimensional positional embedding. The token \mathbf{x}_i for matching is obtained by adding the positional embedding and visual token:

$$\mathbf{x}_i = \mathbf{d}_i + \text{MLP}(p_i). \quad (3)$$

Differing from previous segment matching pipelines [8, 30], our approach matches each line-enclosed segment in the target frame with a specific index label in the reference frame, rather than employing label-to-label matching.

Multiplex Transformer. Building upon [8, 27], we employ multiplex self- and cross-attention layers to aggregate tokenized features. In a self-attention layer, queries, keys, and

values are derived from a single source feature:

$$\text{self-attention}(\mathbf{x}_i) = \text{softmax} \left(\frac{\mathbf{Q}_i \mathbf{K}_i}{\sqrt{D}} \right) \mathbf{V}_i, \quad (4)$$

where \mathbf{Q}_i , \mathbf{K}_i , and \mathbf{V}_i represent \mathbf{x}_i processed by MLPs for query, key, and value. In contrast, the cross-attention layer computes keys and values from another feature:

$$\text{cross-attention}(\mathbf{x}'_i, \mathbf{x}'_j) = \text{softmax} \left(\frac{\mathbf{Q}'_i \mathbf{K}'_j}{\sqrt{D}} \right) \mathbf{V}'_j. \quad (5)$$

Following each self- or cross-attention layer, the output is added to the original \mathbf{x}_i or \mathbf{x}'_i , and then processed with a feed-forward MLP. After repeating these operations N times, as illustrated in Fig. 7, we obtain the final aggregated features $\hat{\mathbf{x}}_i \in \mathbb{R}^{K_i \times C}$ and $\hat{\mathbf{x}}_j \in \mathbb{R}^{N_j \times C}$.

Loss function. After the propagation step, the similarity matrix $\mathbf{S} \in \mathbb{R}^{N_j \times K_i}$ is represented as follows:

$$\mathbf{S}_{mn} = \frac{\exp(\hat{\mathbf{x}}_{im} \cdot \hat{\mathbf{x}}_{jn})}{\sum_{m=1}^{N_j} \exp(\hat{\mathbf{x}}_{im} \cdot \hat{\mathbf{x}}_{jn})}. \quad (6)$$

This matrix can be interpreted as the color probability $\hat{\mathbf{y}}_m \in \mathbb{R}^{K_i}$ for segment m in the target frame. Given that each frame has a one-hot ground truth index label \mathbf{y}_m , the loss function is derived from the cross-entropy loss:

$$L_{ce} = - \sum_{m=1}^{K_i} \mathbf{y}_m \log(\hat{\mathbf{y}}_m). \quad (7)$$

Inference Pipeline. Since real animation’s colors have an obvious domain gap with our randomly generated colors, we adopt a color redistribution strategy to reassign randomly generated colors to replace the original colors during the warping and matching pipeline. This strategy can avoid the mismatching problem when several predefined colors




Figure 9. The visualization results of our proposed methods and other approaches on real hand-drawn animation. Compared with previous methods, our proposed approach demonstrates enhanced robustness in handling occlusion situations (e.g., hand of the running girl and shadow near eyebrows of the girl in the second row.)

Table 2. Quantitative comparison of our method with different methods. ‘Acc’ and ‘Acc-Thres’ denote segment-wise accuracy, providing insights into the potential workload reduction for digital painters. In ‘Acc-Thres,’ segments smaller than 10 pixels are filtered out. ‘Pix-Acc,’ ‘Pix-F-Acc,’ and ‘Pix-B-MIoU’ represent pixel-wise accuracy, foreground pixel-wise accuracy, and pixel-wise background MIoU, respectively, reflecting the visualization performance. All RAFT models undergo training on MPI-Sintel [5], and ‘F/C’ indicates the models fine-tuned on AnimeRun’s colored frames and contours.

Type	Method	3D rendered test set					Hand-drawn test set				
		Acc	Acc-Thres	Pix-Acc	Pix-F-Acc	Pix-B-MIoU	Acc	Acc-Thres	Pix-Acc	Pix-F-Acc	Pix-B-MIoU
Reference-based	ControlNet	0.1709	0.1856	0.4274	0.2653	0.4600	0.1715	0.1812	0.2449	0.2298	0.2260
Segment matching	AnT(AnimeRun)	0.5584	0.5930	0.8734	0.6537	0.8954	0.6779	0.7043	0.9193	0.7377	0.9429
	AnT(Cadmium)	0.6634	0.7713	0.9765	0.9471	0.9801	0.7224	0.7879	0.9792	0.9303	0.9899
Optical flow	RAFT	0.6589	0.7456	0.9724	0.9099	0.9763	0.7528	0.8093	0.9819	0.9190	0.9898
	RAFT (F)	0.7133	0.8054	0.9818	0.9449	0.9852	0.7758	0.8424	0.9860	0.9425	0.9922
	RAFT (C)	0.7028	0.7944	0.9723	0.9411	0.9690	0.7718	0.8356	0.9833	0.9375	0.9916
Inclusion matching	Ours	0.8266	0.8726	0.9905	0.9724	0.9948	0.8593	0.8929	0.9900	0.9638	0.9984

are too closed. Finally, the estimated color for each line-enclosed segment $\hat{c}_j \in \mathbb{R}^{N_j \times M}$ can be derived by:

$$\hat{c}_j = \hat{S}c_i, \quad (8)$$

where $\hat{S} \in \mathbb{R}^{N_j \times N_i}$ is the similarity matrix for each line-enclosed segment in reference and target frames. In this paper, we set the match threshold as 0.2 to filter out all the segments which are unmatchable.

5. Experiments

5.1. Implementation Details

During the training stage, our model undergoes 300,000 iterations with a batch size of 2, utilizing NVIDIA GeForce RTX 3090 GPUs. We employ the Adam optimizer with a learning rate of 10^{-4} and no weight decay. Throughout the training process, we keep the weights of the pretrained optical flow estimation module and the CLIP image encoder fixed. In the feature extraction module, the offset estimation module is structured as a lightweight U-Net with three down-sampling layers and a bottleneck featuring 128 features. Simultaneously, the feature extraction network is implemented as a U-Net with four down-sampling layers and a

bottleneck containing 512 features. The channel number C for our extracted feature is set to 128. Within the multiplex transformer, the layer number N is set to nine, and the number of heads is configured as four. The data augmentation pipeline is shown in the supplementary material.

5.2. Comparisons with Previous Methods

Since there is no existing paint bucket colorization method, we mainly compare our method with reference-based line art colorization [42], segment matching methods [8, 30], and optical flow estimation methods [35]. The quantitative comparison is shown in Table 2. The visual comparison of real hand-drawn animation and rendered 3D test sets are shown in Fig. 9 and Fig. 10.

Reference-based Line Art Colorization. The saturated luminance of the blank regions within line art impedes the application of reference-based methods for natural gray-scale image colorization. Thus, we choose ControlNet [42] as our reference-based line art colorization method. We use Stable Diffusion v1.5 [24] as our base model and employ the ControlNet’s line art checkpoint to encode the line art condition. Then, we adopt the IP-Adapter [38] to encode the reference image. We set the number of steps to 50 and use

Table 3. Quantitative comparison of different datasets, different network structures, and different methods on test data. ‘IM’ signifies the inclusion matching pipeline, which incorporates adjacent index label merging and segment-label matching strategies. Since the inclusion matching pipeline and coarse warping module with RAFT can also be applied to other networks, we also adopt them to AnT for comparison to illustrate the generalization of these methods.

Type	Method	3D rendered test set					Hand-drawn test set				
		Acc	Acc-Thres	Pix-Acc	Pix-F-Acc	Pix-B-MIoU	Acc	Acc-Thres	Pix-Acc	Pix-F-Acc	Pix-B-MIoU
Baseline	Ours	0.8266	0.8726	0.9905	0.9724	0.9948	0.8593	0.8929	0.9900	0.9638	0.9984
Dataset	AnimeRun	0.7359	0.7958	0.9693	0.9201	0.9765	0.8278	0.8634	0.9818	0.9399	0.9906
	w/o CLIP	0.8180	0.8653	0.9873	0.9685	0.9910	0.8489	0.8843	0.9878	0.9570	0.9925
Model	w/o RAFT	0.8001	0.8447	0.9824	0.9481	0.9878	0.8446	0.8750	0.9804	0.9404	0.9851
	w/o DConv	0.8247	0.8685	0.9894	0.9709	0.9928	0.8525	0.8833	0.9875	0.9549	0.9937
	Ours w/o IM	0.7637	0.8352	0.9721	0.9573	0.9688	0.8348	0.8763	0.9777	0.9495	0.9822
Method	AnT	0.7450	0.7930	0.9708	0.9361	0.9758	0.8162	0.8497	0.9817	0.9374	0.9905
	AnT with IM	0.8063	0.8447	0.9837	0.9551	0.9887	0.8388	0.8698	0.9839	0.9388	0.9943
	AnT with IM/RAFT	0.8177	0.8587	0.9879	0.9681	0.9932	0.8483	0.8828	0.9886	0.9598	0.9946

the text prompt ‘A character in the white background’. We calculate the average RGB value per segment and match the average color with the nearest color in the reference image for evaluation.

Optical Flow Estimation. Since optical flow methods cannot produce segment-wise estimation without post-processing, we warp the color and calculate the colors that appear most frequently in each segment as the baseline of optical estimation method. Since finetuning RAFT on AnimeRun can promote the optical flow’s performance effectively, we use the finetuned RAFT in the following sessions.

Segment Matching. We also compare our method with the Cadmium application, the official implementation of AnT [8], and AnimenRun’s reimplemented AnT [30]. The visual comparison shown in Fig. 9 and Fig. 10 demonstrate the robust performance of our proposed method, particularly in scenarios where strict segment correspondence is absent, and challenges with large motion and complex deformations. More visual results are presented in the suppl. material.

5.3. Ablation Study

Dataset. To demonstrate the effectiveness of our proposed dataset, we also train the proposed method on the AnimeRun. In the training process, we adopt the inclusion matching strategy by merging the segments in the reference frames. The quantitative results reported in Table 3 suggest the positive influence of our proposed dataset.

Model. We conduct an ablation study on our model design to examine the necessity of each structure. In our experiments, we explore alternative configurations by replacing the deformable convolution and offset estimation module with the original convolution, eliminating the coarse warping module, and omitting the CLIP image encoder for ablation. The results presented in Table 3 underscore the significance of each module, highlighting the crucial role of the coarse warping module.

Method. The results in Table 3 further highlight the efficacy of our inclusion matching pipeline. Its application to AnT [8] yields a substantial improvement in network performance, suggesting the broad applicability and generalization of our inclusion matching pipeline. Ablation study




Figure 10. Visual comparison of different methods on our 3D rendered test set. Our proposed method produces satisfactory results in the presence of occlusion, significant motion, and substantial changes in viewing angles.

visualizations are available in the supplementary material.

6. Conclusion

We have introduced a new approach for line art paint bucket colorization. Departing from traditional segment matching methods that struggle with one-to-one correspondence, we proposed a novel inclusion matching pipeline that estimates each segment’s inclusion relationship rather than the direct correspondence. To facilitate the learning of this inclusion relationship, we presented a dedicated dataset, *PaintBucket-Character*, and a two-stage colorization pipeline. Our pipeline incorporates a coarse color warping module and an inclusion matching module. The experimental results demonstrate the superior performance of our method in handling occlusion and accommodating large movements within challenging scenarios.

Acknowledgement. This study is supported under the RIE2020 Industry Alignment Fund Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] Aplaybox. <https://www.aplaybox.com/>. 2, 4
- [2] Cadmium. <https://cadmium.app/>. 2
- [3] Mixamo. <https://www.mixamo.com/>. 2, 4
- [4] Kenta Akita, Yuki Morimoto, and Reiji Tsuruno. Hand-drawn anime line drawing colorization of faces with texture details. *Computer Animation and Virtual Worlds*, 2023. 3
- [5] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 1, 3, 7
- [6] Ruizhi Cao, Haoran Mo, and Chengying Gao. Line art colorization based on explicit region segmentation. In *Computer Graphics Forum*, 2021. 3
- [7] Yu Cao, Xiangqiao Meng, PY Mok, Xueting Liu, Tong-Yee Lee, and Ping Li. Animediffusion: Anime face line drawing colorization via diffusion models. *arXiv preprint arXiv:2303.11137*, 2023. 3
- [8] Evan Casey, Víctor Pérez, and Zhuoru Li. The animation transformer: Visual correspondence via segment matching. In *ICCV*, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [9] Shu-Yu Chen, Jia-Qi Zhang, Lin Gao, Yue He, Shihong Xia, Min Shi, and Fang-Lue Zhang. Active colorization for cartoon line drawings. *TVCG*, 28(2), 2020. 3
- [10] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *CVPR*, 2023. 5
- [11] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, 2017. 5
- [12] Trung D. Q. Dang, Thien Do, Anh Nguyen, Van Pham, Quoc Nguyen, Bach Hoang, and Giao Nguyen. Correspondence neural network for line art colorization. In *SIGGRAPH*, 2020. 3
- [13] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *ICLR*, 2022. 3, 4
- [14] Ming-Kuei Hu. Visual pattern recognition by moment invariants. *IRE transactions on information theory*, 8(2), 1962. 3
- [15] Hyunsu Kim, Ho Young Jhoo, Eunhyeok Park, and Sungjoo Yoo. Tag2Pix: Line art colorization using text tag with SE-Cat and changing loss. In *ICCV*, 2019. 3
- [16] Junsoo Lee, Eungyeup Kim, Yunsung Lee, Dongjun Kim, Jaehyuk Chang, and Jaegul Choo. Reference-based sketch image colorization using augmented-self reference and dense semantic correspondence. In *CVPR*, 2020. 3
- [17] Shaolong Liu, Xingce Wang, Zhongke Wu, and Hock Soon Seah. Shape correspondence based on kendall shape space and rag for 2d animation. *The Visual Computer*, 36, 2020. 2, 3
- [18] Shaolong Liu, Xingce Wang, Xiangyuan Liu, Zhongke Wu, and Hock Soon Seah. Shape correspondence for cel animation based on a shape association graph and spectral matching. *Computational Visual Media*, 2023.
- [19] Akinobu Maejima, Hiroyuki Kubo, Takuya Funatomi, Tatsuo Yotsukura, Satoshi Nakamura, and Yasuhiro Mukaigawa. Graph matching based anime colorization with multiple references. In *SIGGRAPH*. 2019. 2, 3
- [20] Koki Ohira. *After Effects For Animation*. 2010. 2
- [21] Kyoto Animation publishing department. *Guide Kyoto Animation version finish, shooting*. Kyoto Animation Co., Ltd, 2010. 1
- [22] Kyoto Animation publishing department. *The Kyoto Animation Guide to Drawing*. Kyoto Animation Co., Ltd, 2010. 1
- [23] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 5
- [24] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 7, 3
- [25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*. Springer, 2015. 2, 3, 5
- [26] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 3
- [27] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *CVPR*, 2020. 3, 6
- [28] Maria Shugrina, Ziheng Liang, Amlan Kar, Jiaman Li, Anagad Singh, Karan Singh, and Sanja Fidler. Creative flow+ dataset. In *CVPR*, 2019. 3
- [29] Li Siyao, Shiyu Zhao, Weijiang Yu, Wenxiu Sun, Dimitris Metaxas, Chen Change Loy, and Ziwei Liu. Deep animation video interpolation in the wild. In *CVPR*, 2021. 3, 4
- [30] Li Siyao, Yuhang Li, Bo Li, Chao Dong, Ziwei Liu, and Chen Change Loy. Animerun: 2d animation visual correspondence from open source 3d movies. *NeurIPS*, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [31] Li Siyao, Tianpei Gu, Weiye Xiao, Henghui Ding, Ziwei Liu, and Chen Change Loy. Deep geometrized cartoon line inbetweening. In *ICCV*, 2023. 1
- [32] Daniel Šỳkora, Jan Buriánek, and Jiří Žára. Unsupervised colorization of black-and-white cartoons. In *International Symposium on Non-photorealistic Animation and Rendering*, 2004. 3
- [33] Daniel Šỳkora, John Dingliana, and Steven Collins. Lazybrush: Flexible painting tool for hand-drawn cartoons. In *Computer Graphics Forum*, 2009. 3
- [34] Daniel Šỳkora, John Dingliana, and Steven Collins. As-rigid-as-possible image registration for hand-drawn cartoon animations. In *International Symposium on Non-photorealistic Animation and Rendering*, 2009. 3
- [35] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 1, 5, 7, 2

- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 3
- [37] Shukai Wu, Xiao Yan, Weiming Liu, Shuchang Xu, and Sanyuan Zhang. Self-driven dual-path learning for reference-based line art colorization under limited data. *IEEE TCSVT*, 2023. 3
- [38] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arxiv:2308.06721*, 2023. 7, 3
- [39] Lei Zhang, Hua Huang, and Hongbo Fu. EXCOL: An EXtract-and-COMplete layering approach to cartoon animation reusing. *TVCG*, 18(7), 2012. 2, 3
- [40] Lvmin Zhang, Chengze Li, Tien-Tsin Wong, Yi Ji, and Chungping Liu. Two-stage sketch colorization. *ACM TOG*, 37(6), 2018. 3
- [41] Lvmin Zhang, Chengze Li, Edgar Simo-Serra, Yi Ji, Tien-Tsin Wong, and Chungping Liu. User-guided line art flat filling with split filling mechanism. In *CVPR*, 2021. 3
- [42] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 3, 7, 4
- [43] Qian Zhang, Bo Wang, Wei Wen, Hai Li, and Junhui Liu. Line art correlation matching feature transfer network for automatic animation colorization. In *WACV*, 2021. 3
- [44] Haichao Zhu, Xueting Liu, Tien-Tsin Wong, and Pheng-Ann Heng. Globally optimal toon tracking. *ACM TOG*, 35(4), 2016. 2, 3
- [45] Changqing Zou, Haoran Mo, Chengying Gao, Ruofei Du, and Hongbo Fu. Language-based colorization of scene sketches. *ACM TOG*, 38(6), 2019. 3