Delft University of Technology

Delaunay Painting

Perceptual Image Colouring from Raster Contours with Gaps

Parakkat, Amal Dev; Memari, Pooran; Cani, Marie Paule

DOI
10.1111/cgf.14517

Publication date
2022

Document Version
Final published version

Published in
Computer Graphics Forum

Citation (APA)
Parakkat, A. D., Memari, P., & Cani, M. P. (2022). Delaunay Painting: Perceptual Image Colouring from Raster Contours with Gaps. *Computer Graphics Forum*, *41*(6), 166-181. https://doi.org/10.1111/cgf.14517

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.
For technical reasons the number of authors shown on this cover page is limited to a maximum of 10.

DOI: 10.1111/cgf.14517

COMPUTER GRAPHICS *forum*
Volume 0 (2022), number 0 pp. 1–16

Delaunay Painting: Perceptual Image Colouring from Raster Contours with Gaps

Amal Dev Parakkat,[1,2] Pooran Memari[3] and Marie-Paule Cani[3]

[1]Delft University of Technology, Delft, Netherlands
[2]LTCI—Telecom Paris, Institut Polytechnique de Paris, Palaiseau, France
memari@lix.polytechnique.fr
[3]LIX, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France
memari@lix.polytechnique.fr,marie-paule.cani@polytechnique.edu

Abstract

We introduce Delaunay Painting, a novel and easy-to-use method to flat-colour contour-sketches with gaps. Starting from a Delaunay triangulation of the input contours, triangles are iteratively filled with the appropriate colours, thanks to the dynamic update of flow values calculated from colour hints. Aesthetic finish is then achieved, through energy minimisation of contour-curves and further heuristics enforcing the appropriate sharp corners. To be more efficient, the user can also make use of our colour diffusion framework, which automatically extends colouring to small, internal regions such as those delimited by hatches. The resulting method robustly handles input contours with strong gaps. As an interactive tool, it minimizes user's efforts and enables any colouring strategy, as the result does not depend on the order of interactions. We also provide an automatized version of the colouring strategy for quick segmentation of contours images, that we illustrate with applications to medical imaging and sketch segmentation.

Keywords: sketch coloring, image processing, assistive interfaces, interaction, computational geometry, modelling, shape completion

CCS Concepts: • Computing methodologies → Image manipulation; Shape analysis; • Applied computing → Fine arts; • Theory of computation → Computational geometry

1. Introduction

Colouring raster contour images, such as line arts, is an essential step in many image manipulation tasks. Flood-fill, applied by the bucket filling tool present in most standard software such as Microsoft Paint or GIMP, is a simple solution for colour filling. Unfortunately, it cannot handle contours with gaps, although gaps are a common error in hand-drawn line art. Professional artists may also voluntarily leave gaps (also called a 'Professional Gaps') as part of their style while sketching line art. These gaps are small breaks in a line mimicking the light reflecting off an object. Gaps are also considered as an easy way to create transitions while drawing curves or long lines. Movies like 'Ernest and Celestine' or 'Le Grand Méchant Renard et Autres Contes' use such minimalist sketching style with gaps, enhanced with colouring. Though they make perfect sense from an artist's perspective, colouring line arts with gaps is a difficult, and time-consuming task for digital artists, which is more hectic when it needs to be done for a full animation movie.

We present a specialized algorithm, which efficiently handles such colouring process, and is—contrary to flood-fill—robust to the absence of well-defined boundaries in the input. To achieve this, we rely on a novel, yet intuitive Delaunay-based framework. The main idea is to use Delaunay triangulation to create a graph between regions delimited by the contours, and then control a colour flow through graph edges, from a few, input colour hints.

We introduce two versions of our colouring procedure, depending on how the colour hints are obtained: an interactive version, and an automatic version. While the interactive colouring procedure is useful for artistic colouring of contour drawings (where the colour hints fully depend on the artist's creativity, and cannot be predicted automatically), the automatic colouring approach can be used for automatically segmenting contour sketches into regions. We illustrate it

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd
This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

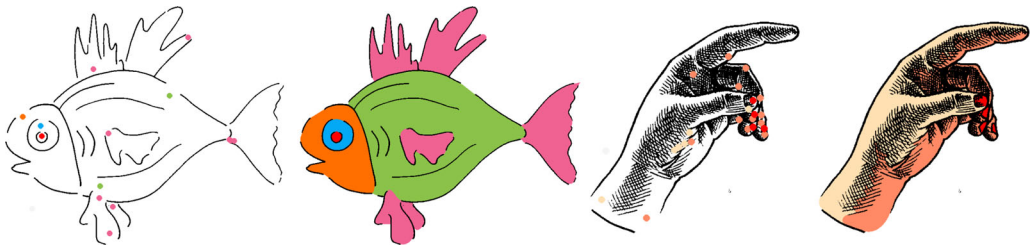

Figure 1: *From left to right: A simple line-art and a more complex sketch along with user hints, and the corresponding results.*

by an application to medical image segmentation. While such automatic colouring is straightforward, making the interactive colouring approach usable for colouring line-arts brings additional challenges such as the response time, intuitiveness and aesthetic beauty of the final results. Hence, the following objectives are taken into consideration while designing the interactive method:

- The tool should be as simple to use as Flood-fill: Everyone is familiar with the bucket-fill tool. It is easy to use and straightforward to understand. Such intuitive interaction is required.
- Visually pleasing results: Unless the gaps are correctly filled, the final coloured results will not look complete. So the border of coloured regions delimited by disconnected contours should reconnect gaps in an aesthetically pleasing way.
- Minimum user-interaction: The user should be able to colour visually salient regions in the contour with minimum interaction. Hence, an appropriate colour flow mechanism should be provided to spread colours.

To achieve the first point, our method uses an intuitive flow-based colour propagation process on the underlying Delaunay triangulation [PCS21], making it as simple to use as Flood-fill. Second, instead of reconnecting gaps in the contours with straight lines, which would affect the aesthetic beauty of the final result, we make use of local contour information around gaps for generating the border of coloured regions. More precisely, we make use of the concept of Scale Invariant Minimum Variation Curves (SIMVC) and a sharp corner detection heuristic to delimit the coloured region in a visually pleasing way (see Figure 1, left). Lastly, we introduce a simple colour diffusion mechanism to achieve the objective of minimal user input, even for drawings with additional information such as texture or hatches. In existing methods, contours sharply bounds colour spreading, requiring the user to paint each region separately, whatever their size. This becomes a cumbersome task when there are many small regions, which commonly occurs with complex raster sketches (see Figure 1, right). In contrast, we use a simple diffusion method based on the perceptual distance between regions (applicable even without any gap between them), enabling colours to seamlessly spread from user coloured areas to the uncoloured parts of the image. This colour spreading procedure is reversible and also can be easily updated.

Overall, the proposed interactive Delaunay colouring method is straightforward and saves the user from many hurdles when colouring raster contour images with gaps: As simple as flood-fill, it only requires a minimal number of colour hints, the colour filling is fast and easily predictable, and the user can easily correct any mistake, since our method is order-independent and reversible.

In summary, while our method reuses the flow mechanism from Parakkat *et al*. [PCS21] in the new context of colouring raster contour images with minimal user inputs, it also makes new technical contributions:

- A colour diffusion mechanism to minimize user efforts by propagating colours to uncoloured small regions (for examples hatches)
- An aesthetic curve completion technique to fill gaps in the input contours in a visually pleasing way
- An automatic colouring algorithm with possible applications in segmentation of medical images and sketches

2. Related Work

Related work can be classified into two categories, namely automatic versus user-guided colourization methods.

Automatic colourization methods make use of prior knowledge to colour an input image without any user intervention [LWCO*07, ISSI16, YBC*19, SCH20]. Most recent works use machine learning techniques for colourizing greyscale photographs, where the colour to be used is predicted through an analysis of the input [ISSI16, YBC*19, SCH20]. This kind of work is not applicable to our line art colouring problem, due to the absence of extra information such as scales of grey in our input. Moreover, results would then be limited to quite standard, learned colouring palettes, with no user control. Another set of similar works [HA17, FHOO17] make use of a single, reference coloured image to identify the appropriate colours to be used, which again restricts the set of colours as well as the colouring style.

User-guided colourization methods were introduced to give more freedom to digital artists. The latter are asked to provide additional hints to guide colourization, such as colour brush-strokes, or even text input such as specifying the region and the colour that should be used to fill it [ZMG*19, KJPY19, CSG*17]. Indeed, this last category of methods can only be used for colouring meaningfully, well-identified shapes enhanced with semantic information, and hence cannot be used for generic inputs. It is worth noting that a few learning-based methods, such as Shi *et al*. [SZC*20] enabled to colour line arts in videos as well. In this paper, we focus on colouring methods based on user-defined coloured brush strokes, which were extensively developed.

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

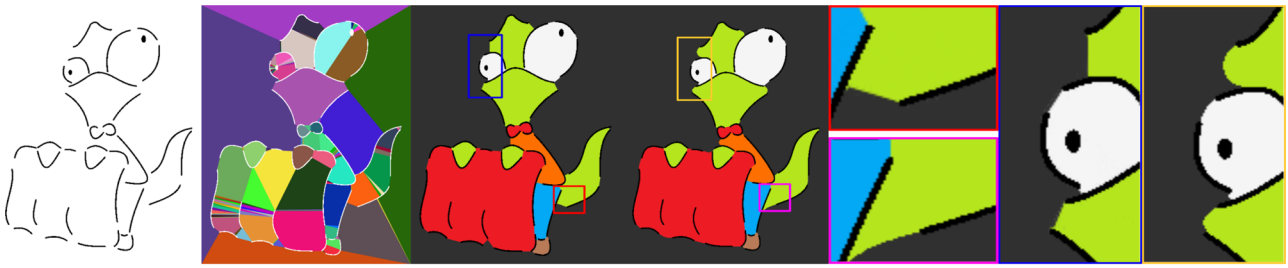

(a) (b) (c) (d) (e) (f) (g)

Figure 2: *Comparison with Parakkat* et al. *[PMGC20]. (a) Input sketch, (b) segmentation and (c) colouring result using [PMGC20], (d) our colouring result. The main differences can be seen in the insets ((e) top: Parakkat* et al. *[PMGC20], bottom: ours, (f) Parakkat* et al. *[PMGC20] and (g) ours).*

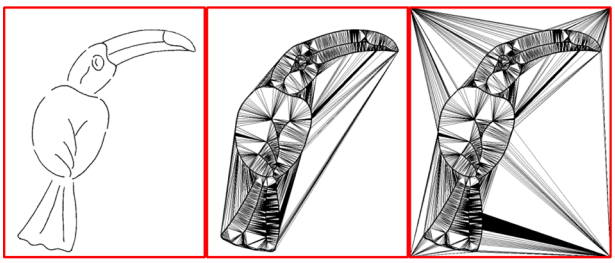

(a) (b) (c)

Figure 3: *(a) Input sketch, (b) Delaunay triangulation of extracted sketch pixel positions, (c) Delaunay triangulation after padding points.*

One of the prominent works in brush-based assisted colouring is Lazybrush [SDC09], where the authors express the problem of colouring based on user hints as an energy minimization problem. Various machine learning-based solutions were also introduced [LWCO*07, ZLW*18, HJRD19, Pai20]. As will be shown in Section 8.1, these methods do not necessarily generate a flat colouring result. Moreover, the context of Luan *et al*. [LWCO*07], in which users apply colour-strokes to a greyscale photograph, is quite different from ours. Closer to our sketch-colouring goals, Zhang *et al*. [ZLW*18], Hati *et al*. [HJRD19], PaintsChainer [Pai20] and Zhang *et al*. [ZLSS*21] address the lack of user control of pure, learning-based approaches by enabling the user to add multiple, additional colour hints to influence the result. However, when not applied to a specific category of drawings on which the model could be trained (e.g. manga characters), such solutions typically require extensive user intervention, since the same region may need to be recoloured multiple times until the desired output is achieved. When the goal is not to address a specific class of drawings, methods that do not use machine learning may thus be more suitable.

Colouring contour sketches may also be considered as a variant of a **contour closing problem**, a problem tackled since 1994 [GTF94], and which was recently investigated using learning-based approaches [SISSI17, LDL*19]. In Parakkat and Muthuganapathy [PM16], the authors pose the classical curve reconstruction problem as a way to connect gaps in a contour sketch, enabling the newly created regions to be coloured using the standard flood-fill tool. Though more generic curve reconstruction methods were later proposed (e.g. [PMM18]), the effectiveness of these algorithms typically depend on sampling density and stroke positions. Moreover, using standard flood-fill after reconnecting contours reduces user's freedom, since only uniform colours can then be applied within the newly closed regions. This prevents the user from giving multiple colours hints within a given region, as we did, for instance, on the wrist of the hand in Figure 1. Instead of creating connections as in curve reconstruction to fill gaps, Fourey *et al*. [FTR18] connected appropriate strokes to create a set of segments in the contours. The regions are then interactively updated and coloured with the help of the user. The method, which was later embedded in Gimp as a line-art colouring tool, is easier to use than the previous ones. However, user freedom is still limited since it can colour only a set of segments, and the boundaries of the resulting sketches are not necessarily well formed, as will be shown through comparison to our solution in Section 8.1.

In this work, our goal is to propose a simple and flexible colouring technique that can be used as a user-friendly toolbox in various sketch colouring scenarios. We, therefore, do not rely on any optimization or machine-learning-based approach, which avoids the need for extensive user input. Our solution belongs to geometry-based methods, but specifically addresses the issues of user freedom and quality of the resulting coloured sketches, by introducing an intuitive and light-weight framework for colouring of raster contour sketches with gaps.

Related Delaunay-based techniques: Note that a former image colouring from the same authors was published as a short paper [PMGC20], but the later used a time-consuming sketch segmentation algorithm, which also brought restrictions on colouring freedom, since only pre-segmented regions could be coloured (e.g. regions shown in Figure 2b). Moreover, this early solution did not generate visually pleasing boundaries for coloured regions, and did not include any colour diffusion mechanism between neighbouring regions. In contrast, user-freedom, visual quality and efficiency are all improved by our new interactive method, which comes as well with an automatic companion method. The latter can be used as an auxiliary or corrective segmentation tool for medical images or freehand sketches, as shown by our preliminary results. In addition, we reused the first step shared by our two methods, the so-called

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

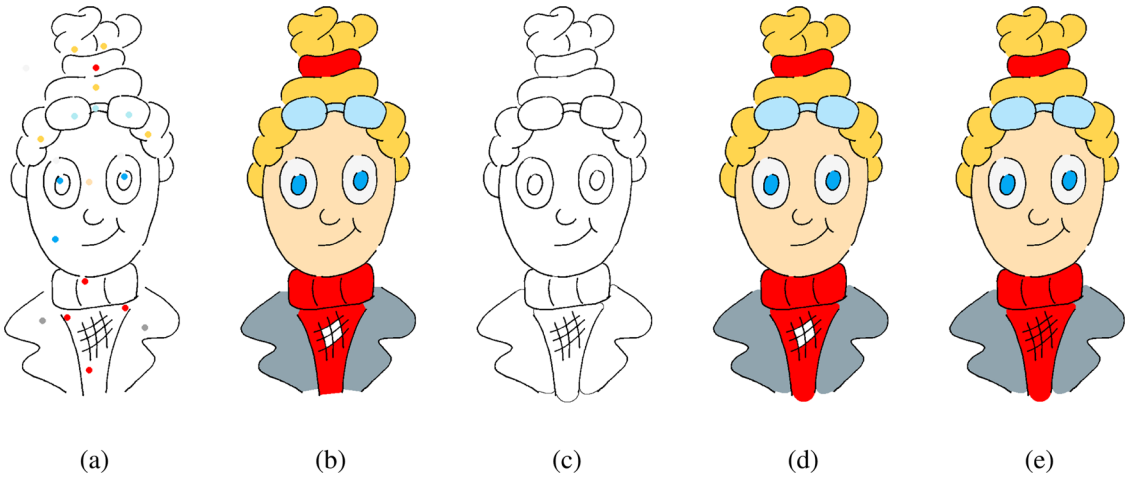

Figure 4: *Overview of our method. (a) Input contours (sketch) with colour hints, (b) result of Delaunay grouping, (c) aesthetically completed curves, (d) result after aesthetic curve completion, (e) result after colour diffusion.*

Delaunay grouping technique detailed in Section 4, in a recent paper aimed at the structuring and vectorization of rough stroke sketches (a quite different purpose compared to colouring, see also Favreau *et al*. [FLB16]), with a focus on user experience [PCS21]. Lastly, let us mention that the colouring flow we introduce is fundamentally different from the flow defined in Dey *et al*. [DGG03], which is based on Delaunay balls radii and designed for shape segmentation.

3. Overall Framework

We aim to propose a flat colouring method for contour images, where contours may have gaps. We will first present an interactive method and then an automatic companion method that we illustrate with a few basic examples in the context of medical image and free-hand sketch segmentation.

For the main interactive method, our goal is to make the interaction process as easy for the user as using the flood-fill tool, while robustly handling non-closed contours. In particular, the method should work for colouring minimalist sketches with open strokes, such as those used in several recent animation movies. In contrast with previous work, our solution relies on computational geometry.

The entire process starts with the user providing a raster, contour drawing with a set of unbounded regions that has to be bounded by colouring or segmenting it (throughout the remaining of this paper, we will be calling these regions *Bounded regions*). Though throughout this paper we used raster images (mostly hand-drawn) as inputs, one can also use vector images along with a curve sampler. Initially, the pixels that are part of the contours are extracted, and their positions are used to compute a Delaunay triangulation. Figure 3 shows a sample contour along with the computed Delaunay triangulation. As the name suggests, our method paints these Delaunay triangles as per the requirement. To facilitate colouring the background, there should be some triangles present in the background as well. To do that, we pad four points on the input image corners, which ensures that the entire image is triangulated. Figure 3c shows the updated Delaunay triangulation after inserting the padding points.

The initial Delaunay triangulation is further used for Delaunay grouping [PCS21] (Section 4). In Delaunay grouping, using a set of hints, the bounded regions are coloured iteratively. The input to this step is a set of contours along with colour hints (which are either interactively defined by the user or computed automatically) representing the bounded regions (as in Figure 4a), and the result is a set of flat coloured bounded regions (as in Figure 4b).

After Delaunay grouping, the bounded regions' boundaries will be represented by a set of straight lines (Delaunay edges)—which are not convenient for some applications, such as colouring line-arts. The user can thus make use of an aesthetic curve completion step (Section 5), which replace straight Delaunay edges (were they connecting the gaps between the contours or adjacent regions not separated by a contour) with visually pleasing curves, using the notion of SIMVC and sharp-corner detection heuristics. Figure 4d shows the effect of aesthetic curve completion. The input is a Delaunay grouping result, and the output is a set of bounded regions with visually pleasing finishing applied on the boundaries (the aesthetically completed curves can be seen in 4c).

To make interactive colouring easier, we also provide the user with an optional, colour diffusion framework (Section 6). From only a few colour hints, the latter helps to spread the colours into uncoloured small regions delimited by hatches or textures, avoiding the burden of individually colouring them. The input is the result after Delaunay grouping (as in Figure 4d), and the output is a full colouring, where colours have been spread to all the uncoloured regions (Figure 4e). The complete pipeline of our colouring framework is shown in Figure 5.

Since the position of colour hints directly impacts the grouping of Delaunay triangles (Section 4), such hints have to be automatically placed at relevant locations to achieve automatic segmentation of an input image. In Section 7, we discuss a strategy to generate such colour hints and automatically generate a meaningful segmentation.

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

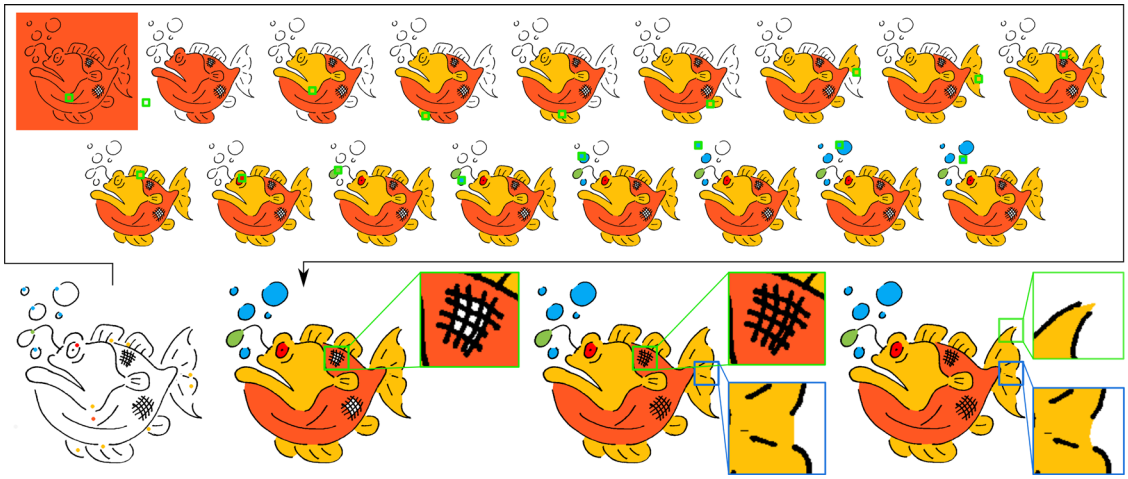

Figure 5: *A complete pipeline of our colouring framework on a complex sketch. From left to right: input sketch with colour hints, result of interactive Delaunay grouping (with steps shown inside the box), result after colour diffusion, final result after adding aesthetic contours.*

Given a set of colour hints and the straightforward pre-process that we have explained, let us now describe two key steps of our method, namely Delaunay grouping and aesthetic curve completion.

4. Delaunay Grouping

The Delaunay grouping process described in this section is based on an intuitive flow mechanism that we have already introduced in the short version of this paper [PMGC20] and further used in Parakkat *et al*. [PCS21], as a part of a sketch simplification algorithm.

A region adjacency graph G is initially computed, in which each Delaunay triangle (seen as a region in the input) is denoted as a vertex, and an edge is created between two vertices if they share a Delaunay edge, which is not part of the input contour (i.e. no edge is created between regions separated by a contour stroke). As we will see, this construction will prevent the merging of regions separated by a closed contour. A **weight** w is assigned to each edge of the graph G, and is set to its Euclidean length. In addition, a ($Color_Strength()$) parameter is associated to each graph vertex to represent the amount of colour already assigned to the corresponding region. This parameter is initially set to zero.

Once this graph is created, regions are iteratively merged based on the successive colour hints given by the user in the form of mouse clicks, as follows.

Once a user picks a colour and clicks (filling colour) at a specific position, the region R and its corresponding graph vertex v (denoted as $vertex(R)$) are identified from the location of the mouse click. All the regions in the graph that are reachable from region R are then filled with the selected colour. Note that this process is initially similar to the *bucket fill* tool (since all vertices inside a closed boundary do have a path between them).

After filling this colour, the $Color_Strength(u)$ of all vertices u, which are reachable from vertex v, is updated to the value of the flow that can reach u when coloured at v, denoted by $Edge_Flow(u, v)$,

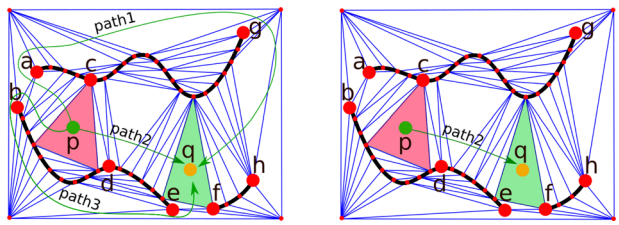

Figure 6: *Priority based path selection*

constrained by the the length of the shortest transition edge along the way—i.e. the smallest weight along a graph path, defined as:

$$\begin{aligned} Edge_Flow(u, v) &= max(f(X) : \forall\, paths\, X \;\; from\; u\; to\; v) \\ f(X) &= min(Weight(u, v) : \forall\, (u, v) \in X) \end{aligned} \tag{1}$$

The user then iteratively picks different colours and clicks on a chosen position in the contour (as in *bucket filling*). Based on the region R the user selected, the colour is recursively spread to the neighbouring regions R_i, but only if the $Color_Strength(vertex(R_i))$ is smaller than or equal to the $Edge_Flow(vertex(R_i), vertex(R))$.

We use a priority queue to ensure that the colour spreads through the largest gap first. For example, as shown in Figure 6, once a colour is applied on a point p, it can spread to another point q through various paths (path1, path2 and path3). Flow-through each of these paths is restricted by the shortest edges in their path, in this case, $|\, ab\, |$, $|\, cd\, |$ and $|\, ef\, |$ for path1, path2 and path3, respectively. Usage of priority queue ensures that the colour spreads from p to q through path2 (since that is the path having the maximum flow value). The priority queue is maintained as follows.

When a user clicks on a region R, the region is filled with the user-selected colour, the $Color_Strength(vertex(R))$ is updated to ∞ and each neighbour u of $vertex(R)$ is inserted into the priority queue with $Edge_Flow(u, vertex(R))$ as priority. After that, vertices v are

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

Algorithm 1. Filling algorithm

```
procedure FILLCOLOR(G,R,C)
    Fill(R,C)
    Initialize currentflow of each vertex to zero
    for each adjacent vertex v of the (w ← Vertex(R)) do
        InsertPQ(v,Weight(v,w))
    while PQIsNotEmpty() do
        B ← RemovePQ()
        v ← B.element(), len ← B.priority()
        if currentflow(v)<=len then
            currentflow(v) ← len
            Fill(Region(v),C)
            for each adjacent vertex u of v do
                if currentflow(u)<=Weight(u,v) then
                    if len>Weight(u,v) then
                        InsertPQ(u,Weight(u,v))
                    else InsertPQ(u,len)
```

iteratively taken from the priority queue and if $Color_Strength(v)$ is smaller than $Edge_Flow(v, vertex(R))$, then the region corresponding to v is coloured accordingly. In particular, the $Color_Strength(v)$ is updated to $Edge_Flow(v, vertex(R))$ and all neighbours of v are inserted to the queue. This procedure runs until the priority queue is empty. The overall procedure is summarized in Algorithm 1.

Figure 7 shows various steps in the iterative colouring procedure. In the first row, from left to right, the results of iterative Delaunay colouring are shown after giving each colour hint. The colour hints, along with the colour maps showing the influence of each mouse click on the underlying Delaunay triangles, are shown in the second row. The impact at a triangle T_i after clicking at T_j is computed as $\frac{Color_Strength(Region(T_i))}{||Longest_Edge(T_j)||}$ (where $Longest_Edge(T_j)$ is the longest edge in the triangle T_j). The largest to least influence is coloured from a spectrum varying from red to yellow, and if there is no influence, then it is coloured in white.

Since each triangle is given independent control, our algorithm is order-independent: indeed, whatever the order in which colours were applied, we end up in the same final result, thanks to the graph flow mechanism, which models how much a transition edge in the contour is likely to be suppressed.

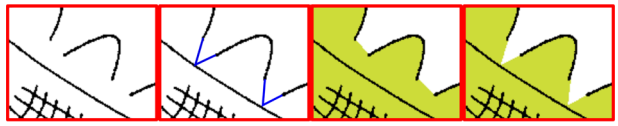

(a) (b) (c) (d)

Figure 8: *Sharp corners: (a) input sketch, (b) expected result, (c) result of Delaunay colouring, (d) result after gap filling procedure.*

(a) (b) (c) (d)

Figure 9: *(a) Input sketch, (b) expected result, (c) result of Delaunay colouring, (d) result after our gap filling procedure.*

5. Aesthetic Curve Completion

Simply applying Delaunay grouping to the Delaunay triangulation of the input separates the resulting coloured regions by straight edges. Using straight lines to connect gaps between curved contours is, however, not perceptually appropriate. Our aesthetic curve completion technique identifies the Delaunay edges connecting contour gaps and if needed, replaces them with appropriate curves, i.e. either sharp corners or smooth cubic curves (see Figures 8 and 9).

Our algorithm starts by evaluating the tangents at the endpoints of each contour gap to be filled. To achieve a robust computation of tangents on raster data, use a one-pixel width skeleton of the input sketch, as follows:

- Once the Delaunay grouping is over, the interior and exterior of the required contour boundary get different colours. The Delaunay edges connecting endpoints of the gap therefore share

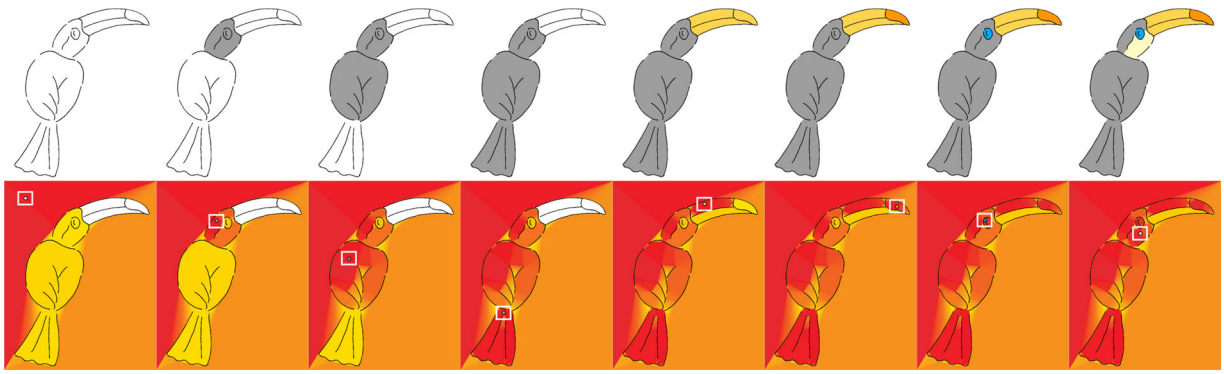

Figure 7: *First row: intermediate results of our Delaunay painting, Second row: cumulative influence of mouse clicks (shown inside white squares.)*

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

two triangles of different colours. We call such edges 'transition edges'.

- For each endpoint of a transition edge, the nearest point p_1 in the one-pixel skeleton of the input contour sketch is identified.
- A flood-fill procedure is initiated from p_1 within the contour, and stopped when we reach a pixel p_2 lying at a given distance (set to 5 pixels in our experiments).
- The contour's tangent at p_1 is set to the vector joining p_2 to p_1.

Once tangents are identified, we conduct three tests (discussed next) to check whether the gap should be completed using a sharp corner, or a smooth curve. If any of the sharp corner test fails, then a smooth cubic curve is used to fill the gap.

5.1. Sharp corners

In this section, we propose a method to identify and handle the boundaries that need to be connected using sharp corners (see Figure 8). To do so, let e_{AB} be the edge to be replaced and T_A and T_B be the corresponding tangents, respectively. We check the following three constraints for a gap to be qualified for getting replaced by a sharp corner:

- Angle constraint, based on the angle between the tangent: The constraint is said to be qualified if the angle between the tangents is less than $\pi/3$;
- Perpendicular constraint, based on the distance from the edge to be replaced, to the point where the tangent directions intersect: The qualifier is satisfied if the perpendicular distance from the intersection point p_{ip} to the edge e_{AB} is less than $2* \mid\mid AB \mid\mid$ (which prevents adding long an thin spikes to the contours);
- Linearity constraint, based on how linear the contours are near the gap. The criterion is said to be qualified if the pixels lying close to the endpoints are linearly arranged (we check whether the maximum distance from the pixels lying near the contour endings to the line defined by the tangent is less than a given threshold).

If these three constraints are qualified (a criteria which can also be tuned), we identify it as a sharp corner and replace e_{AB} by $e_{Ap_{ip}}$ and $e_{p_{ip}B}$.

5.2. Minimum variation curves

If any of the constraints above is not satisfied, we use a smooth cubic curve to connect the gap (see Figure 9). To achieve a perceptually pleasing result, we use SIMVC, which tend to minimize changes of curvature [Mor92]. As in Entem *et al.* [EPB*19], we define the SIMVC curve based on the tangents at the endpoints of the contours to be connected.

Let us T_A and T_B be the two tangents at the end-point of a (initially straight) transition edge. To achieve our objective of generating a perceptually pleasing curve between them, we connect A and B with a cubic Bezier curve defined by four control points (A, P_1, P_2, B), where $P_1 = A + c_1 T_A$ and $P_2 = B + c_2 T_B$ (P_1 and P_2 are the control points in the tangent directions from A and B, respectively). We optimize the free parameters c_1 and c_2 for minimizing the SIMVC energy. The latter was initially introduced [Mor92] as:

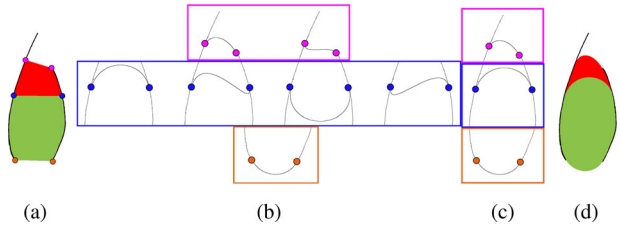

Figure 10: *(a) User input with endpoints circled over, (b) possible configurations of curves connecting the gaps, (c) selected curves, (d) final result.*

$$E_{SIMVC-Moreton} = (\int ds)^3 \int (\frac{d\kappa(s)}{ds})^2 ds \tag{2}$$

and later modified in Entem *et al.* [EPB*19] to create slightly shorter curves connecting the gaps as:

$$E_{SIMVC-Entem} = \frac{(\int ds)^5}{\|B-A\|^2} \int (\frac{d\kappa(s)}{ds})^2 ds \tag{3}$$

Where $C(s)$ is the parametric 2D curve with curvature $\kappa(s)$.

We made use of 7-15 Gauss–Kronrod quadrature rule [KMN89] for numerically solving the integration and used Gradient-descent for minimizing the energy function.

As shown in most of our examples (e.g. the wrist region of the hand at the right of Figure 1), the gaps that have to be filled along transition edges are not always between two ends of contour curves. Indeed, one (or both) of the contours may continue further, which gives two possible tangent orientations at the endpoint(s) of the transition edge (see Figure 10). For selecting the appropriate tangent, we initially identify all possible tangents at each endpoint for each transition edge. Then for each pair of tangents across the edge, we compute the MVC connecting them. Among all these pairs, the pair having minimum SIMVC energy is selected. Figure 10 shows a coloured sketch with endpoints marked as circles, all possible configurations of SIMVC, the selected curves with least SIMVC, and the final result, respectively.

In practice, to achieve interactive results, we restrict the search of the free parameters by a scale of 5 of the tangent length. We also skipped the SIMVC computation if the gap is too small (we quantify too small as gaps smaller than 8 pixels) since the difference will not be visible. Figure A.4 shows a few line-arts coloured using our tool.

6. Colour Diffusion

The input contours sometimes have many small regions (for example, hatches). Although we could use our tool to fill each of these regions one by one, this would be tedious and time consuming. We thus introduce a colour diffusion mechanism, enabling propagate colour from coloured areas to the uncoloured areas.

As shown in Figure 1, while giving additional information like hatching or stippling, the artists usually make use of thin brush

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

Algorithm 2. Colour diffusion

```
procedure DIFFUSECOLOR(IMAGE I)
    RFI = RandomFloodFill(I)
    BoundaryPixels = ExtractBoundaries(RFI)
    Compute Delaunay Triangulation DT(BoundaryPixels)
    (ColRegions, UncolRegions) = BipartateImage(RFI, I)
    DelaunayFilter(DT, RFI)
    while UncolRegions ≠ ϕ do
        e_ij = SmallestValid(Filtered DT(RFI))
        if v_i ∈ ColRegions and v_j ∈ UncolRegions then
            FloodFill(v_j), ColorAt(v_i))
            ColRegions = ColRegions ⋃ Region(v_j)
            UncolRegions = UncolRegions - Region(v_j)
        else
            if v_j ∈ ColRegions and v_i ∈ UncolRegions then
                FloodFill(v_i), ColorAt(v_j))
                ColRegions = ColRegions ⋃ Region(v_i)
                UncolRegions = UncolRegions - Region(v_i)
```

strokes (compared to the stroke size they used to draw the main sketch). The main intention behind this is to make the viewer clearly distinguish the sketched contours from the hatch strokes. Based on this observation, we introduce a simple, iterative colour diffusion framework in which, the colour from the best possible coloured region is iteratively diffused to an adjacent, uncoloured region. To quantify this 'best possible coloured region', we define a diffusion strength parameter between pairs of adjacent regions R_i and R_j, denoting how much strength a colour needs, to diffuse between them. It is defined as the shortest possible distance between any two pixels (p_i, p_j) such that $p_i \in R_i$ and $p_j \in R_j$.

The colour is then recursively spread from coloured region to the uncoloured regions in the order of increasing diffusion strength. First, the regions are classified into two groups based on whether the user coloured them or not. The diffusion strength from the uncoloured regions to the coloured regions are computed. The tuple having the least strength is used to transfer the colour. The regions (coloured and uncoloured) are updated, and the procedure is repeated until all regions are coloured.

Though the procedure is simple and straightforward, identifying the regions with minimum diffusion strength in a brute force way is computationally expensive. To make it fast, we again make use of the underlying Delaunay structure. The complete colour diffusion process, summarized in Algorithm 2, proceeds as follows.

We start with a random flood-fill algorithm to decompose the contour into different regions. Since the pixels contributing to the minimal diffusion strength can only (trivially) lies on the boundaries, we extract the boundary of each decomposed region. Then we compute the Delaunay triangulation of the extracted boundary pixels. Based on the user colouring, the regions are classified into *ColRegions* (regions which user coloured) and *UncolRegions* (where the colour could not reach yet). Since the Delaunay edges connecting pixels inside a region are of no relevance and only results in more computation time, we apply a Delaunay filtering to filter out all such Delaunay edges (by removing Delaunay edges connecting pixels from the same region). Among all unfiltered Delaunay edges, the valid Delaunay edge with the smallest edge length is identified (where a Delaunay edge becomes valid only if it is between one coloured and one uncoloured region). Once a valid edge is identified, the colour is diffused from the coloured region to the uncoloured region. Since the uncoloured region under consideration is now coloured, we update the regions appropriately. The procedure is continued until all the uncoloured regions get a colour through diffusion. Figure 11 shows a toy example demonstrating our colour diffusion procedure. It has to be noted that the method is not limited to hatched regions, as illustrated on Figure 11.

Since the diffusion process ultimately depends on the stroke size between the regions, sometimes, the result might not match the user expectation. Fortunately, thanks to our interactive colour filling and independent colour diffusion, the user can provide extra hints to recolour the regions. Figure 12 shows such an example, after giving a set of colour hints (Figure 12a), using the Delaunay grouping, the sketch is coloured. Since a few wool chunks (represented using circular strokes) are closed, Delaunay grouping will leave those regions uncoloured (Figure 12b). To colour these uncoloured regions, we applied colour diffusion. Unfortunately, as shown in Figure 12c, a couple of chunks on the head mistakenly got the colours diffused from the background. The user then provides extra colour hints (as in Figure 12d) to get a perfectly coloured region. Figure 12e shows the final result of our framework.

A few results of our painting algorithm can be seen in Figures A.5 and A.6.

7. Automatic Colouring Algorithm

Motivated by applications such as automatic colouring and art line completion, we now introduce an automatic variant of the proposed colouring algorithm. Although the automatic prediction of colours that have to be filled inside a contour is nearly impossible since it heavily depends on the artistic imagination, our automatic colouring results can be further adjusted by the artist. Alternatively, the method can be used as a lightweight tool to segment contours of a raster image into plausible regions.

To simplify computations and avoid any unnecessary complications due to stroke width, the automatic colouring starts by thinning the contours (i.e. the black regions in the input image) to make it single-pixel width. After thinning, potential endpoints that can contribute to gaps to be filled are found by identifying the open contours (using the neighbourhood information available at each pixel). As in Delaunay grouping, our main objective for 'Automatic colouring' is to segment the region by appropriately connecting these endpoints with proper Delaunay edges. So, next, we compute the Delaunay triangulation as in Section 4 but using the image with thinned contours as the input. Assuming the contour is Delaunay confirming, the next step is to pair the endpoints. To do so, we pick all Delaunay edges connecting two endpoints. If there are multiple such edges, the smallest edge connecting this pair of vertices is preserved. Depending on the contour, not all endpoints can be paired; and such endpoints might have to be connected to a non-endpoint. To address such cases, if an endpoint is not paired with other endpoints, adapting the idea from NN-Crust [DK99], we choose the shortest adjacent Delaunay edge, which makes an angle $> \pi/2$ with the contour.

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

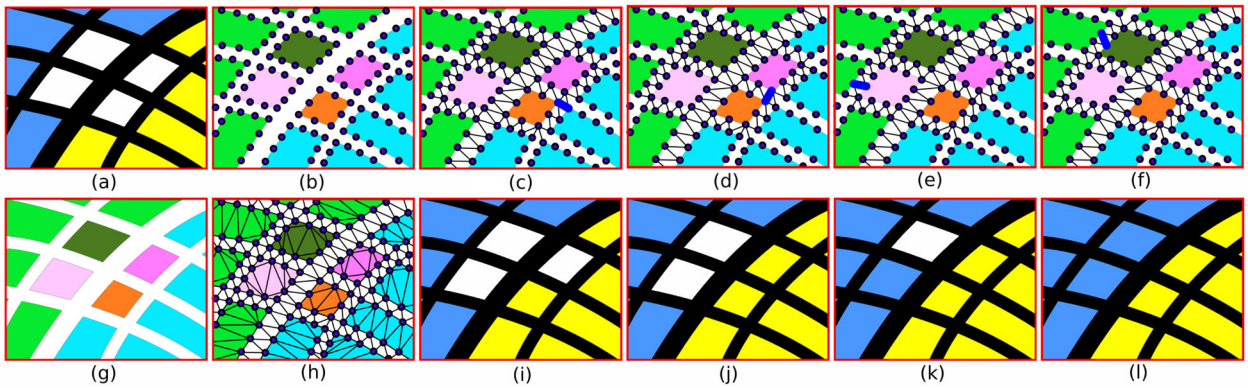

Figure 11: *Representative workflow of colour diffusion in a hashed region: top row: (a) a result after Delaunay grouping with four uncoloured regions (shown in white colour), (b) boundary samples (only a few of them are shown for a clearer illustration), (c–f) filtered Delaunay edges with smallest valid edge shown in blue colour. Bottom row: (g) result of RandomFloodFill() on (a); (h) Delaunay triangulation of (b), (i–l) Respective updated coloured regions from (c–f).*

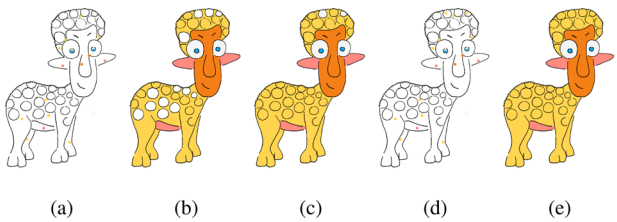

Figure 12: *(a) Input sketch with colour hints, (b) result of Delaunay grouping, (c) result after colour diffusion, (d) adding extra colour hints to the existing hints, (e) our final result.*

Once the right Delaunay edges connecting the endpoints (let them be called segmenting edges) are identified, we use it to split the regions. In other words, we ensure that the triangles (Tri_A and Tri_B) sharing the segmenting edges are labelled differently in the automatic colouring procedure. To facilitate the labelling, depending on the current labels of Tri_A and Tri_B, two situations can arise: In the first one, both Tri_A and Tri_B can be unlabelled, in which we can initiate Delaunay grouping from Tri_A and then Tri_B, ensuring them to have different labels. Another possibility is for Tri_A and Tri_B to have the same labels. In this case, initiating Delaunay grouping from both the triangles will over segment the regions, and randomly initiating Delaunay grouping from one of the triangles might not always assign different labels to both the triangles. To tackle this, we keep track of the flow direction during Delaunay grouping, i.e. the triangle from which the current triangle is getting coloured, and do the following:

- $FLOW(Tri_A) = Tri_B$—means the colour flows from Tri_B to Tri_A. In this case, we can safely initiate the Delaunay colouring with new colour from Tri_A.
- $FLOW(Tri_B) = Tri_A$—case under which the colour flows from Tri_A to Tri_B. In which, the colouring can be initiated from Tri_B.
- $FLOW(Tri_A) \neq Tri_B$ and $FLOW(Tri_A) \neq Tri_B$—special case in which there is a larger path from Tri_A to Tri_B without crossing the segmenting edge under consideration. In this case, we can arbitrarily pick Tri_A or Tri_B and start colouring. The intermediate previously coloured triangle in the path will prevent the colour flowing from Tri_A to Tri_B and vice versa.

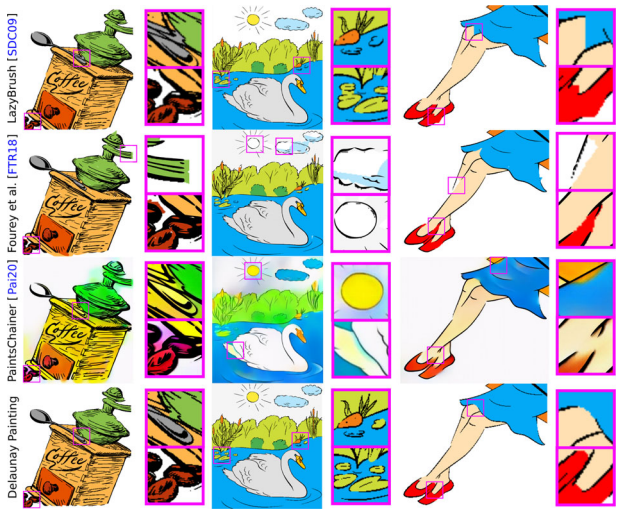

Figure 13: *Various steps in the automatic segmentation algorithm. The order is shown using yellow colour arrows, segmenting edges are shown in red colour, the edge under consideration is marked using orange colour circle.*

The algorithm automatically picks and groups/labels the triangles appropriately in the decreasing order of the edge length to result in the required segmented contour. An input sample contour and the various steps in our automatic segmentation are shown in Figure 13.

7.1. Applications to automatic image segmentation

Our automatic colouring method can be particularly relevant for fast image segmentation. One such, exciting application is medical

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

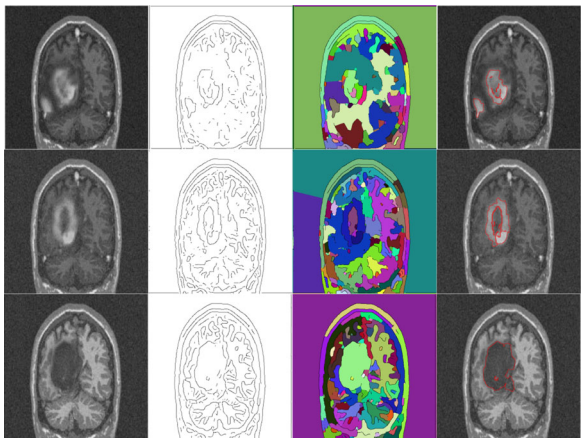

Figure 14: *Tumour labelling (left to right: Input MRI, result after thresholding and inverting Canny edges, result of automatic colouring, labelled MRIs using our method). [Image source: [PBG09].*

image segmentation, which is a very well-studied and challenging area. While efficient, advanced tools have been specifically developed for segmentation tasks, the question of properly segmenting pathological cases is still open. For instance, the medical imaging community has a special interest in identifying tumours in a medical image, especially from MRI scans [PBG09]. In particular to better guide the 3D segmentation input of many recent learning-based methods, manual labelling in 2D is still used for training data. We made a simple test for this application scenario, as a first attempt to alleviate this user-intensive task. Since our algorithm can only deal with contour images, we pre-processed the input MRI grey-scale images using a Canny edge detector, followed by binary thresholding and then image inversion. Figure 14 shows a few MRI scans and various steps in the automatic labelling procedure. These preliminary while promising results show that our technique may potentially be used either for an initial segmentation process, or as a complementary tool to adjust and correct some resulting segmentation, for which the expert could click in the interest regions and highlight the desired regions appropriately. In summary, we hope that some ideas of our Delaunay painting method can inspire further work in automatic or interactive pathological 2D images labelling, and enrich or simplify this time-consuming process.

Another interesting use-case for this automatic segmentation algorithm is the segmentation of freehand sketches into possibly overlapping regions. The problem itself is a non-trivial one [ST16] because of the presence of significant gaps, absence of strokes and boundaries represented by multiple strokes. Luckily, our segmentation algorithm can fill large gaps, group regions bounded by various strokes, and group strokes shared between various regions. To get a meaningful result by appropriately resolving ambiguities and occlusion, the automatically generated segments are manually picked and merged. During this manual picking and merging, the user was also given the flexibility to split the regions (if needed). Figure 15 shows the segmentation generated in this way on different challenging cases.

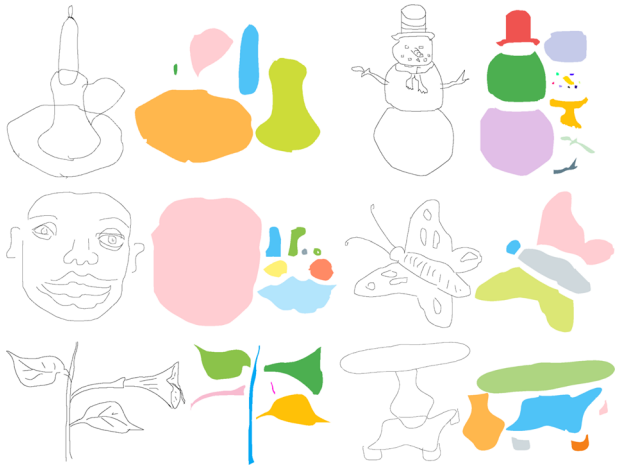

Figure 15: *Few sketches (from [ST16]), segmented using our method and merged manually.*

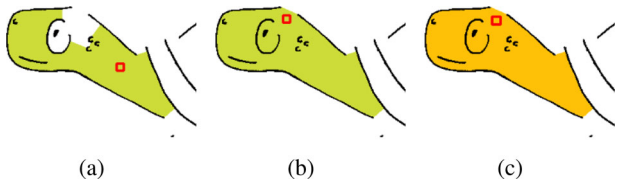

Figure 16: *A sketch coloured using same colours in different orders (the colour hint number and the final result are shown in the right).*

8. Results and Discussion

Implementation: The method was implemented in C++, using the CGAL and OpenCV libraries: we used CGAL for Delaunay triangulation and its Delaunay data-structure for the efficient implementation of Delaunay painting. Whereas, image processing and interaction-related functions were implemented using OpenCV. Zhang-Suen's algorithm [ZS84] was used for computing the skeleton of contour images.

Qualitative results: Though we explored a solution for automatic colourization, which has many applications on its own, the interactive version of our method is the most relevant for artistic colouring, where the artists needs to be in the loop to pick and drop the appropriate colours (irrespective of whether the underlying contour is segmented or not). Figures A.4–A.6 show that artists are able to generate visually pleasing results.

Moreover, a demo of our system is given in the companion video.

Time: Our tool is light-weight and responds at an interactive rate. The experimentation was conducted on a Core i3 processor @ 1.70 GHz speed and 4GB RAM. In average, the Delaunay computation, Delaunay grouping, colour diffusion and aesthetic curve completion took around 0.1, 0.1, 0.5 and 88 s, respectively. The time required for colour filling and colour diffusion is negligible compared to the gap-filling procedure, which fortunately needs to be done only once at the end of the colouring procedure. Colouring order: Since we give independent control to each triangle, we can attain the expected

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

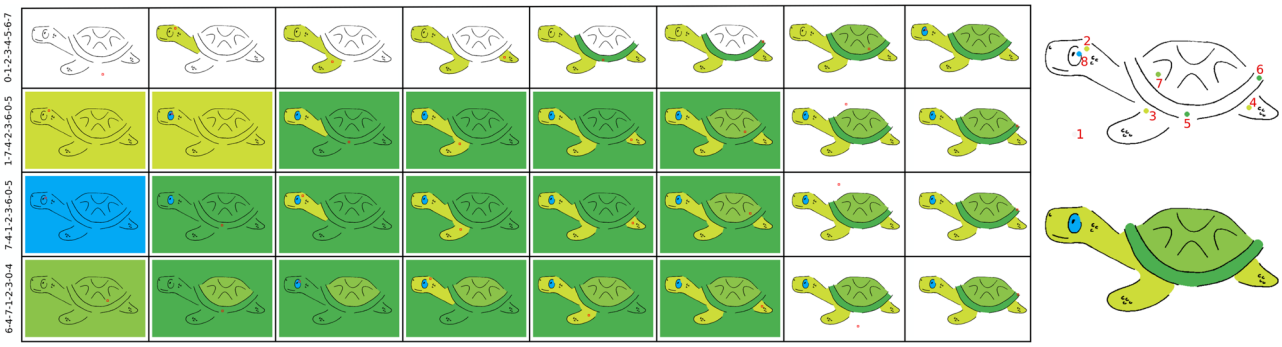

Figure 17: *Influence of colouring near the prominent gap. Case 1: Giving colour hints in two steps as in (a) and then (b), Case 2: Giving colour hint near the prominent gap in the boundary required only one input as in (c).*

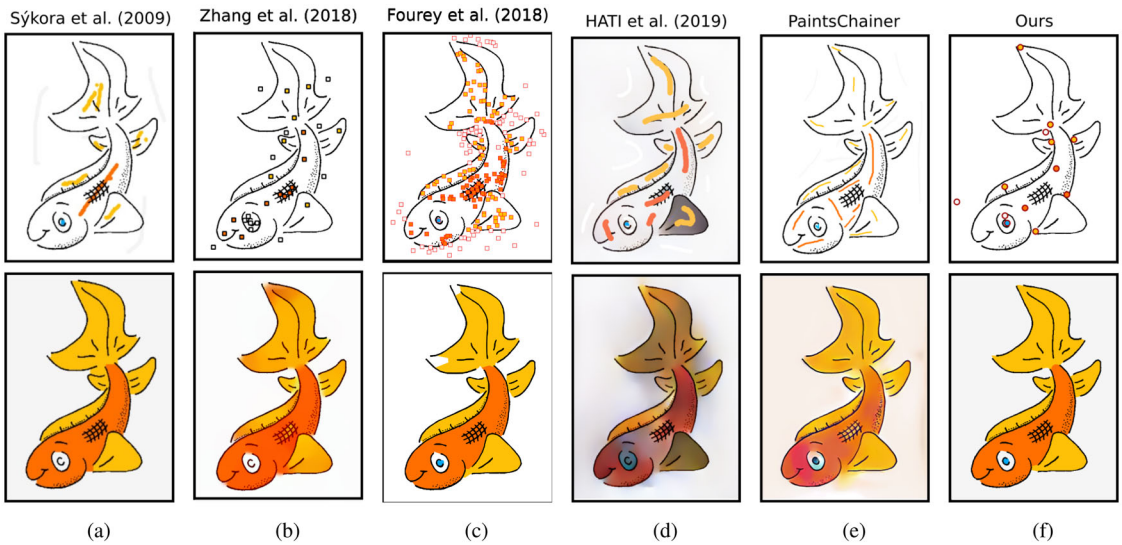

Figure 18: *Delaunay painting versus related work. First row shows the input given to respective algorithms and second row shows the resulting coloured sketches.*

result irrespective of the colouring order. Figure 16 shows a sample sketch coloured in several different orders. The user might need to give more hints depending on the order chosen, but thanks to our simple flow mechanism, the expected result can always be achieved.

Tips and tricks: The results and the time taken demonstrate that our tool is straightforward to use. Still, we can make colouring tasks easier by using a couple of tricks. The first one is to pre-fill the most significant region at the border of the image by a specific background colour.

This avoids the unwanted spreading of other colours to the background, hence reducing the need for recolouring and unnecessary confusions. Second, whenever filling a region with a colour, we encourage the user to click near the largest, unwanted gap in the boundary.

This indeed will help the system to pick the locally largest Delaunay triangle and hence will prevent the need of clicking again on some other part to fill the same colour. Figure 17 shows an example of how the colouring near a prominent gap influenced the final result (Note that, depending on the gap size and on the number of gaps around a region, the user sometimes has to give multiple hints to satisfactorily colour a region).

8.1. Comparative colouring scenarios and results analysis

Let us now compare our colouring framework with related existing works through an example illustrated in Figure 18. First, Figure 18a shows the result of Lazybrush [SDC09] for a sample set of contours. As can be seen in this result and as discussed in their paper, the boundaries are not good-looking in an aesthetic sense. Also, colouring a drawing using a LazyBrush is not necessarily straightforward for a novice user, since the colouring highly depends on the stroke size and strength (and hence is more appropriate for a professional user). The user may have to repaint the same region multiple times to get the desired result. Figures 18(b), 18(d) and 18(e) show the result of colouring using Zhang *et al.* [ZLW*18], PaintsChainer [Pai20] and Hati *et al.* [HJRD19] methods. As seen in

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

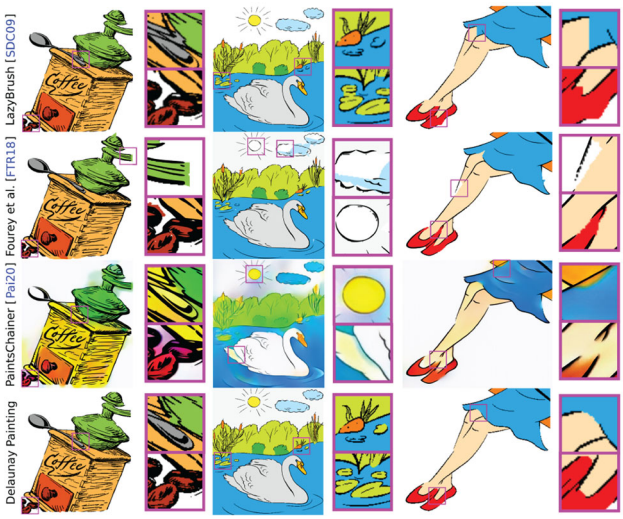

Figure 19: *Visual comparisons (differences are highlighted).*

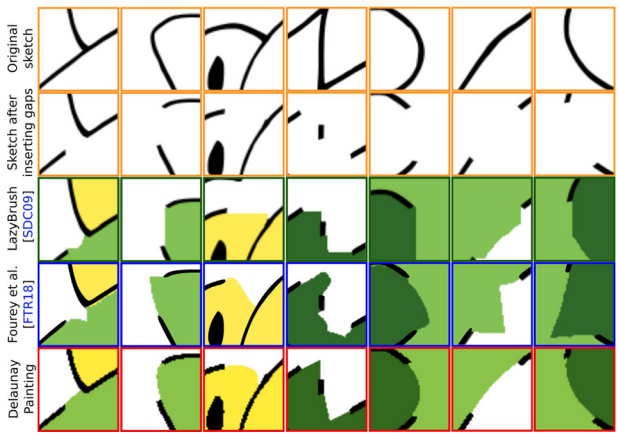

Figure 20: *Curve completion by various algorithms (magnified version of Figure 21). From left to right: increasing gap size.*

the figures, these methods result in relevant and satisfactory colouring. However, learning-based approaches, following relatively time-consuming process, are highly dependent on the training data and cannot work for random artistic drawings (cannot be learned in advance since it depends on artist's imagination). Lastly, Figure 18c shows a sample contour coloured using the line art colouring tool in GIMP (an implementation of Fourey *et al*. [FTR18]). As can be seen, the contour is not coloured properly and multiple user inputs are required to paint even a small area, which might perceptually look like a single region. Also note that the colouring of this sketch required 203 mouse clicks, which is a hectic task for the user. Figure 19 shows further visual comparisons between our solution and previous ones, on a few complex examples, showing that the quality of results can often be improved using our method.

A magnified visual comparison of the curve completion procedure for filling gaps used by various algorithms is provided in Figure 20 , while Figure 21 shows the full images. Note that except

Figure 21: *Curve completion experimentation. From left to right: Input sketch, altered sketch after manually adding gaps of various size, results of colouring using Sykora* et al. *[SDC09], Fourey* et al. *[FTR18] and our tool.*

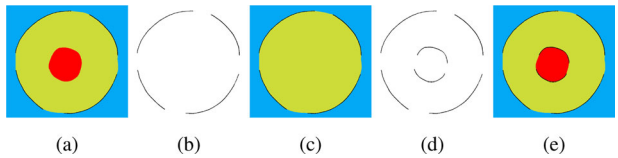

Figure 22: *(a) Expected result, (b) input sketch, (c) result generated using our method, (d) required input sample, (e) result of colouring on required sketch.*

for small gaps (on which contour beautification was ignored to save time), our method was able to generate visually pleasing gap-filling curves, close to the expected ones.

To still go further in this evaluation, we conducted a user study (with a style inspired from Zhang *et al*. [ZLW*18]) that we report in Appendix.

8.2. Limitations

Though the proposed method is easy to use and has many advantages, it also has a few limitations. The first limitation is that, as the name suggests, our method can only be used for flat-colouring and is not designed for applying shading. Shading could, however, be easily applied on top of our flat-colouring results using existing methods such as Johnston [Joh02], Sykora *et al*. [SKv*14] or Hudon *et al*. [HGP*19].

Also, we expect a few strokes to be present to define the boundary that has to be coloured. For example, for colouring a contour as in Figure 22a, the input sketch shown in Figure 22b is not enough since the inner shape does not have any contours. But, once we add a few sketch strokes to define the inner border as in Figure 22d, our method can achieve the expected result.

Failure case: Sometimes, the Delaunay edge connecting a gap in a contour may be missing, and thus this contour cannot be completed. The problem occurs when there is no empty circle containing both endpoints of a gap (see Figure 23). In other words, we expect the input contours to be 'Delaunay confirming'. Also, the heuristics used in our curve completion step are not guaranteed to always give a perceptually correct result, and are also prone to noise since they rely on pixel operations.

Lastly, our method only works for handling contour drawings and was not designed for colouring greyscale images, which are a failure case for us (contrary to the method in Iizuka *et al*. [ISSI16]).

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

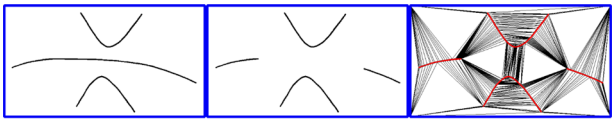

Figure 23: *Limitation of Delaunay-based colouring: (a) expected boundary, (b) input sketch, (c) underlying Delaunay triangulation.*

Figure 24: *Parakkat* et al. *[PCS21] analysis: from left to right: initial Colouring, interactively generated segments, result after merging regions.*

9. Conclusion

In this paper, we introduced a simple framework for colouring raster sketches with gaps in the contours. We first presented a novel Delaunay-triangulation-based colouring method, using an intuitive flow mechanism introduced in Parakkat *et al*. [PCS21] (where there is no notion of colours, but two options as shown in Figure 24 to interactively segment or merge regions, which are tailored for this specific application). Using the SIMVC concept and a few sharp corner heuristics, we introduced a method to connect gaps in the contour, enabling to give an aesthetic finish to the coloured sketch. Finally, a colour diffusion method was used to save user efforts, avoiding them to individually colour each region in areas with dense patterns such as hatches. With minimum effort, the artist/user can modify or adjust the result. Our interactive method is an easy-to-use tool from different aspects, such as comparatively minimal input, independent colour-ordering and quick response time. As our results show, our simple framework can be used to colour a wide variety of hand-drawn sketches in a visually pleasing way.

We have also presented an automatic companion method and applied it to efficient image segmentation. Our pipeline provides a handy tool for automatic pre-segmentation and fast interactive segmentation correction for 2D medical images, which is a well-studied while still challenging application.

In the future, we would like to strengthen the colour diffusion process to appropriate regions by analysing the contour globally, in combination with our current Delaunay-based local analysis. Lastly, we plan to investigate the temporal consistency of our method, by exploring the robust propagation of colour hints through the sketches corresponding to the different frames of an animation.

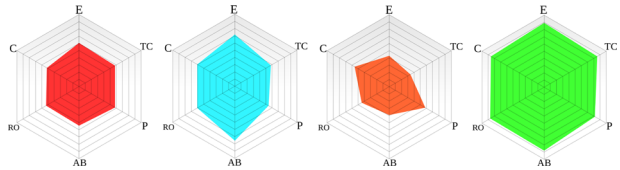

Figure A.1: *Visualization of our user study. From left to right: LazyBrush, PaintsChainer, GIMP and ours.*

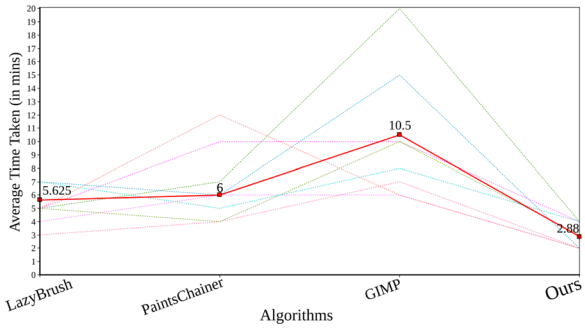

Figure A.2: *Time taken (average time is shown in red bold line) by various users to colour the sketches using different tools*

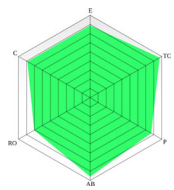

Figure A.3: *Second user study: Evaluation of our solution.*

Acknowledgement

The authors would like to thank Harihara Gowtham, Karan Singh and Prudhviraj Madipally (for their help with our previous related works - [PMGC20] and [PCS21]), Even Entem (for his help with the SIMVC implementation), Yliess Hati and Lvmin Zhang (for helping us with the comparison), the users (for participating in our study), and reviewers (for their valuable comments that helped us to improve the manuscript).

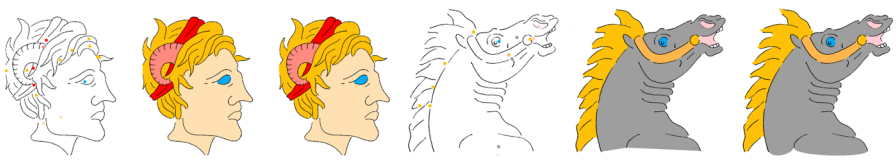

Figure A.4: *A few results of colouring given using our tool.*

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

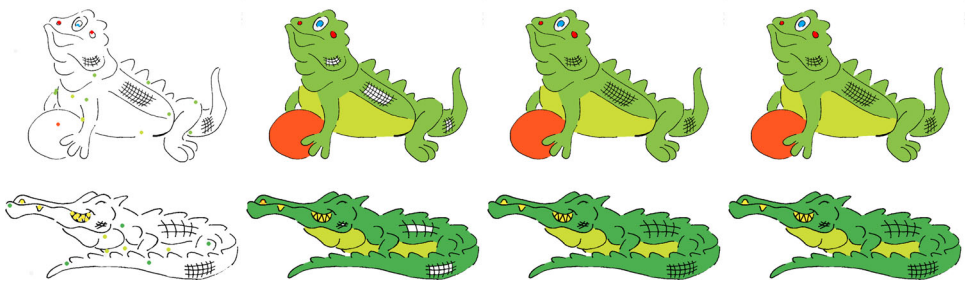

Figure A.5: *A few results of colouring given on simple sketches using our tool.*

Figure A.6: *Various complex sketches coloured using our method (original sketches taken from Pixabay).*

Appendix A: User Study

Inspired by Zhang *et al*. [ZLW*18], we conducted a user study to evaluate the strengths and weaknesses of the proposed system. We invited eight non-experts users, aged between 18 and 45 to participate in the interactive colouring session remotely. The users were given an initial tutorial on how to use the interface and then asked to randomly pick and colour one of the sketches shown in Figure A.6 using our tool and the other three openly available colouring tools (LazyBrush [SDC09], PaintsChainer [Pai20] and GIMP [FTR18]). It has to be noted that the users were explicitly directed to define the background colour first, to make the colouring process intuitive. At the end of the colouring process, the users were asked to participate in a multi-dimensional survey to grade each interface (on a scale of 0–10) based on the following aspects:

- Easiness (E)—How easy it is to give colouring
- Time consumption (TC)—How time-consuming the entire process is
- Predictability (P)—How easy it is to identify where colours should be given to get the expected results
- Aesthetic beauty (AB)—How good looking the final results are
- Region obedience (RO)—Whether the user was able to capture the expected boundary
- Control (C)—How easy it is to control the colour propagation

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

As it can be seen in the results of Figure A.1, our proposed method has largely outperformed the competitors from all these aspects. Let us report some common observations made by various users.

- LazyBrush: Difficult to predict and control the effect of brush size. It was easy to colour small regions and users did not have to worry about precise mouse location.
- PaintsChainer: The final result was having an artistic feeling (like in a painting). But colour bleeding along with random colours made the final result look bad. Also, this method was not able to capture the boundaries correctly.
- Fourey *et al*. (GIMP): Colouring small pieces makes it time-consuming. Difficult to control the colours near borders. Small uncoloured regions are left in some areas.

Users were also asked to comment on the advantages and disadvantages of our tool, and which resulted in the following remarks:

+ Easy to use without worrying about the gaps;
+ Less work since colours are spreading;
+ Captures the borders well;
- Editing/refining small regions or redoing colours is time consuming. A 'lazy' paint brush as in LazyBrush would have been nice;
- Should have included a better colour palette.

In general, though our tool currently does not support functionalities like undo/redo and does not have an advanced colour palette, our simple interface allowed the users to colour the sketches without worrying about contour gaps. Our method could not only capture the user's intended boundaries, but also simplify user interaction and provide aesthetically pleasing results. We also report the time taken by different users to colour the sketch using various tools in Figure A.2. It can be seen that users were able to quickly colour complex sketches using our tool. To get a feedback of how easy our interface is (without comparing to other methods), 10 other users aged between 13 and 50 were asked to evaluate our interface, by grading each attribute on a scale of 0–10. The results are shown in Figure A.3, with the following comments:

- Following the guideline, it is a convenient and light interface.
- Cannot add free-form shapes.
- Very practical as far as we start by background colour toward more detailed ones.
- A pre-visualization of colour change (before click) may be useful.

References

[CSG*17] Chen J., Shen Y., Gao J., Liu J., Liu X.: Language-based image editing with recurrent attentive models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2017). http://arxiv.org/abs/

[DGG03] Dey T. K., Giesen J., Goswami S.: Shape segmentation and matching with flow discretization. In *Algorithms and Data Structures*. F. Dehne, J. -R. Sack and M. Smid (Eds.). Springer, Berlin Heidelberg (2003), pp. 25–36.

[DK99] Dey T. K., Kumar P.: A simple provable algorithm for curve reconstruction. In *SODA'99: Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms* (USA, 1999), Society for Industrial and Applied Mathematics, pp. 893–894.

[EPB*19] Entem E., Parakkat A. D., Barthe L., Muthuganapathy R., Cani M.-P.: Automatic structuring of organic shapes from a single drawing. *Computers & Graphics 81* (2019), 125–139.

[FHOO17] Furusawa C., Hiroshiba K., Ogaki K., Odagiri Y.: Comicolorization: Semi-automatic manga colorization. In *SA'17: Proceedings of the SIGGRAPH Asia 2017 Technical Briefs* (New York, NY, USA, 2017), Association for Computing Machinery.

[FLB16] Favreau J.-D., Lafarge F., Bousseau A.: Fidelity vs. simplicity: A global approach to line drawing vectorization. *ACM Transactions on Graphics 35*, 4 (July 2016), 1–10.

[FTR18] Fourey S., Tschumperlé D., Revoy D.: A fast and efficient semi-guided algorithm for flat coloring line-arts. In *EG VMV'18: Proceedings of the Conference on Vision, Modeling, and Visualization* (Goslar, DEU, 2018), Eurographics Association, pp. 1–9.

[GTF94] Gangnet M., Thong J.-M. V., Fekete J.-D.: Automatic gap closing for freehand drawing. In *SIGGRAPH'94: Proceedings of the Technical Sketch Submission* (Dec. 1994).

[HA17] Hensman P., Aizawa K.: cgan-based manga colorization using a single training image. In *Proceedings of the 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)* (Los Alamitos, CA, USA, Nov. 2017), vol. *3*, IEEE Computer Society, pp. 72–77.

[HGP*19] Hudon M., Grogan M., Pagés R., Ondřej J., Smolić A.: 2DToonShade: A stroke based toon shading system. *Computers & Graphics: X 1* (2019), 100003.

[HJRD19] Hati Y., Jouet G., Rousseaux F., Duhart C.: Paintstorch: A user-guided anime line art colorization tool with double generator conditional adversarial network. In *CVMP'19: Proceedings of the European Conference on Visual Media Production* (New York, 2019), Association for Computing Machinery.

[ISSI16] Iizuka S., Simo-Serra E., Ishikawa H.: Let there be color!: Joint end-to-end learning of global and local image priors for automatic image colorization with simultaneous classification. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2016) 35*, 4 (2016).

[Joh02] Johnston S. F.: Lumo: Illumination for cel animation. In *NPAR'02: Proceedings of the 2nd International Symposium on Non-Photorealistic Animation and Rendering* (New York, 2002), Association for Computing Machinery, pp. 45.

[KJPY19] Kim H., Jhoo H. Y., Park E., Yoo S.: Tag2Pix: Line art colorization using text tag with SECat and changing loss (2019). http://arxiv.org/abs/1908.05840.

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd

[KMN89] Kahaner D., Moler C., Nash S.: *Numerical Methods and Software*. Prentice-Hall, Inc., USA, 1989.

[LDL*19] Liu F., Deng X., Lai Y.-K., Liu Y.-J., Ma C., Wang H.: Sketchgan: Joint sketch completion and recognition with generative adversarial network. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019), pp. 5823–5832.

[LWCO*07] Luan Q., Wen F., Cohen-Or D., Liang L., Xu Y.-Q., Shum H.-Y.: Natural image colorization. In *Rendering Techniques*. J. Kautz and S. Pattanaik (Eds.). The Eurographics Association (2007).

[Mor92] Moreton H.: *Minimum curvature variation curves, networks, and surfaces for fair free-form shape design*. PhD thesis, (1992) Berkeley, CA, USA.

[Pai20] PaintsChainer: Petalica paint, online demo (by preferred networks, inc.). https://petalica-paint.pixiv.dev/index_en.html. Accessed June 2020.

[PBG09] Prastawa M., Bullitt E., Gerig G.: Simulation of brain tumors in MR images for evaluation of segmentation efficacy. *Medical Image Analysis 13*, 2 (2009), 297–311. Includes Special Section on Functional Imaging and Modelling of the Heart.

[PCS21] Parakkat A. D., Cani M.-P., Singh K.: Color by numbers: Interactive structuring and vectorization of sketch imagery. In *Proceedings of the CHI'21*, Association for Computing Machinery (2021), pp. 1–11.

[PM16] Parakkat A. D., Muthuganapathy R.: Crawl through neighbors: A simple curve reconstruction algorithm. *Computer Graphics Forum 35*, 5 (2016), 177–186.

[PMGC20] Parakkat A. D., Madipally P., Gowtham H. H., Cani M.-P.: Interactive flat coloring of minimalist neat sketches. In *Eurographics 2020—Short Papers* (2020), A. Wilkie and F. Banterle (Eds.), The Eurographics Association.

[PMM18] Parakkat A. D., Methirumangalath S., Muthuganapathy R.: Peeling the longest: A simple generalized curve reconstruction algorithm. *Computers & Graphics 74* (2018), 191–201.

[SCH20] Su J.-W., Chu H.-K., Huang J.-B.: Instance-aware image colorization (2020). http://arxiv.org/abs/2005.10825.

[SDC09] Sýkora D., Dingliana J., Collins S.: Lazybrush: Flexible painting tool for hand-drawn cartoons. *Computer Graphics Forum 28*, 2 (2009), 599–608.

[SISSI17] Sasaki K., Iizuka S., Simo-Serra E., Ishikawa H.: Joint gap detection and inpainting of line drawings. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), pp. 5768–5776.

[SKv*14] Sýkora D., Kavan L., Čadík M., Jamriška O., Jacobson A., Whited B., Simmons M., Sorkine-Hornung O.: Ink-and-ray: Bas-relief meshes for adding global illumination effects to hand-drawn characters. *ACM Transactions on Graphics 33*, 2 (Apr. 2014), pp. 1–15.

[ST16] Schneider R. G., Tuytelaars T.: Example-based sketch segmentation and labeling using CRFs. *ACM Transactions on Graphics 35*, 5 (July 2016), 1–9.

[SZC*20] Shi M., Zhang J.-Q., Chen S.-Y., Gao L., Lai Y.-K., Zhang F.-L.: Deep line art video colorization with a few references (2020). http://arxiv.org/abs/2003.10685.

[YBC*19] Yoo S., Bahng H., Chung S., Lee J., Chang J., Choo J.: Coloring with limited data: Few-shot colorization via memory-augmented networks, In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019).

[ZLSS*21] Zhang L., Li C., Simo-Serra E., Ji Y., Wong T.-T., Liu C.: User-guided line art flat filling with split filling mechanism. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021).

[ZLW*18] Zhang L., Li C., Wong T.-T., Ji Y., Liu C.: Two-stage sketch colorization. *ACM Transactions on Graphics 37*, 6 (Dec. 2018), 1–14.

[ZMG*19] Zou C., Mo H., Gao C., Du R., Fu H.: Language-based colorization of scene sketches. *ACM Transactions on Graphics 38*, 6 (Nov. 2019), 1–16.

[ZS84] Zhang T. Y., Suen C. Y.: A fast parallel algorithm for thinning digital patterns. *Communications of the ACM 27*, 3 (Mar. 1984), 236–239.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Data Video S1

© 2022 The Authors. *Computer Graphics Forum* published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd