Wehrli 2.0: An Algorithm for "Tidying up Art"

Nikolai Ufer[1,2], Mohamed Souiai[1], and Daniel Cremers[1]

[1] Department of Computer Science, Technical University of Munich,
Boltzmannstr. 3, D-85748 Garching, Germany
Mohamed.Souiai@in.tum.de, Cremers@in.tum.de
[2] Department of Mathematics, University of Munich,
Theresienstr. 39, D-8033 Munich, Germany
N.Ufer@campus.lmu.de

Abstract. *We propose an algorithm for automatizing the task of "Tidying up Art" introduced by the comedian Wehrli [1]. Driven by a strong sense of order and tidyness, Wehrli systematically dissects famous artworks into their constituents and rearranges them according to certain ordering principles. The proposed algorithmic solution to this problem builds up on a number of recent advances in image segmentation and grouping. It has two important advantages: Firstly, the computerized tidying up of art is substantially faster than manual labor requiring only a few seconds on state-of-the-art GPUs compared to many hours of manual labor. Secondly, the computed part decomposition and reordering is fully reproducible. In particular, the arrangement of parts is determined based on mathematically transparent criteria rather than the invariably subjective and irreproducible human sense of order.*

Key words: Tidying up Art, Image Segmentation, Label Cost Prior, Convex Relaxation, Convex Optimization, Fast Global K-Means

1 Introduction

1.1 Ursus Wehrli's Project of "Tidying up Art"

Starting in 2002, the Swiss comedian Ursus Wehrli developed the project of "Tidying up Art". Wehrli argues that while he likes art he is systematically disturbed by many of history's most famous artworks being highly unordered and chaotic. Wehrli tackles this shortcoming in the works of famous artists by dissecting their works and rearraging respective segments in a well ordered manner – see Figure 1 for an example. His work has become extremely popular and Wehrli's books have been best sellers for many years [1–3]. A closer look at Wehrli's works reveals that his approach has two important shortcomings:

- Manually dissecting a painting and rearranging all parts is an extremely tedious process which can easily take several hours of work. Tidying up the entire art history would take Wehrli years or even decades.

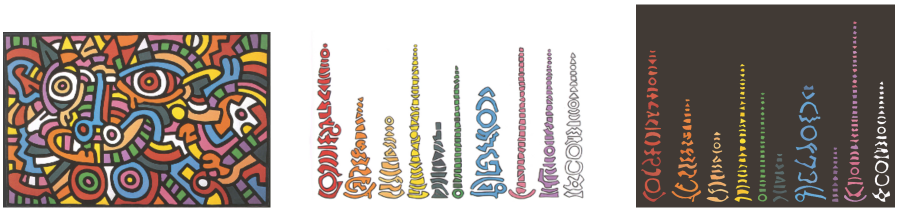

(a) Painting by Haring (b) Wehrli's manual tidying (c) Output of *Wehrli 2.0*

Fig. 1. Given an original artwork (a) *Wehrli 2.0* generates automatically a tidied up version of it (c). The algorithm reproduces Wehrli's notion of tidyness (b).

- The result of the above approach is hardly reproducible. The dissection of the artwork into parts is performed quite heuristically. Moreover, the subsequent ordering is based on a highly intransparent and irreproducible human notion of order.

The contribution of this paper is to introduce the algorithm *Wehrli 2.0* which is designed to alleviate the above shortcomings.

1.2 Related Work

Dissecting the image plane into its constituents is a problem of image segmentation and as such one of the most studied problems in image analysis. There is abundant literature on mathematical models for image segmentation, starting with the pioneering works of Mumford and Shah [4], Blake and Zisserman [5] and Kass et al. [6]. While the length regularization imposed in respective cost functions is desirable for meaningful segmentations, it gives rise to difficult optimization problems, the general multiregion segmentation being NP hard (in its spatially discrete formulation). Nevertheless, over recent years people have developed efficient algorithms for approximate minimization including the graph cut based alpha expansions [7] or various forms of convex relaxation [8–10]. In this work, we will make use of convex relaxation techniques because they do not exhibit any grid bias and are easily parallelized [11].

For a fully unsupervised partitioning, however, respective algorithms also need to estimate appropriate color models associated with each region. In practice, we observed that the commonly suggested alternating estimation of color models and segmentation is computationally demanding and likely to get stuck in suboptimal local solutions. For a robust unsupervised performance, it is therefore important to optimally estimate multiple color models.

1.3 Contribution

We propose an algorithm called *Wehrli 2.0* which aims at fully automatizing Wehrli's work in a manner that we can simply insert an artwork and the computer generates a tidied up version of it – see Figure 1 for an example.

To this end, we developed a fully unsupervised multi-region segmentation method which combines state-of-the-art convex relaxation techniques with fast global k-means color model estimation. Subsequently we propose ordering criteria to optimally rearrange all parts. In numerous experiments, we compare the performance of our algorithm with Wehrli's manual work. While the results are never entirely identical, these experiments show that:

- Our algorithm provides results which are qualitatively similar to those obtained by Wehrli and thus captures the essence of his work.
- The proposed computerized solution to Wehrli's endeavour is substantially faster with computation times of a few seconds on recents GPUs.
- Our algorithm is fully reproducible in terms of the part decomposition and the systematic ordering according to transparent criteria such as color, size or aspect ratio of respective parts.

We believe that the proposed algorithm may help the comedian Ursus Wehrli in his endeavour to systematically tidy up the entire history of art.

2 Image Segmentation

2.1 A Minimal Partition Model

Given the color image $I\colon \Omega \to \mathbb{R}^3$ defined over the image plane $\Omega \subset \mathbb{R}^2$ we propose to segment it into an unknown number n of pairwise disjoint regions Ω_i by minimizing the Mumford-Shah like energy [4, 12]:

$$E_\lambda(n, \Omega_1, \dots, \Omega_n, p_1, \dots, p_n) = \sum_{i=1}^{n} \lambda |\partial \Omega_i| \; - \int_{\Omega_i} \log p_i(I(x))\, dx \; + \; \nu\, n_{eff}\, . \quad (1)$$

The first term penalizes the boundary length $|\partial\Omega_i|$ of each region Ω_i, weighted with $\lambda \geq 0$. The second term is the negative log likelihood for observing a color I given that the respective point is part of region Ω_i. The last term is a penalizer of the number n_{eff} of non-empty regions weighted by positive parameter $\nu \geq 0$. It corresponds to a minimum description length prior [13, 12]. In this paper, we will simply consider isotropic Gaussian color models $p_i(I)$:

$$p_i(I) = \frac{1}{(2\pi\sigma_i^2)^{3/2}} \exp\left(-\frac{\|I-\mu_i\|_2^2}{2\sigma_i^2}\right)\,, \quad (2)$$

with mean μ_i and standard deviation σ_i, because these best reproduce Wehrli's implicit notion of part decomposition. Of course, more sophisticated color models are conceivable.

2.2 Optimization by Fast Global K-Means and Convex Relaxation

The joint optimization of (1) with respect to color analysis reveals that this difficulty arises for two reasons:

- Even for *fixed* color models, the corresponding discrete labeling problem is given by the Potts model [14] which is known to be NP hard. Without the length regularity, however, it would be a trivial problem to solve, namely a direct maximum likelihood assignment of respective pixels to their favorite color model.

- In addition, the alternating estimation of color models and region grouping is in practice prone to local minima. Moreover, the iteration of color estimation and multi-region segmentation is typically very slow and therefore impractical for interactive methods. In the absence of length regularity it is typically tackled by k-means clustering. Yet, the latter approach is known to converge to suboptimal local solutions.

We cannot expect to efficiently and optimally solve an NP hard problem. Yet, we observe that a key computational difficulty enters through the length regularity which couples the optimal decision for each pixel to respective decisions for neighboring pixels. On the other hand, in the application considered in this paper, the length regularity is generally associated with a very small weight λ because the artworks that need tidying up typically do not exhibit high levels of noise. We therefore propose to compute an initial solution by solving (1) for $\lambda = 0$:

$$E_\lambda(\{\Omega_i, \mu_i, \sigma_i\}) = \sum_{i=1}^{n} \int_{\Omega_i} \frac{\|I(x) - \mu_i\|_2^2}{2\sigma_i^2} + 3\log(\sigma_i)\, dx\,. \tag{3}$$

where n is chosen sufficiently large. To solve this problem, we revert to the fast global k-means algorithm [15] which is less prone to local minima than the traditional k-means algorithm. Alternatively, one can also retain the number n of regions in the optimization and solve the joint problem. This corresponds to the *uncapacitated facility location problem* which is known to be NP hard since it can be reduced from the *set-cover problem* – see [16] for details.

Once the initial color models (without length regularity) are estimated, we set λ to its non-zero value and solve the problem

$$E_\lambda(n, \Omega_1, \ldots, \Omega_n) = \sum_{i=1}^{n} \lambda|\partial\Omega_i| + \int_{\Omega_i} \frac{\|I(x) - \mu_i\|_2^2}{2\sigma_i^2} + 3\log(\sigma_i)\, dx + \nu\, n_{eff}\,. \tag{4}$$

2.3 Convex Formulation

The optimization problem (4) is a non-convex problem. Building up on a sequence of recent advances in variational multi-label optimization [10, 17, 9, 18] we can equivalently write it as the minimization of the *convex* energy

$$\min_{u \in \mathcal{U}_b} \sum_{i=1}^{n} \int_{\Omega} u_i(x) f_i(x)\, dx + \lambda \int_{\Omega} |Du_i| + \nu \max_{x \in \Omega} u_i(x) \tag{5}$$

over the *non-convex* set of binary indicator functions:

$$\mathcal{U}_b = \left\{ \big(u_1, ..., u_n\big) \in BV(\Omega; \{0,1\})^n \;\middle|\; \sum_{i=1}^{n} u_i(x) = 1, \quad \forall x \in \Omega \right\} . \tag{6}$$

Here Du denotes the distributional derivative (generalizing the gradient to non-differentiable indicator functions). The term f_i is given by:

$$f_i(x) = \frac{\|I(x) - \mu_i\|_2^2}{2\sigma_i^2} + 3\log(\sigma_i) . \tag{7}$$

It is the nonnegative (local) cost associated with assigning a pixel $x \in \Omega$ the label of region i.

We can relax the problem (5) to a fully convex optimization problem by allowing the functions u_i to take on real values in the interval $[0, 1]$. This amounts to replacing the constraint set $\mathcal{U}_b$ by its convex hull:

$$\mathcal{U} = \left\{ \big(u_1, ..., u_n\big) \in BV(\Omega; [0,1])^n \;\middle|\; \sum_{i=1}^{n} u_i(x) = 1, \quad \forall x \in \Omega \right\} . \tag{8}$$

Albeit convex, the arising problem (5) is highly non-smooth because of the non-differentiability of the Total Variation and the max function. By using Fenchel's duality, we can introduce two auxiliary variables p and v in order to obtain a differentiable formulation for the Total Variation and respectively for the max function. Thus the optimization problem (5) over the constraint set (8) is equivalent to the saddle-point formulation:

$$\min_{u \in \mathcal{U}} \max_{p \in \mathcal{P}} \max_{v \in \mathcal{V}} \sum_{i=1}^{n} \int_{\Omega} u_i(x) \Big(f_i(x) - \operatorname{div} p_i(x) + v_i(x) \Big)\, dx . \tag{9}$$

with respective convex sets for the dual variables:

$$\mathcal{V} = \left\{ v \in \Big(L^2(\Omega, \mathbb{R}_0^+) \Big)^n \;\middle|\; \int_{\Omega} v_i(x)\, dx = \nu; \quad \forall i = 1, \ldots, n \right\} , \tag{10}$$

$$\mathcal{P} = \left\{ p \in \Big(\mathcal{C}_c^1(\Omega, \mathbb{R}^2) \Big)^n \;\middle|\; \|p_i(x)\|_2 \leq \lambda,\ \forall x \in \Omega, \forall i = 1, \ldots, n \right\} . \tag{11}$$

This particular choice of the constraint set $\mathcal{P}$ was introduced in the work of Zach et al. [10]. While a tighter relaxations was suggested in [8, 17], we chose the former representation because the back-projections on $\mathcal{P}$ are faster to compute and because the differences in segmentation were not noticeable in the application considered here.

3 Numerical Optimization

We solve the saddle-point problem (9) by means of a recently proposed algorithm [19] and extensions of it [20]. It consists of a gradient descent in the primal and a gradient ascent in the dual variable. While the constraint on the set $\mathcal{P}$ can be handeled by simple pointwise truncation, we handeled the constraints $\mathcal{V}$ and $\mathcal{U}$ by means of lagrange mulipliers. Since all updates can be done *pointwise*, the method is straight-forwardly parallelized on a GPU allowing speedups of an order of magnitude and runtimes in the range of a few seconds.

4 Reordering of Parts

A major aspect of "Tidying up Art" is to rearrange the individual parts of the dissected painting according to some ordering principle. Our reordering formalism imitates an ordering criterion which seems to be most frequent in Wehrli's work, namely the grouping of parts based on color and size. To this end, we proceed as follows:

- For each color label $k = 1, \ldots, n$, select the region $\Omega_k = \{x \in \Omega \mid u_k(x) = 1\}$.
- For each region Ω_k, determine its connected components by means of the flood-fill algorithm and perform a postprocessing morphological closing (erosion followed by dilatation) for seperating slight connections of one or two pixels width.
- Arrange all parts horizontally according to their color label.
- In each column, arrange all parts of a given color according to their largest principal component, aligned according to their centroid and rotated such that the dominant principal axis is horizontal.

The parts are sorted in descending order with respect to their largest principal component. Thus, larger and elongated segments tend to be at the bottom of the vertical arrangement. The horizontal ordering with respect to color values is done by the hue values of the HSV color space.

5 Experiments

We ran the algorithm *Wehrli 2.0* on several artworks peforming the following steps:

- Run fast global k-means in order to determine the color labels of our color model for a certain artwork.
- Perform image segmentation algorithm introduced in Section 2 using the convex optimization approach presented in Section 3.
- Determine all segments and order them as described in Section 4.

The three steps of our algorithm are illustrated in Figure 2. We use constant deviations for the color model distribution, i.e. $\sigma_k = 1$ for each label, set $n = 15$, and choose the parameters $\lambda = 0.075$ and $\nu = 25$ for the segmentation.

(a) Painting by Jawlensky

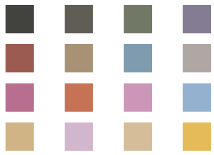

(b) Fast Global K-Means

(c) Segmentation

(c) Reordering

Fig. 2. The three steps of the *Wehrli 2.0* algorithm applied on the artwork "Mystischer Kopf: Galka" by Alexej Jawlensky (a). The color labels (b) are determined using fast global k-means. Image (c) shows the result of the MDL segmentation. The resulting regions are rearranged as in Section 4.

5.1 Fast Global K-Means vs K-Means

A comparison of the output of our algorithm using k-means and the fast global k-means algorithm (Figure 3) shows that the global k-means algorithm gives a more differentiated color model for the subsequent segmentation algorithm which in turn results in more accurate regions. The comparison shows that using fast global k-means assures that more labels are preserved in the tidied-up result.

5.2 Artworks tidied and cleaned

Many classical artworks are already decades old and with time have invariably accumulated dirt and dust. This degradation process can be a major problem for the preservation of art.[3] To account for these unfortunate effects of time, it is therefore of utmost importance that one not only tidies art but also cleans it properly. Figure 4 illustrates that *Wehrli 2.0* can handle even dirty images. A proper cleaning is obtained by simply arranging the pieces computed by our segmentation algorithm (rather than the dirty input segments). We did observe that this drastic cleaning results in a loss of small scale details. On the other hand, small scale details are substantially overrated in the art world.

[3] Rembrandt's famous painting of the Militia Company, for example, was so dimmed and defaced over the years, that later generations are now referring to it as *Night Watch*.

(a) Painting by Haring (b) K-Means (c) Fast Global K-Means

Fig. 3. The fast global k-means algorithm reproduces improved color labels compared to the k-means algorithm. As can be seen (b) we obtain a mixed color stack including yellow and amber segments since the k-means algorithm doesn't identify the amber cluster.

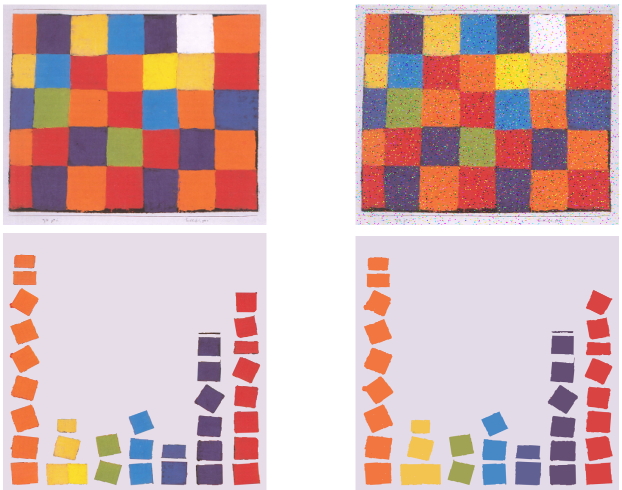

(a) Painting Tidied (b) Painting Tidied and Cleaned

Fig. 4. This figure illustrates the result of our algorithm applied to a high quality copy of "Farbtafel" by Paul Klee on the left side and the results of a noisy version of the same painting on the right side. The results are fairly similar since the segments of the noisy version are filled with the color of the corresponding color label.

5.3 Qualitative Results

Figure 5 shows three examples of artwork which the artist Wehrli has worked on. The direct comparison demonstrates that *Wehrli 2.0* produces very similar result to the manual labor of Wehrli. There is an important difference, though: While Wehrli's heuristic order of parts does not follow any recognizable logic,

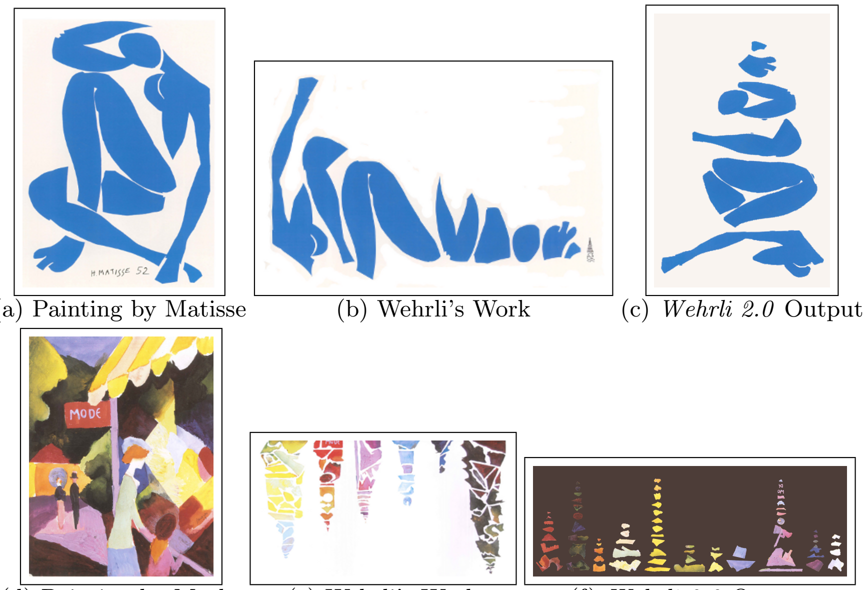

(a) Painting by Matisse (b) Wehrli's Work (c) *Wehrli 2.0* Output

(d) Painting by Macke (e) Wehrli's Work (f) *Wehrli 2.0* Output

Fig. 5. A comparison of our results using the *Wehrli 2.0* algorithm (c,f) with Wehrli's manual tidying up (b,e) of the following artworks: "Nu bleu IV" by Henri Matisse (a) and "Modefenster" by August Macke (d).

the output of *Wehrli 2.0* strictly follows simple ordering criteria and is fully deterministic and reproducible.

5.4 Runtime

All experiments were performed on a desktop PC with a NVIDIA Geforce GTX 480 GPU and a 2.40GHz quadcore CPU. For the image (a) in Figure 1 with 350x229 pixels and 15 labels, for example, the color model estimation using a Matlab implementation of the fast global k-means, took 23 seconds, while the multi-region segmentation, using a GPU implementation of the convex optimization required only 4 seconds.

6 Conclusion

We introduced the algorithm *Wehrli 2.0* to automatize the task of "Tidying up Art" introduced by the Swiss comedian Ursus Wehrli. The algorithm is based on a multi-region segmentation method which combines recent convex relaxation techniques with fast global k-means color model estimation. In contrast to Wehrli's manual work, which is tedious and time consuming, we showed that our algorithm *Wehrli 2.0* produces qualitatively similar results in a matter of seconds on a home computer.

References

1. Wehrli, U., Olenhusen, A.: Kunst aufräumen. Kein & Aber (2002)
2. Wehrli, U., Olenhusen, A.: Noch mehr Kunst aufräumen. Kein & Aber (2006)
3. Wehrli, U., Born, G., Spehr, D.: Die Kunst, aufzuräumen. Kein & Aber (2011)
4. Mumford, D., Shah, J.: Optimal approximations by piecewise smooth functions and associated variational problems. Communications on Pure and Applied Mathematics **42** (1989) 577–685
5. Blake, A., Zisserman, A.: Visual Reconstruction. MIT Press (1987)
6. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. International Journal of Computer Vision **1** (1988) 321–331
7. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE Transactions on Pattern Analysis and Machine Intelligence **23** (2001) 1222–1239
8. Chambolle, A., Cremers, D., Pock, T.: A convex approach for computing minimal partitions. Technical report TR-2008-05, Departement of Computer Science, University of Bonn, Bonn, Germany (2008)
9. Lellmann, J., Kappes, J.H., Yuan, J., Becker, F., Schnörr, C.: Convex multi-class image labeling by simplex-constrained total variation. In: Scale Space and Variational Methods in Computer Vision (SSVM), Voss, Norway (2009) 150–162
10. Zach, C., Gallup, D., Frahm, J.M., Niethammer, M.: Fast global labeling for real-time stereo using multiple plane sweeps. In: Vision, Modeling and Visualization Workshop (VMV), Konstanz, Germany (2008) 243–252
11. Klodt, M., Schoenemann, T., Kolev, K., Schikora, M., Cremers, D.: An experimental comparison of discrete and continuous shape optimization methods. In: European Conference on Computer Vision (ECCV), Marseille, France (2008) 332–345
12. Zhu, S.C., Yuille, A.: Region competition: Unifying snakes, region growing, and bayes/mdl for multi-band image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence **18** (1996) 884–900
13. Leclerc, Y.G.: Constructing simple stable descriptions for image partitioning. International Journal of Computer Vision **3** (1989) 73–102
14. Potts, R.B.: Some generalized order-disorder transformations. Mathematical Proceedings of the Cambridge Philosophical Society **48** (1952) 106–109
15. Likas, A., Vlassis, N., Verbeek, J.: The global k-means clustering algorithm. Pattern Recognition **36** (2003) 451–461
16. Delong, A., Osokin, A., Isack, H.N., Boykov, Y.: Fast approximate energy minimization with label costs. International Journal of Computer Vision **96** (2012) 1–27
17. Pock, T., Chambolle, A., Bischof, H., Cremers, D.: A convex relaxation approach for computing minimal partitions. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Miami, Florida (2009) 810–817
18. Yuan, J., Boykov, Y.: TV-based multi-label image segmentation with label cost prior. In: Proceedings of the British Machine Vision Conference (BMVC), Aberystwyth, UK (2010) 101.1–101.12
19. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: An algorithm for minimizing the piecewise smooth mumford-shah functional. In: IEEE International Conference on Computer Vision (ICCV), Kyoto, Japan (2009) 1133–1140
20. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. Journal of Mathematical Imaging and Vision **40** (2011) 120–145