

Automatic Image Colorization via Multimodal Predictions

Guillaume Charpiat^{1,3}, Matthias Hofmann^{2,3}, and Bernhard Schölkopf³

¹ Pulsar Team, INRIA, Sophia-Antipolis, France
gcharpia@sophia.inria.fr

² Wolfson Medical Vision Lab, Dpt. of Engineering Science, University of Oxford, UK
mh@robots.ox.ac.uk

³ Max Planck Institute for Biological Cybernetics, Tübingen, Germany
bs@tuebingen.mpg.de

Abstract. We aim to color greyscale images automatically, without any manual intervention. The color proposition could then be interactively corrected by user-provided color landmarks if necessary. Automatic colorization is nontrivial since there is usually no one-to-one correspondence between color and local texture. The contribution of our framework is that we deal directly with multimodality and estimate, for each pixel of the image to be colored, the probability distribution of all possible colors, instead of choosing the most probable color at the local level. We also predict the expected variation of color at each pixel, thus defining a non-uniform spatial coherency criterion. We then use graph cuts to maximize the probability of the whole colored image at the global level. We work in the L-a-b color space in order to approximate the human perception of distances between colors, and we use machine learning tools to extract as much information as possible from a dataset of colored examples. The resulting algorithm is fast, designed to be more robust to texture noise, and is above all able to deal with ambiguity, in contrary to previous approaches.

1 Introduction

Automatic image colorization consists in adding colors to a new greyscale image without any user intervention. The problem, stated like this, is ill-posed, in the sense that one cannot guess the colors to assign to a greyscale image without any *prior* knowledge. Indeed, many objects can have different colors: not only artificial, plastic objects can have random colors, but natural objects like tree leaves can have various nuances of green and turn brown during autumn, without a significant change of shape.

The color *prior* most often considered in the literature is the user: many approaches consist in letting the user determine the color of some areas and then in extending this information to the whole image, either by pre-computing a segmentation of the image into (hopefully) homogeneous color regions, or by spreading color flows from the user-defined color points. This last method supposes a definition of the difficulty of the color flow to go through each pixel,

usually estimated as a simple function of local greyscale intensity variations, as in Levin *et al.*[1], in Yatziv and Sapiro [2], or in Horiuchi [3], or by predefined thresholds to detect color edges [4]. However, we have noticed, for instance, that the simple, efficient framework by Levin *et al.* cannot deal with texture examples such as in figure 1, whereas simple oriented texture features such as Gabor filters would have obviously solved the problem⁴. This is the reason why we need to consider texture descriptors. More generally, each hand-made criterion for edge estimation has its drawbacks, and therefore we will *learn* the probability of color change at each pixel instead of setting a hand-made criterion.



Fig. 1. Failure of standard colorization algorithms in presence of texture. Left: manual initialization; right: result of Levin *et al.*'s approach. In spite of the general efficiency of their simple method (based on the mean and standard deviation of local intensity neighborhoods), texture remains difficult to deal with. Hence the need of texture descriptors and of learning edges from color examples.

User-based approaches present the advantage that the user can interact, add more color points if needed, until a satisfying result is reached, or even place color points strategically in order to give indirect information on the location of color boundaries. Our method can easily be adapted to incorporate such user-provided color information. The task of providing first a fully automatic colorization of the image, before a possible user intervention if necessary, is however much harder.

Some recent attempts of *predicting* the colors gave mixed results. For instance the problem has been studied by Welsh *et al.*[5]; it is also one of the applications presented by Hertzmann *et al.* in [6]. However the framework developed in both cases is not mathematically expressed, in particular it is not clear whether an energy is minimized, and the results shown seem to deal with only a few colors, with many small artifacts probably due to the lack of a suitable spatial coherency criterion. An other notable study is by Irony *et al.*[7]: it consists in finding a few points in the image where a color prediction algorithm reaches the highest confidence, and then in applying Levin's approach as if these points were given by the user. Their approach of color prediction is based on a learning set of colored images, partially segmented by the user into regions. The new image to be colored is then automatically segmented into locally homogeneous regions whose texture is similar to one of the colored regions previously observed, and the colors are transferred. Their method reduces the effort required of the user

⁴ Their code is available at <http://www.cs.huji.ac.il/~weiss/Colorization/>.

but still requires a manual pre-processing step. To avoid this, they proposed to segment the training images automatically into regions of homogeneous texture, but fully automatic segmentation (based on texture or not) is known to be a very hard problem. In our approach we will not rely on an automatic segmentation of the training or test images but we will build on a more robust ground.

Irony’s method also brings more spatial coherency than previous approaches, but the way coherency is achieved is still very local since it can be seen as a one-pixel-radius filter, and furthermore it relies partly on an automatic segmentation of the new image.

But above all, the latter method, as well as all other former methods in the history of colorization (to the knowledge of the authors), cannot deal with the case where an ambiguity concerning the colors to assign can be resolved only at the global level. The point is that local predictions based on texture are most often very noisy and not reliable, so that information needs to be integrated over large regions to become significant. Similarly, the performance of an algorithm based on texture classification, such as Irony’s one, would drop dramatically with the number of possible texture classes, so that there is a real need for robustness against texture misclassification or noise. In contrast to previous approaches, we will avoid to rely on very-local texture-based classification or segmentation and we will focus on more global approaches.

The color assignment ambiguity also happens when the shape of objects is relevant to determine the color of the whole object. More generally, it appears that boundaries of objects contain much not-immediately-usable information, such as the presence of edges in the color space, and also contain significant details which can help the identification of the whole object, so that the colorization problem cannot be solved at the local level of pixels. In this article we try to make use of all the information available, without neglecting any low probability at the local level which could make sense at the global level.

Another source for prior information is motion and time coherency in the case of video sequences to be colored [1]. Though our framework can easily be extended to the film case and also benefit from this information, we will deal only with still images in this article.

First we present, in section 2, the model we chose for the color space, as well as the model for local greyscale texture. Then, in section 3, we state the problem of *automatic image colorization* in machine learning terms, explain why the naive approach, which consists in trying to predict directly color from texture, performs poorly, and we show how to solve the problem by learning multimodal probability distributions of colors, which allows the consideration of different possible colorizations at the local level. We also show how to take into account user-provided color landmarks when necessary. We start section 4 with a way to *learn* how likely color variations are at each pixel of the image to color. This defines a spatial coherency criterion which we use to express the whole problem mathematically. We solve a discrete version of it via graph cuts, whose result is already very good, and refine it to solve the original continuous problem.

2 Model for colors and greyscale texture

We first model the basic quantities to be considered in the image colorization problem.

Let I denote a greyscale image to be colored, \mathbf{p} the location of one particular pixel, and C a proposition of colorization, that is to say an image of same size as I but whose pixel values $C(\mathbf{p})$ are in the standard RGB color space. Since the greyscale information is already given by $I(\mathbf{p})$, we should add restrictions on $C(\mathbf{p})$ so that computing the greyscale intensity of $C(\mathbf{p})$ should give $I(\mathbf{p})$ back. Thus the dimension of the color space to be explored is intrinsically 2 rather than 3.

We present in this section the model chosen for the color space, our way to discretize it for further purposes, and how to express continuous probability distributions of colors out of such a discretization. We also present the feature space used for the description of greyscale patches.

2.1 L - a - b color space

In order to quantify how similar or how different two colors are, we need a metric in the space of colors. Such a metric is also required to associate to any color a corresponding grey level, i.e. the closest unsaturated color. This is also at the core of the color coherency problem: an object with a uniform reflectance will show different colors in its illuminated and shadowed parts since they have different grey levels, so that we need a way to define robustly colors against changes of lightness, that is to consider how colors are expected to vary as a function of the grey level, i.e. how to project a dark color onto the subset of all colors who share a particular brighter grey level.

The psychophysical L - a - b color space was historically designed so that the Euclidean distance between the coordinates of any colors in this space approximates as well as possible the human perception of distances between colors. The transformation from standard RGB colors to L - a - b consists in first applying gamma correction, followed by a linear function in order to obtain the XYZ color space, and then by another highly non-linear application which is basically a linear combination of the cubic roots of the coordinates in XYZ . We refer to Lindbloom's website⁵ for more details on color spaces or for the exact formulas. The L - a - b space has three coordinates: L expresses the luminance, or lightness, and is consequently the greyscale axis, whereas a and b stand for two orthogonal color axes.

We choose the L - a - b space to represent colors since its underlying metric has been designed to express color coherency. In the following, by *grey level* and *2D color* we will mean respectively L and (a, b) .

With the previous notations, since we already know the grey level $I(\mathbf{p})$ of the color $C(\mathbf{p})$ to predict at pixel \mathbf{p} , we search only for the remaining 2D color, denoted by $ab(\mathbf{p})$.

⁵ <http://brucelindbloom.com/>

2.2 Discretization of the color space

This subsection contains more technical details and can be skipped in the first reading.




Fig. 2. Examples of color spectra and associated discretizations. For each line, from left to right: color image; corresponding 2D colors; location of the observed 2D colors in the ab -plane (a red dot for each pixel) and the computed discretization in color bins; color bins filled with their average color; continuous extrapolation: influence zones of each color bin in the ab -plane (each bin is replaced by a Gaussian, whose center is a black dot; red circles indicate the standard deviation of colors within the color bin, blue ones are three times bigger).

In section 3 we will temporarily need a discretization of the 2D color space. Instead of setting a regular grid, we define a discretization adapted to the color images given as examples so that each color bin will contain approximately the same number of observed pixels with this color. Indeed, some entire zones of the color space are useless and we prefer to allocate more color bins where the density of observed colors is higher, so that we can have more nuances where it makes statistical sense. Figure 2 shows the densities of colors corresponding to some images, as well as the discretizations in 73 bins resulting from these densities. We colored each color bin by the average color of the points in the bin. To obtain these discretizations, we used a polar coordinate system in ab and cut recursively color bins with highest numbers of points at their average color into 4 parts. Discussing the precise discretization algorithm is not relevant here provided it makes sense statistically; we could have used K -means for instance.

Further, in section 4, we will need to express densities of points over the whole plane ab , based on the densities computed in each color bin. In order to interpolate continuously the information given by each color bin i , we place Gaussian functions on the average color μ_i of each bin, with a standard deviation proportional to the statistically observed standard deviation σ_i of points of each color bin (see last column of figure 2). Then we interpolate the original density

$d(i)$ to any point \mathbf{x} in the ab plane by:

$$d_G(\mathbf{x}) = \sum_i \frac{1}{\pi(\alpha\sigma_i)^2} e^{-\frac{(\mathbf{x}-\mu_i)^2}{2(\alpha\sigma_i)^2}} d(i)$$

We found experimentally that considering a factor $\alpha \approx 2$ improved significantly the distribution aspect. Cross-validation could be used to determine the optimal α for a given learning set.

2.3 Greyscale patches and features

The grey level of one pixel is clearly not informative enough to decide which 2D color we should assign to it. The hope is that texture and local context will give more clues for such a task (see section 3). In order to extract as much information as possible to describe local neighborhoods of pixels in the greyscale image, we compute SURF descriptors [8] at three different scales for each pixel. This leads to a vector of 192 features per pixel. In order to reduce the number of features and to condense the relevant information, we apply Principal Component Analysis (PCA) and keep the first 27 eigenvectors, to which we add, as supplemental components, the pixel grey level as well as two biologically inspired features: a weighted standard deviation of the intensity in a 5×5 neighborhood (whose meaning is close to the norm of the gradient), and a smooth version of its Laplacian. We will refer to this 30-dimensional vector, computed at each pixel, as *local description* in the following, and denote it by \mathbf{v} or \mathbf{w} .

3 Color prediction

Now that we have modelled colors as well as local descriptions of greyscale images, we can start stating the image colorization problem. Given a set of examples of color images, and a new greyscale image I to be colored, we would like to extract knowledge from the learning set to predict colors $C(\mathbf{p})$ for the new image.

3.1 Need for multimodality

One could state this problem in simple machine learning terms: learn the function which associates to any local description of greyscale patches the right color to assign to the center pixel of the patch. This could be achieved by kernel regression tools such as Support Vector Regression (SVR) or Gaussian Processes [9].

There is an intuitive reason why this would perform poorly. Many objects can have different colors, for instance balloons at a fair could be green, red, blue, etc., so that even if the task of recognizing a balloon was easy and that we knew that we should use colors from balloons to color the new one, a regression would recommend the use of the average value of the observed balloons, i.e. grey. The problem is however not specific to objects of the same class. Local descriptions

of greyscale patches of skin or sky are very similar, so that learning from images including both would recommend to color skin and sky with purple, without considering the fact that this average value is never probable.

We therefore need a way to deal with multi-modality, i.e. to predict different colors if needed, or more exactly, to predict at each pixel the *probability* of *every* possible color. This is in fact the conditional probability of colors *knowing* the local description of the greyscale patch around the pixel considered. We will also be interested in the confidence in these predictions in order to know whether some predictions are more or less reliable than others.

3.2 Probability distributions as density estimations

The conditional probability of the color c_i at pixel \mathbf{p} knowing the local description \mathbf{v} of its greyscale neighborhood can be expressed as the fraction, amongst colored examples $e_j = (\mathbf{w}_j, c(j))$ whose local description \mathbf{w}_j is similar to \mathbf{v} , of those whose observed color $c(j)$ is in the same color bin B_i . We thus have to estimate densities of points in the feature space of grey patches. This can be accomplished with a Gaussian Parzen window estimator:

$$p(c_i|\mathbf{v}) = \left(\sum_{\{j: c(j) \in B_i\}} k(\mathbf{w}_j, \mathbf{v}) \right) / \sum_j k(\mathbf{w}_j, \mathbf{v})$$

where $k(\mathbf{w}_j, \mathbf{v}) = e^{-(\mathbf{w}_j - \mathbf{v})^2 / 2\sigma^2}$ is the Gaussian kernel. The best value for the standard deviation σ can be estimated by cross-validation on the densities. With this framework we can express how reliable the probability estimation is: its confidence depends directly on the density of examples around \mathbf{v} , since an estimation far from the clouds of observed points loses signification. Thus, the confidence in a probability prediction is the density in the feature space itself:

$$p(\mathbf{v}) \propto \sum_j k(\mathbf{w}_j, \mathbf{v})$$

In practice, for each pixel \mathbf{p} , we compute the local description $\mathbf{v}(\mathbf{p})$, but we do not need to compute the similarities $k(\mathbf{v}, \mathbf{w}_j)$ to all examples in the learning set: in order to save computational time, we only search for the K -nearest neighbors of \mathbf{v} in the learning set, with K sufficiently large (as a function of the σ chosen), and estimate the Parzen densities based on these K points. In practice we choose $K = 500$, and thanks to fast nearest neighbor search techniques such as kD-tree⁶, the time needed to compute all predictions for all pixels of a 50×50 image is only 10 seconds (for a learning set of hundreds of thousands of patches) and this scales linearly with the number of test pixels. Note that we could also have used sparse kernel techniques such as SVR to estimate, for each color bin, a regression between the local descriptions and the probability of falling into the color bin. We refer more generally to [9] for details and discussions about kernel methods.

⁶ We use, without particular optimization, the TSTOOL package available at <http://www.physik3.gwdg.de/tstool/>.

3.3 User-provided color landmarks

We can easily consider user-provided information such as the color c at pixel \mathbf{p} in order to modify a colorization obtained automatically. We set $p(c|\mathbf{p}) = 1$ and set the confidence $p(\mathbf{p})$ to a very large value. Consequently our optimization framework is still usable for further interactive colorization. A re-colorization with new user-provided color landmarks does not require the re-estimation of color probabilities, and therefore lasts only a fraction of second (see next part).

4 Spatial coherency

For each pixel of a new greyscale image, we are now able to estimate the probability distribution of all possible colors (within a big finite set of colors since we discretized the color space into bins). The interest in such computation is that, if we add a spatial coherency criterion, a pixel will be influenced by its neighbors, and the choice of the best color to assign will be done accordingly to the probability distributions in the neighborhood. Since all pixels are linked by neighborhoods, even if not directly, they all interact with each other, so that the solution has to be computed globally. Indeed it may happen that, in some regions that are supposed to be homogeneous, a few different colors may seem to be the most probable ones at a local level, but that the winning color at the scale of the region is different, because in spite of its only second rank probability at the local level, it ensures a good probability *everywhere* in the whole region. The opposite may also happen: to flip a whole such region to a color, it may be sufficient that this color is considered as extremely probable at a few points with high confidence. The problem is consequently not trivial, and the issue is to find a global solution. In this section we first *learn* a spatial coherency criterion, then find a good solution to the whole problem with the help of graph cuts.

4.1 Local color variation prediction

Instead of picking randomly a prior for spatial coherence, based either on detection of edges, or on the Laplacian of the intensity, or on a pre-estimated complete segmentation, we *learn* directly how likely it is to observe a color variation at a pixel knowing the local description of its greyscale neighborhood, based on a learning set of real color images. The technique is similar to the one detailed in the previous section. For each example \mathbf{w}_j of colored patch, we compute the norm g_j of the gradient of the 2D color (in the $L-a-b$ space) at the center of the patch. The expected color variation $g(\mathbf{v})$ at the center of a new greyscale patch \mathbf{v} is then:

$$g(\mathbf{v}) = \frac{\sum_j k(\mathbf{w}_j, \mathbf{v}) g_j}{\sum_j k(\mathbf{w}_j, \mathbf{v})}.$$

Thus we now have priors both on the colors and on the color variations.

4.2 Global coherency via graph cuts

The graph cut, or max flow, algorithm is a minimization technique widely used in computer vision [10, 11] because of its suitability for many image processing problems, because of its guarantee to find a good local minimum, and because of its speed. In the multi-label case with α -expansion [12], it can be applied to all energies of the form $\sum_i V_i(x_i) + \sum_{i \sim j} D_{i,j}(x_i, x_j)$ where x_i are the unknown variables, with possible values in a **finite** set \mathcal{L} of labels, where the V_i are any functions, and where $D_{i,j}$ are any pair-wise interaction terms with the restriction that each $D_{i,j}(\cdot, \cdot)$ should be a metric on \mathcal{L} .

The reason why we temporarily discretized the color space in section 2 was to be able to use this technique. We formulate the image colorization problem as an optimization problem on the following energy:

$$\sum_{\mathbf{p}} V_{\mathbf{p}}(c(\mathbf{p})) + \lambda \sum_{\mathbf{p} \sim \mathbf{q}} \frac{|c(\mathbf{p}) - c(\mathbf{q})|_{Lab}}{g_{\mathbf{p},\mathbf{q}}} \quad (1)$$

where $V_{\mathbf{p}}(c(\mathbf{p})) = -\log(p(\mathbf{v}(\mathbf{p})) p(c(\mathbf{p})|\mathbf{v}(\mathbf{p})))$ is the cost of choosing color $c(\mathbf{p})$ for pixel \mathbf{p} (whose neighboring texture is described by $\mathbf{v}(\mathbf{p})$) and where $g_{\mathbf{p},\mathbf{q}} = 2(g(\mathbf{v}(\mathbf{p}))^{-1} + g(\mathbf{v}(\mathbf{q}))^{-1})^{-1}$ is the harmonic mean of the estimated color variation at pixels \mathbf{p} and \mathbf{q} . An 8-neighborhood is considered for the interaction term, and $\mathbf{p} \sim \mathbf{q}$ means that \mathbf{p} and \mathbf{q} are neighbors.

The term $V_{\mathbf{p}}$ penalizes colors which are not probable at the local level according to the probability distributions obtained in section 3, with the strength depending on the confidence in the predictions. The interaction term between pixels penalizes color variation where it is not expected, according to the variations predicted in the previous paragraph.

We use the graph cut package⁷ provided by [13]. The solution for a 50×50 image and 73 possible colors is obtained by graph cuts in a fraction of second and is generally already very satisfying. The computation time scales approximately quadratically with the size of the image, which is still fast, and the algorithm performs well even on massively down-scale versions of the image, so that a good initial clue can still be given quickly for very big images too. The computational costs compete with those of the fastest colorization techniques [14] while bringing much more spatial coherency.

4.3 Refinement in the continuous color space

We can now go back to the initial problem in the continuous space of colors. We interpolate probability distributions $p(c_i|\mathbf{v}(\mathbf{p}))$ estimated at each pixel \mathbf{p} for each color bin i , to the whole space of colors with the technique described in section 2, so that $V_{\mathbf{p}}(c)$ is now defined for any color c and not only for color bins. The energy (1) can consequently be minimized in the continuous space of colors. We start from the solution obtained by graph cuts and refine it with a

⁷ available at <http://vision.middlebury.edu/MRF/code/>

gradient descent. This refinement step will generally not introduce huge changes like flipping of whole regions, but will bring more nuances.




Fig. 3. Da Vinci case: Mona Lisa colored with Madonna. The border is not colored because of the window size needed for SURF descriptors. Second line: color variation predicted (white stands for homogeneity and black for color edge); most probable color at the local level; 2D color chosen by graph cuts. Note that previous algorithms could not deal with regions such as the neck or the forehead, where blue is the most probable color at the local level because greyscale skin looks like sky. Surroundings of these regions and lower-probability colors are decisive for the final color choice.

5 Experiments

We now show results of automatic colorization. In figure 3 we colored a famous painting by Leonardo Da Vinci with another painting of his. The paintings are significantly different and textures are relatively dissimilar. The prediction of color variation performs well and helps much to determine the boundaries of homogeneous color regions. The multimodality framework proves extremely useful in areas such as Mona Lisa’s forehead or neck where the texture of skin can be easily mistaken with the texture of sky at the local level. Without our global optimization framework, several entire skin regions would be colored in blue, disregarding the fact that skin color is a second probable possibility of colorization for these areas, which makes sense at the global level since they are surrounded by skin-colored areas, with low probability of edges. We insist on the fact that the input of previous texture-based approaches is very similar to the

“most probable color” prediction (second line, middle image), whereas we consider the probabilities of all possible colors at all pixels. This means that, given a certain quality of texture descriptors, we handle much more information.

In figure 4 we perform similar experiments with photographs of landscapes. The effect of the refinement step can be observed in the sky, where nuances of blue vary more smoothly.




Fig. 4. Landscape example. Same display as in figure 3, plus last image: colors obtained after refinement step. Note that the sky is more homogeneous, the color gradients in the sky are smoother than when obtained directly by graph cuts (previous image).

We compare our method with the one by Irony *et al.*, on *their own example* [7] in figure 5; the task is easier and results are similar. The boundaries of our color regions even fit better to the zebra contour. Grass areas near the zebra are colored according to the grass observed at similar locations around the zebra in the learning image. This bias should disappear with bigger training sets since the color of the background would become independent of zebra’s presence. We would have liked to compare results on a wider scale, in particular on our difficult examples, but we face the problems of source code availability and objective comparison criteria.

In figure 6 we consider a very difficult task: the one of coloring an image from a Charlie Chaplin movie, with many different objects and textures, such as a brick wall, a door, a dog, a head, hands, a loose suit... Because of the number of objects, and because of their particular arrangement, it is unlikely to find a single color image with a similar scene that we would use as a learning image. Thus we consider a small set of three different images, each of which shares a partial similarity with Charlie Chaplin’s image. The underlying difficulty is that each learning image also contains parts which should not be re-used in this target image. The result is however promising, considering the learning set. Dealing with bigger learning datasets should slow the process only logarithmically, during the kD-tree search.

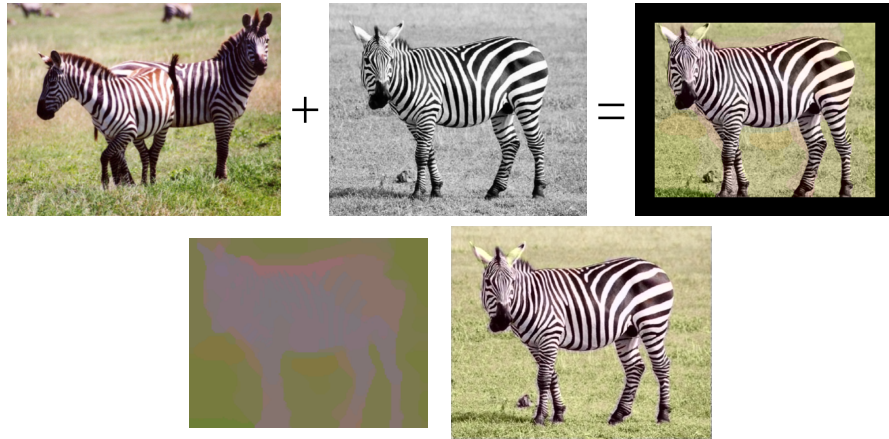


Fig. 5. Comparable results with Irony *et al.* on their own example[7]. First line: our result. Second line, left: our predicted 2D colors; right: Irony *et al.*'s result with the assumption that it is a binary classification problem. Note that this example they chose could be solved easily without texture consideration since there is a simple correspondence between grey level intensity and color. Concerning our result, the color variations in the grass around the zebra are probably due to the influence of the grass color for similar patches in the learning image, and this problem should disappear with a larger training set. Note however, if you zoom in, that the contour of the zebra matches exactly boundaries of 2D color regions (purple and pink), whereas in Irony *et al.*'s result, a several pixel wide band of grass around the zebra's legs and abdomen are colored as if they were part of the zebra. The moral is that we should not try to learn colors from a single overfitted, carefully-chosen training image but rather check how the colorization scales with the number of training images and with their non-perfect suitability.

6 Discussion

Our contribution is both theoretical and practical.

We proposed a new approach for automatic image colorization, which does not require any intervention by the user, except from the choice of relatively similar color images. We stated the problem mathematically as an optimization problem with an explicit energy to minimize. Since it deals with multimodal probability distributions until the final step, this approach makes better use of the information that can be extracted from images by machine learning tools. The fact that we solve directly the problem at the global level, with the help of graph cuts, makes the framework more robust to noise and local prediction errors. It also makes it possible to resolve large scale ambiguities which previous approaches could not. This multi-modality framework is not specific to image colorization and could be re-used in any prediction task on images.

We also proposed a way to learn and predict where color variations are probable and how important they are, instead of choosing a spatial coherency criterion

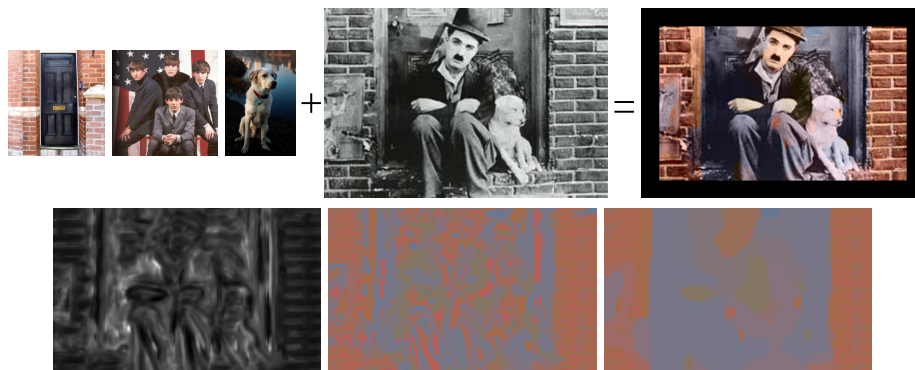


Fig. 6. Charlie Chaplin frame colored by three different images. Second line: prediction of color variations, most probable color at the local level, and final result. This example is particularly difficult because many different textures and colors are involved and because there does not exist color images very similar to Charlie Chaplin’s one. We have chosen three different images, each of which shares a partial similarity with the target image. In spite of the complexity, the prediction of color edges or of homogeneous regions remains significant. The brick wall, the door, the head and the hands are globally well colored. The large trousers are not in the learning set; the mistakes in the colors of Charlie Chaplin’s dog are probably due to the blue reflections on the dog in the learning image and to the light brown of its head. Both would probably benefit from larger training sets.

by hand, and this performs quite well (see figures 3 and 6 again). It relies on many features and adapts itself to the current training set.

Our colorization approach outperforms [5] and [6], whose examples contain only few colors and lack spatial coherency. Our process requires less or similar intervention than [7] but can handle more ambiguous cases and more texture noise. The computational time needed is also very low, enabling real-time user interaction since a re-colorization with new color landmarks lasts a fraction of second.

The automatic colorization results should not be compared to those obtained by user-helped approaches since we do not benefit from such decisive information. Nevertheless, even with such an handicap, our results clearly compete with the state of the art of colorization with few user scribbles.

Acknowledgments

We would like to thank Jason Farquhar, Peter Gehler, Matthew Blaschko and Christoph Lampert for very fruitful discussions.

References

1. Levin, A., Lischinski, D., Weiss, Y.: Colorization using optimization. In: SIGGRAPH '04, Los Angeles, California, ACM Press (2004) 689–694
2. Yatziv, L., Sapiro, G.: Fast image and video colorization using chrominance blending. *IEEE Transactions on Image Processing* **15**(5) (2006) 1120–1129
3. Horiuchi, T.: Colorization algorithm using probabilistic relaxation. *Image Vision Computing* **22**(3) (2004) 197–202
4. Takahama, T., Horiuchi, T., Kotera, H.: Improvement on colorization accuracy by partitioning algorithm in cielaab color space. In Aizawa, K., Nakamura, Y., Satoh, S., eds.: *PCM (2)*. Volume 3332 of *Lecture Notes in Computer Science.*, Springer (2004) 794–801
5. Welsh, T., Ashikhmin, M., Mueller, K.: Transferring color to greyscale images. In: *SIGGRAPH '02: Proc. of the 29th annual conf. on Computer graphics and interactive techniques*, ACM Press (2002) 277–280
6. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.H.: Image analogies. In Fiume, E., ed.: *SIGGRAPH 2001, Computer Graphics Proceedings*, ACM Press (2001) 327–340
7. Irony, R., Cohen-Or, D., Lischinski, D.: Colorization by Example. In: *Proceedings of Eurographics Symposium on Rendering 2005 (EGSR'05, June 29–July 1, 2005, Konstanz, Germany)*. (2005) 201–210
8. Bay, H., Tuytelaars, T., Van Gool, L.: Surf: Speeded up robust features. In: *9th European Conference on Computer Vision, Graz Austria (May 2006)*
9. Schölkopf, B., Smola, A.J.: *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA (2001)
10. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. In: *EMMCVPR '01: Proc. of the Third Intl. Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition, London, UK, Springer-Verlag (2001)* 359–374
11. Kolmogorov, V., Zabih, R.: What energy functions can be minimized via graph cuts? In: *ECCV*. (2002) 65–81
12. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. In: *ICCV*. (1999) 377–384
13. Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M.F., Rother, C.: A comparative study of energy minimization methods for markov random fields. In Leonardis, A., Bischof, H., Pinz, A., eds.: *ECCV*. Volume 3952 of *Lecture Notes in Computer Science.*, Springer (2006) 16–29
14. Blasi, G., Recupero, D.R.: Fast Colorization of Gray Images. In: *Proc. of Eurographics Italian Chapter*. (2003)