Correcting Colours for Aided Recomposition of Fragments

G. Carlomagno, F. Renna, N. Ancona, N. Mosca, G. Attolico and A. Distante
Istituto di Studi sui Sistemi Intelligenti per l'Automazione,
C. N. R. Via Amendola, 122 D-I –70126 - Bari
Italy
[renna, attolico]@ba.issia.cnr.it

Abstract

The paper describes the colour correction approach of a system for the virtual aided recomposition of fragmented frescos, being developed and proved for the reconstruction of the S. Matthew's fresco, painted at the end of the 13th century by Cimabue for the Upper Church of S. Francis in Assisi and broken in more than 140.000 fragments during the earthquake of 1997. One of its key features is the selection of fragments by a query-by-example approach applied to sample images (fragments and/or details extracted from the whole image of the fresco). Unfortunately the only available picture of the whole fresco has been acquired without any colour reference and is chromatically different across its extension and with respect to fragments: its comparison with them requires specific colour transformations to be applied to each of its regions. Two different methods to identify colour correspondences between fresco and fragments (manually and automatically from the already placed fragments) and two approaches to solve the resulting least-squares problem and evaluate the transformation matrices have been used and compared. The solution used in the system provides effective results at a very low computational cost. The restorers do not need to be aware of technical details related to the colour correction problem. The obtained improvement in terms of colour similarity is shown.

1. Introduction

This work describes a component of a system for the virtual aided recomposition of fragmented frescos, whose interest arises from the need to recompose the S. Matthew fresco, painted at the end of the 13th century by Cimabue for the Upper Church of S. Francis in Assisi and broken into more than 140.000 fragments by the earthquake in September 1997. The large extension of the fresco (about 35 squared meters), the huge number of fragments and the technique used by Cimabue (that makes the pictorial film very sensitive to the physical manipulation required by the traditional recomposition) have suggested the application of digital tools to this challenging problem. Moreover, fragments do not cover the whole fresco, partially belong to a neighbour fresco broken during the same event and exhibit contours that do not always match exactly.

The system transposes the traditional recomposition process in a digital way. The operators have the critical role of applying new tools and flexible algorithms of image analysis to increase the efficacy of their work [1].

On the multi-monitor graphical station (Fig. 1) a part of the image of the whole fresco (shown in a scaled version on the central monitor) can be selected as background for the working area displayed at full resolution on the left-side monitor. The operator simultaneously rotates and translates each fragment across the region of interest using a special mouse to find its place. Virtual containers (right monitor) are the digital counterpart of the boxes used in the real lab to organize logically related fragments.

Figure 1. The developed workstation for the virtual aided recomposition of fragmented frescos

To improve the recomposition process, the system supports the retrieval of digital images of fragments from the database using an incremental and iterative query-by-example modality. The operator picks up a set of images (fragments and/or details of the reference image) and the system selects from the database the fragments more similar to them. This process can be repeated with different sets (by adding, removing or changing the examples) until the operator's needs are fulfilled. Colour and texture are the most important features for similarity evaluation [3]: in fact shapes, damaged during the fragmentation process, do not necessarily match perfectly.

Unfortunately, the colours of the reference image, acquired several years before the earthquake under unknown illumination conditions, are very different from those of real fragments. Moreover, the fresco was painted on a vaulted ceiling, inducing different lighting conditions (and appearance of colours) in different areas of the fresco and, after the acquisition, several causes have produced different deteriorations across the extension of the fresco.

Correcting colours of the reference image [9] allows the use of its details to index the database and simplify its comparison with the images of fragments to find their correct place.

Two methods have been proved to identify the colour pairs required by the colour correction: the manual extraction of almost monochromatic patches from both the reference image and the corresponding fragments and the automatic processing of polychromatic fragments and of their corresponding patches from the reference image. Results obtained by these two approaches are compared. SVD and a constrained linear least-squares approach have been compared to derive from these pairs a matrix that corrects colours. Specific transformations are required for the particular chromatic characteristics of each part of the whole reference image. The process automatically applies to each fragment whose position in the whole fresco has been recognized to evaluate/refine the colour correction in its neighbourhood.

2. Data Set Selection

This step identifies the colour correspondences required to evaluate the matrix that should bring the reference image to match the corresponding colours of fragments [2][7]. Picking up corresponding colours is hard because the image of the whole fresco has a resolution that is about five times lower than that of fragments, which generally do not match the orientation of the corresponding patch of the fresco. Two techniques, one manual and one automatic, have been compared to extract colour pairs.

The manual procedure extracts, by hand, a suitable number of patches from both the reference image and the corresponding fragments: they must be mostly homogeneous in colour. These patches are therefore considered as monochromatic and are represented by the mean colour of their pixels. In this way, each patch provides one of the colours used to compute the transformation matrix. This technique takes a significant amount of time and requires the operator to have a greater consciousness of the colour correction process to reach a fruitful distribution of pairs in the colour space.

The automatic extraction of colour pairs requires:

- to extract from the reference image a patch having the same shape of the fragment;
- to spatially register the patch and the fragment image;
- to sample the registered images.

The registration of patches extracted from the reference image, to correctly align them with the corresponding fragments, is mostly based on the moment features [5]. They rely on the evaluation of the scale factor and of the rotation that must be applied to the patch. A binary mask of each image is produced, where black is associated to background and white corresponds to meaningful pixels. The centre of mass and the orientation of each fragment can be evaluated from this mask. The centre of mass of the white region R, with N pixels in m rows and n columns, is:

$$\overline{m} = \frac{1}{N} \sum_{(m,n)\in R} m \qquad \overline{n} = \frac{1}{N} \sum_{(m,n)\in R} n$$

While the orientation is:

$$q = \frac{1}{2} \tan^{-1} \left[\frac{2\, m_{1,1}}{m_{2,0} - m_{0,2}} \right]$$

where the central moments are:

$$m_{p,q} = \sum_{(m,\, n)\in R} (m - \overline{m})\,(n - \overline{n})$$

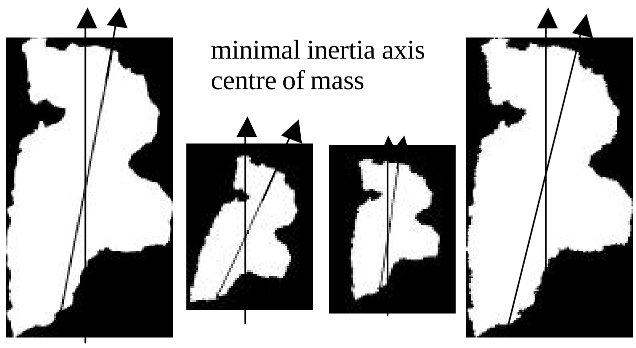

Figure 2. The first image represents the mask of a fragment: the axis of minimal inertia passes through its centre of mass. The second, third and fourth images show the mask of the corresponding patch extracted from the reference image: before and after rotation, and after the application of the estimated scale factor

The ratio between the length of the axes of minimal inertia, evaluated on the masks of both images, returns the scale factor. Moreover, the difference between the orientation angles returns the rotation angle (Fig. 2). After the registration, the two images can be sampled to obtain the colour pairs: the whole process can be applied to every fragment with a simple and fully automatic routine.

3. Colour Correction

Colour correction must increase the chromatic similarity between the fragment F at hand and the corresponding patch P of the image of the whole fresco.

Let us suppose P and F to be correctly aligned, so that the pixel $P(i,j)$ of the patch corresponds to the pixel $F(i,j)$ of the fragment. Moreover, let us indicate with **x** and **y** the triplets of RGB colours relative to the same pixels in P and F respectively. The problem of colour correction can be regarded as the problem of finding a vector valued function **f** such that **y=f(x)**. Note that we are not interested in analysing the causes of colour changes between the images of fresco and fragment, such as illumination, sensor characteristics, modifications of colours in time. On the contrary, we are interested in finding a mapping between colours on the basis of some examples of these changes, each taking the form of an input-output pair **(x,y)**.

This set of examples of size l, called training set, is defined as $S = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{l}$ where $\mathbf{x}_i \in \Re^3$ are the colours of the reference image, $\mathbf{y}_i \in \Re^3$ are those of fragments and $i = 1, 2, \ldots, l$ is the number of considered pairs. This data set can be considered as obtained by a random sampling of an unknown vector valued function **y=f(x)** in presence of noise. In general, **f** is a non linear transformation that is estimated starting from a finite number of examples. In this paper the analysis will be limited to linear approximations of the unknown function **f**. The best linear approximations obtained by two methods are compared: the former solves a constrained minimization problem, the latter uses the Singular Value Decomposition.

Both methods determine twelve elements: at least four colours are needed (each providing three equations) but a larger set of l linear equations, solved using the least-squares method [4], provides a more reliable solution.

3.1. Regularized Linear Least Squares

The aim is to determine the vector valued linear function **y=Ax+b,** where **A** is a q by d matrix (in our case q=d=3) and **b** is a vector in $\Re^q$, which approximates at the best the function **f** in the least squares sense. This is equivalent to solve the following constrained minimization problem:

$$\min_{\mathbf{A},\mathbf{b}} \frac{1}{l}\sum_{i=1}^{l}\left\| \mathbf{y}_i - (\mathbf{A}\mathbf{x}_i + \mathbf{b})\right\|^2$$

subject to $\|\mathbf{A}\|_F^2 \le a$ where $a \in \Re$, $\|\mathbf{A}\|_F^2$ is the Frobenius norm of $\mathbf{A}$ defined as $\|\mathbf{A}\|_F^2 = tr(\mathbf{A}\mathbf{A}^T)$ and $\|\cdot\|$ is the Euclidean norm. At this aim, let us consider the following Lagrangian function:

$$L(\mathbf{A},\mathbf{b}) = \frac{1}{l}\sum_{i=1}^{l}\left\|\mathbf{y}_i - (\mathbf{A}\mathbf{x}_i + \mathbf{b})\right\|^2 + l\left(\|\mathbf{A}\|_F^2 - a\right)$$

where l is the Lagrange multiplier relative to the constraint of the described problem. To determine the solution, the partial derivates of $L(\mathbf{A},\mathbf{b})$ with respect to the unknowns **b** and **A** have been set to zero, so that:

$$\mathbf{A} = \mathbf{S}_{yx}\left(\mathbf{S}_{xx} + l\,\mathbf{I}_d\right)^{-1} \text{ and } \mathbf{b} = \boldsymbol{\mu}_y - \mathbf{A}\boldsymbol{\mu}_x$$

where:

$$\mathbf{S}_{xx} = \frac{1}{l}\sum_{i=1}^{l}(\mathbf{x}_i - \boldsymbol{\mu}_x)(\mathbf{x}_i - \boldsymbol{\mu}_x)^T$$

$$\mathbf{S}_{yx} = \frac{1}{l}\sum_{i=1}^{l}(\mathbf{y}_i - \boldsymbol{\mu}_y)(\mathbf{x}_i - \boldsymbol{\mu}_x)^T$$

$$\boldsymbol{\mu}_x = \frac{1}{l}\sum_{i=1}^{l}\mathbf{x}_i \quad \boldsymbol{\mu}_y = \frac{1}{l}\sum_{i=1}^{l}\mathbf{y}_i$$

and l is the only free parameter that can be estimated on the basis of the leave-one-out error.

3.2. SVD Method

This approach uses Singular Value Decomposition [6]. In this case, the aim is to compute the matrix $\mathbf{A_1}$ such as:

$$\mathbf{A_1} = [\mathbf{A}|\mathbf{b}] \quad \text{and} \quad \left[y_i^r \; y_i^g \; y_i^b\right]^{\mathrm{T}} = \mathbf{A}_1\left[x_i^r \; x_i^g \; x_i^b \; 1\right]^{\mathrm{T}}$$

Let us define:

$$\mathbf{a} = [a_{11}\; a_{12}\; a_{13}\; b_1\; a_{21}\; a_{22}\; a_{23}\; b_2\; a_{31}\; a_{32}\; a_{33}\; b_3]^{\mathrm{T}}$$

$$\mathbf{y} = \left[y_1^r \quad y_1^g \quad y_1^b \quad . \quad . \quad . \quad y_l^r \quad y_l^g \quad y_l^b\right]^{\mathrm{T}}$$

$$\mathbf{X} = \begin{bmatrix} x_1^r & x_1^g & x_1^b & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & x_1^r & x_1^g & x_1^b & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & x_1^r & x_1^g & x_1^b & 1 \\ . & . & . & . & . & . & . & . & . & . & . & . \\ x_l^r & x_l^g & x_l^b & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & x_l^r & x_l^g & x_l^b & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & x_l^r & x_l^g & x_l^b & 1 \end{bmatrix}$$

so that: $\mathbf{y} = \mathbf{X}\mathbf{a}$. The matrix **X** can be decomposed with the well known SVD; its factorization is $\mathbf{X} = \mathbf{U}\mathbf{W}\mathbf{V}^{\mathbf{T}}$ where **U** (3l by 3l) and **V** (12 by 12) are orthogonal matrices and $\mathbf{W} = diag(s_1, s_2, \ldots, s_r)$ with r = min (3l, 12) and $s_1 \ge s_2 \ge \ldots \ge s_r$ singular values of matrix **X**. Therefore, by solving the problem in the least squares sense, the following identity exists:

$$\|\mathbf{X}\mathbf{a} - \mathbf{y}\|_2^2 = \left\|\mathbf{U}^{\mathbf{T}}\left(\mathbf{U}\mathbf{W}\mathbf{V}^{\mathbf{T}}\mathbf{a} - \mathbf{y}\right)\right\|_2^2 = \left\|\mathbf{W}\mathbf{V}^{\mathbf{T}}\mathbf{a} - \mathbf{U}^{\mathbf{T}}\mathbf{y}\right\|_2^2$$

By changing variables $\mathbf{V^T a}=\mathbf{z}$ and $\mathbf{U^T y}=\mathbf{d}$, the previous identity become:

$$\|\mathbf{Xa}-\mathbf{y}\|_2^2 = \|\mathbf{Wz}-\mathbf{d}\|_2^2 =$$
$$=(s_1 z_1 - d_1)^2 + \ldots + (s_r z_r - d_r)^2 + d_{r+1}^2 + \ldots + d_{3l}^2$$

Then the solution $\mathbf{z}$ that minimizes the problem is:

$z_1 = \dfrac{d_1}{s_1}, \ldots, z_r = \dfrac{d_r}{s_r}$ and the residual is $d_{r+1}^2 + \ldots + d_{3l}^2$.

4. Leave -One -Out Error

The leave-one-out error (loo-err) measures the error of generalization of a supervised learning machine, that is its capacity of correctly predicting the output for new input patterns. This rather general procedure leaves out one example $(\mathbf{x}_i, \mathbf{y}_i)$ at a time from the training set, trains the estimator $\mathbf{f}^i$ on the remaining l-1 examples and computes the error between the prediction value $\mathbf{f}(\mathbf{x}_i)$ and the target $\mathbf{y}_i$ given by $\left\|\mathbf{y}_i - \mathbf{f}^i(\mathbf{x}_i)\right\|^2$. The quantity:

$$L = \frac{1}{l}\sum_{i=1}^{l}\left\|\mathbf{y}_i - \mathbf{f}^i(\mathbf{x}_i)\right\|^2$$

is a measure of the generalization error of the learning machine $\mathbf{f}$ trained using all the training examples (see the theorem of Luntz and Brailovsky in [8]). This procedure is time consuming (it requires to train l different predictors, one for each training sample) but it is useful for at least two different tasks: comparing different estimators trained on the same data set (see 5.2) and setting their free parameters (see 3.1).

5. Experimental Results

The experiments aim to point out the performance of the two methods for data set selection and of the two methods for solving the least mean squares problem.

5.1. Comparison between Manual and Automatic Approaches

The manual and automatic approaches, to build the training sets needed to evaluate the transformation matrices, have been compared on two regions of the fresco: Judean and Mantle. In this case, the transformation matrices have been computed using the SVD method. The matrices are JM and JA (evaluated from data extracted from the Judean area with the manual and automatic methods respectively) and MM and MA (evaluated from data extracted from the Mantle area).

The two methods have been compared by estimating the mean m of the Euclidean distances and their standard deviations on the training set $S = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{l}$. The mean m has been computed with the following expressions:

$$m_C = \frac{1}{l}\sum_i \|\mathbf{y}_i - \mathbf{A}\mathbf{x}_i\|^2 \qquad m_{NC} = \frac{1}{l}\sum_i \|\mathbf{y}_i - \mathbf{x}_i\|^2$$

where $\mathbf{A}$ is one of the evaluated matrices. The results show that all distances are sensibly reduced after colour correction. Moreover, the linear predictors trained by Judean area performs well also on Mantle area showing a good generalization capacity to correctly predict unknown coloured areas (Table 1 and 2) and vice versa.

Training set Judean	NC	JM	JA	MM	MA
Mean Error	50,05	14,41	16,94	18,60	15,91
Stand Dev	25,47	9,19	13,19	14,51	8,26
Training set Mantle					
Mean Error	46,81	16,93	17,97	11,85	17,69
Stand Dev	17,84	10,82	11,04	7,92	12,33

Table 1. The column NC reports the mean error m_{NC} and the standard deviation of distances on the training set used to evaluate the JM (above) and MM (below) matrices. The others report the same quantities on the same training set but corrected using the four different matrices

Training set Judean	NC	JM	JA	MM	MA
Mean Error	61,04	31,81	30,17	32,89	30,60
Stand Dev	36,65	24,61	21,99	24,58	22,15
Training set Mantle					
Mean Error	55,04	27,52	24,06	26,77	23,20
Stand Dev	30,19	21,78	19,11	21,33	19,11

Table2. The same measures described in table 1, but obtained on training sets extracted automatically

Afterwards, the Euclidean distance between corresponding colours have been integrated over five fragments and their registered patches, before and after colour correction. Two of these fragments (Frag 3 and Frag 5) belong to the Judean area and have been used to evaluate its transformations. Other two (Frag 1 and Frag 2) belong to the Judean area but have not been used to evaluate the matrices. The last one (Frag 4) belongs to the Mantle and has been used to evaluate its transformations. For each fragment the resulting measure is:

$$\frac{\sum_{(m,n)}\left\|(r_f - r_{pc})^2 + (g_f - g_{pc})^2 + (b_f - b_{pc})^2\right\|}{\sum_{(m,n)}\left\|(r_f - r_{nc})^2 + (g_f - g_{nc})^2 + (b_f - b_{nc})^2\right\|} \qquad (1)$$

The (r, g, b)$_f$, (r, g, b)$_{nc}$ and (r, g, b)$_{pc}$ triplets are the RGB components of pixels in fragments, not corrected and corrected corresponding patches respectively. Figure 3 shows that the whole amount of these differences becomes about less than a half for every correction.

Another measure (Table 3) quantifies the number of colours common to each of the five fragments and its corresponding patch, before and after colour correction. The registered patches exhibit much less colours than the corresponding fragments. This is due to the scale factor that exists between the real fragments and the reference image (about five): the registration step enlarge the patch extracted from the reference image but without adding new colours. Moreover transformation matrices normally cause a further reduction in the number of colours by compressing their distribution in the colour space.

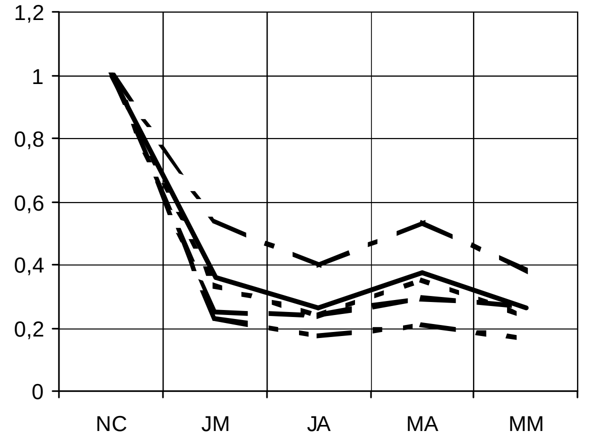

Figure 3. The graphic shows the measure described in (1) applied to five fragments and to their corresponding patches not corrected and corrected using the four correction matrices. Frag1 dashed line, Frag2 dotted, Frag3 solid, Frag4 dash-dot-dot, Frag5 dash-dot

	Frag	NC	JM	JA	MM	MA
Frag1	65933	17380	14938	7585	15573	7480
Frag2	99635	20535	17321	9174	17893	9090
Frag3	86530	14717	13142	8117	13449	8048
Frag4	67859	13759	12063	6732	12496	6657
Frag5	97209	28047	22618	10179	23722	9964

Table 3. The table shows the number of colours common to each fragment and its corresponding registered patch, without and with correction

The intersection and the union of the two histograms (of each fragment and of corresponding patch) have been evaluated: the number of colours belonging to both these images increases with colour correction (Table 4). A larger number of colours of corrected patches is representative of the corresponding fragments. The percentage of colours belonging to the intersection with respect to the total number of colours of the patch has been evaluated (Table5). Corrected patches include a greater number of colours that belong to the corresponding fragment. The correction increases colour similarity between the reference image and the corresponding fragments, making meaningful the retrieval process based on colour similarity with respect to details extracted from the corrected reference image.

Also the percentage of colours of the reference image that are present in the histogram of fragments increases after correction (Table 5) that makes significantly closer the two colour spaces.

	patch NC	patch JM	patch JA	patch MM	patch MA
Frag1	2396	9906	5766	6793	5286
Frag2	2186	11822	8803	9228	8684
Frag3	810	8455	7589	7251	7172
Frag4	430	4294	5726	7748	5786
Frag5	4021	15713	9739	12741	9606

Table 4. The table shows the number of colours shared by the histograms of each fragment and of the corresponding registered patch

	NC %	JM %	JA %	MM %	MA %
Frag1	13,8	66,3	76	43,6	70,1
Frag2	10,6	68,2	95,9	51,6	95,5
Frag3	5,5	64,3	93,5	53,9	89,1
Frag4	3,1	35,6	85	62	86,9
Frag5	14,3	69,5	95,7	53,7	96,4

Table 5. The table shows the percentage of colours of patches extracted from the reference image that are also present in the corresponding fragments. The results refer to patches not corrected and corrected with the JM, JA, MM, and MA matrices respectively

Figure 4 visually represent colour intersection. On the left, there is a patch extracted from the reference image, the second, third and fourth images show in dark grey (white) the pixels whose colour is present (not present) in the corresponding fragment before and after colour correction with the matrices JM and JA respectively.

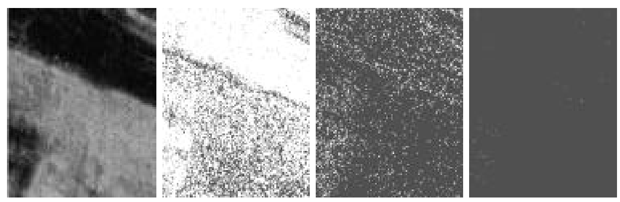

Figure 4. A patch extracted from the reference image (left). The other three images represent in dark grey the pixels whose colour is present in the corresponding fragment: the second refers to the not corrected patch; the third and the fourth to the patch corrected using the JM and JA matrices respectively

The carried out experiments show that the performance of the automatic approach is comparable to manual. Since it is necessary to efficiently evaluate different matrices to correct different areas of the fresco, the automatic approach has been chosen and implemented in the system.

5.2. Comparison between Line ar Approximation Methods

Table 6 shows the results obtained by comparing the SVD-based and the regularized linear least-squares methods. The experiment has been carried out by selecting different training sets consisting of 500, 5.000 and 50.000 RGB triplets. The performance of the two resolution methods have been evaluated using the leave-one-out and the mean error on training sets. Experiments confirm that the regularized linear method provide the same effective results with a sensibly lower computational load, being effective for online colour correction.

Training set	Loo1-Err	Loo2-Err	Mean-Err1	Mean-Err2
500	30.041	30.041	30.138	30.132
5000	30.086	30.087	30.001	30.001
50000	30.044	30.047	30.009	30.012

Table 6. The leave-one-out error and the mean error obtained applying the SVD-based (1) and the regularized linear least-squares (2) methods on the data extracted using the automatic approach to colour pairs selection

Figure 5. An area of the fresco, with some of the fragments correctly located using the system, before and after colour correction: fragments are almost unnoticeable after the colour correction

6. Conclusion

The critical problem of reducing the distance of colours between a reference image and the fragments inside a system for virtual aided recomposition of fragmented frescos has been described. It affects the query-by-example retrieval of fragments from the database, mainly colour-based, and their placement identified by restorers.

The colour transformation has been modelled as a matrix provided by a Least-Squares approach applied on pairs of corresponding colours extracted using two different methods. The former uses the mean RGB values of almost monochromatic patches manually extracted from the reference image and the corresponding real fragments. The latter automatically extracts colours pairs from fragments and the corresponding correctly aligned patches from the fresco: it does not require special interventions by the operator because the system processes any new fragment placed on the fresco to evaluate a suitable transformation for its neighbourhood. In both cases the correction strongly reduces the distance between images.

Two different methods to solve the Least-Squares problem and to determine a suitable linear approximation of the transformation in the colour space have been proved. The first solves a suitably composed over-determined system using the SVD decomposition to identify the twelve elements of the desired matrix. The other one uses a constrained linear least-squares approach that strongly simplify the computational load, without decreasing the efficacy of the obtained results, being more suitable for online colour correction.

The automatic extraction of colours pairs and the constrained linear least-squares compose an automatic tool for colour correction that increases the chromatic similarity between the reference image of the fresco and the fragments. This results has a twofold effect: it improves the visual perception of the operators that look for the right position of fragments and makes more meaningful the retrieval of fragments performed by the query-by-example approach applied to details extracted from the corrected reference image.

7. Acknowledgements

The authors thank Marco Malavasi (Dip. Rapporti con le Regioni – CNR) for its important support; Giuseppe Basile, Lidia Rissotto, Angelo Rubino (ICR), Manuela Viscontini (Università della Tuscia) and Laura Cacchi for many valuable discussions during project development.

References

[1] G. Attolico, A. Distante. Aided virtual recomposition of fragmented frescos, *Dall'utopia alla realtà, Note dal cantiere dei frammenti della Basilica Superiore di S. Francesco in Assisi*, n. 2, April 2001.

[2] M. Barni, V. Cappellini, F. Bartolini. Image Processing for Virtual Restoration of Artworks, *IEEE MultiMedia Virtual Heritage*, June 2000, pp. 34 – 37.

[3] A. Branca, G. Attolico, S. Tornesello, A.Distante. Fragment similarity estimation for the restoration of frescos, *Proc. of the Irish Machine Vision & Image Processing Conference*, September 2001, Maynooth.

[4] G.H. Golub, C.F. Van Loan. *Matrix Computations*, The Johns Hopkins University Press, 1996.

[5] B. Jähne. *Digital Image Processing*, Springer, 2002.

[6] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery. *Numerical Recipes in C*, Cambridge University Press.

[7] Report on Cultural Heritage Project, supported by CNR (Italian National Research Council): http://lci.det.unifi.it/Projects/PFBC/Ripulitura.htm.

[8] V. Vapnik. *Statistical Learning Theory*, John Wiley & Sons, INC., 1998.

[9] J. Yin, J. R. Cooperstock. *Colour Correction Methods with Applications to Digital Projection Environments*, Journal of WSCG, Vol.12, No.1-3, February 2-6, 2004, Plzen, Czech Republic.