

# Using Colour to Model Outliers

*David Hasler<sup>1</sup> and Sabine Süsstrunk<sup>2</sup>*

*<sup>1</sup>GretagMacbeth, Steinfurt, Germany*

*<sup>2</sup>Audiovisual Communications Laboratory  
Swiss Federal Institute of Technology (EPFL)  
Lausanne, Switzerland*

## Abstract

Computer vision applications are able to model and reconstruct three dimensional scenes from several pictures. In this work, we are interested in the group of algorithm that register each image with respect to the model and aim at constructing a model of the scene. At the lowest level, most of these algorithms are comparing the pixel values of the image to the ones predicted by the model to refine the result. As research advances, the models are getting better and better, but no matter how complex they are, there will always be unpredictable situations that cannot be handled by the model. A recurring example is when an object appears in one image of the set, but in none of the others. The situation occurs, for example, when a moving entity crosses rapidly the field of view of the camera. In this work, we study the error generated by such an unexpected object at a pixel level and how colour can improve the estimation. We will derive the expected error distribution that this hypothetical object may cause.

Our model is primarily intended as a basis for outlier removal in scene modelling algorithms. It gives a clear answer to whether, and with which confidence, a part of the image can be considered as part of the model or should be discarded, without using any dedicated thresholding scheme.

The model is demonstrated on a trivial example where we match two images of a scene using a static camera. The example shows that the outlier distribution can be predicted by using the histograms of both images. We also show that by considering not only greyscale information, but also colour information, the outlier detection performance improves. We want to emphasise that the central part of this paper is the outlier modelling and not the outlier rejection scheme, which could be solved—for the trivial examples we are showing—by many other techniques.

## Keywords

Outlier model, outlier rejection, colour segmentation, computer vision.

## 1. Introduction

We are interested in studying the modelling of scenes constructed from several photographs. More specifically, we are building a tool to improve the robustness of computer vision algorithms that work with multiple views of a scene. These models often use bayesian inference where everything fits to a nice theoretical framework until the system has to deal with objects that suddenly appear in one image of the set, which we call *outliers*. The usual techniques to handle this situation are based on optimisation or statistics [1, 2] and are applied to the imaging domain without taking into account all the knowledge that could be extracted from images. Generally speaking, a standard outlier rejection scheme consists in characterising the class of data to be expected from the experiment, and defines the outliers as being the data set that “diverges too much” from this class [3, 4, 5, 6, 7]. This implies the use of a threshold which is chosen in a somewhat arbitrary fashion—for example by deciding that the system should allow a certain amount of false rejection. Other techniques assume that the outliers can be characterised by a uniform distribution [8, 9, 10] or a gaussian distribution [11] and handle the outlier rejection as a standard mixture problem<sup>1</sup>. An outlier model based on image content is introduced in [12], and shows on an image stitching example that such a model improves substantially the robustness of the stitching to outliers.

We propose to extend the work in [12] to take advantage of colour information in the images. We also characterise an *outlier* in an image sequence as a region that undergoes an unexpected motion, or an object that suddenly appears in one of the images. The key idea is in the way the outlier statistical characteristics are predicted: *We assume that by comparing two arbitrary parts of two images, we get an error pattern similar to the one generated by an outlier*. We then compute the expected error distribution generated by the outliers. In a more general context, one image is the scene measured by the camera, and the other the prediction given by the scene model, based on a different image. The scene model should take into account

<sup>1</sup>We are not implying here that these assumptions are wrong or sub-optimal in the context in which the papers are written.

all processing steps that may cause a colour mismatch (for example white balancing), see [13, chap. 5] for a review of the current techniques.

To evaluate the performance of our approach, we chose the simplest scene model we could think of: a two dimensional scene pictured with a static camera. We use a laplacian distribution to characterise the inliers in order to accommodate for acquisition noise. We illustrate our model by computing an outlier rejection scheme that is expressed as a standard mixture of inliers versus outliers [14, 15]. The outlier rejection is carried out using a maximum a-posteriori criterium. Despite being trivial, the example shows two fundamental benefits of outlier modelling: First, the outliers can be detected and rejected without using any arbitrary threshold, and second, the system can handle an arbitrary percentage of outliers in the image. Furthermore, the papers shows that colour can be of great help in segmenting out the outlier. We chose a simple model because a more complex one would be less appropriate to distinguish the influence of colour from the influence of the other parameters on the results.

The paper starts by presenting the colour outlier model, which is the principal focus of this paper. It reviews first the case of greyscale images presented in [12] and then extends it to colour images. The outlier model is verified by comparing two images that have nothing in common and see whether the error distribution is the one predicted by the model. It is followed by a more practical example that shows a scene taken with a static camera and see if the model can predict the error distribution in this case. For this second experiment, a model for the inliers is required, as discussed in Section 3. To illustrate what the model does to the data, we plotted for each pixel the posterior probability of belonging to the outliers using greyscale visualization. The reader should keep in mind that this pseudo-segmentation is solely obtained using individual pixel colours, and that the goal is to show the behaviour of the outlier modelling and not to segment the image (which could be obtained by many other techniques in such a simple example).

## Remark

The model presented here deals with image differences at a pixel level. It has nothing in common with all the work in modelling the statistics of natural images (see [16] for an example).

## 2. The outlier model

Considering a photographic image pair, an error is obtained by subtracting the pixel values of the first image from the pixel values found at the same location in the second im-

age. In a practical computer vision application the image pair is the result of a complex transformation—called *registration*—of the original photographs. Here, we just consider the original photographs, taken with a static camera.

We believe that *outliers* generate an error pattern similar to the error generated by comparing two random regions of the scene. The idea is to characterise the outliers by computing their error distribution and then fit a mixture of inliers and outliers to the error data. The probability that a particular pixel belongs to the outliers can be computed using a Maximum A Posteriori (MAP) estimation.

Let  $I_0$  and  $I_1$  be two images that have nothing in common. We will assume that  $I_0(\mathbf{p})$  and  $I_1(\mathbf{p})$  are two independent stationary random vectors (containing each three colour components),  $\mathbf{p}$  being the position in the image.

### 2.1. Greyscale images

Let us first review the model in [12] for greyscale images. The error distribution is given by

$$P(r) = \sum_{\forall u} \Pr\{I_0(\mathbf{p}) = u, I_1(\mathbf{p}) = u - r\}, \quad (1)$$

where  $u$  denotes all possible values contained in the images, and  $r$  is the error value. Note that for greyscale images  $I_0(\mathbf{p})$  and  $I_1(\mathbf{p})$  are scalars. Since  $I_0(\mathbf{p})$  and  $I_1(\mathbf{p})$  are independent,

$$P(r) = \sum_{\forall u} \Pr\{I_0(\mathbf{p}) = u\} \Pr\{I_1(\mathbf{p}) = u - r\}. \quad (2)$$

By approximating the intensity probability distribution of the pixels in the image with the image histogram  $\mathbf{H}$ , normalised such that  $\sum_u H(u) = 1$ , Equation (2) becomes

$$P(r) = \sum_{\forall u} H_0(u)H_1(u - r). \quad (3)$$

In other words, the outlier distribution is approximated by the cross-correlation of the two image histograms.

### 2.2. Colour images

When matching colour images, the error distribution is given by

$$P(\mathbf{r}) = \sum_{\forall \mathbf{u}} \Pr\{I_0(\mathbf{p}) = \mathbf{u}\} \Pr\{I_1(\mathbf{p}) = \mathbf{u} - \mathbf{r}\}. \quad (4)$$

where  $\mathbf{u}$  denotes all possible values contained in the images, and  $\mathbf{r}$  is the error value. Note that both are vectors with 3 components,  $\mathbf{r} \triangleq [r_\alpha, r_\beta, r_\gamma]^T$ . By approximating the probability mass function of the pixels in the image with the image histogram  $\mathbf{H}$  (normalised such that  $\sum_{\mathbf{u}} H(\mathbf{u}) = 1$ ), Equation (4) becomes

$$P(\mathbf{r}) = \sum_{\forall \mathbf{u}} H_0(\mathbf{u})H_1(\mathbf{u} - \mathbf{r}). \quad (5)$$

This last equation is a three dimensional convolution operation. In order to simplify it, we additionally approximate the three dimensional convolution as the product of three unidimensional ones. This is equivalent to assume that the *errors* in the three colour components are independent<sup>2</sup>. In practice, the colour components are uncorrelated by mean of a simple *ad hoc* transformation. However, the assumption of error independence in the channels reduces drastically the computational burden and, as will be shown in the experimental evaluation of the model, it still provides very good results. Finally, the overall error probability is approximated by

$$\begin{aligned} P(\mathbf{r}) &= \sum_{\forall u_\alpha} H_{0\alpha}(u_\alpha)H_{1\alpha}(u_\alpha - r_\alpha) \cdot \\ &\cdot \sum_{\forall u_\beta} H_{0\beta}(u_\beta)H_{1\beta}(u_\beta - r_\beta) \cdot \\ &\cdot \sum_{\forall u_\gamma} H_{0\gamma}(u_\gamma)H_{1\gamma}(u_\gamma - r_\gamma). \end{aligned} \quad (6)$$

In other words, the outlier distribution is equal to product of the cross-correlations of the images histograms, computed independently on each colour component of the image pair. Equation (6) defines the *outlier Model*.

### 2.3. Experimental evaluation of the modelling

To evaluate the outlier model, we will compare two images that have nothing in common, compute the error histogram, and compare it to the predicted outlier distribution. We call the *error histogram* the histogram of the image obtained by subtracting one image from the other. In the derivation of the outlier model, we assumed that the *errors* in the three colour components are independent<sup>3</sup>. To match at best this assumption, we first decorrelate the colour components with a Singular Value Decomposition (SVD): Let  $\mathbf{I}_{RGB}$  be a  $n \times 3$  matrix containing all the pixels of both images—each column of  $\mathbf{I}_{RGB}$  contains the red, green and blue components of the pixels in the image. The SVD of  $\mathbf{I}_{RGB}$  is defined as

$$\mathbf{I}_{RGB} = \mathbf{U} \cdot \mathbf{S} \cdot \mathbf{V}^T, \quad (7)$$

where  $\mathbf{S}$  is a diagonal matrix and  $\mathbf{U}$  and  $\mathbf{V}$  are unitary matrices. The matrix  $\mathbf{V}$  is used to transform the pixel RGB values into a new colour space as follows

$$\mathbf{I}_{\alpha\beta\gamma}(\mathbf{p}) = \mathbf{V}^T \cdot \mathbf{I}_{RGB}(\mathbf{p})$$

where  $\mathbf{I}_{\alpha\beta\gamma}(\mathbf{p})$  are the colour components of a pixel expressed in the new colour space and  $\mathbf{I}_{RGB}(\mathbf{p})$  is the pixel

<sup>2</sup>This does not imply that the three *colour components* have to be independent.

<sup>3</sup>We did not assume that the colour channels are independent.

in the original image. After having transformed the images in  $\alpha\beta\gamma$  space, we apply the procedure of Section 3. Note that the  $\alpha\beta\gamma$  space has no perceptual meaning. However, with respect to colour image processing, a decorrelation transform as applied here can be interpreted as a transformation into opponent colour image encoding.

Figure 1 shows the fit of the outlier model to the error histogram generated by comparing a cave to a grapevine image. We can see that there is a very good match between the model and the measurements.

## 3. The inlier model

The inlier model is of little importance here, but is required to build a practical application example. For our situation, an inlier is—ideally—an object that appears at the exact same location in the image pair, and has the exact same pixel value. In practice, the pixel values may differ because of acquisition noise, and the image might not be perfectly aligned. We use a laplacian distribution with unknown variance to be able to handle these differences. Since the error values have 3 decorrelated dimensions, we will assume that the error distribution can be written as a separable laplacian:

$$\begin{aligned} P_{\mathcal{I}(\sigma_\alpha, \sigma_\beta, \sigma_\gamma)}(\mathbf{r}) &= \left( \frac{1}{2\sigma_\alpha} \cdot e^{-\frac{|r_\alpha|}{\sigma_\alpha}} \right) \cdot \left( \frac{1}{2\sigma_\beta} \cdot e^{-\frac{|r_\beta|}{\sigma_\beta}} \right) \cdot \\ &\cdot \left( \frac{1}{2\sigma_\gamma} \cdot e^{-\frac{|r_\gamma|}{\sigma_\gamma}} \right). \end{aligned} \quad (8)$$

See Figure 2 for an experimental evaluation of the inlier model. Note that the distribution in (8) is only an approximation of a multivariate laplacian [17].

## 4. The outlier mixture model

By combining the results of the two previous sections, we can build a new model for matching two images. This model is a mixture of the outlier model of Section 2 with the inlier model of Section 3. We can describe the error generated by matching two images as

$$P(\mathbf{r}) = \phi P_{\mathcal{I}}(\mathbf{r}) + (1 - \phi) P_{\mathcal{O}}(\mathbf{r}) \quad (9)$$

$$\mathbf{H}_m = \phi \mathbf{H}_{\mathcal{I}}(\sigma_\alpha, \sigma_\beta, \sigma_\gamma) + (1 - \phi) \mathbf{H}_{\mathcal{O}}, \quad (10)$$

where  $P_{\mathcal{I}}$  stands for the inlier probability density,  $P_{\mathcal{O}}$  is the outlier probability density,  $\phi$  is the proportion of inliers, and  $\mathbf{r}$  the error vector. Equation (10) is equivalent to (9) but using vector notation, where  $\mathbf{H}_m$  denotes the model histogram,  $\mathbf{H}_{\mathcal{I}}(\sigma_\alpha, \sigma_\beta, \sigma_\gamma)$  the inlier histogram and  $\mathbf{H}_{\mathcal{O}}$  the outlier histogram. To fit the outlier mixture model to the measurements, we compute the maximum likelihood

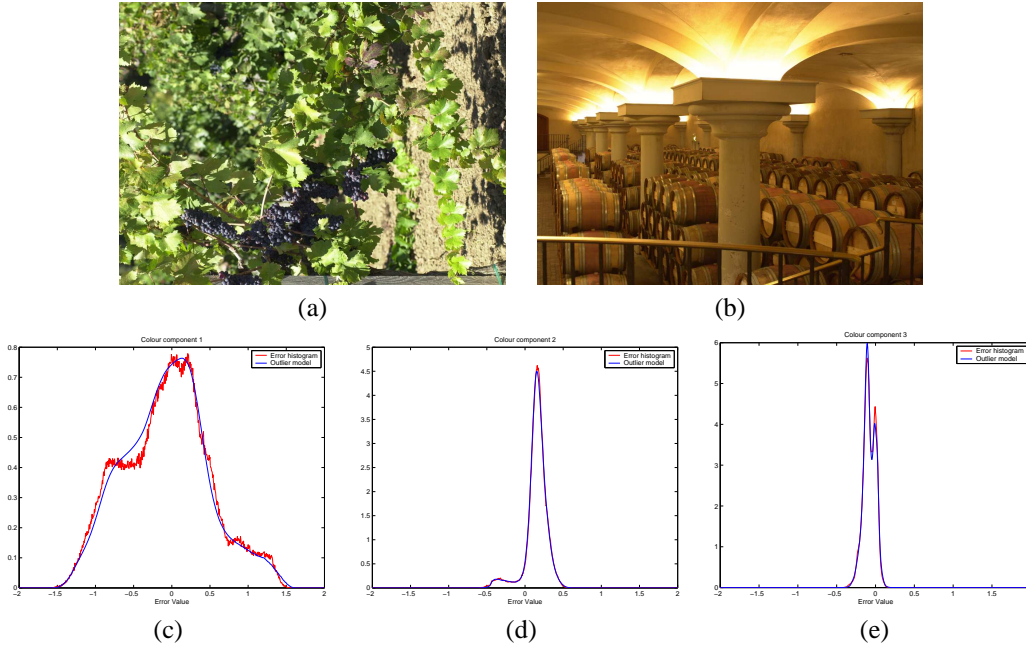


Figure 1: Illustration of the outlier model. The images in (a) and (b) are compared, and the error histogram is plotted versus the outlier model in (c), (d) and (e), for each colour channel. The outlier model exhibits a good match with the measured error histogram.

estimates of the parameters  $\{\phi, \sigma_\alpha, \sigma_\beta, \sigma_\gamma\}$

$$\begin{aligned} \{\hat{\sigma}_\alpha, \hat{\sigma}_\beta, \hat{\sigma}_\gamma, \hat{\phi}\} = \arg \max & \sum_{\forall r_\alpha} H_{e\alpha}(r_\alpha) \log [H_{m\alpha}(r_\alpha)] + \\ & + \sum_{\forall r_\beta} H_{e\beta}(r_\beta) \log [H_{m\beta}(r_\beta)] + \sum_{\forall r_\gamma} H_{e\gamma}(r_\gamma) \log [H_{m\gamma}(r_\gamma)], \end{aligned} \quad (11)$$

where  $H_e$  is the error histogram, i.e. the histogram of the image obtained by subtracting one image from the other.

Note that the outlier *distribution* depends only on the image histograms and not on the error. In other words, we can compute the outlier distribution without doing any image superpositions or any error computation. Nevertheless, to compute the outlier *proportion* (and the inlier standard deviations) we need to compute the error by superimposing the images.

## 5. Where is the outlier?

To get a better feeling of the outcome of the experiment, we propose to show a pseudo-segmentation of the image. This is done, once the inlier characteristics and the outlier proportion are known, by computing the probability that a particular pixel belongs to the inliers, using a MAP criterion:

tion:

$$\begin{aligned} P(\text{inlier} | \mathbf{r}) &= \frac{P(\text{inlier}) \cdot P(\mathbf{r} | \text{inlier})}{P(\mathbf{r})} \\ &= \frac{\phi H_{\mathcal{I}}(\sigma_\alpha, r_\alpha) H_{\mathcal{I}}(\sigma_\beta, r_\beta) H_{\mathcal{I}}(\sigma_\gamma, r_\gamma)}{\left( \phi H_{\mathcal{I}}(\sigma_\alpha, r_\alpha) H_{\mathcal{I}}(\sigma_\beta, r_\beta) H_{\mathcal{I}}(\sigma_\gamma, r_\gamma) + (1 - \phi) H_{\mathcal{O}}(r_\alpha) H_{\mathcal{O}}(r_\beta) H_{\mathcal{O}}(r_\gamma) \right)}, \end{aligned} \quad (12)$$

where  $H_{\mathcal{I}}$  denotes the inlier histogram,  $H_{\mathcal{O}}$  the outlier histogram and  $\phi$  is the proportion of inliers.  $\mathbf{r}$  is the error vector,  $\sigma$  is the inlier standard deviation and  $\alpha, \beta, \gamma$  denote the colour channel. This probability  $P$  is plotted as a grayscale value in Figure 4.

## 6. Results

The outlier mixture model is evaluated using the pictures of Figure 3 showing a fountain; in one of them a lady with an umbrella passes by—the lady being the outlier. The error histogram is compared to the outlier mixture model in Figure 5. The fit has been performed by using all three colour channels, as described in Section 4. Then, for each pixel, the posterior probability that it belongs to the outliers is shown in Figure 4. To show the impact of the colour information, this probability is first computed using only the first colour channel, then using two colour channels and finally using the whole colour information. We can clearly see that each colour component improves the

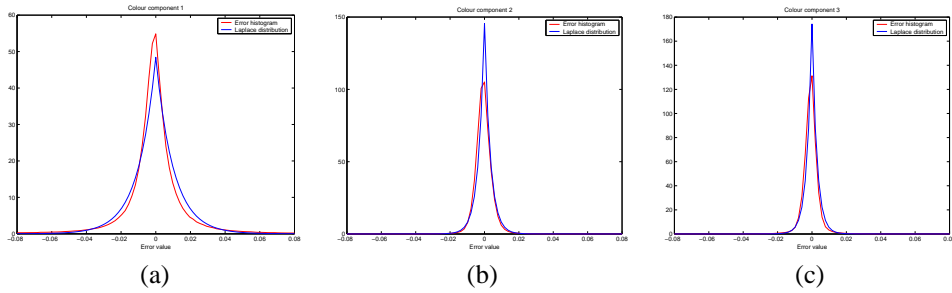


Figure 2: Error distribution computed on two consecutive images of the same scene, shown for the three colour channels. (a)  $\alpha$  channel. (b)  $\beta$  channel. (c)  $\gamma$  channel. Because of acquisition noise, and because of a slight displacement of the camera between the two views (less than one pixel in amplitude), the images are not exactly alike. The histogram of the error is compared to a laplacian distribution, showing that the laplacian is appropriate to model the inliers.

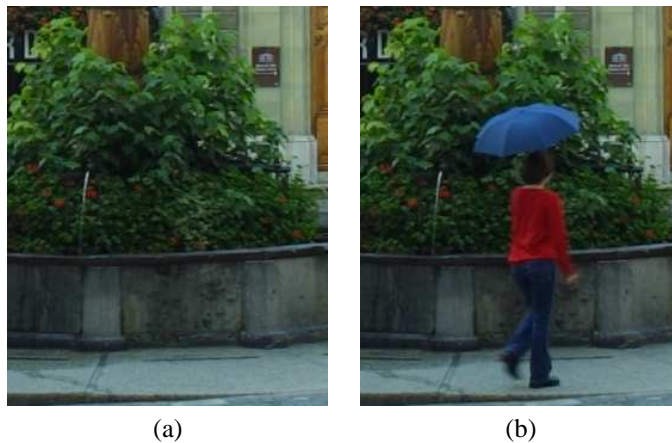


Figure 3: Images used to test the outlier mixture model. The lady with the umbrella forms the outliers of this picture pair.

result. The second example in Figure 6(c) shows the resistance of the model to a large numbers of outliers. Here, a standard outlier rejection scheme would have difficulties to perform a segmentation, as shown in Figure 6(d).

To test the mixture model in an even more extreme situation, we applied it to the images of Figure 1. The model found 99.8% of outliers.

For comparison, we ran the same test by assuming that the outlier distribution is uniform: the system concluded that there were no outliers in the image pair of Figure 1. It accommodated the error by using a huge variance for the inliers. Hence, our model is more robust and precise than the uniform distribution to model the outliers in image pairs.

### 6.1. Iterating the outlier mixture model

To get the outlier model of the mixture in Section 4, we used the histogram of the whole image to compute the outlier distribution. According to our outlier model, we

should have built our histograms using only the pixels that belonged to the outliers—that were unknown at that time. Now, by assuming that the outlier mask of Equation (12) is a better estimate of the outlier than the whole image, we can reiterate the algorithm by constructing the image histograms according to the outlier mask. This is done by making the contribution of each pixel to the histogram equal to its probability to belong to the outliers. We can then recompute the outlier distribution, and estimate the parameters of Equation 11 once again. The improvement in the outlier histograms is depicted in Figure 7. In practice only one iteration is needed, and the fewer outliers are in the scene, the more important this iteration step is. Nevertheless, we found that this iteration has only little impact on the overall outlier proportion and on the outlier mask.

### 6.2. Pushing the model to its limits

Until now, we assumed that the *errors* in the three uncorrelated colour channels were independent. To see the im-




Figure 4: Illustration of the outlier mixture model. The images in Figure 3 are compared. For each pixel, the probability that it belongs to the outliers is shown by its greyscale value. (a) is using only the first colour channel (b) uses the first two colour channels (c) uses all the colour information to compute the model parameters.




Figure 5: Comparison between the outlier mixture model and the error histogram of the images in Figure 3. The logarithm (in base 10) of the error histogram is plotted against the mixture model. (a) First colour component  $\alpha$ —the one with the largest singular value. (b) second colour component  $\beta$  (c) third colour component  $\gamma$ .

part of this assumption on the results, we recomputed the example of Figure 5 by using Equation (5) as the outlier model. Our tests show only a marginal difference between the separable approach (of Equation 6) and the present one. Figure 8 shows the difference between the two approaches.

The computation of the non separable model (in Equation 5) is very expensive. The computation of the outlier histogram requires  $2^{48}$  operations. Hence, to build the outlier histogram, we used a slightly different—but equivalent—technique, which consists in subtracting each pixel of the first image to every pixel of the next image. This requires  $3 \cdot n^2$  operations, where  $n$  is the number of pixel in the image. Unfortunately, it is still too complex to be of any practical use.

## 7. Conclusions

This paper presents a new approach for modelling outliers when comparing colour images on a pixel-to-pixel basis. It extends the work in [12] in two ways: The main contribution is the handling of colour, and a minor contribution is the iteration of the outlier model. Colour is shown to improve the segmentation of a scene in a trivial example. Iterating the model gives more reliable histograms when there are few outliers.

In general, the outlier model can be used with any motion and scene modelling scheme that compares the model to the incoming image at a pixel level. The inlier model should deliver a probability density function for the pixel-to-pixel error and should also allow for a certain amount of mis-registration in order to iteratively obtain a better model of the scene.

Our outlier model can be used in many ways: It can

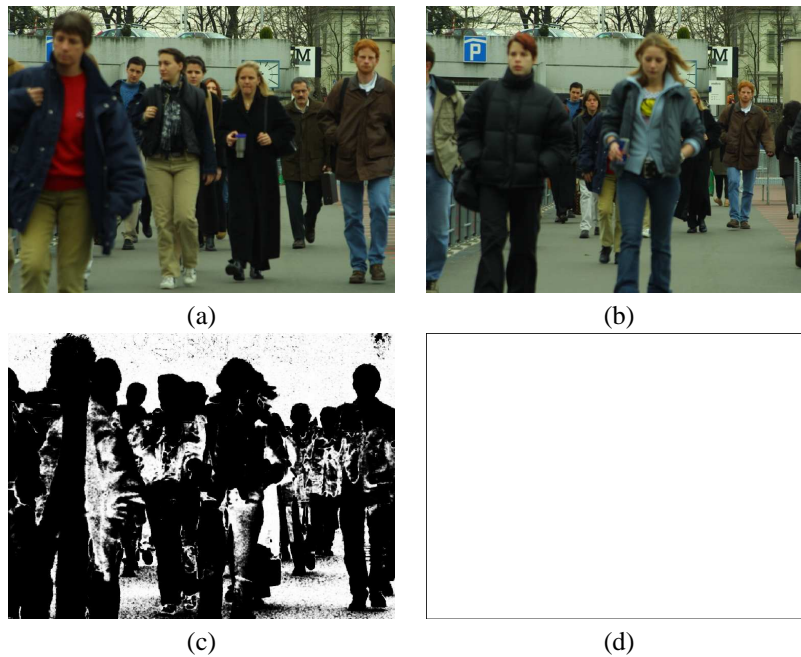


Figure 6: Illustration of the outlier mixture model. The images in Figure (a) and (b) are compared. For each pixel, the probability that it belongs to the outliers is shown by its greyscale value: in (c) using our outlier model and in (d) using a uniform distribution to model the outliers. (d) shows a total failure of the uniform model—the system tagged everything as inlier with a large variance. This shows that by using our model we can substantially improve the outlier resilience at the lowest level of a computer vision algorithm.

be applied to robust motion estimation by using the posterior probability to belong to the inliers as a weight factor in a motion estimation iteration, as was done in [12] for greyscale images. The overall percentage of outlying pixels can also be applied when performing change detection. In general, defining an outlier as a random superposition should enable the construction of more sophisticated algorithms for segmenting out outliers from images.

The model cannot be used in the cases where the inliers are not characterised by identical pixel values in the ideal case (after being processed by an upper level modelling scheme), such as in image retrieval applications, for example.

## References

- [1] P. J. Rousseeuw and A. M. Leroy, *Robust regression and outlier detection*. Applied probability and statistics, New York: John Wiley and Sons, Inc., 1987.
- [2] P. Huber, *Robust Statistics*. New York: Wiley-Interscience, 1981.
- [3] V. L. Brailovsky, “An approach to outlier detection based on bayesian probabilistic model,” in *Proceedings of ICPR*, pp. 70–74, 1996.
- [4] G. Danuser and M. Stricker, “Parametric model fitting: From inlier characterization to outlier detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 2, pp. 263–280, 1998.
- [5] V. Lakshamanan, V. DeBrunner, and R. Rabin, “Texture-based segmentation of satellite weather imagery,” in *ICIP*, vol. 2, pp. 732–735, 2000.
- [6] R. Dahyot, P. Charbonnier, and F. Heitz, “Robust visual recognition of colour images,” in *CVPR*, vol. 1, pp. 685–690, 2000.
- [7] D. Toth and T. Aach, “Improved minimum distance classification with gaussian outlier detection for industrial inspection,” in *Intl. conf. on Image Analysis and Processing*, pp. 584–588, 2001.
- [8] N. S. Netanyahu and I. Weiss, “Analytic outlier removal in line fitting,” in *Proceedings of the 12th IAPR International. Conference on Computer Vision and Image Processing.*, vol. 2B, pp. 406–408, 1994.
- [9] P. Schroeter, J.-M. Vesin, T. Langenberger, and R. Meuli, “Robust parameter estimation of intensity distributions for brain magnetic resonance images,”




Figure 7: Same as Figure 5 but computing a second iteration of the outlier modelling. We can see that the model better predicts the error histogram.




Figure 8: Comparison between the separable and non separable outlier model. (a) outlier mask obtained using 3 separate colour components (Same as Figure 5c). (b) outlier mask using the 3D outlier model of Equation (5). The three-dimensional model is slightly better at separating the outliers from the background, but its computation cost is prohibitive.

- IEEE transactions on Medical Imaging*, vol. 17, no. 2, pp. 172–186, 1998.
- [10] P. S. Torr, R. Szeliski, and P. Anandan, “An integrated bayesian approach to layer extraction from image sequences,” in *ICCV ’99*, 1999.
- [11] M. West, “Outlier models and prior distributions in bayesian linear regression,” *J.R. Statist. Soc. B*, vol. 46, no. 3, pp. 431–439, 1984.
- [12] D. Hasler, L. Sbaiz, S. Ssstrunk, and M. Vetterli, “Outlier modeling in image matching,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 3, pp. 301–315, 2003.
- [13] D. Hasler, *Perspectives on Panoramic Photography*. PhD thesis, Swiss Federal Institute of Technology (EPFL), 2001.
- [14] B. G. Lindsay, *Mixture Models: Theory, Geometry and Applications*, vol. 5 of *NSF-CBMS Regional Conference Series in Probability and Statistics*. 1995.
- [15] G. McLachlan and D. Peel, *Finite Mixture Models*. Wiley series in Probability and Statistics, New York: John Wiley and Sons, 2000.
- [16] J. Huang and D. Mumford, “Statistics of natural images and models,” in *CVPR*, 2000.
- [17] M. D. Ernst, “A multivariate generalized laplace distribution,” *Computational Statistics*, vol. 13, no. 2, pp. 227–232, 1998.