8—1 Curve Matching with Probabilistic Relaxation

Ying Shan
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399, USA
Email: yshan@microsoft.com

Zhengyou Zhang
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399, USA
Email: zhang@microsoft.com

Abstract

Reliable curve matching is a difficult problem but is required in many vision-based applications. It is particularly difficult when the edges in question are not limited to be on straight lines. Our main contribution in this paper is to propose a new algorithm for curve matching (including lines). The method is described within a probabilistic relaxation framework. Novel similarity-invariant unary and binary measurements suitable for curves are developed, and an additional measurement is introduced to model the uncertainty of the binary measurements, which is very important in computing the matching support from neighboring matches. Experiments with complex real scenes show that the rate of correct matching is higher than 98%.

1 Introduction

Information that can be used for curve matching falls into three categories, i.e., the geometrical constraint, the similarity between the curves, and the compatibility among the neighborhood matches. The most important geometric constraint between views is the epipolar constraint, which has been used in all curve-based stereo matching algorithms to reduce the search space. An early example exploiting this constraint was the PMF method [1, 2], where the epipolar lines are horizontal scan lines. The technique described in [3] also uses this constraint, but does not restrict the epipolar lines to be horizontal (the fundamental matrix is used to define the epipolar geometry).

The *similarity function* between two curves is usually defined in a high dimensional feature space. The features may include attributes such as the intensity in the neighborhood [3], orientation [4], and the local shape characteristics [5]. These are the *unary measurements*.

The *compatibility function* among neighboring matches is usually defined by the relationship between neighboring pairs of matches. The compatibility function is usually related to the local affinity or similarity assumption, and the *binary measurements*, defining the relationship, could be the angle and distance between a pair of neighboring curves [6, 7]. Relaxation techniques [8, 7] are useful methods that integrate the similarity function and the compatibility function to progressively reduce the matching ambiguity.

The goal of this paper is to attack the curve matching problem. Curves include both straight and non-straight edges. Many techniques exist for detecting edges. We detect edge points in sub-pixel accuracy by finding the zero-crossing between the integer pixels on the DOG (difference of Gaussian) image. Adjacent edge points are then linked together into connected edge chains (also called curves for simplicity). The linking process does not include any heuristic processing such as one-pixel gap filling or straight line fitting. An edge chain (curve) $\mathcal{C}(s)$ is represented as a linked list, and is parameterized by the index variable s.

Techniques for line matching have been successfully applied to scenes containing mainly planar surfaces [3]. These techniques, however, have two fundamental difficulties when applied to more general scenes. First, they are not suitable for scenes containing curves because the line model is insufficient to describe curves. Second, they are not suitable for scenes that are taken by a camera closed to the scene, where the local affinity or similarity assumption for long line segments is no more valid. Furthermore, the projection of straight lines in 3D onto images may not be straight anymore due to radial lens distortion.

Two central problems related to curve matching are the design of good unary and binary measurements, and the definition of appropriate similarity and compatibility functions. Previous work on curve matching such as [2, 9] gave good examples on the unary measurements and similarity function between curves, but were weak when dealing with binary measurements and compatibility functions. Indeed, their compatibility functions were usually computed from measure-

ments such as disparity [5] or disparity gradient [2], which are only suitable for the description of relationships between two pairs of points, and are not scale invariant. Another important issue to be considered is the uncertainty in the binary measurements. Obviously, the local affinity or similarity assumption is only valid within a limited area of the image. Binary measurements obtained from curves far away from each other have more uncertainty than those from nearby curves. This should be reflected in the computation of the compatibility function.

In this paper, we develop a curved-edge matching algorithm within a probabilistic relaxation framework similar to that in [7]. The framework described in [7] only deals with straight line segments. We adapt the framework to deal with curves and explicitly model the uncertainty in the binary measurements. Epipolar geometry is used to reduce the matching ambiguity. A piecewise linear model is used to approximate a curve with line segments. A set of binary measurements for line segments are proposed, which are similarity-invariant and measured in the same physical space. They are then integrated to provide a set of binary measurements and a compatibility function for the complete curve. A similarity-invariant unary measurement is also proposed.

1.1 Notation

Curve matching is modeled as a labeling problem from the object space (left image) to the model space (right image). There are N curves in the left image, and M in the right. The curves in the left image form the object space, denoted by $A = \{a_1, \ldots, a_N\}$. The curves in the right image form the model space, each labeled as ω_j. We wish to match the object to the model. We therefore assign to each object a_i a label θ_i, which may take as its value any of the $M+1$ model labels that form the set $\Omega = \{\omega_0, \omega_1, \ldots, \omega_M\}$, where ω_0 is the null label used to label curves for which no match in the right image is appropriate.

Four sets of indices are defined for convenience as $N_0 \equiv \{1, \ldots, N\}$, $N_i \equiv \{j \mid j \in N_0, j \neq i\}$, $M_0 \equiv \{0, \ldots, M\}$, and $M_i \equiv \{j \mid j \in M_0, j \neq i\}$. The unary measurement set is defined as $\mathcal{X} = \{\mathbf{x}_i \mid i \in N_0\}$, where $\mathbf{x}_i$ is the unary measurement vector for the object a_i. Unary measurements include intensity similarity, and will be addressed in section 2. The binary measurement set for the object a_i is defined as $\mathcal{A}_i = \{\mathbf{A}_{ij} \mid j \in N_i\}$, where $\mathbf{A}_{ij}$ is the vector of binary measurements between objects a_i and a_j. Binary measurements include distance between two curves, and will be addressed in section 3. A special measurement set $\mathcal{B}_i = \{\mathbf{B}_{ij} \mid j \in N_i\}$ is also defined, where $\mathbf{B}_{ij}$ is the uncertainty measurement to be defined later. To follow the traditional notation, we will use the uppercase P to denote the probability of an event, the lowercase p to denote the probability density function, and $\mathcal{N}_{\mathbf{v}}(\mu, \Sigma)$ to denote the Gaussian probability density function of a random vector $\mathbf{v}$ with the mean μ and the covariance matrix Σ.

1.2 A framework with uncertainty measurement

With the above notation, the matching problem becomes that of finding for object a_i a model label ω_{θ_i} with the highest probability given the measurements $\mathcal{X}$, $\mathcal{A}_i$, and $\mathcal{B}_i$:

$$\begin{aligned} & P(\theta_i = \omega_{\theta_i} \mid \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i) \\ = & \max_{\omega_\lambda \in \Omega} P(\theta_i = \omega_\lambda \mid \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i) \end{aligned} \tag{1}$$

For the convenience of the discussion that follows, we introduce the notation of the event set

$$\mathcal{L}^\lambda = \{\theta_i = \omega_\lambda\} \bigcup \{\theta_j = \omega_{\theta_j} \mid j \in N_i\}$$

to indicate that object a_i is labeled with a given label ω_λ while other objects can change their labels. Obviously, for $\mathcal{L}^{\theta_i}$, the ith event $\mathcal{L}_i^{\theta_i}$ is equivalent to event $\theta_i = \omega_{\theta_i}$. Using Bayes' formula and the theorem of total probability, we have

$$\begin{aligned} & p(\mathcal{L}_i^{\theta_i} \mid \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i) \\ = & \frac{p(\mathcal{L}_i^{\theta_i}, \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i)}{p(\mathcal{X}, \mathcal{A}_i, \mathcal{B}_i)} \\ = & \frac{\sum_{\{\omega_{\theta_j} \in \Omega, j \in N_i\}} p(\mathcal{L}^{\theta_i}, \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i)}{\sum_{\omega_\lambda \in \Omega} \sum_{\{\omega_{\theta_j} \in \Omega, j \in N_i\}} p(\mathcal{L}^\lambda, \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i)} \end{aligned} \tag{2}$$

By applying the product rule for $p(\mathcal{L}^\lambda, \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i)$, we have

$$\begin{aligned} & p(\mathcal{L}^\lambda, \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i) \\ = & p(\mathcal{X} \mid \mathcal{L}^\lambda, \mathcal{A}_i, \mathcal{B}_i) p(\mathcal{L}^\lambda, \mathcal{A}_i, \mathcal{B}_i) \\ = & p(\mathcal{X} \mid \mathcal{L}^\lambda) p(\mathcal{A}_i \mid \mathcal{L}^\lambda, \mathcal{B}_i) p(\mathcal{L}^\lambda, \mathcal{B}_i) \\ = & p(\mathcal{X} \mid \mathcal{L}^\lambda) P(\mathcal{L}^\lambda) p(\mathcal{A}_i \mid \mathcal{L}^\lambda, \mathcal{B}_i) p(\mathcal{B}_i) \\ = & P(\mathcal{L}^\lambda \mid \mathcal{X}) p(\mathcal{A}_i \mid \mathcal{L}^\lambda, \mathcal{B}_i) p(\mathcal{X}) p(\mathcal{B}_i) \end{aligned} \tag{3}$$

where the second equation holds because $\mathcal{X}$ (unary measurement set) does not depend on $\mathcal{A}_i$ (binary measurement set) and $\mathcal{B}_i$ (uncertainty of the binary measurement set), and the third equation holds because $\mathcal{L}^\lambda$ and $\mathcal{B}_i$ are independent. Since there is no knowledge about $\mathcal{X}$ and $\mathcal{B}_i$, $p(\mathcal{X})$ and $p(\mathcal{B}_i)$ are constants. Substituting (3) into (2), we then have

$$P(\mathcal{L}_i^{\theta_i} \mid \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i) =$$

$$\frac{\sum_{\{\omega_{\theta_j}\in\Omega, j\in N_i\}} P(\mathcal{L}^{\theta_i} \mid \mathcal{X}) p(\mathcal{A}_i \mid \mathcal{L}^{\theta_i}, \mathcal{B}_i)}{\sum_{\omega_\lambda\in\Omega} \sum_{\{\omega_{\theta_j}\in\Omega, j\in N_i\}} P(\mathcal{L}^{\lambda} \mid \mathcal{X}) p(\mathcal{A}_i \mid \mathcal{L}^{\lambda}, \mathcal{B}_i)} \quad (4)$$

Assuming that the events in the $\mathcal{L}^\lambda$ are independent, we have

$$\begin{aligned} & P(\mathcal{L}^\lambda \mid \mathcal{X}) \\ &= P(\theta_i = \omega_\lambda \mid \mathbf{x}_i) \prod_{j\in N_i} P(\theta_j = \omega_{\theta_j} \mid \mathbf{x}_j) \end{aligned}$$

Since $\mathbf{A}_{ij}$ only depends on the ith and jth events in $\mathcal{L}^\lambda$ and $\mathbf{B}_{ij}$, and since $\mathbf{A}_{ij}$s are independent from each other, we have

$$\begin{aligned} & p(\mathcal{A}_i \mid \mathcal{L}^\lambda, \mathcal{B}_i) \\ &= \prod_{j\in N_i} p(\mathbf{A}_{ij} \mid \theta_i = \omega_\lambda, \theta_j = \omega_{\theta_j}, \mathbf{B}_{ij}) \end{aligned}$$

With the following simple notation

$$\begin{aligned} & \mathbf{P}_j^\lambda = \\ & P(\theta_j = \omega_{\theta_j} \mid \mathbf{x}_j) p(\mathbf{A}_{ij} \mid \theta_i = \omega_\lambda, \theta_j = \omega_{\theta_j}, \mathbf{B}_{ij}) \end{aligned}$$

we then have

$$\begin{aligned} & \sum_{\{\omega_{\theta_j}\in\Omega, j\in N_i\}} P(\mathcal{L}^\lambda \mid \mathcal{X}) p(\mathcal{A}_i \mid \mathcal{L}^\lambda, \mathcal{B}_i) \\ &= P(\theta_i = \omega_\lambda \mid \mathbf{x}_i) \sum_{\omega_{\theta_1}\in\Omega} \mathbf{P}_1^\lambda \cdots \sum_{\omega_{\theta_N}\in\Omega} \mathbf{P}_N^\lambda \\ &= P(\theta_i = \omega_\lambda \mid \mathbf{x}_i) \prod_{j\in N_i} \sum_{\omega_{\theta_j}\in\Omega} \mathbf{P}_j^\lambda \quad (5) \end{aligned}$$

Substituting (5) into (4) leads immediately to

$$\begin{aligned} & P(\theta_i = \omega_{\theta_i} \mid \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i) \\ &= \frac{P(\theta_i = \omega_{\theta_i} \mid \mathbf{x}_i) \mathbf{Q}(\theta_i = \omega_{\theta_i})}{\sum_{\omega_\lambda\in\Omega} P(\theta_i = \omega_\lambda \mid \mathbf{x}_i) \mathbf{Q}(\theta_i = \omega_\lambda)} \quad (6) \end{aligned}$$

where the *support function* is given by

$$\begin{aligned} \mathbf{Q}(\theta_i = \omega_\alpha) &= \prod_{j\in N_i} \sum_{\omega_\beta\in\Omega} \mathbf{P}_j^\beta \\ &= \prod_{j\in N_i} \sum_{\omega_\beta\in\Omega} P(\theta_j = \omega_\beta \mid \mathbf{x}_j) \\ & \quad p(\mathbf{A}_{ij} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \mathbf{B}_{ij}) \quad (7) \end{aligned}$$

where the first item in the summation is the similarity function, and the second is the compatibility function that takes into account explicitly the uncertainty measurement $\mathbf{B}_{ij}$.

Based on (6), we update the labeling probability in a relaxation scheme according to the following iterative equation

$$\begin{aligned} & P^{(n+1)}(\theta_i = \omega_{\theta_i}) \\ &= \frac{P^{(n)}(\theta_i = \omega_{\theta_i}) \mathbf{Q}^{(n)}(\theta_i = \omega_{\theta_i})}{\sum_{\omega_\lambda\in\Omega} P^{(n)}(\theta_i = \omega_\lambda) \mathbf{Q}^{(n)}(\theta_i = \omega_\lambda)} \quad (8) \end{aligned}$$

where n is the iteration number, and $P^{(0)}(\theta_i = \omega_{\theta_i}) = P(\theta_i = \omega_{\theta_i} \mid \mathbf{x}_i)$. The iteration process terminates if the change in the probabilities is less than a predefined small value or the maximum number of iterations is reached. The reader is referred to [7, 10, 8] for more details.

As has been mentioned in section 1, candidate curves in the right view for the current curve a_j in the left view are found by the epipolar geometry. This is equivalent to reducing the model label set for a_j from Ω to Ω_j in the above derivations. Since the size of Ω_j is usually much smaller than that of Ω, we increase considerably the computational efficiency of the relaxation process.

1.3 Adapted framework with combined measurements

In our case, the unary and binary measurements depend on information in both object and model space (i.e., left and right image). Therefore, we should consider measurements for combined object-model pairs. Let $\mathbf{x}_i^{(\alpha)} (\alpha \in M_0)$ be the combined unary measurement defined for the pair of the ith object and the αth model. The unary measurement of object a_i, $\mathbf{x}_i$, is then itself a set of combined unary measurements, i.e. $\{\mathbf{x}_i^{(\alpha)} \mid \alpha \in M_0\}$. Let $\mathbf{A}_{ij}^{(\alpha\beta)} (\alpha \in M_0, \beta \in M_0)$ be the combined binary measurement defined for two object-model pairs (i, α) and (j, β). Then, the binary measurement $\mathbf{A}_{ij}$ is itself a set of combined binary measurements, i.e., $\{\mathbf{A}_{ij}^{(\alpha\beta)} \mid \alpha \in M_0, \beta \in M_0\}$.

In order to adapt the framework for working with the combined measurements, consider the similarity function in (7) first. Since the event $\theta_j = \omega_\beta$ does not depend on the combined unary measurements other than the $\mathbf{x}_j^{(\beta)}$, we have

$$P(\theta_j = \omega_\beta \mid \{\mathbf{x}_j^{(\beta)} \mid \beta \in M_0\}) = P(\theta_j = \omega_\beta \mid \mathbf{x}_j^{(\beta)}) \quad (9)$$

Similarly, $A_{ij}^{(\alpha\beta)}$'s are independent of each other, and we have

$$\begin{aligned} & p(\mathbf{A}_{ij} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \mathbf{B}_{ij}) \\ &= p(\{\mathbf{A}_{ij}^{(kl)}\} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \mathbf{B}_{ij}) \end{aligned}$$

$$= \prod_{\substack{k \in M_0 \\ l \in M_0}} p(\mathbf{A}_{ij}^{(kl)} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \mathbf{B}_{ij})$$

For $\mathbf{A}_{ij}^{(kl)}$ which are not equal to $\mathbf{A}_{ij}^{(\alpha\beta)}$ (i.e., k and l are not matched with i and j), there is no information about how to compute the binary measurement $\mathbf{A}_{ij}^{(kl)}$. Hence, we assume the density function to be a uniform distribution, i.e., a constant, within its domain. In consequence, we have

$$\begin{aligned} & p(\mathbf{A}_{ij} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \mathbf{B}_{ij}) \\ &= \epsilon \; p(\mathbf{A}_{ij}^{(\alpha\beta)} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \mathbf{B}_{ij}) \end{aligned} \quad (10)$$

where $\{\mathbf{A}_{ij}^{(kl)}\}$ is the abbreviation of $\{\mathbf{A}_{ij}^{(kl)} \mid k \in M_0, l \in M_0\}$, and $\epsilon = \prod_{\{kl \neq \alpha\beta\}} p(\mathbf{A}_{ij}^{(kl)} \mid \mathbf{B}_{ij})$ is a constant. By substituting (10) and (9) into (7), (7) into (6), and eliminating the constant ϵ, (6) can be rewritten as

$$\begin{aligned} & P(\theta_i = \omega_{\theta_i} \mid \mathcal{X}, \mathcal{A}_i, \mathcal{B}_i) \\ &= \frac{P(\theta_i = \omega_{\theta_i} \mid \mathbf{x}_i^{(\theta_i)}) \mathbf{Q}(\theta_i = \omega_{\theta_i})}{\sum_{\omega_\lambda \in \Omega} P(\theta_i = \omega_\lambda \mid \mathbf{x}_i^{(\lambda)}) \mathbf{Q}(\theta_i = \omega_\lambda)} \end{aligned} \quad (11)$$

and (7) as

$$\begin{aligned} \mathbf{Q}(\theta_i = \omega_\alpha) &= \prod_{j \in N_i} \sum_{\omega_\beta \in \Omega} \mathbf{P}_j^\beta \\ &= \prod_{j \in N_i} \sum_{\omega_\beta \in \Omega} P(\theta_j = \omega_\beta \mid \mathbf{x}_j^{(\beta)}) \\ & \quad p(\mathbf{A}_{ij}^{(\alpha\beta)} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \mathbf{B}_{ij}) \end{aligned} \quad (12)$$

2 The combined unary measurement and the similarity function

In this section, a combined unary measurement invariant to similarity (scaled Euclidean) transformation is developed and the similarity function is computed based on it. This part of the algorithm is similar to the PMF [2] algorithm. However, a major change has been made to ensure that the score computation is less sensitive to the image transformation, as we detail below.

2.1 Seed score and curve score

Referring to Fig. 1, the seed score is computed as the correlation between these two corresponding neighborhoods, and is given by

$$S(k_L, l_R) = \sum_{\substack{\xi \in N(k_L) \\ \eta \in N(l_R)}} \frac{[I_L(\xi) - \bar{I}(k_L)]\,[I_R(\eta) - \bar{I}(l_R)]}{N_n\, \sigma(I_{k_L})\, \sigma(I_{l_R})} \quad (13)$$

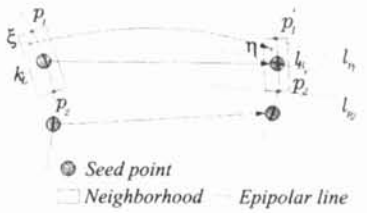

Figure 1: The neighborhood of the seed point in the left image and its correspondence in the right image

where $I_i(x)$ is the intensity value on a point x in the ith image, $N(x)$ is the neighborhood for seed point x, η is a point in the second image corresponding to point ξ in the first image according to the similarity transformation that can be computed from the pair of matched points [11], and N_n is the number of points in the neighborhood. The details of computing the mean ($\bar{I}(x)$) and the standard deviation ($\sigma(x)$) in $N(x)$ can be found in [12]. It is obvious that the correlation score thus computed is invariant under the similarity transformation.

The curve score $L(i_L, j_R)$ for the ith curve in the left image and the jth curve in the right image, which is the combined unary measurement in our context, is the average of all possible seed scores. That is,

$$L(i_L, j_R) = \frac{1}{N_s} \sum_{\substack{k_L \in Sd(i_L) \\ l_R \in Sd(j_R)}} S(k_L, l_R) \quad (14)$$

where N_s is total number of seed points on i_L, and $Sd(i_L)$ and $Sd(j_R)$ are the seed point sets of the left and right curves, respectively.

2.2 Similarity function

We can now compute the similarity function. The combined unary measurement $\mathbf{x}_i^{(\alpha)}$ in (11) is a scalar in our case, and will be denoted by $x_i^{(\alpha)}$. It is equal to $L(\theta_i, \omega_\alpha)$ computed above. According to Bayes' theorem, we have

$$\begin{aligned} & P(\theta_i = \omega_\alpha \mid x_i^{(\alpha)}) \\ &= \frac{p(x_i^{(\alpha)} \mid \theta_i = \omega_\alpha) \hat{P}(\theta_i = \omega_\alpha)}{\sum_{\omega_\lambda \in \Omega} p(x_i^{(\alpha)} \mid \theta_i = \omega_\lambda) \hat{P}(\theta_i = \omega_\lambda)} \end{aligned} \quad (15)$$

where $\hat{P}(\theta_i = \omega_\lambda)$ is the prior probability equal to a prefixed value ζ if $\lambda = 0$ (i.e., no match), and to $(1 - \zeta)/M_0$ otherwise. We assume that $p(x_i^{(\lambda)} \mid \theta_i = \omega_\lambda) = \mathcal{N}(\mu_l, \sigma_l)$. The value of μ_l and σ_l can be computed from the histogram of $x_i^{(\lambda)}$ using the initial matches. The initial matches could be obtained by se-

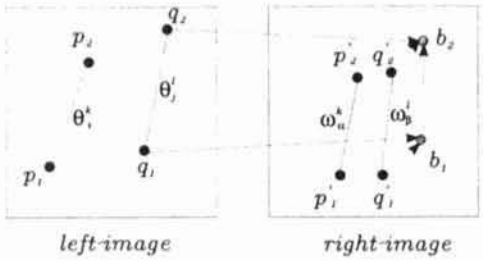

Figure 2: Configuration of two pairs of line segments in both images

lecting only the labeling with $x_i^{(\lambda)}$ that is the highest in both $\{x_i^{(\beta)} \mid \beta \in M_0\}$ and $\{x_{(j)}^{(\lambda)} \mid j \in N_0\}$.

The unary measurements are used in (15) to compute $P(\theta_i = \omega_\alpha \mid x_i^{(\alpha)})$. It is then used to initialize the relaxation scheme, that is: $P^{(0)}(\theta_i = \omega_\alpha) = P(\theta_i = \omega_\alpha \mid x_i^{(\alpha)})$.

It is interesting to note that since our unary measurement is similarity invariant, there is no need to introduce a motion related term as in [7] for the computation of the similarity function.

3 The combined binary measurement and the compatibility function

A set of combined binary measurements for the line segments is proposed. These measurements are grouped together in an optimal way to form the combined binary measurements for curves. The compatibility function for curves can then be computed. The uncertainty measurements for the combined binary measurements are also given in this section.

3.1 Combined binary measurements for line segments

As we mentioned before, the similarity transformation is a reasonable mapping from a neighborhood in one image to the corresponding neighborhood in the other. Consider two pairs of line segments illustrated in Fig. 2. The similarity transformation $\mathbf{x}' = s\mathbf{R}_\theta\mathbf{x}+\mathbf{t}$ can be computed from $(\mathbf{p}_1, \mathbf{p}_1')$ and $(\mathbf{p}_2, \mathbf{p}_2')$, as described in [11]. We then compute $\mathbf{b}_1 = s\mathbf{R}_\theta\mathbf{q}_1 + \mathbf{t}$, $\mathbf{b}_2 = s\mathbf{R}_\theta\mathbf{q}_2 + \mathbf{t}$, and form a random vector $\mathbf{z}$ as the follows

$$\mathbf{z} = [\mathbf{v}_1^{(x)}, \mathbf{v}_1^{(y)}, \mathbf{v}_2^{(x)}, \mathbf{v}_2^{(y)}]^\top \tag{16}$$

where $\mathbf{v}_1 = \mathbf{b}_1 - \mathbf{q}_1'$, and $\mathbf{v}_2 = \mathbf{b}_2 - \mathbf{q}_2'$. Ideally, if the local similarity transformation is valid and the point coordinates are noise-free, we have $\mathbf{z} = \mathbf{0}$. These conditions are not satisfied in practice, and we assume that the components of $\mathbf{z}$ are i.i.d, and have the same standard deviations σ, i.e., $\mathbf{z} = \mathcal{N}_\mathbf{z}(\mathbf{0}, \Sigma)$, where $\Sigma = diag(\{\sigma\})$.

3.2 Combined binary measurement for curves and the compatibility function

Now consider a pair of curves a_i and a_j. The numbers of line segments are k_i and k_j, respectively. For each line segment k on the shorter curve, the closest segment from the other curve is selected, and measurement $\mathbf{z}_k$ as defined in (16) is computed. The rationale of choosing the closest line segments is that the similarity transformation better applies to a small neighborhood than to a larger one. For the vector set $\{\mathbf{z}_j \mid j \in [1, \ldots, K]\}$, where $K = \min(k_i, k_j)$, consider the following measurement vector

$$\widetilde{\mathbf{z}} = \zeta_1\mathbf{z}_1 + \cdots + \zeta_K\mathbf{z}_K \tag{17}$$

where ζ_k's are coefficients and $\sum_{k=1}^K \zeta_k = 1$. It is obvious that $\widetilde{\mathbf{z}}$ is also a joint Gaussian with $\widetilde{\mathbf{z}} = \mathcal{N}_{\widetilde{\mathbf{z}}}(\mathbf{0}, \widetilde{\Sigma})$, where $\widetilde{\Sigma} = diag[\{\widetilde{\sigma}\}]$, and

$$\widetilde{\sigma}^2 = \zeta_1^2\sigma_1^2 + \cdots + \zeta_K^2\sigma_K^2 \tag{18}$$

We then select ζ_j to minimize (18), which yields

$$\zeta_i = [\sigma_i^2 \sum_{j=1}^K \sigma_j^{-2}]^{-1} \tag{19}$$

Equation (17) is an unbiased estimate of the measurement error with minimal variance. The compatibility function is then given by

$$p(\mathbf{A}_{ij}^{(\alpha\beta)} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \mathbf{B}_{ij}) = p(\widetilde{\mathbf{z}} \mid \theta_i = \omega_\alpha, \theta_j = \omega_\beta, \widetilde{\Sigma}) = \mathcal{N}_{\widetilde{\mathbf{z}}}(\mathbf{0}, \widetilde{\Sigma}) \tag{20}$$

where $\mathbf{A}_{ij}^{(\alpha\beta)} = \widetilde{\mathbf{z}}$ is the combined binary measurements, and $\mathbf{B}_{ij} = \widetilde{\Sigma}$ is the uncertainty measurements.

3.3 The distance between segments and the standard deviation

The standard deviation σ of vector $\mathbf{z}$ in (17) should be a function of the distance d between the line segments. This is because the local similarity assumption becomes weaker when d becomes larger. The function that we use is

$$\sigma(d) = \frac{\rho}{\sqrt{2\pi}}[(1-\gamma)(1-e^{-\frac{d^2}{\tau^2}}) + \gamma] \tag{21}$$

where $\rho = \min(W, H)$, W and H are respectively the image width and height, τ is a positive scalar that controls the range within which the contribution from neighboring segment is effective, and $\gamma = \sqrt{2\pi}\sigma^0/\rho$ where σ^0 is the desired standard deviation when the neighboring segment is very close to the segment under consideration.

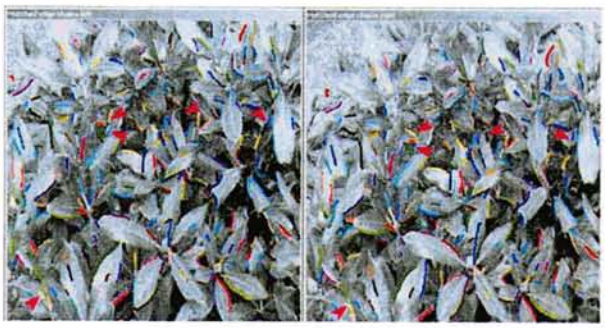

Figure 3: Matching results for a nature scene. 305 out of 310 matches are correctly matched and the rate is 98%. All bad matches are marked with pink arrows in the images for reader's convenience.

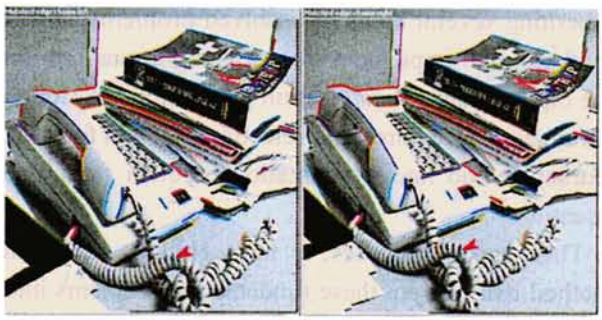

Figure 4: Matching results for an indoor scene taken by a camera closed to the scene. 107 out of 108 matches are correctly matched and the rate is above 99%. The bad match is marked with a pink arrow in the images for reader's convenience.

4 Experimental results

Fig. 3 is the result for an outdoor scene with many curves. The program found 310 pairs, 305 of which are correctly matched. Fig. 4 is the result for an indoor scene which is taken by a camera close to the scene. There is a significant perspective distortion between the left and the right image. The program found 108 matches, and only one is bad. Fig. 5 is another indoor scene extracted from an INRIA stereo sequence. Despite the significant camera motion between the two views, the program found 110 pairs, and only one is bad. The overall rate of correct matching is higher than 98%.

Figure 5: Matching results for an indoor scene. 109 out of 110 matches are correctly matched and the rate is above 99%. The bad match is marked with a pink arrow in the images for reader's convenience.

References

[1] Stephen B Pollard, John E. W. Mayhew, and John P Frisby. PMF: A stereo correspondence algorithm using a disparity gradient limit. In John E. W. Mayhew and John P. Frisby, editors, *3D model recognition from stereoscopic cues*, pages 11–25. MIT, 1991.

[2] Stephen B Pollard, John E. W. Mayhew, and John P Frisby. Implementation details of the PMF stereo algorithm. In John E. W. Mayhew and John P. Frisby, editors, *3D model recognition from stereoscopic cues*, pages 33–41. MIT, 1991.

[3] Cordelia Schmid and Andrew Zisserman. Automatic line matching across views. In *CVPR'97*, pages 666–671, 1997.

[4] G. Médioni and R. Nevatia. Segment-based stereo matching. *CVGIP*, 31:2–18, 1985.

[5] Nasser M. Nasrabadi. A stereo vision technique using curve-segments and relaxation matching. *IEEE PAMI*, 14(5):566–572, 1992.

[6] Z. Zhang and O. Faugeras. *3D Dynamic Scene Analysis*. Springer-Verlag, 1992.

[7] William J. Christmas, Josef Kittler, and Maria Petrou. Structural matching in computer vision using probabilistic relaxation. *IEEE PAMI*, 17(8):749–764, aug 1995.

[8] J. Kittler and E.R. Hancock. Combining evidence in probabilistic relaxation. *Int'l J. Pattern Recognition and Artificial Intelligent*, 3:29–51, 1989.

[9] Frank Candocia and Malek Adjouadi. A similarity measure for stereo feature matching. *IEEE Trans. IP*, 6(10):1460–1464, oct 1997.

[10] R.A. Hummel and S.W.Zucker. On the foundations of relaxation labeling process. *IEEE Trans. PAMI*, 5(3):267–286, may 1983.

[11] Ying Shan and Zhengyou Zhang. Corner guided curve matching and its application to scene reconstruction. In *CVPR'97*, pages 796–803, 2000.

[12] Gang Xu and Zhengyou Zhang. *Epipolar geometry in stereo, motion and object recognition*. Kluwer Academic Publishers, 1996.