

Robust Registration of 2D and 3D Point Sets

Andrew W. Fitzgibbon
Department of Engineering Science
University of Oxford.

Abstract

This paper introduces a new method of registering point sets. The registration error is directly minimized using general-purpose nonlinear optimization (the Levenberg-Marquardt algorithm). The surprising conclusion of the paper is that this technique is comparable in speed to the special-purpose ICP algorithm which is most commonly used for this task. Because the routine directly minimizes an energy function, it is easy to extend it to incorporate robust estimation via a Huber kernel, yielding a basin of convergence that is many times wider than existing techniques. Finally we introduce a data structure for the minimization based on the chamfer distance transform which yields an algorithm which is both faster and more robust than previously described methods.

1 Introduction

A common problem in computer vision is the registration of 2D and 3D point sets [1, 4, 6, 7, 19, 26]. Applications include the integration of range datasets [12, 23], and alignment of MRI/CAT scans[8, 20]. Typically, a cloud of point samples from the surface of an object is obtained from two or more points of view, in different reference frames. The task of *registration* is to place the data into a common reference frame by estimating the transformations between the datasets. What makes the problem difficult is that correspondences between the point sets are unknown *a-priori*. A popular approach to solving the problem is the class of algorithms based on the Iterated Closest Point (ICP) technique introduced in [1, 26]. ICP is attractive because of its simplicity and its performance. Although the initial estimate does need to be reasonably good, the algorithm converges relatively quickly.

This paper abandons one of the basic characteristics of ICP—its closed-form inner loop—and employs instead a standard iterative nonlinear optimizer, the Levenberg-Marquardt algorithm [17]. This approach, perhaps surprisingly, incurs no significant loss of speed, but allows the extension of ICP to use truly robust statistics, with a concomitant reduction of dependence on the initial estimate. In contrast, existing ways of introducing robustness [3, 22, 26] are often many times slower than our proposal.

The paper is in several sections. Section 2 defines the problem and the notation used in the rest of the paper. In section 3 we review existing work on point set registration, ICP and otherwise, and also summarize the key ideas in nonlinear optimization that will be required. Section 4 redefines registration as a nonlinear optimization problem. This approach is compared with traditional ICP in 2D and 3D experiments. The second half of the paper, beginning in Section 5 exploits the simplicity of the new approach to develop a robust error function which greatly increases the radius of convergence on supplied examples, again with no significant loss in speed.

n -D Transformation	p
2D Euclidean	3
2D Affine	6
2D Projective	8
3D Euclidean	6
3D Similarity	7
3D Affine	12
3D Projective	15

Notation summary:

N_m model points \mathbf{m}_j
 N_d data points \mathbf{d}_j
 Data point \longleftrightarrow closest model point: $i \longleftrightarrow \phi(i)$.
 Transformation $T(\mathbf{a}; \mathbf{x})$, parameters $\mathbf{a} \in \mathbb{R}^p$
 Error $\Delta_i(\mathbf{a}) = \mathbf{m}_{\phi(i)} - T(\mathbf{a}; \mathbf{d}_i)$
 Residuals $\mathbf{e} \in \mathbb{R}^{N_d}$, elements $E_i = \epsilon^2(|\Delta_i(\mathbf{a})|)$
 Jacobian \mathcal{J} , size $N_d \times p$, ij^{th} element = $\frac{\partial e_i}{\partial a_j}$

Table 1: (a) Transformations commonly occurring in registration problems. (b) Notation summary for the paper.

In all the above examples, closest-point computations are based on explicit Delaunay simplicization of the point set, allowing $O(\log n)$ closest-point computations. In section 6, we show how both the new and traditional procedures can be modified to use a fast lookup based on the distance transform. This speeds up both algorithms significantly, but the proposed technique benefits more, resulting in an algorithm which is faster and more accurate than traditional ICP, with a wider basin of convergence.

More generally, the message of the paper is that specialized “home-grown” strategies for function minimization (of which ICP is one example) do not necessarily outperform more general—but more sophisticated—nonlinear optimization algorithms.

2 Problem statement and definitions

The paper deals exclusively with the two-frame case, although multiple-frame approaches [7, 19] should immediately benefit. We are given two sets of points in \mathbb{R}^n , which for convenience we shall denote by *model* and *data*, with their elements being denoted by $\{\mathbf{m}_i\}_{i=1}^{N_m}$ and $\{\mathbf{d}_i\}_{i=1}^{N_d}$. The task of registration is to determine the parameters of a transformation T which, when applied to the data points, best aligns model and data. The parameters of T are represented by a p -vector \mathbf{a} . Common transformations and corresponding values of p are listed in table 1.

For 2D registration (see figure 1), $p = 3$ and the parameters of T are rotation angle θ and translation vector (t_x, t_y) . Collecting the parameters into a parameter vector $\mathbf{a} = [\theta, t_x, t_y]$, we define

$$T_{2D}(\mathbf{a}; \mathbf{x}) = T(\theta, t_x, t_y; \mathbf{x}) = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \mathbf{x} + \begin{pmatrix} t_x \\ t_y \end{pmatrix} \quad \text{for } \mathbf{x} \in \mathbb{R}^2$$

Alignment is measured by an error function $\epsilon^2(|\mathbf{x}|)$, and a typical choice is to define

$$\epsilon^2(|\mathbf{x}|) = \|\mathbf{x}\|^2$$

In order to measure alignment, we require that correspondence between the model and data points is specified. This correspondence is denoted by the function $\phi(i)$ which selects, for each data point, the corresponding model point. In order to cope with data points for which

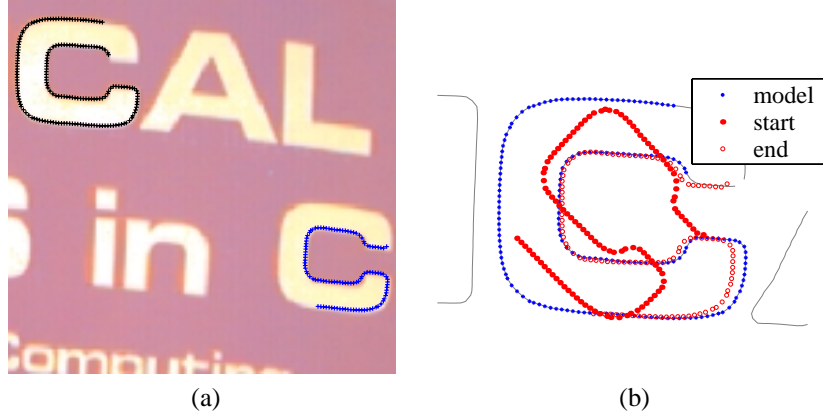


Figure 1: (a) 2D curves to be registered. The curves are a subset of the output of an edge detector applied to the image. Synthetic rotations and translations of the lower-right ‘C’ are registered to the upper left one. (b) Convergence of LMICP with Lorentzian kernel $\epsilon(r) = \log(1 + r^2)$ from a 40° rotation.

no correspondent is found, we also introduce *weights* w_i , which are set to zero for points with no match, and one otherwise. Thus, the error to be minimized is

$$E(\mathbf{a}, \phi) = \sum_{i=1}^{N_d} w_i \epsilon^2(\| \mathbf{m}_{\phi(i)} - T(\mathbf{a}; \mathbf{d}_i) \|) \quad (1)$$

In general, the function ϕ is considered part of the minimization process: in ICP-like algorithms $\phi(i)$ is chosen as the point which minimizes the distance between model and data, yielding the error function

$$E(\mathbf{a}) = \sum_{i=1}^{N_d} w_i \min_j \epsilon^2(\| \mathbf{m}_j - T(\mathbf{a}; \mathbf{d}_i) \|) \quad (2)$$

and finally the estimate of the optimal registration is given by minimizing over \mathbf{a} :

$$\hat{\mathbf{a}} = \operatorname{argmin}_{\mathbf{a}} \sum_{i=1}^{N_d} w_i \min_j \epsilon^2(\| \mathbf{m}_j - T(\mathbf{a}; \mathbf{d}_i) \|)$$

2.1 The ICP algorithm

In its simplest form, the ICP algorithm iterates two steps. Beginning with an initial estimate of the registration parameters, \mathbf{a}_0 , the algorithm forms a sequence of estimates \mathbf{a}_k which progressively reduce the error $E(\mathbf{a})$. Each iteration of the algorithm comprises the following two steps, labelled C and T :

C . Compute *correspondences*, ϕ : Set

$$\phi(i) = \operatorname{argmin}_{j \in \{1 \dots N_m\}} \epsilon^2(\| \mathbf{m}_j - T(\mathbf{a}_k; \mathbf{d}_i) \|) \quad i = 1 \dots N_d$$

so that $\mathbf{m}_{\phi(i)}$ is the closest model point to the datum \mathbf{d}_i transformed by the current estimate \mathbf{a}_k .

- T.* Update *transformation, a*: Set $\mathbf{a}_{k+1} = \operatorname{argmin}_{\mathbf{a}} \sum_{i=1}^{N_d} \epsilon^2(\|\mathbf{m}_{\phi(i)} - T(\mathbf{a}; \mathbf{d}_i)\|)$ In many common cases, step *T* can be performed in closed form, or via well understood, numerically stable, procedures such as singular value decomposition.

It is easy to see that both steps must reduce the error, and that the error is bounded below. Thus convergence to a local minimum is guaranteed. Furthermore, it is straightforward to discern a termination criterion: when the set of correspondences does not change in step *C*, the value of \mathbf{a}_{k+1} will be set equal to \mathbf{a}_k in the *T* step, so no further change is possible.

2.2 The Levenberg-Marquardt algorithm

The proposal of this paper is to directly minimize the model-data fitting error (2) via nonlinear minimization. The Levenberg-Marquardt (LM) algorithm is an optimization procedure which is particularly suited to functions such as E which are expressed as a sum of squared *residuals*, but alternative procedures such as conjugate gradients or even a pure Gauss-Newton algorithm could be used, and similar results would be expected. In this paper, only LM was tested.

The question that will immediately arise regarding the application of LM to the ICP problem is how the derivatives of E may be computed. Indeed the requirement for first derivatives might be seen as immediately disqualifying E from LM optimization, given the discrete minimization over j within the summation. However, it will be seen in section 6 that in fact the derivatives of E may be easily and efficiently obtained, and that they are—or can be made—smooth. For the moment, we shall assume they behave smoothly, and compute them via finite differencing [17, 11], at a cost of p extra function evaluations per inner loop. This means that each iteration's cost is increased by a factor of $1 + p$, for the simplest form of LM-ICP. However, in typical cases, where LM requires fewer iterations to achieve a certain accuracy, this factor is an upper bound. In many cases, the reduction in number of iterations will exceed the increase in per-iteration cost. Finally, it must be emphasised that an implementer concerned with speed should read section 6.

There now follows a derivation of the Levenberg-Marquardt algorithm. Readers who are familiar with its operation may wish to skim this description in order to synchronize notation, and proceed to the next section. The error function $E(\mathbf{a})$ can be written as the sum of N_d residuals as follows:

$$E(\mathbf{a}) = \sum_{i=1}^{N_d} E_i^2(\mathbf{a}), \quad E_i(\mathbf{a}) = \sqrt{w_i} \min_j \epsilon(\|\mathbf{m}_j - T(\mathbf{a}; \mathbf{d}_i)\|)$$

An important concept in the derivation of LM will be the vector of residuals $\mathbf{e}(\mathbf{a}) = \{E_i(\mathbf{a})\}_{i=1}^{N_d}$ so that $E(\mathbf{a}) = \|\mathbf{e}(\mathbf{a})\|^2$.

The Levenberg-Marquardt algorithm combines the gradient descent and Gauss-Newton approaches to function minimization. Using the notation above, the goal at each iteration is to choose an update to the current estimate \mathbf{a}_k , say \mathbf{x} , so that setting $\mathbf{a}_{k+1} = \mathbf{a}_k + \mathbf{x}$ reduces the error $E(\mathbf{a})$. Expanding $E(\mathbf{a} + \mathbf{x})$ around \mathbf{a} , we obtain

$$E(\mathbf{a} + \mathbf{x}) = E(\mathbf{a}) + (\nabla E(\mathbf{a}) \cdot \mathbf{x}) + \frac{1}{2!}((\nabla^2 E(\mathbf{a}) \cdot \mathbf{x}) \cdot \mathbf{x}) + \text{h.o.t.}$$

Expressing this in terms of \mathbf{e} , we have

$$\begin{aligned} E(\mathbf{a}) &= \mathbf{e}^\top \mathbf{e} \\ \nabla E(\mathbf{a}) &= 2(\nabla \mathbf{e})^\top \mathbf{e} \\ \nabla^2 E(\mathbf{a}) &= 2(\nabla^2 \mathbf{e})\mathbf{e} + 2(\nabla \mathbf{e})^\top \nabla \mathbf{e} \end{aligned}$$

We shall denote the $N_d \times p$ *Jacobian* matrix $\nabla \mathbf{e}$ by \mathcal{J} , with ij^{th} entry $J_{ij} = \frac{\partial E_i}{\partial a_j}$. Introducing the Gauss-Newton approximation [17], i.e. neglecting $(\nabla^2 \mathbf{e})\mathbf{e}$, we arrive at

$$E(\mathbf{a} + \mathbf{x}) \approx \mathbf{e}^\top \mathbf{e} + \mathbf{x}^\top \mathcal{J}^\top \mathbf{e} + \mathbf{x}^\top \mathcal{J}^\top \mathcal{J} \mathbf{x}$$

The task at each iteration is to determine a step \mathbf{x} which will minimize $E(\mathbf{a} + \mathbf{x})$. Using the approximation to E which we have just derived, we differentiate with respect to \mathbf{x} and equate with zero, yielding

$$\nabla_{\mathbf{x}} E(\mathbf{a} + \mathbf{x}) = \mathcal{J}^\top \mathbf{e} + \mathcal{J}^\top \mathcal{J} \mathbf{x} = 0$$

solving this equation for \mathbf{x} yields the Gauss-Newton update, and gives the algorithm for one iteration of Gauss-Newton ICP:

1. Compute the vector of residuals $\mathbf{e}(\mathbf{a}_k)$, and its $N_d \times p$ matrix of derivatives \mathcal{J} with respect to the components of \mathbf{a} . (For a 2D rigid-body transformation, \mathbf{a} has 3 components, and \mathcal{J} is $N_d \times 3$).
2. Compute the update $\mathbf{x} = -(\mathcal{J}^\top \mathcal{J})^{-1} \mathcal{J}^\top \mathbf{e}$.
3. Set $\mathbf{a}_{k+1} = \mathbf{a}_k + \mathbf{x}$.

Of course, the above strategy does not guarantee that the step taken will result in a reduced error at $E(\mathbf{a}_{k+1})$. Whether or not it does so depends on the accuracy of the second-order Taylor series expansion at \mathbf{a}_k , and on the validity of the Gauss-Newton approximation. However, it can be shown that when these approximations are good, as they tend to be when near the minimum, convergence is rapid and reliable.

By comparison, an accelerated gradient descent approach as used by some previous registration algorithms [1, 7, 18] is obtained by replacing step 2 with

2. Compute the update $\mathbf{x} = -\lambda^{-1} \mathcal{J}^\top \mathbf{e}$.

where the value of λ controls the distance travelled along the gradient direction. For small λ , the iteration moves a long way along the downhill direction; large λ implies a short step. In contrast to Gauss-Newton, gradient descent does guarantee to reduce E , providing λ is sufficiently large. However, its convergence near the optimum is dismally slow.

The Levenberg-Marquardt algorithm combines both updates in a relatively simple way in order to achieve good performance in all regions. Step 2 is replaced by

2. Compute the update $\mathbf{x} = -(\mathcal{J}^\top \mathcal{J} + \lambda \mathbf{I})^{-1} \mathcal{J}^\top \mathbf{e}$.

Now large λ corresponds to small, safe, gradient-descent steps, while small λ allows fast convergence near the minimum. The art of a good Levenberg-Marquardt implementation is in tuning λ after each iteration to ensure rapid progress even where the Gauss-Newton approximations are poor. Details of such strategies may be found in [14, 17].

We have therefore derived this paper's first proposal, the LM-ICP algorithm, summarised in Figure 2.

```

function  $\mathbf{a} = \mathbf{lmicp} ( \{\mathbf{m}_j\}_{j=1}^{N_m}, \{\mathbf{d}_i\}_{i=1}^{N_d}, \mathbf{a}_0 )$ 
Set  $\lambda$  to an initial value —see Press[17], p.684
Set  $\mathbf{a} = \mathbf{a}_0$ 
repeat
  Compute  $\mathbf{e}_k = \mathbf{e}(\mathbf{a})$  —One closest-point computation
  Compute  $\mathcal{J}$  — $p$  closest-point computations
  Modify  $\lambda$  until  $\mathbf{a}_k = \mathbf{a} - (\mathcal{J}^\top \mathcal{J} + \lambda \mathbb{I})^{-1} \mathcal{J}^\top \mathbf{e}_k$  reduces the error  $\|\mathbf{e}(\mathbf{a}_k)\|^2$ .
  —One or more closest-point computations

  Set  $\mathbf{a} = \mathbf{a}_k$ 
until  $\lambda$  is large —So only a small gradient descent step reduced the error.

```

Figure 2: The basic LM-ICP algorithm. No data structure is implied for the closest point computations, see §6 for a fast LM-ICP.

3 Previous work

Having defined the problem and thus introduced our notation, we are in a position to compare existing work on point-set registration. Besl [1] introduced the name ICP, and provided enough examples of the performance of the basic algorithm to ensure its enduring popularity.

Zhang [26] enhanced the basic technique by replacing the error function $\epsilon^2(|x|) = x^2$ by a robust kernel [10, 17]. Chen and Medioni [4] assume the model points \mathbf{m}_i can be supplied with surface normals \mathbf{n}_i , which allows the point-to-point distance to be replaced by a point-to-tangent plane. Both of these extensions confer improved convergence properties on the basic algorithm without a significant increase in computational cost.

The problem of multiple local minima has been addressed by using colour [13] and curvature [8] properties to improve the correspondence step—points are allowed to match based not just on proximity, but also on similarity of surface shape or texture. These extensions are as easily included in this paper’s algorithm as in the original, but are not investigated here as they would be expected to favour both old and new approaches equally. Extensions to non-Euclidean and non-rigid registration [8, 20], replace the rigid-body computation in the T step with more general parametrized transformations. Again, the modifications to LMICP which would incorporate these extensions are straightforward.

The extensions to multiple-view reconstructions [7, 19] necessitated the introduction of nonlinear optimization strategies, although these were essentially limited to variations on gradient descent. Grimson [9] explicitly minimizes $E(\mathbf{a})$ again using a nonlinear optimizer. However, the important distinction between these efforts and the current work is that they retain *lock-step*: the separation of the correspondence and update steps, and the concomitant sluggish convergence near the minimum.

The separation of correspondence and update remains true of EM-based approaches [15, 25, 5]. Although the correspondence is “soft”, it is still computed in a separate step, meaning that there is no opportunity to simultaneously adjust correspondence and transform parameters. In medical imaging problems [20, 25, 24], particularly the registration of MRI or CT scans, explicit minimization over registration parameters is commonplace, as the problems solved there (maximization of mutual information for example), do not admit a closed-form optimization step. However, there the problems do not generally allow for the computation of derivatives, so simple adaptation to 3D point-set registration would not yield the improvements

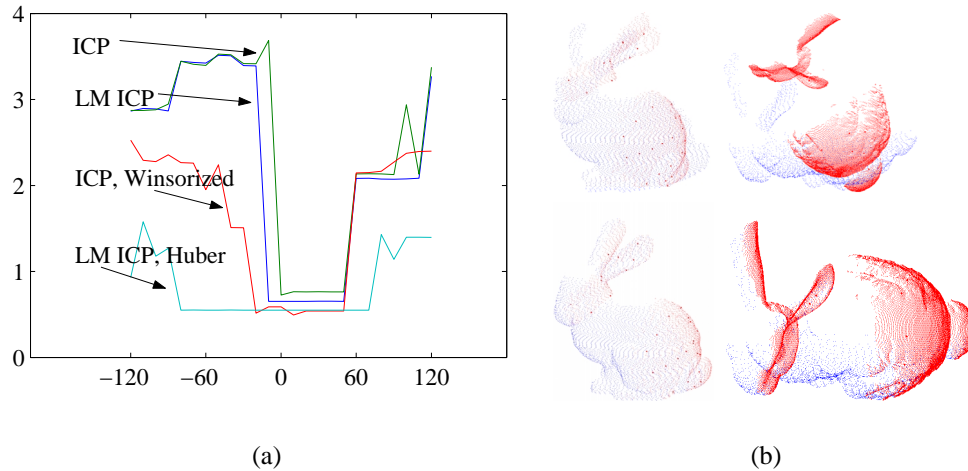


Figure 3: (a) 2D Basins of convergence. Algorithms are initialised at one-degree intervals between $\pm 120^\circ$ of the true solution, and the value of the minimum is plotted as a function of initial guess. LM-Huber is significantly wider than all others, at a cost of a factor of 2 in number of iterations. (b) 3D Example. Top: Initial alignment for 3D registration. Bottom: LM-Huber optimum.

shown in this paper.

The closest work to that reported here is the paper of Champleboux *et. al.* [3], who use Levenberg-Marquardt and the distance transform. However, the work assumes no outliers in the data, reverting to an explicit lock-step in order to reject gross errors. However, for the non-robust case, part of this paper’s contribution is the comparison of the Champleboux and Besl approaches with the result that the former is found to be significantly faster and more accurate.

4 Experiments: LM-ICP

In this section we compare the simplest case of point registration: minimizing the distance between two pointsets where the data are almost entirely a subset of the model. This is the case to which the simplest forms of ICP and LM-ICP apply. In this simple case, LM derivatives are computed using finite differences, which means each LM iteration is a factor of p more expensive than the ICP iteration. However, in §6, this cost will disappear, so that iteration counts are indeed the correct abscissae in this section.

The first experiment investigates the registration of 2D curves under Euclidean transformations. Figure 1 shows a section of an image of a book cover, with overlaid edge-detected curves. The task is to register two curves. The curves to be registered were chosen from different parts of the image in order to obtain a realistic variation in arc-length and sampling artefacts. The curves were subjected to synthetic rotations in order to test the algorithms over a range of initial conditions. Results are displayed graphically in figure 3. The summary of these results is that LM-ICP has a slightly larger basin of convergence, and can find an op-

timum with slightly reduced error. The Levenberg-Marquardt algorithm requires 50% fewer iterations on average, but unless speedups such as those in §6 are employed, this advantage will be lost in the computation of finite-difference derivatives. Even in this case, it might be thought surprising that the general-purpose Levenberg-Marquardt approach is even comparable with ICP’s closed-form inner optimization, and this is really the key message of this paper.

The extension of the algorithm to 3D is extremely straightforward, and an example is shown in figure 3b. Frames bun000 and bun090 of the Stanford “bunny” dataset [23] were presented to LM-Huber (see §5), with the identity transformation as initial guess. ICP failed to converge, while LM-Huber produced the visually correct solution in the figure.

5 Robust estimation

Many attempts have been made to widen the basin of convergence of the ICP algorithm, and these largely amount to introducing robust estimation. This is difficult with standard ICP, as no closed-form robust estimate of the T step is known, so authors have used either a nonlinear or RANSAC-based estimator [26, 18, 22], or exclude (*Winsorise*) points with large errors at the C step [3]. With LM-ICP, it is trivial to modify the error function to include a robust kernel. One must take a little care to ensure that the Levenberg-Marquardt algorithm behaves well if the kernel chosen is not smooth, but without going into the details, the examples here use either of the following kernels

$$\text{Lorentzian: } \epsilon(r) = \log\left(1 + \frac{r^2}{\sigma}\right) \quad \text{or} \quad \text{Huber: } \epsilon(r) = \begin{cases} r^2 & r < \sigma \\ 2\sigma|r| - \sigma^2 & \text{otherwise} \end{cases}$$

LM-ICP using these kernels is compared against ICP with Winsorised residuals [3], as this is the most common way of robustifying ICP. Figure 3 shows that LMICP with the Huber kernel has a basin of convergence twice as large as that of Winsorised ICP.

6 Fast ICP using the distance transform

The distance transform [2] of a set of points $\mathcal{M} = \{\mathbf{m}_j\}_{j=1}^{N_m}$ is defined as $D(\mathbf{x}) = \min_j \|\mathbf{m}_j - \mathbf{x}\|$. Algorithms exist for its computation on a discrete grid [2], which are extremely efficient, taking time which is a small constant number of machine instructions times the number of points plus the resolution of the grid. In 3D, the signed distance transform is the preeminent data structure for merging 3D models [12], meaning that it is readily available for the registration step.

We can analogously define the ϵ -distance transform D_ϵ by

$$D_\epsilon(\mathbf{x}) = \min_j \epsilon^2(\|\mathbf{m}_j - \mathbf{x}\|) \tag{3}$$

and if the mapping $\|\mathbf{x}\| \mapsto \epsilon^2(\|\mathbf{x}\|)$ is monotonic, we obtain that $D_\epsilon(\mathbf{x}) = \epsilon^2(\|D(\mathbf{x})\|)$, so that existing algorithms to compute D may be used to compute D_ϵ , without requiring knowledge of the form of ϵ .




Figure 4: Distance transform $D(x, y)$ and derivatives. Precomputing D and $\nabla_{\mathbf{x}}D$ allows fast computation of the error $E(\mathbf{a})$ and its derivatives $\nabla_{\mathbf{a}}E$ within the Levenberg-Marquardt iteration.

Now, the relationship between the distance transform and the registration problem is easily noted. Combining equations (3) and (2) we obtain

$$E(\mathbf{a}) = \sum_{i=1}^{N_d} w_i D_\epsilon(T(\mathbf{a}; \mathbf{d}_i))$$

Furthermore, we are now in a position to compute derivatives of E . We compute a discretization of D (see illustration, Figure 4a), from which we can immediately compute finite difference derivatives (figure 4a,b). Thus, we have ready access to $\nabla_{\mathbf{x}}D$, which remains constant throughout the minimization. Application of the chain rule yields $\nabla_{\mathbf{a}}E(\mathbf{a}) = \sum_{i=1}^{N_d} w_i \nabla_{\mathbf{x}}D_\epsilon(T(\mathbf{a}; \mathbf{d}_i)) \nabla_{\mathbf{a}}T(\mathbf{a}; \mathbf{d}_i)$.

Implementation of 2D LMICP bilinearly interpolating a discrete distance transform reduced the elapsed time for registration (average in fig 3) from 13 seconds to 0.37 seconds (MATLAB on a 650Mhz Pentium III). Traditional ICP can benefit from the distance transform by storing, at each point, the integer label of the closest point rather than the point itself. This makes each iteration of ICP almost exactly the same cost as those of LM-ICP—the time taken to compute T is similar to the cost of the Levenberg-Marquardt update. However, in this case, the Levenberg-Marquardt algorithm’s superior convergence means the overall runtime is reduced.

7 Discussion

We propose that point-set registration is better performed using a general-purpose nonlinear optimization procedure than via the popular ICP algorithm. The general-purpose routine is faster, and much simpler to program. Because it is simpler to program, it may be enhanced to incorporate robust estimation, without loss in speed. In contrast, standard ICP suffers a significant speed penalty when robust metrics are introduced [22]. In fact, standard ICP cannot minimize a robust kernel unless an iterative approach is used in the T step. This paper shows that pulling the iteration outside both C and T steps leads to a faster algorithm.

It can be shown, although it is omitted here, that LM-ICP must require, at worst, p times as many function evaluations as regular ICP. In practice this limit was never met.

The more general conclusion of this work is that specialized algorithms such as ICP are not always to be preferred to general-purpose techniques. This is true in this paper, and concurs with similar observations which have been made in neural network learning [16] and photogrammetry [21]. MATLAB source code for the algorithm is available from [27].

Acknowledgements. I would like to thank David Capel, Frederik Schaffalitzky and Andrew Zisserman for discussions relating to this work, and the Royal Society for its generous funding.

References

- [1] P. J. Besl and N. McKay. A method for registration of 3-D shapes. *IEEE PAMI*, 14(2):239–256, Mar 1992.
- [2] G. Borgefors. Hierarchical chamfer matching: A parametric edge matching algorithm. *IEEE PAMI*, 10(6):849–865, 1988.
- [3] G. Champleboux, S. Lavallée, R. Szeliski, and L. Brunie. From accurate range imaging sensor calibration to accurate model-based 3D object localization. In *Proc. CVPR*, pages 83–89, 1992.
- [4] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Image and Vision Computing*, 10(3), Apr 1992.
- [5] H. Chui and A. Rangarajan. A new algorithm for non-rigid point matching. In *Proc. CVPR*, pages 40–51, 2000.
- [6] S. J. Cunnington and A. J. Stoddart. N-view point set registration: A comparison. In *Proc. BMVC.*, pages 234–244, 1999.
- [7] D. Eggert, A. W. Fitzgibbon, and R. B. Fisher. Simultaneous registration of multiple range views satisfying global consistency constraints for use in reverse engineering. *CVIU*, 69(3):253–272, 1998.
- [8] J. Feldmar and N. Ayache. Rigid and affine registration of smooth surfaces using differential properties. In *Proc. ECCV*, 1994.
- [9] W. E. L. Grimson, T. Lozano-Pérez, W. Wells, G. Ettinger, S. White, and R. Kikinis. An automatic registration method for frameless stereotaxy, image-guided surgery, and enhanced reality visualization. In *Proc. CVPR*, pages 430–436, 1994.
- [10] J. P. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel. *Robust Statistics: An Approach Based on Influence Functions*. Wiley, New York, 1986.
- [11] R. I. Hartley. Euclidean reconstruction from uncalibrated views. In J. Mundy, A. Zisserman, and D. Forsyth, editors, *Applications of Invariance in Computer Vision*, LNCS 825, pages 237–256. Springer-Verlag, 1994.
- [12] A. Hilton, A. J. Stoddart, J. Illingworth, and T. Windeatt. Marching triangles: Range image fusion from complex object modelling. In *Proc. ICPR*, volume 2, pages 381–384, 1996.
- [13] A. Johnson and S. Kang. Registration and integration of textured 3D data. In *Proc. 3DIM '97, Ottawa*, pages 234–241, 1997.
- [14] J. E. Dennis Jr, D. M. Gay, and R. E. Welsch. Algorithm 573: NL2SOL—An adaptive nonlinear least-squares algorithm. *ACM TOMS*, 7:369–383, 1981.
- [15] B. Luo and E. R. Hancock. Matching point-sets using procrustes alignment and the EM algorithm. In *Proc. BMVC.*, pages 43–52, 1999.
- [16] T. Masters. *Advanced Algorithms for Neural Networks: A C++ Sourcebook*. Wiley, 1995.
- [17] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C (2nd Ed.)*. Cambridge University Press, 1992.
- [18] D. A. Simon, M. Hebert, and T. Kanade. Techniques for fast and accurate intra-surgical registration. *Journal of image guided surgery*, 1(1):17–29, 1995.
- [19] A. J. Stoddart and A. Hilton. Registration of multiple point sets. In *Proc. ICPR*, pages 40–44, 1996.
- [20] C. Studholme, D. Hill, and D. Hawkes. Automated 3D registration of truncated MR and CT images of the head. In *Proc. BMVC.*, pages 27–36, 1995.
- [21] W. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. Bundle adjustment: A modern synthesis. In W. Triggs, A. Zisserman, and R. Szeliski, editors, *Vision Algorithms: Theory and Practice*, LNCS. Springer Verlag, 2000.
- [22] E. Trucco, A. Fusiello, and V. Roberto. Robust motion and correspondence of noisy 3D point sets with missing data. *Pattern recognition letters*, 20:889–898, 1999.
- [23] G. Turk and M. Levoy. Zipped polygon meshes from range images. In *Proc. ACM SIGGRAPH*, pages 311–318, 1994.
- [24] P. Viola. *Alignment by Maximization of Mutual Information*. Aitr 1548, MIT, AI Lab, Jun 1995.
- [25] W. Wells. Statistical approaches to feature-based object recognition. *IJCV*, 21:63–98, 1997.
- [26] Z. Zhang. On local matching of free-form curves. In *Proc. BMVC.*, pages 347–356, 1992.
- [27] <http://www.robots.ox.ac.uk/~awf/lmicp>.