ROBUST SHAPE FITTING AND SEMANTIC ENRICHMENT

Torsten Ullrich and Dieter W. Fellner

Institute of Computer Graphics and Knowledge Visualization
Graz University of Technology
Inffeldgasse 16c, A-8010 Graz, Austria
t.ullrich@cgv.tugraz.at
http://www.cvg.tugraz.at

KEY WORDS: Algorithms, Classification, Geometry, Knowledge Base, Mathematics, Point Cloud, Reconstruction, Registration

ABSTRACT:

A robust fitting and reconstruction algorithm has to cope with two major problems. First of all it has to be able to deal with noisy input data and outliers. Furthermore it should be capable of handling multiple data set mixtures.
The decreasing exponential approach is robust towards outliers and multiple data set mixtures. It is able to fit a parametric model to a given point cloud. As parametric models use a description which may not only contain a generative shape but information about the inner structure of an object, the presented approach can enrich measured data with an ideal description. This technique offers a wide range of applications.

1 INTRODUCTION

Since the beginning of computer-aided design, reverse engineering has been an important field of research and it has become a viable method to record and archive digital descriptions of artifacts. The acquisition of an artifact starts measuring an object using 3D scanning technologies based on laser scanners, computer tomography, or photogrammetry. The measured data is the basis to re-engineer a high-level geometry description suitable for various uses. This reverse engineering process comprehends fitting, approximation and numerical optimization techniques.

Reverse engineering forms the link between recording techniques on the one hand and modeling and visualization on the other hand. This connection is well established for conventional modeling approaches based on Bézier patches and NURBS surfaces for example.

In contrast to conventional model descriptions, a procedural model does not store an object's geometry (control points, vertices, faces, etc.) but a sequence of operators and parameters in order to create a model instance.

All models with well organized structures and repetitive forms benefit from procedural model descriptions. In these cases generative modeling is superior to conventional approaches.

Its strength lies in a compact description which does not depend on the counter of primitives but on the model's complexity itself. As a result especially large scale models and scenes can be created efficiently which has promoted generative modeling within the last few years.

Another advantage of procedural modeling techniques is the included expert knowledge within an object description; e.g. classification schemes used in architecture, civil engineering, etc. can be mapped to procedures. For a specific object only its type and its instantiation parameters have to be identified. This identification process – needed in reverse engineering – can be done manually or user-assisted by the algorithm presented in this paper.

2 RELATED WORK

The creation of an accurate high-level description for a given model is known as reverse engineering. The tradition of reverse engineering is quite long and even a short overview about the main techniques would go beyond the scope of this article (Farin, 1990).

Existing fitting methods can be classified into two main groups. Algorithms resulting in a complete model description can be classified as *Complete Fitting* methods. Such algorithms result among other things in polyhedral surfaces (Hoppe et al., 1992), (Amenta et al., 1998), radial basis functions (Carr et al., 2001), constructive solid geometry (Rabbani and van den Heuvel, 2004), or subdivision surfaces (Cheng et al., 2004) – just to name a few.

The most important work related to the proposed algorithm is "Creating Generative Models from Range Images" by R. Ramamoorthi and J. Arvo (Ramamoorthi and Arvo, 1999). They also use generative models to fit point clouds. The main difference is, that they modify the generative description during the fitting process. Starting with the same generative description to fit a spoon as well as a banana does not allow to generate or preserve semantic data.

The second main group is herein after referred to as *Subpart Fitting*. These algorithms analyze an object and describe a part of the object's geometry. Depending on the algorithm's need to partition the input data in advance, all subpart fitting algorithms can be categorized into two subgroups.

The various random sample consensus (RANSAC)-based methods (Fischler and Bolles, 1981), (Wahl et al., 2005) are examples for *Subpart Fitting* without preceding segmentation. RANSAC algorithms compute free parameters of a subpart for a randomly selected adequate number of samples e.g. points of a point cloud. The samples then vote if they agree with the suggested parameters. This process is repeated until a sufficiently broad consensus is achieved. Advantages of this approach are its ability to ignore outliers without explicit handling and the fact that it can be extended to extract multiple instances in a data set.

The majority of subpart fitting algorithms need a preceding segmentation. A simple least squares fitting for example is unable to fit a plane to a point cloud, if the point cloud consists of point samples from more than one plane. Filtering tests (Benko et al., 2002), feature detection algorithm (Gumhold et al., 2001), (Pauly et al., 2003) structure (Vosselman and Sithole, 2004), (Hofer et al., 2005) and symmetry (Martinet et al., 2006) recognition techniques solve this problem. Reconstruction algorithms based on non-uniform rational b-splines (NURBS) (Wang et al., 2004), developable surface (Peternell, 2004) or least squares techniques (Shakarji, 1998) and many more operate on segmented data sets.

Except for a few algorithms such as RANSAC, which use statistical computations, many reverse engineering algorithms base on a numerical, global optimization. The objective of global optimization is to find the globally best solution of (possibly nonlinear) models, in the presence of multiple local optima. An overview on global optimization techniques can be found in (Pinter, 2004).

The fitting algorithm in this paper uses a numerical routine to minimize an error function. The minimization itself is performed by a modified differential evolution method (Storn and Price, 1997) combined with conjugate gradients techniques (Fletcher and Reeves, 1964).

3 GENERATIVE RECONSTRUCTION

The presented algorithm belongs to the category of subpart fitting algorithms without preceding segmentation. Algorithm which do need a preceding segmentation have the chicken and egg dilemma: without segmentation the fitting is complicated and vice versa.

Our method as well as RANSAC-based methods accept unprocessed point cloud data. RANSAC algorithms have to compute free parameters of a subpart for randomly selected samples. This inverse problem is hard to solve – even for simple geometric objects (cylinders (Beder and Förstner, 2006), cones, . . .).

Regarding the used error functions and metrics involved decreasing exponential fitting is similar to least squares techniques (LSQ), which minimizes the error function

$$f_{LSQ}(\vec{x}) = \sum_{i=1}^{n} d^2\left(\mathbf{M}(\vec{x}),\ \mathbf{p_i}\right) \stackrel{!}{=} \min_{\vec{x}}, \qquad (1)$$

whereas

- $\mathbf{P} = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$ denotes the given point cloud,
- $\mathbf{M}$ represents a generative model, and
- $\mathbf{M}(\vec{x})$ is a model instance using the parameters $\vec{x}$.
- The function d measures the Euclidean distance between two objects.

As the squares of the distances are used, outlying points have a disproportionate effect on the fit. Even worse, if a model shall be fitted to a data set, least squares methods implicitly assume that the entire set of data can be interpreted by one parameter vector. Therefore the fitting result of a plane fitted to a sampled cube is not one of the six cube sides, but one plane containing the cube's diagonal. In the context of statistics and computer vision generalized maximum likelihood estimators (Zhang, 1997) are of big importance in order to overcome the sensitivity of least squares estimates to outliers. The main idea is to introduce a weighting function ψ on top of the distance.

The Gaussian weighting function

$$\psi_{EXP}(x) = 1 - e^{-x^2/\sigma^2} \qquad (2)$$

is illustrated in Figure 1. The result is called *decreasing exponential* and its objective function to minimize is

$$\begin{aligned} f_{EXP}(\vec{x}) &= \sum_{i=1}^{n} \psi_{EXP}(d(\mathbf{M}(\vec{x}), \mathbf{p}_i)) && (3)\\ &= \sum_{i=1}^{n} 1 - e^{\frac{-1}{\sigma^2}\cdot d^2(\mathbf{M}(\vec{x}),\mathbf{p}_i)} \stackrel{!}{=} \min_{\vec{x}}. && (4)\end{aligned}$$

This minimization process can be interpreted geometrically. If σ tends to zero, the limit function is

$$\psi(x) = \begin{cases} 0, & x = 0 \\ 1, & \text{otherwise.} \end{cases} \qquad (5)$$

The globally best solution resp. the global minimum of this limit function is reached, if and only if the maximum number of points have distance zero to the model to fit; e.g. the model parameters will be determined so that most of the input data points belong to the model's surface. Therefore, the returned error value of $\lim_{\sigma\to 0} f_{EXP}(\vec{x})$ equals the number of points, which do not belong to the model's surface.

In practice σ should have a positive value. It should correspond to the noise level of the input data set. An adequate σ ensures that points whose distance to a model is $\pm\varepsilon$ have only a small contribution to the error function. The heuristic to select σ such that a point with distance d and noise ε will be weighted

$$\psi(d + \varepsilon) = 1/2,$$

yields good results.

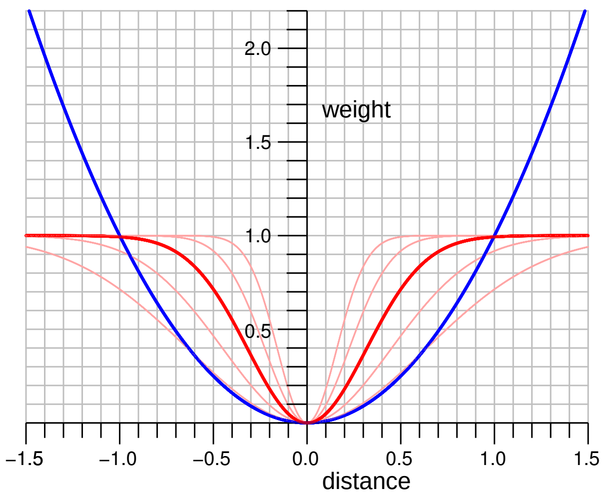

Figure 1: The least squares approach minimizes the sum of squared distances, $\psi_{LSQ}(x) = x^2$, which weights outlying points (blue). The function $\psi_{EXP}(x) = 1 - e^{-x^2/\sigma^2}$ used by the decreasing exponential approach (red) reduces this effect significantly. The choice of σ depends on the noise level of the input data.

The big advantage of decreasing exponential fitting is the fact that it only needs very few information about a model. It only needs a distance function measuring the Euclidean distance between an arbitrary point in 3D and the surface of a model. The usage of such a limited interface allows to combine this approximation technique with explicit as well as generative shape descriptions.

An explicit shape description follows the composition of primitives approach. It describes 3D objects and whole scenes as an agglomeration of elementary geometric objects: Points, triangles, NURBS-Patches, and many others. In the composition-of-primitives approach all elements are specified and listed individually. Without further information there is not much difference between a Greek temple and a statue. As James Kajiya points out, it's all a "matter of positioning the control points in the right places" (Snyder, 1992). But the inner logic of an object should be reflected in its construction, and its description of the result should reflect this process.

The importance of further information and semantic meta data becomes obvious in the context of mandatory services required by a digital library: markup, indexing, and retrieval. The promising approach of procedural and generative modeling is a key to solve this problem. Due to the naming of functions and algorithms as well as possible markup techniques, procedural model libraries are the perfect basis for digital library tasks. The advantages of procedural modeling arise from the generative approach. Expert knowledge about the inner structure of an object and about its composition can be formulated in a generative description. The enrichment of measured data with an ideal description enhances the range of potential applications – not only in the field of cultural heritage. A nominal/actual value comparison may indicate wear and tear effects as well as changes in style.

4 OPTIMIZATION

The vector $\vec{x}$ of free parameters of a Model $\mathbf{M}$ may have various geometric interpretations:

- A fixed sized, static object may vary its position and orientation. These six degrees of freedom can be combined to one vector. The result of this isometric optimization process is the best fit position and orientation of the static object according to a given point cloud.
- A parametric model may be described by a set of parameters; e.g. a sphere can be described by a center point and its radius. The result of this optimization process – called model optimization – is the sphere description that fits best to a given point cloud.

The fitting of a static 3D model to an acquired point cloud is an isometric optimization. The free parameters, which have to be optimized, form a bijective map that preserves distances – an isometry. The parameters of the map form a six-dimensional vector and may be interpreted among other things as three Euler angles and a three-dimensional translation.

Due to the fact that the weighting function ψ used by the decreasing exponential approach has an upper limit

$$\forall\, x \in \mathbb{R} \,:\, \psi(x) < 1$$

and converges relatively fast to one, the objective function (4) does not have to be evaluated completely during the optimization. Holding the point cloud in a regular grid structure of equally sized cubes (illustrated in Figure 2) allows to estimate ψ for all points in a cube by the number of points, if the cube is beyond a distance threshold. Thus the evaluation of the objective function $f(\vec{x})$ can be made efficiently; especially if the points are distributed uniformly all over the cubes, the evaluation of $f(\vec{x})$ only depends on the size of the model to fit – respectively on the number of cubes that have to be processed – and not on the number of all points.

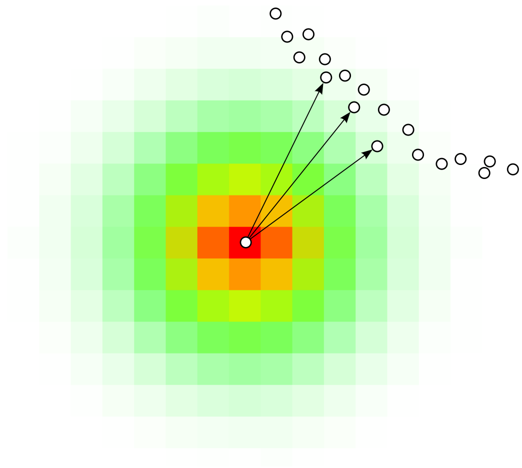

Figure 2: The storage of a point cloud in a regular grid structure of equally sized cubes allows a fast evaluation of a sum of ψ_{EXP}-weighted distances. For a given model only those cubes whose distance to the model is within a threshold may contain points with a weighted distance significantly smaller than one. These points have to be evaluated separately, the weighted distance of all other points can be approximated.

Beside static model fitting, decreasing exponential approaches can be used in the field of model optimization. In such a process not the position of a static model but its dynamic form defined by free parameters is of interest.

Especially so-called parametric objects can be identified easily. Only an Euclidean distance function d is needed. If the overall fitting process uses a derivation free minimization algorithm, no derivation of the distance function is needed. Therefore, d does not have to be in closed form. Due to the possibility to evaluate the distance between a point and a specific instance of a parametric object algorithmically, arbitrary parametric object descriptions are feasible.

As a matter of course this optimization task also profits from a cube grid-structured point cloud and the bounded weighting function ψ as mentioned above. For a given model only those cubes whose distance to the model is within a threshold may contain points with a weighted distance significantly smaller than one. While these point distances have to be calculated exactly, the sum of weighted distances of all other points can be approximated by the number of all other points. The complexity of evaluating $f(\vec{x})$ is therefore in most cases sub-linear in the number of points.

5 IMPLEMENTATION

As the optimization task using the decreasing exponential approach can be formulated as the most common form of global optimization: e.g. as a typical minimization problem, a wide range of optimization strategies solving this problem exists.

To use state-of-the-art optimization algorithms in a flexible environment the first implementation has been realized using

Maple^{TM} (Version 10). The implementation can be scaled down to the definition of the weighting function ψ and some calls of Maple's *minimize* function.

In addition to estimated problems of a slow execution due to script interpretation without run-time optimization, it reveals some numerical problems:

1. As the weighting function is almost constant beyond the origin, the gradients are extremely small in most cases. Many newton-like optimization routines have problems in handling such situations.
2. The convergence is even worse, if the initial values are not set adequately.

These problems have been addressed in an optimized, native implementation. To speed up the native implementation the point cloud is stored in a hash-mapped regular grid structure which allows a fast evaluation of the objective function using distance approximations based on clustering.

The minimization itself is performed by Differential Evolution (Storn and Price, 1997). Differential Evolution is a heuristic approach for minimizing nonlinear, non differentiable, continuous space functions. Based on evolutionary computing the algorithm does not use a single iteration vector but a population of vectors spread all over the domain. It does not need adequately chosen initial values.

Furthermore, by running several vectors simultaneously, superior parameter configurations can help other vectors escape local minima. Having found several solution candidates, a postprocessing for local corrections using conjugate gradients according to Fletcher/Reeves (Fletcher and Reeves, 1964) is done.

6 RESULTS & CONCLUSION

The proposed method of decreasing exponentials allows a generalized object fitting, which can easily applied to various model descriptions – only a point-to-object distance function is needed.

Figure 3: The CAD data set which has been used to generate a disturbed, uniformly sampled point cloud.

Beside the tests on fitting geometric primitives to point clouds the presented approach has also been tested on a CAD data set: The model of a Greek temple ($28m$ length, $18m$ width, $11m$ height) has been sampled uniformly. The resulting point cloud describes the object's surface and has been disturbed by a $5cm$ offset in normal direction (normally distributed with std-derivation of 1.0 and zero mean value).

The temple (see Figure 3) consists of 28 columns. Using a parametric, cylindrical column description with three free parameters (column position in 2D and its radius)[1] the decreasing exponential fitting is able to determine all columns with an accuracy at an average of $1cm$ ($= 0.03\%$ of the largest bounding box edge).

The second example object is the Theseus Temple in Vienna, Austria (see Figure 4). Is was build between 1820 and 1823 by Peter Nobile as a smaller imitation of the Theseion Temple in Athens.

Figure 4: Approximately one hundred photos of the Theseus Temple have been taken in order to generate a point cloud – the input data set of the first fitting process.

The network of Excellence in Processing Open Cultural Heritage (EPOCH) (http://www.epoch.eu) offers its partners the EPOCH 3D web service. The web service provides a photogrammetrical reconstruction and returns point clouds or textured meshes. In combination with the open source tool MeshLab (http://meshlab.sourceforge.net) 3D models have been generated from approximately one hundred photos taken with a consumer camera. The sequence includes only photos taken in an arc around the temple. Parts not visible in the photos (Figure 4) are not present in the point cloud. The resulting model consists of 2 million points. Due to the low quality of the photos the data set contains a lot of noise.

To detect the temple columns a row of cylinders has been used as generative model. The model parameters were

- the number of columns,
- the center point on the xy-plane of the first column,
- the center point on the xy-plane of the last column, and
- the cylinder radius.

[1] Fitting the column's height is non-applicable as a bigger height value always leads to columns with more points on it.

The parameter domain has been split. In this way it has been possible to find all of the column rows in parallel threads. A second generative model has been a stairway. The free parameters were

- starting point,
- direction of the stairway,
- riser height and tread length.

The number of steps and the width of the treads are considered to be infinite. These infinite structures have been cut along the bounding box of the relevant points near their surface.

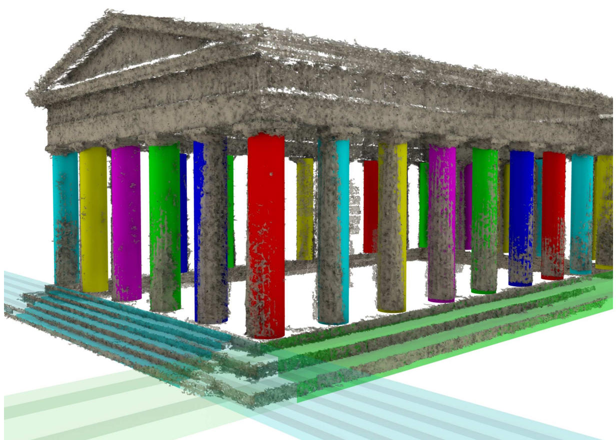

Figure 5: The presented algorithm is able to fit generative descriptions such as columns and stairways to the Theseus Temple point cloud. Even structures, which are not present in the input data set (e.g. the treads), are realized as they are part of the generative description. The visualization has been blended with a low resolution set of the input model.

The most important feature demonstrated in this example is the ability to recognize structures which are only partly present in the input data set (see Figure 5). The point cloud does neither contain the columns' surfaces turned away from the camera nor the stairways' treads.

In combination with optimization strategies to avoid being trapped in local minima the decreasing exponential fitting approach produces reasonable results:

- In contrast to least squares approaches it is able to manage outliers and compositions of multiple data sets.
- Compared to RANSAC algorithms, it can easily be generalized to fit all kinds of objects without having to interpret a set of points and without having to know how many points are needed to define a certain object.
- Furthermore, the decreasing exponential approach is a minimization process based on a error function whose differentiability does only depend on a distance function, which allows to benefit from a wide variety of well-studied, numerical methods.

Further studies will therefore investigate the best combination of numerical methods to solve the decreasing exponential minimization.

ACKNOWLEDGMENTS

The authors would like to thank Volker Settgast for valuable assistance and support – especially for the visualizations.

REFERENCES

Amenta, N., Bern, M. and Kamvysselis, M., 1998. A New Voronoi-Based Surface Reconstruction Algorithm. Computer Graphics 32, pp. 415–421.

Beder, C. and Förstner, W., 2006. Direct Solutions for Computing Cylinders from Minimal Sets of 3D Points. Proceedings of the European Conference on Computer Vision 2006 3951, pp. 135–146.

Benko, P., Kós, G., Várady, T., Andor, L. and Martin, R., 2002. Constrained Fitting in Reverse Engineering. Computer Aided Geometric Design 19(3), pp. 173 – 205.

Carr, J. C., Beatson, R. K., Cherrie, J. B., Mitchell, T. J., Fright, W. R., McCallum, B. C. and Evans, T. R., 2001. Reconstruction and Representation of 3D Objects with Radial Basis Functions. Proceedings of the 28th annual conference on Computer graphics and interactive techniques 28, pp. 67 – 76.

Cheng, K.-S. D., Wang, W., Qin, H., Wong, K.-Y. K., Yang, H. and Liu, Y., 2004. Fitting Subdivision Surfaces to Unorganized Point Data using SDM. Proceedings of 12th Pacific Conference on Computer Graphics and Applications 1, pp. 16 – 24.

Farin, G., 1990. Curves and Surfaces for Computer Aided Geometric Design. Academic Press Professional, Inc.

Fischler, M. A. and Bolles, R. C., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM 24(6), pp. 381–395.

Fletcher, R. and Reeves, C. M., 1964. Function minimization by conjugate gradients. The Computer Journal 7, pp. 149–154.

Gumhold, S., Wang, X. and MacLeod, R., 2001. Feature Extraction from Point Clouds. Proceedings of 10th Int. Meshing Roundtable 1, pp. 1–13.

Hofer, M., Odehnal, B., Pottmann, H., Steiner, T. and Wallner, J., 2005. 3D shape recognition and reconstruction based on line element geometry. Proceedings of the Tenth IEEE International Conference on Computer Vision 2, pp. 1532–1538.

Hoppe, H., DeRose, T., Duchamp, T., McDonald, J. and Stuetzle, W., 1992. Surface reconstruction from unorganized points. Proceedings of the 19th annual conference on Computer graphics and interactive techniques 1, pp. 71 – 78.

Martinet, A., Soler, C., Holzschuch, N. and Sillion, F., 2006. Accurate Detection of Symmetries in 3D Shapes. ACM Transactions on Graphics 25, pp. 439 – 464.

Pauly, M., Keiser, R. and Gross, M., 2003. Multi-scale Feature Extraction on Point-Sampled Surfaces. Computer Graphics Forum 22, pp. 281–289.

Peternell, M., 2004. Developable Surface Fitting to Point Clouds. Computer Aided Geometric Design 21, pp. 785–803.

Pinter, J. D., 2004. Global Optimization. From MathWorld – A Wolfram Web Resource, created by Eric W. Weisstein (http:// mathworld.wolfram.com/ Global Optimization.html).

Rabbani, T. and van den Heuvel, F., 2004. Methods for Fitting CSG Models to Point Clouds and their Comparison. Proceedings of Computer Graphics and Imaging - 2004 1, pp. 1–6.

Ramamoorthi, R. and Arvo, J., 1999. Creating Generative Models from Range Images. Proceedings of SIGGRAPH 99 1, pp. 195–204.

Shakarji, C. M., 1998. Least-Squares Fitting Algorithms of the NIST Algorithm Testing System. Journal of Research of the National Institute of Standards and Technology 106(6), pp. 633–641.

Snyder, J. M., 1992. Generative modeling for computer graphics and CAD. Academic Press Professional, Inc.

Storn, R. and Price, K., 1997. Differential Evolution: A simple and efficient heuristic for global optimization over continuous spaces. Journal of Global Optimization 11, pp. 341–359.

Vosselman, G. and Sithole, G., 2004. Recognizing structure in laser scanner point clouds. Proceedings of Conference on Laser scanners for Forest and Landscape assessment and instruments 1, pp. 1–6.

Wahl, R., Guthe, M. and Klein, R., 2005. Identifying Planes in Point-Clouds for Efficient Hybrid Rendering. Proceedings of The 13th Pacific Conference on Computer Graphics and Applications 1, pp. 1–8.

Wang, W., Pottmann, H. and Liu, Y., 2004. Fitting B-spline curves to point clouds by curvature-based squared distance minimization. ACM Transactions on Graphics 25, pp. 214 – 238.

Zhang, Z., 1997. Parameter Estimation Techniques: A Tutorial with Application to Conic Fitting. Image and Vision Computing Journal 15(1), pp. 59–76.