Computers in Industry 64 (2013) 1201–1208

Contents lists available at SciVerse ScienceDirect

Computers in Industry

journal homepage: www.elsevier.com/locate/compind

Recognising 3D products and sourcing part documentation with scanned data

Frank Mill [a,*], Andrew Sherlock [b], Qi Pan [c], Esme Anderson [a]

[a] University of Edinburgh, School of Engineering, King's Buildings, Edinburgh EH9 3JL, UK
[b] ShapeSpace Ltd, UK
[c] University of Cambridge, UK

ARTICLE INFO

Article history:
Received 10 September 2012
Received in revised form 19 March 2013
Accepted 28 March 2013
Available online 29 May 2013

Keywords:
Search by shape
Networks
3D scanning
PLM
CAD
PDM

ABSTRACT

Searching databases of 3D models is a crucial yet difficult problem that has been studied by the academic community for a considerable time. A useful and robust method for finding engineering parts remains difficult however. Previous work typically describes finding the best match in a single search. Work described in this paper uses scanning techniques allied to shape similarity measures to produce a system that successfully allows search by browsing. We also describe some new shape descriptors and methods of identifying and dealing with chirality. The technique is evaluated in the context of the part search applications. The use of the techniques is applied to large (80,000 + parts) databases of real world engineering components in use in automotive and aerospace companies. The methods employed are applicable to a wide range of scenarios in engineering, as well as the arts, archaeology, medicine and commerce.

1. Introduction

The scenario where a user holds an object in front of a computer and asks 'what can you tell me about this?' is where the work described here starts. More specifically it is aimed primarily at engineers who need to find data related to a product on a company network or intranet in situations where exact part names or numbers are not easily to hand. To perform such a search, or in fact a wide variety of searches, it is common to first formulate a query and then analyse the results of the query action and then perform refining further search.

In the system described in this paper fast efficient scanning is combined with a novel search engine that is based on part geometries and this allows the user to find files related to a physical hand held part. A number of input methods have evolved to formulate the query model and different strategies tried for the subsequent search, the best combination being dependent on the specific application.

The work described in this paper relates to applications in a wide area of product management situations, e.g. part information retrieval for design re-use, maintenance, marketing or user support. Rather than aiming to increase the fidelity of 3D scanned models the aim of the work was to enable fast and accurate identification of part data already stored in local, intranet or even internet based file stores. This may be in the form of a wide number of possible representations such as those found in CAD repositories, PLM systems or catalogues and these may be difficult to navigate due to the lack of exact part data.

The background of the work is based on previous studies developing systems to characterise shape [1–3] for applications in part classification and search. These systems calculate many key parameters of parts such as their surface area to volume ratio or their aspect ratio and these in turn are used to group or cluster part collections so that they can be easily searched. This enables rapid part retrieval without the need for exact part names or numbers. Shape based searches are useful for simply finding parts but they may also aid part database management by identifying duplicates or multiple similar shapes or they can be used to assist re-use of existing designs [4–6]. In general they can be used where downstream (from design) users require 3D part representations, e.g. manufacturing, maintenance or non-engineering functions such as marketing and customer support.

Searches for part data are often performed to find 3D models, drawings or other associated documentation such as manuals, analysis results or manufacturing plans.

Shape based search allows parts to be found when only approximate ideas of a part's shape are known. Fig. 1 shows some typical search strategies that are in use in the system described in

* Corresponding author. Tel.: +44 131 337 4779.
E-mail address: Frank.Mill@ed.ac.uk (F. Mill).

0166-3615/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.compind.2013.03.019

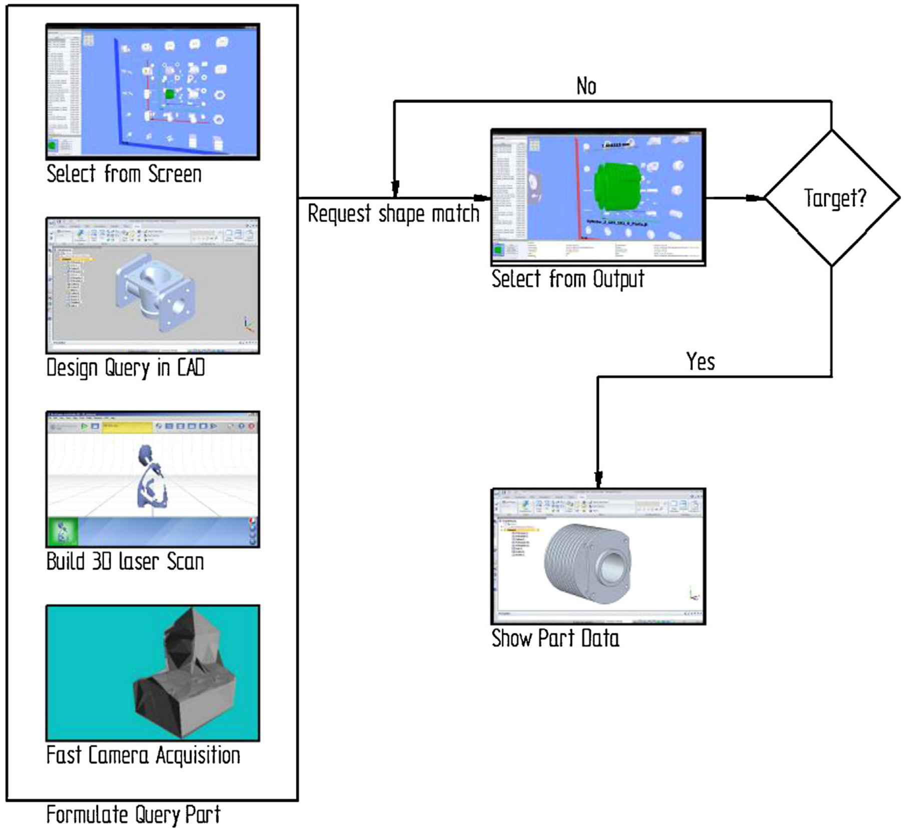

Fig. 1. Search flowchart showing how different query methods can be used to access a recursive search (browsing) process.

this paper. User searches can be carried out by picking from a bank of candidate on-screen shapes (select from screen) and subsequently selecting better matches from further screens of suggested examples or by designing individual 3D model queries (Design Query in CAD). Recently new methods have been developed whereby users employ rapid 3D scanning techniques (Build 3D Laser Scan and Fast Camera Acquisition) to build a query model from a physical part or part model.

Most research in 3D scanning use is aimed at producing high resolution scans and making high fidelity models from these, typical problems being interpreting actual geometry from the resultant point clouds. In the research presented here the emphasis is on fast scanning. The aim is to scan with just enough resolution to get a satisfactory query model and as a further development ordinary white light camera scans (e.g. from laptop webcams) have been used. The current system can search for part models and associated data that is stored in most 3D formats including STL, VRML, IGES, STEP and the majority of vendor specific formats for CAD and finite element analysis.

There are a great many applications for searching file systems and networks based on part shape, partly because the techniques allow access to data in a wide range of storage media including web or cloud based repositories.

In CAD environments where most of the users have considerable 3D modelling skills it often quickest to harness these and allow the user to design a query part from within his familiar CAD system. Even complex components can have very rough models constructed in seconds by experienced modellers. However, for users without CAD facilities, commonly the vast majority, this is not a favourable option and instead scanning methods may be employed. Although 3D laser scanned (and point probe) data have been successfully tried by the authors, the strategy suffers from the fact that in order to create useful scans specialised suitable equipment is needed and before the scan can be started is often necessary to spend time setting up a part. Research has been carried out by Pan et al. [7] that seeks instead to build 3D models direct from simple camera images that can be constructed using common devices such as those available on most computers and laptops. What is described in this paper is a system that optimises the general search process for a user. It is our view that much of the previous work described in the literature, e.g. that reviewed by Tangelder [2] works on the assumption that search will be performed on a general set of shapes. These are frequently general objects that are to be located and differentiated between (e.g. aeroplanes, cars, furniture and animals). This classification based view has seen the development of test part databases of general shapes which often exhibit shapes that are not relevant in manufacturing environments and which often have subsequent problems that are rare in the real world. It is rare for example to find a CAD designed engineering part that is not designed on a major x, y or z axis.

2. ProFORMA

The ProFORMA system takes solely a live video feed from a webcam as input, and contrary to many other reconstruction systems, aims to build a coarse 3D reconstruction for immediate use, rather than an accurate 3D reconstruction for later use. This makes it ideally suited to performing as the front-end for reconstructing a 3D model of a query object in a search environment, where it is desirable to get immediate results. Additionally, the ProFORMA system is designed to be used with the camera in a fixed position, with the query object being rotated in front of it. This has the natural benefit of being able to segment the object of interest from the background, something which is still an issue for systems which involve the inverse scenario of a moving

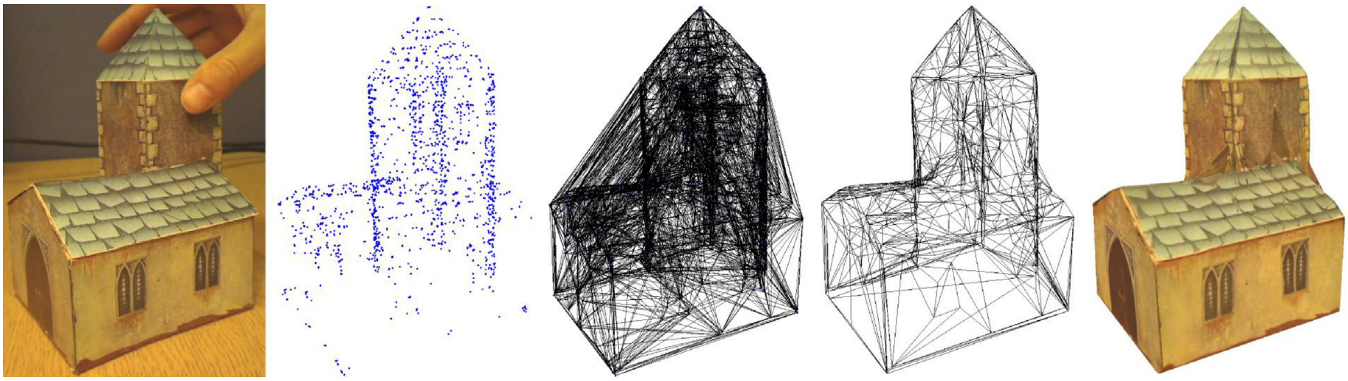

Fig. 2. The ProForma scanning system, reprinted with permission from (Pan et al. [7]). A video showing the scan development can be seen at: http://www.youtube.com/watch?v=vEOmzjImsVc.

camera and stationary object. Fig. 2 shows the sequence of activities used to generate a 3D model.

The model build process for the system takes around 1 min to complete and feedback is given to the user throughout the process so that he can see what has been done so far and what areas need further scanning. Rapid construction of the model is enabled by using a novel probabilistic tetrahedron carving algorithm which uses the visibility of observed features to quickly create a surface model of the object. Parts of the object that are occluded from the camera are shown in red so that the user can take corrective action to complete the views. In a search environment it is not always necessary to complete the scan however as the user only needs to build a model that is 'good enough' for use as a query. No assumptions are made about the object or its shape, however, some texturing is necessary to provide known points. For metallic objects such as those of interest in this paper this meant marking the object with pens but temporary stickers or labels can also be used.

Use of the system starts with the user showing the object to the camera and the software begins tracking the object. A two threaded keyframe based system is used, as described by [8], whereby separate threads are used for tracking and for reconstruction. Smooth continuous motion of the object is best, with the video sequence providing small distances between points both temporally and spatially which is used for tracking but which provides little 3D information. The tracking thread consists of three trackers, a robust point tracker that follows transient features with 3D location from frame to frame and which is robust to large motions, a second tracker which suffers from less drift, and a 2D tracking function.

The reconstruction thread creates a rendered 3D model from a list of landmarks, keyframes and keyframe camera poses. Because

Table 1
Typical metrics used in generation of shape signature.

Informal descriptor	Definition	Key
Compactness	SA/PA	CHA – convex hull surface area
ConvexHullCompactness	SA/CHA	CHCoG – convex hull centre of gravity or centre of area (assuming constant density)
Crinkliness	PA/PV	CHV – convex hull surface volume
Packing	SV/PV	Dmax – maximum identifiable dimension
AspectRatio0	Lp1/Lp3	Dmin – minimum identifiable dimension
AspectRatio1	Lp1/Lp2	FC – number of facets in triangulated mesh
AspectRatio2	Lp2/Lp3	L1 – shortest distance to a point wrt first principal axis
XYAspectRatio	Lx/Ly	L2 – shortest distance to a point wrt second principal axis
XZAspectRatio	Lx/Lz	L3 – shortest distance to a point wrt third principal axis
YZAspectRatio	Ly/Lz	Ld
SurfaceArea	PA	Lp1 – part length along first principal axis
ConvexHullSurfaceArea	CHA	Lp2 – part length along second principal axis
Volume	PV	Lp3 – part length along third principal axis
ConvexHullVolume	CHV	Lx – part length along *X* axis
DiagonalLength	Ld	Ly – part length along *Y* axis
SmallestDim	Dmin	Lz – part length along *Z* axis
MiddleDim	(Dmax – Dmin/2) + Dmin	PA = part surface area
LargestDim	Dmax	PCoG – part centre of gravity or centre of area (assuming constant density)
XDim	Lx	PV = part volume
YDim	Ly	SA = surface area of sphere with SV = PV (or CHV)
ZDim	Lz	SV = volume of sphere that minimally bounds part
CentreOfAreaRadius	PCoG	
Principal Moment of Inertia0	Sum(L1^2)	
CentreOfAreaRadiusConvex Hull	CHCoG	
Principal Moment of Inertia2	Sum(L2^2)	
Principal Moment of Inertia1	Sum(L3^2)	
Spikeness0	Sum(L1^4)/(L1^2)	
Spikeness1	Sum(L2^4)/(L3^2)	
Spikeness2	Sum(L3^4)/(L3^2)	
FaceCount	FC	

it builds full 3D meshes this method does not require solving the problem of generating 2D views in order to match them with 2D images as described in Refs. [9–11]. As is shown in Fig. 2 a point cloud is created which is then converted into a mesh through a Delaunay tetrahedralisation. Tetrahedra are then carved away based on visibility and probabilistic carving algorithm is used to smooth the resulting surfaces of the model. Finally, textures are added to the model. For the purposes of search however the carved tetrahedral mesh is sufficient to calculate parameters that characterise the object.

On completion of the scan phase the built model is read into the ShapeSpace package for analysis. This involves calculating a number of characteristics of the sample part and using these as a query to search through the part network(s) that are thought to contain the target part. There are multiple methods that have been developed by the research community to characterise shapes [2,3] and these have been used to judge the similarity between parts in databases. Some methods rely on recognising features and on the distribution of these. Whilst these methods appear to be good for fast general purpose searches they do require that features can be identified and they are criticised for being insensitive to feature location within models.

A common alternative to feature methods makes use of spatial maps or functions that typically describe spherical harmonic or a wide variety of methods that are mathematically similar. In general these apply spheres of decreasing size around a voxelised representation of a part and measure the proportion that is on or inside the surface for a given radius. They thus produce signatures that can be compared but they do not, in general, work well with mechanical features such as small threaded holes e.g. skeleton models and other 3D graph representations (e.g. topological graphs, which are similar to feature graphs) are also used as reduced-data models and these can then be compared, but in general use these methods are less sensitive than those previously discussed, especially for typical engineering parts.

The above, and other methods, are typically used for attempts to find closest matching parts. These strategies have had some success, however, they are ultimately limited because they are usually evaluated against some concept of how good they are at recognizing similarity. Since there can be no standard definition of what similarity actually is then there can be no technique that is superior to others except in a practical sense of how well it meets the users' expectations in a particular application. In different contexts users will often have a different concept of what is meant by similar, i.e. similar in what way?

In the approach used by the authors a flexible method is adopted which uses multiple methods of shape characterisation and aggregates these in a way that can support the concept of different types of similarities. The system can therefore be tuned to be more sensitive to some measures and hence be better in particular application areas, e.g. sheet metal or extruded shapes than any general technique might. The system can also adopt methods for partial representations of parts as described in Ref. [12].

3. ShapeSpace

The ShapeSpace program works by initially crawling through a database of parts which might be represented in almost any CAD format and produces STL meshes of these. A wide range of parameters are calculated (in general use 30 different values) and these form the signature of each part. Some of the measures are relatively simple such as the aspect ratio but others are more complex or make use of specific commercially protected algorithms.

Using a 30 entry shape characterising signature allows the use of many previously developed algorithms that are described in the literature. The exact choice used is selected for an individual part environment. This approach also allows the system to readily adopt new measures that can be developed for specific applications. A typical list of measures in use is shown in Table 1.

This table gives a brief description of the individual measures that are used to form the shape signature for any individual 3D model. Thus the shape signature S is a vector of these quantities as follows:

$$S(i) = \{w_1M_1 + w_2M_2 + \ldots\ldots\ldots\ldots\ldots\ldots w_nM_n\}$$

where M_j is a shape descriptor as shown in Table 1, W_j is a weighting factor between 0 and 9, i is the part identifier and n is the total number of shape descriptors in use.

There is clearly overlap between some of these measures, however because they can be pre-calculated, there is little cost in generating them. A principal component analysis (PCA) can be executed for a particular part collection to estimate the extent of this overlap. For example Fig. 2 shows the results of a PCA carried on the use of the measures in a collection 250 CAD files that were generated for various mechanical machine designs (Fig. 3).

The system generates a shape signature for each part stored in the database and when a query is entered it generates a new signature for that part ($S(q)$). A pseudocode version of the search strategy can be as follows:

```
setup{
for each model i;
    generate shape signature S(i);
next i;
cluster parts according to k-means}
loop{
given a query part q;
    generate shape signature S(q);
    do until q = target part;
      select cluster for q;
      do until screen is full;
        select nearest neighbours;
      allow user to select best guess at target part}
```

Tests with users have shown that in the 'virtual warehouse' 3D environment that the parts are presented in, 256–512 models can readily be viewed and understood. The parts are displayed most likely first (in banks of 25 – see Fig. 5) and then less so the further back they are on screen.

There are various strategies that can be used to generate clusters and these can even be mixed to form complex networks of

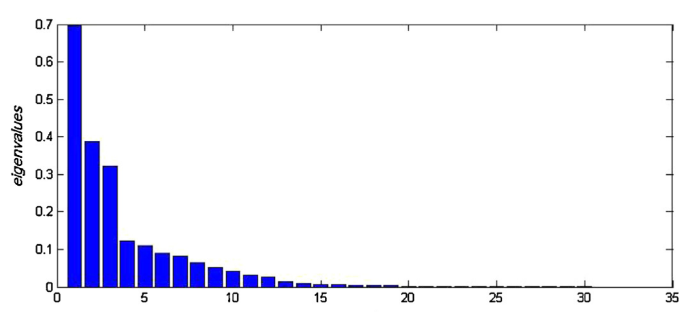

Fig. 3. Results of PCA showing the influence of the various eigenvalues.

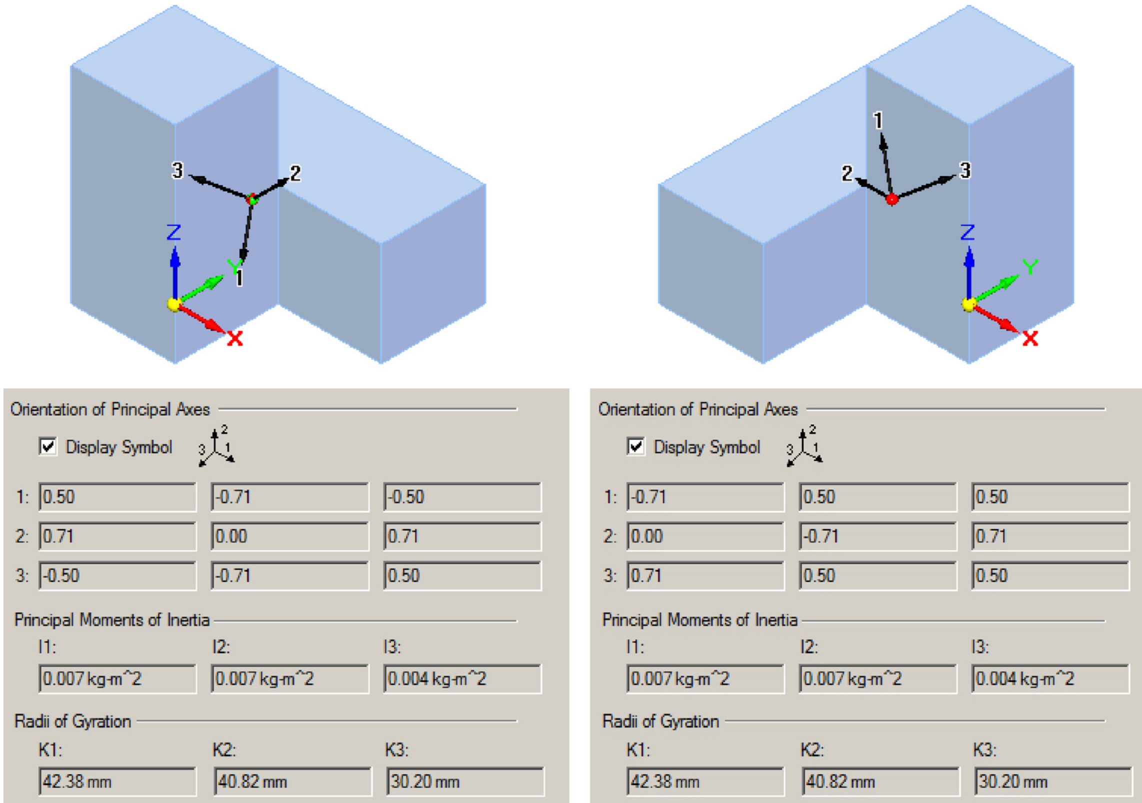

Fig. 4. Principal axes calculation for mirror image parts showing how orientation results differ between two parts that can be thought of as right and left handed.

parts but standard measures based on Manhatten distance are found to be useful. This method will work well where parts are geometries vary 'evenly' across the search space. A problem that is frequently encountered is that of generating false positive chiral parts, finding the left hand version of a right handed part or vice versa. Chiral parts have been found to be very common in automotive, aerospace and many other industry sectors and differentiating between mirror images of parts can be exceedingly complex (we have not found a general solution). In practice the number of false positives can be reduced substantially in nearly all

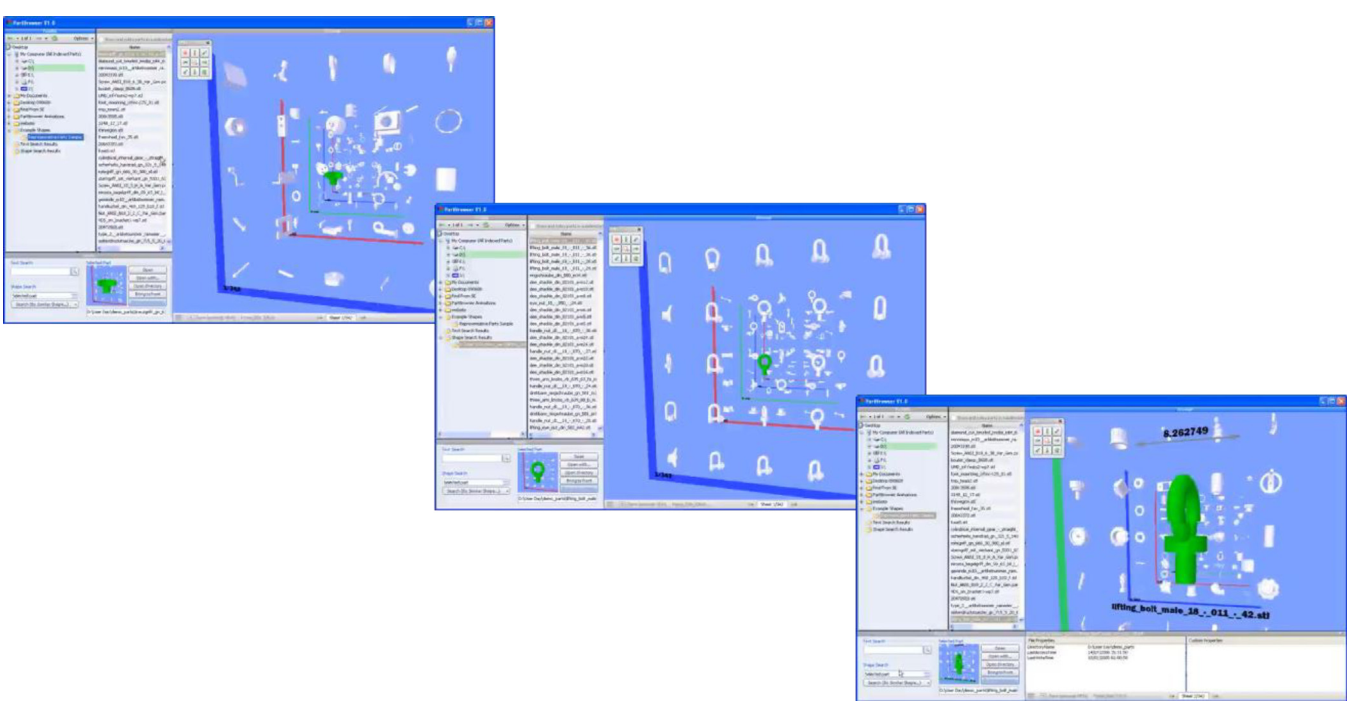

Fig. 5. A typical search sequence using a ShapeSpace. An animation of this can be seen at: http://www.youtube.com/watch?v=YeW7vnaPk7k.

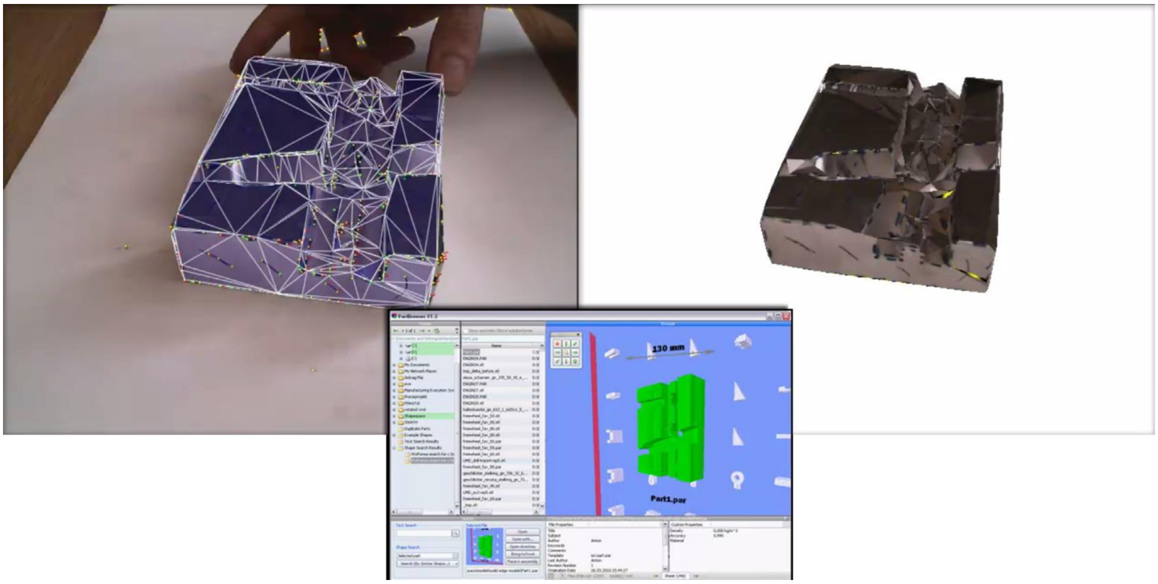

Fig. 6. Integrated scan and search.

cases by making use of the generation of principal axes. Fig. 4 shows the calculation of principal axes for two simple mirror image parts.

If the direction of the first two principal axes for each part are mapped onto each other it is usual to find that the direction of the third axes are opposing and thus chiral parts can be detected. The choice of whether these parts are displayed or not can be left to the user. This problem is particularly severe when duplicate searching is being undertaken rather than simple single part search.

The aim of the search strategy described is not simply to find the best match for a part in one step but to produce a large number of suggestions to the user that represent the search space in a way that allows him to navigate based on his idea of similarity, similar because a part is 'wavy' or similar because a part is long and thin. This means that the idea of closeness is not simply based on an aggregated measure of all the parameters. Suggestions are presented to the user on the basis of globally similar parts (from aggregated measures) but also on the basis of parts that are only similar in specific or limited domains. In this way user intent can be employed to guide the search process. The suggested parts or links are given to the user in a 3D warehouse format and the user can easily manipulate the screen to move forward or back through suggestions. Fig. 4 shows a series of screen captures that show a typical search process through a series of screens from Shape-Space.

The example given in Fig. 5 shows a common search through 3 screens. The user starts in this case with an initial screenful of suggestions and picks one most similar candidate to the one being searched for. The part is identified on the 2nd screen and its details given, in this case from a database of 40,000 parts. Longer searches do take place and larger databases have been used but with positive manipulation of the search algorithm, as outlined earlier, and with good suggested parts being offered to the user, based on judicious clustering, it has been found that even in databases of 80,000 plus parts the typical search length is usually around 4 at most and almost always less than 6. Exact statistics do not exist because the databases in use are constantly changing and the purpose and types of search vary continuously.

Searches based on scanned data are typically shorter than those starting with a general screen of parts. In some cases the system will find the required part immediately, however, this is not always possible because scanned parts may be merely similar to those being searched for, for a variety of reasons. Firstly, scanned parts are often worn and damaged and are therefore not perfect representations of the original data version of the part. Secondly, sometimes the target part is not actually the part being used as a query because it is a newer version or replacement part and is therefore ultimately preferred.

4. Integrated scanning and search

By joining the two techniques described so far, the development of a fast system of identifying 3D components is made possible thereby providing the user with whatever linked information is available. A typical search through a network of 3D models (in this case a database of around 40,000 parts) is shown in Fig. 6.

Fig. 6 shows 3 views, the upper left picture represents what the camera sees in terms of recognition of the surfaces of a part presented to it and the triangulation being applied to it. The second view in the upper right shows the generated query model and finally the third bottom centre view shows the part being immediately recognised and identified by the shape matching algorithms.

For subsequent searching and part selection use is made of a network based model of the part database. Although use can sometimes be made of networks built from relations based on common design features, these are unsuitable for the application described in this paper and instead the networks are derived solely from shape measures. The shape measuring and characterisation is based on a crawler that works its way through the database performing calculations and posts the results to a central searchable location. This data reduction means that files can be readily searched without access to the original data, thereby ensuring security of the original data, which may be especially important in cloud based implementations. A wide variety of measures (normally around 30) are calculated and presented as a vector for each part. Common values that are evaluated are the volume, surface area or aspect ratio of the part. There is overlap between some of the measures but the dual aims of assessing similarity but also at times trying to differentiate between parts means that all calculated parameters are currently

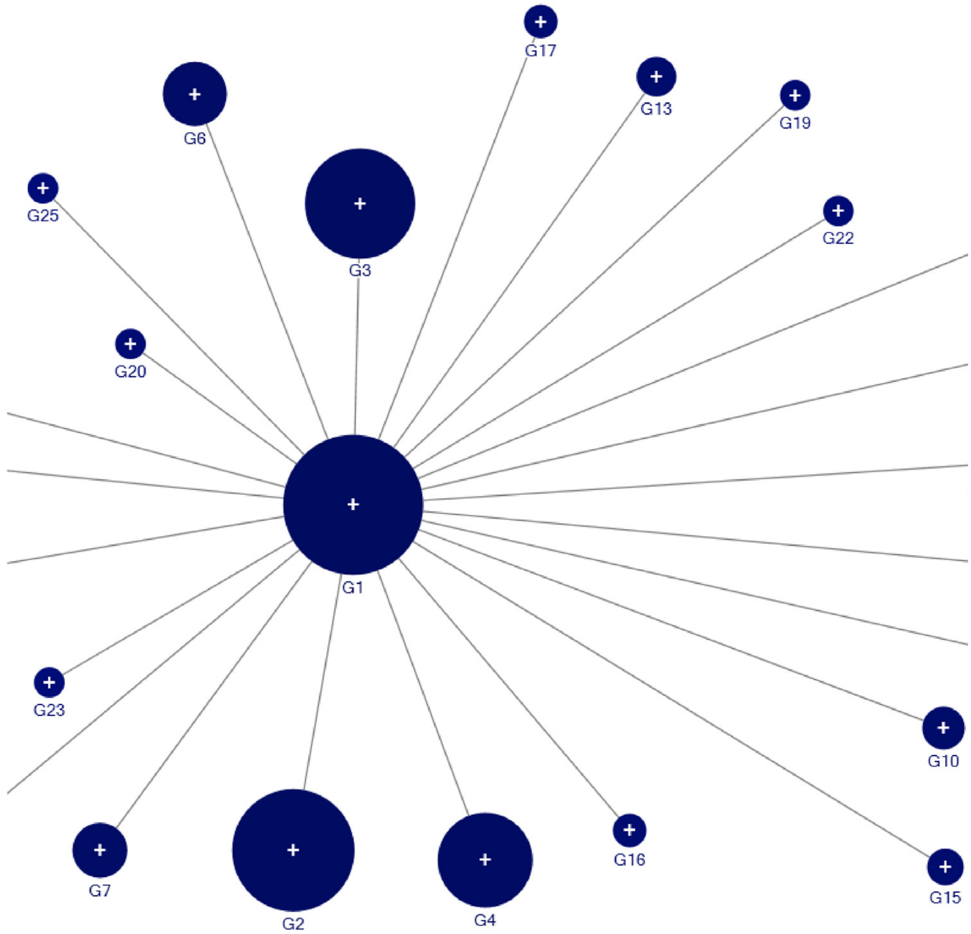

Fig. 7. Network view of part clusters. The groups shown consist of relatively small numbers of parts grouped together by the measures generated from shape signatures.

retained. E.g. in a trial database it was found that the correlation between volume and surface area was close to 0.4 as might be expected from a random population of engineered shapes. It is possible that a principal component analysis (PCA) could reduce the number of measures used however this would save very little time or storage and all of the data can be useful in developing dissimilarity matrices during the clustering phases of the network.

5. Network

Whilst performing post-query searches of a large database the ShapeSpace system does not use a fixed network structure. Instead the links in the network are generated each time the user chooses a part as being similar in some way to the part he is looking for. Thus the parts that are shown on screen may be thought of as the nodes in a network in which every part on screen is joined to every other part on screen. This corresponds to a dynamic clustering approach where all the parts in the database are plotted in a 30 dimensional space. When the user selects a part, the distances from that part to the others can be readily obtained using Euclidean or Manhattan measures and clustering applied so that parts are selected for display to the user on the resulting refreshed view, typically as follows. The closest n parts (according to the aggregate total distance) are selected for viewing front and central to the user on screen. Thereafter a further set of m suggestion parts are selected if they lie close to the chosen part in one particular dimension (i.e. they may be similar because they have similar relative surface areas despite being dissimilar in other ways e.g.). Finally a few (l) parts are chosen at random from parts of the search space not yet sampled.

It is tempting to try to optimise the values of l, m and n that should be chosen to minimise the search, however, this is not possible because the search process is dependent on the navigability of the network rather than any single simple structural aspect of it. Thus only experience with users and the actual network being searched allows some adjustment of the values to be made. However in general the approach does allow the development of a network view that corresponds to complete connectedness, some highly connected nodes and a good degree of navigability for the user. Although these networks cannot be said to be small world they have been found to be relatively efficient from a search perspective. It is also possible in the search tool that previously unused relations can be generated to augment the shape based ones. Previously it was stated that feature data would not be used because it would not be available from the scanned data. Once the user makes a choice to search further there may in fact be such data stored in the database of existing parts and this can then be used. Further information that enhances the searchability of the network could be information regarding features or assembly relationships. Fig. 6 shows a simple network view of a part database with suitable clustering applied based on a dissimilarity matrix Fig. 7.

The size of each node represents the number of parts contained in it. The graph shows how a user can in theory go between any two part models in 2 steps (the geodesic distance) in this database of 500 parts. Larger databases of tens of thousands of parts exhibit geodesic distances of around 4–6.

As has been pointed out navigability or searchability of a network of parts cannot be formally defined because it is dependent on user skill and every search is for a particular purpose and therefore uses different user domain knowledge. In general however it has been found that users are likely to only search for a given amount of time but the system described successfully allows the user to readily find parts with an upper limit of around 5 clicks or 4 choices for well over 90% of cases in large databases of real company data where 80,000 parts have been used.

Although the work and examples reported here was aimed primarily at applications in engineering product management there a considerable uses for this technology in other areas, e.g. medicine, the arts [13] and archaeology [14].

In the setup used in this work, it typically took around 30 s to 1 min to generate sufficient data through laser scanning or fast camera exposure to generate a suitable model for search. It would be useful to further integrate the two systems so that the triangulated data could be continuously sent to the search engine so that the user would become aware as soon the engine found a match in real time, thus minimising the generation of any redundant scanned data.

6. Conclusions

The paper has described a successful method of locating engineering parts in real world databases of 80,000+ parts and which combines various query building methods with a shape browsing strategy. We use as a measure of success the fact that 90% of parts can be located within 6 'clicks' and that computation time is not an issue. The system successfully combines existing and novel methods of shape description in a composite weighted vector that allows flexibility and can be readily adapted to new environments. We argue that search strategies must be able to be customised to specific collection types. Also presented is a practical means of recognising and using (or removing) data relating to left or right handed parts.

7. Equipment

In order to generate the models, searches and graphs presented in this paper a number of tools were adopted as follows. ProFORMA was written as a bespoke application in Linux with C++. Similarly the ShapeSpace system was written in C++ and C# under Windows. The test models were taken from several industrial sites. Other programs were written to clean up and format the data so that network analysis tools could be used to further condition the results and draw the networks. Pajek was used for most of the analysis tasks with NodeXL also being employed, particularly for drawing the network.

References

[1] H. Rea, R. Sung, J. Corney, D. Clark, N. Taylor, Interpreting three-dimensional shape distributions, Proceedings of the Institution of Mechanical Engineers Part C: Journal of Mechanical Engineering Science 219 (June) (2005) 553–566.

[2] J. Tangelder, R. Veltkamp, A survey of content based 3D shape retrieval methods, Multimedia Tools and Applications 39 (2008) 441–471.

[3] A. Cardone, S.K. Gupta, M. Karnik, A survey of shape similarity assessment algorithms for product design and manufacturing applications, Journal of Computing and Information Science in Engineering 3 (2003) 109–118.

[4] C.-F. You, Y.-L. Tsai, 3D solid model retrieval for engineering reuse based on local feature correspondence, International Journal of Advanced Manufacturing Technology 46 (2010) 649–661.

[5] J. Bai, S. Gao, W. Tang, Y. Liu, S. Guo, Design reuse oriented partial retrieval of CAD models, Computer-Aided Design 42 (2010) 1069–1084.

[6] D.E.R. Clark, J.R. Corney, F. Mill, H.J. Rea, A. Sherlock, N.K. Taylor, Benchmarking shape signatures against human perceptions of geometric similarity, Computer-Aided Design 38 (September) (2006) 1038–1051.

[7] Q. Pan, G. Reitmayr, T. Drummond, ProFORMA: probabilistic feature-based on-line rapid model acquisition, Presented at the 20th British Machine Vision Conference, London, UK, 2009.

[8] G. Klein, D. Murray, Parallel tracking and mapping for small AR workspaces, Presented at the Int'l Symposium on Mixed and Augmented Reality, 2007.

[9] T.F. Ansary, J.P. Vandeborre, M. Daoudi, A framework for 3D CAD models retrieval from 2D images, Annales des telecommunications (Annals of Telecommunications), 60 (2005) 1337–1359.

[10] Y. Wang, R. Liu, T. Baba, Y. Uehara, D. Masumoto, S. Nagata, An images-based 3D model retrieval approach, Lecture Notes in Computer Science 4903 (2008) 90–100.

[11] L. Liang, W. Hanzi, C. Tat-Jun, D. Suter, Z. Shusheng, Retrieving 3D CAD models using 2D images with optimized weights, in: 3rd International Congress on Image & Signal Processing (CISP), 01/04/2010, 2010, p. 1586.

[12] I. Cheuk Yiu, S.K. Gupta, Retrieving matching CAD models by using partial 3D point clouds, Computer-Aided Design & Applications 4 (2007) 629–638.

[13] S. Philipp-Foliguet, M. Jordan, L. Najman, J. Cousty, Artwork 3D model database indexing and classification, Pattern Recognition 44 (2011) 588–597.

[14] A. Itskovich, A. Tal, Surface partial matching & application to archaeology, Computers & Graphics 35 (2011) 334–341.

Frank Mill is chartered engineer and works as a senior lecturer at The University of Edinburgh. He graduated with a B.Sc. (honours: first class) Technology with Industrial Studies in 1983 from Napier College and with a Ph.D. in Shape Optimisation in 2004 from the University of Edinburgh. He has worked with Redcastle Systems and ShapesSpace Ltd. on technical aspects of CAD. He is a co-founder and director of ShapeSpace Ltd and has carried out extensive research, consultancy and teaching in CAD.

Andrew Sherlock is the CEO of ShapeSpace Ltd and Technical Director of Actify (UK) Ltd. He graduated with a first class B.Eng. honours degree in electrical and mechanical engineering (1995) and with a Ph.D. in shape optimisation (2004) from the University of Edinburgh. He has held positions in various companies involved with optimisation and CAD/PLM work and is a co-founder of ShapeSpace Ltd.

Qi Pan is a senior engineer at Qualcomm Austria Research Centre GmbH. He graduated in 2007 with a B.A./M.Eng. with distinction in electrical and information sciences and with a Ph.D. in computer vision in 2010 from the University of Cambridge.

Esme Anderson is currently pursuing a Ph.D. in computer aided design at the University of Edinburgh. She graduated with an upper second class Master of Engineering in mechanical engineering in 2011 from the University of Edinburgh.