CHEMICAL ENGINEERING TRANSACTIONS

VOL. 44, 2015

A publication of AIDIC

The Italian Association
of Chemical Engineering
Online at www.aidic.it/cet

Guest Editors: Riccardo Guidetti, Luigi Bodria, Stanley Best

ISBN 978-88-95608-35-8; **ISSN** 2283-9216

DOI: 10.3303/CET1544009

Building a Statistical Shape Model of the Apple from Corresponded Surfaces

Femke Danckaers*[a], Toon Huysmans[a], Mattias Van Dael[b], Pieter Verboven[b], Bart Nicolai[b], Jan Sijbers[a]

[a] iMinds Vision Lab, Dept. of Physics, University of Antwerp, Belgium
[b] Division of Mechatronics, Biostatistics and Sensors (MeBioS), KU Leuven, Belgium
femke.danckaers@uantwerpen.be

In this paper, a method for building a 3D statistical shape model of the apple is described. The framework consists of two parts. First, a reference surface is registered to each apple surface, derived from 3D CT scans of apples, of the population to obtain meaningful correspondences between the shapes. In the second part, the corresponded surfaces are used to build a statistical shape model from the population of apples. This model maps out the variability within the population and by adapting the shape model parameters, new, realistic surfaces can be obtained. By parameterizing the surface, an apple can be described with a compact set of basis functions, which has applications in surface fitting description, recognition, or meshing, e.g. for storage simulation. The constructed apple shape model is tested on performance and has proven to be a good representation of the population and can be used in many applications.

1. Introduction

Capturing the variability of fruit is useful in many ways, such as object detection (Rakun, et al., 2012) and shape prediction. A statistical shape model is useful to predict the final size of an apple (Zadravec, et al., 2013) or to search for a correlation between the stages of growth of an apple (Stajnko, et al., 2013). Another application is to use the shape variability to estimate the volume of an apple from a single view (Iqbal, et al., 2011). Realistic shape models of Conference pears have been used to develop nondestructive methods for measuring fruit firmness (Jancsok, et al., 2001). Fruit package designers may use fruit shape models to evaluate the effect of the fruit on airflow characteristics and thus cooling uniformity (Ghulam, 2015).

Current techniques are mostly based on 2D contour models or simplified 3D models based on contours (Ho, et al., 2011). This leads to data loss in the resulting shape model. With our suggested approach, the entire shape of the apple is characterized. Therefore, a shape model contains much more information, which may lead to better or more accurate decisions in the applications.

To build a shape model, we need to know the correspondences between the fruit surface instances in the population. One option is to annotate the corresponding points manually, but this is time-consuming and error prone (Bromiley, et al., 2014). In this proposed framework, the corresponding points are automatically found by registering each apple surface with a template surface. By doing so, each registered instance will have the same correspondences as the template surface.

3D surface registration is an elegant approach to obtain correspondences. The goal of surface registration is to minimize the geometric distance between the reference and target apple surface, while maintaining the correspondences. (Amberg, et al., 2007) presented an algorithm in which each vertex is displaced separately by an affine transformation matrix. They introduced a stiffness parameter in the registration procedure, causing a vertex to be displaced along with its neighbors. The stiffness value decreases during the iteration, allowing a more elastic deformation, which resulted in a good geometric fit, but often suboptimal correspondences.

Please cite this article as: Danckaers F., Huysmans T., Van Dael M., Verboven P., Nicolai B., Sijbers J., 2015, Building a statistical shape model of the apple from corresponded surfaces, Chemical Engineering Transactions, 44, 49-54 DOI: 10.3303/CET1544009

Our main goal was to develop a surface registration framework that provides an accurate geometric fit while maintaining the correspondences. With the correspondences, we want to obtain a model - that is compact,- (i.e. has few parameters), is highly specific, (i.e. only describes apples of a certain class), but also with sufficient generalization ability to be able to describe new instances of that class.
In the first part of this paper, we describe our method for surface registration with automatic transfer of correspondences from the reference surface to the target surface. In each registration step, we realign the reference surface with the target surface, in order to maintain the correspondences.
In the second part of our framework, a shape model is built from the corresponded surfaces (Cootes, 1995). Thereby, it is important that the surfaces are superimposed by optimally translating and rotating the surfaces. When desired for the application the, the surfaces can also be scaled in this step. The optimal poses are determined by Procrustes analysis. The model is built by performing principal components analysis (PCA) on the corresponding points of the population. In this model, the mean surface and the main variations are incorporated.

2. Methods

In this section, the developed framework is described. The reader is referred to (Danckaers, et al., 2014) for further details on the generic algorithm. The first part of the framework is surface registration. The registered surfaces are used in the second part of the framework, where a shape model is built.

2.1 Surface registration

In the surface registration part, a reference surface is registered to a target surface, such that the geometric distance between those surfaces becomes minimal while maintaining correspondences (Danckaers, et al., 2014). First, an initial global rigid registration is executed. Then, a global rigid registration and an elasticity modulated registration are iteratively repeated. During the iterations, the stiffness gradually decreases, such that the surface will become more elastic through the iterations. The framework is illustrated in **Errore. L'origine riferimento non è stata trovata.**
The first step of surface registration is applying a rigid alignment. To that end, in both surfaces corresponding points are identified. This is done by casting a normal ray from each vertex of the reference surface to the target surface. When the normal of an intersection point is in the same direction (within a tolerance) as the normal of the point on the reference surface, that points can be considered corresponding. Another restriction for corresponding points is that the normal may not intersect the surface multiple times before reaching the corresponding point.
In the elastic part of the registration the vertices are allowed to translate separately, while motion is restricted by a stiffness parameter β that regulates the strength of the connection with the neighboring vertices and which decreases during the iterations. In this way, the movement of neighboring vertices is constrained, resulting in similar movements for nearby vertices, as displayed in **Errore. L'origine riferimento non è stata trovata.**. By applying weights to each vertex the importance of this vertex can be set. If no corresponding point for a vertex of the target mesh can be found, its weight is set to zero. In that case, this vertex simply

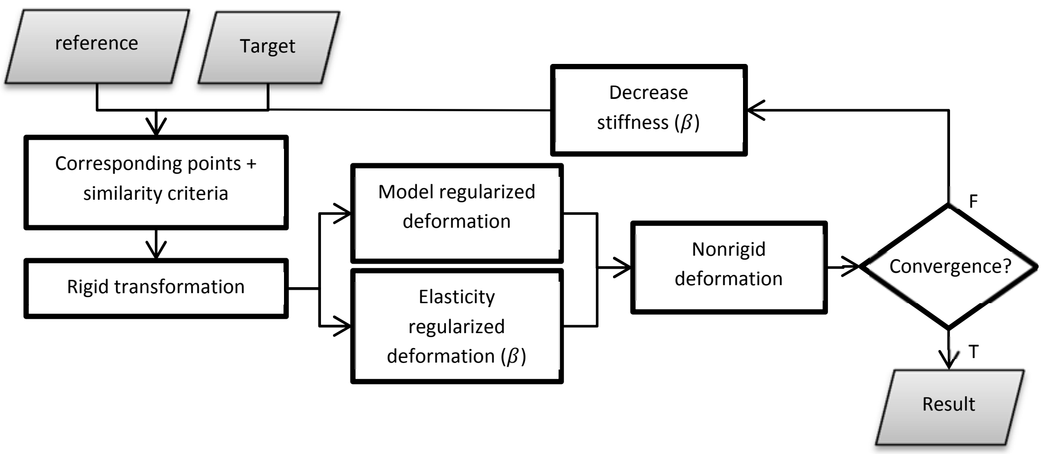

Figure 1 The surface registration framework

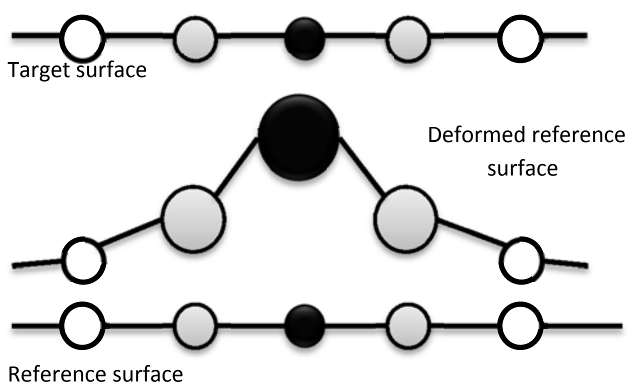

Figure 2 Schematic representation of the stiffness. When a vertex of the reference surface is translated towards its corresponding vector on the target surface, the neighbouring are forced to gradually move along.

translates along with its neighboring vertices.

2.2 Building a shape model

The second part of our framework consists of building a statistical shape model based on the correspondences that resulted from the surface registration (Cootes, 1995). To build a shape model, it is important that the surfaces are superimposed by optimally translating and rotating the surfaces. The optimal poses are determined by Procrustes analysis. The model is built by performing principal components analysis (PCA) on the corresponding points of the population. In this model, the mean surface and the main variances are incorporated. The population of *n* apples is represented by an *n*-dimensional point cloud, where each point represents an apple. This cloud can be represented by *n-1* eigenmode vectors, where the first eigenmode is the largest variance in the population, the second eigenmode is the second largest variance perpendicular to the first, etc. This means that a new, realistic surface can be formed by adapting the shape model parameters.

2.3 Surface registration with shape model prior

For surface reconstruction, e.g. in the case of occlusions or partial surface data, prior knowledge about the shape of an apple is needed. A shape model contains information of the class of apples under consideration allowing one to complete apples from partial data with the most plausible shape.

The PCA model is fitted to the target surface effectively calculating the contributions for each of the shape modes. The goal is to adjust the shape model so that the shape of the model approaches the shape of the target apple surface. A new surface $\boldsymbol{X}$ can be formed by multiplying the weights $\boldsymbol{b}$ of the instance with the principal components $\boldsymbol{P}$ of the shape model and adding this up to the mean apple shape $\overline{\boldsymbol{X}}$.

$$\boldsymbol{X} = \overline{\boldsymbol{X}} + \boldsymbol{Pb} \quad (1)$$

2.4 Parameterization

Parameterization of a surface is the task of defining a map between the surface and a simple parameter domain, like the plane, sphere or cylinder (Huysmans, et al., 2005). Such a map equips each point of the apple surface with a coordinate in the space of the parameter domain. In this paper, the apples are represented with a triangle mesh and the map is only defined explicitly for the vertices. Parameterization can be seen as the result of a continuous deformation of the surface into the parameter domain. By parameterizing the shape model, each apple in the model can be easily described using basis functions, like spherical harmonics or b-splines.

We manually create a hole in the top and bottom of the apple, so we can work in the cylindrical parameter domain. A mapping from the cylinder to the triangle mesh of the apple is needed. Therefore the apple mesh will be represented by a progressive mesh. With this representation, the number of triangles is reduced until the simplest shape, an open prism with six vertices, is left. This simple shape can be easily parameterized by equidistant placement of its six vertices on the two boundaries of the cylindrical domain. The next levels in the progressive mesh are parameterized by inserting the removed vertices one at a time and optimizing their positions on the cylinder in a way that the mapping between the cylinder and the apple introduces a minimum of distortion. When all vertices are re-inserted, the parameterization of the original apple surface is obtained and each vertex has also a (u, v) coordinate in the cylindrical coordinate system.

Figure 3 Left: mean geometric error (in mm) displayed on reference apple. The distance between the original apple and the registered apple is displayed by a color map. Right: shape deformation of a reference apple towards a target apple.

The statistical shape model is equipped with a parameterization by parameterizing the mean surface. Then,

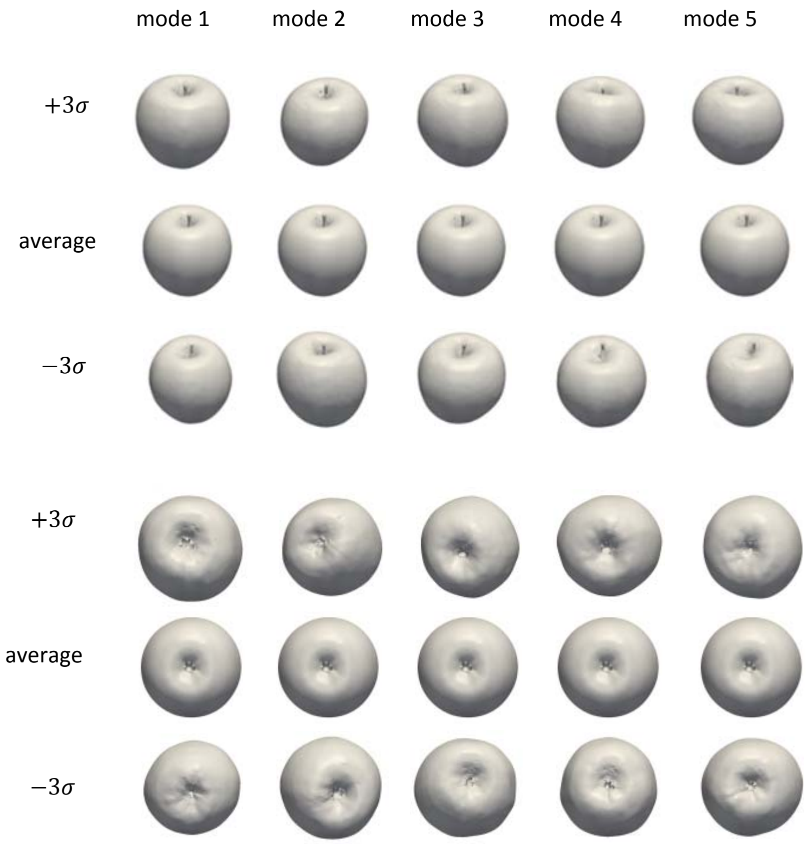

Figure 4 Two views (left: front view, right: bottom view) of the first five shape eigenmodes of the apple shape model, plus and minus three standard deviations. The model was created from 30 Jonagold apple surfaces.

through the correspondence, all instances of the model also have parameter coordinates. With this technique, the point-based models can be described by B-splines. This is a very compact representation and is usable in CAD and finite-element environments, so the models can be used for simulations.

3. Experiments and results

For the experiments, we used 3D CT scans of 30 Jonagold apples. All scans were obtained in a Philips HOMX 161 X-ray system (Department of Materials Engineering (MTM), KU Leuven, Belgium), operating at a voltage of 85 kV and a current of 0.41 mA for 180° rotation with a 0.5° scan step and 16 frames averaging. The image reconstruction was performed with NRecon [Bruker micro-CT, Kontich, Belgium] and is based on a modified Feldkamp cone-beam algorithm.

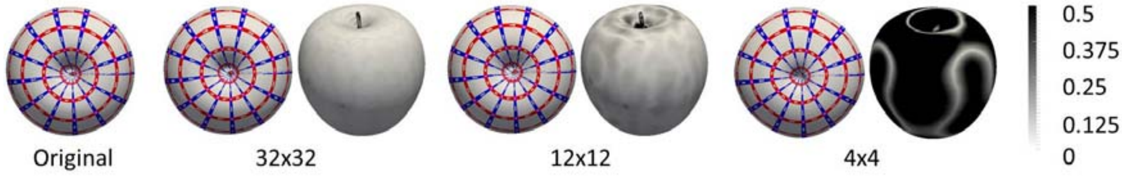

Figure 5 Left: top view of parameterized apple shape model with iso-parametric curves. Right: B-spline approximations with different number of control points on the B-spline grid and the distance, in mm, between the original surface and the approximation.

3.1 Surface registration

To obtain meaningful correspondences, each apple must be registered with the same reference surface. We registered a target apple from the Jonagold population to every other apple in that population. The first apple from the population is chosen as template surface. This template surface is uniformly resampled with 40.000 vertices. Each apple of the population is registered with this reference apple. Each registered apple was compared to the original apple and the distance between these surfaces was calculated. The mean absolute distance between the reference apple and target apple was $0.051mm$. In **Errore. L'origine riferimento non è stata trovata.**, the geometric registration error is displayed on the reference apple. The darker the area is, the larger the geometric error in that area. The largest errors occur in the regions of the peduncle and the stamen. The remainder of the surface has a very low surface registration error.

3.2 Shape model

For these experiments, we built a statistical shape model from 30 registered surfaces of Jonagold apples. The first five shape modes of the apple model are displayed in **Errore. L'origine riferimento non è stata trovata.**4. The largest variations of the shape of the apple are described in the first modes.

3.3 Parameterization

The shape model is parameterized so it can be described with basis functions. In **Errore. L'origine riferimento non è stata trovata.**5, the apple surface with iso-parametric curves is shown. The apple is approximated by B-splines with different numbers of control points. A 4x4 approximation reassembles an apple, but the difference between the original and the approximation is clearly visible. An approximation with 32x32 control points is nearly identical to the original apple model. Therefore it is usable in CAD and finite-element environments, so the models can be used for simulations.

3.4 Model Performance

The Jonagold model is tested on performance (Taylor, et al., 2008). It is tested on compactness, generalization, and specificity.
The compactness measure describes how the model captures the variation. Compactness is the cumulative

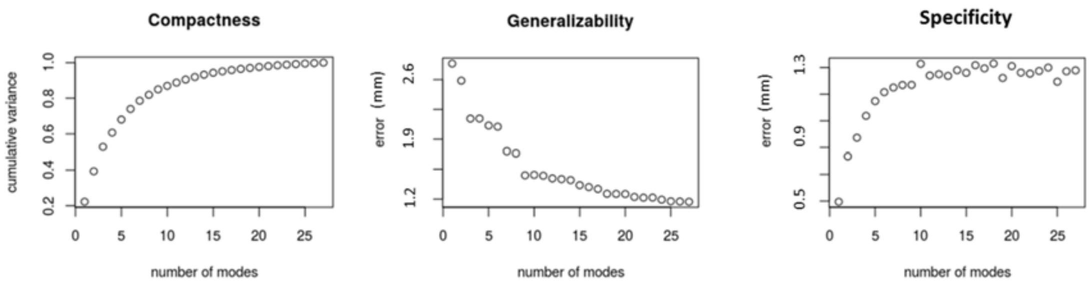

Figure 6 The different model performance measures. The cumulative variance is normalized. The generalizability and specificity error are the average distance calculated in mm per vertex.

sum of the standard deviations of the principal components. This test shows that the model captures more than 80% of the shape variation within the first 10 modes.

The generalization ability of a model determines how well the model generalizes to unseen instances of the class. It is measured as the average approximation error after fitting leave-one-out versions of the model to the left-out surface. From the generalizability test can be concluded that the error of fitting the model to an unseen surface decreases drastically from 10 shape modes.

The model specificity measures to what extent the random samples, generated by the model, resemble the original apple surfaces. For each principal component, a random principal component weight is generated, following the multidimensional Gaussian distribution of the PCA model. From each sample, the distance to the apple from the training set that reassembles the sample the most is calculated. The mean error in mm for 1000 trials is calculated. The test is performed with an increasing number of shape modes. The specificity test proves that the apple model is able to generate instances that differ a lot from the training surfaces.

In **Errore. L'origine riferimento non è stata trovata.**, the different model performance measures are visualized. The values of compactness are normalized.

4. Conclusion

Experiments on the surface registration technique resulted in a good geometric fit and good correspondences.

By parameterizing the surface, the apple model can also be described by other parameters than PCA parameters. Our approach of modelling and subsequent parameterization is also applicable to other fruit and vegetable shapes. Specifically for elongated shapes, like pears, bananas, and cucumbers cylindrical parameterization could be useful.

The model performance tests prove that our apple shape model is a good representation of the population, is able to generate realistic apples with different shapes than the apples form the training set, and is employable in many applications.

The developed surface registration and modelling techniques are also applicable to other fruit types of any topology.

Acknowledgements

This work was supported by the Agency for Innovation by Science and Technology in Flanders (IWT SBO Tomfood) and the iMinds B-Slim project.

References

Amberg, B., Romdhani, S. & Vetter, T., 2007. *Optimal Step Nonrigid ICP Algorithms for Surface Registration.* Minneapolis, MN, IEEE, pp. 1-8.

Cootes, T., 1995. Active shape models-their training and application. *Computer vision and image understanding,* pp. 38-59.

Danckaers, F. et al., 2014. *Correspondence Preserving Elastic Surface Registration with Shape Model Prior.* s.l., IEEE.

Ghulam, M., 2015. Date Fruites Classification Using Texture Descriptors and Shape-Size Features. *Engineering Applications of Artificial Intelligence,* pp. 361-367.

Ho, Q. et al., 2011. A Three-Dimensional Multiscale Model for Gas Exchange in Fruit. *Plant Physiology,* pp. 1158-1168.

Huysmans, T., Sijbers, J. & Verdonk, B., 2005. Parameterization of Tubular Surfaces on the Cylinder. *Journal of the Winter School of Computer Graphics,* pp. 97-104.

Iqbal, S., Gopal, A. & Sarma, A., 2011. *Volume Estimation of Apple Fruits Using Image Processing.* s.l., s.n.

Jancsok, P., Clijmans, L., Nicolaï, B. & De Baerdemaeker, J., 2001. Investigation of the Effect of Shape on the Acoustic Responce of Conference Pears by Finite Element Modelling. *Postharvest Biology and Technology,* pp. 1-12.

Rakun, J. et al., 2012. *Detecting Natural Objects by Means of 2D and 3D Shape Analysis.* Optija, Croatia, s.n., pp. 345-354.

Stajnko, D. et al., 2013. Modeling of 'Gala' Apple Fruits Diameter for Improving the Accuracy of Early Yield Prediction. *Scienta Horticulturae,* pp. 306-312.

Taylor, C., Twining, C. & Davies, R., 2008. *Statistical Models of Shape: Optimisation and Evaluation.* s.l.:Springer Publishing Company.

Zadravec, P. et al., 2013. Fruit Size Prediction of Four Apple Cultivars: Accuracy and Timing. *Scientia Horticulturae,* pp. 177-181.