Integrated Modeling of Shape Semantics for Industrial Design

István T. Hernádvölgyi, Giuliana Ucelli, Giuseppe Conti, Christian Faisstnauer and Raffaele de Amicis
Fondazione Graphitech
Via F. Zeni 8, Rovereto (TN) 38068, Italia
http://www.graphitech.it

Abstract—**Modeling with 3D shapes plays a significant role in Industrial Design. Therefore semantic processing and interpretation of shapes and surfaces in CAD environments is an important problem. In this paper, we outline our architecture which integrates Knowledge Management and shape geometry for the purposes of cataloging, archiving and querying shapes. Unlike previous work, in our framework the textual annotation and geometry are closely integrated. We use domain knowledge encoded in a knowledge base to infer the geometric context of individual design elements and use this information to better analyze the corresponding shape. Our proposed architecture could be adapted in CAD applications where there is a well understood conceptual layering of the product assembly and the individual components are of similar geometry.**

I. INTRODUCTION

In a number of applications, such as Industrial Design, Architecture and Medical Imaging, digital 3D models and shapes play a significant role. Therefore it is important that these shapes are properly cataloged so they can be retrieved and analyzed. Semantic information in a 3D digital model, however, is inherently implicit, hence it has to be extracted by some dedicated algorithm. In reality, even free form shape recognition turned out to be a hard problem. Systems that mine the web for VRML models that are similar to some query object (for example [18], [3]) are not accurate enough to properly distinguish between objects which are similar by design (for example surfaces which represent specific panels making up the body of cars). Most model search engines were specifically designed to retrieve objects of which none or very little information is known. Some systems consider textual information as well, which is usually obtained from the file name and the context of the web page where the model was found [18], [12]. The basis of geometrical comparison is usually a shape descriptor or fingerprint which is extracted from the digital format and serves as the search key. For generic shape retrieval the descriptor should be invariant under translation, scaling and rotation to account for the unrestricted versatility of models posted on the Web. On the other hand, industrial models created for a specific purpose in CAD environments represent well defined classes of objects. A shape in an actual design is unlikely to occur by itself but to be part of a more complex object and this context could provide useful information to interpret the shape more precisely. While the same type of element created in two different CAD tools would almost certainly be oriented and scaled differently, we can use the inferred geometric context to establish a common frame of reference. To date there has been little effort to design and develop systems where shape comparison makes use of the actual context where the shape occurs. The notable exceptions are systems specifically designed for medical [17] and archaeological [16] analysis. These applications use specific geometric domain knowledge to derive compact and telling shape descriptors which facilitate accurate analysis. Inspired by these ideas, we propose a novel and generic architecture which takes advantage of the fact that most 3D models created by CAD tools are made up of unique elements which are further grouped into layers representing well defined semantic objects. We go further than just merely relying on naming conventions and consider using *geometric* knowledge from the context of a shape to derive more specific descriptors for the class of objects it belongs to. We also describe how this can be achieved in the framework of the MPEG-7 standard [13].

II. RELATED WORK

Storing and using semantics is a major concern in modern design products. This is also the very purpose of the new MPEG-7 standard [13], [11] which encourages designers to describe their models in a standard way so it can be cataloged and therefore shared by other systems. The use of Knowledge Management (KM) tools has already been recognized as a means to enhance competitiveness of business companies [5], [8], [7]. The currently running WIDE [22] and Wise [2] projects are specifically targeting knowledge management in the Engineering domain. Besides describing our architecture of an intelligent shape catalogue, our primary focus in this paper is to demonstrate how some basic knowledge to interpret the geometric context can provide useful information to derive shape descriptors. In the general framework, when the purpose is to browse models scattered across the Internet, the shape descriptor has to have certain invariant properties to account for arbitrarily scaled and oriented objects. The descriptors suggested by Osada et al [15], Kazhdan et al [6] and Novotni et al [14] are scale and orientation invariant. It is also common to establish orientation considering only the model itself disregarding its context. These are usually the center of mass and the Principal Component Analysis (PCA) axes [21]. They are used to align the object first and the shape signature is derived with respect to this coordinate system [20]. Using domain knowledge to align the object for better analysis has been considered to compare medical images [10], bones [17] and archaeological artifacts [16]. In these systems the feature points are selected by comparing the sample to other images or by semi automatically identifying regions of interest on the object. Semantics inferred from certain geometric features of the object were considered in [4], but without the context of the objects. As far as we can tell, we were the first to propose context geometry inferred with the help of knowledge management tools to be considered in the derivation of shape descriptors.

III. PROPOSED ARCHITECTURE

In this section we introduce our proposed architecture in detail. First we outline a high-level design showing the top modules of our system. Next we concentrate on our feature extraction module which takes geometric shape context into account and illustrate how it is integrated with the textual annotation of the design element. We also identify existing technologies which facilitate a relatively straight-forward realization. Finally we describe what sort of tools are needed to aid the integration of existing designs created by other tools.

A. High-level Design

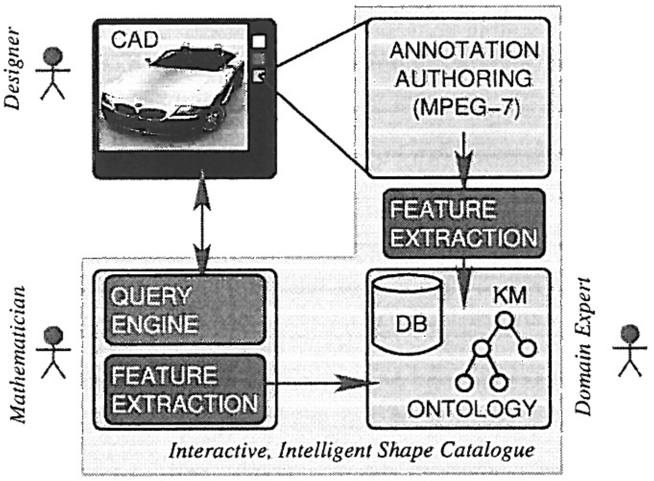

Fig. 1. High-level design for shape matching

Figure 1 depicts the outline of our architecture. We propose no major modifications to the usual CAD environment. In fact, we would like to emphasize that our interactive shape catalogue could be considered as an "add on" to a CAD tool. The emerging MPEG-7 standard allows annotation of multimedia content and therefore we will use it as the basis format to store design elements in our database (marked as DB). The annotation part will be described in more detail later, but we would like to mention that it is semi automated and requires no special training for a designer. The ontology is designed by a domain expert with the aid of a knowledge manager. It can be any third party knowledge management (KM) system and its purpose is to encode specific knowledge about the domain of design. The main purpose of the KM module is to provide a resolution engine for high level queries. This module will also be detailed later. The designer interacts with the Query Engine to retrieve cataloged elements. If the query involves geometry (like an example shape), then it first will pass through a Feature Extraction routine, which creates a shape descriptor that could also incorporate geometry from the context markers provided with the query. The shape descriptors are also stored with the actual models in the database. This does not cause bloating since a shape descriptor is much more compact then the original model. One of our innovations is the incorporation of geometric context into the shape descriptors, and therefore the following section contains a thorough description.

We would also like to highlight the user interaction with this architecture. The designer of the product is mainly interacting with the CAD environment and his access to the shape catalogue is only through the Query Engine and the annotation tool. The shape descriptors are designed by experts in geometry (identified as "mathematician") and do not concern the designer. The context knowledge specific to a certain line of industrial product is established by a domain expert. In the next section we describe how some node labels of the ontology also correspond to efficiently identifiable portions of the model's geometry which will aid the "mathematician" to design more distinguishing shape descriptors.

B. Novelty

In this section, we detail how the topology and layering of a complex element can provide geometric context which could be exploited in order to customize shape descriptors for shapes of particular functionality or purpose. Figure 2 shows the wheel *layer*. It is made up of the bolts, a rim and the tire, with their own separate geometries. Separate objects which make up the same logical component can be grouped together in CAD systems and can be given their own names. When the model is exported the names of the layers as well as the names of the components together with their geometry are saved. Figure 3 depicts surfaces and layers which form the outer hull of a generic car. While there are significant geometrical differences between surfaces representing the body of cars (consider a Ferrari vs. a Fiat Panda), the topology shown in Figure 3 is representative for a number of cars. In fact, there

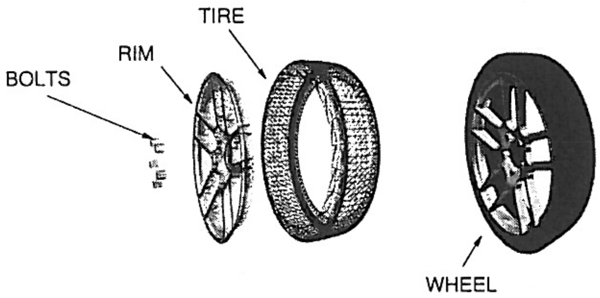

Fig. 2. Objects on the Same Layer

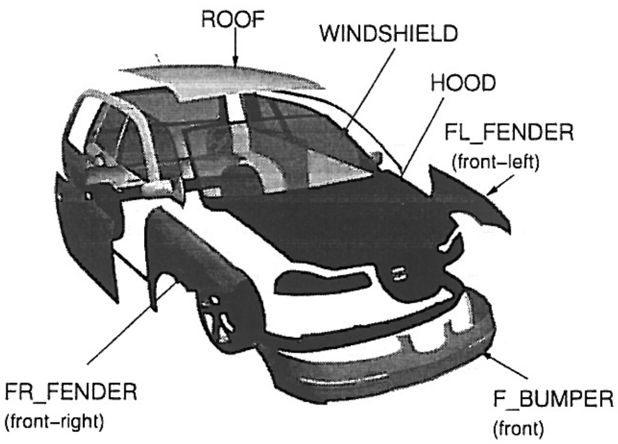

Fig. 3. Topology of panels constituting the outer hull of a car

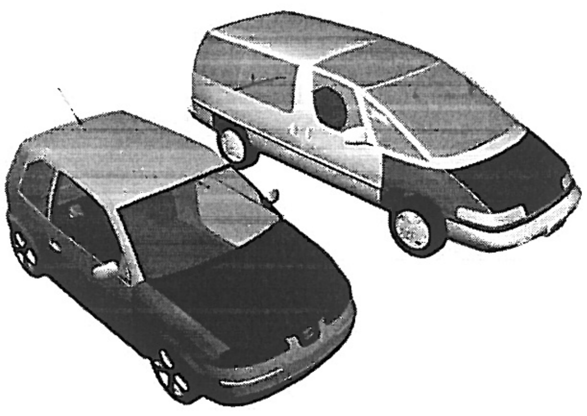

Fig. 4. The hood of two cars (also highlighted the front fenders)

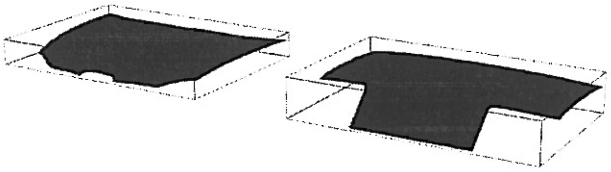

Fig. 5. Bounding boxes determined from the fender alignments.

are only a few different body part topologies across car designs (hatchback, cabriolet, sedan, station wagon, etc). These can easily be expressed as logical sentences. For example, the following rules (in Prolog) encode that the *hood* is between the left and right front fenders for any model and the windshield is between the hood and the roof for cars which have roofs (non cabrio).

```
...
between(_,fl_fender,fr_fender,hood).
between(X,hood,roof,windshield) :-
   has(X,roof).
has(sedan,roof).
...
```

A complete knowledge base which encodes car body part topologies can be used to infer the exact geometric context of any single part. This geometric context also provides a geometric frame of reference for the object. Most shape descriptors [6], [15], [14], [9] are designed specifically to be invariant under some affine transformations (rotation, translation, scale) because the intended scale and orientation cannot be easily determined from the geometry alone. For a layered CAD model, however, the accurate orientation and relative scale can be determined (with the help of a small knowledge base) from the geometry of surrounding elements. Figure 4 shows two cars with geometrically different hoods. The hood on any car is located between the front fenders. Observe that from the geometry of the fenders a precise coordinate system can be determined for the hoods. A plane is fitted between the top line of the fenders and this is taken to be the $y = 0$ plane. The x and z axes are also clear, since the fenders are symmetric and the mid-line between the front fenders toward the front bumper is the "forward direction" (x-axis) which will halve the $y = 0$ plane and the z-axis is perpendicular to these. The origin can either be taken at the centroid's projection on the $y = 0$ plane or at the midpoint between the fenders. Figure 5 shows the aligned bounding boxes. Note that the hoods are now nearly parallel as opposed to their different slopes as they are oriented in the cars. An alternative and common method of establishing a coordinate system without considering the context geometry is performing principal component analysis (PCA) [21]. In this case, however, the protruding grill of the van's hood tilts the principal axes considerably which makes the resulting PCA coordinate systems inadequate to use as a common frame of reference. In Figure 6, the angle α shows the actual "tilt" between corresponding axes.

Now we derive two actual shape descriptors which use the reference frame (Figure 5) established from the geometry of related objects from the context. This is meant to be an example and would clearly not work for all type of shapes. Our aim is to show the opportunities provided by some domain knowledge and context geometry. This construction is loosely inspired by [20], we will point out differences as we go along. Let us subdivide the intersection of the xz-plane with the bounding box into an $N \times N$ grid. Let $f^1_{ij} \in R$ be defined as the average y value of surface points on the hood in the i, j grid region. $i, j \in [-N/2, N/2]$. If there is no surface point, $f^1_{i,j}$ is set to a suitably large positive or negative value. The shape descriptor ω^1_{uv} is obtained from a 2D Discrete Cosine

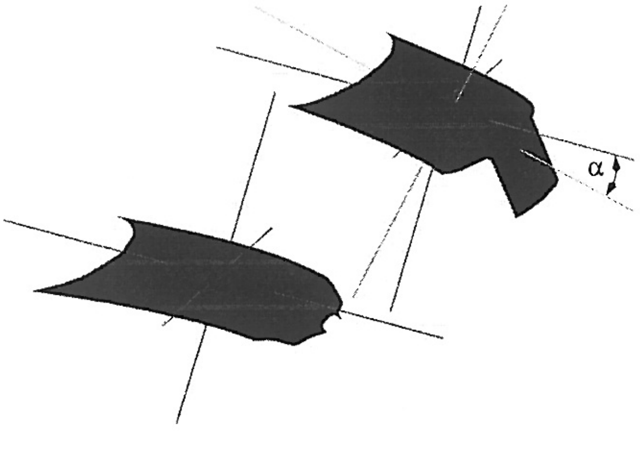

Fig. 6. PCA axes from the tessellated hood surfaces.

Transform of the surface

$$\omega^1_{uv} = c_u c_v \sum_{i=0}^{N_x-1} \sum_{j=0}^{N_z-1} f^1_{ij} \cos\left(\frac{\pi(2i+1)u}{2N_x}\right) \cos\left(\frac{\pi(2j+1)v}{2N_z}\right) \quad (1)$$

where

$$c_u = \begin{cases} \sqrt{\frac{1}{N_x}} & \text{for } u = 0 \\ \sqrt{\frac{2}{N_x}} & \text{for } u = 1, 2, ..., N_x - 1 \end{cases} \quad (2)$$

The shape descriptor is composed of the low harmonics ω^1_{uv} for $0 \leq u < K_x, 0 \leq v < K_z$, where $K_x << N_x$ and $K_z << N_z$.

The other shape descriptor is derived from contour points on the surface. We take n cylinders with increasing radii with base on the xz plane and intersect the surface. $f^2(\phi, r)$ is the y value of the intersection point with the cylinder whose radius is r measured at the positive angle ϕ from the x-axis. The shape descriptor is composed of the first few low harmonics of the discrete 1D cosine transform of the n contours.

$$\omega^2_{u,r} = c_u \sum_{i=0}^{N-1} f^2\left(\frac{2i\pi}{N}, r\right) \cos\left(\frac{\pi(2i+1)u}{2N}\right) \quad (3)$$

We verified with our models that, indeed, the low frequency components are the major constituents. Figure 7 shows an example surface reconstructed from the low harmonics. It is clear that the reconstructed objects (both the surface and the contour lines) are indicative of the original shape. We also have generated a random database of 5,000 NURBS surfaces (Non-uniform Rational B-splines). Figure 8 shows the 5 shapes retrieved from this database which were deemed most similar to the framed query shapes. Similarity is measured as the Euclidean distance between the shape descriptors. It is clear from the figure that the shapes retrieved resemble the query object well. However, the human perception of similarity is not necessarily based on the overall shape of an item. Our shape descriptors (and most in the literature) are derived from samples of the surface or the volume, in essence, treating each region equally. A human, on the other hand, may be more perceptive to notice peculiar details. Arguably, the items labeled A and B in Figure 8 could be judged less similar to the query objects because of the presence (or lack) of a noticeable detail (encircled in the Figure). These details may be identified by performing a more detailed analysis (eg. by considering higher frequency components) but it is clear that a reference frame established by context geometry could also help in locating and reasoning about geometric details.

In [20], the coordinate system is established by PCA analysis, which, as we pointed out earlier, could be problematic. We also used Discrete Cosine Transforms rather than Discrete Fourier Transforms because we had better surface reconstruction. Moreover, [20] used a 3D transform where f is defined as the model's intersection with a N^3 voxel grid. Since we "know" that hoods are pressed from sheet metal (and therefore not well represented as a volume), our approximation of the surface will be more appropriate for comparing hoods (and single ply surfaces, such as other body elements). This knowledge also resulted in lower dimensionality descriptors (there is no need to voxelization).

These examples show how geometric information from the context (determination of the object's orientation) and some general knowledge (hoods are bent sheet metal) can be easily incorporated into a shape descriptor. Again, it is important to realize that our descriptors are descriptors for hoods which could also work well for other surfaces of car body parts; our target domain. It is not appropriate to make comparisons between objects that are better represented as volumes (like the entire surface structure of a car).

C. *Enabling Technologies*

1) MPEG-7: MPEG-7 [13] is developed specifically to provide textual context annotation to accompany multimedia data. Moreover it was one of its driving principles that it is "object based". In essence, this means that the content can be described as a composition and that the individual objects can be independently accessed. MPEG-7 also includes a variety of shape descriptors [11]. These are region-based, contour-based and 3D spectrum descriptors. For us the latter one is of more importance to compare 3D models, although both region and contour-based descriptors can be applicable for characteristic projections of 3D models. The "built-in" 3D shape descriptor of MPEG-7 is based on the principal curvatures of the entire object (shape index) [9], [11]. Considering the MPEG-7 standard was designed to be media independent, the shape descriptor must be very generic. MPEG-7, however, is also extensible in a standard way (XML). Since we suggest that shape descriptors should also take advantage of the context of other shapes associated with the model, we only need to be able to isolate the individual parts in the model with the help of annotation and then our own processor can take over. The "built-in" descriptors are not lost, but we can provide a more informed processor within the framework of MPEG-7.

2) Knowledge Management: On one hand, we use ontologies to describe high-level conceptual and topological knowledge about certain assemblies. The example we have

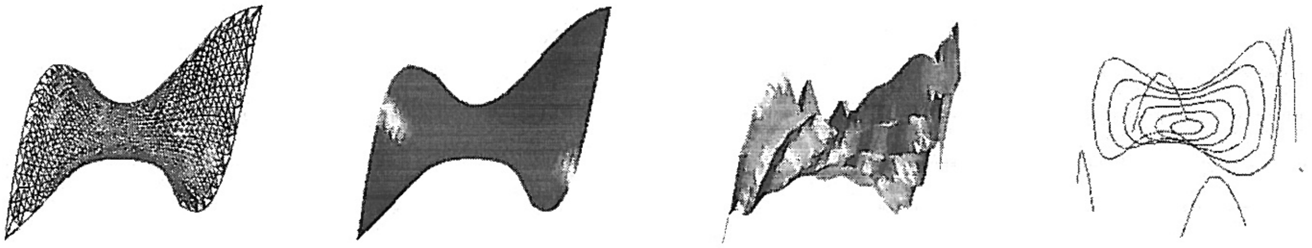

Fig. 7. Surface reconstruction from the shape descriptors. Wire-frame and tessellated rendering of a NURBS surface and the reconstructions from the low harmonics from the surface and contour cosine transforms.

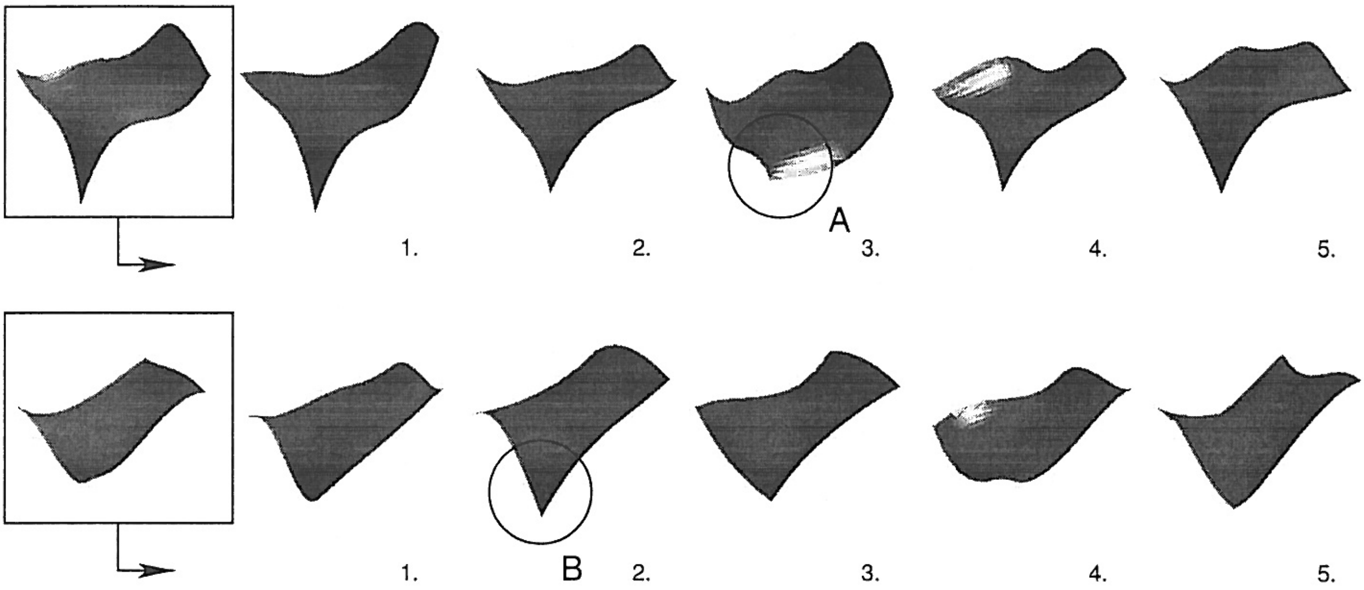

Fig. 8. The 5 "most similar" shapes (to the framed query shapes) retrieved from a database of 5,000 random NURBS surfaces.

carried in this paper is the well known layout of car body parts. This knowledge is obvious for engineers, but from the geometry of an assembly it would be extremely hard to reverse engineer. As we have shown in the previous section, having this knowledge at hand is quite useful to perform geometric reasoning and analysis. A model can be textually annotated and described, but at some point, especially when it comes to describing fine details, geometry is expectedly more descriptive. To get to these details, however, the ontology is used to isolate the object and to establish a precise frame of reference from the context objects in which it can be analyzed in more detail.

In our system, we also use knowledge management (KM) for another purpose. To be able to link concepts to their geometry, there is a need to have clear meaning of terms in the target domain. This knowledge must be encoded in an ontology which is used to infer how to obtain the geometrical entities from the repository. The label of car parts suggested in Figure 3, of course, are just suggestions and not a proposed standard. Nevertheless, a KM system can be employed to perform term resolution to ensure interoperability between proprietary systems or peer knowledge could be shared as in [19]. For example, it is expected that the model parts created by a CAD will have suggestive names, and with the help of a KM system, simple linguistic preprocessing and some term enrichment from a tool like WordNet [23], a good level of correspondence can be automatically established between the layer and part names of models designed by different authors.

3) High-level Interfaces: We certainly anticipate that fully automatic conceptual alignment of assemblies developed by different designers is a very hard problem. Using a knowledge management tool (as described above) will establish some correspondences, but the rest will require some human help. For this purpose, we are designing multi-modal interfaces to ease and simplify the task as much as possible. This includes a suite of selection tools that can be used to isolate part of a geometry (even if it is not detachable in the original design) and can be easily tagged by dragging the label from a tree view of the assembly ontology.

The query engine also needs a new kind of interface. For example, when the designer is looking for hoods by giving an example, schematic fenders will pop-up which will help him (and our feature extraction routine) to sketch the query object. In a Virtual Reality environment, this can be done very nicely by tracing a pen. Providing a frame of reference for the sketch as well, will also result in a more accurate sketch.

D. Progress and Future Work

The intelligent shape catalogue component outlined in this paper is being realized by two currently running projects in Graphitech and it is also related to the European project, AIM@SHAPE [1].

SIMI-Pro (Semantic Information Management System for Innovative Product Design) is focusing on the knowledge management component. Its research goals include automatic term resolution and distributed knowledge sharing for engineering design, development and realization. The understanding and encoding of product assemblies to aid geometric reasoning was singled out in this paper to demonstrate a novel approach how KM tools can provide useful knowledge to empower traditional CAD systems.

MoSeS (Modeling Shape Semantics) is a new a project whose primary focus is the extraction and modeling of semantics of shapes making up industrial design components. Proper shape analysis requires more than just simple classification (such as this element is a door rather than a roof). As we outlined in this paper, currently we are investigating the viability of using object specific shape descriptors. We already have preliminary evidence based on our recent results that there exist clear cases where establishing a common coordinate system and knowing some characteristics of the shapes associated with the type of objects can result in efficiently computable accurate shape descriptors.

We also prefer the derivation of the common frame of reference from context geometry and domain knowledge from ontologies, rather than requiring the human designer to provide actual reference points. However the annotation model and the query engine require high level user interfaces, such as a "context sketch-pad" and 3D tracing. These modules are being realized within the SpaceDesign-Pro and InSIDe projects.

IV. CONCLUDING REMARKS

In this paper, we have described our architecture which integrates Knowledge Management with Industrial Design. While there are many on-going projects and approaches to use KM tools with CAD, ours stands out in respect to using domain knowledge to infer context geometry and in turn using this to provide a better framework for analysis. We have shown how one can use the context to derive specific shape descriptors for the class of objects. We also believe that there are a lot of application domains (other than cars) where a similar approach can be adapted.

It is important, however, to remember that shape analysis and matching do not have to rely on the geometric context and we are not offering a substitute for generic descriptors that search for 3D models on the Web. We believe it is important to open up the possibility and make the geometric context available for product specific analysis.

V. ACKNOWLEDGMENTS

Part of this work has been done within the projects MoSeS and SIMI-Pro, financed by the Provincia Autonoma di Trento. The authors would also like to thank the consortium of Aim@Shape for support in this research activity. The models in this paper were created by Stefano Agostini.

REFERENCES

[1] AIM@SHAPE. http://www.mpi-sb.mpg.de/ scherbaum/aim@shape.

[2] Y. Barnard and A. Rothe. Knowledge management in engineering: Supporting analysis and design processes in innovative industries. In *Building the Knowledge Economy, Issues, Applications, Case Studies*, pages 931 – 938, 2003.

[3] Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. On visual similarity based 3D model retrieval. *Computer Graphics Forum (EUROGRAPHICS'03)*, 22(3):223–232, September 2003.

[4] J. Corney, H. Rea, D. Clark, J. Pritchard, M. Breaks, and R. MacLeod. Coarse filters for shape matching. *IEEE Computer Graphics and Applications*, pages 65–74, 2002.

[5] T. D. Davenport and L. Prusak. Working knowledge: How organizations manage what they know. Harvard Business School.

[6] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Symposium on Geometry Processing*, 2003.

[7] B. Koch and U. Jasnoch. DEPOT: An internet-based knowledge management system for design. In *Proceedings of ISE, Las Vegas*, pages 49 – 51, 2001.

[8] B. Koch and M. Koch. An internet-based information service to support the product development process in design offices. In *Proceedings of FAIM, Maryland*, pages 77–82, 2000.

[9] J. Koenderink. Solid shape, 1990. MIT Press, Cambridge, Mass.

[10] G. Krell, H. R. Tizhoosh, T. Lilienblum, and C. J. Moore. Fuzzy image enhancement and associative feature matching in radiotherapy. In *International Conference on Neural Networks (ICNN), Houston, TX*, June 1997.

[11] B. S. Manjunath, P. Salembier, and R. Sikora. *Introduction to MPEG-7*. John Wiley & Sons, LTD, 2002.

[12] Patrick Min. *A 3D Model Search Engine*. PhD thesis, Princeton University, 2004.

[13] MPEG. http://www.chiariglione.org/mpeg/.

[14] M. Novotni and R. Klein. 3D Zernike Descriptors for Content Based Shape Retrieval. In *The 8th ACM Symposium on Solid Modeling and Applications*, June 2003.

[15] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Matching 3D models with shape distributions. In *International Conference on Shape Modeling and Applications*, 2001.

[16] A. Razdan, M. D. Liu Bae, M. Zhu, G. Farin, A. Simon, and M. Henderson. Using geometric modeling for archiving and searching 3D archaeological vessels. In *CISST*, 2001.

[17] Anshuman Razdan, Jeremy Rowe, Matthew Tocheri, and Wilson Sweitzer. Adding semantics to 3D digital libraries. In *5th International Conference on Digital Libraries (ICADL 2002), Singapore*, 2002.

[18] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. The princeton shape benchmark, June 2004. To appear in Shape Modeling International, Genova, Italy.

[19] G. Ucelli, R. De Amicis, G. Conti, F. Giunchiglia, and S. Noll. A semantic-based information management system to support innovative product design. In *Atti del Congresso Annuale AICA, Trento*, pages 77–85, 2003.

[20] D. V. Vranic and D. Saupe. 3D shape descriptor based on 3D Fourier transform. In *Proceedings of the EURASIP Conference on Digital Signal Processing for Multimedia Communications and Services (ECMCS 2001), Budapest, Hungary*, September 2001.

[21] D. V. Vranić, D. Saupe, and J. Richter. Tools for 3D-object retrieval: Karhunen-Loeve transform and spherical harmonics. In *IEEE Workshop Multimedia Signal Processing, Cannes*, 2001.

[22] WIDE. http://www.ist-wide.info.

[23] WordNet. http://www.cogsci.princeton.edu/ wn/.