Recognition of 3D-Objects: an Algorithm for the Inexact Matching of High-Level Polyhedral Representations

Hans Peter Amann, Heinz Hugli

University of Neuchâtel, Institute of Microtechnology
Rue A.-L. Breguet 2, CH-2000 Neuchâtel, Switzerland

Abstract

This paper describes a powerful inexact matching algorithm that has been applied with success to high-level 3D-object representations in a 3D object recognition system. In a promising way, the algorithm combines several approaches proposed in the last couple of years: an extension to the backtrack strategies for inexact matching of attributed relational subgraphs, error-correction isomorphism, determination of local attribute similarity and global transformation fitting, features which are efficiently used for search-tree pruning. The algorithm was tested successfully in a series of experiments involving scenes with single and multiple objects.

Index terms: 3D-object recognition, scene analysis, inexact subgraph matching, machine vision

1. Introduction

Vision is the most important manner for sighted humans to estimate shapes, to localize and recognize objects in the real world. Therefore, quite naturally, efforts have been made in the last couple of years to integrate similar facilities into automatic vision systems. Two typical examples of applications are assembly lines and robot vision. Solid results have been obtained for two-dimensional objects and for specifically defined problems in artificial, human-made environments. For three-dimensional objects in turn, the efforts made, up to now, have essentially shown the high complexity of the problem. Nevertheless, researchers are now in the position to test particular 3D-object recognition algorithms, to realize prototypes with limited features and use them in well-defined, relatively simple implementations.

The problem

The task of 3D-object recognition using machine vision can roughly be described as follows:

> Given a scene containing one or more 3D objects and a library of reference objects, identify one (or more) of the 3D-scene objects either partially or completely with respect to the given reference library object's elements.

A frame for an object recognition system

This rough task definition gives rise to a frame for an object-recognition system, consisting of the following elements:

i) Acquisition of rough 3D data and extraction of a convenient high-level object representation.
ii) Matching of the acquired and extracted high-level representation with some reference or model representations, resulting in a list of hypotheses on the best correspondences.
iii) Verification and classification of the resulting hypotheses, the selection of compatible partial results, the extraction of useful information on orientation, the localization of recognized elements , …

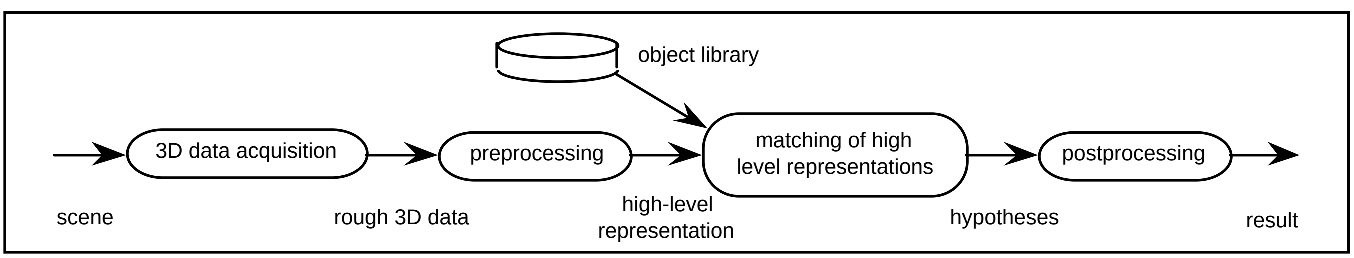

Figure 1.1: a frame for an object-recognition system

Content

In the present paper, we show a solution to the inexact high-level matching problem in the context of 3D-object recognition. Given two high-level object representations in the form of attributed relational graphs, we extract the most promising (partial) matches so that we can establish hypotheses on scene objects. The match of the two subparts of reference and scene objects is characterized by a similarity measure that indicates match quality and a transformation the reference object has to undergo so

H.P. Amann & H. Hugli, " An algorithm for the inexact matching of high level 3D polyhedric representations", Selected papers on model-based vision, Ed. Hatem Nasr, SPIE Milestone Series, Vol. MS 72, SPIE, Bellingham, WA, USA, pp. 313-321, 1992

that it fits scene data best .

We treat the subject of graph matching as a tree-search problem, using a "best-first" search. The present work is distinctive in the sense that it uses a very weak rule for tree expansion in order that a wide spectrum of topological inexactitude can be accepted. Usually, this inexact matching feature requires an important increase of computational complexity. Therefore, criteria were introduced that efficiently enable pruning the search tree, the most important condition being rigidity. Despite of the sound results obtained with this single condition, additional tree-pruning conditions based on the remaining graph attributes had to be added in order to deal with cases where the rigidity condition is not sufficient, e.g., for man-made objects where perpendicular and parallel faces are common.

Various experiments with the algorithm described show that convincing results can be obtained, i) for object recognition with single-object scenes and ii) for object recognition and object separation for multiple-object scenes, including objects which occlude themselves mutually.

Outlines

First, we describe the preprocessing applied to rough-scene data and create a high-level representation called PFRG (§2) that enables redefinition of the object-recognition problem as a (sub)graph matching problem. Next, we briefly analyze some characteristics of exact (sub)graph matching, an approach we consider too restrictive for a typical application. Before going into details of our algorithm, we rapidly present some of the related previous work (§4). Many of these ideas were integrated in our matching algorithm discussed precisely in paragraph 5. Moreover, some remarks on postprocessing are included. The remaining paragraphs are devoted to results (§6), concluding remarks and an overview on further subjects of interest.

2. Acquisition of data and extraction of high-level representations

In order to give a complete view of the environment for which the algorithm was developed, we quickly present the data-acquisition device and the preprocessing applied (figure 2.1).

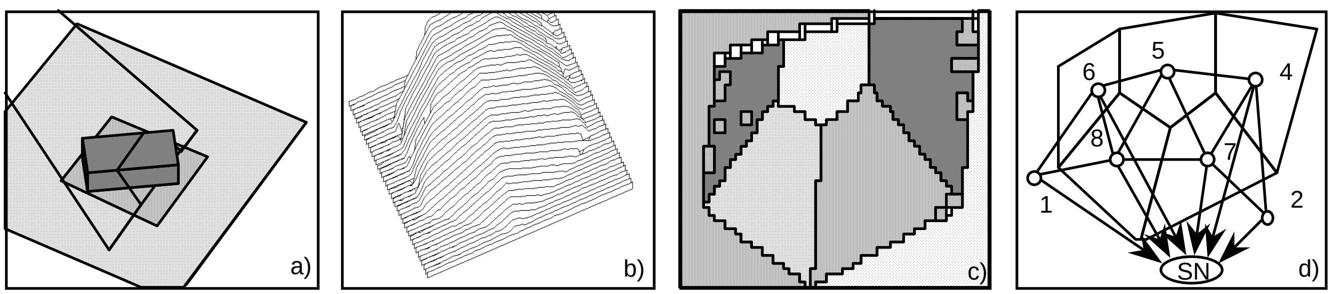

Figure 2.1: from 3D-data acquisition to a high-level representation:
a) laser range finder, b) 2.5D depth map, c) region/border representation, d) attributed relational graph representation

2.1 Data acquisition and preprocessing

The input device is a laser range finder using light planes [6], delivering a large set of 3D measurements of the visible surfaces. A range-image segmentation process is used in order to approximate the 3D point set by surfaces in our implementation planar surfaces [7,8,9]. The resulting set of surface patches that represents visible scene-object surfaces, can now be expressed at a higher level of abstraction, e.g., in the form of attributed relational graphs (ARGs). These ARG representations are neither complete nor precise representations of the scene objects due to: occlusion, limited resolution of the acquisition device, preprocessing, the choice of attributes, …

2.2 High-level representation: our choice

Two constraints were specified for our problem: i) limitation to polyhedral objects and ii) use of a high-level representation based on planar surfaces. While both constraints do not modify the nature of important problems in 3D-object recognition, they significantly reduce the complexity of the implementation. We use a high-level representation called a planar face representation graph (PFRG), a special form of an ARG. Each node stands for an object's region, approximated by a planar surface patch and each arc for a neighborhood relation between two regions (table 2.2). The node attributes are the area and vector perpendicular to the represented surface. The only arc attribute is the border length between adjacent regions.

region + node	attributes:	surface area of the region normal vector of the region

H.P. Amann & H. Hugli, " An algorithm for the inexact matching of high level 3D polyhedric representations", Selected papers on model-based vision, Ed. Hatem Nasr, SPIE Milestone Series, Vol. MS 72, SPIE, Bellingham, WA, USA, pp. 313-321, 1992

border + arc	attribute:	border length

Table 2.2: planar face representation graph PFRG

2.3 PFRG preprocessing

PFRGs of reference objects will be rather concise, PFRGs of acquired scene data are likely to be cluttered due to reasons mentioned before (§2.1). Hence, some simple algorithms were implemented who i) merge regions that are likely to have been separated and ii) which remove dummy regions, essentially the ones situated on the borders of real surfaces. Figure 2.3 shows schematically one of them, performing the removal of small border regions whose adjacencies are known but too small for proper determination of attributes "surface area" and "surface normal". Obviously, these algorithms depend on the actual range-image segmentation algorithms used and their details must therefore be discussed in the corresponding context [10].

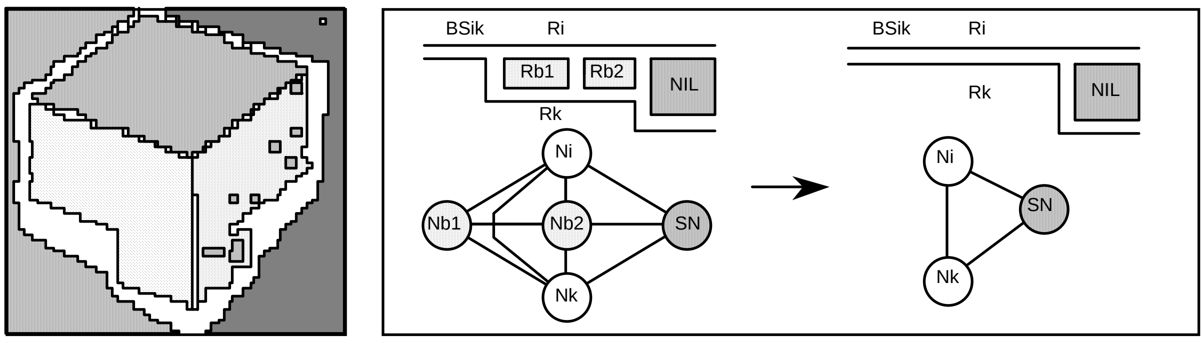

Figure 2.3: examples of "cluttering" on the scene object "Cube1" (left) and a preprocessing algorithm eliminating small border elements with nondeterminable attributes (right)

2.4 Task redefinition

Once the high-level representation and its relation to real-world data are defined, we can give a more exact task description:

> Establish an ordered list of hypotheses on similar (sub)PFRGs, with both scene and reference objects represented in the form of PFRGs, while taking into account both structural and attributal similarities.

Until now, task defining is rather general. We have not yet explained the type of graph matching we will do. The next paragraphs discuss this point. First, we briefly analyze exact graph matching, then refer to some promising related work and finally, present our selected approach in detail: inexact matching.

3. Exact graph matching

The process of 1-1 node matching of non-attributed graphs has a long history in mathematics; the tools are therefore well known for both graph matching (identical number of nodes) and subgraph matching (different number of nodes) [14]. The extension to attributed graphs can be done through introduction of similarity measurements.

These methods are based on the assumption that all of the nodes and arcs of the smaller graph have their counterparts in the bigger graph: each node and each arc can be matched 1-1 to a node or arc respectively of the second graph (isomorphism). They can be implemented either with an "exhaustive-search" algorithm or as "back-track" searches where existing solutions are extended to bigger ones, considering some pruning criteria at each extension step [2,3].

It is clear that in a general case of 3D-object recognition, the assumption made above does not hold. Occlusion, shadowing, nonidealities of the data-acquisition equipment and of the preprocessing algorithms can lead to inaccurate high-level representations. Therefore, the success of "exact matching" approaches will be heavily limited.

H.P. Amann & H. Hugli, " An algorithm for the inexact matching of high level 3D polyhedric representations", Selected papers on model-based vision, Ed. Hatem Nasr, SPIE Milestone Series, Vol. MS 72, SPIE, Bellingham, WA, USA, pp. 313-321, 1992

Figure 3.1: scene object (left), partial reference object (middle) and an extract of a possible search tree (right)

Figure 3.1 shows a simple example. For reasons explained before, the border between regions B and C of the scene object were not detected correctly, arc A(B,C) representing the adjacency of regions B and C is missing. An exact correspondence-search algorithm, using a "backtrack" approach and starting from the 1-1 node match $M_1 = \{(A,1)\}$ could be extended to $M_2 = \{(A,1),(B,2)\}$. The attempt to match (C,3) would be rejected since the sub-graph **ABC** is nonisomorphic with respect to reference sub-graph **123**.

Even in this relatively simple situation, "exact matching" behaves in an insufficient way. The drawbacks of this method would have to be compensated by a subsequent recombination of small partial solutions. Therefore, a method covering "inexact matching" (and obviously also exact matching) has been investigated and implemented.

4. Related work

Many of the ideas implemented in our PFRG inexact matching algorithm come from or were triggered by work done in the last couple of years.

In the field of exact graph matching, Ullmann [12] treated the subject of subgraph isomorphism. Tsai and Fu [1] established a first base for error-correcting isomorphic matching and it was continued by Eshra and Fu [2,3] in their backtrack-based inexact matching scheme using small sub-graphs called Basic Attributed Relational Graphs (BARGs). Shapiro's paper [4] on the use of metrics was very helpful for the determination of dissimilarities and their combination into a global cost measure. Oshima and Shirai [11] showed us the interesting feature of "root matches" that enables determination of interesting starting points for backtrack tree searches. Furthermore, they introduced reliability measurements for graph attributes. Finally, the important concept of object rigidity (§5) and interesting mathematical tools necessary for its convenient implementation as a global search tree heuristic are due to Faugeras and Hebert [5].

5. Inexact matching of high-level representations

The present paragraph is devoted to our algorithm for inexact matching of two high-level 3D polyhedral representations. Note that for object-recognition purposes, this matching step has to be repeated for each element of a set of reference objects. Given a scene and a reference PFRG, we look for a list of promising matches M_k which can be both partial and inexact. For each of the matches M_k, we also measure the overall similarity as well as transformation T_k that the reference object would have to undergo so that it fits "best" to the scene object (figure 5.1).

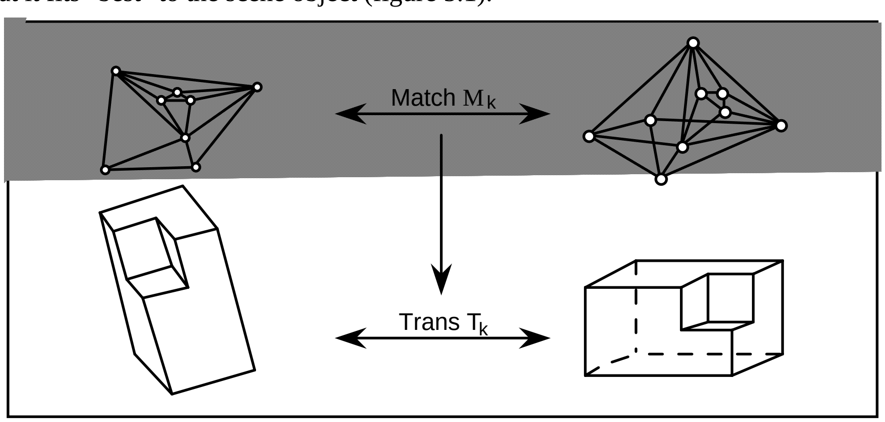

Figure 5.1: the "matching" process: left, the scene object and its PFRG; right, the reference object and its PFRG

5.1 Notations and definitions

Scene and reference PFRGs are denoted with $G_S = (N_S, A_S)$ and $G_R = (N_R, A_R)$ respectively, where N_S and N_R stand for the respective node sets and A_S and A_R for the respective sets of arcs. $A(N_{Sm}, N_{Sn})$ and $A(N_{Rm}, N_{Rn})$ stand for nondirected arcs interconnecting adjacent scene and reference nodes respectively. Finally, $M(N_S, N_R)$ stands for an established correspondence between a scene node N_S and a reference node N_R. M_k is a set of 1-1 node correspondences $M(N_S, N_R)$: $M_k = \{\ldots, M, \ldots\}$.

With these notations, we can express the problem definition more precisely:

> Given two planar face representation graphs (PFRGs) G_S and G_R, find the "best" set M_k of 1-1 scene-node / reference-node correspondences according to an overall "cost" function.

Note that using this approach, the matching process is guided by the establishment of node-node correspondences. While arc-arc correspondences will be established implicitly where possible (inexact matching), and therefore will participate in the global "cost" function, missing nodes will neither be detected nor considered.

5.2 State-Space lattice for the "best" match search

In order to find the "best" match between reference PFRG G_R and scene PFRG G_S, we apply a "best-first" search algorithm in a state-space lattice where each state S_k corresponds to a set of 1-1 node correspondences M_k. In our implementation, this match is represented by the core-node sets $C_{Rk} = \{\ldots, N_{Ri}, \ldots\}$ and $C_{Sk} = \{\ldots, N_{Sj}, \ldots\}$ respectively and the set of 1-1 correspondences $M_k = \{\ldots, (N_{Ri}, N_{Sj}), \ldots\}$ of their nodes.

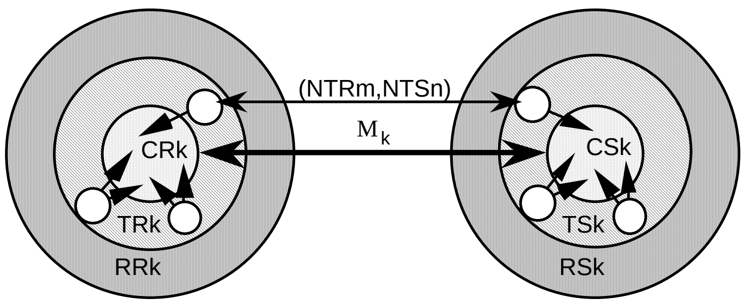

Figure 5.2: reference node set (left) and scene node set (right)

Furthermore, each state S_k has an associated set of terminal nodes T_{Rk} resp. T_{Sk}. These sets contain all nodes which have at least one arc to any of the matched core nodes of the respective core-node set. The remaining nodes are contained in the rest-node sets R_{Rk} and R_{Sk} respectively. Each state S_k of the search tree also has an associated set of attributes:

i) overall match cost M_k,
ii) "best-fitting" transformation T_k and
iii) number of arcs inserted or deleted to obtain a virtually isomorphic match.

With the "best-first" search strategy selected, the leaf state with the minimum overall cost is expanded first. Obviously, a nonrestricted expansion of the search tree would lead to a nonaffordable amount of computation. Therefore, we introduce a set of search-tree pruning criteria. The most important ones will be presented, for details see [10].

5.3 Search-tree pruning criteria

Connectivity condition

The first search-tree pruning criterion, the connectivity condition, was already introduced implicitly: a subgraph match M_k can only be extended by couples of terminal nodes. With respect to the rules introduced in previous work [3], this one is particularly weak in the present design. It allows for good flexibility in inexact matching while still limiting the expansion of the search tree. If ever the application of this rule leads to multiple subgraph matches, postprocessing could help to compensate this problem (§5.5).

Rigidity condition

For common objects, it is straightforward to introduce the rigidity hypothesis: the angles between the surface normals are constant and invariant to rotation and translation. (While the approach is more general, the use of PFRGs limits us currently to rotations). We therefore determine for subgraph matches M_k "best fitting" rotations Rot_{Sk} for the set of unitary normal vectors of the respective core-node sets [5,12]. Knowing rotation Rot_{Sk}, we can restrict the set of possible additional node matches $M(NT_{Sm}, NT_{Rn})$ to the couples of nodes that fulfill the rigidity condition; the orientation difference between a rotated reference-surface normal and the corresponding scene-surface normal must be smaller than a given threshold.

Visibility condition

With conventional data-acquisition equipment, a single view of the scene is taken. Thus, visible faces have an orientation so that the angle between the observation vector and surface normals is smaller than $\pi/2 \pm \varepsilon$. Transformation of the scene-observation vector back to the reference space hence enables pruning the set of reference terminal nodes according to their surface normals.

Neighbor match condition

The candidate couple (NT_{Rm}, NT_{Sn}) of terminal nodes is a candidate for a new match if and only if there exists at least one adjacency $A(NT_{Rm}, NC_{Ri})$, one adjacency $A(NT_{Sn}, NC_{Sj})$ and a core-node correspondence (NC_{Ri}, NC_{Sj}) as elements of previous match set M_k. This condition guarantees that terminal nodes have adjacencies to at least one matched pair of core-nodes: neighboring reference-nodes will be matched with neighboring scene nodes.

Cost conditions

The transition of one state S_k in the search tree to a next state S_i corresponds to the extension of the match M_k to match M_i by a new couple of nodes (NT_{Rm}, NT_{Sn}), modifying the overall matching cost of paragraph 5.4 by $\Delta cost(S_i, S_k)$. With the idea in mind that search branches - that would lead to an important increase in the global cost - should be discarded, this "modification cost" can be used as an additional pruning condition. Moreover, particular similarity conditions can be individually established for each attribute.

5.4 Cost function

In order to determine an overall cost function, dissimilarities determined for each type of attribute and for structural differences are combined so that search-tree states of different levels, representing matches of different sizes, can be compared: weighting must depend on the number of elements involved.

cost(S_k) = **Error!**+ ε $e_{rotfit}(S_k)$,number of nodes(S_k)) + **Error!**

Figure 5.3: cost function

Cost function is composed of two parts. First, the weighted sum of accumulated dissimilarities of the area of the matched surface patches $\Sigma_{area}(S_k)$ plus the mean squared rotation fitting error $e_{rotfit}(S_k)$ divided by the number of matched nodes. The second part consists of an arc-dissimilarity measure: the sum of arc attribute dissimilarities $\Sigma_{border}(S_k)$, increased by a dissimilarity measure for inserted arcs (A_{ins}) and deleted arcs (A_{del}) divided by the number of arcs involved (matched, inserted and deleted). The weights α, β, γ, δ and ε have been determined heuristically by a series of experiments on real data.

5.5 Postprocessing

Postprocessing, an essential step in the 3D object recognition process, covers all the steps subsequent to the match, necessary to complete the recognition task. Let us just mention the two most interesting ones among them:

i) verification of resulting hypotheses and
ii) combination of partial matches.

Verification

Match hypotheses should be verified before accepted as solutions, using information that has not been exploited for the match setup. In our implementation, the pruning process takes into account all elements of the high-level representation used; it is thus necessary to perform the verification task at a lower level of data abstraction.

Compatible solutions

Due to the connectivity condition of paragraph 5.3, the final match between a single scene object and a reference object can be composed of more than one partial match from the list of (verified) solutions. Obviously, they have to fulfill some compatibility conditions, e.g., the rotations fitted must be very similar, and each element of both scene and reference PFRGs appears just once. If more than one object is involved, the situation is even more complicated since a single-reference object could have more than one counterpart in the scene. This subject needs further discussion in a wider context.

6. Experiments and results

Weighting parameters determination

In a first series of tests, weighting factors for the overall cost were determined. It showed up that for all subsequent experiments, a single, well-chosen set of parameters was sufficient (α=0.05, β=0.04, γ=0.01, δ=0.01, ε=0.9).

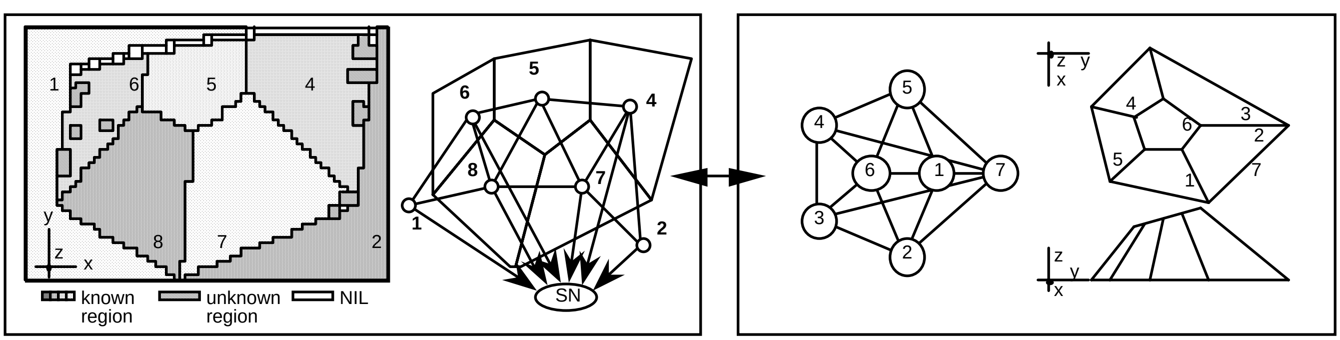

H.P. Amann & H. Hugli, " An algorithm for the inexact matching of high level 3D polyhedric representations", Selected papers on model-based vision, Ed. Hatem Nasr, SPIE Milestone Series, Vol. MS 72, SPIE, Bellingham, WA, USA, pp. 313-321, 1992

Figure 6.1: weighting parameters determination (test object "P" left, and reference object "Pyramid2" right)

Cost function selectivity

The following tests were run in order to check selectivity achieved with defined cost function. Results show that a system that exclusively uses surface orientations as attributes, e.g., in the form of gaussian images, does not behave sufficiently well. With the use of the surface-area attribute corresponding to a representation of the form of extended gaussian images EGIs [13], and the border-length attribute, selectivity becomes sufficient despite simplicity of the chosen representation. Obviously, both the surface-area attribute and the border-length attribute are sensitive to occlusion. Therefore, a reliability measure should be introduced accordingly that enables weighting the influence of the node and arc matches in the overall cost function [11].

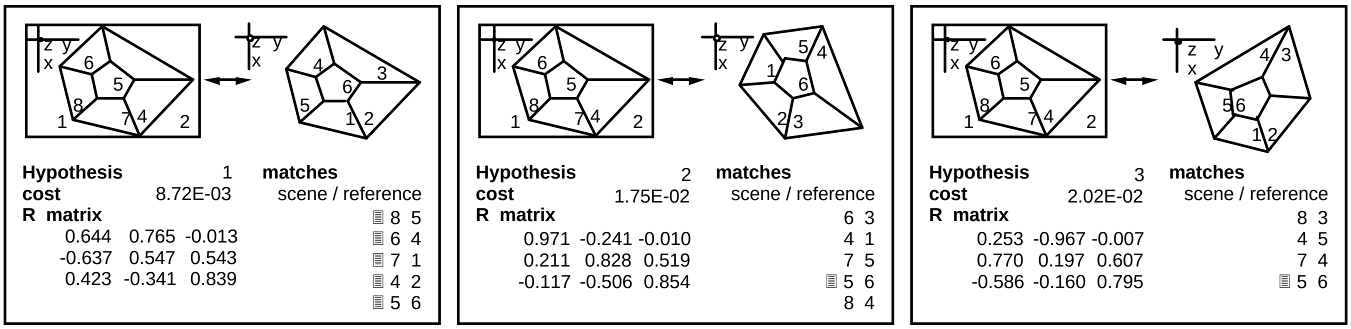

Figure 6.2: the first three hypotheses for the problem of figure 6.1

Figure 6.2 shows an example for good selectivity. From left to right, the first three hypotheses obtained for a matching attempt between scene object "P" and the reference object "Pyramid2" are presented.

The first hypothesis is the correct one, the second and the third - having considerably higher costs - represent matches made with the scene object rotated about the z-axis "one face left and right", respectively. Cost evolution the number of node correspondences and correct node correspondences are shown in figure 6.3.

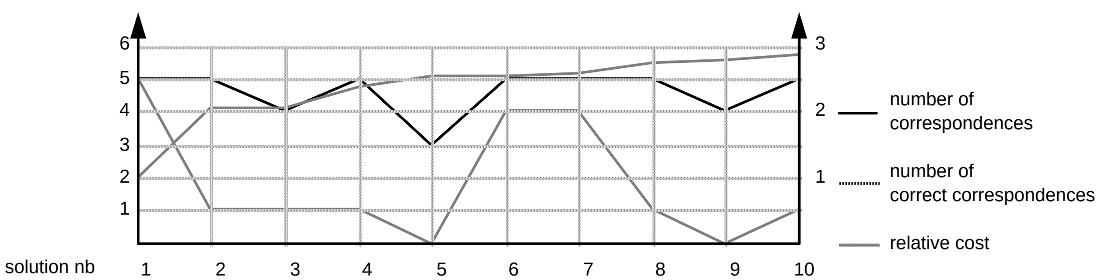

Figure 6.3: statistics for the first ten solutions for the problem of figure 6.1

Object recognition: single scene-object / multiple references

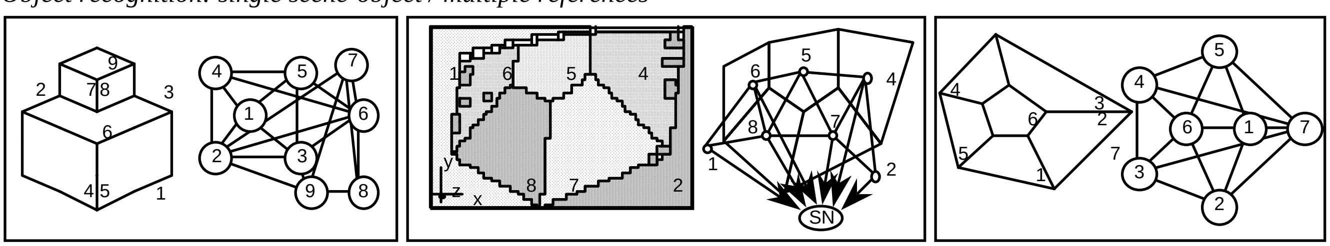

Figure 6.4: single-object recognition: test object "P" (middle), reference objects "Cube1" (left) and "Pyramid2" (right)

Next, the initial problem of 3D object recognition is dealt with: a single 3D-scene object is to be recognized with respect to a set of reference objects. Figure 6.4 shows a simple example: scene object "P" (middle) was tested against the two reference objects, "Cube1" (left) and "Pyramid2" (right). Ordering the merged set of hypotheses for both reference objects leads to a list partially shown in figure 6.5.

H.P. Amann & H. Hugli, " An algorithm for the inexact matching of high level 3D polyhedric representations", Selected papers on model-based vision, Ed. Hatem Nasr, SPIE Milestone Series, Vol. MS 72, SPIE, Bellingham, WA, USA, pp. 313-321, 1992

1	2	3
cost 8.724E-03	**cost** 1.753E-02	**cost** 2.024E-02
R matrix 0.644 0.765 -0.013 -0.637 0.547 0.543 0.423 -0.341 0.839	**R matrix** 0.971 -0.241 -0.010 0.211 0.828 0.519 -0.117 -0.506 0.854	**R matrix** 0.253 -0.967 -0.007 0.770 0.197 0.607 -0.586 -0.160 0.795
5 of 5 faces matched correctly (6 visible)	1 of 5 faces matched correctly (6 visible)	1 of 4 faces matched correctly (6 visible)

Figure 6.5: the first three hypotheses out of the merged set for the problem of figure 6.4

For the first ten hypotheses found, the match between scene object "P" and reference object "Pyramid2" always delivers a lower global cost than the match with the second reference object "Cube1": the algorithm behaves as expected.

Object recognition and object separation

The more ambitious task of object recognition in multiple-object scenes was also worked out. The PFRG of a scene "CP", showing a cube that partially occludes a pyramid, was tested vs. the reference PFRGs "Cube1" and "Pyramid2".

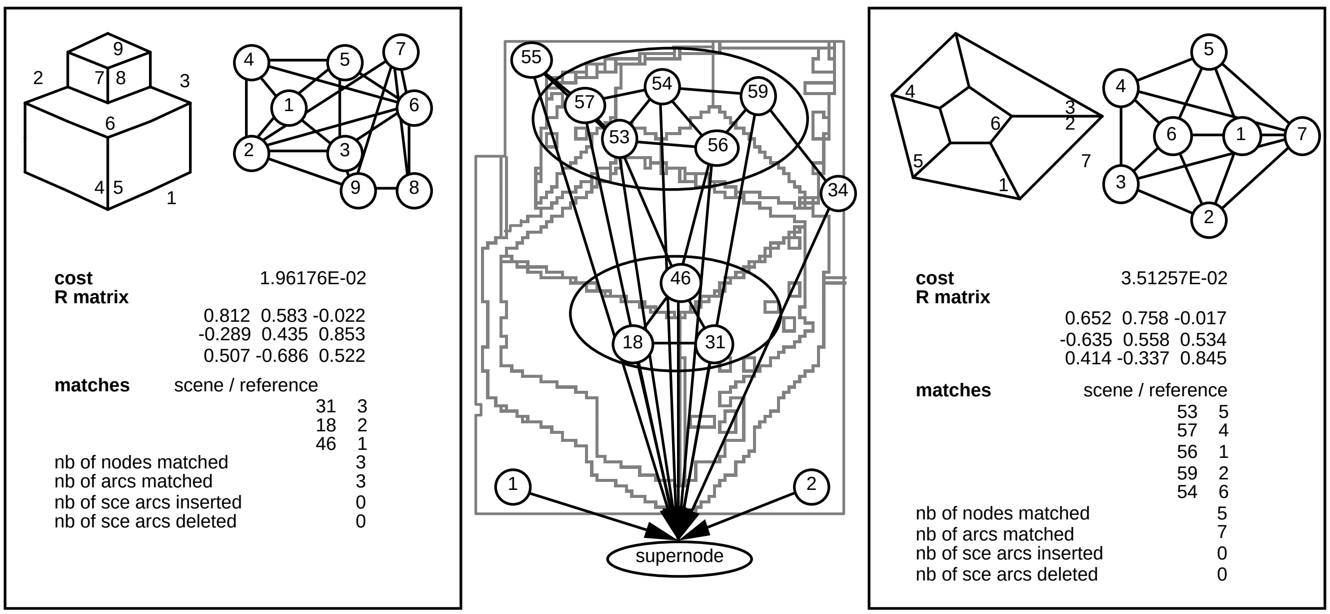

Figure 6.6: two compatible solutions: scene “CP” (center) vs. references "Cube1" (left) and "Pyramid2" (right)

With the algorithm described before, and a simple test of solution compatibility, the desired result, represented in figure 6.6, was obtained. The "best" set of compatible solutions contains the correct match between reference object "Cube1" and the part of scene "CP" representing the cube and the reference object "Pyramid2" with the part of the scene representing the pyramid.

Nevertheless, with our choice of high-level representation and correspondence-search criteria, the reliability of the correspondences decreases, in particular for the partially occluded objects. Match results must be completed, either by subsequent (low-level) verification or by the use of richer data in both the high-level representation and the correspondence-search criteria.

7. Conclusions and overview

In the present paper, we have shown a solution to the inexact high-level matching problem in the context of 3D-object recognition. Given two high-level object representations in the form of attributed relational graphs, we extract the most promising (partial) matches so that we can set up hypotheses on the scene objects. The match of the two subparts of reference and scene objects is characterized by a similarity measure that indicates match quality and a transformation - the geometrical rotation the reference object has to undergo so that it fits the scene data best.

We treat the subject of graph matching as a tree-search problem, using a "best-first" search. The present work is distinctive in the sense that it uses a very weak rule for tree expansion so that a wide spectrum of topological inexactitude can be accepted. Usually, this inexact matching feature requires an important increase of computational complexity. Therefore, criteria were introduced that enable pruning the search tree efficiently. The most important condition we use is rigidity. Despite fair results obtained with it, additional tree-pruning conditions based on the remaining graph attributes had to be added in order to deal with cases where the rigidity condition is not sufficient, e.g., for man-made objects where perpendicular and parallel faces are common.

H.P. Amann & H. Hugli, " An algorithm for the inexact matching of high level 3D polyhedric representations", Selected papers on model-based vision, Ed. Hatem Nasr, SPIE Milestone Series, Vol. MS 72, SPIE, Bellingham, WA, USA, pp. 313-321, 1992

A final match, consisting of a set of 1-1 node matches, is characterized by an overall similarity measure as well as by the geometrical transformation that best maps a reference object to the scene. Various experiments show sound results obtained with the described algorithm, i) for object recognition with single-object scenes and ii) for object recognition and object separation for multiple-object scenes, including objects which occlude themselves mutually.

Finally, the presented algorithm should not be considered a complete solution. Rather it should be considered as a building block for a more elaborate recognition system that also considers, among other things, verification of the generated hypotheses at a lower description level. A more elaborate system would also include object representations using additional information, e.g., descriptions of edges and surface shape.

8. Acknowledgements

We wish to thank our colleague Gilbert Maître who implemented the data-acquisition device as well as the range image segmentation algorithms used for this project.

This work has been supported by the Swiss National Foundation for Scientific Research under project numbers FN2265, FN2000-5.138 and FN 20-25572.

9. References

[1] W. H. Tsai, K. S. Fu, "Subgraph error-correction isomorphism for syntactic pattern recognition" IEEE Trans. on Systems, Man, Cybernetics, Vol 13, No 1, 1983

[2] M. A. Eshera, K. S. Fu, "A Similarity Measure between Attributed Relational Graphs for Image Analysis" 7th Conf. Pattern Recognition, Montreal, 1984

[3] M. A. Eshera, K. S. Fu, "An Image Understanding System Using Attributed Symbolic Representation and Inexact Graph Matching", IEEE PAMI-8, Vol 5, Sept. 1986

[4] L. Shapiro, R. M. Haralick, "A Metric for Comparing Relational Descriptions" IEEE PAMI-7, No 1, Jan. 1985

[5] O. D. Faugeras, M. Hebert, "The Representation, Recognition and Positioning of 3D-Shapes from Range Data" in "Three-Dimensional Machine Vision", ed T. Kanade, Kluwer Academic Press, 1987, pp 301

[6] M. Rioux, "Laser range finder based on synchronized scanners" Applied Optics, Vol 23, No 21, Nov. 1984, pp 3837-3844

[7] G. Maître, H. Hugli, F. Tièche, H.P. Amann, "Range Image Segmentation based on Function Approximation" Proceedings ISPRS-90, Zurich, Sept. 1990

[8] G. Maître, H. Hugli, "Range Image Segmentation by controlled-continuity Spline Approximation for Parallel Computing", Proc. SPIE Conf. on Intelligent Robotic Systems, Boston, Nov. 1991

[9] P. J. Besl, R. C. Jain, "Segmentation through variable-order surface fitting" PAMI-10, No 2, March 1988

[10] H. P. Amann, "3D Object Recognition Based on Surface Representations" PhD thesis, University of Neuchâtel (Switzerland), Oct. 1990

[11] M. Oshima, Y. Shirai, "An object recognition system using three-dimensional information" in "Three-Dimensional Machine Vision", ed T. Kanade, Kluwer Academic Press, 1987, pp 355

[12] I. L. Kantor, "Hypercomplex numbers: An elementary introduction to algebras" Springer, New York, 1989

[13] B. K.Horn, "Extended Gaussian Images" Proc. IEEE, Vol 72, No 12, Dec. 1984

[14] J. R. Ullmann, "An Algorithm for Subgraph Isomorphism" ACM, Vol 23, No 3, Jan. 1976, pp 31-42

About the authors

Hans Peter Amann received his MS in electrical engineering from the Swiss Federal Institute for Technology, Zurich, and PhD in signal processing from the University of Neuchâtel in 1982 and 1990 respectively. He has left this domain and now works on modeling and simulation of digital electronic devices.

Heinz Hugli received his MS in electrical engineering and PhD in image processing from the Swiss Federal Institute of Technology in 1971 and 1980 respectively. He is currently research director at the Institute of Microtechnology, University of Neuchâtel, where he teaches pattern recognition and robot vision and leads a research group.

H.P. Amann & H. Hugli, " An algorithm for the inexact matching of high level 3D polyhedric representations", Selected papers on model-based vision, Ed. Hatem Nasr, SPIE Milestone Series, Vol. MS 72, SPIE, Bellingham, WA, USA, pp. 313-321, 1992