

MODEL-BASED THREE DIMENSIONAL INTERPRETATIONS OF TWO DIMENSIONAL IMAGES

Rodney A. Brooks

Artificial Intelligence Laboratory
Stanford University, Stanford, Ca. 94305 USA

ABSTRACT

ACRONYM is a comprehensive domain independent model-based system for vision and manipulation related tasks. Many of its sub-modules and representations have been described elsewhere. Here the derivation and use of invariants for image feature prediction is described. We describe how predictions of image features and their relations are made and how instructions are generated which tell the interpretation algorithms how to make use of image feature measurements to derive three dimensional sizes and structural and spatial constraints on the original three-dimensional models. Some preliminary examples of ACRONYM's interpretations of aerial images are shown.

1. Introduction

At the previous IJCAI we reported [7] on the design and development of a model-based vision system called ACRONYM, which could identify instances of modeled objects in images. Since then the scope of ACRONYM has been increased to include extraction of three dimensional information from images (including monocular images) [6], reasoning about how to grasp objects (Binford [2]), and real time simulation of multiple manipulator work stations for purposes of off-line programming and the design and analysis of new manipulators (Soroka [12]).

To support these developments we have added a class and subclass relation representation scheme to the geometric modeling system ([9], [6]). This is based on the use of symbolic algebraic constraints. In support of this a constraint manipulation systems which includes a partial decision procedure on consistency of sets of non-linear inequalities was formulated and implemented [6]. A geometric reasoning system which can deal with underconstrained spatial relations was developed [8]. A new matcher which could manipulate the constraint systems was built for interpretation [9]. All of these systems were implemented in a mixture of MACLISP and a new rule system built for the purpose.

We have thus moved from a purely geometric representation and qualitative geometric reasoning system to a system with a combined algebraic and geometric representation and a geometric reasoning system which can make precise deductions about partially specified situations. The geometric and algebraic aspects of the representation complement each other during interpretation.

In this paper we deal with the techniques developed for image feature and feature-relation prediction, and then give some first examples (February 1981) of the performance of the new incarnation of ACRONYM on some images. The low level processes we currently use provide either little or noisy data. Nevertheless ACRONYM makes strong and accurate deductions about the objects appearing in the images. We expect even better performance when more accurate low level descriptive processes become available.

This work was supported by the Defense Advanced Research Projects Agency under contract MDA903-80-C-0102, by the National Science Foundation under contract DAR78-15914 and by a grant from the ALCOA corporation.

2. Prediction

In the ACRONYM system generic object classes and specific objects are represented by volumetric models based on generalised cones along with a partial order on sets of non-linear algebraic inequalities relating model parameters. Image features and relations between them which are invariant over variations in the models and camera parameters are identified by a geometric reasoning system. Such predictions are combined first to give guidance to low level image description processes, then to provide coarse filters on image features which are to be matched to local predictions. Predictions also contain instructions on how to use noisy measurements from identified image features to construct algebraic constraints on the original three dimensional models. Local matches are combined subject both to consistently meeting predicted image feature relations, and the formation of consistent sets of algebraic constraints derived from the image. The result is a three dimensional interpretation of the image.

This section describes some of the invariants that are identified by the reasoning system, and gives examples of how the back constraints are set up giving three dimensional information about the instances of the models which appear in images.

2.1 Constraints

To illuminate the discussion in succeeding sub-sections we briefly describe the uses and capabilities of ACRONYM's constraint mechanism and the allowed structure of constraints themselves.

ACRONYM's three-dimensional models are represented by *units* and *slots* (e.g. Bobrow and Winograd [4]). Any slot which admits numeric *fillers* also admits *quantifiers* (predeclared variable names) and expressions over quantifiers using the operators $+$, $-$, $\times$, $/$ and $\sqrt{}$.

Constraints can be put on quantifiers. They take the form of inequalities between expressions as defined above, along with the possibility of including max and min (on the left and right of $\leq$, respectively). Equality can be encoded as two inequalities. For instance suppose a cylinder is represented as a generalized cone whose straight spine has its length defined by the quantifier CYL_LENGTH and whose cross section is a circle with radius CYL_RADIUS. Then the class of all cylinders of volume 5 (in some units) can be represented by the two constraints:

$$\begin{aligned} 5 &\geq \text{CYL_LENGTH} \times \text{CYL_RADIUS} \times \text{CYL_RADIUS} \times \pi \\ 5 &\leq \text{CYL_LENGTH} \times \text{CYL_RADIUS} \times \text{CYL_RADIUS} \times \pi \end{aligned}$$

The ACRONYM constraint manipulation system (CMS), described in detail in [6], operates on sets of constraints. A set of constraints (implicitly conjunctive) defines a subset of n -dimensional space for which all constraints are true (where n is the number of quantifiers mentioned in the constraint set). This is called the satisfying set, and is empty if the constraints are inconsistent. The CMS is used for three tasks related to this constraint set.

1. Given a set of constraints partially decide whether their satisfying set is empty. The outcomes are "empty" or "I don't know".

2. Find numeric (or $\pm\infty$) upper and lower bounds on an expression in quantifiers over the satisfying set of a constraint set. This uses procedures called SUP and INF.

3. (A generalisation of 2.) For an expression E and a set of quantifiers V find expressions L and H in V such that $L \leq E \leq H$ identically over the satisfying set of the constraint set.

In 2 and 3 the expressions being bounded can include trigonometric functions such as sin, cos and arcsin. The CMS we have implemented in ACRONYM is a non-linear generalisation of the linear SUP-INF method described by Bledsoe [3], and Shostak [11]. It behaves identically to that described by the latter for purely linear sets of constraints and linear expressions. In addition it can often produce good bounds (numeric and expressions) on highly non-linear expressions in the presence of many non-linear constraints.

2.2 Shape prediction

We predict shapes as *ribbons* (the two dimensional analogue of three dimensional generalised cones) and *ellipses*. These are also the features which are found by the low level descriptive process we are temporarily using in ACRONYM.

Ribbons are a good way of describing the images generated by generalised cones. Consider a ribbon which corresponds to the image of the swept surface of a generalised cone. For straight spines, the projection of the cone spine into the image would closely correspond to the spine of the ribbon. Thus a good approximation to the observed angle between the spines of two generalised cones is the angle between the spines of the two ribbons in the image corresponding to their swept surfaces. We do not have a quantitative theory of these correspondences. Ellipses are a good way of describing the shapes generated by the ends of generalised cones. The perspective projections of ends of cones with circular cross-sections are exactly ellipses.

Shape prediction involves deciding what shapes will be visible, predicting ranges for shape parameters (to be used as a coarse filter during interpretation and also to guide the low level descriptive processes) and deriving instructions about how to locally invert the perspective transform and hence use image measurements to generate constraints on the original three dimensional models.

To predict the shapes generated by a single generalised cone, we do not explicitly predict all possible qualitatively different viewpoints. Rather we predict what shapes may appear in the image, and associate with them methods to compute constraints on the model that are implied by their individual appearance in an the image. For example, identification of the image of the swept surface of a right circular cone constrains the relative orientation of the cylinder to the camera (we call these back constraints). Identification of an end face of the cylinder provides a different set of constraints. If both the swept surface and an end face are identified then both sets of constraints apply. We also predict specific relations between shapes that will be true if they are both observed correctly. For more complex cones, the payoff is even greater for predicting individual shapes rather than exhaustive analysis of which shapes can appear together.

At other times during prediction invariant cases of obscuration are noticed. For instance it may be noticed that one cone abuts another so that its end face will never be visible. The consequences of such realisations are propagated through the predictions.

Prediction of shapes proceeds in five phases. First, all the contours on a generalised cone which could give rise to image shapes are identified by a set of special purpose rules. These include occluding contours and contours due purely to internal cone faces. Thus for instance a right square cylinder will generate contours for the end faces, the swept faces, and contours generated by the swept edges at diagonally vertices of the square cross section. The contours are generated independently of camera orientation, and in terms of object dimensions rather than image quantities.

Second, the orientation of the generalised cone relative to the camera (this is done by the geometric reasoning system, see [8], [6]) is then examined to decide which contours will be visible and how their image shapes will be distorted over the range of variations in the model parameters which appear in the orientation expressions.

The third phase predicts relations between contours of a single generalised cone (see section 2.3).

Fourth, the actual shapes are then predicted. The expected values for shape parameters in the image are estimated as closed intervals (see below).

Finally the back constraints which will be instantiated during interpretation are constructed.

2.2.1 Back constraints

Suppose that we wish to predict the length of an observable feature which is generated by something of length l lying in a plane parallel to the camera image plane, at distance d from the camera. Furthermore suppose the camera has a focal ratio of f . Then the measured length of the observed feature is given by $p = (l \times f)/d$. Any or all of l , f and d may be expressions in quantifiers, rather than numbers. Using the CMS we can obtain bounds on the above expression for image feature length, giving that it will lie in some range $P = [p_l, p_h]$ where p_l and p_h are either numbers or $\pm\infty$. For more complex geometries the expression for p will be more complex, but the method is the same (trigonometric functions are usually involved).

Now, given an image feature, which is hypothesized to correspond to the prediction we have to decide whether it acceptable on the basis of its parameters. The low level descriptive processes are noisy and provide an error interval, rather than an exact measurement for image parameters. Suppose the interval is $M = [m_l, m_h]$ for a feature parameter predicted with expression p . Then the parameter is acceptable if $P \cap M$ is non-empty. This is the coarse filtering used during initial hypothesis of image feature to feature prediction matches.

But note also that it must be true that the true value of p for the particular instance of the model which is being imaged must lie in the range M . Thus we can add the constraints:

$$\begin{aligned} m_l &\leq (l \times f)/d \\ m_h &\geq (l \times f)/d \end{aligned}$$

to the instance of the model being hypothesized, where l , f and d are numbers or expressions in quantifiers.

2.2.2 Trigonometric back constraints

When the expression p involves trigonometric functions the above method of generating back constraints will not work. It would generate constraints involving trigonometric functions, which our CMS can not handle.

One approach to this problem is to bound expression p above and below by expressions involving no quantifiers contained in arguments to trigonometric functions, and then use these expressions in setting up the back constraints. This has the unfortunate side effect of losing all information implied by the image feature about the quantifiers eliminated from the bounds.

A second approach is sometimes applicable. If a trigonometric function has as its argument e , an expression, and if the CMS determines that e is bounded to lie within a region of the function's domain where it is strictly monotonic and hence invertible, then specific back constraints on e can be computed at interpretation time (as distinct from during prediction). We illustrate with an example. A cylinder with length CYL_LENGTH is sitting upright on a table. A camera with unknown but constrained pan and tilt (the latter is constrained to lie in the interval $[\pi/12, \pi/6]$) is looking across from the side of the table, and it is elevated above table top height. The geometric details and numeric

constants are not important here. Suffice it to say that the geometric reasoning system deduces that the pan of the camera is irrelevant to the prediction of the length of the ribbon corresponding to the swept surface of the cylinder. It predicts that the length of the ribbon in the image will in fact be:

$$\frac{-2.42 \times \text{CYL_LENGTH} \times \cos(-\text{TILT})}{\text{CYLINDER.CAMZ}}$$

where 2.42 is the focal ratio of the camera and CYLINDER.CAMZ is an internal quantifier generated by the prediction module.

Both of the above approaches are used to generate back constraints to ensure coverage of all the relevant quantifiers. They are:

$$\begin{aligned} m_h &\geq -2.096 \times \text{CYL_LENGTH} \times (1/\text{CYLINDER.CAMZ}) \\ m_l &\leq -2.338 \times \text{CYL_LENGTH} \times (1/\text{CYLINDER.CAMZ}) \\ -\text{TILT} &\leq -\arccos(\sup(-0.413 \times m_h \\ &\quad \times \text{CYLINDER.CAMZ} \times (1/\text{CYL_LENGTH}))) \\ -\text{TILT} &\geq -\arccos(\inf(-0.413 \times m_l \\ &\quad \times \text{CYLINDER.CAMZ} \times (1/\text{CYL_LENGTH}))) \end{aligned}$$

The first two are non-trigonometric back constraints and at interpretation time a simple substitution of the measured numeric quantities for m_l and m_h is done. The latter two require further computation at interpretation time. After the substitution, expressions must be bounded over the satisfying set of all the known constraints, and the function $\arccos$ applied to give numeric upper and lower bounds on the quantifier TILT.

The techniques described here work for a more general class of functions than trigonometric functions (in the current implementation of ACRONYM we use it for functions $\sin$, $\cos$ and $\arcsin$). The requirement is that the domain of the function (e.g. the interval $[-\pi, \pi]$ for $\sin$ and $\cos$), can be subdivided into a finite number of intervals over which the function is strictly monotonic, and hence locally invertible.

2.3 Feature relation prediction

Image feature (shape) predictions are organised as the nodes of the prediction graph. The arcs of the graph predict image-domain relations between the features. During interpretation pairs are constructed which match image features and prediction nodes. They are coarsely checked for consistency by attempting to instantiate the prediction arcs. Some arcs also include back constraints which the instantiation of the arc implies about the model. These are treated in exactly the same manner as those associated with image feature predictions.

Prediction arcs are generated to relate multiple shapes predicted for a single cone. For instance a right circular cylinder prediction includes shapes for the swept surface and perhaps each of the end faces (depending on whether the camera geometry is known well enough to determine *a priori* exactly which faces will be visible). It can be predicted that a visible end face will be co-incident at at least one point in the image with a visible swept surface. (In fact a stronger prediction can be made: the straight spine of the swept surface image ribbon can be extended through the center of mass of the elliptical image of the end face.)

Prediction arcs are also generated between shapes associated with predictions for different generalised cones. These are actually of more importance in arriving at a consistent global interpretation of collections of image features as complex objects.

The semantics of the arc types we currently use are as follows.

2.3.1 Exclusive

If a generalised cone has a straight spine, and during sweeping, the cross section is kept at a constant angle to the spine, then at most one of the cone's end faces can be visible in a single image. *Exclusive* arcs relate image features which are mutually exclusive for this or other

reasons. (Note that in this case, instantiations of the two end faces would probably result in inconsistent back constraints being applied to the spatial orientation of the original model, so that eventually the CMS would detect an inconsistency. However checking for the existence of a simple arc at an early stage is computationally much cheaper than waiting to invoke the decision procedure.)

2.3.2 Colinear

If two straight line segments in three-space are *colinear* then any two-space image of them will either be a single degenerate point or two colinear line segments. As was pointed out earlier, the spine of the image shape corresponding to the swept surface of a cone is usually a good approximation to the projection of the spine of the cone into the image. Thus if two cones are known to have colinear spines in three dimensions, a colinear spine arc between the predictions of their swept surfaces can be included.

2.3.3 Coincident

If two cones are physically *coincident* at some point(s) in three-space, then for any camera geometry, if they are both visible then their projections will be coincident at some point(s) (except for some cases of obscuration). Failure to match predicted coincident arcs turns out to be the strongest pruning process during image interpretation.

2.3.4 Angle

If the *angle* between the spines of two generalised cones as viewed from the modeled camera is invariant over all the rotational variations in the model, or if an expression for the observed angle can be symbolically computed and is sufficiently simple, then a prediction of the observed angle can be made. For example wing-wing and wing-fuselage angles are invariant when an aircraft is viewed from above — this is because the only rotational freedom of an aircraft on the ground is about an axis parallel to the direction of view of an overhead camera. Again the fact that the projections of model spines correspond to image spines is used here. This arc type includes (trigonometric) back constraints which make use of the observed angle. Some such constraints constrain relative spatial orientations of generalised cones. Others provide constraints on the orientation of the plane of rotation, which generated the angle, relative to the camera, and hence constraints on an object's orientation relative to the camera.

2.3.5 Approach-ratio

Suppose a cone B is affixed at one end of its spine to another cone A , with a straight spine, somewhere along its length. The spines need not be coincident, but the cones must be. Suppose the spine of cone A has endpoints a_1 and a_2 , and let a_3 be the point on the spine of A closest to the end of the spine of B . Then the *approach-ratio* is the ratio of the length of the spine segment from a_1 to a_3 and the length of the complete spine from a_1 to a_2 . If the spines of A and B are both observable, then the approach-ratio is invariant under a normal projection for all camera geometries. Thus it is a *quasi-invariant* for a perspective projection for a camera sufficiently far from the object. For example, the ratio of the distance from the rear of the fuselage to the point of wing attachment, to the length of the fuselage, is almost invariant over all viewing angles for objects sufficiently far from the camera. Again this relies on the correspondences between the projection of a cone spine and the spine of the ribbon generated by the image of its swept surface. Approach-ratio arcs are only generated for pairs of image features which have a coincident arc. They provide back constraints on the model via the symbolic expression which describes the modeled spine approach ratio.

2.3.6 Distance

Sometimes symbolic expressions for the image distance between two image features can be computed. *Distance* arcs are only generated for pairs of image features which also have an angle arc, but no coincident arc. Distance arcs generate back constraints on the original model.




Fig. 3.1: Instances of class models of Boeing-747s and L-1011s.

2.3.7 Ribbon-contains

This is a directed arc type which two dimensionally relates two predicted ribbons, one of which will contain the other in the image. For instance, *ribbon-contains* arcs are built between the ribbon predicted from the occluding contour of a generalized cone with rectangular cross-section, and each of the ribbons generated by the two visible swept faces.

3. Some Image Interpretations

The image interpretations reported here are of a rather preliminary nature. They were carried out when the various sub-systems had only been running together for about two weeks. Further experimentation has been hampered by address space limitations — the current system occupies two 256K address spaces on a DEC-10. An effort is underway to transport the system to a VAX.

In the examples to be described here ACRONYM was given a generic model of wide-bodied passenger jet aircraft, along with class specializations to L-1011s and Boeing-747s. The Boeing-747 class had further subclass specializations to Boeing-747B and Boeing-747SP. The subclasses are not completely partition their parent classes. The classes are described by sets of constraints on some 30 quantifiers. Figure 3.1 shows instances of the two major modeled classes of jet aircraft. These diagrams were drawn by ACRONYM from the models given it to carry out the image interpretations. The constraints for the generic class of wide bodied jets are given in figure 3.2. Units are meters.

The camera was modeled as being between 1000 and 12000 meters above the ground. Thus there is little *a priori* knowledge of the scale of the images. A specific focal ratio was given: 20. (Similar interpretations have been carried out with a variable focal ratio, but then the final constraints on camera height and focal ratio are coupled, and not as clear for illustrative purposes — no accuracy is lost due to the non-linearities that are introduced into the constraints, although both computation time and garbage collection time are increased.)

The aircraft models, the camera model and the number of pixels in each dimension of the image (512×512 in these examples) were the only pieces of world knowledge input to ACRONYM. It has no special knowledge of aerial scenes: all its rules are about geometry and algebraic manipulation. These were applied to the particular generic models it was given, to make predictions and then to carry out interpretations.

Figures 3.3 through 3.5 show three examples of interpretations carried out by ACRONYM. In each case part *a* is a half-tone of the original grey level image. The *b* version is the result of applying the line finder of Nevatia and Babu [10]. That line finder was designed to find linear features such as roads and rivers in aerial photos. Close examination of results on these images indicate many errors, and undue enlargement in width of narrow linear features. It also produces many noise edges in smooth brightness gradients (not visible at the resolution of the reproductions of these figures). These edges are the lowest level input to ACRONYM.

An edge linker [5] is directed by the predictions to look for ribbons and ellipses. In this case there is very little *a priori* information about the scale of the images. The *c* versions of each figure show the ribbons fitted to the linked edges when it is searching for candidate matches for the fuselage and wings of aircraft. There is even further degradation of image information at this stage. This is the only data which the ACRONYM reasoning system is given to interpret. Notice that in the figure 3.5 almost all the shapes corresponding to aircraft are lost. Quite a few aircraft in 3.4 are lost also. Besides losing many shapes, the combination of the edge finder and edge linker conspire to give very inaccurate image measurements. We assume all image measurements have a $\pm 30\%$ error, except that for very small measurements, we assume that pixel noise swamps even those error estimates. Then the error is estimated to be inversely proportional to the measurement with a 2 pixel measurement admitting a 100% error. Thus the data which ACRONYM really gets to work with is considerably more fussy than indicated by the the *c* series of figures.

We intend to make use of new and better low level descriptive processes being developed in our laboratory by other researchers as soon as they become robust enough for every day use (e.g. Baker [1] whose descriptions from stereo will also include surface and depth information).

Despite this very noisy descriptive data ACRONYM makes good interpretations of the images. The *d* series of figures show its interpretations with the ribbons labeled by what part of the model they were matched to. (The numbers which may be unreadable in 3.3d show the groupings into individual aircraft.)

ACRONYM first uses the most general set of constraints, those associated with the generic class of wide-bodied jets, when carrying out initial prediction and interpretation. Interpretation adds additional constraints for each hypothesized aircraft instance. For example in finding the correspondences in figure 3.4d constraints were added which eventually constrained the WING_WIDTH (the width of the wings where they attach to fuselage) to lie in the range [7, 10.5677531] compared to the modeled bounds of [7, 12]. The height of the camera, modeled to lie in the range [1000, 12000] is constrained by the interpretation to the range [2199, 3322].

Once a consistent match or partial match to a geometric model has been found in the context of some set of constraints (model class), it is easy to check whether it might also be an instance of a subclass. We need only add the extra constraints associated with the subclass and check for consistency with those already implied by the interpretation using the CMS as described in section 2.1. The aircraft located in 3.4d is consistent with the constraints for an L-1011, but not for a Boeing-747. Examination of the images by the author had previously indicated that the aircraft was an L-1011. The additional symbolic constraints implied by accepting that the aircraft is in fact an L-1011 propagate through the entire constraint set. Although the constraints describing an L-1011 do not include constraints on camera height, the back constraints deduced during interpretation relate quantifiers representing such quantities as length of the wings to the height (and focal ratio in the more general case). Thus the height of the camera is further constrained in 3.4d to lie in the range [2356, 2489]. Recall that all image measurements were subject to $\pm 30\%$ errors, and that this estimate has taken all such errors into account.

Figure 3.3d indicates matches were found for three airplanes. Examination of the data in 3.3c indicates that this is the best that could be expected. Note however that only partial matches were found in all three cases. For such small ribbons errors were apparently larger than the generous estimate used. The fuselage ribbon in the leftmost aircraft (number 1) for instance fails to pass the coarse filtering stage. Despite the partial match, this particular aircraft is found to be consistent with the constraints for an L-1011, but not consistent with those of a Boeing-747. Again this is correct.

The other two aircraft identified are even more interesting. The author had thought from casual inspection of the grey level image

that they were instances of Boeing-747s. They both gave matches consistent with the class of wide-bodied jets. As expected neither was consistent with the extra constraints of an L-1011. However, although each individual parameter range from the interpretation constraint sets was consistent with the individual parameter value or range for the class of Boeing-747s, neither set of constraints was consistent with that subclass (the constraints contain much finer information than just the parameter ranges — in the same manner as in the example above where constraints on wing length propagate to constrain the camera height). On close examination of the grey-level image it was determined that the aircraft were not in fact Boeing-747's. The author used the fact that they were much smaller than the L-1011 to make that deduction, but ACRONYM made the deduction at the local level before considering comparisons between aircraft.

The aircraft (probably Boeing-707's) are in fact too small to be wide-bodied jets of any type. Since the scale of the image is unknown *a priori* this can not be deduced locally. However it is reflected in the height estimates derived at the local level — [5400, 8226] interpreting the L-1011 just as a generic wide-body, ([5786, 6170] as an L-1011), and [9007, 11846] for the rightmost aircraft. Thus ACRONYM deduces that either the left aircraft is a wide-body and the others are not, or the right two are wide-bodies and the left one is not (it is too big).

Finally note that geometrically there were other candidates for aircraft in the ribbons of figure 3.3c. For instance the wing of the aircraft just to the right of those identified and a ribbon found for its passenger ramp could be the two wings of an aircraft with a fuselage missing between them. In fact these two ribbons were instantiated as an aircraft on the basis of the coarse filters on the nodes and arcs. However the set of back constraints they generated were mutually inconsistent.

Thus we can see from the examples that even with very poor and noisy data the combined use of geometry and symbolic algebraic constraints can lead to accurate image interpretations. The system should be tested on more accurate low level data to fully evaluate the power of this approach.

REFERENCES

- [1] Baker, H. Harlyn; *Edge Based Stereo Correlation*, Proceedings ARPA Image Understanding Workshop, College Park, MD, April 1980, 168-175.
- [2] Binford, Thomas O.; *Computer Integrated Assembly Systems*, Proceedings NSF Grantees Conference on Industrial Automation, Cornell Univ., Sep. 1979.
- [3] Bledsoe, W. W.; *The Sup-Inf Method in Presburger Arithmetic*, Memo ATP-18, Dept. of Math. and Comp. Sci., University of Texas at Austin, Austin, Texas, Dec. 1974.
- [4] Bobrow, Daniel G. and Terry Winograd; *An Overview of KRL, a Knowledge Representation Language*, Cognitive Science 1, 1977, 3-46.
- [5] Brooks, Rodney A.; *Goal-Directed Edge Linking and Ribbon Finding*, Proceedings ARPA Image Understanding Workshop, Menlo Park, April 1979, 72-76.
- [6] Brooks, Rodney A.; *Symbolic Reasoning Among 3-D Models and 2-D Images*, AI Journal, to appear.
- [7] Brooks, Rodney A., Russell Greiner and Thomas O. Binford; *The ACRONYM Model-Based Vision System*, Proceedings IJCAI-6, Tokyo, August 1979, 105-113.
- [8] Brooks, Rodney A. and Thomas O. Binford; *Representing and Reasoning About Partially Specified Scenes*, Proceedings ARPA Image Understanding Workshop, College Park, MD, April 1980, 95-103.
- [9] Brooks, Rodney A. and Thomas O. Binford; *Interpretive Vision and Restriction Graphs*, Proceedings of the First Annual National Conference on Artificial Intelligence, Stanford, August 1980, 21-27.
- [10] Nevatia, Ramakant and K. Ramesh Babu; *Linear Feature Extraction and Description*, Computer Graphics and Image Processing 13, 1980, 257-269.
- [11] Shostak, Robert E.; *On the SUP-INF Method for Proving Presburger Formulas*, JACM 24, 1977, 529-543.
- [12] Soroka, Barry I.; *Debugging Manipulator Programs with a Simulator*, Autofact West Conference, Society of Manufacturing Engineers, Anaheim, Nov. 1980.

```

ENG-DISP-GAP e [6, 10]
ENG-DISP e [0, 4]
ENG-GAP e [7, 10]
STAB-ATTACH e [3, 5]
R-ENG-ATTACHMENT e [3, 5]
ENG-OUT e [5, 12]
WING-ATTACHMENT e [20, 40]
    WING-ATTACHMENT ≥ 0.4*FUSELAGE-LENGTH
    WING-ATTACHMENT ≤ 0.6*FUSELAGE-LENGTH
STAB-RATIO e [0.2, 0.65]
STAB-SWEEP-BACK e [3, 7]
STAB-LENGTH e [7.6, 13]
STAB-THICK e [0.7, 1.1]
STAB-WIDTH e [5, 11]
RUDDER-RATIO e [0.3, 0.4]
RUDDER-SWEEP-BACK e [3, 9]
RUDDER-LENGTH e [8.6, 14.2]
RUDDER-X-HEIGHT e [7, 13]
RUDDER-X-WIDTH e [0.7, 1.1]
WING-RATIO e [0.35, 0.45]
WING-THICK e [1.5, 2.5]
WING-WIDTH e [7, 12]
    WING-WIDTH ≤ 0.5*WING-LENGTH
WING-LIFT e [1, 2]
WING-SWEEP-BACK e [13, 18]
WING-LENGTH e [22, 33.5]
    WING-LENGTH ≥ 2*WING-WIDTH
    WING-LENGTH ≥ 0.43*FUSELAGE-LENGTH
    WING-LENGTH ≤ 0.65*FUSELAGE-LENGTH
REAR-ENGINE-LENGTH e [6, 10]
ENGINE-LENGTH e [4, 7]
ENGINE-RADIUS e [1, 1.8]
FUSELAGE-RADIUS e [2.5, 4]
FUSELAGE-LENGTH e [40, 70]
    FUSELAGE-LENGTH ≥ 1.66666666*WING-ATTACHMENT
    FUSELAGE-LENGTH ≥ 1.53846154*WING-LENGTH
    FUSELAGE-LENGTH ≤ 2.5*WING-ATTACHMENT
    FUSELAGE-LENGTH ≤ 2.3256814*WING-LENGTH
R-ENG-QUANT e [0, 1]
    R-ENG-QUANT ≤ 2 + -1*F-ENG-QUANT
F-ENG-QUANT e [1, 2]
    F-ENG-QUANT ≤ 2 + -1*R-ENG-QUANT

```

Fig. 3.2



Fig. 3.3a



Fig. 3.4a



Fig. 3.5a



Fig. 3.3b




Fig. 3.4b



Fig. 3.5b




Fig. 3.3c




Fig. 3.4c




Fig. 3.5c




Fig. 3.3d




Fig. 3.4d




Fig. 3.5d