Planning dextrous hand precision grasps from range data, using preshaping and finger trajectories

David Wren and Robert B. Fisher

Department of Artificial Intelligence, University of Edinburgh,
5 Forrest Hill, Edinburgh EH1 1JG, Scotland
{davidw,rbf}@aifh.ed.ac.uk

Inspired by human grasping behaviour, use of hand preshapes has for some time been recognised as a useful way of reducing the complexity of planning dextrous hand grasps. It is common for preshapes to be chosen and fitted to the object using coarse object characteristics. In this paper we use range data to provide the more detailed knowledge of surface geometry required when planning precision grasps. We define sets of preshapes, with associated modes of fingertip closure, and give criteria for the choice and positioning of the preshapes, based on intersecting fingertip trajectories with the graspable surfaces of the object. Grasps are planned in simulation using real range data and an anthropomorphic hand model. Our results show that complex objects can be grasped using relatively simple, fast algorithms when the hand movements are constrained in this way.

1 Introduction

In this paper we are concerned with the planning of robot grasps with a dextrous hand. Parallel jaw grippers, with one degree of freedom, can pick up a wide range of objects. We should therefore have good reasons before using dextrous hands instead of such grippers. Clearly, dextrous hands are essential if the robot is to manipulate the object between its fingers. However, dextrous hands have big advantages over parallel jaw grippers even when the aim is just to acquire the object in a stable grasp:

- Due to the increased number of hand-object contacts, dextrous hand grasps can be more stable than parallel jaw gripper grasps. Put another way, the issue of stability becomes less crucial as the number (and area, in the case of soft-fingered hands) of hand-object contacts increases.
- A dextrous hand can grasp an object from a much wider range of wrist positions and approach directions than can a parallel jaw gripper. This is useful in a cluttered environment, and also facilitates easier integration of arm motion planning and grasp planning.

Napier [8] introduced the concept of *precision* and *power* grasps. Lyons [7] further refined these definitions for robotics, by introducing the *lateral grasp* in which the insides of the distal (*i.e.* final) link of each finger contacts the object. This lies between the precision and power grasp in terms of manipulability and stability. We will refer to the lateral grasp as a subset of the precision grasps, since there is still just one region of contact per finger.

Robot grasp planning can be very complex. When mounted on a 6 degree of freedom arm/wrist combination, the two most widely used dextrous hands, the Salisbury [9] and the Utah/MIT [3] hands, have 15 and 22 degrees of freedom respectively. In recent years the preshaping paradigm has been widely recognised as a useful way to ease the complexity of the problem of finding satisfactory values for the degrees of freedom. A hand preshape is the finger posture adopted as the wrist moves towards the object. The grasp is then executed by placing the wrist into a position that encompasses the object, and then flexing (*i.e.* closing) the fingers until they make contact with the object.

Stansfield [11] preshaped a Salisbury hand using a knowledge-based system, which is used to grasp polyhedral and simple curved objects, using data from two CCD cameras and haptic exploration to give 3D edges and 2D regions. Lyons [6, 7], given the desired distances between fingertips, used potential fields to shape the hand, but does not derive the desired distances from actual data. Bard *et al* [1] used preshape to plan power grasps on objects. The objects are modelled using *elliptical cylinders*, which are especially suitable for planning power grasps. The preshapes are planned using a set of heuristics based on the properties of the elliptical cylinders.

When using the precision paradigm, it is not enough to simply make a decision about which preshape to use; we must also define the trajectories taken by the fingers as they are flexed from the preshape to form the grasp. This is much more important for precision grasps than it is for power grasps, because power grasps have more possible hand-object contacts, whereas here we need to ensure that the fingers contact graspable surfaces with a good contact orientation.

1.1 Grasp Strategies

A *grasp strategy* consists of a preshape and a set of finger trajectories, from which a grasp can be formed without movement of the robot wrist. The preshape is a prescribed hand configuration and the finger trajectories are the motions of the fingers after the preshape is formed and the wrist position has been fixed.

Grasp strategies constrain the range of possible finger movements whilst still allowing a sufficient number of degrees of freedom to be able to cope with a wide range of objects (i.e. though the finger trajectories are specified, each finger can be stopped at any point along its trajectory by contact with the object). The grasping problem can thus be decomposed into two stages:

1. Choice of grasping strategy given the task, object geometry and range of possible wrist approaches.
2. Placement of the wrist such that the grasping strategy's finger trajectories intersect with graspable features on the object.

This approach reduces the complexity of the problem whilst preserving the flexibility of a dextrous hand. Rather than viewing a dextrous hand as a completely general device, it is viewed as a set of flexible tools, where each grasping strategy is a different "tool". The decomposition also allows us to examine human grasping behaviour as an successful example. Such research can lead to useful representations of hand configurations (*e.g.* [5]). The use of a limited number of prescribed trajectories means that this approach can also be utilised with dextrous robot hands that are simpler and cheaper than the Salisbury or Utah/MIT hands.

It should be noted that this approach de-emphasises stability analysis, which is the focus of much research into robot grasping. The grasp strategies are devised to make it likely that any grasp thus formed is stable. The preshapes are symmetrical, in that all digits are flexed identically, and the fingers are abducted by the same amount in opposite senses. In our hand model there are three digits, so the fingertip contacts lie in a plane. Assuming there is enough friction to resist gravity, we can usefully limit stability analysis to the plane. For contact normals perpendicular to the fingertip trajectories, in the preshape formation the grasp will be in equilibrium and stable. As the contact normals deviate from this ideal, or as the fingers move along the fingertip trajectories, friction is required to make the grasp stable in the plane. It is difficult to usefully analyse the relative stabilities of different grasping strategies, since the final stability depends very much on object geometry. However, if the robot hand has soft, deformable fingers [10] then with suitable force control a wide range of finger positions and contact orientations can be made stable. In this work we therefore concentrate on planning kinematically accessible grasps; however, we do address stability from the side in our measure of quality of preshape fits.

1.2 Hand Model

Figure 1 shows the hand model and the two sets of precision grasping strategies used in our experiments. The hand is roughly anthropomorphic, in that the finger dimensions are similar to that of the human hand, and certain joints are coupled.

The hand has a thumb and two fingers, all with identical dimensions and, in the absence of external forces, each with equal angles of flexion (α, β, γ at the proximal, middle and distal joints respectively). The fingers can *abduct* by an angle δ (*i.e.* they can rotate about an axis perpendicular to the palm and through the proximal joint). Within each finger, the distal and middle joints are coupled, and the angles of abduction of each finger are coupled. This coupling of joints could be hard-wired into the hand design, in which case they would have a strong influence on how the object could be manipulated, or just used as modes of movement to facilitate the planning of the grasps with complex hands such as the Salisbury or Utah/MIT hand.

In the two strategies, the preshape is defined by just two variables — one controlling the flexion of the fingers, another controlling the abduction of the fingers. The finger closure trajectories are defined by just one variable. In the case of strategy set $P1$ this is the proximal joint; in the case of set $P2$ it is the proximal joint for the preshape, and the middle joint for finger closure.

In the grasp strategies, the joint motions are constrained to give a kinematics similar to that of the human hand when precision grasping. In strategy set $P1$, the fingers are flexed at the proximal joint; i.e. it might be used when grasping a book by the sides. In strategy set $P2$ the proximal joint is fixed, and the fingers are flexed at the middle and distal joints; i.e. this strategy might be used for a lifting grasp, with hooking of the fingers under the edges of the object. Set $P1$ can grasp a wider range of objects than can set $P2$, due to an increased distance between tips and palm and longer finger trajectories during closure. However, set $P2$ provides two important functions: it can grasp wider objects than can set $P1$, and it tends to produce lateral grasps, *i.e.* grasps formed with the inside distal segments. Contacts formed with strategy set $P1$ tend to produce more grasps at the very end of the fingertips. Choice of grasping strategy can therefore be, to some extent, task driven, though ultimately which particular type of fingertip grasp is produced is determined by object geometry, not task requirements.

The angle of the distal link is constrained so that it cannot point away from the "centre" of the preshape. This is done to make it less likely that collisions occur between the object surface and the hand — indeed, in our experiments no collision checking is implemented, and all the grasps planned are feasible.

Strategy Set P1	preshape	$\pi/12 < \alpha_0 < \pi/2$	$\beta_0 = \pi/4$	$\gamma_0 = \pi/6$	$0 < \delta_0 < \pi/2$
	trajectory	$\alpha_0 < \alpha < \pi/2$	$\beta = \beta_0$	$\gamma = \gamma_0$	$\delta = \delta_0$
Strategy Set P2	preshape	$0 < \alpha_0 < \pi/2$	$\beta_0 = \frac{3}{10}\pi - \frac{3}{5}\alpha$	$\gamma_0 = \frac{2}{3}\beta_0$	$0 < \delta_0 < \pi/2$
	trajectory	$\alpha = \alpha_0$	$\beta_0 < \beta < \pi/2$	$\gamma = \frac{2}{3}\beta$	$\delta = \delta_0$

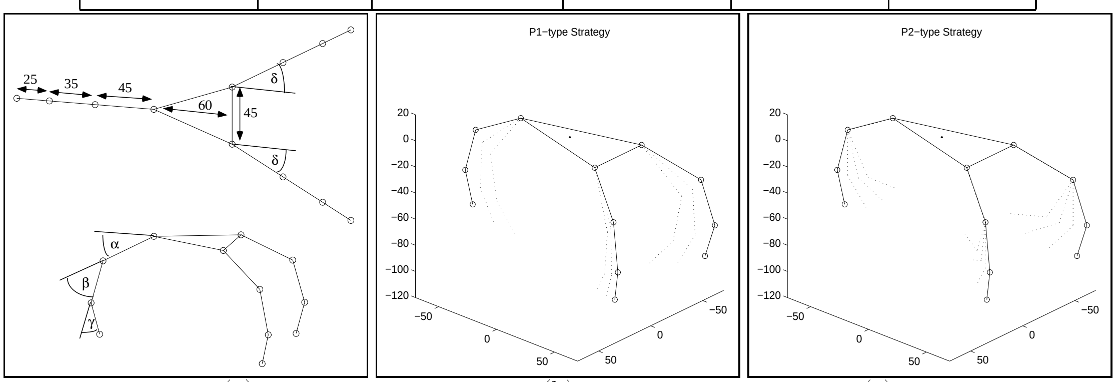

(a) (b) (c)

Fig. 1.: (a) Hand model showing the link distances and the location of angular degrees of freedom, (b) $P1$-type grasp strategy (clasping), (c) $P2$-type grasp strategy (grabbing). All distances shown in *mm*. Refer to text for more details.

The decision of which precision grasping strategy to use cannot be made using coarse volumetric descriptions alone (which can be sufficient to plan *power* grasps), but requires knowledge of the surface geometry of the object. We use a range sensor — a laser striper — to obtain the surface data (consisting of a set of quadric surface patches with both intrinsic shape parameters and surface extent), which is then processed to obtain *graspable features*, namely sufficiently large surface features with curvature characteristics that provide good fingertip contact for soft (*e.g.* rubber-coated) fingers [13].

2 Algorithm

Experiments were conducted in simulation using object data acquired from real objects, and the hand model shown in Figure 1. The current implementation works on polyhedral objects; the principles, however, can be and will be applied to general curved objects.

2.1 Input

A range image of the object is taken from two known viewpoints using a laser striper. The images are taken at a resolution of $1mm$ in the x and y directions, with 256 different z values, the resolution of which depends of the height of the object.

The range data is then segmented by fitting planes and quadric surfaces, grown from seeds acquired by an (H, K) curvature classification procedure (e.g., following [2]). The raw data, the analytical descriptions of the surface patches and the patch boundaries for each view are rotated into the global coordinate frame to provide the visual input to the grasp planner. No fusion of patches occurs, since oversegmentation does not affect the grasp planner other than to make it slower. Nor are there any checks for patches overlapping because, again, this does not affect the planning process. Building precise visual models is difficult; fortunately the type of models required for grasping have much laxer requirements than those required for object recognition. For example, repeatability of segmentation is not required.

The graspable features are then grouped into candidate grasping sets. These sets comprise of 2 or 3 features, according to whether the fingers are placed on the same or a separate feature respectively. We find all sets where opposition between candidate thumb and finger patches exists. Some of these sets will not be reachable by the hand; however, the subsequent algorithm is fast enough not to make this a problem.

2.2 Determination of preshape

Given a candidate grasping set, we calculate the desired *preshape aperture*. We define the *preshape aperture* to be the distance of the thumb-tip to the point midway between the two fingertips.

The preshape aperture is chosen such that the graspable features will be reachable and hence graspable over their whole extent, in order to maximise the leeway in wrist placement. We use three observations based on human grasping behaviour in making this choice.

Firstly, the minimum size of the aperture is proportional to the expected distance between the graspable features (see [4]). Secondly, the aperture is kept as small as possible, taking into account the uncertainty in wrist and object position (see [12]). Thirdly, the plane of the preshape aperture is constrained to lie parallel to the "aperture plane", which is determined from the object position with respect to the robot arm and the environment. For many objects and tasks the aperture plane is extremely simple to determine, *e.g.* for objects lying on the ground the best aperture plane will almost always be parallel to the ground, in order to give maximum leeway with respect to hand-ground collisions. In the experiments described here, therefore, the aperture plane is always parallel to the ground plane.

However, just using these constraints would lead to overly large apertures, as it does not take into account the angle of the preshape with respect to the contact normals. We therefore introduce another constraint, which is to keep the projection of the thumb-tip trajectory in the aperture plane normal to the intersection of the contact surface with that plane. This assumption is made throughout the analysis. This recognises the fact that the thumb is the single most important digit in terms of grasp stability and, as such, to avoid slippage should be kept normal to the contact surface.

Once the preshape aperture has been determined, a grasp strategy is then chosen from each of the strategy sets suitable for the task in hand (in our case, a precision grasp). The abduction of the preshape is determined from the contact normals, so as to minimise the angle between finger trajectories and surfaces in the plane of the aperture (assuming thumb contact normal to the surface). The flexion parameter of the preshape can then be uniquely determined by the preshape aperture.

2.3 Fitting of preshape

Each grasping strategy is then fitted to the graspable features, such that the finger trajectories intersect the features and are free of collisions as they move to do so. To fit the grasping strategy, we minimise a distance metric. The ideal metric would be the variance of the distances, along fingertip trajectories, from preshape to contact. However, for ease of computation, we use a slightly modified distance metric — the variance of the distances along the *projections* of the fingertip trajectories in the aperture plane. This simplification brings with it the danger of having the trajectories miss the graspable surfaces, so the distance of the aperture plane from the origin is bounded to maximise the overlap, normal to the aperture plane, of the fingertip trajectories and the targeted graspable surfaces. This uniquely determines the wrist position for each grasping strategy.

The minimisation procedure is performed by sampling hand positions at various points over the thumb patch. It starts with the thumb-tip at the centre of the patch, and the hand will only be shifted in order to decrease the distance metric by an amount deemed "useful". This is to avoid the hand shifting to the edge of a patch in order to lose only a very small amount of the distance metric. We wish to keep the contact points close to the centres of the patches to make the grasp robust with respect to positioning error.

It is then verified that the fingertip trajectories do intersect with the desired grasping surface. If they do not intersect within the boundary of the surface patch, the aperture is shifted first down and then up until it does intersect. If there are still no intersections, the algorithm reports a failure to plan a grasp.

The distance metric helps give stable grasps, because it tends to equalise the distances travelled along fingertip trajectories. This tends to pull the hand into a symmetrical configuration, which increases the stability of the grasp, subject to surface geometry.

It should be noted that the better the initial estimate of the preshape, the faster and more reliable the determination of wrist position.

3 Results

We show results for a selection of polyhedral objects, using the two precision grasping strategies shown in Figure 1. The object patch boundaries are displayed. The joints and fingertips of the hand are drawn as circles, the hand segments as straight lines. Curves joining fingertips to surface patches show the trajectory of the fingertips from the preshape position to the final contact point.

Figure 2 shows the preshape for each strategy fitted to a simple cuboid. The fingertip trajectories are traced, showing that they do indeed intersect the object. This is not surprising, given that the object has the simplest geometry possible; indeed this is the "ideal" object for the preshape planning procedure we use, in that it enables the preshape to get as close as possible to the desired grasping regions - the small gap between preshape and object is due to the leeway measurement to allow for finger radius and wrist position uncertainty. It can be seen that the $P2$ strategy is useful for implementing a lateral grasp, whereas the $P1$ strategy produces a grasp of higher precision.

Figure 3 shows the ability of the preshape to *abduct* in order to give normal finger-surface contacts. Again, the $P2$ strategy gives a more lateral grasp.

Figure 4 shows a grasp planned on a more complex object. There are a variety of possible two-feature and three-feature grasps on this object. There are 6 planes visible, and from these there are 8 two-feature sets and 4 three-feature sets. A preshape can only be fitted to one of the two-feature sets, because the algorithm can only abduct the fingers in response to patch normals, not in order to access small patches. This leaves the 4 three-feature sets, of which three are graspable.

Figure 4 shows the results for one of the three-feature sets: (a) shows the range data from two different views combined, (b) shows the planned $P1$-type grasp, (c) the same grasp viewed from above. For this grasp it is crucial to consider the fingertip trajectories: Figure 4(d) shows an attempted grasp with the same preshape but *different* fingertip trajectories (the fingers are actuated at the middle and distal links, as in the $P2$-type strategy). Note how the fingertip trajectories actually fail to intersect with the object — this strategy therefore fails.

The quality of this grasp, in terms of both stability and task-specificity, is lowered by the large deviation of the final grasp configuration from the task-specific preshape and by the relative orientations of the grasping features. Analysis of the static mechanics of the grasp shows that it requires a minimum coefficient of friction $\mu = 0.32$.

4 Conclusions

We have introduced the notion of associating finger closure modes with preshapes. Firm criteria, based on human grasping and robustness with respect to positioning, have been given for deciding on preshape parameters. Preshapes are then fitted with the aim of minimising the variance of the fingertip distances from preshape to contact. This helps give stable grasps, and helps avoid collisions.

The algorithm described has been implemented in simulation, and typically takes a few seconds on a Sun Sparc 10 to process a candidate set of grasping features, at a thumb-tip position sampling resolution comparable to the striper resolution. While the examples in the previous section used polyhedral objects, the algorithm can be applied to non-polyhedral objects. In fact, all that is required is to have enough locally smooth surface patches (which may be connected by rather irregular surface elsewhere).

The current algorithm's only major failing is that it does not allow the fingers to abduct in response to the relative positions of finger contacts, which means that it fails to grasp some small patches. This can probably be overcome by providing some criteria to weigh up the relative requirements of stability and accessibility.

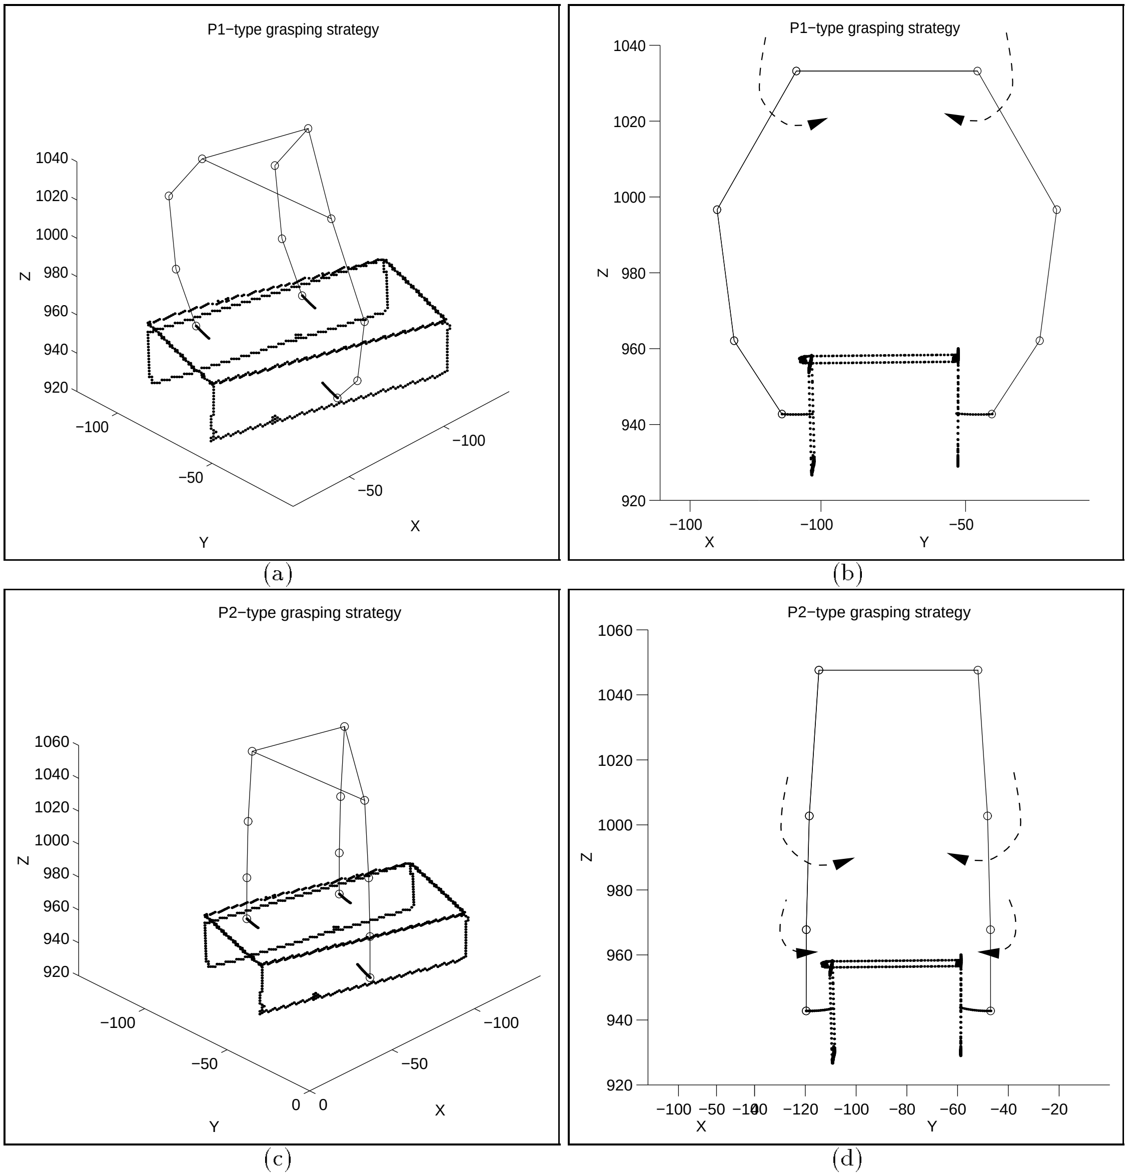

Fig. 2.: Two grasps on a cuboid. (a) $P1$-type grasp , (b) viewed side-on, (c) $P2$-type grasp, (d) viewed side-on. The small segment at the end of each finger is the trajectory of the fingertip from the preshape position to the point of contact.

Further work will also examine the choice of aperture plane. In the work here the plane is fixed to be parallel to the ground. For some objects, *e.g.* cylinders with their axes vertical (such as mugs), there are other good aperture planes to be used — in the case of the vertically aligned cylinder, any plane which goes through the axis of the cylinder is a good aperture plane.

Work is currently underway to extend the algorithm to deal with objects described by general quadric surfaces; the only difficulties in this are implementational — the criteria used to form and fit the preshapes remain the same.

References

1. C. Bard, C. Bellier, C. Laugier, J. Troccaz, G. Vercelli, and B. Triggs. Achieving dextrous grasping by integrating planning and vision based sensing. Technical report, LIFIA-IRIMAG, 1993.
2. P.J. Besl. *Surfaces in Range Image Understanding*. Springer-Verlag, 1988.

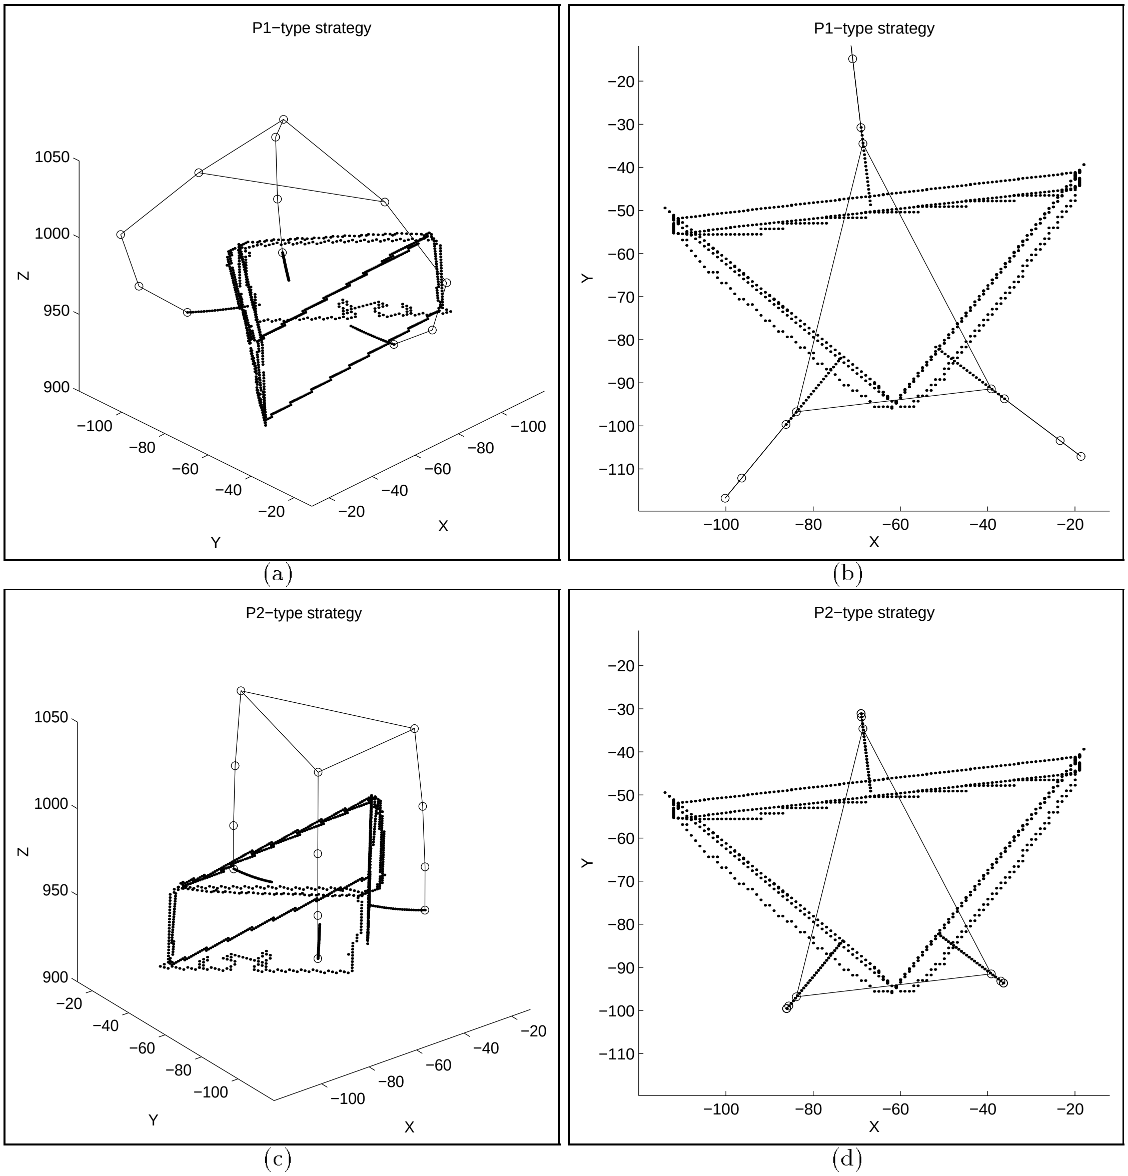

Fig. 3.: Two grasps on a triangular prism. (a) $P1$-type grasp , (b) viewed from above, (c) $P2$-type grasp, (d) viewed from above. The small segment at the end of each finger is the trajectory of the fingertip from the preshape position to the point of contact.

3. S.C. Jacobsen, E.K. Iverson, D.F. Knutti, R.T. Johnson, and K.B. Biggers. Design of the utah/m.i.t. dextrous hand. In *IEEE International Conference on Robotics and Automation*, volume 3, pages 1520–1532, 1986.
4. M. Jeannerod. Intersegmental coordination during reaching at natural visual objects. In *Attention and Performance IX*. Erlbaum, Hillsdale, NJ, 1981.
5. S.B. Kang and K. Ikeuchi. A framework for recognising grasps. Technical Report CMU-RI-TR-91-24, Carnegie Mellon University, November 1991.
6. D. Lyons. Tagged potential fields: An approach to specification of complex manipulator configurations. In *IEEE International Conference on Robotics and Automation*, volume 3, pages 1749–1754, 1986.
7. D.M. Lyons. A simple set of grasps for a dextrous hand. In *IEEE International Conference on Robotics and Automation*, pages 588–593, 1985.
8. J.R. Napier. The prehensile movement of the human hand. *J. Bone Joint Surgery*, 38B:902–913, 1956.
9. K. Salisbury and S. Craig. Articulated hands: Force control and kinematic issues. *International Journal of Robotics Research*, 1(1):4–17, 1982.
10. K.B. Shimoga and A.A. Goldenberg. Soft materials for robotic fingers. In *IEEE International Conference on Robotics and Automation*, volume 2, pages 1300–1305, 1992.

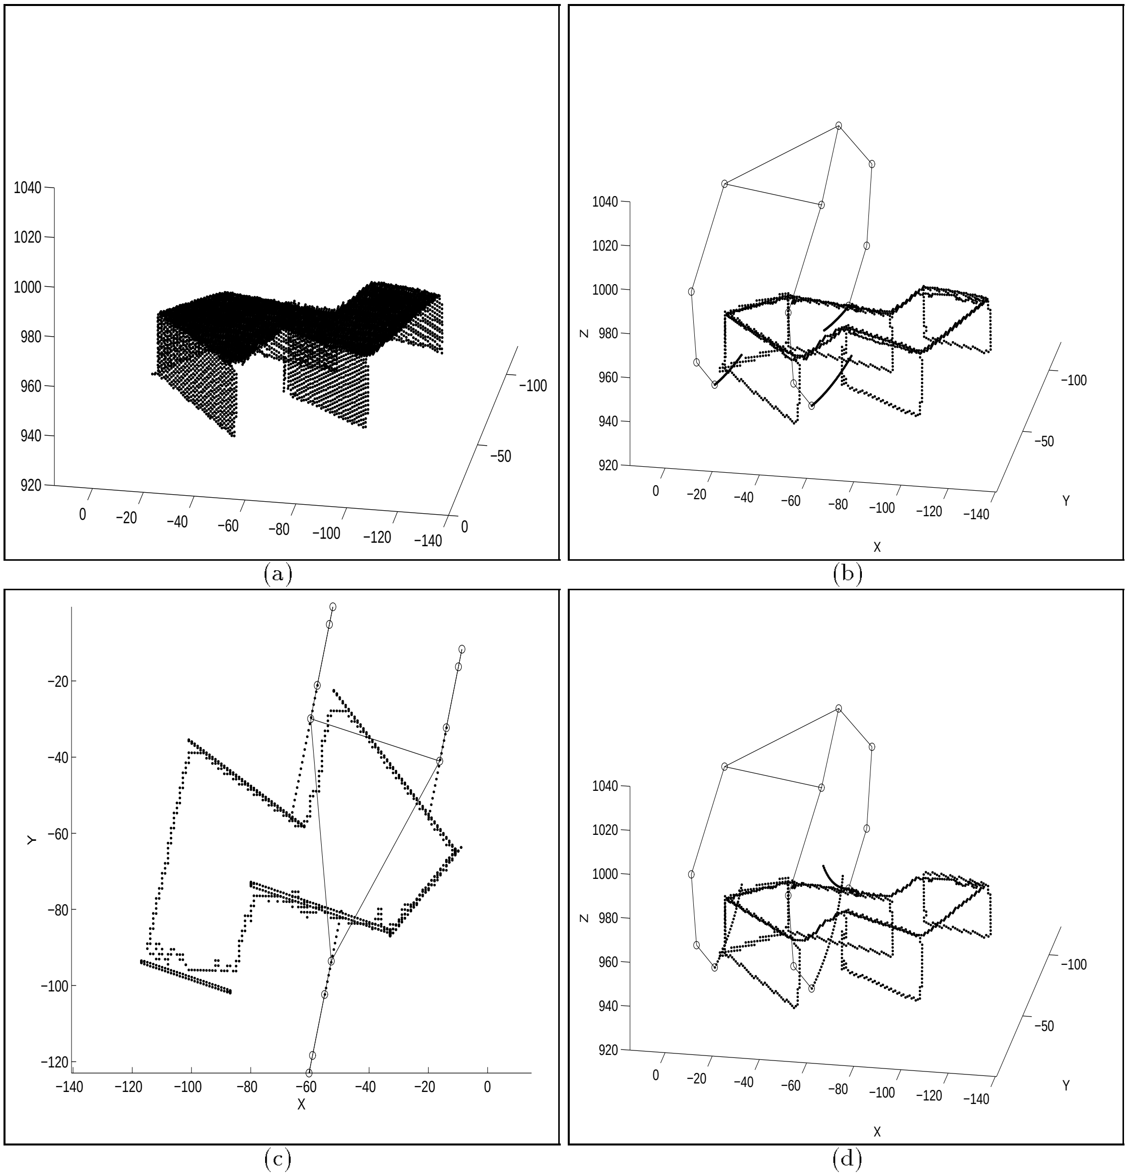

Fig. 4.: (a) range data, (b) $P1$-type grasp, (c) viewed from top, (d) the same preshape, with *different* finger trajectories, does intersect with the grasping features. The small segment at the end of each finger is the trajectory of the fingertip from the preshape position to the point of contact.

11. S.A. Stansfield. Robotic grasping of unknown objects: A knowledge-based approach. *Int. J. of Robotics Research*, 10(4), August 1991.
12. A.M. Wing, A. Turton, and C. Fraser. Grasp size and accuracy of approach in reaching. *Journal of Motor Behaviour*, 18(3), 1986.
13. D. Wren and R.B. Fisher. Identifying robot finger grasping points from range data. In *Sixth Intl. Conf. on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, pages 522–530, 1993.