Shape from Chebyshev Nets

Jan Koenderink[1] and Andrea van Doorn[2]

[1] Helmholtz Instituut, Universiteit Utrecht, PO box 80000, 3508 TA Utrecht, The Netherlands, j.j.koenderink@fys.ruu.nl
[2] Andrea J. van Doorn, Laboratory for Form Theory, Fac. of Industrial Design Eng., Technical University Delft, Jaffalaan 9, 2628 BX Delft, The Netherlands

Abstract. We consider a special type of wiremesh covering arbitrarily curved (but smooth) surfaces that conserves length in two distinct directions at every point of the surface. Such "Chebyshev nets" can be considered as deformations of planar Cartesian nets (chess boards) that conserve edge lengths but sacrifice orthogonality of the parameter curves. A unique Chebyshev net can be constructed when two intersecting parameter curves are arbitrarily specified at a point of the surface. Since any Chebyshev net can be applied to the plane, such nets induce mappings between any arbitrary pair of surfaces. Such mappings have many desirable properties (much freedom, yet conservation of length in two directions). Because Chebyshev nets conserve edge lengths they yield very strong constraints on the projection. As a result one may compute the shape of the surface from a single view if the assumption that one looks at the projection of a Chebyshev net holds true. The structure of the solution is a curious one and warrants attention. Human observers apparently are able to use such an inference witness the efficaciousness of fishnet stockings and bodysuits in optically revealing the shape of the body. We argue that Chebyshev nets are useful in a variety of common tasks.

1 Introduction

The already well established field of photogrammetry has recently made remarkable progress[1993], largely because of innovative methods developed by the computer vision community. Modern methods make it viable to dispense largely with extensive camera (pre–)calibration and yet to obtain projective or affine solutions from two or more views. Such solutions can then be post–calibrated on the basis of known or assumed metric properties of the scene. Examples of such properties are parallelity, orthogonality and equipartition of length. The ideal scene would contain a Cartesian 3D–lattice. Of course, *if* such a fiducial structure were available one could actually dispense with the first step and unravel a single perspective view. In this paper we consider a generalization of this latter possibility: Shape from a single view on the basis of prior information concerning the metrical structure. We will focus on smooth, general *surfaces* on which a network of fiducial curves has been drawn.

Various authors have remarked upon the observation that the projection of such networks often allows the human observer to obtain a vivid 3D impression of the surface[1981, 1983, 1986]. However, this clearly need not hold for arbitrary networks: As a counterexample one could pick any network in the projection (e.g., a Cartesian grid) and use the inverse projection to put it on any surface. In such cases the impression is always that of a *flat*, frontoparallel object[1986], *even if the actual surface is highly curved.* One has speculated that such so called "shape from contour" is enabled by nets of principal curvature directions[1981]. For instance, Stevens, 1981 has: ... *to conclude anything about the sign of Gaussian curvature, the physical curves must be lines of curvature.* On the face of it this seems unlikely, and in this paper we consider more general possibilities.

2 Curvilinear "Cartesian grids"

Consider the problem of how to generalize the notion of "Cartesian grid" (chess board) on a general, curved surface. Clearly one has to do some concessions in order to be able to apply the net to the surface. Two obvious possibilities are, 1^{stly}, to keep all angles at $\pi/2$, or, 2^{ndly}, to preserve the equality of edge lengths. In the 1^{st} case one obtains *conformal* nets. In general their edge lengths will vary from place to place. The nets of principal curvature directions are one possibility (since the principal directions are orthogonal), but infinitely many others exist. In the 2^{nd} case one obtains the so called "Chebyshev nets". In that case the angles vary from place to place, *i.e.*, the mazes become parallelograms. In this paper we consider this 2^{nd} possibility.

Chebyshev nets occur in real life (among more) as basket ball nets (part of a pseudosphere), nets used as hammocks, food containers, stretched over balloons, *etc.*, fishnet stockings and bodysuits. In figure 1 we show an example: A dancer's legs dressed in fishnet stockings. Notice the clarity with which the 3D shape is revealed, especially in the (complicated) knee region. Since the final category was presumably designed to bring out the body shape particularly clearly, one guesses that "shape from (Chebyshev) nets" will be possible. When we consider the edges of the net as freely rotatable about the vertices (the knots in a real wire net), such networks are evidently *deformable*. Indeed, the fishnet stockings probably started out as *planar* Cartesian nets (flat pieces of very tenuous cloth) which were then stretched over the body.

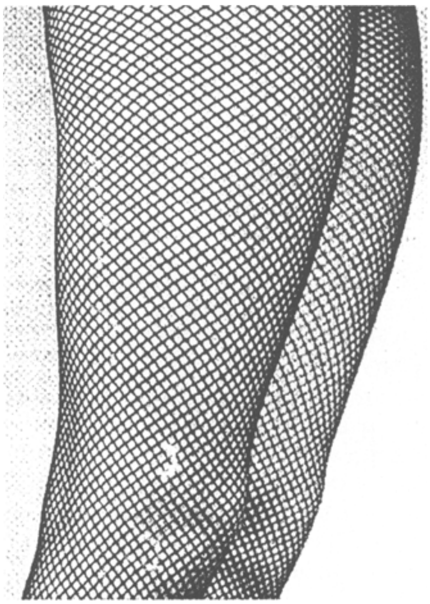

1. *Photograph of a dancer's legs clad with fishnet stockings. Notice how well the 3D shape is visually revealed from the projection. Notice especially the (geometrically) extremely complicated knee area.*

3 Chebyshev Nets

The main facts on Chebyshev nets are that 1^{stly}, they can be constructed with great freedom: One may specify two arbitrary (transversal) curves through any point of the surface and proceed to construct a unique Chebyshev net on them, maze by maze[1927]. Then, 2^{ndly}, each Chebyshev net can be applied to the plane (and hence to any other smooth surface). There are some limits to the validity of these statements though: They apply typically only to finite regions. Outside these regions the net "collapses" and needs to be "overstretched" which is forbidden by the constant edge length constraint[1882]. Indeed, the length of no diagonal can exceed double the edge length. In such cases one has to patch pieces together in other to "clothe" the surface. This introduces the notion of (tailor-like) "cutting and sewing". The classical reference is Chebyshev's lecture "On the cutting of our clothes". Chebyshev[1878] demonstrated how to construct a tight, sexy suit for the unit sphere from two pieces of "cloth". We illustrate how to cut the cloth for such a suit in figure 2. We are not aware of any such illustration in the literature, but presumably this mimicks Chebyshev's solution for the problem. Two identical pieces of cloth should be cut—minding the weft and warp directions—and sewn together along corresponding points. The seam will run along the equator of the unit sphere (see figure 3). Of course the net will be in tension when applied over the sphere and a "crooked seam" will result in uneven tension. (Our example is exact, the edges being geodesic arcs of length 6°. Presumably Chebyshev used infinitesimal edge length; the difference should be slight though.) Notice that this method yields a nice (piecewise) "Cartesian" coordinate system for the unit sphere: This suggests another application for Chebyshev nets.

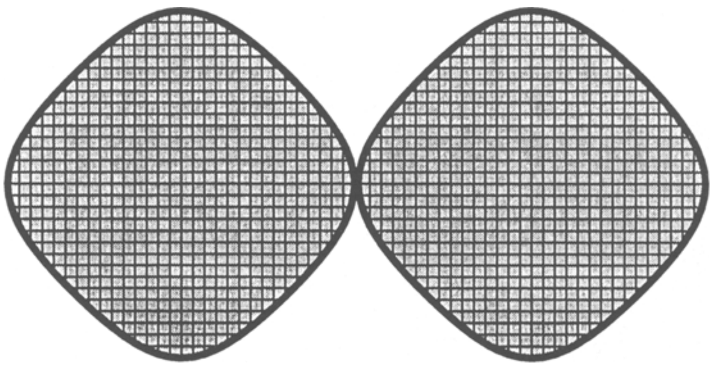

2. *How to cut the cloth in order to sew a tight suit for the unit sphere. The two identical pieces should be sewn together along corresponding points. In clothing the unit sphere the seam should run straight along the equator. The resulting fit should be perfect, though perhaps slightly uncomfortable since the cloth will be in tension.*

The 1^{st} fundamental form (metric) for a Chebyshev net with parameters u, v (thus u =constant and v =constant are the "wires" of the net) is simply

$$ds^2 = du^2 + 2\cos\zeta du\, dv + dv^2,$$

where ζ denotes the angle between the wires (thus we obtain the Pythagorean theorem for $\zeta = \pi/2$). One easily shows that the Gaussian curvature is

$$K(u,v) = \frac{-1}{\sin\zeta}\frac{\partial^2\zeta}{\partial u\,\partial v}.$$

Notice that the diagonals of the mazes are orthogonal: If we use $\alpha = u+v$, $\beta = u-v$ as new parameters the metric becomes

$$ds^2 = \cos^2\frac{\zeta}{2}d\alpha^2 + \sin^2\frac{\zeta}{2}d\beta^2.$$

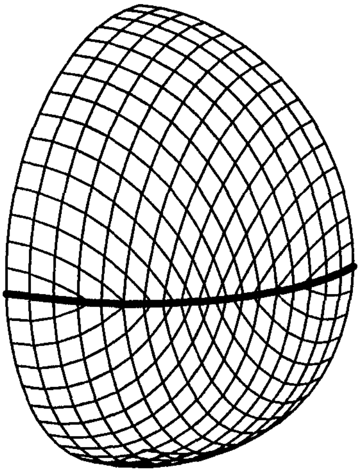

3. A quadrant of the spherical suit. The "seam" is the equator of the sphere, the sides are the primeval weft and warp threads..

In order to find examples of Chebyshev nets one may derive a system of 2^{nd} order differential equations, essentially[1882]

$$\frac{\frac{\partial^2 x}{\partial u \partial v}}{\frac{\partial y}{\partial u}\frac{\partial z}{\partial v} - \frac{\partial z}{\partial u}\frac{\partial y}{\partial v}} = \frac{\frac{\partial^2 y}{\partial u \partial v}}{\frac{\partial z}{\partial u}\frac{\partial x}{\partial v} - \frac{\partial x}{\partial u}\frac{\partial z}{\partial v}} = \frac{\frac{\partial^2 z}{\partial u \partial v}}{\frac{\partial x}{\partial u}\frac{\partial y}{\partial v} - \frac{\partial y}{\partial u}\frac{\partial x}{\partial v}},$$

and solve for $(x(u,v), y(u,v), z(u,v))$. One easily checks that the class of surfaces of translation are a particular set of solutions. Solutions are only simple to obtain in special coordinate systems. For instance, for the unit sphere one finds[1882]

$$\begin{aligned} x(u,v) &= \sin\mathrm{am}(u+v)\cos(u-v)k, \\ y(u,v) &= \sin\mathrm{am}(u+v)\sin(u-v)k, \\ z(u,v) &= \cos\mathrm{am}(u+v), \end{aligned}$$

with *am* the Jacobi amplitude and k the modulus of the elliptic functions.

In figure 4 we illustrate a Chebyshev covering of the plane and in figures 5 and 6 of the sphere. Of course neither of these is unique. The planar case is an interesting one: In this case both the weft and warp families of threads are parallel curves. According to Stevens' speculations[1981] such families should be interpreted as (locally) *cylindrical*. Of course the "shape from Chebyshev net" solution will be planar.

Stevens would probably argue (our speculation) that the Chebyshev net assumption is not a natural one for human observers to make.

Evidently, it is a different problem to find Chebyshev nets on a *given* surface whereby two initial parameter curves can be specified freely. This leads to a 2^{nd} order partial differential equation. We will not consider the problem here. The reader will find the necessary material in Bianchi[1927].

In practice one will often revert to finite methods because of several reasons. For instance, one might be interested in finite edge lengths to start with and consider the edges as rigid rods, freely rotatable about the vertices[1970]. (One may well speak of "finitesimal" nets, whereas the true Chebyshev nets are infinitesimal.) Or one might be interested in piecewise geodesic edges (a stretched wire *has* to be a geodesic). It turns out to be the case that (true) Chebyshev nets with geodesic wires are necessarily developable surfaces. In that case only one family (either the weft or the warp) can be geodesic[1882]. Since the developables form too restricted a class one has to consider

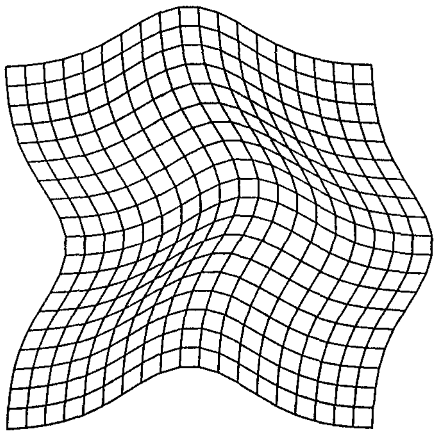

4. A planar Chebyshev net. The weft and warp threads form parallel families.

finitesimal nets with only piecewise geodesic wires (the edges). In such cases one most easily constructs the net maze by maze. In figure 5 we show a covering of the sphere obtained with such a piece–by–piece construction. Such purely geometric methods of course yield nets that will typically be mechanically unstable (statically that is). Examples are the planar nets of translation: Obviously only the subset of Cartesian nets is stable. One produces such stable nets by applying isotropic tension on the wires. In order to construct stable nets one has to consider the detailed static mechanical constraints. An example is the basket ball net: It assumes the form of a hyperbolically curved surface of revolution (constant Gaussian curvature), such that the wires are its asymptotic curves[1942].

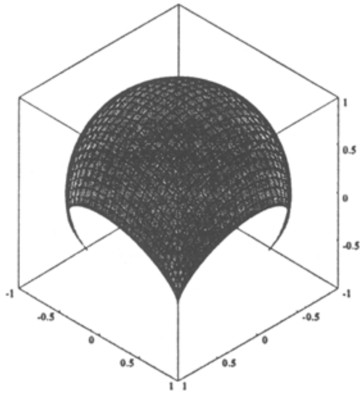

5. *A finitesimal spherical Chebyshev net. The edges are geodesic arcs of fixed length on the sphere. The net has been extended to its natural limits. This net was the basis for the computation of the cut needed to "dress" the sphere illustrated in figure 2.*

4 Shape from (Chebyshev) Nets

First consider an orthogonal projection of a finite Chebyshev net with straight edges on a plane. How does one compute the shape from the projection? One immediate observation is that the observed edge lengths λ_i (say) are related to the (unknown) true length Λ as $\lambda_i^2 + \Delta z^2 = \Lambda^2$, where Δz denotes the depth difference over the edge.

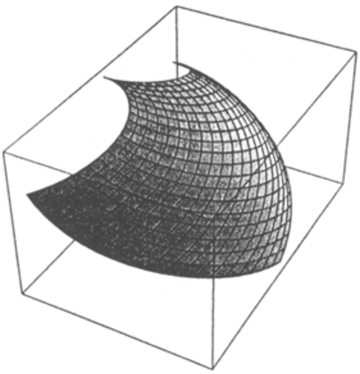

6. *A true (infinitesimal) spherical Chebyshev net: Of course only a few weft and warp threads at fine spacing could be drawn.*

Thus *if the true edge length were known* one could immediately regain the depth differences up to a sign. Since we don't assume the true edge length to be known this doesn't immediately apply. One has an additional constraint though: Surface consistency requires that the (algebraic) sum of the four depth differences over the edges of each maze vanishes. Thus we have (for any maze)

$$\sigma_1\sqrt{\Lambda^2 - \lambda_1^2} + \sigma_2\sqrt{\Lambda^2 - \lambda_2^2} + \sigma_3\sqrt{\Lambda^2 - \lambda_3^2} + \sigma_4\sqrt{\Lambda^2 - \lambda_4^2} = 0,$$

where $\sigma_j = \pm 1$. Clearly $\Lambda^2 \geq \max(\lambda_k)$. There is such an equation for each maze, the unknown Λ^2 being the same for all mazes. These equations generically determine Λ^2 and $\omega_i\{\sigma_1^i, \sigma_2^i, \sigma_3^i, \sigma_4^i\}$, with $\omega_i = \pm 1$, for all i (index i identifies the maze). In order to obtain the full solution we need to establish $\Omega\{\omega_1, \ldots, \omega_n\}$, where $\Omega = \pm 1$ and n denotes the number of mazes. Since Ω clearly represents an essential ambiguity (a "depth reversal") one may set one of the ω_i arbitrarily to ± 1.

In order to find the ω_i one compares adjoining mazes p and q (say): The depth interval over the common edge has to be the same for both mazes, this establishes $\omega_p\omega_q$. If one specifies $\omega_1 = 1$ (say), all the other ω_i generically follow from such a pairwise comparison. In case the data are essentially noise free and the case is generic such a scheme always lead to the solution. The pairwise comparison can for instance be implemented as a painting algorithm. One obtains the true edge length Λ and the true shape up to a depth displacement and a depth inversion. We have implemented the required synchronization of depth reversals for $2^N \times 2^N$ nets via the sequential synchronization of non-overlapping 2×2 subnets, each with elements of size $2^k \times 2^k$ $(k = 0, \ldots N - 1)$. One easily handles even large nets this way.

With perfect data such reconstructions work very well and one obtains both the true edge length and the true shape (up to a global depth reversal of course) from any single view of the net, the solution is essentially exact and depends only on the number of digits carried in the calculations. (See figure 7.)

If the data are noisy it is the pairwise comparison that tends to break down first: At a certain noise level the distinction between a surface attitude and its depth reversal becomes insignificant. The synchronization of local depth reversals is then no longer feasible. (In order to obtain a robust solution one might attempt to find the set of ω_i's that globally minimizes the total failure at the common edges. Such a solution may be obtained via a simulated annealing procedure. However, the expected gain is perhaps not worth the effort.) When the noise level is just slightly bad, one expects failure due to synchronized depth reversal of larger areas. When the noise level is really bad one expects that the individual mazes will be depth reversed at random.

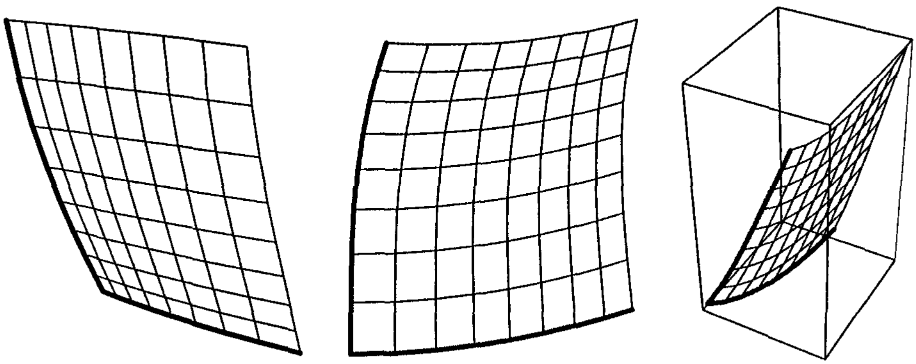

7. *Typical example of "depth from nets": On the left two projections of the same net and on the right a reconstruction. With perfect data the reconstructions from one view are essentially exact. This net covers a paraboloid, the initial curves (drawn in bold line) are general curves, not geodesics.*

In order to study the behavior of the solution in noisy conditions we considered a small (3x3 mazes) net. For such sizes one can still search for the global optimum of local depth reversals with an exhaustive search procedure. (Notice that the effort scales exponentially with the number of mazes.) We constructed a generic, finite

Chebyshev net (straight edges, vertices on the fiducial surface) on a paraboloid, using non-geodesic initial curves (or rather: edge progressions). The net was constructed exactly, maze by maze. We selected a projection in which the mazes had mutually quite different aspects. The projection was perturbed in the following manner: First we found the r.m.s. projected edge length. Next we added normally distributed random numbers to the Cartesian coordinates of the projected vertices, with zero mean and spread proportional to the r.m.s. projected edge length. The constant of proportionality will be referred to as the "noise level". This type of noise would be typical for data obtained by measurements of limited accuracy in the image (projection). We find the expected result, namely that:

— for low noise levels the reconstructions are essentially perfect. This regime extends to noise levels up to about 10^{-2};

— for high noise levels the synchronization of depth reversals breaks down completely. It is still possible to obtain reasonable estimates of the true edge length. This regime starts from noise levels roughly in the range 10^{-2}–10^{-1};

— in an intermediate region (here noise levels in a narrow range of about $3\,10^{-3}$ to $3\,10^{-2}$) one obtains mixed results. Sometimes the solution will be essentially good, merely somewhat deformed. Other times one witnesses depth reversal of local areas (larger than single mazes). In such cases a more intelligent algorithm (using prior information concerning smoothness for instance) might be expected to be able to "mend the damage". Sometimes one obtains really bad results, the depth reversals seem essentially random on the local (single maze) level.

Such behavior is indeed to be expected for a method that depends critically on the assessment of (often small) differences between Euclidean lengths (the true and the projected edge lengths). There is no way such a method could "deteriorate gracefully", rather, one is confronted with sudden breakdowns when critical noise levels are exceeded. Fortunately, the dangerous noise regime is pretty obvious from the image data themselves (histogram of projected edge lengths compared with tolerance).

We present an example in figures 8, 9 and 10. These illustrate three cases with noise level 1%. *Case A* illustrates an essentially correct reconstruction, the result is

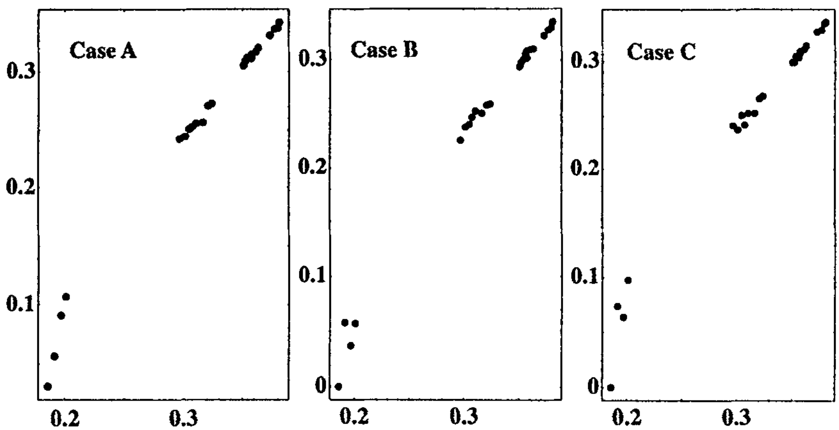

8. *Scatterplots of the recovered absolute differences over the (projected) edges against the veridical values. Notice that the results are quite acceptable in all three cases.*

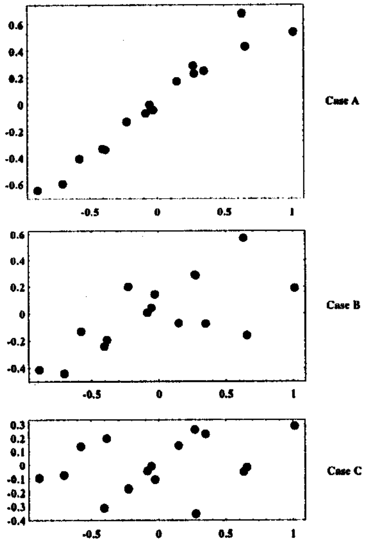

9. *Scatterplots of the recovered depths at the vertices (the average depth has arbitrarily set to zero) against the veridical values (average depths also set to zero). In case A the result is acceptable, in case B there is some correlation, but regional depth reversals make the result less than acceptable, though a smart algorithm might still be able to amend the problems. In case C we are left with a crumpled mess and the depths are essentially unrecoverable.*

merely somewhat distorted as compared with the fiducial shape. *Case B* illustrates the effect of local (but at a larger scale than the single mazes) depth reversals: The reconstruction is at least piece–wise correct. *Case C* illustrates a thoroughly crumpled reconstruction, this result is useless. In all three cases the (3D) edge length was estimated near to veridical (case A a deviation of 9.9%, case B a deviation of 11.5% and in case C a deviation of 11.0% from veridical). From figure 8 it is evident that the absolute depth differences over the edges are also well recovered in all cases (in case A we find a correlation of 99.2%, in case B of 99.1% and in case C of 98.7%). Because of the noise there are surface inconsistencies in all cases of the order of 10% of the true edge length (true edge length was 0.4 whereas the maximum surface consistency violation in case A was 0.029, in case B 0.056 and in case C 0.073). After the depth reversal synchronization procedure we were left with 1 sign violation (out of a total of 24) in case A and 9 sign violations in both cases B and C. Sometimes there were several solutions (with the same number of unresolved depth reversals), namely 2 in case A and a single solution in cases B and C. The unresolved depth reversals caused the major differences in the correlation of the recovered depths at the vertices with the veridical depths. These correlations were 97.5% in case A, 67.4% in case B and 31.1% in case C.

5 Conclusion

We have presented a method that lets us compute the shape of curved surfaces, covered with a wire mesh, from single views under the assumption that the net is the orthogonal, planar projection of a Chebyshev net. Such a reconstruction is robustly possible (given sufficiently precise data), explaining the informal but generally agreed upon fact that human observers can visually appraise the 3D shapes of such items

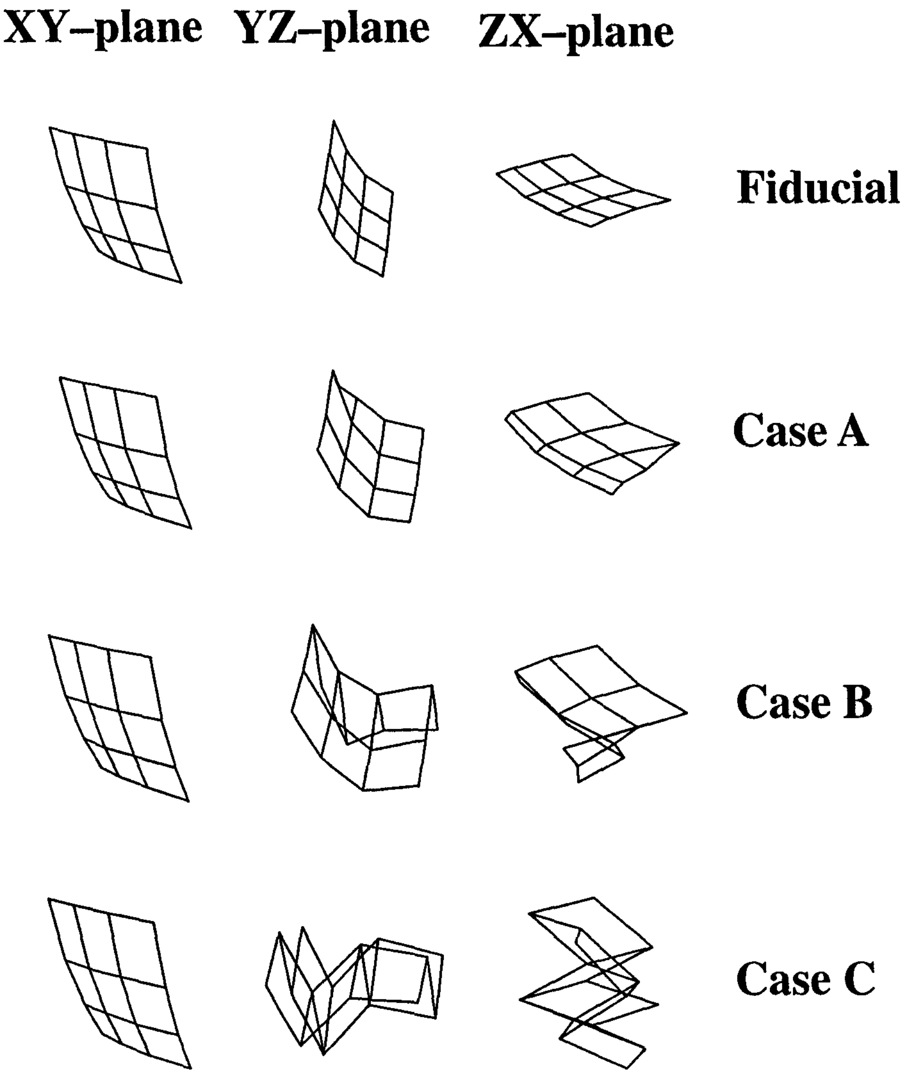

10. *Projections of the fiducial and the recovered nets on the coordinate planes. In this coordinate system the Z–axis is the depth dimension. Thus the projections on the XY–plane are very similar: Essentially the input image with 1% (projected) edgelength perturbations. Clearly case A represents an acceptable solution, case B nicely shows the result of regional depth reversals and case C is a crumpled mess due to essentially random local depth reversals.*

as basket ball nets, wire frames draped over balloons, used as packaging for foods, hammocks or fishnet stockings and bodysuits. The procedure radically differs from the established photogrammetic methods in that it is inherently Euclidean (it requires that the image plane carries a Euclidean metric). As such it is of interest as one extreme item in the toolbox of computer vision methods.

Although we have presented the case of orthogonal, planar projection, it is straightforward to generalize to the case of central projection provided the camera is fully calibrated. One simply changes to polar coordinates. Since the solution will be up to a scaling one may set the true edge length to unity and introduce the distance to the maze as the new unknown. Different from the present case this distance will vary from maze to maze, but the fact that Λ is the same for all mazes was not used in the solution anyway: Thus the solution essentially proceeds as described above and *conceptually* nothing new is gained, though such a solution may well prove to be of value in applications.

We envisage at least these three applications of Chebyshev nets in computer vision: 1^{stly}, they enlarge the set of fiducial objects on which one may draw to post–calibrate projective photogrammetic reconstructions, 2^{ndly}, the fact that "shape from (Chebyshev) nets" is viable and that humans seem able to perform this feat suggests their use in computer graphics. Renderings of Chebyshev nets may be used to provide the spectator with powerful depth cues. Finally, and 3^{rdly}, Chebyshev nets provide versatile parameterizations of surfaces for purposes of object representation. (For instance, we have presented a rather attractive parameterization of the hemisphere.) They are attractive because they are rather immediate generalizations of the planar Cartesian meshes (chess boards), yet they allow for much freedom, *e.g.*, can be naturally rotated about a point, deformed in various ways and applied to arbitrary other surfaces as—for instance—the plane. Thus any surface appears as a deformation of *e.g.* the plane such that lengths in two directions are conserved. This suggests many possible applications in CAD–CAM and object representation.

References

[1927] Bianchi, L.: Lezioni di geometria differenziale. 3^{rd} ed., Vol. I, Part I, Nicola Zarichelli, Bologna (1927) 153–162

[1993] Faugeras, O.: Three–dimensional computer vision. The MIT Press, Cambridge, Mass. (1993)

[1986] Koenderink, J. J.: Optic Flow. Vision Research **26** (1986) 161–180

[1970] Sauer, R.: Differenzengeometrie. Springer, Berlin (1970)

[1981] Stevens, K. A.: The visual interpretation of surface contours. Artificial Intelligence **17** (1981) 47–73

[1983] Stevens, K. A.: The line of curvature constraint and the interpretation of 3–D shape from parallel surface contours. 8th Int.Joint Conf. on Artificial Intelligence (1983) 1057–1062

[1986] Stevens, K. A.: Inferring shape from contours across surfaces. In: From pixels to predicates. Ed. A. P. Pentland, Ablex Publ. Corp., Norwood, NJ. (1986)

[1942] Thomas, H.: Zur Frage des Gleichgewichts von Tschebyscheff–Netzen aus verknoteten und gespannten Fäden. Math.Z. **47** (1942) 66–77

[1878] Tschebyscheff, P. L.: Sur la coupe des vêtements. Association francaise pour l'avancement des sciences. Congrès de Paris (1878) 154

[1882] Voss, A.: Über ein neues Prinzip der Abbildung krummer Oberflächen. Math.Ann. **XIX** (1882) 1–25