The technique of using three consecutive shears to rotate a 2-D image—illustrated by the following diagram—has been known for a while.

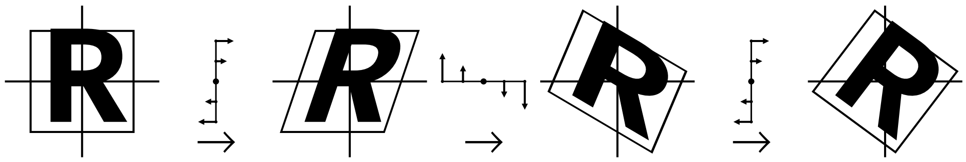

For 3-D images, there is an obvious *nine*-shear solution, obtained by treating the 3-D rotation as the product of three 2-D rotations. Our algorithm, which uses only *three* shears, is a sizable improvement over that bound, and is actually the best that can be attained—just two shears won't do. The algorithm, illustrated by the following figure,

Three-dimensional rotation (Euler angles: $\alpha = -10°$, $\beta = 30°$, $\gamma = -15°$) achieved by three shears.

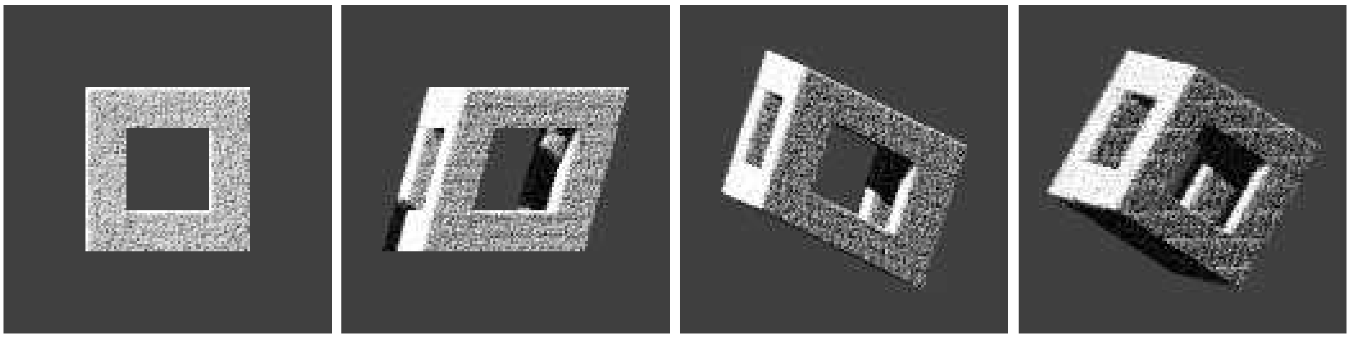

proceeds much as in the two-dimensional case. Namely, a rotation $R(\alpha, \beta, \gamma)$, where α, β, γ are the rotation's Euler angles, is decomposed as the product of three shears $R = S_3S_2S_1$. Each of the three shears S_i is given by matrix of the form $S = \begin{bmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{bmatrix}$ (or a transpose of this form) and is thus characterized by three real parameters (rather than just one as in the two-dimensional case), which are in turn functions of the Euler angles.

This result may be stated more abstractly as a linear algebra identity. In $n \leq 3$ dimensions, any orthogonal matrix O can be written in the form $O = ULU'$, where U and U' are *upper triangular* matrices and L is a *lower triangular* matrix—all three matrices having all 1's along the main diagonals. The question of whether this identity generalizes to $n > 3$ is still open.