

The MATLAB code below demonstrates how we can use bilinear interpolation to transform an input image  $A$  according to a given  $3 \times 3$  projective transformation matrix  $H$ . Every pixel  $(x, y)$  in  $A$  is transformed in homogeneous coordinates as  $[u \ v \ w]^T = H[x \ y \ 1]^T$ , so that the corresponding pixel in the output image is given by  $(x', y') = (u/w, v/w)$ .

```

1 function B = applyhomography(A,H)
2 % cast the input image to double precision floats
3 A = double(A);
4 % determine number of rows, columns and colour channels of A
5 m = size(A,1);
6 n = size(A,2);
7 c = size(A,3);
8 % determine size of output image by forward-transforming the four corners of A
9 p1 = H*[1; 1; 1]; p1 = p1/p1(3);
10 p2 = H*[n; 1; 1]; p2 = p2/p2(3);
11 p3 = H*[1; m; 1]; p3 = p3/p3(3);
12 p4 = H*[n; m; 1]; p4 = p4/p4(3);
13 minx = floor(min([p1(1) p2(1) p3(1) p4(1)]));
14 maxx = ceil(max([p1(1) p2(1) p3(1) p4(1)]));
15 miny = floor(min([p1(2) p2(2) p3(2) p4(2)]));
16 maxy = ceil(max([p1(2) p2(2) p3(2) p4(2)]));
17 nn = maxx - minx + 1;
18 mm = maxy - miny + 1;
19 % initialize the output with white pixels
20 B = zeros(mm,nn,c) + 255;
21 % pre-compute the inverse of H (we'll be applying that to the pixels in B)
22 Hi = inv(H);
23 % loop through B's pixels
24 for x = 1:nn
25     for y = 1:mm
26         % compensate for the shift in B's origin, and homogenize
27         p = [x + minx - 1; y + miny - 1; 1];
28         % apply the inverse of H
29         pp = Hi*p;
30         % de-homogenize
31         xp = pp(1)/pp(3);
32         yp = pp(2)/pp(3);
33         % perform bilinear interpolation
34         xpf = floor(xp); xpc = xpf + 1;
35         ypf = floor(yp); ypc = ypf + 1;
36         if (xpf > 0) && (xpc <= n) && (ypf > 0) && (ypc <= m)
37             B(y,x,:) = (xpc - xp)*(ypc - yp)*A(ypf,xpf,:) ...
38                 + (xpc - xp)*(yp - ypf)*A(ypc,xpf,:) ...
39                 + (xp - xpf)*(ypc - yp)*A(ypf,xpc,:) ...
40                 + (xp - xpf)*(yp - ypf)*A(ypc,xpc,:);
41         end
42     end
43 end
44 % cast the output image back to unsigned 8-bit integers
45 B = uint8(B);
46 end

```

This function follows the  $(x, y)$  convention for pixel coordinates, which differs from the  $(row, column)$  convention. You will note that  $y$  corresponds to row index and  $x$  to column index. The matrix  $H$  must be set up accordingly.

The size of the output is determined automatically, and the output will contain the entire transformed image on a white background. This means that the origin of the output image may no longer coincide with the top-left pixel. In fact, after executing this function, the true origin  $(1, 1)$  will be located at point  $(2 - minx, 2 - miny)$  in the output image (see if you can figure out why!).