Computer Vision Technology

Subpixel Precision
resolution and accuracy

Scope

Vision systems are extremely accurate measurement devices. They exhibit a far greater resolution than the frame store and picture format suggests. By using different integration methods, vision systems can utilize statistical methods and perform very accurate measurements.

Integration techniques were looked upon with scepticism when introduced in the beginning of 1980, but today they are widely accepted and proven in actual production over millions of service hours and billions of inspected parts.

The name for some of these techniques is subpixel precision.

Resolution / Accuracy

Resolution in a vision system is defined as the smallest detectable change in size or intensity of an object.

The resolution gives the maximum accuracy achievable by a given system.

In all vision system applications the accuracy is determined by the properties of the object, the illumination, the property of the atmosphere, the lens system and the environment.

The accuracy is depending on how well these factors can be controlled, compensated and calibrated under real production use.

Vision system accuracy is defined according to the rules of standard deviation, similar to all other measurement devices.

Resolution

The resolution of a vision system is primarily given by the number of pixels in the row and column structure. Typical format used in industry is 512 X 512. This equates to a quarter of a million measurement points.

If all these measurement points could be used the resolution would be enormous in fact 1 / 256.000.

A simple example illustrates this:

If the vision system should measure the area of a part placed under the camera, the system could count the amount of pixels belonging to the part, and the number of pixels belonging to the background. If the part takes up the total picture minus 1 pixel, the area can be determined with the above resolution

Further to this, each pixel describes the gray level in a scale going from typically 0 for black to 255 for maximum illumination. The grey level can be used to further improve the resolution.

Line integration

One commonly used method is line integration. An object is placed in the picture and the vision program detects the coordinates of the edges by using different strategies depending on the nature of the surface of the object.

The result is then a number of coordinates for say the right and the left side of the object. By performing a simple average on these coordinates, the resolution is improved by a factor determined by the square root of the number of measurement points.

By integrating over 100 pixels the resolution is improved by a factor 10.

Plane integration

Plane integration is used where the area of an object is to be measured. The main challenge of the vision program is to determine the number of pixels belonging to the background and the number belonging to the object.

Grey level integration

The normal edge detection techniques determines to one pixel accuracy where the boundary of the object lies on each point on the edge of the object.

When an object is projected onto the CCD camera chip, the edge of the object does not lie on the boundary of the CCD element, but will in most cases lie somewhere in the CCD cell. This means that the CCD pixel holding part of the edge will have a grey level somewhere between the background and the object.

This intensity measurement in one pixel can be used to determine the location of the edge of the object within the one pixel. The resolution in the ideal case will be the range of grey levels or 1 : 256. If the object is covering the pixel completely the intensity will be 0. If the edge is 21 % into the pixel the grey level will be 201. etc.

Grey level integration is difficult to use in general systems because of focussing and the CCD chip row and column interference problems, but it has been shown that in very controlled environments, when the object always has the same intensity, an improvement in resolution of a factor 15 can be achieved.

Time integration

If the resolution in one picture is not sufficient, it is possible to integrate over time.

If the object going to be measured, takes up a very small part of the picture, it may not be possible to make integration over many pixels. When only a small number of pixels are used, the program will often finish the job in a very short time. This opens the possibility of integrating over several pictures. The measurements of each consecutive picture is averaged, and thereby improving the resolution.

Plane and grey level integration

The different methods of integration can be combined to achieve resolutions far beyond the accuracy achievable in a final system.

It is shown in theory and in practical applications that the improvements in resolution can be relied upon. The theory also holds in the real world of steel works.

Standard deviation

All measurement systems except a very few follows the rules of standard deviation.

When using statistical methods the single measurements will group together in such a way as to give a deviation from the mean value.

This will happen to a certain extend with a dependable frequency.

It can be shown that the following applies:

68.3%	probability for SD =	1
95.4%	probability for SD =	2
99.7%	probability for SD =	3
99.994%	probability for SD =	4

This means that 68.3 % of all measurements will fall within the standard deviation, and that 0.006 %, or 1 in 16.666 measurements will deviate 4 times the standard deviation.

It is a good rule of thumb, that the measurement device used for a job, should be ten times better than the tolerances of the object. If this is followed the standard deviation and the probabilities above will not be a problem.

In the table below a vision system is measuring the width of a steel slab. This system was commissioned at the Danish Steel Work in 1983, and proved the theory in a practical application for the first time.

The picture field was 5.000 mm wide and the system measured slabs up to 4.000 mm.

68.3 %	probability for SD =	1	0.8 mm
95.4 %	probability for SD =	2	1.6 mm
99.7 %	probability for SD =	3	2.4 mm
99.994 %	probability for SD =	4	3.2 mm

It was important that the slabs were not too short. The probability of making too short slabs was half the above, as these are ± values.

Accuracy

In order to exploit the resolution and build accurate measurement systems a lot of parameters must be understood, controlled and compensated.

The accuracy achieved in different applications depends on the object, the illumination, the environment and the optical parts of the camera system. Also electrical interference may play a role.

The primary goal when illuminating an object is to make the edges stand crisp and clear in the picture.

The edge of the object is projected in through the lens to the CCD chip.

The lens is not a very accurate part. Most lenses are made to produce sharp images, but the geometrical liniarity may be of secondary importance. This means that a standard 12.5 mm lens may show an unlinearity of 2 - 3 %.

Fortunately this is a constant for the specific lens, and as long as the lens can be adjusted in focus and aperture and sealed, it is possible to map the unlinearities and use a correction table when calculating the coordinates of the features in the picture.

CCD cameras cannot produce a picture of a bright and a dark boundary without creating some errors.

The white part of the picture will eat its way into the dark part as the intensity is increased. This effect is called blooming and cannot be avoided.

The way to control this is by either measuring the light levels in the pixels at the boundary, or to control the illumination of the object, so it will be kept in close tolerances.

A phenomena can appear when using line or plane integration. If the CCD camera produce a very sharp picture it is possible for an object boundary to align exactly with the row or column in the CCD chip. If this happen the resolution may suffer. If the edge is moved from left to right through the picture all the pixels at the edge may shift at the same time.

This is a very rare phenomena as there are always unlinearities and other factors creating noise, helping the pixels shift individually.

One method used commonly when there is a danger for this happening, is to tilt the edges slightly in relation the CCD row and grid direction.

The distance from the camera to the object is of great importance. If the distance from the object to the camera varies with 1 % the size of the object will also vary with the same 1 %.

Several schemes exists to compensate for this. A widely used method is to use two cameras mounted under an angle of 90°. Other methods exists and depends upon the application.

Conclusion

It has been shown in many steel works applications that the sub pixel precision can be relied upon in actual production use.

Sub pixel methods are used extensively in all sorts of applications throughout the industry, and has improved the accuracy of vision systems far beyond early expectations.

Jørgen Læssøe

JLI vision a/s
DENMARK

20. Februar 1993