

Real-time vision-based detection of wire gapping on a rotating transformer core

C.J. Taylor, E. Evans, J.M. Dolan, P.K. Khosla
The Robotics Institute, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213

ABSTRACT


High-quality transformer winding requires precise measurement and control of the gapping between adjacent wires. We take a vision-based approach to the measurement subtask of determining the gaps between copper wires wound onto an oval transformer. The oval core shape, which can have an eccentricity as high as 2-to-1, leads to significant variations in surface normal and viewing distance. We use special lighting, a secondary mandrel shape sensor, and the specular reflection off the wires to give us an accurate model of the experimental geometry. We further exploit the vertical symmetry of the viewed region to condense our 2-D image to a simpler 1-D signal containing reflectance peaks. After utilizing pattern recognition and some additional safety features to separate the wire peaks from background noise, we perform a least squares curve fit of the peaks to determine the subpixel maxima. The final algorithm is computationally fast and yields the desired wire gap in an absolute metric.

Keywords: transformer winding, gap control, gap measurement, real-time vision, peak detection.

2.0 THE NATURE OF THE PROBLEM

2.1 The manufacturing machine and winding process

The overall goal in this and related work is to more fully automate the manufacturing process of dry-type transformers. The process begins with an iron core. A low-voltage winding composed of aluminum or copper foil is added. The core is then sent to the winding machine shown schematically in Figure 2.1 for the high-voltage winding.



The machine can rotate the transformer core at a variable rate in either direction on a central axis. The filament winding support is shown on the right. Fiberglass insulating strands are sent through a resin bath and then guided by the filament winding support onto the transformer core, where they provide additional insulation and structural support between conductor layers. The secondary mandrel shape sensor, which is mounted on the filament support, provides an accurate measurement of the shape of the mandrel.

The conductor wagon is shown on the left side of the schematic. Spools of wire rest on the wagon chassis. While winding a conductor layer, wire is fed from these spools, sent through a number of passive tension control devices and then guided onto the transformer core by a boom arm that extends from the wagon. This arm can be rotated manually around its axis mount on the wagon, and the length of the arm can be changed. The location of the end point of this arm, a fixed parameter during winding, is important in the vision setup.

The high-voltage conductor windings are added layer by layer, with a resin/fiberglass layer separating each conductor layer. The goal of the current work is to increase the quality of the conductor layer by obtaining precise measurements of the gapping between adjacent wires. Such information should aid in keeping this gap as small and as uniform as possible, thereby diminishing the required length of the core for a specified field strength. A desired future use of this information is to wind wires close to one another, but not touching. This would allow uninsulated wire to be used, providing a substantial saving in manufacturing cost.

2.2 Motivation for using vision

The high eccentricity of the initial core shape creates difficulties for any sensor attempting to detect the wire gap from a fixed angle. The sensor will have to cope with significant variations in surface normal as well as variations in the distance between the sensor mount and the wires as the mandrel rotates. These variations occur at a high frequency, as the mandrel may spin as fast as 100 RPM during the winding process. The sensor must be flexible enough to work on a variety of wire types in a variety of configurations. The wires themselves can be extremely small (the largest round wire is 3mm in diameter), so localizing them to ± 1 mm is not acceptable. Goal values are on the order of ± 0.1 mm.

A touch-based sensor system is the least suited for the task at hand. The aforementioned eccentricity of the core would require servoing of both the position (to compensate for variations in distance) and the orientation of the sensor (to compensate for variations in surface normal). The sensor tip of such an actuator would need sufficient compliance to maintain contact with the mandrel while staying fixed in the direction along the mandrel axis. The sensor device would also have to be able to adjust its measuring antennae to sense a wide variety of wire sizes. Worst of all, the surface of the mandrel around the fresh wire is covered with sticky wet resin. A touch-based sensor system is impractical for these reasons.

Laser rangefinding systems are commonly used to scan for the distance and placement of objects. The scan rate required for our application would have to be very fast: by the time the system had completed a linear scan, the mandrel could not have moved very much. Aside from this technical problem, the idea of a laser based system was discarded due to the expense involved. A production line laser system with its accompanying hardware was estimated to cost 5 times as much as a vision-based system. So long as the vision system could provide the accuracy required, it was the economically preferred system. Still, our algorithm can be implemented with a laser-based system with minimal modifications.

A vision system appears to be the most robust solution. The two main advantages with the vision system are that it captures information (the image) fast enough for our high speed problem and it involves no physical contact with the mandrel. Its disadvantages are that the vision calibration process is tedious, and the current implementation requires controlled lighting. Nonetheless, we believe that a vision based approach is the best industrial solution.

2.3 Notation and axes orientation

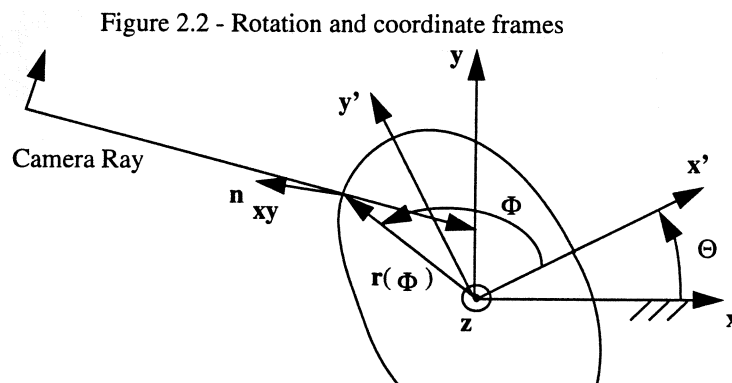


Figure 2.2 illustrates the notation that will be used in this paper. A cross-section of the mandrel from the same perspective as that of Figure 2.1 is shown with the filament support on the right and the conductor wagon on the left. The x-axis is parallel to the floor, the y-axis is perpendicular to the floor and the z-axis (coming out of the page) lies along the mandrel's axis of rotation. The wire gap is measured along the direction of the z-axis. The origin lies on the far side of the mandrel, such that the z-coordinate of any position on the mandrel is positive.

Points on the mandrel are described by two angles (Θ and Φ), a radius $r(\Phi)$, and z , the distance along the mandrel. Θ gives the rotation of the mandrel, i.e., the counterclockwise angle between the x -axis and an axis "inscribed" on the mandrel, x' . This choice of an inscribed axis is arbitrary and determined during the shape sensing stage of the winding process. Φ describes the angle between any surface point on the mandrel and x' , while $r(\Phi)$ is the distance between this point and the z -axis. \mathbf{n}_{xy} is the normal at the mandrel observation point projected onto the x - y plane. As Θ changes, this normal changes. This fact greatly complicates the vision problem.

2.4 Vision system hardware

Our current vision system employs a 35mm lens with a focal length of 135mm. This lens is connected to a CCD video camera through a C-mount to 35mm lens adapter. The camera output is a standard TV signal. This signal is digitized by a frame grabber which resides on the AT bus of a 486 I.B.M. clone. The video signal can be captured at a sampling rate of 640x480 pixels per 1/30sec by the frame grabber, though this rate is typically lower in practice.

3.0 PHYSICS OF THE VISION SETUP

3.1 Reflection and the use of specular reflection

The standard model states that reflection is composed of three dominant components: the diffuse lobe, the specular lobe and the specular spike. The diffuse lobe represents the internal scattering mechanism, sometimes called Lambertian reflection. The light that is re-emitted radiates uniformly over all angles. The specular spike is a result of mirror-like reflection: $\Theta_i = \Theta_r$ about the normal of the surface. The specular lobe, like the specular spike, is a result of reflection off a single surface. This lobe is a by-product of microscopic roughness. On the microscopic scale, the surface normal of the object is not uniformly in the same direction as the macroscopic normal. These microscopic mirrors yield specular reflection within a small solid angle around the specular spike.⁹

Some researchers¹⁰ focus on avoiding specularities, preferring a Lambertian (diffuse) reflector. Others^{6,13} focus on recovering shape information via specular analysis. We are forced to use the specularities in our analysis. Our objects of interest are made out of metal, and reflection from a metal is dominated by the interaction between the light and free electrons. These electrons scatter light effectively, and so the light penetrates only the surface skin of the conductor.⁵ This results in a weak diffuse component. Our conductors do have an insulating layer; however, this layer appears to act merely as a filter, thereby reinforcing the copper's own natural tendency to absorb blue and green light. The wire surface is also very smooth, so that the specular lobe is much less intense than the specular spike.

The use of specular reflection has a number of advantages. The foremost of these is that there is a precise geometric relationship between the position of the light source, the position of the camera, and the surface normal from which the reflection originates. We will show that, to a certain level of approximation, comparing the specular reflection from two adjacent wires yields the gapping between these wires. The high intensity of reflected specular light is also an advantage because our narrow aperture, which is needed to achieve the required depth of field, greatly reduces the amplitude of the incoming light. Because of the nature of our light sources (fluorescent tubes in an x - y plane with constant z), the viewed specular region is composed of a finite width line of highlight points in a direction near-perpendicular to \mathbf{n}_{xy} and z . As will be discussed in section 4, this allows us to condense a two-dimensional picture to two 1-dimensional signals, thereby greatly speeding up the algorithm.


The described specular analysis problem is similar to an earlier study:¹³ determining the radius, orientation and location of a cylinder via highlight information. Our problem has two simplifying aspects: (1) we are dealing with a cylindrical wire of fixed radius and (2) we are only interested in knowing the location of the wire relative to an identical twin wire that is only slightly offset from the first. Constraints on the radius and the location of the light source allow us to validate the approximation that the distance between the wires is proportional to the distance between their specular highlights.

It is easiest to first consider a cylindrical mandrel. If the camera incident ray is aimed directly at the cylindrical mandrel axis, the mandrel normal in the x - y plane, \mathbf{n}_{xy} , is pointing directly back along the incident ray. With this configuration, any point light source whose position can be described by $\mathbf{P}_L = k_1 \mathbf{n}_{xy} + k_2 \mathbf{z}$ with $k_1 > (2r_w + |r(\Phi)|)$, where k_1 and k_2 are scalars and r_w is the wire radius, will produce a specular spikes from wires wrapped around the mandrel (see figure 3.1A).

If we determine the z-component of the specularities (P_{z_i} and $P_{z_{i+1}}$) and use them to locate the wire positions, we expect a margin of error in our measurement, since the true positions are given by W_{z_i} and $W_{z_{i+1}}$.

From the diagram, it looks as if this approach fails completely. Even for a light source much farther away, the discrepancy between our measured wire location and the actual wire location will be significant with respect to our desired accuracy. Fortunately, we are interested not in the absolute location of the wires, but rather in their location relative to each other.

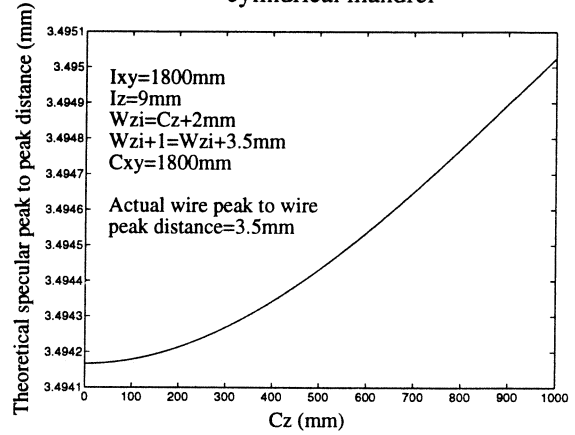
Figure 3.1A - Specular reflection from two adjacent wires (exaggerated differences)



$$\begin{aligned} \text{gap measured from specular spike} &= |P_{z_{i+1}} - P_{z_i}| \\ \text{actual gap} &= |W_{z_{i+1}} - W_{z_i}| \end{aligned}$$

P_{z_i} = z-component of the point on wire i where specular spike originates
 W_{z_i} = z-component of the location of wire i
 V_i = unit vector in the direction of the viewer (camera)

Figure 3.1B - Specular spike reflection off a cylindrical mandrel



n_i = unit surface normal at P_{z_i}
 L_i = unit vector from P_{z_i} in direction of source
 r_w = wire radius

If we assume that the light source and the observer are in the n_{xy} -z plane, ensuring that P_i and P_{i+1} are in the n_{xy} -z plane, which is nearly true anyway, then we can describe n_i in terms of the other wire parameters.

$$n_i = \frac{|W_{z_i} - P_{z_i}|}{r_w} z + \frac{\sqrt{r_w^2 - (W_{z_i} - P_{z_i})^2}}{r_w} n_{xy} \quad (1)$$

For our numerical simulation, we assume values for the position of the illumination source, $(I_{xy}, I_z, I_{xy\perp}=0)$, and the camera, $(C_{xy}, C_z, C_{xy\perp}=0)$. We also assume an offset between C_z and the z-location of the first wire of interest, W_{z_i} . We also assume a fixed wire spacing, $W_{z_{i+1}} - W_{z_i} = \text{constant}$, and a fixed wire radius, r . We can then calculate P_{z_i} and $P_{z_{i+1}}$ by using equation (1) and

$$\frac{\mathbf{I}}{\|\mathbf{I}\|} \cdot \mathbf{n} = \frac{\mathbf{V}}{\|\mathbf{V}\|} \cdot \mathbf{n} \quad (2)$$

though we need the additional physical constraint:


$$0 \leq W_{z_i} - P_{z_i} \leq r \quad (3)$$

This is algebraically tedious to simplify, and thus we use numerical techniques to search for P_{z_i} . Figure 3.1B shows one possible configuration with the light and camera 1.8 meters away (in the direction of n_{xy}). The theoretical correction factor for this simulation is on the order of a few thousandths of a millimeter. Since this is smaller than the observed noise arising from other sources (picture flicker, white speckle, camera calibration), we neglect this correction factor.

3.2 Flat wire

For flat wire, the normal of the wire rapidly spans 90° in the n_{xy} -z plane at both edges, $+90^\circ$ for one edge and -90° for the other. It is from these edges that we observe specular reflection (figure 3.2).

Figure 3.2- Reflection from only one Edge of Flat Wire



We utilize additional lighting (a duplicate light bar on the other side of the conductor wagon) to produce both highlight lines. This suggests that we would have two possible measurements for the gap. In flat wire winding, we have observed two peaks for each wire. However, for tighter winding, the two adjacent peaks would blur together into one peak. Thus there is a certain minimum gap that can be measured. There are other concerns with flat wire. The reflectance intensity from the edge is not very strong because the surface normal rapidly passes into and out of the requisite angle for specular reflection. There are also special wire-segmentation approaches required which we will not discuss in this paper.

3.3 Lighting

In order to use specular highlights to locate the mandrel, we require adequate lighting for all possible orientations of the mandrel. For a cylindrical mandrel, it is a simple task to provide this lighting, since all orientations are equivalent. The location and the \mathbf{n}_{xy} of the viewed wires do not change during mandrel rotation.

An elliptical mandrel greatly complicates lighting requirements. In our off-line simulation, we first determine the point on the mandrel closest to the camera ray. We calculate the component of the surface normal, \mathbf{n}_{xy} , at this point. We then reflect the camera ray about this normal. This reflected ray corresponds to a light incident ray required for a specular spike. Our simulations have shown that aiming the camera at a point on or below the core axis results in problematic reflection directions, requiring light from the ground for all but the most unusual camera mounts. Aiming the camera slightly higher up on the y-axis produces “better” reflection directions, i.e., it is easier to originate light from these directions. For the elliptical mandrel, all of the configurations require a range of incident light angles to provide constant specular spike reflection (see figure 3.3).

In our setup, we utilize two end-to-end fluorescent tubes (four for flat wire) that provide light for a large number of the required ray directions. We calculate an on-line simulation of the reflectance map for all theta by utilizing our knowledge of the provided lighting, the camera position and the mandrel shape (from the shape sensor). We can then avoid analyzing inadequately illuminated frames (resulting from a problematic reflection direction as shown in figure 3.3) by waiting for a valid theta value.

Figure 3.3 - Specular reflection directions off an elliptical mandrel





Figure 3.4 - Wire tangent points for an elliptical mandrel



3.4 Depth of field

Camera depth of field is very important for our system. As an elliptical mandrel core rotates, the effective distance from the surface to the camera optics is changing. For the highest quality picture, we want this entire range to be “in focus”, necessitating a large depth of field. We also want to eliminate any manual adjustment of the camera system between mandrel layers or even for different mandrels, since such an adjustment would spoil our camera calibration.

With our 135mm Nikon lens, we found that an aperture setting of f-stop=22 provides us with sufficient depth of field for our experiments. Of course, there are drawbacks to our narrow aperture window. As we decrease the aperture, we decrease the intensity of the light falling on the sensor plane. It has been necessary to greatly increase the camera gain, to turn on the camera’s gamma compensation, and to alter the framegrabber’s A/D settings. As a result, our pictures exhibit a fair amount of white speckle.

3.5 Ability to view freshly wound wire flush against mandrel

In order to get high resolution wire-gap measurements, we are using a lens with a high degree of magnification. This produces the pixel resolution we need, but it also places constraints on our field of view (FOV), i.e., the area we can observe with our camera. For true closed loop control of the wire gap, we would like to be observing the gap between the freshly wound wire and the adjacent, 2π -out-of-phase wire at the “wire contact point”. The wire contact point is the Φ location on the mandrel where the fresh wire is just beginning to come to rest on the mandrel surface. If we measure the gap at the wire contact point, our wire gap feedback signal will be a measure of the “instantaneous gap” - there will be minimum delay (frame delay) between the current gap and that which we are measuring.

Observing this wire contact point is trivial for a perfectly cylindrical mandrel - the point remains stationary regardless of the Θ orientation of the mandrel because the mandrel radius is a constant. For an elliptical mandrel, the situation is more complicated. The wire contact point follows an elliptical path in the x-y plane (see figure 3.4). Unfortunately, for even moderate mandrel eccentricities, the entire range of contact points cannot be simultaneously viewed with our current vertical FOV from our stationary mount. Since it is impractical (without an extremely specialized lens-sensor system) to increase the vertical FOV without diminishing the horizontal resolution, our wire gap measurement will have a delay. In any event, the wire feeder should be located in such a manner that the contact point never passes above the upper limit of the camera’s vertical FOV for top winding.

3.6 Camera calibration

Camera calibration is necessary to establish the correspondence between the measured image pixel gap and the actual physical gap on the mandrel surface. We employed the simple pinhole perspective, recognizing that more accurate techniques do exist.^{7,11} Perspective projection uses the camera geometry of an ideal lens to determine the world frame visual area projected onto a single camera pixel. As found in any optics textbook, the fundamental equation of perspective is:

$$\frac{y}{f} = \frac{Y}{X} \quad (4)$$

where y is the size of the image on the sensor plane, Y is the size of the world object, f is the focal length of the lens and X is the distance between the world object and the lens. We are interested in determining Y . We can calculate offline a lookup table relating $X'=X+\Delta X$ to Θ from our known camera location and the shape sensor data. Our gap-determination algorithm returns the number of pixels to which y is proportional, yielding our working equation:

$$Y=(mX'+b)(\text{number of pixels}) \quad (5)$$

where we have grouped all the constants into m and b . Both m and b are dependent on f and thus change with the focal setting for our camera system (we avoid changing it). About 30 pixels are used in a typical measurement. An error in the sampling rate ($mX'+b$) of 0.02 mm/pixels will result in a gap measurement error of 0.6mm. We have observed such errors before adjusting ΔX . We have observed radial distortion effects when the measured gap is at the edge of the image which have influenced an otherwise constant gap measurement by as much as 0.5mm, though these effects are typically minimized because the two wires of interest should never stray far from the center of the image. We intend to upgrade our method of camera calibration by employing Tsai’s techniques.^{7,11}

3.7 The shape sensor

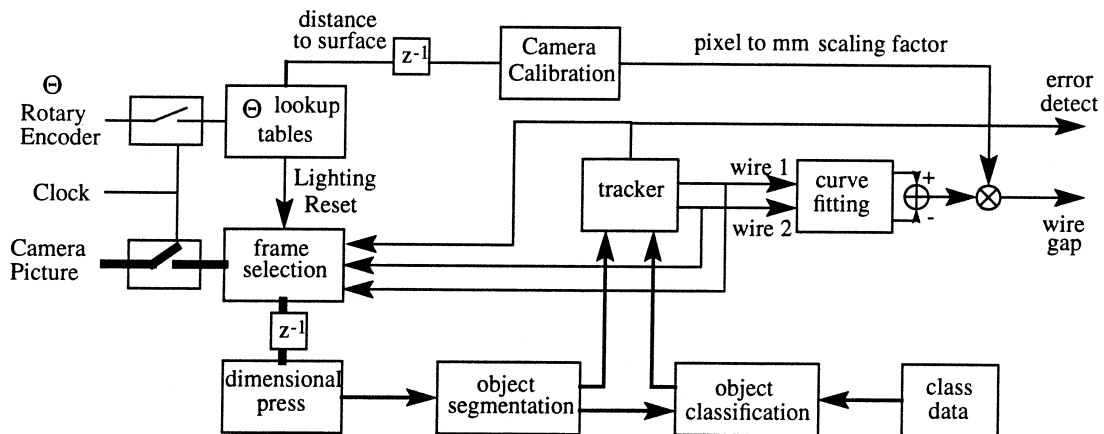
We determine the shape of the mandrel through a secondary mandrel shape sensor which is mounted on the filament support. Protruding from the support are two sensors - a coupled LED/photodiode sensor and a capacitance sensor. By controlling the distance between these sensors and the mandrel surface during rotation (via the filament support), we obtain an accurate model of the mandrel shape. These values are utilized to control the filament winding process, to determine the distance between the camera lens and the mandrel surface, and to calculate the normal at the mandrel surface for use in reflection calculations.

4.0 THE VISION ALGORITHM

4.1 Overview

Figure 4.1 shows the overall block diagram of the algorithm. When operating on a 486 66MHz machine with a real-time DOS kernel simultaneously controlling the winding machine, our algorithm runs at 8Hz, or about one fourth the theoretical limit of the hardware.

Figure 4.1 - Block diagram of gap measurement system



4.2 Vertical averaging - the dimensional press

We take advantage of the inherent vertical symmetry of the problem. Within the selected window from the frame selection stage, the two wires of interest are nearly vertical. They are slightly tilted due to their winding pitch, but this tilt is typically on the order of 0.05mm for 3mm wire. The amplitude deviation between pixels in column i is significantly less than the horizontal deviation across row j . Thus, we perform the computationally fast process of generating a vertically condensed signal with minimal loss of relevant information:

$$CSIGNAL(i) = \sum_{j = wintop; j = j + 2}^{wintop + height} Pixel(i, j) \quad (6)$$

This averaging limits the effect of the noise that might have resulted from observing only one row. It accentuates those objects with vertical symmetry while smoothing and diminishing those without. Our raw image consists of 8-bit pixels. Averaging 15 rows produces a tall and narrow pixel whose resolution is approximately 12-bit. This significantly reduces the quantization noise of what is inherently an 8-bit signal. The peak does tend to get smoothed out in the horizontal direction, but this helps to make the subpixel peak determination routine stable. In addition, we have produced the obvious advantage of reducing our 2-D signal to two 1-D signals. Filtered versions of these signals are used for the segmentation stage. Figure 4.2 shows a representation of this vertical averaging process.

We also calculate TSIGNAL, the moving average of CSIGNAL:

$$\text{TSIGNAL}(i) = T_{\text{scale}} \sum_{j=i-W'_T}^{i+W'_T} \text{CSIGNAL}(j) + T_{\text{offset}} \quad (7)$$

$W_T = 2W'_T + 1$ is set equal to the wire diameter (in pixels) plus a small incrementor. A window of this size, or even one slightly wider, will produce a TSIGNAL that will be lower than CSIGNAL at the wire peak after the appropriate scaling, thus it is an appropriate signal to segment CSIGNAL with.

In the dimensional press stage, we also calculate the derivative of CSIGNAL. Since we are working with discrete signals, we use the difference equation:


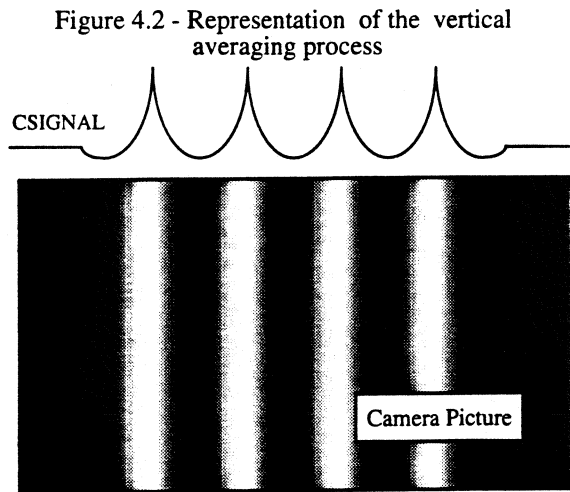
$$\text{DSIGNAL}(i) = \text{CSIGNAL}(i) - \text{CSIGNAL}(i-1) \quad (8)$$

As is usually the case when derivatives are taken, this signal is fairly noisy. Nonetheless, the derivative around the specular peak of a wire follows a standard pattern: the derivative slowly ramps up to a maximum value, drops precipitously, passing through zero to a minimum (negative) value before returning to a signal close to zero, just as we would expect to observe with a vertical edging mask on the raw image. We use the locations of the minimum, maximum and zero crossing values of DSIGNAL to calculate some of our classification features.

4.3 Signal segmentation

Now that we have successfully reduced our picture to a set of one dimensional signals, we proceed to locate intermediate objects, a process called segmentation. Typical segmentation is accomplished by defining a threshold level and defining an object as a region where the signal (2-D or 1-D) rises in amplitude above the threshold. The choice of a threshold level is crucial. A threshold set too high can cause objects to be missed by the algorithm, while a threshold set too low can result in two separate objects blurring into one. A low threshold also increases the number of noise objects observed.

This trade-off is the motivation behind calculating TSIGNAL. TSIGNAL acts as an adaptive threshold level for our 1-D CSIGNAL. Because TSIGNAL adaptively changes according to the local CSIGNAL amplitude, it captures objects that stick out of the background. A careful choice of W_T (see above) ensures that TSIGNAL will be optimally calculated to make the CSIGNAL wire peaks stand out sharply. The troughs on either side of the wire will average out against the specular crest resulting in a signal well below CSIGNAL at the crest. See figure 4.3.



4.4 Wire classification through pattern recognition - overview of Bayesian learning

It is possible to assume that the first two objects detected during the segmentation stage are wires. Such an assumption would frequently be wrong, even though we have attempted to bias the segmentation stage in favor of true wires. We require a technique to isolate true wires from noise objects like wet filaments or mandrel edges.

Bayesian Learning is a standard method of determining discriminant functions by assuming a multivariate normal distribution of clusters of objects called classes. Once obtained, the discriminant functions define a decision boundary between classes. This boundary is automatically computed from training data (feature vectors) that have been assigned to particular classes by an outside expert. We will present the main equations for implementing the approach, referring the reader to any one of the excellent references on the subject.^{8,12}

We assume that we have selected and devised a means of extracting appropriate discriminating features from the objects that were found during the segmentation stage. Currently all classes (the wire class and the various noise classes) are taught by a human. We intend to use automated clustering (ISODATA) to assign the noise classes in a future version.

The program estimates the mean feature vector and the variance-covariance matrix for each taught class. The mean feature vector, μ_c , corresponds to the center of the taught class cluster c . It is computed by taking the average of each feature vector x_i assigned to that class:

$$\mu_c = \frac{1}{N_c} \sum_{i=1}^{N_c} x_i \quad (9)$$

where N_c is the number of instances of class c and x_i is the i th feature vector for class c . The sample variance-covariance matrix for class c is calculated as

$$\Sigma_c = \frac{1}{N_c - 1} \sum_{i=1}^{N_c} (x_i - \mu_c) (x_i - \mu_c)^T \quad (10)$$

Having calculated μ_c and Σ_c for all of our classes, we can then utilize the Bayesian evaluation function:

$$g_c(x) = -\frac{1}{2} (x - \mu_c)^T \Sigma_c^{-1} (x - \mu_c) - \frac{n}{2} \log(2\pi) - \frac{1}{2} |\Sigma_c| + \log P(c) \quad (11)$$

where $P(c)$, the a priori probability of class c , is determined by the relative populations of the various classes. To determine which class an object belongs to, $g_c(x)$ is calculated for all classes. $g_{\text{wire}}(x)$ is then normalized over the sum of $g_c(x)$ for all c to determine the probability that this frame object is a wire. Each object with its corresponding probability is passed to the tracking block.

4.5 Tracking from frame to frame

While the classifier produces quite good results, there are occasionally frames where it will incorrectly identify a wire as noise or noise as a wire. We therefore use the cumulative information from previous frames to help characterize objects from frame to frame. We can then build up cumulative statistics pertaining to objects tracked from frame to frame.

We construct a tracked object when we detect a wire object within the most recent frame that doesn't "match" any of the tracked objects in memory. We destroy any tracked object which hasn't been observed for some time or that has dropped below a certain cumulative probability of being a wire.

We employ a simple 1-D tracker that defines a "match" between a screen object and a tracked object. A screen object is said to match a tracked object when the number of overlapping pixels is more than the number of non-overlapping pixels. If two objects are said to match the same object, then the object with the highest overlap factor takes precedence.

Tracked objects go through a "trial period" during which the tracker compiles their cumulative probability of being a wire. After two or more tracked objects have passed this trial period, the program begins calculating the gap between the two

wires. The program also calculates a smaller picture window surrounding these objects which is used in the next iteration of the algorithm.

4.6 Wire reflectance peak localization and gap determination

After the tracker has determined that it has found two wire objects, the true centers of these objects must be found. We locate the wire by fitting a 2nd order polynomial curve to the data and determining the critical point. We also tried a Gaussian curve and found little discrepancy from our polynomial fit.

5.0 EXPERIMENTAL RESULTS

5.1 Still frame accuracy

To determine still frame accuracy, we first adjusted the vision sampling rate to yield a peak to peak distance equal to the width of a single wire (3.1mm) for two wires that were physically touching. We then used calipers to measure the total distance spanned by a collection of wires. Each gap between the wires (ranging from 0mm to 2.0mm) was measured with the algorithm multiple times and then averaged. These averages were then summed across the entire interval and compared with the result obtained from the calipers. The implied average gap error was 0.023mm/gap.

5.2 Still frame precision

The standard deviation of the subpixel gap measurement for a stationary mandrel was 0.020mm. This favorably compares to our pixel resolution: 0.114mm/pixel.

5.3 Moving precision - cylindrical mandrel

Figure 5.1A and 5.1B respectively show polar and linear plots of data collected from a cylindrical mandrel under dry conditions. The polar angle (Φ) is the position on the mandrel as defined in section 2 and the radius is the amplitude of the respective signals. The algorithm correctly observed the wire winding "tighten up" to a small constant gap. The standard deviation for the steady state gap reading was +/-0.023mm (coincidentally, this is the same as the still frame accuracy), which was not significantly different from the stationary mandrel data.

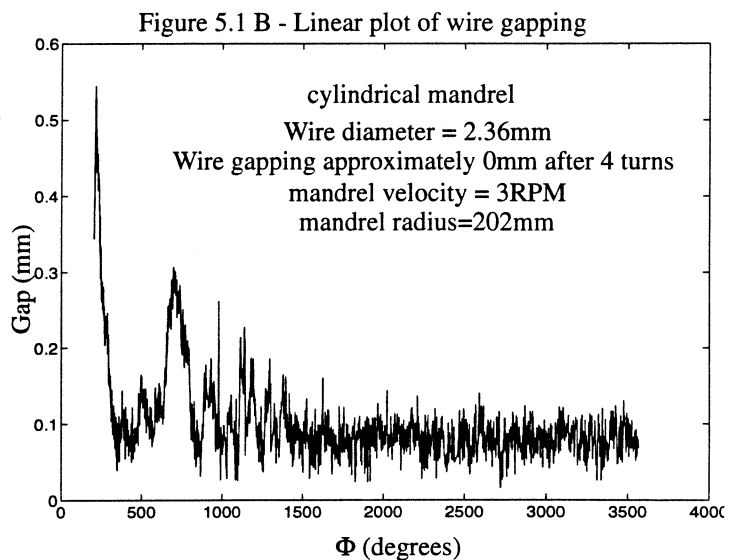



Figure 5.2 shows a data sample from a cylindrical mandrel under wet resin conditions. Note that, when the mandrel stops rotating, there is a significant drop in the noise level. The standard deviation for the moving and stationary regions are 0.056mm and 0.034mm respectively, which is significantly higher than the dry mandrel.

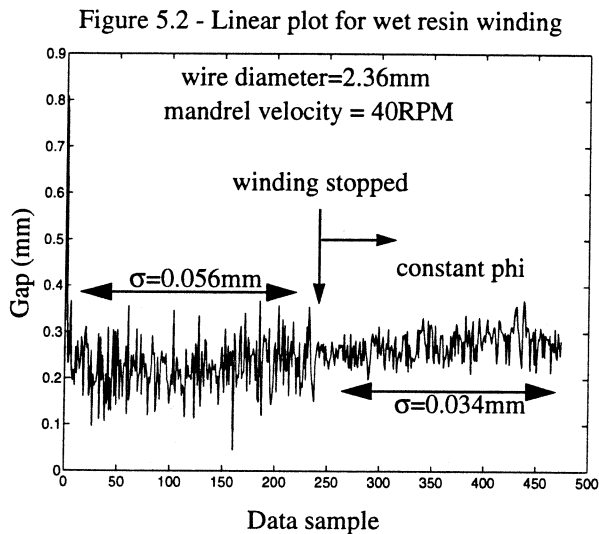
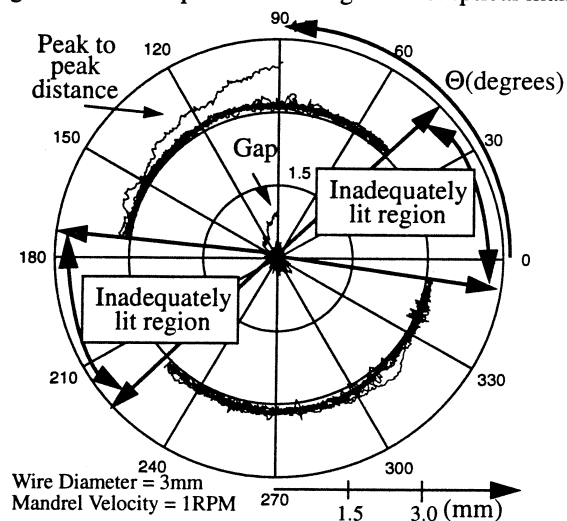


Figure 5.3 - Polar plot for winding on an elliptical mandrel



5.4 Moving precision - elliptical mandrel

Figure 5.3 shows the polar plot of data taken from an elliptical mandrel under dry conditions with controlled lighting. The Inadequately lit regions are mandrel orientations where the light ray would have to originate from one of the problematic ray directions discussed in section 3.3. As with the cylindrical winding, the program detects as the sensed gap diminishes from its initial value.

The standard deviation for the gap signal after it had tightened up was 0.062mm, however there are a few spikes which extend as much as 0.25mm away from the average value. Some of the smoother spikes could be the result of a true gap change. However the sharper, single-sample spikes are clearly noise considering the slow mandrel speed. Note that the standard deviation for this signal is around three times that of the cylindrical mandrel sample shown in figure 5.1 even though the elliptical mandrel is rotating at 1/3 the speed. We suspect this is a combination of calibration errors due to the changing mandrel distance, lighting factors due to the changing normal and actual gap variations.

6.0 CONCLUSION

We have developed an algorithm that determines the gap between wires on a rotating transformer core in an absolute metric. This measurement can be used to passively observe the wire gapping to determine the effectiveness of open loop control or to act as a feedback signal in a closed loop system. The algorithm can be extended to determine the separation between any two objects with prominent, localized, vertical specular highlights.

We have shown that it is possible to measure gaps at speeds as high as 60RPM to within 0.15mm. On a x486, 66MHz machine, while sharing the CPU with the winding control algorithm, our algorithm operates at 8Hz. A faster machine, such as a Pentium, will push this figure closer to the 30Hz limit.

Gap detection for a moving cylindrical mandrel during wet winding results in a standard deviation of 0.056mm. For elliptical mandrels, there are additional errors that increase this deviation to 0.062mm.

The utility of this work arises from the ability to implement the vision system in the control of the gapping between wires. One method is to use the difference between the desired and the measured gap signals in the conductor arm feedback loop. There is another, perhaps simpler way to control the gapping. In order to wind the wires tightly, it is necessary to have the conductor arm slightly behind the winding point (in the z-direction). This applies tension in the z-direction as the

fresh wire presses against the adjacent wire. Too little tension results in an appreciable gap. Too much tension causes the fresh wire to "jump" the adjacent wire, breaking down the winding process. It is possible to roughly determine the amount of tension by measuring the distance between the conductor arm and the location of the fresh wire. With the camera mounted on the conductor wagon, this distance will be the same as the distance between the fresh wire and an arbitrary pixel location plus an offset factor. Thus we can use the absolute position of the wire to determine an indirect measurement of the tension.

7.0 ACKNOWLEDGEMENTS

This research was funded in part by ABB's transformer winding factory in Brilon Germany. The authors wish to thank Mr. A. Davidsson, Mr. T. Rendenbach, Mr. R. Mengerhausen and Mr. V. Mengerhausen of ABB-Brilon for their help in conducting this work.

8.0 REFERENCES

1. Cook, R.L., and Torrance, K.E. "A Reflectance Model for Computer Graphics" *Computer Graphics* 15 (3): 307-316, 1981
2. Dolan, J.M., Khosla, P.K., Talukdar, S.N., "Surface-Closure Algorithms for Filament Winding of Non-Axisymmetric Cylindrical Parts", *Proceedings of the 25th International SAMPE Technical Conference*, Vol 25, pp.680-691, Oct. 1993
3. Efremov, E.D., "The Equation of Yarn Winding on a Surface of Rotation." *Technology of Textile Industry U.S.S.R.* 1965 No. 2
4. Harding, K.G., Bieringer, R., "High Intensity Light Line Using High Pressure Sodium Lamps", *SPIE- Optics, Illumination and Image Sensing for Machine Vision VII* 1992, Volume 1822
5. Healey, G.E. "Using Color for Geometry-Insensitive Segmentation." *Journal of the Optical Society of America* A6(6): 920-937, June 1989
6. Healey, G., and Binford, T. "Local Shape from Specularity." *Computer Vision, Graphics and Image Processing* 42: 62-86 1988
7. Lenz, R.K., Tsai, R.Y., "Techniques for Calibration of the Scale Factor and Image Center for High Accuracy 3-D Machine Vision Metrology", *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol 10, No. 5, September, 1988
8. Nadler, M., Smith, E., 1993, *Pattern Recognition Engineering* (John Wiley and Sons)
9. Nayar, S.K., Ikeuchi, K., and Kanade, T. "Surface Reflection: Physical and Geometrical Perspectives." *IEEE Transactions on Pattern Analysis and Machine Intelligence*
10. Ray, R., Birk, J., and Kelley, R.B. "Error Analysis of Surface Normals Determined By Radiometry." *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-5 (6): 631-645 November 1983
11. Tsai, R., "A Versatile Camera Calibration Technique for High-Accuracy 3D Machine Vision Metrology Using Off-the-Shelf TV Cameras and Lenses", *IEEE Journal of Robotics and Automation*, Vol RA-3, No. 4, August 1987, 323-344
12. Therrien, C.W., 1989 *Decision, Estimation and Classification - An Introduction to Pattern Related Topics* (John Wiley and Sons)
13. Thrift, P., Lee, C.-H., "Using Highlights to Constrain Object Size and Location." *IEEE Transactions on Systems, Man and Cybernetics* TSMC-13 (3) 426-431 1984
14. Torrance, K., Sparrow, E., "Theory for off-specular reflection from roughened surfaces" *Journal of the Optical Society of America* 57 1105-1114 1967