

Design of a Vision-Based Sensor for Autonomous Pig House Cleaning

Ian Braithwaite

*Automation Section, Ørsted-DTU, Technical University of Denmark, 2800 Kongens Lyngby, Denmark
Email: idb@oersted.dtu.dk*

Mogens Blanke

*Automation Section, Ørsted-DTU, Technical University of Denmark, 2800 Kongens Lyngby, Denmark
Email: mb@oersted.dtu.dk*

Guo-Qiang Zhang

*Danish Institute of Agricultural Sciences, Research Centre Bygholm, P.O. Box 536, 8700 Horsens, Denmark
Email: guoqiang.zhang@agrsci.dk*

Jens Michael Carstensen

*Informatics and Mathematical Modelling, Technical University of Denmark, 2800 Kongens Lyngby, Denmark
Email: jmc@imm.dtu.dk*

Received 24 February 2004; Revised 15 February 2005

Current pig house cleaning procedures are hazardous to the health of farm workers, and yet necessary if the spread of disease between batches of animals is to be satisfactorily controlled. Autonomous cleaning using robot technology offers salient benefits. This paper addresses the feasibility of designing a vision-based system to locate dirty areas and subsequently direct a cleaning robot to remove dirt. Novel results include the characterisation of the spectral properties of real surfaces and dirt in a pig house and the design of illumination to obtain discrimination of clean from dirty areas with a low probability of misclassification. A Bayesian discriminator is shown to be efficient in this context and implementation of a prototype tool demonstrates the feasibility of designing a low-cost vision-based sensor for autonomous cleaning.

Keywords and phrases: computer vision, spectral characterisation, Jeffreys-Matusita, Bayesian discriminant analysis, robotic cleaning.

1. INTRODUCTION

Manual cleaning of livestock buildings, using high-pressure cleaning technology, is a tedious and health-threatening task conducted by human labour in intensive livestock production. To remove this health hazard, recent development has resulted in cleaning robots, some of which have been commercialised. The working principle of these robots is to follow a pattern initially taught to them by the operator. Experience shows that cleaning effectiveness is poor and utilisation of detergent and water is higher than for manual cleaning. Furthermore, robot cleaning entails subsequent manual cleaning as robots are unable to detect the cleanness of sur-

faces. The essence of experience is that the key to success of autonomous cleaning will be a sensing system that can determine where cleaning effort should be concentrated.

In the cleaning problem, the first issue to be solved is to define the level of cleanness required. A subsequent issue is to develop methods to discriminate effectively between remains to be removed and the background. The possible ways to categorise remains include chemical and optical composition and shape, which differ from those of the building materials. Sensing remains with specific characteristics could call for vision or ultrasound or laser-based principles.

In this paper, we discuss the spectral properties of residues of organic materials, mainly manure, and report how these alone can be used for discrimination from building materials used in livestock buildings, even though change in optical properties caused by wear of the background is a major issue. Other properties, such as surface texture, are

not considered here as it is believed that spectral properties alone, if shown sufficient, will provide a robust method for cleanness measurement. The paper shows that a multispectral vision technique shows good promise to solve the discrimination problem and shows how feature extraction from clean and nonclean surfaces can effectively be used to characterise, with high probability, areas of a surface that need intensive cleaning. The paper shows how statistical classification methods can be adopted to this application and used with promising results. The remainder of this paper is organised as follows.

Section 2 formulates more precisely the problem addressed in this work. Section 3 presents the model of computer vision used in the design. Section 4 describes the procedure used to characterise the material surfaces relevant for the sensor. Section 5 provides insights into the possibilities presented by multidimensional classification, and measures available to analyse multidimensional data. Section 6 presents the design of the sensor and the discrimination procedure used. Section 7 describes a prototype used to demonstrate the design, and Section 8 presents the conclusions of the work.

2. REQUIREMENTS

The requirements for cleanness detection and to the context in which an intelligent sensor should operate have been considered in detail by Strøm et al. [1]. Pig houses consist of pig pens: the floor and wall areas to clean in each pig pen are typically 26 m². Wall surfaces could be steel or plastic, and floor areas are made of concrete.

The main purpose of the cleaning process in livestock buildings is to reduce the risk of infection between batches of animals. Experience from real-life pig production shows that a visually clean house reduces infection pressure to an acceptable level. The absence of visible contamination with dirt and/or manure may thus be considered to be a proper definition for remote, online detection of cleanness of housing equipment after washing.

Batch production with cleaning between batches is used in growing and finishing pig houses, and most of the cleaning time and effort is spent in the finishing houses. The project therefore focuses on cleaning in finishing pig houses.

A fundamental requirement is to be able to distinguish between clean and nonclean conditions of a surface. The probability of misclassification is a crucial parameter to consider, noting that the required highest misclassification rates for clean and dirty surfaces are not equal. A clean surface being characterised as unclean has the consequence of adding an additional round of cleaning of an area. Misclassifying a dirty area as clean would leave specks of dirt uncleaned.

Noting that the remains of manure are inhomogeneous and unevenly distributed, and that specks can be expected anywhere over the surface, it is required that any area segment of 2 mm in diameter can be reliably characterised as clean or dirty.

In a final implementation, the sensor is envisaged to move around in the pig pen with the robot arm, record

the cleaning level, and pass on the obtained data to the cleaning robot. The level of embedded functionality of the cleaning sensor would be high enough to let it be used for autonomous operation with a cleaning robot. Several complex features could be considered for a final system, including correlation between area segments, texture properties, cleaning history, records of cleanness conditions, and records from past cycles of cleaning.

However, the fundamental issue is whether a clean condition can reliably be discriminated from the dirty one for individual segments of size 2 mm in diameter.

2.1. Clean surfaces

The requirement for cleaning in livestock buildings is that all visible dirt be removed. This means to remove visible organic contaminants, in our case down to a size specified above. This would enable subsequent chemical disinfection of the surfaces, should this be desired. The higher levels of cleaning, where contamination by microorganisms needs to be removed, is not within the scope of this paper.

2.2. Properties of contamination

Contamination of housing equipment in pig buildings is expected mainly to consist of bedding and faecal materials, but it is likely also to contain traces of skin and feed. It was therefore important that the contaminants in pig houses were thoroughly characterised in view of alternative sensor principles at an early stage of the project. According to Møller et al. [2], pig manure is characterised as an average content of dry matter 242 g/l with content of volatile solids 838 g/kg dry matter. Thus, more than 80% of manure is organic material and less than 20% is inorganic. The contamination on housing surfaces in pig houses may be slightly different from fresh manure characterised in this reference.

The chemical composition of the organic contents of the residues left on the surfaces in the pig environment indicate there should be a possibility to discriminate residues from building materials based on spectral properties.

A complication is that parts of buildings are made of inorganic materials, in particular concrete and steel, while other areas are made of plastic coated elements or wood. Figure 1 shows an inhabited compartment before cleaning.

2.3. Functional requirements

The requirements for the sensor were defined in terms of what the combined sensor and robot system must achieve. The cleanness sensor will be able autonomously to

- (i) identify selected surface types in finishing pig houses;
- (ii) distinguish a nonclean from a clean surface in a raster size of less than 2 mm;
- (iii) specify position and area of nonclean parts of the surface;
- (iv) function reliably in the environment of the empty pig house during cleaning;
- (v) define surfaces to clean, specify cleaning parameters, and report results.



FIGURE 1: A pig pen with solid and slatted floor in a finishing pig building.

Based on the requirements, the first issue to investigate is the fundamental question of which properties of clean and dirty surfaces would best inform on cleanness of the individual segments.

3. VISION MODEL

The basic elements in a vision-based measurement system consist of three components: illumination, subject, and camera. The complexity of the system is to a large degree determined by the extent to which the relative placement of the three can be controlled and constrained. Particularly in the case of illumination, control is often critical—external (stray) light can be a seriously limiting factor for system effectiveness.

The light collected by a camera lens is determined by the colour of the viewed object, the spectra of the illuminating light sources, and the relative geometries of these to the camera. The influence of geometry on measured colour can be seen clearly with reflective materials: when viewed such that a light source is directly reflected on a surface, the reflected light is almost entirely determined by the light source rather than by the reflecting material. Since this behaviour makes measurement of surface colour impossible, measurement systems attempt to avoid this geometry. In the measurement system described in Section 4, for example, surfaces are illuminated at 45° and observed at 90° relative to the surface plane.

If direct reflections can be avoided, the light reaching a viewer can be considered to be independent of the measurement system geometry. In this case, the spectrum of the light entering the camera is simply a function of the spectra of the light sources and the colour of the material being viewed. Uniform (homogeneous) materials have single colour, while composite (inhomogeneous) materials have varying colours across a surface.

The camera itself has a wavelength-dependent sensitivity across a range of wavelengths. For a CCD camera, this sensitivity covers the visible wavelengths, and some of the adjoining ultraviolet and near-infrared bands. A typical sensitivity curve is shown in Figure 2; sensitivities vary considerably with CCD type and coating treatments.

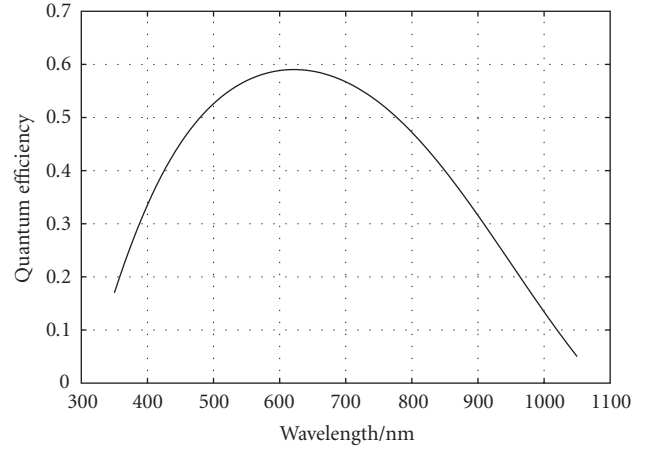


FIGURE 2: CCD quantum efficiency $\eta(\lambda)$.

The sensor designer has some control over the subject illumination. Since different illumination spectra can reveal different features in the viewed object, a number of images can be captured under different lighting conditions. So, the sensor operates with a number of *channels*, where each channel is defined by the spectra of the corresponding light source. Thus, a sensor pixel measurement consists of a vector of readings, one for each channel:

$$x = (I_1 \cdots I_p). \quad (1)$$

Then, all the image pixels can be assembled into an image array, giving a single camera measurement of the form

$$X = \begin{bmatrix} x_{11} & \cdots & x_{w1} \\ \vdots & \ddots & \vdots \\ x_{1h} & \cdots & x_{wh} \end{bmatrix}, \quad (2)$$

where w and h are the width and height of the image in pixels, respectively.

4. PROPERTIES OF CLEAN AND DIRTY SURFACES

How to catch the cleanness information on the different types of surfaces is the major issue in the design of an intelligent sensor. One of the hypotheses is that the reflectance of building materials and contamination differs in the visual or the near-infrared wavelength range. To validate the hypotheses, the optical properties of surfaces to be cleaned and the different types of dirt found in finishing pig units were investigated in the VIS-NIR optical range. If this method is validated, an ordinary CCD camera with defined light sources could be used for cleanness detection.

Measurements were conducted on materials of varying ages and conditions taken from pig houses. The selected housing elements of inventory materials were placed in a pig production building for 4 to 5 weeks in real pig pens, and then removed for study. In all, four surface materials

were considered: concrete, plastic, wood, and metal, in each of four conditions: clean and dry, clean and wet, dry with dirt, and wet with dirt. In each measurement condition, spectral data were sampled at 20 randomly determined positions, in order to avoid the effect caused by the nonhomogeneous properties of the measured surfaces. At each measurement position, spectral outputs were sampled 5 times with an integration time of 2 seconds for each. The average of the five spectra was recorded for analysis.

The spectrometer used in the characterisation was a diffraction grating spectrometer,¹ incorporating a 2048-element CCD (charge-coupled device) detector. The spectral range 400 nm–1100 nm was covered using a 10 μ m slit, giving a spectral resolution of 1.4 nm.

The light source used was a Tungsten-Krypton lamp² with a colour temperature of 2800 K, suitable for the VIS/NIR applications from 350 nm to 1700 nm.

A Y-type armoured fibre optic reflectance probe, with six illuminating fibres around one read fibre (400 μ m) specified for VIS/NIR, was used to connect the light source, the spectrometer, and the measurement objective aided with a probe holder. The probe head was maintained at 45° to the measured surface and a distance of 7 mm from the surface.

The primary results on the reflectance of the different materials under the measurement conditions are showed in Figures 3a, 3b, 3c, and 3d. The curves show the data from the 20 random measurement points under each measurement set-up.

The spectral analysis system has its highest sensitivity in the range 500 to 700 nm, but the entire range from 400 to 1000 nm is useful to provide reflection as function of wavelength. The results suggest that it will be able to make a statistically significant discrimination and hence classify areas that are visually clean. A scenario with multispectral analysis, combined with appropriate illumination or camera filters, is therefore being pursued.

Concrete, the predominant material used for floors, is an inorganic material. The manure and the contaminants may thus be spotted as organic materials on an inorganic background. Under wet conditions, a significant difference may be seen in wavelengths of 750–1000 nm; see Figure 3a. However, the clear differences for steel (stainless) are shown in 400–500 and 950–1000 nm; see Figure 3b. For the brown wood plate, the reflectance under dirty-wet conditions was higher in the wavelengths of 500–700 nm and lower in 750–1000 nm compared with clean-wet conditions. For the green plastic plate, the reflectance under dirty-wet conditions was lower for wavelengths lower than 550 nm and higher than 800 nm, but higher in wavelengths between 600–700 nm compared with clean-wet conditions.

5. CLASSIFICATION

Classification of a surface part as clean or not clean has obvious consequences in the application. For the clean surface, misclassification as not clean will call for another round of cleaning by the robot. Misclassification of the unclean surface as clean has consequences for the quality of the cleaning result. Subsequent manual inspection and cleaning should be avoided if possible, but it could be acceptable for a user to have certain areas characterised as uncertain, as long as these do not constitute a large part of the total area to clean.

With a clear relation between cost and the probability of misclassification, methods to extract features of the observed spectra would be preferred, that could minimise the probability of misclassification, constrained by the complexity of the vision system.

In our context, the number of frequency bands to be analysed has an impact on both the cost of computer-vision equipment and the time needed to capture and analyse the pictures taken.

Several classical methods exist that provide measures of misclassification and separability between the clean and unclean cases.

Let a frequency band in the spectrum be chosen for analysis. A set of measurements on a surface will have a distribution of reflectivity due to differences in the clean surface itself and due to the uneven distribution of residues to be cleaned. Let the distribution function for an ensemble of measurements given the case is θ_i (clean or not clean) be

$$f(r|\theta_i) = f_i(r). \quad (3)$$

A convenient first assumption for analytical discussion is that the population has a multivariate normal distribution of dimension n :

$$\begin{aligned} f_i(r) &= N(\mu_i, \Sigma_i) \\ &= \frac{1}{\sqrt{2\pi}^n \sqrt{\det \Sigma_i}} \exp \left(-\frac{1}{2} (r - \mu_i)^T \Sigma_i^{-1} (r - \mu_i) \right). \end{aligned} \quad (4)$$

With two such distributions for the clean and unclean cases, respectively, the problem is to determine, from one or more measurements of reflectance, whether a given measurement represents a clean or an unclean area. This is illustrated in Figure 4 where a measurement would be characterised as representing a clean area when reflectance is below the borderline between the two distributions. The figure also shows the probability for misclassification. In signal processing, measurement noise is often the prime source of misclassification and repeated measurements would be used to increase the likelihood that the right decision is made.

When noise is the prime nuisance, optimisation based on the Kullbak divergence is often used, for example, in a symmetric version of the divergence [3]:

$$D_{ij} = E \left\{ \ln \frac{f_i(r)}{f_j(r)} | \theta_i \right\} + E \left\{ \ln \frac{f_j(r)}{f_i(r)} | \theta_j \right\}. \quad (5)$$

¹EPP2000-VIS, StellarNet Inc., USA.

²SL1, StellarNet Inc., USA.

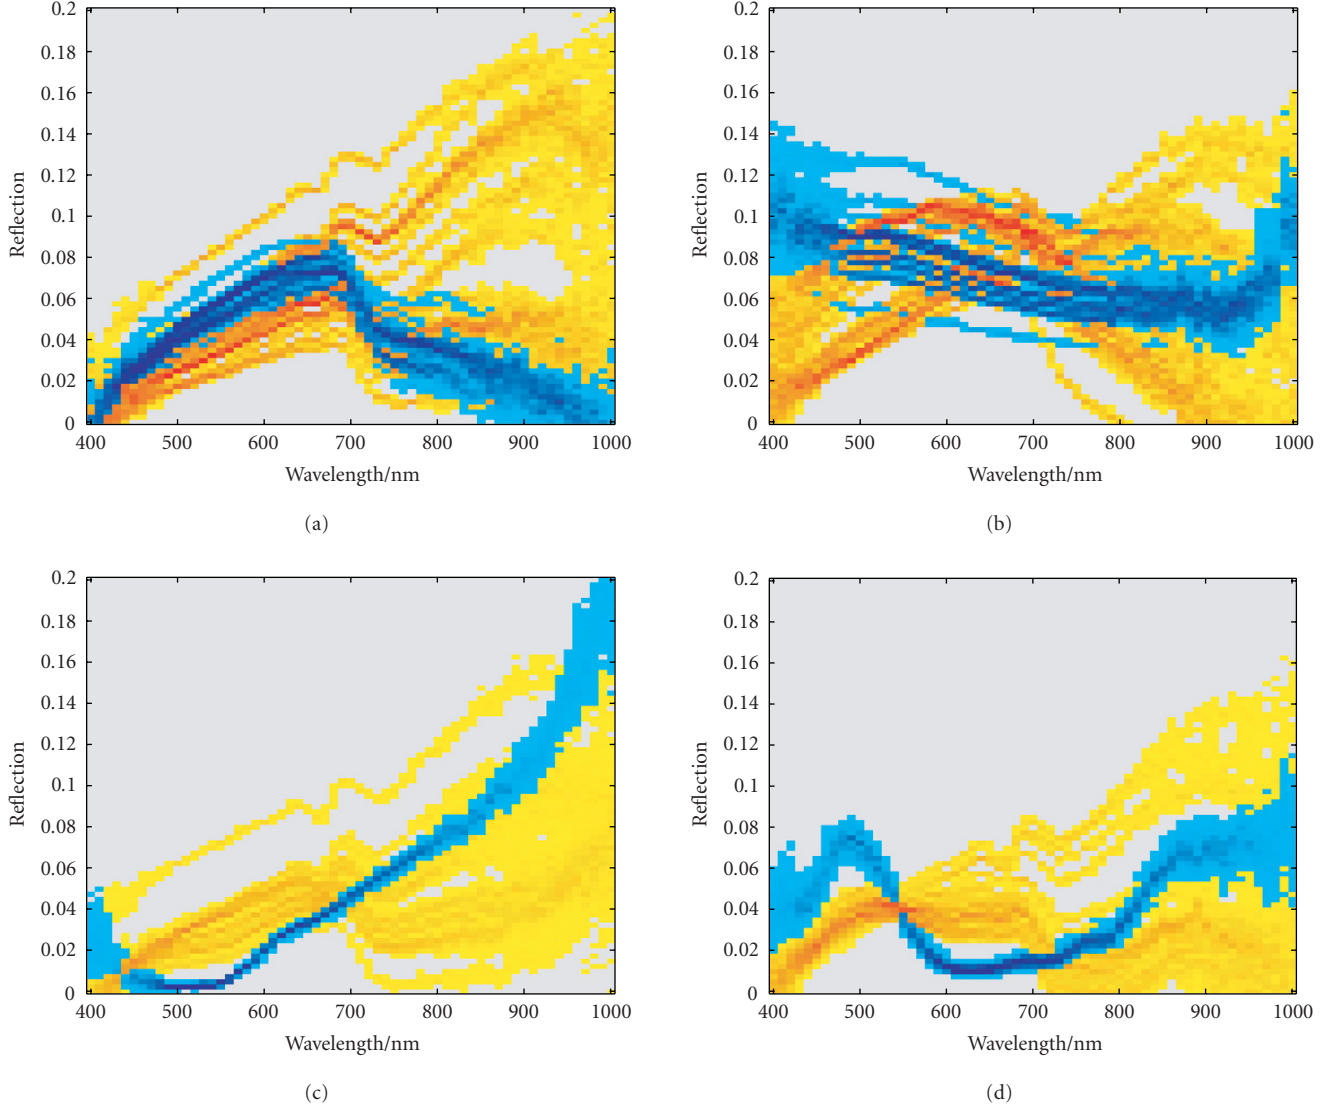


FIGURE 3: Spectral data histograms for four selected materials: (a) concrete slab, (b) steel slab, (c) brown wood plate, (d) green plastic plate. Blue points indicate readings from clean surfaces, yellow/red points indicate readings from dirty surfaces. Darker colours correspond to a higher density of measurement data.

The problem at hand is peculiar, however, as the reason for uncertainty in data is not the measurement noise. As seen from the data plotted in Figures 3a to 3b, considerable variance is caused by inhomogeneities in reflection from the individual area segments of the clean surface. Simultaneously, there is even larger variation in the reflection caused by inhomogeneous composition of the remains of dirt. One area may comprise specks with tiny remains of straw and solid particles, others are covered by a thin layer of uniform contamination. Taking a number of identical pictures of each area element does not enhance information about the area and does not reduce the likelihood of misclassification.

For this reason, it would be advantageous to find a technique of analysis by which the probability of misclassification P_e is minimised for each single area segment treated individually. As apparent from Figures 3a to 3d, the method

should cope with unequal dispersion matrices of distributions for the clean and dirty cases.

The classical Mahalanobis measure [4] is a distance measure between normal multidimensional distributions with equal dispersion. The Jeffreys-Matusita distance (JM distance) [5] is a generalisation that applies to distributions with nonidentical dispersion and it does not require distributions to be normal. Further, it is possible to derive lower and upper bounds for misclassification. Distance measures are discussed in [3, 6, 7].

The Jeffreys-Matusita distance between distributions f_i and f_j is

$$J_{ij} = \left(\int_{\Omega} \left(\sqrt{f_i(r)} - \sqrt{f_j(r)} \right)^2 dr \right)^{1/2}. \quad (6)$$

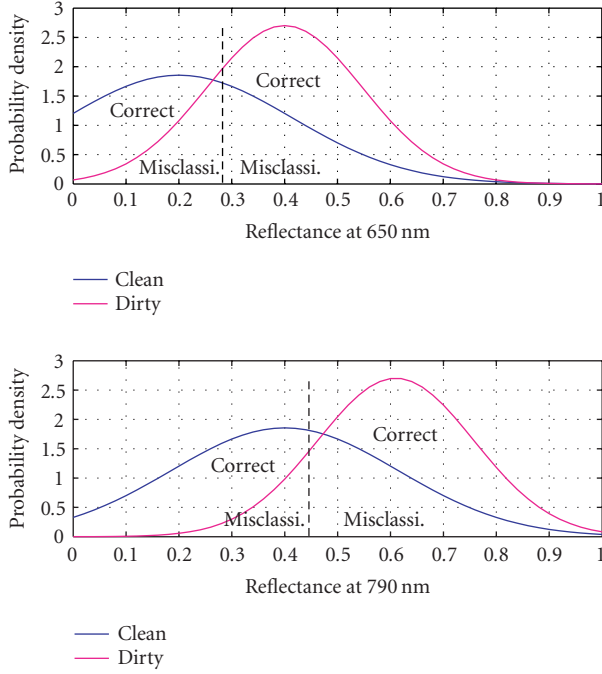


FIGURE 4: Distributions obtained at two distinct wavelengths for clean and dirty cases.

The JM distance measure has been found to be useful in a range of fields, from mineral classification [8] to identification of fungal colonies [9].

The JM distance is $J_{ij} = 0$ when the distributions $f_i(r)$ and $f_j(r)$ are equal. The JM distance takes the value $J_{ij} = \sqrt{2}$ when the two distributions are totally separated. Bhat-tacharya introduced the coefficient

$$\rho_{ij} = \int_{\Omega} \sqrt{f_i(r)} \sqrt{f_j(r)} dr \quad (7)$$

and used the negative logarithm, α_{ij} , of this quantity:

$$\alpha_{ij} = -\ln \rho_{ij}. \quad (8)$$

These have the obvious relation to the JM distance

$$J_{ij}^2 = 2(1 - \rho_{ij}) = 2(1 - e^{-\alpha_{ij}}) \quad (9)$$

Kailath [3] showed that when the two distributions are normal multivariate of degree n : $f_i(r) = N(\mu_i, \Sigma_i)$ and $f_j(r) = N(\mu_j, \Sigma_j)$, then

$$\alpha_{ij} = \frac{1}{8}(\mu_i - \mu_j)^T \Sigma_{ij}^{-1}(\mu_i - \mu_j) + \frac{1}{2} \ln \frac{\det(\Sigma_{ij})}{\sqrt{\det \Sigma_i \det \Sigma_j}}, \quad (10)$$

where

$$\Sigma_{ij} = \frac{1}{2}(\Sigma_i + \Sigma_j); \quad (11)$$

the probability of misclassification P_e is bounded by

$$\frac{1}{8}\rho_{ij}^2 \leq P_e \leq \rho_{ij} \quad (12)$$

which is equivalent to

$$\frac{1}{8}\left(1 - \frac{1}{2}J_{ij}^2\right)^2 \leq P_e \leq 1 - \frac{1}{2}J_{ij}^2. \quad (13)$$

The lower bound of P_e is reached when $J_{ij} = \sqrt{2}$ [3, 6].

Salient features for the present context are the applicability of the JM measure to arbitrary distributions and the bounds for misclassification, although the analytic result of (13) is only an approximation when the distributions are not normal.

5.1. Multispectral techniques to reduce probability of misclassification

The JM distance measure will express the quality of a chosen technique to distinguish between the clean and nonclean surface cases. The complete spectra presented in Section 4 were obtained using a dedicated spectrometer. For commonplace computer-vision techniques to be applied, we need to limit the number of frequencies analysed.

Figure 4 illustrates the theoretic distribution of reflectance for the clean and nonclean cases if monochromatic light is used. The result is two normal distributions with a large overlap. If the discrimination was based on a single wavelength, the area would correspond to misclassification given the surface was clean:

$$P_e(\theta_d|\pi_c) = \int_{r_{\text{sep}}}^1 p(r|\pi_c) dr, \quad (14)$$

where r_{sep} is shown as the dashed line in Figure 4. Misclassification that the dirty surface was declared clean is shown:

$$P_e(\theta_c|\pi_d) = \int_0^{r_{\text{sep}}} p(r|\pi_d) dr. \quad (15)$$

Using monochromatic or narrowband light at a single wavelength was found to give a rather large overlap between distributions and hence a large probability of misclassification for all wavelengths when the surface is made of concrete.

If two monochromatic measurements are used, the distribution is two-dimensional, as illustrated in Figure 5. The two dimensions correspond to observing the reflectance of the surface at two different wavelength bands. The x -axis is reflectance obtained for band λ_1 , the y -axis is that obtained for band λ_2 . While misclassification is large if either of the two wavelength bands are used individually, combining the two observations results in the two-dimensional probability distribution functions shown in Figure 5. The curves shown in Figure 4 are the projections of the same distributions shown in Figure 5. It is indeed possible to discriminate using a separation in the x - y plane of reflectance as the boundary for classification.

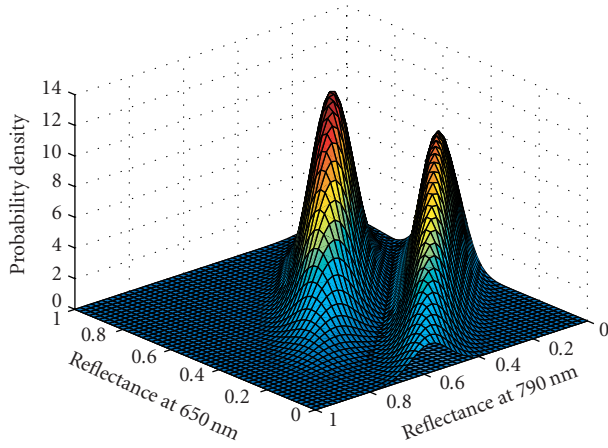


FIGURE 5: The two-dimensional normal distributions of clean and dirty areas are quite well separated and a low probability of misclassification is achieved.

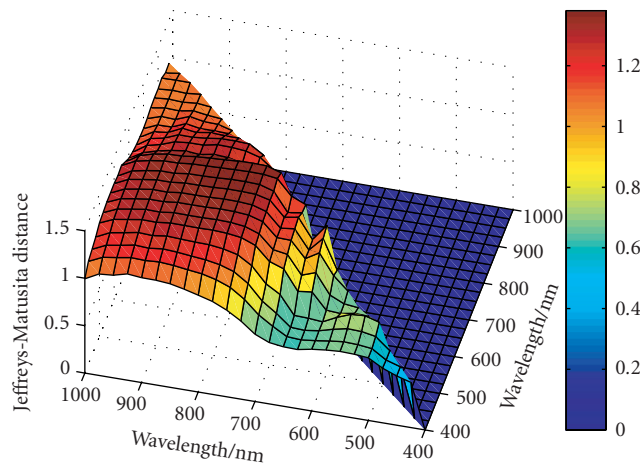


FIGURE 6: Jeffreys-Matusita distance computed for a range of candidate wavelength pairs.

With this approach being promising, the formal approach to design a sensor system should start with choosing two wavelengths, or more, that together give a desired low level of misclassification. Second, a method is needed to find the discriminator function to be used.

5.2. Best choice based on pig pen data

Comprehensive measurement data were obtained from a pig pen where different materials had been in use. The pen had been in use for three months when emptied. Part of the floor was made of newly finished concrete, other parts were 15 years old. Samples of the different materials were collected for analysis.

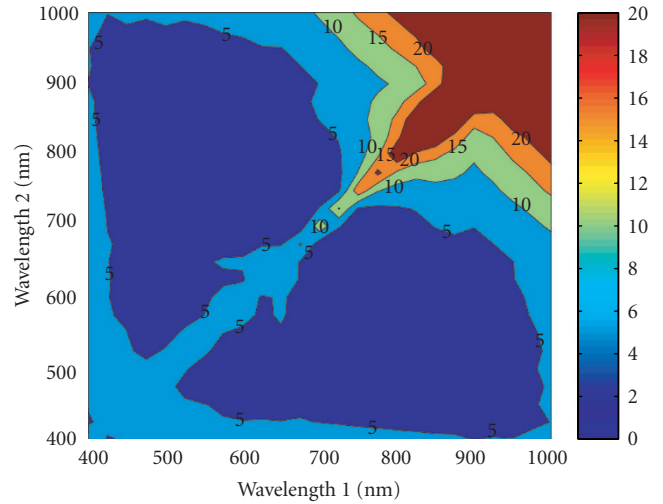


FIGURE 7: Upper bound of classification error as a function of wavelength pairs. The case is a concrete floor after 3-month use of the pen. Measurements were taken in a dry condition.

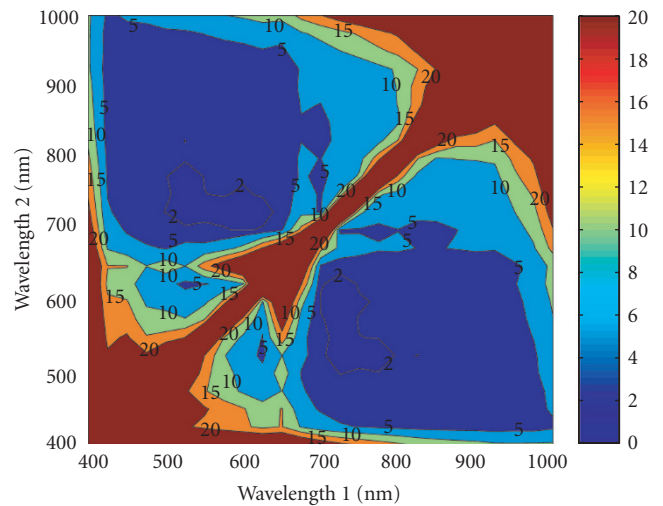


FIGURE 8: Upper bound of classification error. The case is a 15-year-old concrete floor after 3-month contamination in the pen. Measurements were taken in a dry condition.

It was first investigated which pair of two frequency bands would be optimal based on minimising the probability of misclassification. This is equivalent to maximising the JM distance. Figure 6 shows the JM distance measure using measurements from a pig pen that was emptied after three months of use. Data for concrete in dry condition is shown in Figure 7. The upper bound for misclassification probability using two narrowbands around wavelengths 780 nm (infrared) and 650 nm (orange) will have a misclassification probability below 2%. Data for a fifteen-year-old part of the concrete floor in the same pen are shown in Figure 8. Classification can be obtained from multispectral data with a misclassification likelihood below 2%. After cleaning, the floor

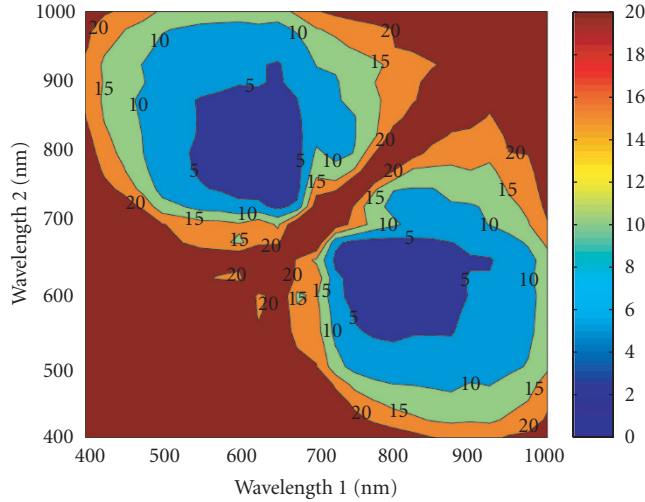


FIGURE 9: Upper bound of classification error as a function of wavelength pairs. The case is a concrete floor after 3-month use of the pen. Measurements were taken in a wet condition.

will be wet, and it is thus essential that discrimination can also be achieved in a wet condition. Figure 9 shows the upper bound for misclassification for the new floor in wet condition. The available ranges for illumination wavelength have now narrowed but good classification is still possible.

The analysis of large sets of experimental pig pen data have thus enabled a selection of illumination parameters to make good classification possible using multispectral discrimination.

5.3. Possibilities for improvement

In order to reduce the misclassification likelihood even further, it may be necessary to consider additional features to aid discrimination. Textural features, for example, could definitely add relevant information under highly controlled circumstances. However, variability in the image acquisition geometry—scale and orientation—for a practical scenario will contribute a significant amount of noise in textural features. In addition to this, it has been seen that the texture itself shows a high variation, and that there is a combined effect of this variation and the abovementioned variation in image acquisition geometry. This could develop into a highly complex texture study, which could be interesting, but which will almost certainly end up with a system less robust and general than a purely spectral system. The aim of this paper has been to check the potential of such a purely spectral system.

6. SENSOR DESIGN

The sensor is pixel based: each pixel is classified either as “clean” or “dirty.” The classification procedure is the Bayesian discriminant analysis, which assigns pixel measurements to classes from which they are most likely to have been produced. The method relies on adequate knowledge of the statistics of the possible classifications, in this work the mea-

surements presented in Section 4 form the basis of the discriminator.

The spectrographic characteristics of pig house surfaces provide much more data than is expected from the camera-based sensor. The camera sensor provides a small number of channels, each described by the pair of filter/illumination characteristics for the respective channel. In the design presented here, each channel is restricted to a narrowband of frequencies, produced, for example, by a number of powerful light-emitting diodes. The sensor collects images corresponding to each channel in turn, by sequencing through the light sources for each channel. By synchronising the light sources to the camera’s frame rate, a set of images corresponding to a single two-dimensional measurement can be acquired in a relatively short time.

6.1. Wavelength selection

From the spectra presented in Section 4, a number of wavelengths must be selected such that classification into clean and dirty classes for pig house surface materials is possible. While some materials may be amenable to classification based on a single light colour (consider, for example, the green plastic in Figure 3d at around 490 nm or 620 nm), concrete, the most important material, is clearly not. However, as illustrated previously, multidimensional analysis can reveal structure that is sufficient to discriminate classes.

Selecting the wavelengths 800 nm and 650 nm, for example, the surface characteristics can be illustrated in the scatter plot shown in Figure 10a. Four populations are shown, corresponding to wet concrete and steel, in both clean and dirty conditions. As can be seen, with these wavelengths, clean and dirty concrete are well separated, whereas clean and dirty steel share a significant overlap. Selecting 650 nm and 450 nm on the other hand, as shown in Figure 10b, separates clean from dirty steel, but fails for concrete. Using all three wavelengths, a discriminator can be constructed to handle both material types.

6.2. Bayesian discrimination

From the training data and the choice of wavelengths determined as just described, a number of populations π_i are modelled as normally distributed, multidimensional random variables:

$$\pi_i \longleftrightarrow N(\mu_i, \Sigma_i). \quad (16)$$

Using the experimental data, x_{ij} , for each class i , estimates of the mean vectors and variance-covariance matrices for each population can be derived:

$$\begin{aligned} \hat{\mu}_i &= \frac{1}{n} \sum_j x_{ij}, \\ \hat{\Sigma}_i &= \frac{1}{n-1} \sum_j (x_{ij} - \hat{\mu}_i)(x_{ij} - \hat{\mu}_i)'. \end{aligned} \quad (17)$$

A Bayesian classifier assigns new measurements to the population to which the measurement is most likely to be

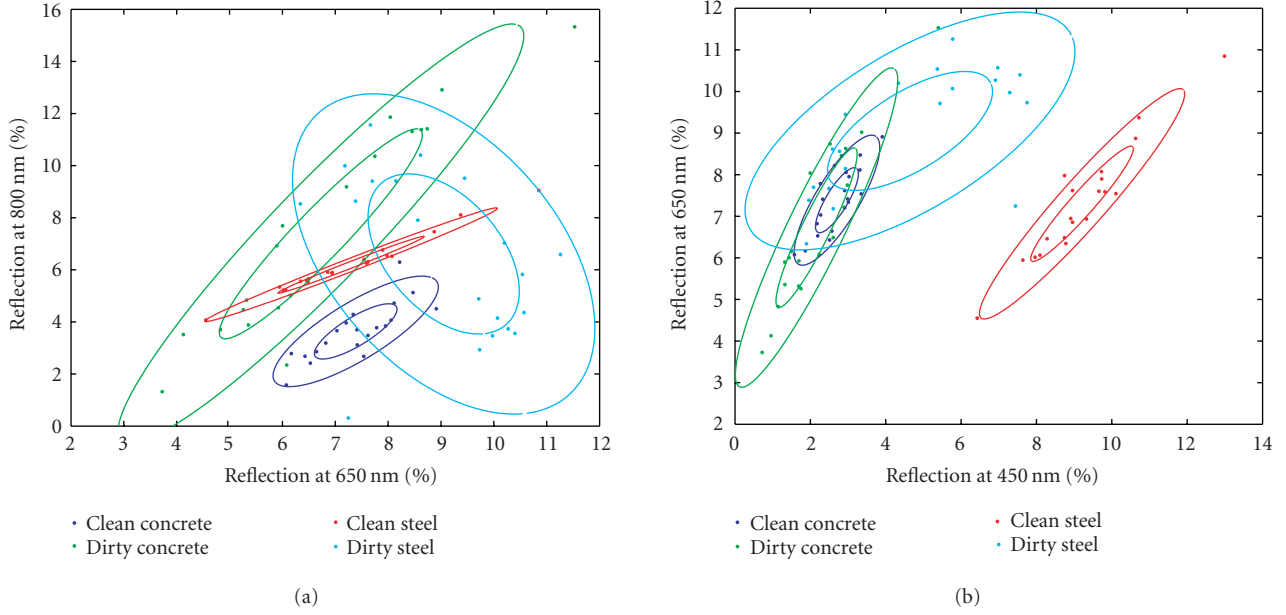


FIGURE 10: Scatter plots showing characteristics of concrete and steel in both clean and dirty states. Each plot shows two selected wavelengths, with each point corresponding to a single spectral measurement. Ellipses show one and two standard deviations for each material state. (a) Possible discriminator for concrete but not for steel—800 versus 650 nm. (b) Possible discriminator for steel but not for concrete—650 versus 450 nm.

associated. Bayes' rule states that the probability that a measurement x is associated with the class π_i is given by

$$P(\pi_i | x) = \frac{P(\pi_i)P(x | \pi_i)}{P(x)}. \quad (18)$$

A Bayesian classifier assigns a measurement to the class for which the probability calculated in (18) is the greatest. The term $P(x | \pi_i)$ is simply the probability distribution function for class i , which can be written as $f_i(x)$, and $P(\pi_i)$ is the prior probability of class i , which can be written as p_i . The denominator $P(x)$ is independent of the class, and is therefore irrelevant with respect to maximising (18) across classes. Thus, a Bayesian classifier chooses the class maximising

$$S_i = f_i(x)p_i, \quad (19)$$

where S_i is being referred to as a *discriminant value* or *score*. The term p_i is the *a priori* probability of a measurement corresponding to the population π_i , and reflects knowledge of the environment prior to the measurement being taken.

In the case of multidimensional normal classes, the probability density is given by

$$f_i(x) = \frac{1}{\sqrt{2\pi}^n} \frac{1}{\sqrt{\det \hat{\Sigma}_i}} e^{-(1/2)(x - \hat{\mu}_i)' \hat{\Sigma}_i^{-1} (x - \hat{\mu}_i)}. \quad (20)$$

Substituting this into (19) gives the discriminant value for the normally distributed case. In practise, the same decision

rule is achieved by applying a monotonic transformation to (19), so that

$$S_i(x) = \log \frac{p_i}{\sqrt{\det \hat{\Sigma}_i}} - \frac{1}{2} (x - \hat{\mu}_i)' \hat{\Sigma}_i^{-1} (x - \hat{\mu}_i) \quad (21)$$

is minimised instead. Since this function is quadratic in x , it is known as a quadratic discriminant function.

7. A PROTOTYPE SENSOR

In order to demonstrate the proposed vision-based classification method, a prototype sensor has been constructed. First, a colour digital video camera was used with white light in a seminatural environment. Subsequently, a monochrome camera with controlled multispectral lighting was used in a real pig house setting.

7.1. Seminatural environment

A prototype vision-based classifier was constructed using readily available components—a desktop computer and colour digital video camera. The prototype demonstrates multivariable statistical classification using two of the three colour channels available in a normal colour image.

In an initial training phase, statistical spectral properties of each type of object to be recognised are measured. This is done by selecting areas of images corresponding to each object type, and calculating mean and variance measures of the pixels in these areas. Based on these statistics, subsequent images are automatically classified, pixel by pixel, to the learned object classes. Each pixel is allocated to the most likely object class, based on the learned statistical properties of each class.

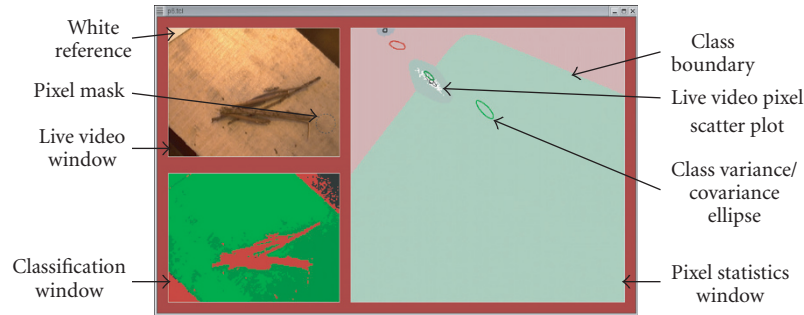


FIGURE 11: Screen shot of prototype classifier. Top left: original camera image. Right: colour map—each dot corresponds to a pixel in the original image, increasing red intensity to the right and increasing blue intensity down. Bottom left: result of classification—each image pixel is now coloured to show its classification. The red area shows what will be cleaned. Ellipses in the right-hand window show trained class definitions.

A continuous lighting calibration is carried out by ensuring that one corner of the image has a known colour. The importance of this calibration underlines the necessity of controlled lighting for the actual sensor.

Figure 11 shows a screen shot of the prototype in action. Live video from the camera is displayed in the top-left window, two-dimensional colour statistics are shown in the right-hand window, and final pixel classifications are shown in the bottom-left window. The display is updated in real time, with a frame rate of 5 Hz.

A pixel mask is displayed in the video window, indicated with a circular, dashed black line. The user is free to resize and move the mask, in order to select regions of interest. The large square window displays two types of information: statistics about the selected live video pixels and definitions of previously learnt object classes. Video pixels are displayed as a scatter plot of white pixels with mean and variance drawn on top in black. The background shows the defined object classes as strongly coloured ellipses, lighter versions of the same colours show the regions the Bayesian classifier associates with each class. Finally, the classification window shows, for each pixel in the original image, which class the pixel is assigned to by the Bayesian classifier. The same colours are used as in the statistics window.

7.2. Pig house environment

In order to test the sensor design under realistic conditions, images of concrete surfaces from a working pig house were collected using a monochrome camera with controlled, strobe lighting. Two wavelengths were selected, one in the visible range—590 nm and one from the infrared one—850 nm. Two images of each of two materials, clean and dirty concrete, were obtained. Each individual image contains 400 by 400 pixels, providing 160 000 individual measurements.

In order to test the method, a classifier was trained using the first clean/dirty image pair and tested using the second image pair. Subsequently the pairs were swapped and the process repeated—twofold cross-validation. Figure 12 shows

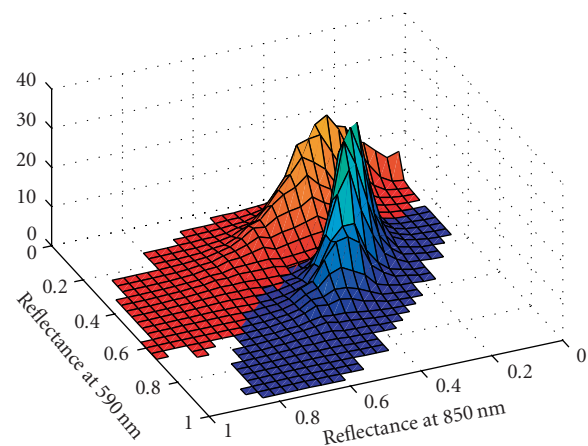


FIGURE 12: Histogram of clean and dirty pixel measurements from two concrete samples: the red surface shows the distribution of dirty pixel measurements, while the blue surface shows the distribution of clean pixel measurements.

the distributions obtained from the first image pair, which can be compared with the theoretical case presented earlier and illustrated in Figure 5.

The images obtained and the resulting classifications are shown in Figure 13. A simple median filter has also been applied to the final classification to remove some of the classification “noise,” particularly in the clean case. Misclassification in the clean case is expensive in the application, since it results in wasted cleaning effort. Classification accuracy, both with and without the filter, is presented in Table 1. The error rates are in good agreement with the earlier analyses.

Several approaches to deal with the misclassification are possible. Firstly, further postprocessing of the classified image, perhaps considering textural properties, could be used to identify misclassified areas. For areas falsely classified as dirty, this would be very helpful. Small areas falsely classified as clean are probably less problematic—repeated cleaning close by will in any case be required.

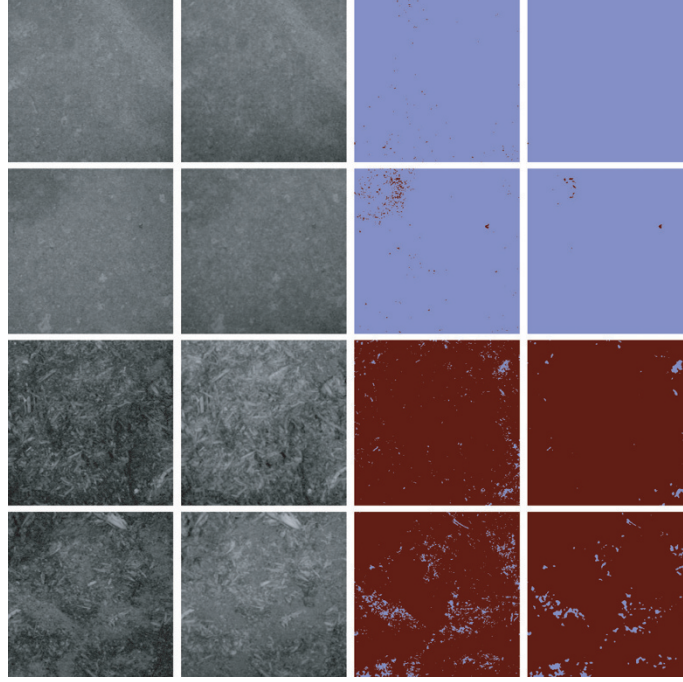


FIGURE 13: Results of image classifications for two clean and dirty concrete samples. The first two columns show the raw images captured with light with wavelengths 590 nm and 850 nm, respectively. The third column shows the pixel-by-pixel classification, where blue pixels have been classified as clean and red as dirty. The fourth column shows the result of applying a five-by-five median filter to the classified image. The first two rows correspond to clean 1 and clean 2, respectively, whereas the third and fourth rows correspond to dirty 1 and dirty 2, respectively.

7.3. Algorithm

The program operation can be summarised more formally as follows.

Given

- (1) A circular image mask with centre (m_x, m_y) and radius m_r .
- (2) A list of classes, $c_k = (\mu_k, \Sigma_k, p_k, \gamma_k)$, with mean vectors, variance-covariance matrices, prior probabilities and display colours, respectively.

Initialisation

- (1) Draw the scatter plot background showing class discrimination boundaries—calculate for each combination of class c_q and possible pixel measurement $x = (I_1 I_2)'$,

$$S_{qx} = \frac{P_q}{\sqrt{\det \Sigma_q}} e^{-(1/2)(x-\mu_q)' \Sigma_q^{-1} (x-\mu_q)}, \quad (22)$$

and choose colour γ_k for (I_1, I_2) such that

$$S_{kx} = \max_q S_{qx}. \quad (23)$$

- (2) Draw the class ellipses defined by μ_k and Σ_k using colour γ_k .

Repeat

- (1) Read new image from camera:

$$X = \begin{bmatrix} x_{11} & & \\ & \ddots & \\ & & x_{hw} \end{bmatrix}. \quad (24)$$

- (2) Update live image display with X .
- (3) Compute mean μ_m and variance-covariance Σ_m of the set of mask pixels given by

$$x_m = \{x_{ij} \mid (i - m_y)^2 + (j - m_x)^2 \leq m_r^2\}. \quad (25)$$

- (4) Update mask statistics in scatter display by drawing the ellipse defined by μ_m and Σ_m .
- (5) Update scatter plot. Compute

$$s_{\alpha\beta} = |\{x_{ij} = (I_1 I_2)' \mid I_1 = \alpha \wedge I_2 = \beta\}|. \quad (26)$$

$s_{\alpha\beta}$ is the number of pixels measuring $(\alpha\beta)'$. For $s_{\alpha\beta} > 0$, plot a scatter plot pixel (α, β) with intensity increasing with s .

- (6) Calculate, for each pixel x_{ij} and class c_q ,

$$S_{qx} = \frac{P_q}{\sqrt{\det \Sigma_q}} e^{-(1/2)(x-\mu_q)' \Sigma_q^{-1} (x-\mu_q)}. \quad (27)$$

TABLE 1: Classification results for test images, with and without a median filter.

Sample	Classification		With filter	
	Clean	Dirty	Clean	Dirty
Clean	0.996	0.004	0.999	0.001
Dirty	0.044	0.956	0.025	0.975

Update class display with

$$Y = \begin{bmatrix} y_{11} & & \\ & \ddots & \\ & & y_{hw} \end{bmatrix}, \quad (28)$$

choosing $y_{ij} = y_k$ such that

$$S_{kx} = \max_q S_{qx}. \quad (29)$$

The computational requirements of the algorithm are modest, and the scenario envisioned for the cleaning robot does not require real-time performance. The prototype described here can process multiple frames per second with unoptimised desktop computer hardware, which is far beyond that which is actually required for a cleaning robot.

A partially offline procedure would also be acceptable, since it is expected that the cleaning and inspection phases will be staggered. The practical problems associated with protecting the sensor during cleaning, and capturing a suitable image immediately after cleaning, mean that some time will elapse between a picture being taken and cleaning resuming.

8. CONCLUSIONS

Based on strong incentives to replace manual labour in the cleaning of pig houses by a fully autonomous robot system, this paper analysed the key factors in the design of a vision-based sensor system to classify surfaces as clean or dirty. Spectral properties of floor and walls in pig houses were characterised using spectrometer measurements on clean and dirty surfaces. The raw spectral data showed a fairly high variation in reflection from dirty surfaces and direct discrimination was found to be impossible for surfaces made of concrete using any single wavelength in the visible or near-infrared ranges, 400–1100 nm, which were covered by the experimental study. The key problem was shown to be variation in reflection caused by the differences in surface material, and contamination, as a function of position, whereas fluctuation over time caused by measurement noise was minor. Obtaining a low probability of misclassification for an observation of an element of the surface was therefore essential for the success of the concept. The paper employed the Jeffreys-Matusita distance measure and bounds on the misclassification probability to assess different scenarios for sensor design. An optimal choice of wavelengths for the illumination was calculated using actual field data and it was demonstrated that a probability below 2% was obtainable with just two different illumination wavelengths.

Design of a prototype algorithm to discriminate clean from dirty areas of a surface based on a Bayesian design of the multivariate classifier was demonstrated using a low-cost camera and standard computer interface. The results demonstrate the potential for designing a reliable and inexpensive vision system for autonomous pig house cleaning.

ACKNOWLEDGMENTS

The support of this research from the Sustainable Technology in Agriculture programme is gratefully acknowledged. The programme is funded jointly by the Danish Ministry of Science, Technology and Innovation and the Ministry of Food Agriculture and Fisheries under Grant number 2053-01-0021.

REFERENCES

- [1] J. S. Strøm, G.-Q. Zhang, P. Kai, M. Lyngbye, I. Braithwaite, and S. Arrøe, "Intelligent sensor for autonomous cleaning of pig houses—system in context," DJF Internal Rep. 182, Danish Institute of Agricultural Sciences, Tjele, Denmark, 2003.
- [2] H. B. Møller, S. G. Sommer, and B. K. Ahring, "Methane productivity of manure, straw and solid fractions of manure," *Biomass and Bioenergy*, vol. 26, no. 5, pp. 485–495, 2004.
- [3] T. Kailath, "The divergence and Bhattacharyya distance measures in signal selection," *IEEE Transaction on Communication Technology*, vol. 15, no. 1, pp. 52–60, 1967.
- [4] B. K. Ersbøll and K. Conradsen, *An Introduction to Statistics*, vol. 2, Informatics and Mathematical Modelling, Technical University of Denmark, Kongens Lyngby, Denmark, 2003.
- [5] K. Matusita, "A distance and related statistics in multivariate analysis," in *Multivariate Analysis*, P. R. Krishnaiah, Ed., pp. 187–200, Academic Press, New York, NY, USA, 1966.
- [6] B. K. Nielsen, *Transformations and classifications of remotely sensed data*, Ph.D. thesis, Informatics and Mathematical Modelling (IMM), Technical University of Denmark, Kongens Lyngby, Denmark, 1989.
- [7] J. Lin, "Divergence measures based on the Shannon entropy," *IEEE Trans. Inform. Theory*, vol. 37, no. 1, pp. 145–151, 1991.
- [8] H. Flesche, A. A. Nielsen, and R. Larsen, "Supervised mineral classification with semiautomatic training and validation set generation in scanning electron microscope energy dispersive spectroscopy image of thin sections," *Mathematical Geology*, vol. 32, no. 3, pp. 337–366, 2000.
- [9] M. E. Hansen and J. M. Carstensen, "Density based retrieval from high-similarity image databases," *Pattern Recognition*, vol. 37, no. 11, pp. 2155–2164, 2004.

Ian Braithwaite was born in London, England, in 1965. He received his B.A. (honors) degree in control engineering from the University of Cambridge in 1988 and, after moving to Denmark in 1993, received an M.S. degree from the Technical University of Denmark in 1999. Mr. Braithwaite is currently employed as a Research Assistant at the Technical University of Denmark, within the research project "Intelligent Sensor for Autonomous Cleaning." His other research interests include real-time systems and embedded software development.



Mogens Blanke is a Professor of automation at the Technical University of Denmark. He graduated from the Technical University of Denmark (DTU) in 1974, was with DTU in 1974–1975, at the European Space Agency in The Netherlands in 1975–1976, and again with DTU from 1976 till 1985. He received the Ph.D. degree from DTU in 1982. He was in the marine automation industry from 1985 to 1990 when he obtained the chair as a Professor in control engineering at Aalborg University, Denmark. He started his present position at DTU in 2000. His primary fields of interest are fault-tolerant control and diagnosis, modeling, and system identification. Application areas include autonomous robots, spacecraft, and surface ships. He has applied his research results for autonomous attitude control for the Danish Ørsted satellite, for autopilot, platform, and maneuvering control, for merchant and naval vessels, and for adaptive diesel engine control. He has published results in about 150 articles and conference papers and in a recent textbook. He is an active member of IFAC, the International Federation of Automatic Control, where he served as a Technical Committee Chair for several years, and was a coordinating Chair and Member of the IFAC Council. He is also a Senior Member of the IEEE.



Guo-Qiang Zhang received the B.S. degree in mechanical engineering from Jilin University of Technology, Changchun, China, in 1982, and the Ph.D. degree in agricultural engineering from the Royal Veterinary and Agricultural University, Copenhagen, Denmark, in 1989. He is currently a Senior Scientist at the Danish Institute of Agricultural Sciences and an Adjunct Professor in China Agricultural University. His research interests include indoor environmental and climatic control for livestock buildings, process modeling, and experimental fluid dynamics. Dr. Zhang has over 70 international publications in the areas of agricultural and biosystems engineering. In his research career, he also evolved in many research activities on automation, remote sensing, and information technology in agricultural and biosystems process. He is a Member of the Danish Society of Engineers, Danish Society of Agricultural Engineering, and European Agricultural Engineering.



Jens Michael Carstensen received the M.S. degree in engineering from the Technical University of Denmark, Lyngby, in 1988, and the Ph.D. degree in statistical image analysis from the Technical University of Denmark in 1992. He joined the Technical University of Denmark faculty in 1992 as an Assistant Research Professor and is currently an Associate Professor. Since 1999, he has been the Chief Technical Officer and founder of Videometer A/S, a vision technology R&D company. His current research interests include statistical methods for image analysis, Markov random fields, texture analysis, shape and appearance models, and industrial and biotechnological applications.

