Full Title

A Machine Vision System for Robust Sorting of Herring Fractions

Name(s) of Author(s)

Erik Guttormsen[a], Bendik Toldnes[b], Morten Bondø[b], Aleksander Eilertsen[b], Jan Tommy Gravdahl[a], John Reidar Mathiassen[b,*]

Author Affiliation(s)

[a] Norwegian University of Science and Technology (NTNU), Department of Engineering Cybernetics, 7491 Trondheim, Norway

[b] SINTEF Fisheries and Aquaculture AS, Brattørkaia 17C, 7010 Trondheim, Norway

Contact information for Corresponding Author

Full name:	John Reidar Mathiassen
Mailing address:	SINTEF Fisheries and Aquaculture AS
	Brattørkaia 17C
	7010 Trondheim
	Norway
Telephone:	(+47) 93453696
E-mail:	John.Reidar.Mathiassen@sintef.no

ABSTRACT:

Among the rest raw material in herring (*Clupea harengus*) fractions, produced during the filleting process of herring, there are high value products such as roe and milt. As of today there has been little or no major effort to process these by-products in an acceptable state, except for by manual separation and mostly mixed into low-value products. Even though pure roe and milt fractions can be sold for as much as ten times the value of the mixed fractions, the separation costs using manual techniques render this economically unsustainable. Automating this separation process could potentially give the pelagic fish industry better raw material utilization and a substantial additional income. In this paper, a robust classification approach is described which enables separation of these by-products based on their distinct reflectance features. The analysis is conducted using data from image recordings of by-products delivered by a herring processing factory. The image data is divided into three respective classes: roe, milt and waste (other). Classifier model tuning and analysis is done using multiclass support vector machines (SVMs). A grid search and cross-validation is applied to investigate the separation of the classes. Two-class separation was possible between milt/roe and roe/waste. However, separation of milt from waste proved to be the most difficult task, but it was shown that a grid search maximizing the precision – the true positive rate of the predictions – results in a precise SVM model that also has a high recall rate for milt versus waste.

Keywords:

machine vision, support vector machines, herring, sorting

Introduction

In 2014 a total of 162 000 tons of rest raw material was produced by herring fileting industries in Norway. This number is much lower than in the previous years, due to regulation of the herring quota. A relatively large part of landed herring is fileted in Norway (70 %). For the time being herring rest raw material is exclusively utilized as meal for the salmon feed industry and not for human consumption [16]. The greater part of the material is sold to other industries which process it into oil, flour or a product called silage. Unless the rest raw material is separated - the most valuable byproducts being milt and roe, with belly flap also being a valuable fraction - it is worth relatively little. If separated, the by-products can be worth ten times as much. However, the separation costs using manual techniques render this uneconomical. The potential for better utilization of these products is large, and in a survey conducted by Nofima AS [6], the potential utility value of milt from herring is described, in part due to the high nutritional value. Both milt and roe from herring contains the important fatty acids EPA, DHA, ARA and DPA along with large amounts of proteins (22-25 % and 24-26 % respectively), [15]. The gonads have a fat content of around 4-5 % of which 65-75 % consists of phospholipids – important lipids that are a major component of all cell membranes. In addition, the utility value of the belly flaps has been studied [13], along with the other filleting by-products (bits and pieces). New product and market possibilities were discovered, regarding the utility value of these products for human consumption.

In previous work, a system for sorting herring roe has been developed [11] They used a fuzzy classifier and 2D features to grade the roe as being either ‘good’ or ‘poor’. Due to the uncertainties in the classifier performance, Hu et al. [11] also propose a general grading framework that includes manual regrading of the fraction of roe that lies near the classification border between 'good' and 'poor'. Later versions of the roe grading system included color features, and 3D imaging using multiple laser stripes was added in order to enable automated weighing of the roe and detection of deformed (3D deformations) class of roe called 'henkei' [14]. This previous work is topically close to ours, even though it does not distinguish between herring fractions. There is machinery available for high speed sorting of other types of food such as nuts, fruit and vegetables. Examples of such systems are the Opus free-fall camera/laser sorting machine (TOMRA Systems ASA, Asker, Norway) – an optical food sorting solution for IQF (frozen) fruit and vegetable processors. Machines such as these do not directly solve our problem, but the techniques they use are worth considering, and are quite similar in some ways to that which we present.

The work presented draws some inspiration from research [17] demonstrating a significant difference in NIR absorbance in herring roe and milt. We also investigated several wavelengths in the visible and NIR regions in previous work [8] and found that a wavelength of 785 nm enabled the best distinction between milt and roe. This paper takes the previous research a large step closer to practical industrial application, by demonstrating a proof-of-concept machine vision system for robust sorting of herring fractions. The hope is that new research and sorting machine development, which the work in this paper is a part of, might help give birth to a whole new consumer market for herring products and enable a better raw material utilization. Implementation of a sorting machine of this kind might generate new income for the processing industry, and also has the potential of giving both the market and herring processing industry more flexibility and choices in terms of product assortment from herring fractions.

Materials and Methods

Herring and the filleting process.

In Norway, whole herring is filleted using filleting machines such as Baader 221 (Nordischer Maschinenbau Rud.Baader GmbH, Lübeck, Germany) that output fillets and other herring fractions. These herring fractions consist of heads, tails, belly flaps, back bones, skins, gonads (roe or milt) and other internal organs. An overview of the filleting process can be seen in Figure 1.

After the herring has been sorted according to size and distributed to the filleting machines, it is oriented head first and with the belly pointing downwards, and then the head and tail is cut. The fish is then brought to the first set of knives where the belly flap is cut and removed, thereby opening the abdominal cavity. The gonads, along with the rest of the intestines, are separated from the rest of the fish by a spinning wheel that scoops out the contents of the abdominal cavity. The content falls directly down through a vertical shaft and drops onto a conveyor belt. Most of the content that drops through this vertical shaft is either milt or roe, normally with just a minimal amount of intestines and other organs. The other content usually drops down at other locations, before and after the milt and roe. The work in this paper focuses on the herring fractions that fall down the vertical shaft where the milt and roe drops.

With high processing speeds of up to 5 fish per second, equivalent to 250-300 fish processed every minute, some will inevitably get stuck and some will avoid the filleting knives and pass intact through the entire machine, ending up among the rest-raw material falling down the vertical shaft where the milt and roe drops. This is something that needs to be taken into consideration when designing the machine vision system, in order to make it robust. It is imperative that unknown or unwanted waste material does not mix with the pure fractions of milt and roe that have been extracted. The different rest raw material fractions are shown in Figure 2.

For the image acquisition in this paper, the herring processor sent us four different herring fractions – milt, roe, belly flap and backbone. Belly flap and backbone are categorized as waste. The fractions were hand-sorted at Nergård Sild AS, vacuum packed fresh in bags and frozen, and then shipped in frozen state to our lab. The day before the image acquisition, the bags were taken out of the freezer and thawed in water at room temperature for 2 hours, before being put into a refrigerated room for thawing at 4˚ C over night. The herring fractions in thawed condition are shown in Figure 2.

Imaging system and image acquisition

The image acquisition system is illustrated in Figure 3, and the concept is based on imaging of herring fractions in free fall, as they drop down out of the filleting machine and onto a rest raw material moving conveyor. The camera is a NIR[1]-enhanced CMOS imager model MQ013RG-E2 (Ximea s.r.o., Slovakia) with an imaging resolution of 1280×1024 pixels. The camera images a

[1] NIR – Near infra-red

reduced-row region of interest as the rest raw material drops through a laser line sheet of light. The laser used is a Z80M18SF785LP30 (Z-LASER GmbH, Germany), emitting an 80 mW near infrared laser line with wavelength 785 nm and fan half-angle of 15 degrees. Imaging is done at a frame rate of 250 images per second at a bit depth of 8 bits per pixel. An angle of 15 degrees between the camera and the laser ensures that the laser line is outside the region of interest (ROI) unless it intersects with a herring fractions falling through the drop zone. This enables us to easily detect the presence or absence of herring fractions.

Laser line reflectance features

The laser line reflectance is different for milt and roe, as can be seen in the image in Figure 4. Since milt and roe are the fractions we are focused on sorting in this paper, the wavelength has been optimized for the purpose of distinguishing these two fractions. Milt has a higher peak reflectance, and less laser line scattering than roe.

Several laser line reflectance features are computed, in order to compactly describe the laser reflectance as it varies with the distance from the laser line. The image has m_{row} rows and m_{col} columns. Let x denote the column index and y denote the row index in the image acquired by the camera, and let $r(x,y)$ be the reflectance corresponding to the image intensity in column x on row y. Let $y_{peak}(x)$ be the row with peak reflectance in column x. Then for each image column x, the following laser line reflectance features are computed:

$$Reflectance(x) = \sum_{y=1}^{m_{row}} r(x,y),$$

$$Direct(x) = r\left(x, y_{peak}(x)\right),$$

$$Scatter(x, y_{offset}) = r(x, y_{peak}(x) + y_{offset}),$$

$$ScatterDirectRatio(x, y_{offset}) = \frac{Scatter(x, y_{offset})}{Direct(x) + 1}.$$

A scatter offset of $y_{offset} = 10$ pixels is selected for the work in this paper, as it was found to optimally separate milt and roe.

The laser line reflectance features in the above equations are essentially feature scan profiles along the x direction of the image. These scan profiles are computed for all the image frames, thereby accumulating feature scan profiles over time which are represented as feature images with x as one dimension and frame number as the other, hence providing a $Reflectance$ image, a $Direct$ image, a $Scatter$ image and a $ScatterDirectRatio$ image.

Feature vector

The image columns containing herring fractions are segmented from the background, based on $y_{peak}(x)$ being valid and within the ROI, since the absence of any falling herring fractions results in an image with no laser line within the ROI. Herring fraction features are computed for each segmented herring fraction, and for each laser line reflectance feature, by taking the mean of the feature image over the segmented area. In addition to the reflectance features, we also include the width (in pixels) and the height (in number of scans) of the herring fractions. Thus, for each segmented herring fraction we get the six-dimensional feature vector

$$\mathbf{x} = [Width \quad Height \quad Reflectance \quad Direct \quad Scatter \quad ScatterDirectRatio]^T.$$

Support vector machine classifier

Despite all the popularity as an industrial machine learning and classification technique, the support vector machine (SVM) has one major drawback – it is designed for two-class binary classification. Most SVM algorithms are built on the work of Cortes and Vapnik [5] developed for binary classification (two classes). Though new methods for multiclass SVMs have been proposed, many have the drawback of being computationally expensive. Although not directly related to SVMs, an early documented method where a multiclass classification problem is broken down to pairwise binary classifications is in Hastie and Tibshirani [18]. They suggest a one-vs.-one (OVO) scheme which involves estimating class probabilities for each pair of classes, and then coupling the estimates together. The OVO technique is also reviewed in Friedman [9], where Bayes optimal two-class decision rule is used.

For a general k-class decision problem, they train a series of $k(k-1)/2$ Bayes classifiers, each separating two of the classes. These boundaries are then used to assign an unknown sample to one of its two respective classes. A voting scheme then selects the class with the most winning two-class predictions as the final prediction for the sample. Although the method might be less sensitive to imbalanced dataset, it suffers from being computationally expensive as the number of classes increases. For a general k-class classification problem, the one-vs.-one method would need $k(k-1)/2$ separate binary classifiers.

In our paper we consider $k = 3$, with the classes *milt*, *roe*, and *waste*. Using the OVO scheme for multi-class SVM requires training of three binary SVM classifiers: 1) *milt* vs. *roe*, 2) *roe* vs. *waste*, and 3) *milt* vs. *waste*.

Assuming we have l samples, each sample indexed by i having a feature vector $\mathbf{x}_i$ and a binary class label $y_i \in \{-1, +1\}$, the support vector machine (SVM) [1,5] requires solving the following optimization problem:

$$\begin{aligned} \underset{\mathbf{w},b,\boldsymbol{\xi}}{\text{minimize}} \quad & \frac{1}{2}\mathbf{w}^T\mathbf{w} + C\sum_{i=1}^{l}\xi_i \\ \textit{subject to} \quad & y_i(\mathbf{w}^T\varphi(\mathbf{x}_i) + b) \geq 1 - \xi_i \\ & \xi_i \geq 0 \\ & i = 1, \dots, l. \end{aligned}$$

Given $\mathbf{w}$ and b, the discriminant function can be written as

$$\hat{y}(\mathbf{x}) = \mathbf{w}^T\varphi(\mathbf{x}) + b.$$

If the discriminant function is a positive value, the SVM classifies the sample as belonging to the positive (label +1) class, and similarly for a negative value. The mapping $\varphi(\mathbf{x}_i)$ is an implicit mapping that depends on the kernel $K(\mathbf{x}_i, \mathbf{x}_j) = \varphi(\mathbf{x}_i)^T\varphi(\mathbf{x}_j)$. For the linear SVM, the kernel is $K(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i{}^T\mathbf{x}_j$, and when using nonlinear SVM the radial basis function (RBF) kernel is $K(\mathbf{x}_i, \mathbf{x}_j) = e^{-\gamma\|\mathbf{x}_i - \mathbf{x}_j\|}$. In practice, the optimization problem is solved in its simpler dual form (Bottou and Lin 2007), since this ensures that the implicit mapping only occurs in the form of the kernel $K(\mathbf{x}_i, \mathbf{x}_j)$ in the optimization problem and the discriminant function. For the SVM implementation in this paper, we use the LIBSVM [4] library, and follow the usage recommendations outlined by its authors [12]. The recommended model selection technique is a grid-search on the RBF kernel hyper-parameters γ and C using cross-validation. Various pairs of these hyper-parameters are tested, and the pair returning the best cross-validation accuracy is selected. For *milt* vs. *waste* we also select the hyper-parameter pair with the best precision for *milt*.

For handling unbalanced classes and to adjust the relative importance of each class, we use the asymmetric soft margin penalty formulation as described by Ben-Hur and Weston [2], and where we use separate soft margins C_+ and C_- with a relative weighting of 1 for the positive class and w_- for the negative class.

Evaluating classifier performance

There are several methods for evaluating a binary classifier. Assuming one class is designated as the positive and the other class is designated the negative, we may illustrate the performance of a binary classifier by the four numbers in the confusion matrix in Figure 5. The numbers TP, FP, TN and FN are the number of samples belonging to each specific location in the confusion matrix. TP denotes the number of positive samples predicted to be positive, FP denotes the number of negative samples predicted to be positive, and similarly for TN and FN. With that notation we can define the following performance metrics for a classifier.

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

Another performance metric for measuring the performance of a binary classifier is the AUC – the area under the receiver operator curve [7], which is sometimes used as an alternative to accuracy.

Accuracy, AUC and precision are three metrics that will be used during the cross-validation and selection of kernel hyper-parameters γ and C using a grid search.

Results and Discussion

Image acquisition and feature extraction was done on herring fractions (n = 814). The different fraction types are milt, roe, belly flap and backbone, shown in Figure 2. The mean and standard deviations of these features, as well as the number of each herring fractions, are listed in Table 1. All features are in units output directly from the feature extractor algorithm, and depend on the image resolution, gain and other image acquisition parameters. The $Width$ and $Height$ features have substantial overlap. One can see that milt and roe are very well separated with respect to the features $Reflectance$, $Scatter$, $Direct$ and $ScatterDirectRatio$, and that roe is also well separated from belly flap and backbone in these features, and that milt has some overlap with belly flap and backbone.

The desired outcome of a sorting machine for herring fractions is to have pure milt and roe fractions – i.e. as close to 100 % precision as possible for these two fractions. The other fractions, such as belly flap and backbone, are to be categorized as waste. When maximizing the sorting precision for milt and roe, it is of less importance whether some of the milt and roe is classified as waste. In the case of e.g. a classifier where roe is the positive class and waste is the negative class, one may want to increase the precision with the consequence of a lowered recall. Unless the classifiers are perfect, there will be such a tradeoff between precision and recall.

Classifier performance is evaluated for each of the three possible one-vs-one classifiers. For each of the three classifiers, a 10-fold cross-validated grid search is done on 70 % of the samples, and the classifier performance is evaluated on a validation set consisting of the remaining 30 % of the samples. The classifier performance results are summarized in Table 2. Referring to this table, the kernel used is either a linear SVM kernel or a nonlinear SVM kernel of the radial basis function (RBF) type. The objective column describes the objective used in cross-validated grid search over the hyper-parameters. A further parameter $w_\_$ is also varied in order to adjust the classifier performance balance between the two classes. The classifier performance is measured by accuracy, precision and recall. The waste class consists of belly flap and backbone.

The classifiers were visualized in a normalized feature space consisting of three of the laser-based features. The features are normalized to the range between 0 and 1, as a preprocessing step for the SVM classifier. A linear SVM classifier was sufficient to perfectly distinguish between milt and roe, as can be seen in Figure 6. A nonlinear SVM classifier, of the RBF type, can perfectly distinguish between roe and waste, as seen in Figure 7. A linear classifier also worked in this case, but had a tight maximal margin [10]. A nonlinear classifier for milt vs. waste is shown in Figure 8, and there is some overlap between the classes.

Several grid search objectives and negative class weights (w_{-}) were tested, with the goal of getting as close to 100 % precision for milt vs. waste. Referring to Table 2, we see that the use of accuracy, as the grid search objective, does not enable perfect precision. The use of AUC increases the precision up to 98.7 % at a recall of 93.1 %. Using precision as the grid search objective enables a 100% precision, at a recall rate of 77%.

In summary, the analysis showed that the milt was perfectly separable from roe, and roe was perfectly separable from waste. Separation of milt and waste on the other hand proved difficult, and the accuracy depended highly on the grid search objective and negative class weight (w_{-}). When the objective of the grid search was to maximize precision, perfect precision was possible at a relatively high recall rate.

Based on the positive results from the work in this paper, the natural next steps are to implement the machine vision system in an industrial setting. The herring fractions used in this paper were shipped in frozen condition, and are not in the same state as when they exit the filleting machine. Also, the work in this paper focuses on four types of herring fractions. Preliminary work [10] suggests that the machine vision system and classifiers may be applicable for other herring fraction types. As future work, it is suggested to perform image acquisition at the rest raw material exit points of a filleting machine, in order to obtain as fresh and as varied herring fractions as possible.

Conclusion

The rest raw material in herring fractions can be accurately sorted by using machine vision in combination with a robust classification approach. Illuminating the herring fractions with a single laser line at 785 nm enables the extraction of laser direct and indirect reflectance features that sufficiently distinguish between roe, milt and waste. A support vector machine classifier, with a radial basis function kernel, is trained on these reflectance features and the classifier hyper-parameters are selected through a grid search that maximizes classification accuracy and precision. Distinguishing between roe and milt, and roe and waste, has 100 % classification accuracy. When distinguishing between milt and waste, milt can be classified with 100 % precision, at a recall rate of 77 %.

Acknowledgments

The work in this paper was financed by the Norwegian Research Council through project grant #219204. We thank the herring processing plant Nergård Sild for providing us with vacuum packed herring fractions that were used in the experiments in this paper. We thank Henning Grande and Halgeir Hansen, Nergård Sild AS, for being the industry contacts for the project of which this paper is a part. We thank Cecilie Salomonsen for making the 3D illustration in Figure 3.

References

1. Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In Proceedings of the fifth annual workshop on Computational learning theory. ACM. p 144-152. (1992

2. Ben-Hur, A., Weston, J.: A user's guide to support vector machines. In: Data mining techniques for the life sciences. Humana Press. p 223-239. (2010)

3. Bottou, L., Lin, C.J.: Support vector machine solvers. Large scale kernel machines. p 301-320 (2007)

4. Chang, C.C., Lin, C.J.: LIBSVM: A library for support vector machines. ACM Transactions on Intelligent Systems and Technology 2(3). P 27:1-27. Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm. (2010)

5. Cortes, C., Vapnik, V.: Support-vector networks. Machine learning 20(3). p 273-297. (1995)

6. Egede-Nissen, H., Vogt, K.G., Haugen, J-E., Høstmark, Ø., Oterhals, Å.: Utvikling av høykvalitets pulverprodukt fra sildemelke. Sensorisk kvalitet på sildemelkepulver testet ved akselererte lagringsbetingelser - Fagrapport 2. Nofima Report 14/2013. ISBN: 978-82-8296-144-8. (2013)

7. Fawcett, T.: An introduction to ROC analysis, Pattern Recognition Letters 27. p 861–874. (2006)

8. Fossum, J.A., Mathiassen, J.R., Toldnes, B., Salomonsen, C.: Teknologi for fraksjonert uttak og sortering av restråstoff fra sild – Fase 1, SINTEF Report A23065. ISBN: 978-82-14-05437-8. (2012)

9. Friedman, J.H.: Another approach to polychotomous classification, Technical report, Department of Statistics, Stanford University. URL: http://www-stat.stanford.edu/ jhf/ftp/poly.ps.Z (1996)

10. Guttormsen, E.: Robust classification approaches to industrial sorting of herring fractions. Masters Thesis, NTNU. (2015)

11. Hu, B-G., Gosine, R.G., Cao, L.X., de Silva, C.W.: Application of a fuzzy classification technique in computer grading of fish products. IEEE Transactions on Fuzzy Systems 6(1), p 144-152. (1998)

12. Hsu, C.W., Chang, C.C., Lin, C.J.: A practical guide to support vector classification. Online: https://www.csie.ntu.edu.tw/~cjlin/papers/guide/guide.pdf. (2010)

13. Kjerstad, M., Larssen, W.E., Nystrand, B.T.: Produkt- og markedsutvikling for restråstoff fra NVG-sild til konsum. Møreforsking Report MA 14-18. ISSN: 0804-54380. (2014)

14. Lee, M.F.R., de Silva, C.W., Croft, E.A., Wu, Q.J.: Machine vision system for curved surface inspection. Machine Vision and Applications 12(4). p 177-188. (2000)

15. Østvik, S.O., Grimsmo, L., Jansson, S., Dauksas, E., Bondø, M.: Biråstoff fra filetering av sild - Kartlegging og analyse av råstoff og utnyttelsesmuligheter. Rapport nr. 164, RUBIN. (2009)

16. Richardsen, R., Nystøl, R., Strandheim, G., Viken, A.: Analyse marint restråstoff. SINTEF report A26863, ISBN 978-82-14-05877-2. (2014)

17. Wold, J.P.: Individbasert kvalitetssortering og kvalitetsmerking av pelagisk fisk: Automatisk sortering basert på indre kvalitetsparametre. Nofima Report 35/2013. ISBN: 978-82-8296-112-7. (2013)

18. Hastie, T., Tibshirani, R.: Classification by pairwise coupling. In Advances in Neural Information Processing Systems, volume 10. MIT Press. (1998)

Table 1 – Number of samples, and the mean and standard deviations of the feature values for each herring fraction used in the classification experiments. SDR is short for the feature ScatterDirectRatio.

	n	Width	Height	Reflectance	Direct	Scatter	SDR
Milt	288	9.03±4.24	6.67±3.02	2178.7±467.0	185.7±31.1	19.8±9.6	0.20±0.09
Roe	236	8.89±3.64	6.77±3.50	364.0±161.6	20.1±7.9	7.8±4.8	0.49±0.17
Belly flap	201	13.7±5.88	6.16±2.63	1923.0±439.8	176.6±27.8	25.5±14.7	0.20±0.12
Backbone	89	16.5±5.83	7.53±2.14	1674.4±487.7	130.5±27.0	29.0±13.1	0.25±0.09

Table 2 - Evaluation of classifier performance, with classification accuracy, precision and recall measured on validation sets.

Pos.	Neg.	Kernel	Objective	w_-	Accuracy (%)	Precision (%)	Recall (%)
Milt	Roe	Linear	Accuracy	1	100.0	100.0	100.0
Roe	Waste	RBF	Accuracy	1	100.0	100.0	100.0
Milt	Waste	RBF	Accuracy	1	93.4	92.9	93.8
				2	94.5	97.5	91.8
				4	91.3	92.8	89.5
Milt	Waste	RBF	AUC[2]	1	92.6	91.9	93.5
				2	93.6	97.5	89.5
				4	93.1	98.7	93.1
				8	90.8	98.6	82.6
				16	87.3	98.5	75.6
Milt	Waste	RBF	Precision	1	93.2	93.8	92.4
				2	90.8	98.0	83.5
				4	88.5	100.0	77.0
				8	85.0	100.0	70.0
				16	82.5	100.0	65.0

[2] Area under the receiver operator curve (ROC).

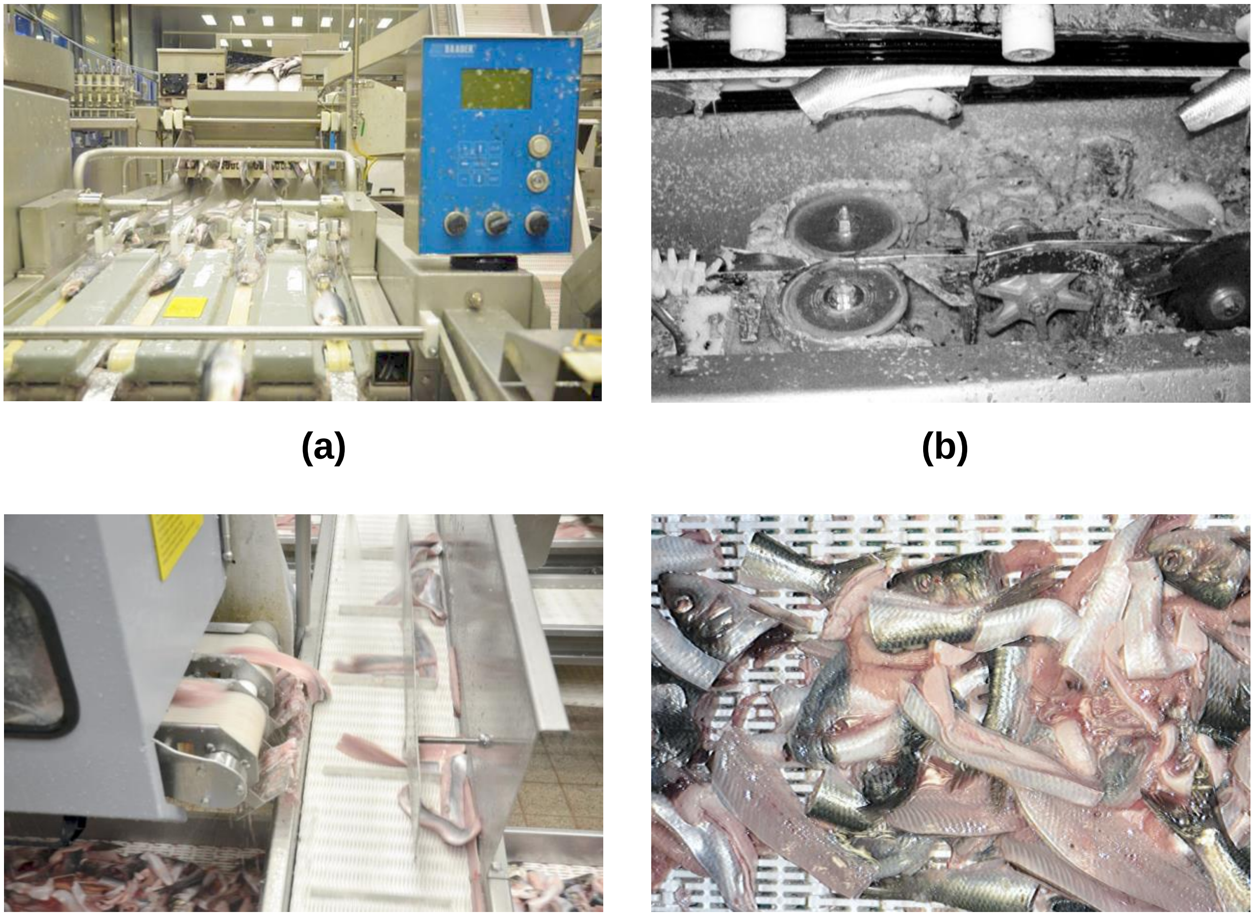

(a) (b) (c) (d)

Figure 1 - Overview of the filleting process, showing the singulation and orienting of herring (a), internal components of the filleting machine (b), fillets (c) and rest raw material herring fractions (d) exiting the filleting machine in separate streams.

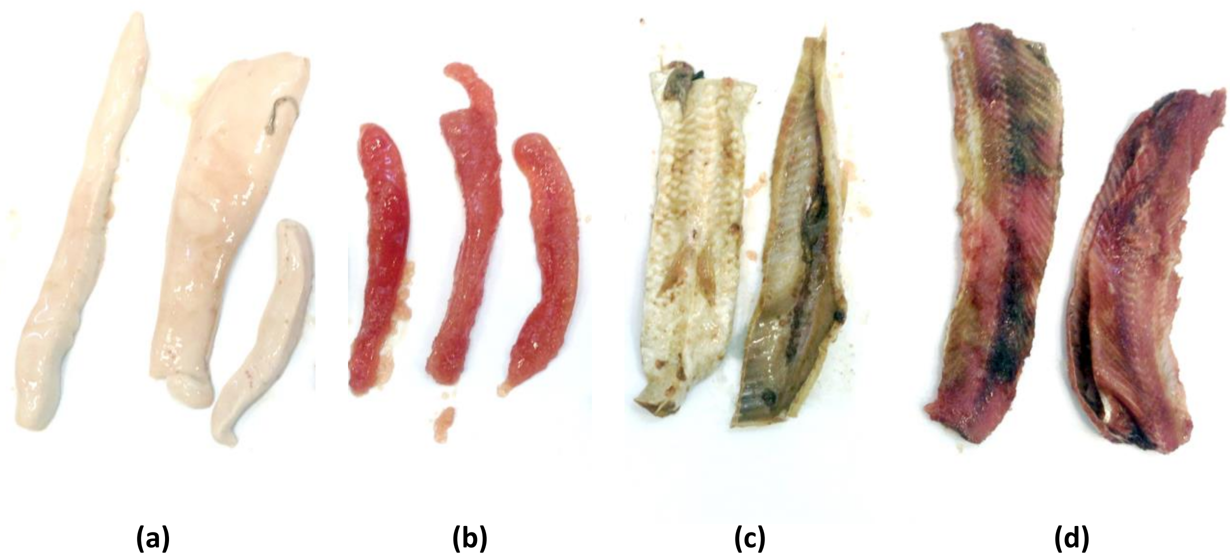

Figure 2 - Herring fractions considered in this paper. Milt (a), roe (b), belly flap outside (c, left) and inside (c, right), and backbone (d).

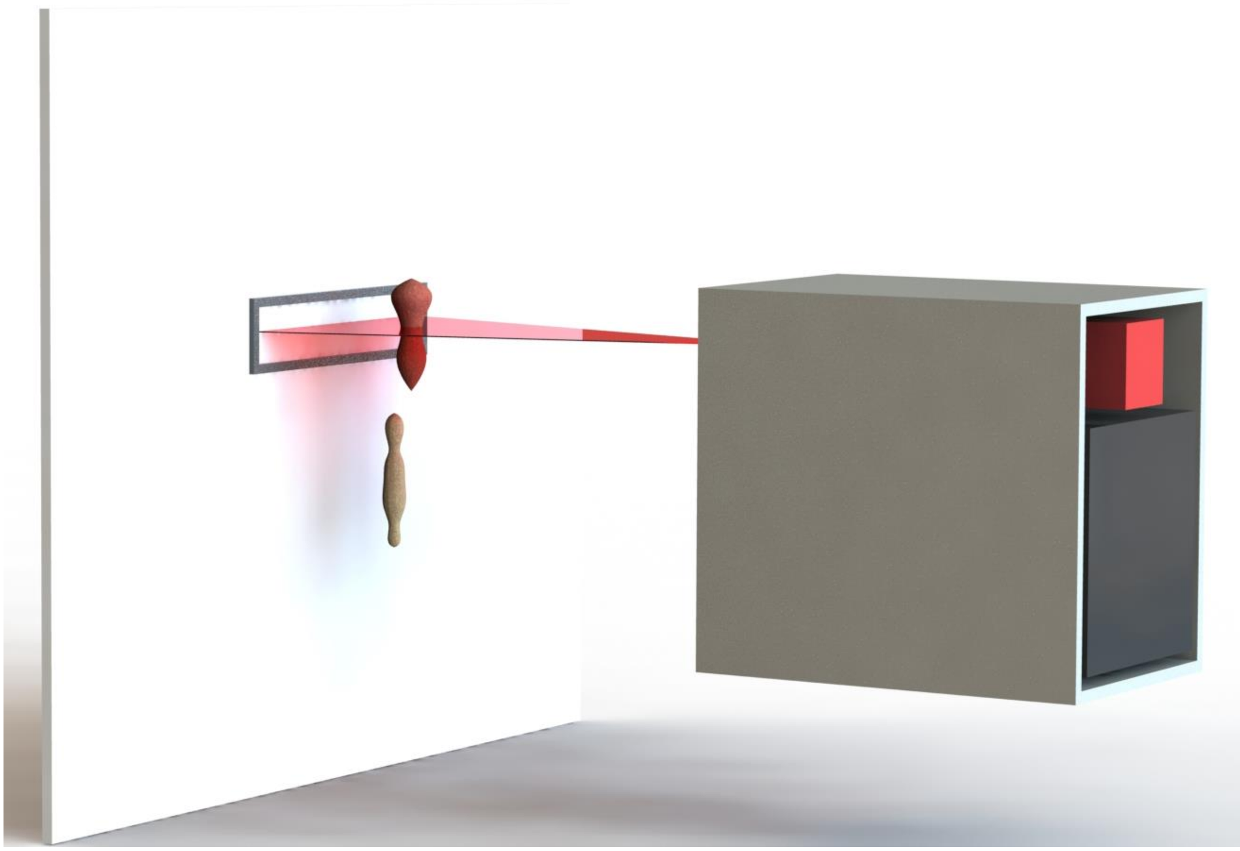

Figure 3 - Illustration of the imaging setup and the principle of dropping the fractions through a laser beam, and imaging a local region of interest.

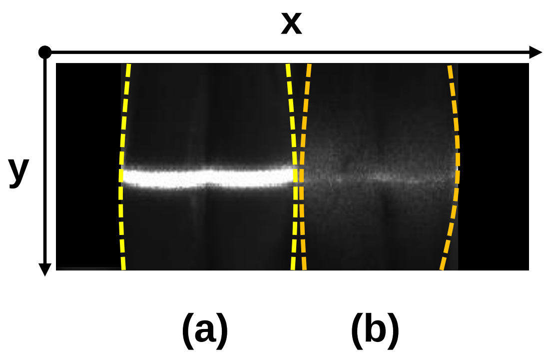

Figure 4 – Image of a laser line (785 nm) illuminating a milt (a) and a roe (b), with indicated image x and y axes.

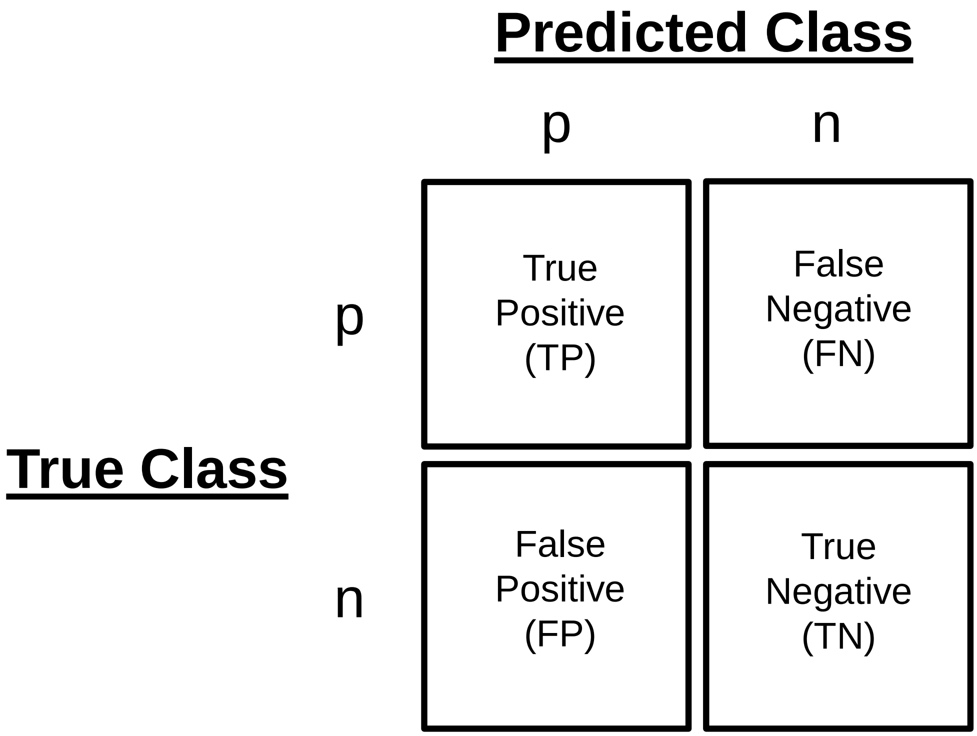

Figure 5 - The confusion matrix for a binary classifier.

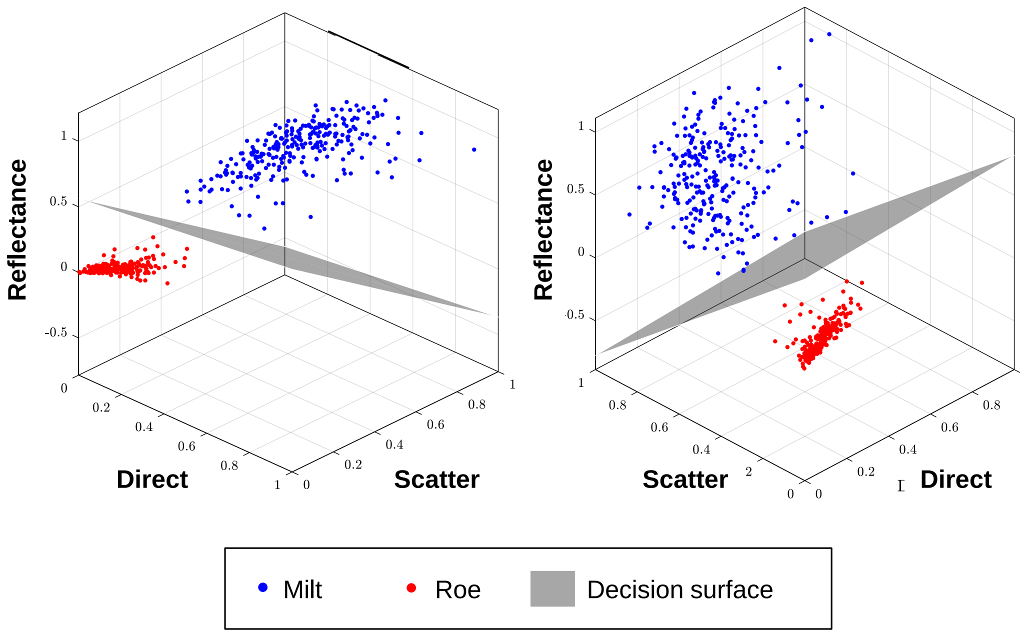

Figure 6 - Linear SVM classifier for milt vs. roe, plotted into the normalized feature space spanned by three of the features.

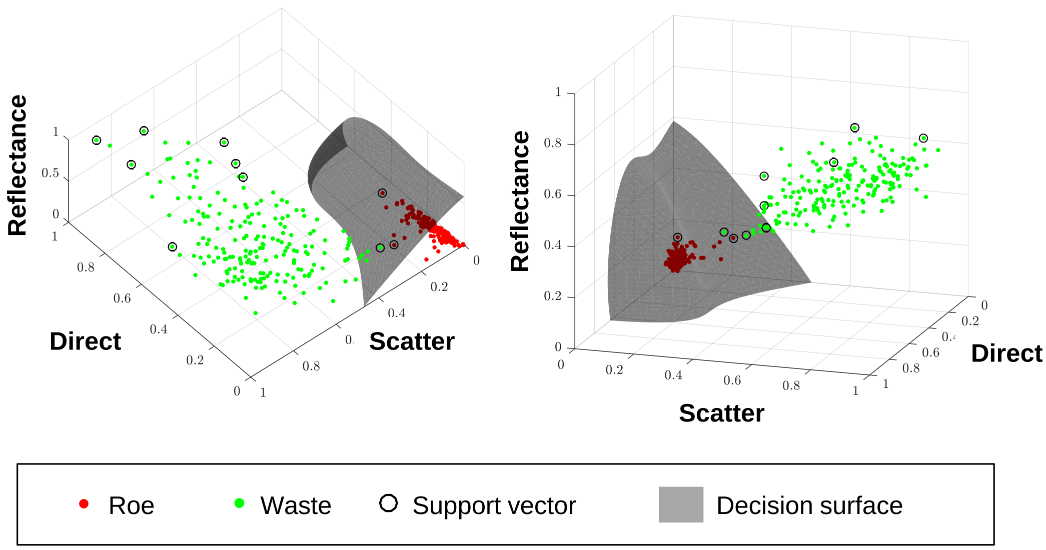

Figure 7 - Radial basis function SVM classifier for roe vs. waste, plotted into the normalized feature space spanned by three of the features.

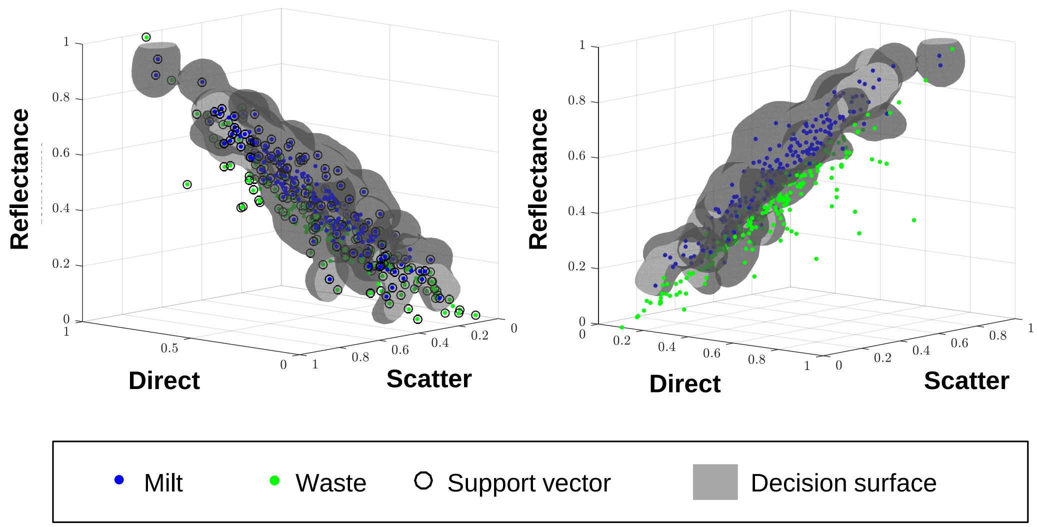

Figure 8- Radial basis function SVM classifier for milt vs. waste, plotted into the normalized feature space spanned by three of the features.