

available at www.sciencedirect.comjournal homepage: www.ejconline.com

Detection of colorectal cancer using MALDI-TOF serum protein profiling

Mirre E. de Noo^{a,*}, Bart J.A. Mertens^b, Aliye Özalp^c, Marco R. Bladergroen^c,
Martijn P.J. van der Werff^c, Cornelis J.H. van de Velde^a, Andre M. Deelder^c,
Rob A.E.M. Tollenaar^a

^aDepartment of Surgery, K6-R, Leiden University Medical Center, Albinusdreef 2, P.O. Box 9600, 2300 RC Leiden, The Netherlands

^bDepartment of Medical Statistics and Bioinformatics, Leiden University Medical Center, Leiden, The Netherlands

^cBiomolecular Mass Spectrometry Unit, Department of Parasitology, Leiden University Medical Center, Leiden, The Netherlands

ARTICLE INFO

Article history:

Received 25 November 2005

Accepted 19 December 2005

Available online 17 April 2006

Keywords:

Protein profiling

Colon cancer

Serum biomarkers

MALDI

Bioinformatics

Clinical proteomics

Diagnosis

ABSTRACT

Serum protein profiling is a promising approach for classification of cancer versus non-cancer samples. The objective of our study was to assess the feasibility of mass spectrometry based protein profiling for the discrimination of colorectal cancer (CRC) patients from healthy individuals. In a randomized block design, pre-operative serum samples obtained from 66 colorectal cancer patients and 50 controls were used to generate MALDI-TOF protein profiles. After pre-processing of the spectra, linear discriminant analysis with double cross-validation was used to classify protein profiles. A total recognition rate (92.6%), sensitivity (95.2%) and specificity (90.0%) for the detection of CRC were shown. The area under the curve of the classifier was 97.3%, and demonstrated the high, significant separation power of the classifier. Double cross-validation shows that classification can be attributed to information in the protein profile. Although preliminary, the high sensitivity and specificity indicate the potential usefulness of serum protein profiles for the detection of colorectal cancer.

© 2006 Elsevier Ltd. All rights reserved.

1. Introduction

Colorectal cancer (CRC) is among the most common malignancies and remains a leading cause of cancer-related morbidity and mortality. It is well recognized that CRC arises from a multistep sequence of genetic alterations that result in the transformation of normal mucosa to a precursor adenoma and ultimately to carcinoma. Given the natural history of CRC, early diagnosis appears to be the most appropriate tool to reduce disease-related mortality.^{1,2} Currently, there is no early diagnostic test with high sensitivity, specificity and positive predictive value, which can be used as a routine screening tool. Therefore, there is a need for new biomarkers

for colorectal cancer that can improve early diagnosis, monitoring of disease progression and therapeutic response and detect disease recurrence. Furthermore, these markers may give indications for targets for novel therapeutic strategies.

Proteomic expression profiles generated with mass spectrometry have been suggested as potential tools for the early diagnosis of cancer and other diseases. Different protein profiles may be associated with varying responses to therapeutics. It has been postulated that on the basis of the presence/absence of multiple low-molecular-weight serum proteins using time-of-flight (TOF) mass spectrometry technologies, such as SELDI-TOF and MALDI-TOF, biomarkers can be identified.^{3–6} Although the data from these studies

* Corresponding author: Tel.: +31 71 5261278; fax: +31 71 5266750.

E-mail address: m.e.de_noo@lumc.nl (M.E. de Noo).

0959-8049/\$ - see front matter © 2006 Elsevier Ltd. All rights reserved.

doi:10.1016/j.ejca.2005.12.023

are encouraging, critical notes have been made on both study design and experimental procedures for proteomic profiling.^{7–9} In addition, the importance of avoiding confounding biological variables, as well as technological factors that may bias the results, have previously been stressed by several authors.^{10,11} Another recurrent topic for debate is the use of independent validation sets for the classification of diseased versus healthy individuals. A specific problem in the discovery-based research field of clinical proteomics is overfitting. Overfitting may occur in the analysis of large datasets when multivariate models show apparent discrimination that is actually caused by data over-interpretation, and hence give rise to results that are not reproducible.^{9,12,13} The chance of overfitting however can be reduced by appropriate application of validity estimation and assessment, such as through application of double cross-validation, when properly implemented.

The objective of this study was to assess the feasibility of mass spectrometry based protein profiling for the discrimination of colorectal cancer patients from healthy individuals. In addition to standardizing technical factors and biological variations, we performed blinded tests and employed a randomized block design experimentation to minimize impact of potential confounding factors and to avoid bias. To minimize danger of overfitting, among other reasons, we used a fairly inflexible classification method based on first-and-second-order statistics only. Specifically, Fisher linear discriminant analysis was employed with double cross-validated integrated estimation and validation of error rate on the entire dataset to calculate an unbiased error rate assessment.

2. Patients and methods

2.1. Patients

Serum samples were obtained from a total of 66 colorectal cancer patients one day before surgery. All patients with stage IV disease had synchronous metastatic disease confined to the liver. Colorectal cancer was histologically confirmed on surgical specimens and preoperatively assessed with abdominal CT scan and carcinoembryonic antigen (CEA) levels. The extent of tumour spread was assessed by TNM classification based on histological examination of the resected specimen. All stages of colorectal cancer were represented in the patient group. The median age of the patient group was 62.8 years (range 32.6–90.3) and the male to female ratio was 31/35. Patients were included from October 2002 till December 2004 in our centre. The control group consisted of 50 healthy volunteers. The median age of the healthy symptom-free control group was 49.7 years (range 25.9–76.6) and the male to female ratio was 21/29. The controls were included from October till December 2004 (Table 1).

2.2. Study design

Having identified plate-to-plate and day-to-day variation as important potential batch effects, we used a randomized blocked design.^{14,15} All the available 116 samples from both groups (controls and colorectal cancer) were randomly distributed across three plates in roughly equal proportions

(Table 2). For colorectal cancer, the distribution of stadia across plates was again in random fashion and in approximately equal proportions (Table 3). The position on the plates of samples allocated to each plate was randomized as well. Each plate was then assigned to a distinct day, which completes the design. Analysis was carried out on 3 consecutive days, Tuesday to Thursday, processing a single plate each day. A duplicate of this randomized blocked study was performed in the following week.

2.3. Serum samples

Informed consent was obtained from all patients and the Medical Ethical Committee approved the study. All blood samples were drawn while the patients or healthy controls

Table 1 – Patient characteristics

	CRC patients		Controls
	Inclusion	Results	
n	66	63	50
Age (mean)	62.6	62.2	49.7
Age (range)	(32.6–90.3)	(32.6–90.3)	(25.9–76.6)
Male/female ratio	34/32	31/32	21/29

Table 2 – Distribution and randomization of serum samples of colorectal cancer patients with different TNM stage before and after the MALDI-TOF experiment

	Plate 1	Plate 2	Plate 3	Total
Colorectal cancer	22	22	19	63
Controls	17	17	16	50
Total	39	39	35	113

The distribution of stadia across plates was performed randomly and in approximately equal proportions.

Table 3 – Distribution and randomization of serum samples of different groups over the three MS target plates

		TNM stage	Plate 1	Plate 2	Plate 3
Inclusion	I		4	4	3
	II		10	10	8
	III		4	4	4
	IV		4	4	4
	0		4	3	3
	Total		26	25	22
Exclusion	I		0	0	1
	II		0	0	1
	III		0	0	1
	IV		0	0	0
	0		4	3	3
	Total		4	3	6

were seated and non-fasting. The samples were collected in a 10 cc Serum Separator Vacutainer Tube and centrifuged 30 min later at 3000 rpm for 10 min. The serum samples were distributed into 1 ml aliquots and stored at -70°C until the experiment.¹⁶

2.4. Isolation of peptides

The isolation of peptides from serum was performed using magnetic beads based hydrophobic interaction chromatography (MB-HIC) kit from Bruker, mainly according to the manufacturers instructions, and adapted for automation on an 8-channel Hamilton STAR[®] pipetting robot (Hamilton, Martinsried, Germany). Magnetic beads with C8-functionality (MB-HIC8) were divided in 5 μl -aliquots in a 96-well microtiter plate, which was placed on the magnetic beads separation device (MPC[®]-auto96, Dynal, Oslo, Norway), with the magnet down. MB-HIC binding solution (10 μl) and 5 μl serum sample were added to the beads and carefully mixed using the mixing feature of the robot. The sample was incubated for 30 s and the magnet was lifted, followed by a 30 s waiting interval to settle the magnetic beads. The supernatant was removed and the magnet was lowered again. The magnetic beads were washed three times with MB-HIC washing solution (also provided with the kit) lifting and lowering the magnet as needed. The peptides were eluted from the beads using 10 μl 50% acetonitrile and 2 μl of this eluate was transferred to a fresh 348-well microtiter plate (Greiner). Most of the remaining eluate (6 μl) was transferred to an auto sampler vial containing 54 μl water and stored for later use. Fifteen microliters of α -cyano-4-hydroxycinnamic acid (0.3 g/l in ethanol:acetone 2:1) was added to the 1 μl eluate in the 348-well microtiter plate and mixed carefully. One microlitre of this mixture was spotted in quadruplicate on a MALDI AnchorChip[™] (Bruker Daltonics, Bremen, Germany).

2.5. Protein profiling

Matrix assisted laser desorption ionisation time-of-flight (MALDI-TOF) mass spectrometry measurements were performed using an Ultraflex TOF/TOF instrument (Bruker Daltonics, Bremen, Germany) equipped with a SCOUT ion source, operating in linear mode. Ions formed with a N2 pulse laser beam (337 nm) were accelerated to 25 kV. With this specific serum preparation peptide/protein peaks in the m/z range of 960–11,169 Da were measured. An independent mass spectrometer operator performed the experiments at 3 consecutive days after cleaning of the instrument. One week later, the experiment was duplicated in exactly same order. Hereafter, the entire process of capturing and concentrating serum proteins using C8 magnetic beads including the generation of readouts of the MALDI-TOF spectra will be designated as the protein profiling procedure.

2.6. Data processing

All unprocessed spectra were exported from the Ultraflex in standard 8-bit binary ASCII format. They consisted of

approximately 45,000 mass-to-charge ratio (m/z) values, covering a domain of 1160–11,600 Da. To increase robustness, the average of four spots was used to represent one serum sample. Subsequently, we lightly smoothed the spectra using the Whittaker¹⁷ smoother. Due to the quadratic nature of the TOF-equation, the high-resolution spectra were binned using a linear scaling at the time scale, resulting in bin widths of approximately 1 Da at the beginning of the spectrum and 3 Da at the end at the mass/charge scale. The resulting spectra generally showed strong baseline effects. These were removed using an asymmetric least squares algorithm. To normalize the spectra, we calculated the median intensity of every spectrum and subtracted it from the original spectrum. Each of the thus normalized spectra was then also divided by the interquartile range of intensity within that spectrum. We consider this more robust than normalization of the spectra on the average, as it is less sensitive to the most extreme intensities. Finally, prior to classification and evaluation of error rate, the logarithm was taken of all intensity measurements (predominantly to ensure numerical stability of computations).

2.7. Statistical data-analysis

Fully validated classification error rates were estimated based on a classical Fisher linear discriminant analysis through complete double cross-validatory joint estimation and assessment of class predictions, as is further explained in [Appendix 1](#).^{18–20} Instead of ordinary leave-one-out cross-validatory choice of k , we employed double cross-validation. This was an extension of leave-one-out cross-validation which combined validatory “choice of model” (the parameter k in this case) with “predictive assessment” (of the same model, through use of error rate or other suitable summary statistic). The reason for this additional “technical complication” was that we did not wish to incur the bias inherent in the assessment, which would normally result from a model choice based on ordinary leave-one-out validation only. In a double cross-validatory evaluation, we removed each individual in turn from the data (just as in ordinary leave-one-out cross-validation), after which the discriminant rule was fully recalibrated and optimised for prediction on the leftover data (now of size $n = 1$, where n is the total initial sample size) and used the same procedure in each case. The choice of the calibration rule (i.e. choice of k in this case) to classify the left-out observation was then again based on a leave-one-out cross-validatory estimation (hence the name “double-cross”) within the leftover set of size $n = 1$. The resulting classification rule was then applied to the left-out datum to obtain an unbiased allocation for this sample. This procedure was then repeated across all individuals and for each person separately, after which we calculated a truly unbiased estimate of the misclassification rates on the basis of the thus validated (and calibrated) classifications. In other words, “double-cross” is actually “leave-one-out cross-validation within leave-one-out cross-validation” and it is precisely because of this that we can avoid bias in error rate estimation that an ordinary application of standard leave-one-out choice would imply.

3. Results

In the first week, three different randomized target plates were successfully measured on 3 consecutive days in the middle of the week. A duplicate experiment was performed in the second week on the same days. Fig. 1 shows a raw data spectrum, directly obtained from the MALDI-TOF mass spectrometer. Before pre-processing and further analysis a mean spectrum of each sample was calculated over all four spots that were measured for each sample. In case all four spots from one sample showed spectra of poor quality due to a technical problem, the sample was left out of the analysis. This was the case for three CRC patients' samples. The above-described pre-processing steps resulted in a sequence

of 4483 normalized m/z values ranging from 1160 to 11,600 Da, for each individual.

3.1. Detection of colorectal cancer

Double cross-validatory analysis and evaluation carried out on the protein spectra measured in week 1, correctly classified 45 of the 50 controls as not cancer. Sixty of the 63 cancer samples were correctly classified as malignant, including 9 of 10 TNM stage I patients (Table 4). The remaining two misclassified patients had stage II disease. All patients with stage III and IV disease were correctly recognised as malignant within the double cross-validatory evaluation. These validated results thus yielded a total recognition rate of

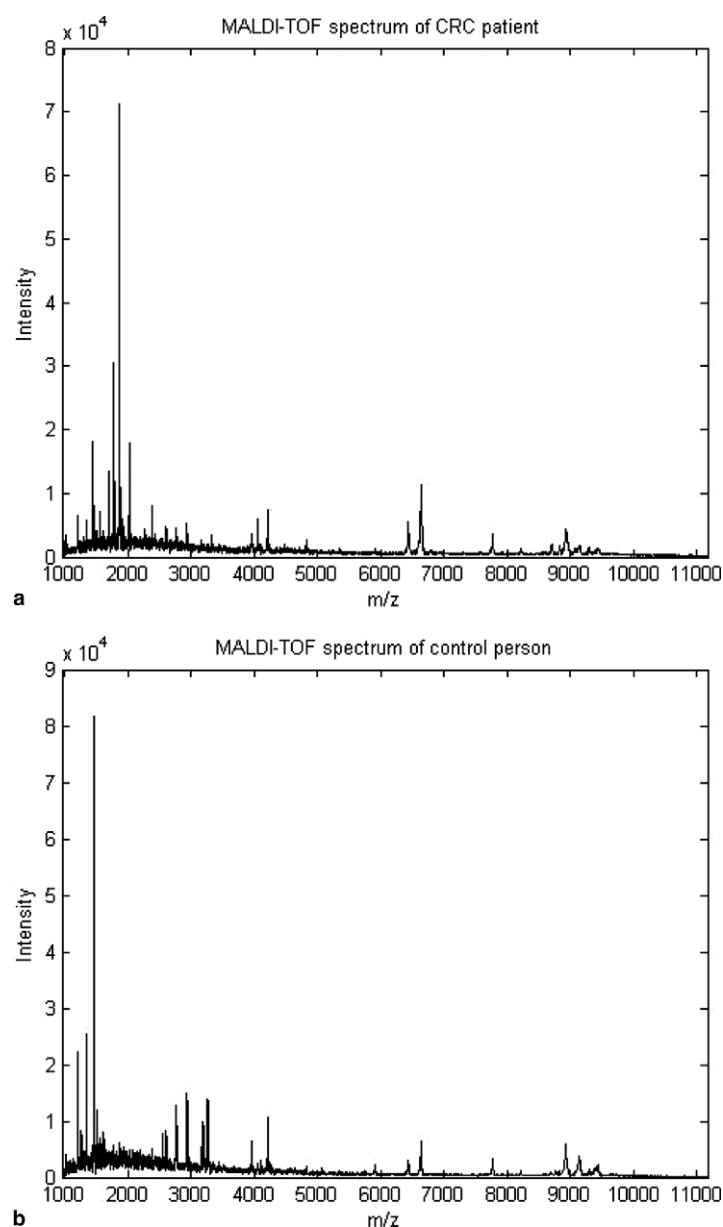


Fig. 1 – (a) MALDI-TOF spectrum of a colorectal cancer patient and (b) a healthy subject after peptide isolation with C8 magnetic beads. On the Y-axis the relative intensity is shown. The mass to charge ration (m/z) is demonstrated on the X-axis in Dalton.

Table 4 – Double cross-validated classification of serum samples

	Test results for detection of CRC		
	Positive	Negative	Total
Controls	45	5	50
CRC patients	3	60	63
	48	65	113

A positive test result assigns subjects to the CRC group and a negative to the controls. In the horizontal plane the actual histologically confirmed diagnosis is stated.

92.6%, a sensitivity of 95.2% and a specificity of 90.0% for the detection of CRC (Table 5). To analyze the actual discriminative power of the classifier, we produced an ROC-curve (again based on the double cross-validated classification probabilities), visualizing the performance of the two-class classifier in Fig. 2. The (double c.v.) AUC of the classifier was 97.6.

We repeated the entire double cross-validated evaluation executed with the week 1 data using the duplicate measured spectra from week 2. This procedure was identical to that carried out in week 1 and used the same calibration spectra. However, prior to classifying each left-out datum in the outer “shell” of the double cross-validated procedure, we substituted the week 1 data with the corresponding measured spectra from the same sample in week 2. In this manner, we could calculate a double cross-validated error rate, which takes the effect of replicate measurement of the spectrum (and thus also recalibration of the equipment) into account. The effect of classifying the remaining replicate data was that the recognition rate dropped to 88.8%. The sensitivity and specificity for the detection of CRC for the second week data was 80.6% and 97.1%, respectively (Table 4). The associated AUC of this repeat double cross-validated estimation on week 2 was 96.8%.

It is of interest to evaluate bias of the double cross-validated calculations. Hence, we performed a permutation exercise, which randomly permutes and reassigns the class labels across subjects and then repeats the entire double cross-validation procedure. Carrying out this procedure more than 600 times resulted in a median recognition rate of 50.0% (95% confidence interval is [36.3, 72.7]). The median AUC was 49.4% with confidence interval of [24.8, 64.2]. As both median recognition rates and AUC's equalled 50%, there was no substantial evidence of bias remaining within the cross-validated calculation.

Having executed the above-described validity evaluation, we explored the nature of the classification through a post hoc analysis. We found that the first two principal com-

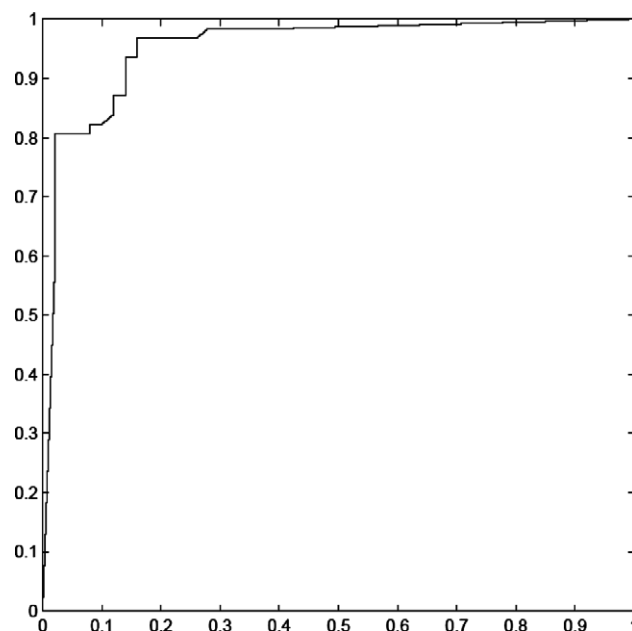


Fig. 2 – ROC-curve for the double cross-validated two-group classifier. The true positive recognition rate (sensitivity) is demonstrated on the y-axis against the false negative recognition rate (1-specificity) on the x-axis of the classifier.

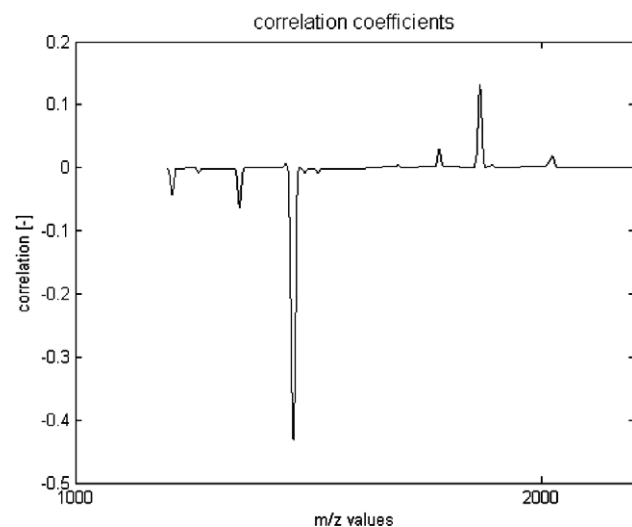


Fig. 3 – Correlation coefficients of two first principal components with the class indicator. The correlation coefficients were calculated from the linear discriminant weightings. The negative correlation of the first peak is an indicator for the control group and the positive correlation of the second peak points out the cases.

Table 5 – Cross-validated classification results for the detection of CRC

Method	First week				Second week			
	TRR	Sensitivity	Specificity	AUC	TRR	Sensitivity	Specificity	AUC
PCA selection	92.6	95.2	90.0	97.3	88.8	80.6	97.1	96.8

TRR is the total recognition rate; Sens and Spec are sensitivity and specificity, respectively. AUC is the estimated area under the ROC curve.

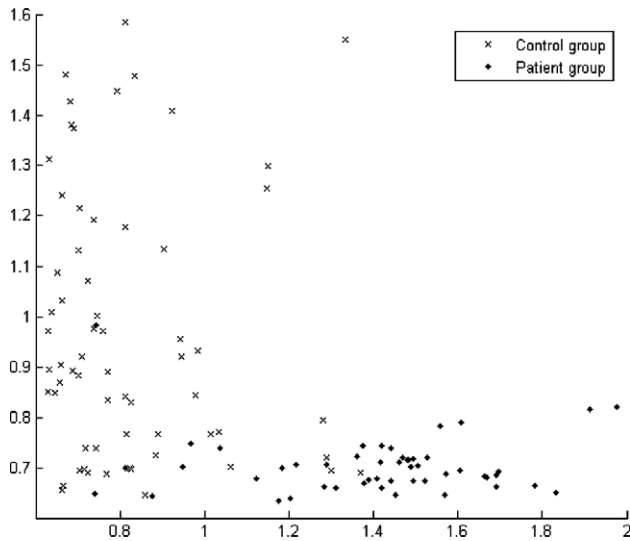


Fig. 4 – Scatter plot of the first two principle components on basis of which the classification patient-control group was made.

ponents provide most of the between-group separation. Fig. 3 shows a plot of the correlation coefficients, with the class indicator, which can be calculated from the linear discriminant weightings^{20,21} in the region between 1160 and 11,600 Da. The remainder of the plot is not shown, as the coefficients are effectively zero in that range. As can be seen, the classification was achieved primarily through a contrast in peak intensities between the first and second principal component. This can also be seen from the scatter plot shown in Fig. 4: low intensities at the first peak for cases separated cases from controls. Likewise, a small contribution for con-

trols at the second peak separated controls from patients. To illustrate these results further, we simply calculated the contrast between the two peak intensities directly across all subjects and constructed a simple one-dimensional summary of the data, as shown in the histogram displayed in Fig. 5, showing overlapping histograms of this (ad hoc) contrast for each group separately. The separation is clearly visible and quantified the significance of this difference by performing a two-sample Student t-test on this contrast ($t = 14.0$, $P < 0.0001$).

4. Discussion

Our study supports the hypothesis that serum protein profiles can discriminate a normal from a malignant state of organs, in our case of the colon. Here, we show that, based upon information in MALDI-TOF serum spectra, a classifier could be constructed for the detection of CRC. This classifier, calibrated and validated on spectra of week one¹ demonstrated a sensitivity and specificity of 95.2% and 90.0%, respectively. Thirty-four patients out of thirty-seven with early stage disease (stage 1 and 2) and all patients with stage 3 or 4 disease were correctly classified as having cancer. For the misclassified control subjects it was not possible to retrieve the current physical state as it concerned anonymous healthy controls.

Sensitivity and specificity of 80.6% and 97.1%, respectively, was achieved when the entire double cross-validated evaluation was repeated for the data of week 2. The latter evaluation, through use of replicate measurements within the double cross-validation, is likely to provide the more realistic assessment of true error rates and appears to better represent possible diagnostic potential as will be discussed further in this paper.

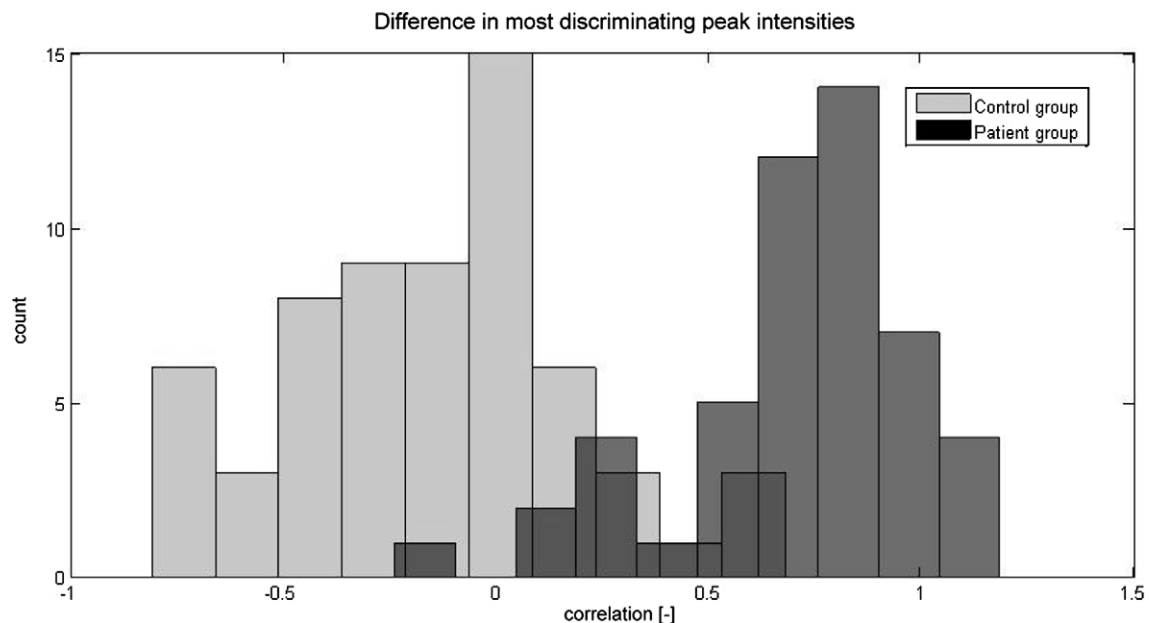


Fig. 5 – Histogram showing the difference between the normalized intensities of the two most discriminating “peaks” (bins). The X-axis shows the difference between the normalized intensities of the peaks. On the Y-axis the number of subjects is displayed.

Although previous studies have reported similar high classification results for various solid tumours, we prefer evaluation through a thorough study design and double cross-validation of classification as proposed in this study.^{3–6,12,22,23} As a great variety of different discriminating peaks for the same malignancy have been described^{3,4,24} caution with proteomic data has been stressed before.^{7,8} The discrepancies in discriminating protein profiles, found by different research groups, lead to serious concerns regarding biological variations and technological reproducibility issues. Therefore, we used a standardized and well-documented sample collection and a thorough study design, matching biological variables and pre-analytical conditions.¹⁶ Still, patient samples from all stages of CRC were equally distributed over the different target plates, as was the male/female ratio between the two groups, excluding these factors as a discriminator in the detection classifier. Unfortunately, there was significant difference in age; the control group being younger than the CRC patients. Ideally, the control group should consist of age-matched symptom-free individuals undergoing a colonoscopy showing no aberrations. However, due to the nature of the intervention, ethical legislation and the increasing disease burden with ageing, this is difficult to realize in clinical practice. Notwithstanding, we performed an analysis to examine the differences in intensity of most discriminating peaks based on age, gender and sample age. In the present study, there was no significant contribution of one of these factors on the most discriminating peaks of our classification model (data not shown).

A source of bias may be the presence of batch effects, such as day-to-day variation or plate-to-plate variation. The presence of batch effects is unavoidable and – rather than to eliminate them from the design – a better approach is to account for and accommodate these effects, in such a way that they do not lead to errors of artificially induced group separation. Consequently, we randomly distributed the available samples from each group across the batches such that proportions were equal across batches within group. The so-called randomized block design ensured that the batch effect – if it materialized – would not induce an artificial between-group effect.^{14,15}

A crucial point of discussion in the evolving field of clinical proteomics is validation of classification.^{9,25} Given the sample size achievable within the experiment, use of a separate (possibly set-aside) validation set was precluded. The other problem is “predictive optimisation”. However, as evaluation of predictive performance of the classifier is our primary focus, it is crucial that calibration is not carried out on the same data used for validation, which in turn would require an additional tuning set. Again, this would greatly increase the burden of collecting sufficient samples. For these reasons, other studies often carry out predictive optimisation on the full data in practice – which results in optimistically biased error rate evaluations, particularly with high-dimensional data such as in mass spectrometry proteomics.²⁶ As we have already suggested, another option is to reduce the available calibration data prior to optimisation, so as to set aside data, both for a training and validation set. However, this “solution” is not as innocent as would appear at first sight, since it typically reduces the calibration set beyond the point of what is

needed for reasonable calibration. Once more, this is particularly the case in high-dimensional cases such as clinical proteomics, where samples of malignancies are relatively difficult to obtain. Both problems may be avoided by carrying out a double-cross-validatory approach, which avoids the need for separate test and validation sets to yield unbiased error rate estimates. The double validatory aspect of the procedure results from the fact that the discriminant rule constructed to classify the left-out data was optimised through a secondary cross-validatory evaluation within the first cross-validatory layer (i.e. full cross-validation again on each “leftover” set after removal of an observation). In this manner, we were able to integrate predictive optimisation and predictive unbiased validation in the same procedure, without loss of data – which is a crucial requirement to get realistic estimates of error rate with high-dimensional data while reducing the risk of overfitting (first proposed by Stone²⁷). Although the principle is sound and understood, this procedure has until recently not been applied in practice due to the considerable computational cost and (algebraic) complexity of the method.

Our classifier is based on Fisher linear discrimination, which has been derived and may be justified based on a variety of principles of inference, such as maximization of the between-group separation relative to within-group error in the two-group case or the likelihood principle for normally distributed within-group populations. The methodology has been amply studied and has been established as reliable and robust form of classification and discriminant analysis. Furthermore, Fisher discrimination does not require an assumption of within-group normal dispersion.^{21,28,29} Hastie and colleagues contain an up-to-date account of many new applications that demonstrate the continuing success of the approach.^{30,18} Much similar and confirmatory experience has accumulated in related fields of application,^{19,31} which identifies this classification method as most reliable in high-dimensional analysis. For proteomic mass spectra, principal components are attractive as it provides a means of non-parametrically smoothing and pooling information across peaks.

The controversy about the use of protein profiles as a pattern diagnostic without analysis of the diagnostic biomarkers remains to be solved for its clinical application. Identification and functional analysis of these discriminating proteins/peptides might render new insights on tumour development and environmental responsiveness, which could eventually be translated in new diagnostic and prognostic insights for the clinician. Unfortunately, little success has been booked so far in assigning reproducible discriminating biomarkers.^{12,25} Though this study showed two most discriminating mass values of MALDI-TOF based protein profiling analysis to be low-molecular weight fragments, we have not identified these potential biomarkers yet.

In the present study we used patterns of proteomic signatures from high dimensional mass spectrometry data to generate a diagnostic classifier for the detection of CRC. To our knowledge, this is the first double cross-validatory study in a randomised block design in this field of research. Although independent validation would strengthen the observations and follow-up studies are now underway, we obtained

maximal reliability in classification in this study while maintaining protection against overfitting. Due to the relatively small sample size we have chosen to use our entire dataset for a within-study validation to avoid optimistic biased (error) misclassification rates. To assess the performance of our classifier a further independent validation study will be necessary. In addition, in future studies the specificity of discriminating protein profiles for colorectal cancer have to be assessed in comparison with other cancer types. Nevertheless, we are currently able to detect CRC accurately on the basis of differences in actual information in the serum protein profiles with a rigorously standardized approach and exclusion of batch effects. Thus, although introduction in a routine clinical setting may take longer than originally hoped for, this study is an initial proof for a successful evolution of the potentially great use of discriminating protein profiles in the detection of CRC.

Conflict of interest statement

None declared.

Acknowledgements

The authors thank all employees at the CKCL and especially Dr. M. Frohlich and Wil van der Bent for their contribution to the serum collection. Gratitude is also expressed to the secretaries at our Department of Surgery and all participating patients and healthy control subjects. Dr. G.J. Liefers is thanked for his contributions to the manuscript.

Appendix 1

Fisher linear discriminant analysis may be defined as assigning an observation to the group for which the smallest within-group distance $D_g(\mathbf{x}) = (\mathbf{x} - \mu_g)\Sigma^{-1}(\mathbf{x} - \mu_g)^T$ is found for the corresponding observed feature vector $\mathbf{x} = (x_1, \dots, x_p)$ with respect to the g th group ($g = 1, 2$ here, for either cases or controls), where p is the dimensionality of the problem, μ_g denotes the population within-group sample mean for the g th group and Σ is the (common) within-group dispersion matrix. We may estimate the population means through the within-group sample means. When the dimensionality of the problem is greater than the sample size, as is the case in this problem, the observed within-group pooled covariance matrix S will typically not be of full rank and hence special measures are called for before we can apply the above paradigm in this context. This can be achieved through an initial principal components decomposition of the observed within-group pooled covariance matrix $S = Q\Lambda Q^T$, where Q and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_r)$ are the matrices of principal component weightings and variances, respectively (r is the rank of the pooled covariance matrix). We then re-estimate the within-group covariance matrix by only retaining the first k components only: $S_{(k)} = Q_{(k)}\Lambda_{(k)}Q_{(k)}^T$, which account for most of the variation in the spectra. The discriminant rule may now be expressed as assigning an observation to the group for which we observe the smallest sample estimate $\hat{D}_g(\mathbf{x}) = (\mathbf{x} - \bar{\mathbf{x}}_g)S_{(k)}^{-1}(\mathbf{x} - \bar{\mathbf{x}}_g)^T$.

In the two-group case, this is also equivalent to least-squares regression analysis using the Moore–Penrose inverse of the pooled covariance matrix when $k = r$ (all components kept, also known as the shortest least squares regression), or else is equivalent to the so-called shrunken least-squares regression.^{20,21} When choosing $k < r$, the choice may be made through appeal to a (cross-) validity evaluation of the performance of the respective possible choices for the parameter k . The above methodology has been described and compared to other methods in the recent paper by Mertens,¹⁸ which shows this method to be competitive in the closely related high-dimensional setting for classification with microarrays. Much similar and confirmatory experience has accumulated in related fields of application, which identifies this classification method as reliable and stable in high-dimensional analysis, as has been described by Stone and Jonathan, among others.^{19,31}

REFERENCES

1. Ruo L, Gougoutas C, Paty PB, Guillem JG, Cohen AM, Wong WD. Elective bowel resection for incurable stage IV colorectal cancer: prognostic variables for asymptomatic patients. *J Am Coll Surg* 2003;**196**(5):722–8.
2. Gill S, Sinicrope FA. Colorectal cancer prevention: is an ounce of prevention worth a pound of cure. *Semin Oncol* 2005;**32**(1):24–34.
3. Adam BL, Qu Y, Davis JW, Ward MD, Clements MA, Cazares LH, et al. Serum protein fingerprinting coupled with a pattern-matching algorithm distinguishes prostate cancer from benign prostate hyperplasia and healthy men. *Cancer Res* 2002;**62**(13):3609–14.
4. Petricoin III EF, Ornstein DK, Pawletz CP, Ardekani A, Hackett BA, Hitt BA, et al. Serum proteomic patterns for detection of prostate cancer. *J Natl Cancer Inst* 2002;**94**(20):1576–8.
5. Rai AJ, Zhang Z, Rosenzweig J, Shih I, Pham T, Fung ET, et al. Proteomic approaches to tumour marker discovery. *Arch Pathol Lab Med* 2002;**126**(12):1518–26.
6. Yanagisawa K, Shyr Y, Xu BJ, Massion PP, Larsen PH, White BC, et al. Proteomic patterns of tumour subsets in non-small-cell lung cancer. *Lancet* 2003;**362**(9382):433–9.
7. Hu J, Coombes KR, Morris JS, Baggerly KA. The importance of experimental design in proteomic mass spectrometry experiments: some cautionary tales. *Brief Funct Genomic Proteomic* 2005;**3**(4):322–31.
8. Coombes KR, Morris JS, Hu J, Edmonson SR, Baggerly KA. Serum proteomics profiling – a young technology begins to mature. *Nat Biotechnol* 2005;**23**(3):291–2.
9. Ransohoff DF. Rules of evidence for cancer molecular-marker discovery and validation. *Nat Rev Cancer* 2004;**4**(4):309–14.
10. Boguski MS, McIntosh MW. Biomedical informatics for proteomics. *Nature* 2003;**422**(6928):233–7.
11. Villanueva J, Philip J, Entenberg D, Chaparro CA, Tanwar MK, Holland EC, et al. Serum Peptide profiling by magnetic particle-assisted, automated sample processing and maldi-tof mass spectrometry. *Anal Chem* 2004;**76**(6):1560–70.
12. Diamandis EP. Analysis of serum proteomic patterns for early cancer diagnosis: drawing attention to potential problems. *J Natl Cancer Inst* 2004;**96**(5):353–6.
13. Baggerly KA, Morris JS, Coombes KR. Reproducibility of SELDI-TOF protein patterns in serum: comparing datasets from different experiments. *Bioinformatics* 2004;**20**(5):777–85.

14. Box GEP, Hunter WG, Hunter JS. *Statistics for experimenters*. New York: John Wiley & Sons Inc.; 1978.
15. Cox DR, Reid N. *The theory of the design of experiments*. London: Chapman/Hall CRC; 2000.
16. de Noo ME, Tollenaar RAEM, Ozalp A, Kuppen PJK, Bladergroen AM, Deelder AM. Reliability of human serum protein profiles generated with C8 magnetic beads assisted MALDI-TOF mass spectrometry. *Anal Chem* 2005;**77**(22):7232–41.
17. Whittaker ET. *On a new method of graduation*. 41 ed.; 2005.
18. Mertens BJA. Microarrays, pattern recognition and exploratory data analysis. *Stat Med* 2003;**22**:1879–99.
19. Mervyn Stone, Philip Jonathan. Statistical thinking and technique for QSAR and related studies. Part II: specific methods. *J Chemometr* 1994;**8**(1):1–20.
20. Ripley BD. *Pattern recognition and neural networks*. Cambridge: Cambridge University Press; 2005.
21. Seber GAF. *Multivariate observations*. New York: John Wiley & Sons Inc.; 2005.
22. Yu JK, Chen YD, Zheng S. An integrated approach to the detection of colorectal cancer utilizing proteomics and bioinformatics. *World J Gastroenterol* 2004;**10**(21):3127–31.
23. Diamandis EP. Point: proteomic patterns in biological fluids: do they represent the future of cancer diagnostics? *Clin Chem* 2003;**49**(8):1272–5.
24. Qu Y, Adam BL, Yasui Y, Ward MD, Cazares LH, Schellhammer PF, et al. Boosted decision tree analysis of surface-enhanced laser desorption/ionization mass spectral serum profiles discriminates prostate cancer from noncancer patients. *Clin Chem* 2002;**48**(10):1835–43.
25. Somorjai RL, Dolenko B, Baumgartner R. Class prediction and discovery using gene microarray and proteomics mass spectroscopy data: curses, caveats, cautions. *Bioinformatics* 2003;**19**(12):1484–91.
26. Baggerly KA, Morris JS, Wang J, Gold D, Xiao LC, Coombes KR. A comprehensive approach to the analysis of matrix-assisted laser desorption/ionization-time of flight proteomics spectra from serum samples. *Proteomics* 2003;**3**(9):1667–72.
27. Stone M. Cross-validatory choice and assessment of statistical predictions. 36 ed.; 1973.
28. McLachlan GJ. *Discriminant analysis and statistical pattern recognition*. New York: John Wiley & Sons Inc.; 2004.
29. Hand DJ. *Construction and assessment of classification rules*. New York: John Wiley and Sons Inc.; 1997.
30. Hastie T, Tibshirani R, Friedman J. *The elements of statistical learning*. Berlin: Springer-verlag; 2001.
31. Stone M, Jonathan P. Statistical thinking and technique for QSAR and related studies. *J Chemometr* 1993;**7**:455–75.