A Comparative Study of NPI and PLANN-ARD by Prognostic Accuracy and Treatment Allocation for Breast Cancer patients

Ian Jarman
P. J. G. Lisboa

This paper presents a retrospective cohort study of breast cancer patients following surgery, to investigate how a new prognostic model, PLANN-ARD and an accepted prognostic index, the Nottingham Prognostic Index (NPI), compare not just in modelling survival but also in relation to the choice of treatment for particular patient cohorts.

2nrich Project Public Release 1.0
http://www.cms.livjm.ac.uk/2nrich/docs/papers/studynpi.pdf

A comparative study of NPI and PLANN-ARD by prognostic accuracy and treatment allocation for breast cancer patients

I.H. Jarman and P.J.G. Lisboa

School of Computing and Mathematical Sciences, Liverpool John Moores University, Byrom Street, Liverpool L3 3AF, England.

Abstract There is considerable interest in the use of Artificial Neural Networks (ANN) in medical applications, including prognostic modelling through survival analysis. This paper presents a retrospective cohort study of breast cancer patients following surgery, to investigate how a new prognostic model, PLANN-ARD and an accepted prognostic index, the Nottingham Prognostic Index (NPI), compare not just in modelling survival but also in relation to the choice of treatment for particular patient cohorts. The clinically used index, NPI, is found to correlate particularly well with treatment received. Both methods are fitted to routinely acquired clinical records with 5-year follow-up from 917 patients recruited during 1983-89, and tested on a subsequent cohort of 559 patients recruited over 1990-93. In addition, a new methodology for prognostic risk assessment was developed by combining the two models already considered. The prognostic groups identified with the combined model show a consistent progression in severity of illness characterised by 5-year survival and correlate also with a progression in the type of treatment received. The methodologies are further evaluated for 269 patient records from a second clinical centre.

Keywords: Artificial Neural Networks, proportional hazards model, survival analysis, prognostic modelling.

1. Introduction

Survival analysis is a large area of interest in medical statistics, where the proportional hazard model [4] is the most widely used method to model censored data one of which, the NPI [7], we use as our benchmark prognostic model and also in our analysis by integrating into our new methodology.

This interest has inevitably increased research into artificial neural networks for censored data and we use an ANN, as proposed by Biganzoli *et al* [1] and modified by H. Wong *et al* [2-5] to form PLANN-ARD (Partial Logistic Artificial Neural Network – Automatic Relevance Determination). This method uses a Bayesian framework to carry out Automatic Relevance Determination (ARD) in a multi-layer perceptron neural network to model censored data. Censored data is a particular feature of survival data and arises when an individual drops out of a study for reasons other than the event of interest, in this analysis death attributed to breast cancer.

By comparing PLANN-ARD with NPI in terms of survival in a cross-tabulation of their respective prognostic indices led to the new methodology for a second prognostic model by grouping patients into risk groups defined by similar survival within the matrix.

Although these models are intended as a prediction of survival for a group of patients with similar clinical and pathological measures, the question arises which of these statistically similar models should be the preferred one.

Altman and Royston [4] went further than just selecting a model based on the above criteria and stated several reasons for wishing to predict the outcome of future patients, which included:

- To inform treatment or other clinical decisions for individual patients – can a model indicate the type of treatment an individual may receive?
- To inform patients and their families – this is the survival prediction for the patient, this seems to be where clinicians mainly use prognostic models.
- To create clinical risk groups for informing treatment or for stratifying patients by disease severity in clinical trials – we want to correlate the risk groups with the severity of the illness.

These reasons should act as a point of reference to guide our selection of the best prognostic model, that is, does the model address each of the above requirements and of course is the reason we are interested in the first of these points, the treatment patients receive who have similar predicted survival outcome.

2. Patients and Data

We have two sets of data from Christie Hospital in Manchester, the data set that the new models were derived, and the Linda McCartney Centre in Liverpool to enable the two centre comparison.

For consistency, data from both hospitals went through 2 stages of a filtering process. The first stage, described in table 1, was applied to allow a comparison to be made within Christie Hospital between prognostic models. The second stage is a refinement of the inclusion criteria, table 2, to enable a fair comparison between the two centres.

The need for the second stage is due to different patient referral procedures that apply at the two centres, further explained below.

Any records with tumour stage 0 were not included.	
Inclusion criteria	
Follow up NPI [5] Metastasis Node Stage Pathological size	≥ 5 years Not missing = 0 = 0 or 1 < 2cms or 2-5cms

Table 1: The first stage of filtering gives us the data set, from which comparisons between different models within Christie's can be made.

1. 2. 3.	NPI = 3 or 4 **OR** Node stage positive **OR** ER negative

Table 2. Second filter gives us the data set from which comparisons between centres can be made.

2.1 Christie Hospital

The first set of data to be used for training the neural network consists of 1616 records of women patients referred to the Manchester Christie Hospital between 1983 and 1989. The event of interest is the death attributed to breast cancer, as determined from the coronary report or specialist consultant. All other causes of death and other losses to follow up is censored data, the study is for a five year period so surviving patients are censored after this time. We wanted to model the data for PLANN-ARD on a low-risk cohort that used the criteria set out in table 1 but omitting NPI from the restrictions, which was not necessary for the modelling process, this left 917 patient records.

The second set of data to be used for validation of PLANN-ARD, consist of 1266 records referred to Christie Hospital between 1990 and 1993. To validate the new model for PLANN-ARD on the low-risk cohort we followed the convention of omitting NPI from the selection criteria in table 1, this left 931 patient records. There were 559 patients remaining for the full inclusion criteria in table 1.

The further adjustment to the data, table 2, for comparison purposes, left a total of 353 patients.

2.2 Linda McCartney Centre

This data consists of 808 patients who were on the patient list of one oncologist at the centre, who recorded details from date of diagnosis. (The patients were referred to the oncologist based on a poorer prognosis, hence the need for a second filtering of data for comparison). These dates ranged from one patient diagnosed in 1957 the rest from 1975 to 2001, from this 269 were consistent with the inclusion criteria. These patients being first diagnosed between 1984 and 1998.

With the final filter applied we were left with 248 patients for comparison purposes.

2.3 Data Description

PLANN-ARD was modelled on the first data set using six fields from the training data, found to have been predictive of survival in a preliminary monthly cohort study with 5 year follow up, using a proportion hazards model with a forward selection stepwise procedure recommended by Collett [10] and Akaike's information criterion (AIC) [11], to measure the significance of adding each variable to the model. Six variables were selected, namely, *age, clinical stage nodes, histology, node ratio, pathological size and ER status*, which agree with the variables selected in a previous study [2] of the same data but adding two more, age and ER status. These were selected from an original list of 15 categorical variables, recorded for each patient, table 3.

For the model selection there were inevitably records with missing variables. Previous analysis on data from Christie Hospital, for patients referred from 1983 to 1989 suggested that missing variables might be informative. For this reason a separate attribute was used for the missing data (variables affected – pathological size, nodes involved, nodes ratio, histology and ER status).

3. PLANN-ARD

PLANN-ARD the Partial Logistic Artificial Neural Network with Automatic Relevance Determination, is the prediction of smooth estimates of the discrete time hazard, which allows for non-proportionality of the hazard function to be visualised directly, as well as potential non-linear interactions between covariates. It is implemented within a MLP neural network, which is a feed forward non-linear neural network with a sigmoidal activation. This is an extended form of a logistic regression model with extra layers comprised of hidden nodes, one hidden layer for our MLP network structure.

Once the network weights are estimated, the survivorship is calculated from the estimated discrete time hazard by multiplying the conditionals for survival over successive time intervals treated as independent events, this gives:

$$S(t_k) = \prod_{l=1}^{k} \left(1 - h(t_l)\right) \quad (1)$$

Variable	Categories	Coded Attributes
Menopausal Status	Pre-menopausal Peri-menopausal Post-menopausal	1 2 3
Age Group	20-39 40-59 60+	1 2 3
Predominant Site	Upper Outer Lower Outer Upper Inner Lower Inner Subareolar	1 2 3 4 5
Side	Right Left	1 2
Maximum Diameter of Tumour	<2cm 2-5cm 5+cm Unknown	1 2 3 9
Clinical Stage Tumour	T0 (No Tumour) T1 (Tumour < 2cm) T2 (2-5cm) T3 (5+cm) T4 (any size but fixed on the rib cage)	0 1 2 3 4
Clinical Stage Nodes	N0 (no nodes found clinically, or node negative by histology N1 (ipsilateral and mobile axillary nodes) N2 (nodes fixed) N3 (nodes fixed and cannot be removed)	0 1 2 3
Metastasis Stage	M0 (no distant metastasis) M1 (positive)	0 1
Clinical Stage	0 1 2 3 4	1 2 3 4 5
Histology	Inv. Duct Inv. Lob/Lob in situ In Situ / Mixed / Medullary / Ucoid / Papillary / Tubular / other Mixed in Situ Unknown	1 2 3 9
Number of Nodes involved	0 1-3 4+ 98 (too many to count) Unknown	1 2 3 4 5
Number of Nodes removed	0-9 10-19 20+ 98 (too many to count) Unknown	1 2 3 4 5
Node Ratio	0-20% 20-40% 40-60% 60+% Unknown	1 2 3 4 5
Pathological Size	<2cm 2-5cm 5+cm Unknown	1 2 3 4
ER Status (Oestrogen)	0-10 10+ 8888 (high positive value) Unknown	1 2 3 4

Table 3. Variables recorded with their attribute description and corresponding code

Estimating the weights requires a likelihood term for the status of one patient at time t_k, by using an indicator label 0 if a patient is alive at time t_k and 1 for the event of interest, death caused by breast cancer in this analysis. This generic non-linear model is called the Partial Likelihood Artificial Neural Network (PLANN) [1]. In contrast to a proportional hazards model [6], PLANN does not require proportionality of the hazards over time and predicts a smooth hazard function.

These generic non-linear models are also prone to over-fitting the data unless regularised to control complexity. A Bayesian framework is used as a robust method for estimating weight parameters. This takes on a three step sequence. First, a penalisation term is added to the objective function. Second, the regularisation parameters are estimated to control the penalty term. The third step involves the interpretation of the whole network as the evidence in favour of a candidate network structure, enabling model selection to be carried out. This third step is not utilised in this analysis as we had identified our model selection in an earlier proportional hazards analysis, stated earlier, this will be incorporated in a future analysis.

Due to the nature of survival analysis the distribution of the event, death caused by breast cancer, is very skewed due to the very few occurrences of the event within the large number of time steps used. This skewness leads to inaccurate regularisation because equal numbers are assumed for the ‘zeros’ and ‘ones’. To overcome this problem an approximation for predicted hazards is used.

Once the modelling process is complete it is necessary to create a mortality risk score or prognostic index for each patient. This allows a ranking system of mortality, which can then be used to separate patients into distinct risk groups in terms of survival. This score is made possible by taking the logit of the hazard prediction of the PLANN-ARD. However, as this is time dependent, a cumulative index is obtained by averaging it over the number of time intervals of the study, this then gives:

$$\Pr ognosticIndex = \frac{\sum_{i=1}^{T} \log it(y_i)}{T} \qquad (2)$$

4. Training and Validating PLANN-ARD

For this data we found that 12 hidden nodes, 20 iterations for early stopping and all the control parameters set to 0.1 to tackle the problem of over-fitting the network was sufficient to model the data accurately. The analysis was then carried out using a time step of 1 month over the 5-year study period giving us 60 time intervals.

A risk score from the hazard prediction was calculated using eq. (2) for each individual record. Groups of patients were then separated into distinct statistically significant prognostic risk groups by means of a pairwise log-rank test.

We can now evaluate these risk groups by plotting survival using Kaplan-Meier (KM) curves with 95% confidence intervals. From these plots we are looking for distinct survival groups indicated by very little overlap of the error bars between risk groups.

From this we identified 4 significantly different risk groups that had distinct survival plots based on the pairwise log-rank test.

Having established our model from the training data, validation of the PLANN-ARD prognostic model is extremely important in order to check the generality of the model on an independent data set. That is to say, if the results of the training data are satisfactorily reproduced on an independent data set we can say this is a valid prognostic model.

Looking at the KM curves for the validation dataset in figure 1 we can see that patients fall into well-defined prognostic groups with very little overlap of the error bars, so we can conclude that the PLANN-ARD is a satisfactory model

5. Comparison with NPI

Now that we have our PLANN-ARD model we need to compare this with our benchmark proportional hazards model, the NPI. We applied the first filter in table 1 leaving 559 patients for the analysis.

The results for NPI and PLANN-ARD are shown in figures 2 and 3 respectively. Both sets of survival curves are very similar and well separated with little overlap of their respective error bars. The only apparent difference between the two, by inspection of the plots, is group 4 where PLANN-ARD has greater overlap with risk group 3 in the first 30 months after surgery.

When we look at the populations and the specific confidence intervals in table 4, NPI has higher survival to five years in each of the 4 corresponding PLANN-ARD risk groups and very different patient numbers in three of the four groups, group 4 being the comparable group in terms of patient numbers.

To discover how well the two models are correlated we produced a scatterplot of the two prognostic indices, figure 4. In addition, a matrix was overlaid onto the plot in order to cross-match PLANN-ARD and NPI groups. The highlighted areas are the sectors with the same risk group for both models.

Although there is some relationship, we expected a much more correlated plot concentrated along the highlighted sectors where the risk groups agree and have similar survival as shown in figures 4 and 5.

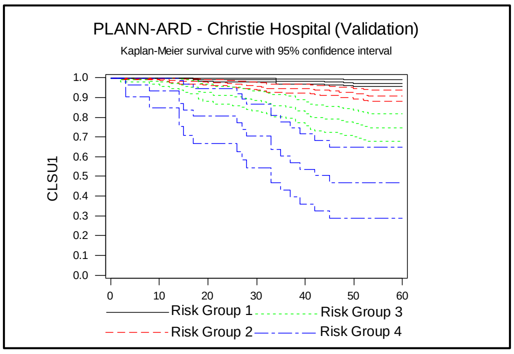

Fig.1 KM curves of the whole validation data for PLANN-ARD show a satisfactory separation of each risk group

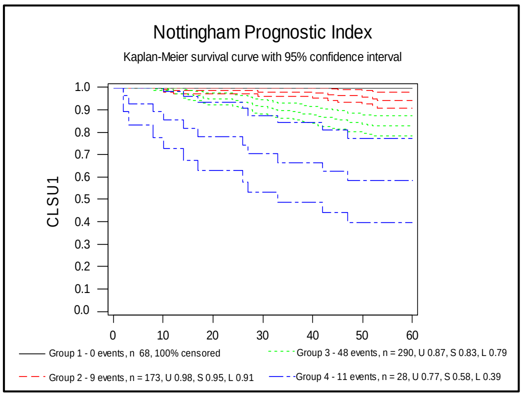

Fig.2 KM curves for NPI on the filtered validation data set with 559 patients. This confirms NPI as a good prognostic model with very little overlap between error bars.

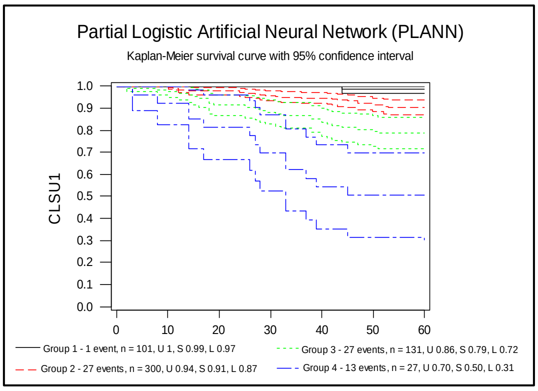

Fig.3 KM curves for PLANN-ARD on the filtered validation data set with 559 patients. The curves for each risk group are very similar to the full low-risk cohort in figure 3.

This led us to examine the survival of patient groups within each matrix point in order to discover which of the prognostic models had homogeneous groups of patients in terms of survival, indicated by consistent survival curves from one matrix plot to the next in either the matrix rows or the columns.

Group 1	PLANN-ARD	NPI
Upper range	1	100% censored Unable to calculate
Survival	0.99	
Lower range	0.97	
n =	101	68
Events	1	0

Group 2	PLANN-ARD	NPI
Upper range	0.94	0.98
Survival	0.91	0.95
Lower range	0.87	0.91
n =	300	173
Events	27	9

Group 3	PLANN-ARD	NPI
Upper range	0.86	0.87
Survival	0.79	0.83
Lower range	0.72	0.79
n =	131	290
Events	27	48

Group 4	PLANN-ARD	NPI
Upper range	0.70	0.77
Survival	0.50	0.58
Lower range	0.31	0.39
n =	27	28
events	13	11

Table 4. Comparisons of each risk group reveal that NPI in each case has higher survival than the comparable PLANN-ARD group. Apart from group 4, patient numbers are very different in each comparable group.

For instance, if the KM survival curves were more consistent in a matrix row, then this would indicate a homogeneous group of patients in terms of survival for a risk group in PLANN-ARD. And alternatively, if a column shows this consistency then NPI would be a homogeneous risk group for survival.

Ultimately this led to the development of the new methodology, of a prognostic model based on observed survival for the two prognostic models within the survival matrix.

5. Cross-Matching of Prognostic Models

From the cross-tabulation table of patient numbers in table 5, there are discrepancies between patients in similar surviving groups. For instance, only 19 patients fall in risk group 1 for both prognostic models.

Having produced a matrix of KM survival curves, figure 5, to mirror the same points in figure 4, we can see that for both models the higher and lower survival groups, 1 and 4, have homogeneous survival curves. Similarly for both models, risk groups 2 and 3 present a more heterogeneous set of survival curves with survival decreasing as the risk group in the corresponding prognostic index increases.

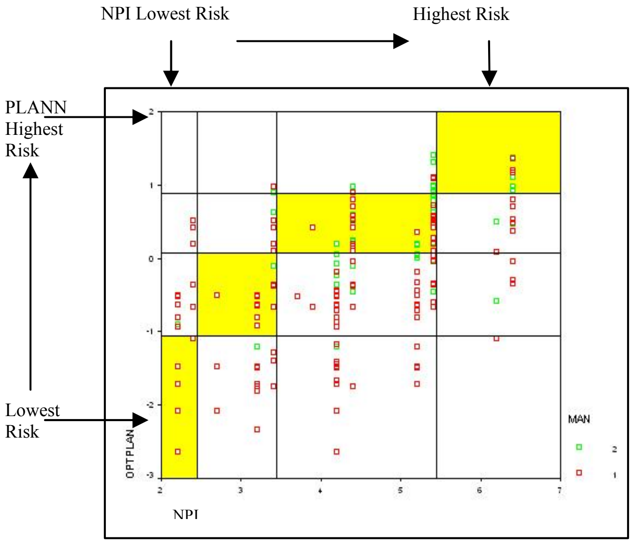

Fig. 4 Scatterplot of NPI vs. PLANN-ARD prognostic scores, overlaid with matrix to cross-match the respective risk groups. The highlighted areas are were risk groups agree for each index.

PLANN * NPI Crosstabulation

Count

		NPI				
		1.00	2.00	3.00	4.00	Total
PLANN	1	19	35	47		101
	2	41	110	142	7	300
	3	8	27	89	7	131
	4		1	12	14	27
Total		68	173	290	28	559

Table 5. A cross-tabulation table representing the number of patients within each matrix sector in figure 6, demonstrating the large disparity between NPI and PLANN-ARD patients in similar survival risk groups.

What actually does emerge is a configuration of patients with similar survival as represented in figure 6.

This pattern of survival curves is the basis of our proposed new prognostic model with patients being allocated into a risk group based on their position within the matrix of NPI vs. PLANN-ARD as opposed to a risk score.

Having allocated patients into risk groups based on survival pattern we produced KM curves, figure 7, which show a completely new mix of patients with very similar survival characteristics as the previous survival curves in figures 2 and 3, population from group 1 to 4 being 150, 287, 89 and 33 respectively.

We now have three prognostic models and, depending which model is used, a different survival prognosis for over half of the 559 patients in this analysis.

This inevitably leads to the problem of which is the correct model to apply. The solution may lie in the analysis of treatment profiles for the prognostic models.

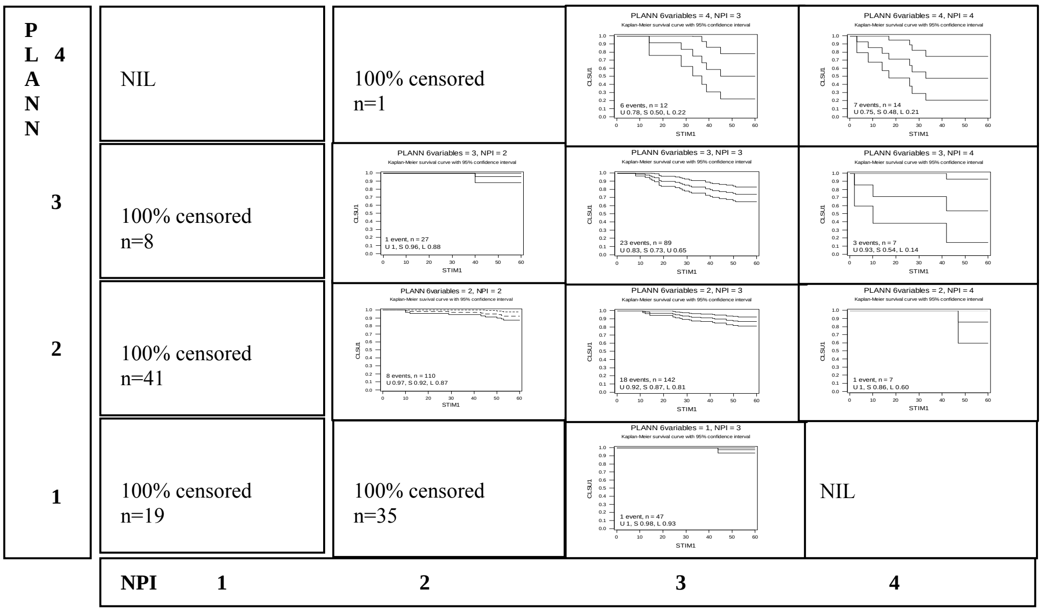

Fig. 5 Matrix of KM survival curves for NPI vs. PLANN-ARD prognostic groups. Each row represents survival curves for a PLANN-ARD risk group and each column survival curves for NPI.

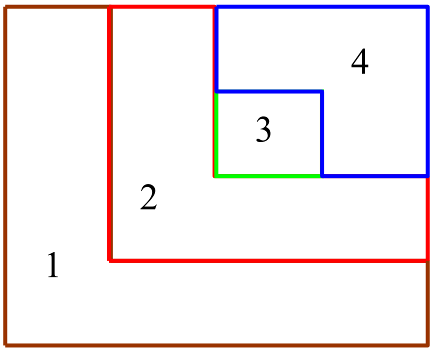

Fig. 6 The pattern of similar survival curves in figure 5 that form the basis of the new prognostic model.

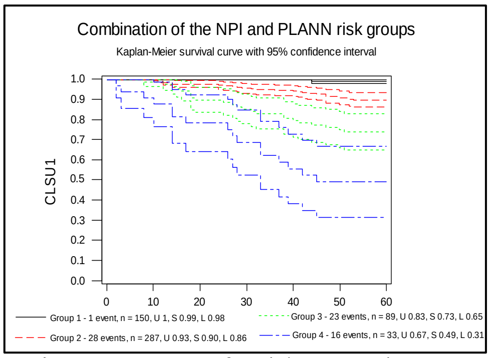

Fig. 7 KM curves for risk groups in a new prognostic model based on survival patterns, in a comparison of NPI vs. PLANN-ARD, shown in figure 6 for the survival matrix in figure 5.

6. Treatment Profiles

Guided by Altman and Royston [4] as stated in the introduction, each model was correlated with the choice of treatment.

Inspecting the treatment profiles for each risk group within the three models in figure 8 NPI and the combined model display a change in treatment profile as survival decreases. This change also makes sense as the lower risk groups have a propensity of patients with no treatment decreasing as survival falls with chemotherapy increasing as survival falls.

This pattern is less apparent for hormone treatment, although there is a tendency for hormone treatment to increase as survival drops. These patterns of change are not as clear-cut in PLANN-ARD but the changes in treatment do alter in a logical manner although the first two risk groups do have the same profile.

To gain further insight we employed the matrix design replacing KM survival curves with pie charts characterizing patient treatment, figure 9. Once again a homogeneous set of charts relating to a single risk group would be an indication of a better model, this time in terms of treatment.

As we can see by breaking down each prognostic group into this matrix format, NPI risk groups have consistent behaviour in each column that represent treatments for a specific risk group. For instance, NPI 1 have patients who only receive hormone treatment or more likely no treatment at all, this profile changes in a consistent manner for plots in each column until NPI 4, where patients now all receive either chemotherapy or hormone treatment. Although, this last column representing risk group 4 does show an increasing proportion of chemotherapy as we progress from low to high-risk PLANN-ARD groups.

This is only a small change when compared to the other models, these show a considerable change in profile within their respective risk groups, rows for

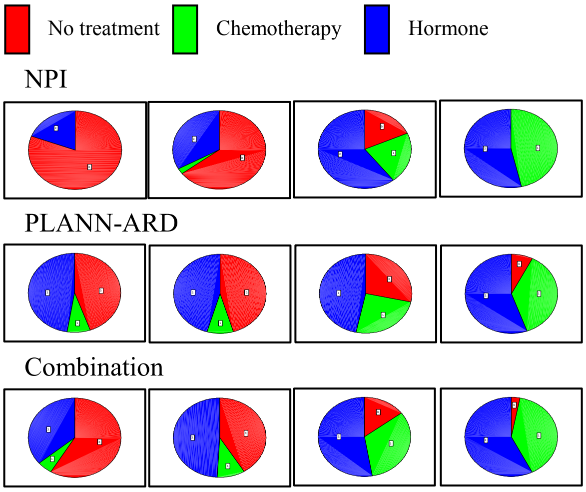

Fig.8 Treatment profiles for each prognostic model, risk increasing from left to right.

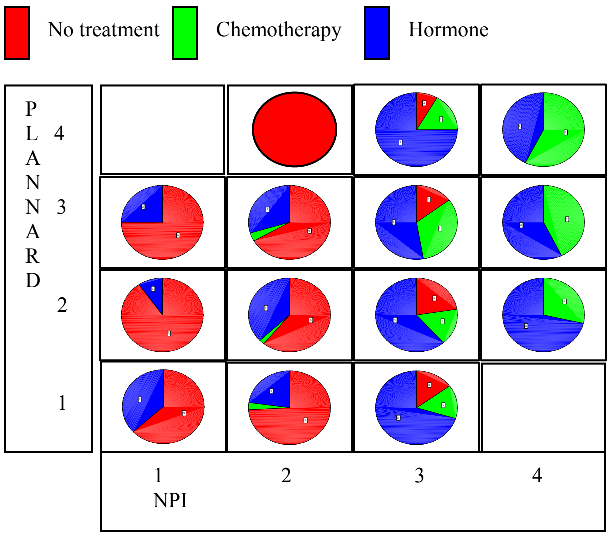

Fig. 9 Treatment profiles for Christie Hospital in a matrix of NPI vs. PLANN-ARD prognostic risk groups.

PLANN-ARD and the same pattern as figure 6 for the Combination Model. The profiles changing within a single risk group in the same manner NPI progressed from risk group to risk group.

From this we can conclude that, on the basis of homogeneous treatment profiles within each of their risk groups, NPI would be a more suitable model.

The decision on choice of treatment is addressed by a rule extraction algorithm OSRE [13].

7. Two-centre comparison

At this point we wanted to determine whether these new models would be well-defined in another centre, the Linda McCartney Centre, and if survival and treatment for the two centres were comparable for patients with the same prognostic risk group for a particular prognostic model. So we applied the second filter in table 2 in order to have similar patient profiles for each centre.

Looking at the KM curves in figure 10 for Christie Hospital and Linda McCartney Centre respectively, we can see that with this reduced set of patients risk groups 3 and 4 have well defined survival ranges, for each prognostic model, with little overlap of error bars and are in all cases comparable between centres.

However for each model, risk group 2 patients for Linda McCartney Centre wholly overlap group 3 patients, this may be due to patients with poorer prognosis referred to the oncologist for reasons that are not recorded for this data set. But nevertheless this observation is the same for each model, so both new models compare well in terms of survival with NPI, as the benchmark prognostic model.

When considering the treatment profiles for Linda McCartney Centre, there is an additional treatment category were a patient has received both chemotherapy and hormone treatment. This difference in treatment regime will be looked at in another paper, here we are concerned with the consistencies in treatment for our different prognostic models within the NPI vs. PLANN-ARD matrix

Inspecting the matrix of treatment profiles in figure 11, we can see that NPI, as for Christie hospital has consistent treatment for each column, although due to the few numbers in group 1 this cannot be stated for this group.

From this evidence, we can conclude that although all the models show the same consistencies within centre and same discrepancies between centres in terms of survival, and so making each a viable prognostic model, the homogeneity that NPI presents for treatment of patients within the individual risk groups, suggest that this may still be a more appropriate prognostic model than the new models presented in this analysis.

8. Conclusions

Two prognostic models were compared, the clinically used NPI score and a neural network PLANN-ARD. While the grouped survival for the prognostic groups generated by each method were very similar, the correlation with treatment received was much closer to the NPI index.

It is also of interest that even while the group survival statistics for the two models were similar, yet the assignment of individual patients to prognostic groups was very different. A combined prognostic index was also defined, for which there is a gradual and clear progression in survival.

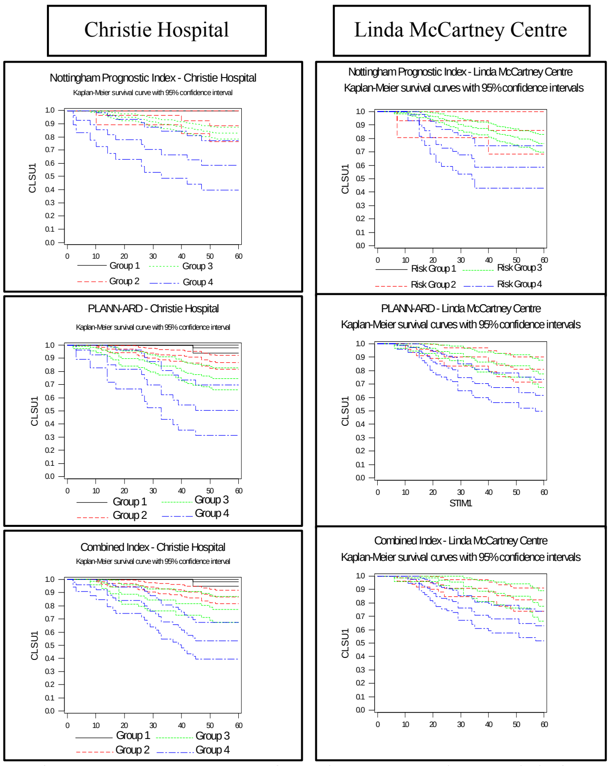

Fig. 10 KM curves of each prognostic model for both centres.

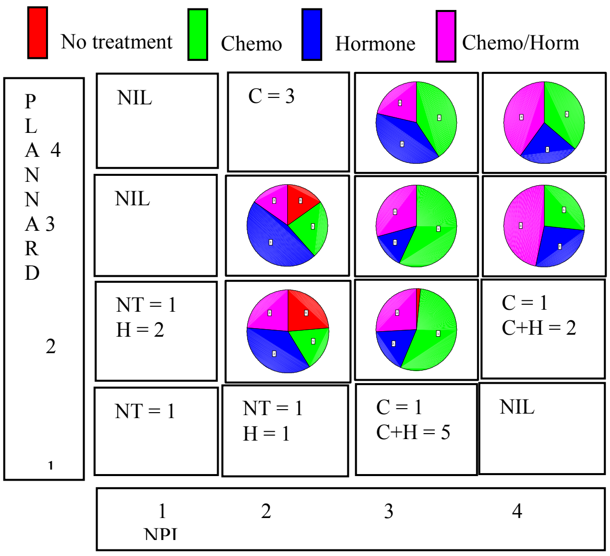

Fig.11 Treatment profiles for Linda McCartney Centre in a matrix of NPI vs. PLANN-ARD prognostic risk groups.

Bearing in mind that a prognostic model should also inform about treatment [4], the results show that NPI is still the preferred prognostic model as it presents homogeneous treatment profiles for subsets of individual risk groups within a matrix, split by the NPI and PLANN-ARD. This consistency was evident across both clinical centres.

Neither PLANN-ARD nor the Combination model displayed this homogeneity within the cross-matching matrix, although there was evidence of a consistent and gradual progression in survival and treatment received.

9. Acknowledgments

The authors gratefully acknowledge EPSRC for funding this research, as well as R. Swindell of Christie Hospital and S. O'Reilly of the Clatterbridge Centre for Oncology for granting access to patient data.

10. References

1. Biganzoli E, Boracchi P, Mariani L, Marubini E, (1998), Feed forward neural networks for the analysis of censored survival data: A partial logistic regression approach, Statist. Med, 17, 1169-1186.

2. Lisboa P.J.G., Wong H., Harris P. and Swindell R. (2003): A Bayesian neural network approach for modelling censored data with an application to prognosis after surgery for breast cancer, Artif Intell Med, **28**, pp. 1-25.

3. Lisboa, P.J.G., Wong, H., Vellido, A., Kirby, S.P.J., Harris, P. and Swindell, R. 'Survival of Breast Cancer Patients Following Surgery: a Detailed Assessment of the Multi-Layer Perceptron and Cox's Proportional Hazard Model', World Congress on Computational Intelligence, Anchorage, Alaska, IJCNN, 112-116, 1998.

4. P.J.G. Lisboa, H. Wong, P. Harris and R. Swindell A retrospective study of breast cancer prognosis using artificial neural networks, *in* Papadourakis, G.M. (ed.) Proc. 4th International Conference on Neural Networks and Expert Systems in Medicine and Healthcare (NNESMED), Milos, Greece, pp. 125-131, 20-22 June, 2001.

5. Lisboa, P.J.G. and Wong, H. Are neural networks best used to help logistic regression? An example from breast cancer survival analysis. Proc. International Joint Conference on Neural Networks, Washington. D.C., paper 577, 4-19 July, 2001.

6. Cox DR, Regression models and life tables, Journal of the Royal Statistical Society, B, 74, (1972), 187-220.

7. Haybittle JL, Blamey RW, Elston CW, Johnson J, Doyle PJ, Campbell FC, Nicholson RI, Griffiths K, A prognostic index in primary breast cancer, Br. J. Cancer, (1982), 45, 3621.

8. Lisboa, P.J.G. 'A review of evidence of health benefit from artificial neural networks in medical intervention', Neural Networks, Invited Paper, 15, 1, 9-37, 2002

9. Altman D G, Royston P. What do we mean by validating a prognostic model? Statist. Med. 2000; 19: 453-473.

10. D. Collett, Modelling Survival Data in Medical Research Chapman and Hall, 1994.

11. Akaike, H. (1974). A new look at the statistical model identification. *IEEE Trans. Appl. Comp*., AC-19, 716-723.

12. Kaplan EL, Meier P. Nonparametric estimation from incomplete observations. J. Am Stat Assoc 1958; 53: 457-81.

13. Etchells, T.A., Jarman, I.H. and Lisboa, P.J.G. 'Empirically derived rules for adjuvant chemotherapy in breast cancer treatment' *submitted to this conference*.