

UC Irvine

UC Irvine Previously Published Works

Title

A threshold-free summary index of prediction accuracy for censored time to event data

Permalink

<https://escholarship.org/uc/item/1hc811c0>

Journal

Statistics in Medicine, 37(10)

ISSN

0277-6715

Authors

Yuan, Yan
Zhou, Qian M
Li, Bingying
[et al.](#)

Publication Date

2018-05-10

DOI

10.1002/sim.7606

Peer reviewed

A Threshold-free Prospective Prediction Accuracy Measure for Censored Time to Event Data

YAN YUAN[†]

School of Public Health, University of Alberta, Edmonton, AB T6G1C9, Canada

QIAN M. ZHOU[†]

Department of Statistics and Actuarial Science

Simon Fraser University, Burnaby, B.C. V5A1S6, Canada

Department of Mathematics and Statistics

Mississippi State University, Starkville, Mississippi 39762, USA

BINGYING LI

Department of Statistics and Actuarial Science

Simon Fraser University, Burnaby, B.C. V5A1S6, Canada

HENGRUI CAI

School of Public Health, University of Alberta, Edmonton, AB T6G1C9, Canada

ERIC J. CHOW

Fred Hutchinson Cancer Research Center

Seattle Children's Hospital, University of Washington, Seattle, Washington, USA

GREGORY T. ARMSTRONG

Department of Epidemiology and Cancer Control, Division of Neuro-Oncology

St. Jude Children's Research Hospital, Memphis, TN 38105, USA

October 15, 2018

[†] To whom correspondence should be addressed. † Dr. Yuan and Dr. Zhou contributed equally to this work.

Abstract

Prediction performance of a risk scoring system needs to be carefully assessed before its adoption in clinical practice. Clinical preventive care often uses risk scores to *screen* asymptomatic population. The primary clinical interest is to predict the risk of having an event by a pre-specified *future* time t_0 . Prospective accuracy measures such as positive predictive values have been recommended for evaluating the predictive performance. However, for commonly used continuous or ordinal risk score systems, these measures require a subjective cut-off threshold value that dichotomizes the risk scores. The need for a cut-off value created barriers for practitioners and researchers. In this paper, we propose a threshold-free summary index of positive predictive values that accommodates time-dependent event status. We develop a nonparametric estimator and provide an inference procedure for comparing this summary measure between competing risk scores for censored time to event data. We conduct a simulation study to examine the finite-sample performance of the proposed estimation and inference procedures. Lastly, we illustrate the use of this measure on a real data example, comparing two risk score systems for predicting heart failure in childhood cancer survivors.

Censored event time; Positive predictive value; Precision-recall curve; Risk prediction; Screening; Time-dependent prediction accuracy

1 Introduction

Clinical medicine is facing a paradigm shift from current diagnosis and treatment practices to prevention through earlier intervention based on risk prediction [1]. Diagnosis and treatment approaches help individual patients seek relief from their symptoms. However, evidence is mounting that health interventions may be more effective in improving long-term health outcomes when they target asymptomatic individuals who are predicted to be at high risk for the condition of interest [2, 3]. The condition of interest typically has the following characteristics: 1) its seriousness may result in a high risk of mortality or significantly affect the quality of life; 2) early detection/intervention can make a difference in disease prognosis; and importantly but subtly 3) its event rate is low. A prevention approach to medicine relies on the development of risk scores to stratify individuals into different risk groups. Early intervention strategies are typically recommended to subjects who are in the high-risk group.

[†]Dr. Yuan and Dr. Zhou contributed equally to this work.

In the prevention paradigm, the use of risk scores as population *screening* tools is increasingly advocated in clinical practices, e.g. [4]. For example, one systematic review identified forty-six algorithms that predict the risk of type 2 diabetes [5]. Another study established several risk score systems to predict congestive heart failure for childhood cancer survivors who are at an elevated risk due to treatment toxicity [6]. One of the defining characteristics of screening is a low event rate in the targeted asymptomatic population. Taking the aforementioned two diseases as an example, the crude prevalence of undiagnosed type 2 diabetes, a common disease, was low at 3.5% in 1987 and 5.7% in 1992 [7], while the cumulative event rate of congestive heart failure by 35 years post childhood cancer diagnosis was 4.7% [6]. The event rate is much lower for other notable conditions such as cancer, multiple sclerosis, AIDS, dementia, COPD and perinatal conditions. A low event rate and a focus on prevention necessitate the development of *screening* tools such as risk scores.

Before a risk scoring system is adopted for clinical screening, evaluation of its predictive accuracy is critical. The most popular accuracy metric used in the clinical literature is the area under the receiver operating characteristic (ROC) curve (AUC). The AUC is a summary index of two accuracy metrics - sensitivity and specificity - which are both retrospective metrics. Thus, the AUC does not reflect the prospective predictive accuracy of risk score systems. Indeed, one influential article criticized these retrospective metrics as being of little use for clinicians because clinical interest almost always focuses on prediction [8]. In contrast, a prospective accuracy measure, such as positive predictive value (PPV), can prospectively answer the question: "Can risk scores be trusted?" Unfortunately, a risk score with high sensitivity and specificity, and thus a high AUC, can have poor PPV when applied to low-prevalence populations. This limitation is often overlooked by clinicians and biomedical researchers. Despite its popularity, studies confirm that the AUC is insensitive in evaluating risk prediction models. For example, including a marker with a risk ratio of 3.0 showed little improvement on the AUC, while it could shift the predicted 10-year disease risk for an individual patient from 8% to 24% [9]. This range would result in different recommendations on follow-up/intervention strategies.

Compared to the AUC, the PPV provides a more appropriate assessment of the prospective prediction performance of the risk score [10], making the PPV a superior metric for risk score systems used as screening tools. The PPV is calculated with data from a prospective cohort, where the risk scores are computed using baseline information and the outcome is followed prospectively.

Originally, the PPV was defined for a dichotomous test. Moskowitz and Pepe (2004) extended the definition of PPV for a continuous risk score [10]. Assuming that the higher the risk score, the greater the individual risk, the PPV is defined as the probability of having the disease when the risk score value is larger than a given cut-off value z ,

$$\text{PPV}(z) = \Pr\{D = 1 \mid Z \geq z\} \quad \text{and} \quad \text{NPV}(z) = \Pr\{D = 0 \mid Z < z\}, \quad (1)$$

where $D = 1$ indicates the presence of the disease, and $D = 0$ indicates the absence of the disease. Zheng et al. (2008) further generalized the definition to accommodate the censored event time outcome [11]. Since the PPV is threshold dependent, as seen in (1), it is often evaluated at either several fixed specificities or several fixed quantiles of the . Such evaluations allow the comparison across different risk score systems [10, 12]. The selection of specificities or quantiles can be subjective, and it is possible that different systems could outperform others, depending on the cut-off points selected [13].

For the above reasons, a threshold-free summary metric for the PPV is needed to facilitate its clinical usage. Two curves of PPV have been investigated in the literature. Raghavan et al. (1989) and Zheng et al. (2010) considered a curve of PPV versus quantiles of the risk score [13, 14]. However, they did not provide a summary index of the proposed PPV curve. A second curve is called the precision-recall (PR) curve, which was proposed in the information retrieval community [14, 15], where precision is equivalent to the PPV and recall is equivalent to the sensitivity. The relationship of PR and ROC curves and the area under them has been discussed in Davis and Goadrich [16]. They showed that the PR curve of a risk score system dominates that of another system if its ROC curve is also dominant. However, such a relationship does not exist for the area under these two curves [16, 17]. Two recent papers illustrated the advantage of using the area under the PR curve over the AUC for predicting low prevalence diseases [18, 19]. We refer to the summary metric for the area under the PR curve as the average positive predictive value (AP) [18]. These previous research on the area under the PR curve have only considered binary outcomes. However, for many clinical applications, the outcome is time to event.

We make three contributions in the assessment of risk scoring systems for clinical screening. First, we define a time-dependent AP, AP_{t_0} for censored event time outcomes. We propose a robust nonparametric estimator of AP_{t_0} without modeling assumptions on the relationship between the

risk score and event time. Secondly, we provide a statistical inference procedure to compare the AP_{t_0} two risks scores regarding . Thirdly, we provide an R package to implement our method. The paper is organized as follows. In Section 2, we introduce the definition and interpretation of AP_{t_0} . In Section 3, we present the inference procedures for estimating AP_{t_0} of a single risk score as well as the comparison between two competing risk scores. In Section 4, we conduct a simulation study to investigate the performance of the proposed estimation and inference procedures in finite samples. In Section 5, we illustrate the proposed metric AP_{t_0} by analyzing two risk score systems with data from the Childhood Cancer Survival Study [20]. We conclude with a discussion and suggestions for future work in Section 6.

2 Time-dependent Average Positive Predictive Values

Consider a continuous risk score Z . Let T be the time to the event of interest. Time-dependent PPV and TPF [11, 21] are defined as

$$\text{PPV}_{t_0}(z) = \Pr\{T < t_0 \mid Z \geq z\} \quad \text{and} \quad \text{TPF}_{t_0}(z) = \Pr\{Z \geq z \mid T < t_0\}. \quad (2)$$

In the above setting, the event status is time-dependent, i.e., $D_{t_0} = I(T < t_0)$, where $I(\cdot)$ is an identity function. Consequently, the PPV and TPF are also functions of t_0 .

Following [18], we define AP_{t_0} , as the area under the time-dependent PR curve $\{(\text{TPF}_{t_0}(z), \text{PPV}_{t_0}(z)), z \in \mathcal{R}\}$,

$$AP_{t_0} = \int_{\mathcal{R}} \text{PPV}_{t_0}(z) d\text{TPF}_{t_0}(z). \quad (3)$$

Note that the TPF describes the distribution function of Z in "cases" who experience the event by time t_0 , i.e. $T < t_0$. It can be shown that $AP_{t_0} = E_{Z_1} \{\text{PPV}_{t_0}(Z_1)\}$, where Z_1 denotes the risk score in cases. In the real data example of Section 5, we will show that AP is estimated to be 0.114 at $t_0 = 35$ years for a risk score system. That is, by 35 years post diagnosis, we expect that on average 11.4% of the subjects with a high risk score (compared to the risk score of a randomly selected case) will experience the event of interest.

In addition, $\text{PPV}_{t_0}(z)$ can be written as $\text{PPV}_{t_0}(z) = P(Z \geq z \mid T < t_0)P(T < t_0)/P(Z \geq z) = \pi_{t_0} \{1 - F_1(z)\} / \{1 - F(z)\}$, where $F_1(z) = \Pr(Z < z \mid T < t_0) = P(Z_1 < z)$ is the distribution function of the risk score Z_1 for cases, $F(z) = P(Z < z)$ is the distribution function of the risk

score Z for the target population, and $\pi_{t_0} = Pr(T < t_0)$ is the event rate by time t_0 in the target population. Thus, the AP can be written as

$$AP_{t_0} = \pi_{t_0} \int_{\mathcal{R}} \frac{1 - F_1(z)}{1 - F(z)} dF_1(z). \quad (4)$$

A perfect risk score system would always assign higher values to cases, individuals with $T < t_0$, compared to those controls, individuals with $T \geq t_0$, i.e. $P(Z \geq Z_1 | T \geq t_0) = 0$. This leads to $AP_{t_0} = 1$ from equation (4). A non-informative risk score system would randomly assign risk scores to both cases and controls. i.e., $P(Z \geq z | T \geq t_0) = P(Z \geq z | T < t_0)$ for each z , which leads to $AP_{t_0} = \pi_{t_0}$. Thus, the theoretical range of AP_{t_0} is $[\pi_{t_0}, 1]$.

3 Estimating and Comparing AP_{t_0}

3.1 Nonparametric Estimator of AP_{t_0} for a single risk score

Often, the event times of some subjects are censored due to the end of the study or loss to follow up. Due to censoring, one can only observe $X = \min\{T, C\}$ where C is the censoring time, and $\delta = I(T < C)$. Let $\{(X_i, \delta_i, Z_i), i = 1, \dots, n\}$ be n independent realizations of (X, δ, Z) .

In the presence of censoring, event status at t_0 , $I(T_i < t_0)$, may not be observed for some subjects. We suggest using the inverse probability weighting (IPW) [22, 23] to account for censoring. The time-dependent PPV and TPF are estimated by

$$\widehat{PPV}_{t_0}(z) = \frac{\sum_{i=1}^n \widehat{w}_{t_0,i} I(Z_i \geq z) I(X_i < t_0)}{\sum_{i=1}^n I(Z_i \geq z)} \quad \text{and} \quad \widehat{TPF}_{t_0}(z) = \frac{\sum_{i=1}^n \widehat{w}_{t_0,i} I(Z_i \geq z) I(X_i < t_0)}{\sum_{i=1}^n \widehat{w}_{t_0,i} I(X_i < t_0)},$$

where $\widehat{w}_{t_0,i}$ is the inverse of the estimated probability that the time-dependent event status $I(T_i < t_0)$ is observed, specifically

$$\widehat{w}_{t_0,i} = \frac{I(X_i < t_0)\delta_i}{\widehat{\mathcal{G}}(X_i)} + \frac{I(X_i \geq t_0)}{\widehat{\mathcal{G}}(t_0)},$$

where $\widehat{\mathcal{G}}(c)$ is a consistent estimator of the survival function of the censoring time, $\mathcal{G}(c) = Pr(C \geq c)$. Note that the proposed estimator does not impose any assumptions on the relationship between the risk score Z and the event time T . Under the assumption of independent censoring, i.e., the censoring time C is independent of both the event time T and the risk score Z , $\mathcal{G}(c)$ can

be obtained by the nonparametric Nelson-Aalen or Kaplan-Meier estimator. If the censoring time C depends on the risk score Z , additional model assumptions might be required. For example, a proportional hazards (PH) model could be fit to estimate $\mathcal{G}_z(t) = Pr(C \geq c \mid Z = z)$.

Based on the estimated $PPV_{t_0}(z)$ and $TPF_{t_0}(z)$, AP_{t_0} can be estimated by

$$\widehat{AP}_{t_0} = \frac{\sum_{j=1}^n I(X_j \leq t_0) \widehat{w}_{t_0,j} \widehat{PPV}_{t_0}(Z_j)}{\sum_{j=1}^n I(X_j \leq t_0) \widehat{w}_{t_0,j}}. \quad (5)$$

[22] shows that $\widehat{PPV}_{t_0}(z)$ and $\widehat{TPF}_{t_0}(z)$ are both consistent estimators, and thus, \widehat{AP}_{t_0} is also a consistent estimator of AP_{t_0} for any given value of t_0 .

In practice, we often deal with discrete risk scores, where tied risk scores are common. To accommodate risk scores with ties, following [24], we modify the above estimator (5) by replacing $\widehat{PPV}_{t_0}(Z_j)$ with

$$\widetilde{PPV}_{t_0}(Z_j) = \frac{\sum_{i=1}^n \widehat{w}_{t_0,i} \{I(Z_i > Z_j) + \frac{1}{2}I(Z_i = Z_j)\} I(X_i < t_0)}{\sum_{i=1}^n \{I(Z_i > Z_j) + \frac{1}{2}I(Z_i = Z_j)\}}.$$

To construct confidence intervals, we suggest the nonparametric bootstrap [25] method. Specifically, let $\widehat{AP}_{t_0}^{\text{B}} = \{\widehat{AP}_{t_0}^b, b = 1, 2, \dots, B\}$ denote the estimated AP_{t_0} obtained from B bootstrap resamples. A 95% confidence interval (CI) for the AP_{t_0} is given as $(\widehat{AP}_{t_0}^{\text{B},0.025}, \widehat{AP}_{t_0}^{\text{B},0.975})$, where $\widehat{AP}_{t_0}^{\text{B},0.025}$ and $\widehat{AP}_{t_0}^{\text{B},0.975}$ are the 2.5% and 97.5% empirical percentiles of the $\widehat{AP}_{t_0}^{\text{B}}$, respectively.

3.2 Comparing two risk scores

We consider comparing two risk scores Z_1 and Z_2 on the AP_{t_0} scale. In many studies, both risk scores Z_1 and Z_2 are calculated for each individual. With the paired data, we can quantify the relative predictive performance of Z_1 vs. Z_2 , using the ratio of their respective time-dependent AP,

$$rAP_{t_0} = AP_{Z_1,t_0} / AP_{Z_2,t_0},$$

where AP_{Z_1,t_0} and AP_{Z_2,t_0} denote the time-dependent AP for Z_1 and Z_2 at t_0 respectively. For a single risk score Z , the ratio $AP_{Z,t_0} / \pi_{t_0}$ can be regarded as the relative predictive performance of Z compared to a non-informative risk score.

The AP ratio rAP_{t_0} can be estimated by $r\widehat{AP}_{t_0} = \widehat{AP}_{Z_1,t_0} / \widehat{AP}_{Z_2,t_0}$, where \widehat{AP}_{Z_1,t_0} and \widehat{AP}_{Z_2,t_0} are

the the nonparametric estimator \widehat{AP}_{t_0} in (5) of Z_1 and Z_2 respectively. The standard percentile bootstrap method described in Section 3.1 can be used to construct a CI for rAP_{t_0} or test $H_0 : rAP_{t_0} = 1$ for any given time point t_0 . Specifically, the CI could be obtained based on the empirical percentiles of the B bootstrap counterparts of $r\widehat{AP}_{t_0}$, denoted by $r\widehat{AP}_{t_0}^b = \widehat{AP}_{Z_1, t_0}^b / \widehat{AP}_{Z_2, t_0}^b$, where $\widehat{AP}_{Z_1, t_0}^b$ and $\widehat{AP}_{Z_2, t_0}^b$ are the estimated AP_{t_0} for Z_1 and Z_2 based on the same bootstrap resample, $b = 1, \dots, B$.

4 Simulation study

We conducted a simulation study to examine the performance of the time-dependent AP estimator in finite samples. In this simulation study, we considered two risk scores U_1 and U_2 . They were generated from a standard normal distribution $N(0, 1)$. The event time associated with both risk scores for the i th subject was generated from the following model

$$\log(T_i) = 7.2 - 1.1U_{i1} - 2.5U_{i2} - 1.5\log(U_{i1}^2) + \epsilon_T,$$

where $\epsilon_T \sim N(0, 1.5)$. This setting provides an example where the ROC curves of the two risk scores cross at time $t_0 = 8$, shown in Figure 1, with AUC_{U_1, t_0} and AUC_{U_2, t_0} are similar in values. On the other hand, the PR curve of U_1 dominates that of U_2 over the most range of the TPR with AP_{t_0} of U_1 greater than that of U_2 .

The censoring time C_i was generated following $C_i = \min(A_i, B_i + 1)$ where, $A_i \sim Uniform(0, 50)$, and $B_i \sim Gamma(25, 0.75)$. This configuration results in about 50% of censoring overall. Let $X_i = \min(T_i, C_i)$, $\delta_i = I(T_i \leq C_i)$. In this setting, the censoring time is independent of both the event time and risk scores.

We considered three prediction time points t_0 where the corresponding event rates are, $r = P(T_i < t_0)$, 0.01, 0.05 and 0.1, respectively. To allow a reasonable number of events by t_0 , we generated the data $\{(X_i, \delta_i, U_{1i}, U_{2i}), i = 1, \dots, n\}$ with sample size n being 2000 and 5000 (Tables 1 and 2). In each table, we report the following summary statistics based on 1000 repetitions: bias, empirical standard error (ESE) of the estimator, average standard errors from bootstrap (ASE^b), and the empirical coverage probability ($ECOV^b$) of 95% confidence intervals obtained from 1000 bootstrap resamples as described in Section 3.




Figure 1: The ROC curves in the left panel and the precision-recall curves in the right panel for the two risk scores U_1 and U_2 at $t_0 = 8$ when the event rate is 5%. The numbers shown in graph correspond to the AUC and the AP values.

These results show that the estimators of both time-dependent AP and ratio of AP, AP_1/AP_2 , between the two risk scores had small biases in all t_0 values and different sample sizes. The bias decreases with increasing event rate and increasing sample size. Also, the standard errors ASE^b obtained from bootstrap were close to the empirical standard errors. Thus, the confidence intervals attained the nominal coverage probabilities for both smaller sample size 2000 and larger sample size 5000.

We remark that this simulation provides an illustrative example of the relationship between ROC curve and PR curve as well as the relationship between the AUC and the AP investigated in [16]. When the ROC curves of two competing risk scores cross, the PR curves cross too. In situations like this, the AUC and the AP may rank the risk scores differently. In our simulation setting, U_2 outperforms U_1 according to the AUC, which indicates that U_2 is better at discriminating between cases and controls. On the other hand, U_1 outperforms U_2 according to the AP, which suggests that U_1 is a better screening tool for stratifying subjects into different risk groups.

5 Data Analysis

In this section, we illustrate the use of AP_{t_0} metric with a data set from the Childhood Cancer Survivor Study [20]. This cohort follows children who were initially treated for cancer at 26 US

Table 1: Results of simulation with sample size 2000.

t_0	Event rate		TRUE	BIAS	ESE	ASE^b	$ECOVp^b(\%)$
0.5	0.0101	AP ₁	0.182	0.0365	0.0810	0.0795	92.3
		AP ₂	0.124	0.0339	0.0689	0.0678	93.0
		rAP	1.47	0.4890	1.5300	1.7600	95.1
8	0.0495	AP ₁	0.364	0.0096	0.0527	0.0516	92.5
		AP ₂	0.266	0.0129	0.0452	0.0450	93.4
		rAP	1.37	0.0140	0.3290	0.3320	95.7
36	0.0991	AP ₁	0.462	0.0098	0.0534	0.0558	95.9
		AP ₂	0.375	0.0118	0.0493	0.0501	94.5
		rAP	1.23	0.0135	0.2310	0.2420	94.9

Table 2: Results of simulation with sample size 5000.

t_0	Event rate		TRUE	BIAS	ESE	ASE^b	$ECOVp^b(\%)$
0.5	0.0101	AP ₁	0.182	0.0185	0.0500	0.0504	93.1
		AP ₂	0.124	0.0155	0.0416	0.0417	94.8
		rAP	1.47	0.1550	0.7060	0.7600	93.8
8	0.0495	AP ₁	0.364	0.0042	0.0337	0.0333	92.9
		AP ₂	0.266	0.0049	0.0291	0.0288	93.7
		rAP	1.37	0.0060	0.2160	0.2100	95.4
36	0.0991	AP ₁	0.462	0.0034	0.0354	0.0346	95.5
		AP ₂	0.375	0.0037	0.0310	0.0313	94.1
		rAP	1.23	0.0051	0.1490	0.1510	95.0

and Canada institutions between 1970 and 1986 and who survived at least 5 years after their cancer diagnosis. Among the survivors, cardiovascular disease has been recognized as a leading contributor to morbidity and mortality [26]. To inform future screening and intervention strategy for congestive heart failure (CHF) in this population, [6] developed several risk score systems using the CCSS data and validated them on external cohorts. For the purpose of illustration, we chose two of these risk scores and evaluated their predictive performance using the proposed AP_{t_0} .

We included 11,457 subjects in our analysis from the CCSS study who met the original study inclusion criteria and had both risk scores. In this data, a total of 248 subjects experienced the CHF. Between the two risk scoring systems we focused on in this data analysis, the simpler model used information on age at cancer diagnosis, sex, whether the patient was exposed to chest radiotherapy, and whether the patient was exposed to a particular chemotherapy agent. We refer to this model as the simple model. The more elaborate model, known as the heart dose model, in-

cluded detailed clinical information on the average radiation dose to the heart and the cumulative dose of the specific chemotherapy agent, along with age at diagnosis and sex. This is an example where a simple risk score system utilizes minimum treatment information and can be used for any patient by virtually all clinicians, while the more complex risk score system demands specific dose information which may not be readily available to clinicians providing long-term follow-up care. We used the original risk scores of the simple model and the heart dose model from the reference study [6]. Briefly, these scores were constructed via linear combinations of the corresponding covariates, where the regression coefficients were obtained from Poisson regression models.

Table 3 reports the estimated AP_{t_0} with 95% CIs for both the simple model (denoted by AP_{s,t_0}) and heart dose model (denoted by AP_{h,t_0}) at $t_0 = 20$ and 35 years post-diagnosis where the corresponding estimated event rates were 1.3% and 4.7% respectively. These two models were compared using the ratio of AP, i.e. $rAP_{t_0} = AP_{h,t_0} / AP_{s,t_0}$. In addition, we also provided the estimated time-dependent AUC (AUC_{t_0}) at these two time points as well as the difference of AUCs between these two models $\Delta AUC_{t_0} = AUC_{h,t_0} - AUC_{s,t_0}$. To illustrate the time-varying performance for each model as well as the comparison between these two models over time, AP_{t_0} , AUC_{t_0} , rAP_{t_0} and ΔAUC_{t_0} versus t_0 were plotted in Figure 2. Note that we assumed independent censoring in estimating the AP and the AUC.

The results in Table 3 show that the heart dose model outperforms the simple model at both time points. For example, the estimated AP_{20} of the heart dose model is 0.075, which indicates that by 20 years post-diagnosis, using the risk score from the heart dose model, we expect that on average 7.5% subjects with a high risk score (compared to the risk score of a randomly selected case) will experience heart failure. This AP is almost six times of the event rate 1.3%, which corresponds to the AP of a non-informative risk score system. In contrast, the estimated AP_{20} for the simple model is 0.039, roughly half of that of the heart dose model ($rAP_{20}=1.96$, 95%CI:1.41-2.88). At 35 years post diagnosis, the heart dose model is significantly better than the simple model with $rAP_{35} = 1.45$ (95%CI: 1.26 - 1.70). Indeed, The plots (c) and (d) in Figure 2 show that in terms of the AP_{t_0} , the heart dose model outperforms the simple model at identifying the high risk subjects from the targeted population at all time points considered. On the other hand, the AUCs are similar between these two models, as seen in plot (d). Especially, ΔAUC is not significantly different from 0 at the beginning and towards the end of the time period that were considered,

For example $\Delta\text{AUC}_{35} = 0.01$ (95% CI: -0.02 - 0.03, $p=0.47$), shown in Table 3. It suggests that according to the AUC, there is not much performance difference between the two models. If, due to incorporating more information, the heart dose model is indeed superior to the simple model in terms of identifying the high risk individuals, the results in Table 3 and Figure 2 implies that the AP is a better metric for discriminates the risk prediction performance than the AUC does.

Table 3: Estimated AP_{t_0} and AUC_{t_0} with 95% CIs for two risk scoring systems at $t_0 = 20$ and 35 years, respectively. The comparison of APs is measured by rAP and the comparison of AUCs is measured by ΔAUC .

t_0	Event rate	Risk score system	AP	AUC
20 years	0.013	Simple	0.039 (0.029, 0.052)	0.79 (0.75, 0.82)
		Heart dose	0.077 (0.050, 0.127)	0.82 (0.78, 0.86)
		Comparison	1.96 (1.41, 2.88)	0.03 (0.01, 0.05)
35 years	0.047	Simple	0.079 (0.066, 0.095)	0.81 (0.78, 0.84)
		Heart dose	0.114 (0.092, 0.144)	0.82 (0.78, 0.85)
		Comparison	1.45 (1.26, 1.70)	0.01 (-0.02, 0.03)

6 Discussion

One of the main goals of clinical risk prediction is to screen the asymptomatic population and to stratify them for tailored intervention. *Prospective* accuracy measures such as PPV_{t_0} is preferred for this purpose. However, the calculation of PPV_{t_0} demands a threshold for continuous risk scores, which can create practical difficulties for evaluating risk score systems, especially when more than two systems are compared. In this paper, we defined and interpreted AP_{t_0} , which is the area under the time-dependent precision-recall curve, for event time data. We proposed a nonparametric estimator of AP_{t_0} and a ratio estimator of AP_{t_0} for comparing two competing risk score systems. We suggested the use of the bootstrap method for inference, which is broadly applicable in practical settings. We also developed an R package `APtools` for download available in CRAN which implements our method for binary and survival outcomes.

The AUC has been the most widely used performance metric in the clinical research community. A number of authors have pointed out that the AUC is informative on the classification performance and discrimination power [27, 11, 18], but not an appropriate metric for assessing the prospective accuracy performance [10, 11]. Consistent with the criticism on the insensitivity

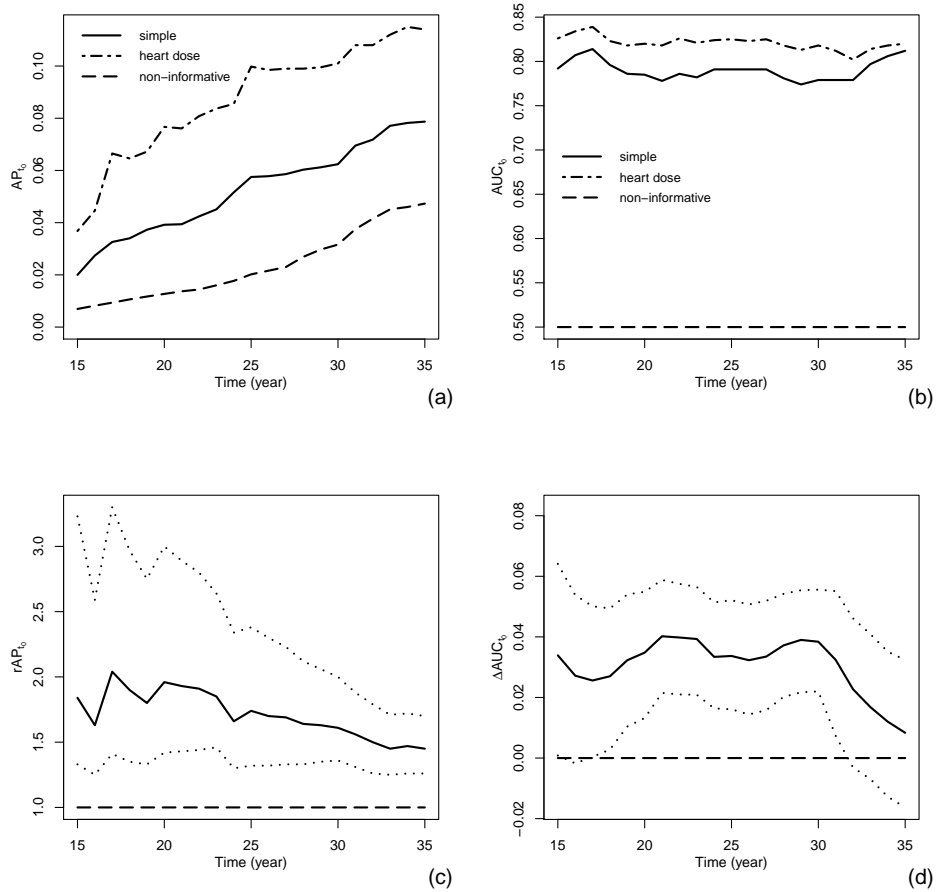


Figure 2: CCSS Data analysis: panel (a) shows the estimates of the time-dependent AP for the simple model AP_{s,t_0} and heart dose model AP_{h,t_0} ; panel (b) shows the estimates of the time-dependent AUC for the simple model AUC_{s,t_0} and heart dose model AUC_{h,t_0} ; panel (c) shows the estimates of $rAP_{t_0} = AP_{h,t_0} / AP_{s,t_0}$, the ratio of the time-dependent AP of the heart dose model over that of the simple model; panel (d) shows the estimates of $\Delta AUC_{t_0} = AUC_{h,t_0} - AUC_{s,t_0}$, the difference of the time-dependent AUC between the heart dose model and the simple model. The dotted lines in panels (c) and (d) represent the pointwise 95%CI for rAP_{t_0} and ΔAUC_{t_0} , respectively.

of the AUC in evaluating risk prediction models (See [9]), our data analysis illustrated that using the AUC as the metric, the performance of the simple model and the heart dose model appears close. However, based on the AP, the heart dose model outperforms the simple model and could be preferred in clinical screening. Thus the conclusion reached based on the AUC in this example may mislead researchers and clinicians.

It should be noted that when comparing different risk score systems, the ranking of their AUCs

and APs are not necessarily concordant; our simulation study in Section 4 gives such an example. As the purpose of this article is to introduce a time-dependent AP_{t_0} , we refer readers to [16] for insight on the relationship between the ROC curve and PR curve and to [17] where the relationship between the AUC and the AP is illuminated. Thus, risk scores which perform well in separating cases from controls may perform poorly in identifying a higher risk subpopulation, which is the goal of screening. Compared to the AUC, the AP as the summary metric of PPV is better suited in evaluating the usefulness of the risk scores and comparing the *prospective* prediction performance among competing risk scores, when the objective is *screening* through risk stratification.

[11, 13] proposed to use the curves of PPV_{t_0} versus risk score quantiles as an assessment tool for quantifying prospective prediction accuracy. One curve corresponds to one particular value of t_0 , which limits its ability to assess the accuracy across time points. In contrast, plotting AP_{t_0} against time could facilitate visualizing the performance of different risk score systems over time in one single plot.

Unlike the AUC, the AP is event rate dependent and should be estimated in a prospective cohort or population-based study. AP cannot be estimated from a case-control study; the estimate will be of very little use because the prevalence rate is artificially fixed by the study design. While the range of the AUC is always between 0.5 and 1, the range of AP is between the event rate and 1. While AP's wide range could be advantageous in differentiating risk score systems, caution is needed when re-evaluating risk score systems in other study populations for the same outcome. This is because the underlying event rate may differ among populations. Thus, it is possible that AP will select different risk score systems as superior for the same outcome in different study populations.

For future work, we will consider estimating the time-dependent AP with multiple markers, as well as the incremental value of AP by adding new markers on top of an existing risk profile. In addition, similar to the partial AUC, partial AP could be defined as the area over a certain range of interest, such as those at the low values of TPF where PPV is typically high. Competing risk in the context of AP is another topic that needs to be addressed for event time data.

Acknowledgments

The authors would like to thank Yan Chen for data support. Dr. Zhou's research is supported by the Natural Sciences and Engineering Research Council of Canada. Dr. Yuan's research is supported by the Canadian Institutes of Health Research. The CCSS was supported by the National Cancer Institute (CA55727, G.T. Armstrong, Principal Investigator).

References

- [1] Wright CF, Zimmern RL. Conceptual issues for screening in the genomic era-time for an update? *Epidemiology, biostatistics, and public health* 2014; **11**(4).
- [2] Espeland M. Reduction in weight and cardiovascular disease risk factors in individuals with type 2 diabetes: one-year results of the look ahead trial. *Diabetes care* 2007; .
- [3] James MT, Hemmelgarn BR, Tonelli M. Early recognition and prevention of chronic kidney disease. *The Lancet* 2010; **375**(9722):1296–1309.
- [4] Goff DC, Lloyd-Jones DM, Bennett G, Coady S, D'Agostino RB, Gibbons R, Greenland P, Lackland DT, Levy D, O'Donnell CJ, *et al.*. 2013 acc/aha guideline on the assessment of cardiovascular risk: a report of the american college of cardiology/american heart association task force on practice guidelines. *Journal of the American College of Cardiology* 2014; **63**(25_PA).
- [5] Buijsse B, Simmons RK, Griffin SJ, Schulze MB. Risk assessment tools for identifying individuals at risk of developing type 2 diabetes. *Epidemiologic reviews* 2011; :mxq019.
- [6] Chow EJ, Chen Y, Kremer LC, Breslow NE, Hudson MM, Armstrong GT, Border WL, Feijen EA, Green DM, Meacham LR, *et al.*. Individual prediction of heart failure among childhood cancer survivors. *Journal of Clinical Oncology* 2015; **33**(5):394–402.
- [7] Lindström J, Tuomilehto J. The diabetes risk score. *Diabetes care* 2003; **26**(3):725–731.
- [8] Grimes DA, Schulz KF. Uses and abuses of screening tests. *The Lancet* 2002; **359**(9309):881–884.
- [9] Cook NR. Use and misuse of the receiver operating characteristic curve in risk prediction. *Circulation* 2007; **115**(7):928–935.

- [10] Moskowitz CS, Pepe MS. Quantifying and comparing the predictive accuracy of continuous prognostic factors for binary outcomes. *Biostatistics* 2004; **5**(1):113–127.
- [11] Zheng Y, Cai T, Pepe MS, Levy WC. Time-dependent predictive values of prognostic biomarkers with failure time outcome. *Journal of the American Statistical Association* 2008; **103**(481):362–368.
- [12] Wald N, Bestwick J. Is the area under an roc curve a valid measure of the performance of a screening or diagnostic test? *Journal of medical screening* 2014; :0969141313517 497.
- [13] Zheng Y, Cai T, Stanford JL, Feng Z. Semiparametric models of time-dependent predictive values of prognostic biomarkers. *Biometrics* 2010; **66**(1):50–60.
- [14] Raghavan V, Bollmann P, Jung GS. A critical investigation of recall and precision as measures of retrieval system performance. *ACM Transactions on Information Systems (TOIS)* 1989; **7**(3):205–229.
- [15] Manning CD, Schütze H. *Foundations of statistical natural language processing*, vol. 999. MIT Press, 1999.
- [16] Davis J, Goadrich M. The relationship between precision-recall and roc curves. *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, ACM: New York, NY, USA, 2006; 233–240, doi:10.1145/1143844.1143874. URL <http://doi.acm.org/10.1145/1143844.1143874>.
- [17] Su W, Yuan Y, Zhu M. A relationship between the average precision and the area under the roc curve. *Proceedings of the 2015 International Conference on The Theory of Information Retrieval*, ACM, 2015; 349–352.
- [18] Yuan Y, Su W, Zhu M. Threshold-free measures for assessing the performance of medical screening tests. *Frontiers in Public Health* 2015; **3**.
- [19] Ozenne B, Subtil F, Maucort-Boulch D. The precision–recall curve overcame the optimism of the receiver operating characteristic curve in rare diseases. *Journal of clinical epidemiology* 2015; **68**(8):855–859.

- [20] Robison LL, Armstrong GT, Boice JD, Chow EJ, Davies SM, Donaldson SS, Green DM, Hammond S, Meadows AT, Mertens AC, *et al.*. The childhood cancer survivor study: a national cancer institute–supported resource for outcome and intervention research. *Journal of Clinical Oncology* 2009; **27**(14):2308–2318.
- [21] Heagerty PJ, Lumley T, Pepe MS. Time-dependent ROC curves for censored survival data and a diagnostic marker. *Biometrics* 2000; :337–344.
- [22] Uno H, Cai T, Tian L, Wei L. Evaluating prediction rules for t-year survivors with censored regression models. *Journal of the American Statistical Association* 2007; **102**(478).
- [23] Lawless JF, Yuan Y. Estimation of prediction error for survival models. *Statistics in medicine* 2010; **29**(2):262–274.
- [24] Pepe MS. *The statistical evaluation of medical tests for classification and prediction*. Oxford University Press, USA, 2003.
- [25] Efron B. Bootstrap methods: another look at the jackknife. *The Annals of Statistics* 1979; **7**(1):1–26.
- [26] Oeffinger KC, Mertens AC, Sklar CA, Kawashima T, Hudson MM, Meadows AT, Friedman DL, Marina N, Hobbie W, Kadan-Lottick NS, *et al.*. Chronic health conditions in adult survivors of childhood cancer. *New England Journal of Medicine* 2006; **355**(15):1572–1582.
- [27] Gail MH, Pfeiffer RM. On criteria for evaluating models of absolute risk. *Biostatistics* 2005; **6**(2):227–239.