

Cohort heterogeneity and competing risk in cancer survival

ACC Coolen
King's College London



Umeå, May 10th 2017

Cohort heterogeneity and competing risks

- Regression for time-to-event data

- Consequences and fingerprints

- Informative censoring: the intuition

Bayesian latent class models

- Rationale and definition

- Tests on synthetic data

Applications in cancer research

- Epidemiological cancer data

- Data from failed cancer trials

New directions

- Prospective latent class prediction

- Overfitting in multivariate survival analysis

Summary

Cohort heterogeneity and competing risks

- Regression for time-to-event data

- Consequences and fingerprints

- Informative censoring: the intuition

Bayesian latent class models

- Rationale and definition

- Tests on synthetic data

Applications in cancer research

- Epidemiological cancer data

- Data from failed cancer trials

New directions

- Prospective latent class prediction

- Overfitting in multivariate survival analysis

Summary

Regression for time-to-event data


► Data

$$D = \{(\mathbf{z}_1, t_1, r_1), \dots, (\mathbf{z}_N, t_N, r_N)\}$$

$\mathbf{z}_i = (z_{i1}, \dots, z_{id})$: d covariates (measured at $t = 0$)

$t_i > 0$: first failure time (death, onset of disease, ...)

$r_i \in \{0, 1, \dots, R\}$: failure type (or 'risk')



► Heterogeneity

visible: variability in the available covariates

latent: variability in host or disease, not visible in the covariates
(individuals with same covariates \mathbf{z} are not clones ...)

Competing risks, identifiability and interpretation



► *Competing risks*

Informative censoring, i.e. event times of risks

are *statistically dependent*: $p(t_1, \dots, t_R | \mathbf{z}) \neq \prod_{r=1}^R p(t_r | \mathbf{z})$

reported time: $t = \min\{t_1, \dots, t_R\}$

► *Interpretation of crude hazard rates*

Eliminating one risk can change hazard rate of others ...

if hazard rate for risk 1 is low:

- (i) event 1 is intrinsically unlikely?
- (ii) or it is often preceded by event 2?

to disentangle risks: need $p(t_1, \dots, t_R | \mathbf{z})$

► *Tsiatis' identifiability problem (1975)*

Joint event time distribution $p(t_1, \dots, t_R | \mathbf{z})$

cannot be inferred from survival data alone ...

Conventional methods

for analysing time-to-event data

Kaplan-Meier estimators

Cox regression

.....

- ▶ not designed to handle disease/host heterogeneity, beyond variability in covariates
- ▶ to allow interpretation:
have to assume different risks are uncorrelated,
dangerous when many censoring events ...

random effects models, frailty models,
latent class models

- ▶ usually constructed for primary risk only,
so still cannot handle correlated risks
- ▶ do not exploit the link between latent heterogeneity
and competing risks ...

Consequences and fingerprints of latent heterogeneity



- ▶ *Violation of proportional hazards assumption*

- ▶ *Interpretation of time dependencies tricky*

even if all *individual* hazard rates h_i are time-independent, cohort hazard rate will be time-dependent:

$$h(t) = \frac{\sum_{i=1}^n h_i e^{-h_i t}}{\sum_{i=1}^n e^{-h_i t}}$$

- ▶ *Interpreting cause-specific survival curves (KM, Cox) no longer possible ...*




If in interpreting our data we assume
censoring risks uncorrelated with primary risk

censoring by competing risks
can give nonsensical results ...

- harmful drugs look beneficial
- beneficial drugs look harmful
- false protectivity of covariates

(ULSAM prostate cancer data)




would we have spotted this
problem if the covariate represented
the expression of a specific gene?

Link between cohort heterogeneity and informative censoring




Say 1000 people,
two risks, hazard rates h_A and h_B

- ▶ homogeneous cohort:
all *individuals* have (h_A, h_B)



- ▶ heterogeneous cohort,
but *non-informative censoring*




Heterogeneity and informative censoring




Say 1000 people,
two risks, hazard rates h_A and h_B

- ▶ homogeneous cohort:
all *individuals* have (h_A, h_B)



- ▶ heterogeneous cohort,
informative cohort filtering
result:
underestimation of h_A



Heterogeneity and informative censoring




Say 28 people,
binary covariate: $z=0,1$

associations risk **A**: β_A

associations risk **B**: β_B


(B: competing risk, strong)



without risk B:

as many **A** events for
 $z=0$ as for $z=1$,

overall association $\beta_A = 0$




Heterogeneity and informative censoring




Say 28 people,
binary covariate: $z = 0, 1$

associations risk **A**: β_A
associations risk **B**: β_B
(B: competing risk, strong)



Effect of risk B:



*what will we now
observe for risk **A**?*

Heterogeneity and informative censoring




what will we now
observe for risk A ?

A cases with $z=0$: 4

A cases with $z=1$: 1

overall association $\beta_A < 0$,

false protectivity



Cohort heterogeneity and competing risks

- Regression for time-to-event data

- Consequences and fingerprints

- Informative censoring: the intuition

Bayesian latent class models

- Rationale and definition

- Tests on synthetic data

Applications in cancer research

- Epidemiological cancer data

- Data from failed cancer trials

New directions


- Prospective latent class prediction

- Overfitting in multivariate survival analysis

Summary

Bayesian latent class methods: rationale and definition

- ▶ model all risks simultaneously
- ▶ individuals with *same* covariates can have *distinct* associations and *distinct* base hazard rates
- ▶ risks are assumed independent only *at the level of individuals* (this removes Tsiatis' identifiability problem)
- ▶ competing risks, informative censoring:
reflect correlated association parameters of different risks



proportional hazards within classes \nRightarrow proportional hazards at cohort level
independent risks within classes \nRightarrow independent risks at cohort level

Personalised cause-specific hazard rate model variants		
$M = 1$	Heterogeneous frailties	$h_r^i(t) = \lambda_r(t)e^{\beta_r^{\ell 0} + \sum_{\mu} \beta_r^{\mu} z_i^{\mu}}$
	Homogeneous associations	
	Homogeneous base hazard rates	
$M = 2$	Heterogeneous frailties	$h_r^i(t) = \lambda_r(t)e^{\beta_r^{\ell 0} + \sum_{\mu} \beta_r^{\ell \mu} z_i^{\mu}}$
	Heterogeneous associations	
	Homogeneous base hazard rates	
$M = 3$	Heterogeneous frailties	$h_r^i(t) = \lambda_r^{\ell}(t)e^{\beta_r^{\ell 0} + \sum_{\mu} \beta_r^{\ell \mu} z_i^{\mu}}$
	Heterogeneous associations	
	Heterogeneous base hazard rates	

- Bayesian analysis and model selection:
reliable error bars, and multiple classes *only if data demand it*
- reduces to standard Cox regression if no heterogeneity
(Occam's Razor action of Bayesian model selection)
- non-primary events all contribute to latent class inference
- fully transparent interpretation,
unlike some other competing risk approaches ...
- formulae for survival curves *decontaminated* for informative censoring,
and *retrospective class allocation* of individuals

(Rowley et al, SIM, 2017)

Technicalities ...

- ▶ *censoring*
modelled as 'risk' $r=0$ with no associations
- ▶ *data likelihood*

$$p(t, r | \mathbf{z}) = \sum_{\ell=1}^L w_{\ell} p(t, r | \mathbf{z}, \ell), \quad p(t, r | \mathbf{z}, \ell) = \lambda_r^{\ell}(t) e^{\boldsymbol{\beta}_r^{\ell} \cdot \mathbf{z} - \Lambda_0(t) - \sum_{r'=1}^R \exp(\boldsymbol{\beta}_{r'}^{\ell} \cdot \mathbf{z}) \Lambda_{r'}^{\ell}(t)}$$

- ▶ *base rates*
spline construction for $\{\lambda_r^{\ell}(t)\}$, with K spline points
- ▶ *Bayesian model selection*
 K : baserate complexity
 L : number of latent classes
 M : heterogeneity complexity
- ▶ *numerical implementation*
curvature estimation near parameter boundaries ...
avoiding local minima in high-dim searches ...
CPU efficiency ...

Upon determining parameters and hyper-parameters

explicit formulae for e.g.

- ▶ covariate-conditioned survival curves and hazard rates:

$$\text{crude :} \quad h_r(t|\mathbf{z}) = \frac{\sum_{\ell} w_{\ell} \lambda_r^{\ell}(t) e^{\boldsymbol{\beta}_r^{\ell} \cdot \mathbf{z} - \sum_{r'=1}^R \exp(\boldsymbol{\beta}_{r'}^{\ell} \cdot \mathbf{z}) \Lambda_{r'}^{\ell}(t)}}{\sum_{\ell} w_{\ell} e^{-\sum_{r'=1}^R \exp(\boldsymbol{\beta}_{r'}^{\ell} \cdot \mathbf{z}) \Lambda_{r'}^{\ell}(t)}},$$

$$\text{decontaminated :} \quad \tilde{h}_r(t|\mathbf{z}) = \frac{\sum_{\ell} w_{\ell} \lambda_r^{\ell}(t) e^{\boldsymbol{\beta}_r^{\ell} \cdot \mathbf{z} - \exp(\hat{\boldsymbol{\beta}}_r^{\ell} \cdot \mathbf{z}) \Lambda_r^{\ell}(t)}}{\sum_{\ell} w_{\ell} e^{-\exp(\hat{\boldsymbol{\beta}}_r^{\ell} \cdot \mathbf{z}) \Lambda_r^{\ell}(t)}}.$$

- ▶ cause-specific cumulative incidence function:

$$F_r(t|\mathbf{z}) = \int_0^t dt' e^{-\Lambda_0(t')} \sum_{\ell} w_{\ell} \lambda_r^{\ell}(t') e^{\boldsymbol{\beta}_r^{\ell} \cdot \mathbf{z} - \sum_{r'=1}^R \exp(\boldsymbol{\beta}_{r'}^{\ell} \cdot \mathbf{z}) \Lambda_{r'}^{\ell}(t')}.$$

- ▶ class membership probabilities:

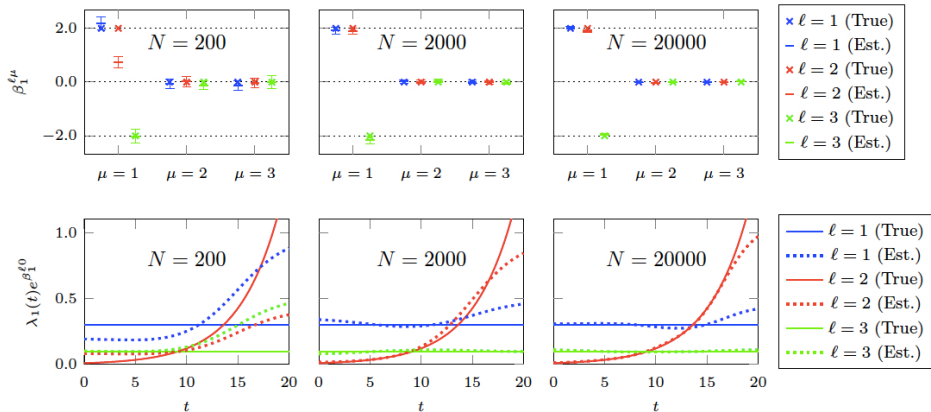
$$p(\ell|t, r, \mathbf{z}) = \frac{w_{\ell} p(t, r|\mathbf{z}, \ell)}{\sum_{\ell'=1}^L w_{\ell'} p(t, r|\mathbf{z}, \ell')}$$

Tests on synthetic data

inference of classes and parameters

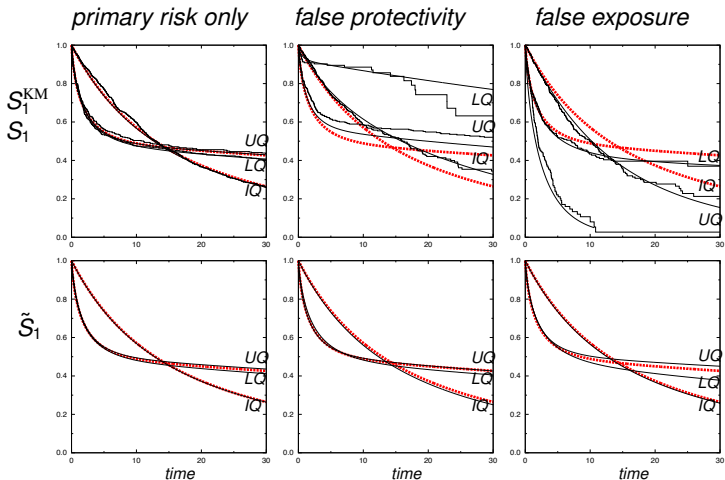
3 classes:

red, blue, green



Tests on synthetic data

decontaminating survival curves for informative censoring



S_1^{KM} : Kaplan-Meier
 S_1 : crude survival curve

red dashed: true survival curves
 \tilde{S}_1 : decontaminated curves

Cohort heterogeneity and competing risks

- Regression for time-to-event data

- Consequences and fingerprints

- Informative censoring: the intuition

Bayesian latent class models

- Rationale and definition

- Tests on synthetic data

Applications in cancer research

- Epidemiological cancer data

- Data from failed cancer trials

New directions

- Prospective latent class prediction

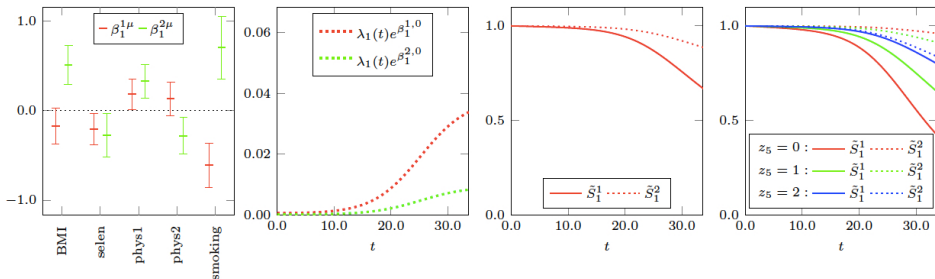
- Overfitting in multivariate survival analysis

Summary

Prostate cancer data

(ULSAM data base, $n = 2047$)

Cox regression:
smoking is protective against PC



negative association with smoking *only* in
extremely frail subgroup of patients

red class: high overall frailty
green class: low overall frailty


(Rowley et al, SIM, 2017)

Breast cancer data


(AMORIS data base, $N = 1798$)

Cox regression finds no significant associations
(proportional hazards violated)



BC death



CV death



*solid: BC
dashed: CV*



red class: predominantly younger women

green class: predominantly older women

(Wulaningsih et al,
BMC Cancer 2015)

Applications to failed cancer trials



- ▶ *failed clinical trials*

often some drug benefit, but not enough in view of costs ...
(in the absence of a biomarker to select patients)

- ▶ *two possibilities*

1. there exist measurable differences between individuals that explain response variation, we just don't know what they are ...
subgroups with distinct quantitative characteristics,
cohort is in principle **stratifiable**
2. there are no measurable differences between individuals to explain response variation: cohort **not stratifiable**

- ▶ *Bayesian Latent class analysis*

- rational method for determining whether cohort is stratifiable
- retrospective class assignment: tool for identifying latent classes

Addition of cetuximab to oxaliplatin-based first-line combination chemotherapy for treatment of advanced colorectal cancer: results of the randomised phase 3 MRC COIN trial



Timothy S Maughan, Richard A Adams, Christopher G Smith, Angela M Meade, Matthew T Seymour, Richard H Wilson, Shelley Idziaszczyk, Rebecca Harris, David Fisher, Sarah L Kenny, Edward Kay, Jenna K Mitchell, Ayman Madi, Bharat Jasani, Michelle D James, John Bridgewater, M John Kennedy, Bart Claes, Diether Lambrechts, Richard Kaplan, Jeremy P Cheadle, on behalf of the MRC COIN Trial Investigators

Summary

Background In the Medical Research Council (MRC) COIN trial, the epidermal growth factor receptor (EGFR)-targeted antibody cetuximab was added to standard chemotherapy in first-line treatment of advanced colorectal cancer with the aim of assessing effect on overall survival.

Lancet 2011; 377: 2103-14

Published Online
June 4, 2011

outcome:

Interpretation This trial has not confirmed a benefit of addition of cetuximab to oxaliplatin-based chemotherapy in first-line treatment of patients with advanced colorectal cancer. Cetuximab increases response rate, with no evidence of benefit in progression-free or overall survival in *KRAS* wild-type patients or even in patients selected by additional mutational analysis of their tumours. The use of cetuximab in combination with oxaliplatin and capecitabine in first-line chemotherapy in patients with widespread metastases cannot be recommended.

The COIN trial (colorectal cancer)



$n = 398$

HR [95% CI]	$\beta(0)$	FRET eff	Her2-Her3	Cetux	KRASmut
Cox (M1L1K5), $\ln Z = -2419.82$					
	-1.89	0.9 [0.7-1.0] p=0.3	1.1 [0.9-1.5] p=0.4	<u>0.8 [0.7-0.9]</u> p=0.03	1.3 [1.1-1.7] p=0.006
Model M2L2K4(R), $\ln Z = -2418.064$					
class I, W=31%	-2.57	1.8 [0.8-4.6] p=0.2	0.8 [0.4-1.7] p=0.6	<u>0.5 [0.3-1.0]</u> p=0.05	1.5 [0.9-2.6] p=0.2
alloc[p ₁ > 0.5]: N=59					
class II, W=69%	-1.56	0.5 [0.4-0.8] p=0.006	1.4 [0.9-2.1] p=0.1	<u>1.0 [0.7-1.4]</u> p=0.8	1.3 [0.9-1.9] p=0.1
alloc[p ₂ > 0.5]: N=339					

- ▶ two sub-cohorts, with similar base hazard rates, but distinct overall frailties and associations.
- ▶ method provides retrospective class assignment
- ▶ new tools to identify *a priori* the responders to Cetuximab?

The TOPICAL trial (lung cancer)

$n = 580$



Cohort heterogeneity and competing risks

- Regression for time-to-event data

- Consequences and fingerprints

- Informative censoring: the intuition

Bayesian latent class models

- Rationale and definition

- Tests on synthetic data

Applications in cancer research

- Epidemiological cancer data

- Data from failed cancer trials

New directions

- Prospective latent class prediction


- Overfitting in multivariate survival analysis

Summary

Prospective latent class prediction

If any of the covariates
correlate with retrospective
class membership:
(e.g. Amoris)

replace



$$p(t, r | \mathbf{z}) = \sum_{\ell=1}^L w_{\ell} p(t, r | \mathbf{z}, \ell) \quad \rightarrow \quad p(t, r | \mathbf{z}) = \sum_{\ell=1}^L w_{\ell}(\mathbf{z}) p(t, r | \mathbf{z}, \ell)$$

- ▶ suitable parametrisation $w_{\ell}(\mathbf{z})$
- ▶ *prospective* class prediction,
i.e. objective data-driven stratification to rescue failed trials
- ▶ but increasingly complex models,
many parameters: danger of overfitting ...

overfitting in Cox regression

ML method ...


p-values, z-scores,
confidence intervals,
don't measure overfitting!

rule of thumb: $d_{\max} = \text{events}/10$

- ▶ too optimistic ...
- ▶ must depend on β ...
- ▶ covariate correlations ...


What happens in overfitting regime?

Can we predict the optimal point?




$N = 500$,
predicted versus true regression coefficients
(synthetic data, no censoring)

$$d/N = 0.002$$




$N = 500$,
predicted versus true regression coefficients
(synthetic data, no censoring)

$$d/N = 0.10$$




$N = 500$,
predicted versus true regression coefficients
(synthetic data, no censoring)

$$d/N = 0.20$$




$N = 500$,
predicted versus true regression coefficients
(synthetic data, no censoring)



$$d/N = 0.30$$



$N = 500$,
predicted versus true regression coefficients
(synthetic data, no censoring)

$$d/N = 0.40$$





Base hazard rates underestimated for short times,
and over-estimated for large times ...


$d/N = 0.05, 0.15, 0.25, 0.35, 0.45, 0.55$
(lower to upper curves)

Gaussian association pars, $\langle \beta_\mu^2 \rangle = 0.25$, $N = 400$

Bad news

Overfitting *more dangerous*
than finite sample noise ...

*we always inflate associations
(whether positive or negative)*



Good news

Unlike pure noise,
deterministic bias may be predictable ...

Roadmap for research ...

- ▶ Predict asymptotic impact of overfitting, in terms of
 - ratio d/N
 - correlations among covariates
 - true association strengths β
- ▶ Overfitting correction of Cox parameters
 - reliable regression at ratios $d/N \sim 0.5$ or more?

Cohort heterogeneity and competing risks

- Regression for time-to-event data

- Consequences and fingerprints

- Informative censoring: the intuition

Bayesian latent class models

- Rationale and definition

- Tests on synthetic data

Applications in cancer research

- Epidemiological cancer data

- Data from failed cancer trials

New directions

- Prospective latent class prediction

- Overfitting in multivariate survival analysis

Summary

Summary

- ▶ Latent host or disease heterogeneity
 - can cause informative censoring
 - can cause violation of proportional hazards
 - invalidates interpretation of cause-specific quantities in Cox and KM
- ▶ Bayesian latent class regression
 - models all risks simultaneously
 - reliable discovery and mapping of latent cohort heterogeneity
 - Bayesian formulation: only reports substructure demanded by data
 - computationally feasible
- ▶ Application to epidemiological data with competing risks
 - decontamination formulae to undo competing risk effects
 - formulae for cumulative incidence functions
 - fully transparent interpretations
- ▶ Application to clinical trial data
 - data-driven stratification to rescue ‘fallen angels’

TACK SÅ MYCKET

King's College London

Paul Barber, James Barrett, Hans Garmö, Lars Holmberg, Katherine Lawler,
Tony Ng, Mark Rowley, Mieke van Hemelrijck, Wahyu Wulaningsih

Umeå University

Christel Häggström

Karolinska

Niklas Hammar, Göran Walldius

Uppsala University

Birgitta Grundmark, Björn Zethelius

Waseda University

Masato Inoue

Funding:

EPSRC, MRC, Prostate UK,
Ana Leaf Foundation

papers, seminars, notes:

<https://nms.kcl.ac.uk/ton.coolen>

