## Introduction to Survival Analysis

Statistical Physics Approaches to Systems Biology, Havana, Feb 2019
ACC Coolen
King's College London and Saddle Point Science


Introduction
Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

1982:
Commodore 64


## next generation data

 previous generation analysis ...

## Regression Models and Life-Tables

D. R. Cox

Journal of the Royal Statistical Society. Series B (Methodological), Volume 34, Issue 2 (1972), 187-220.

Stable URL:

Introduction
Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

Biomedicine has changed drastically in recent decades modern biomedical data


- volume of data ...
- diversity of data sources
(clinical, genomic, biomarkers, health records, imaging, ...)
- complexity of experimental pipelines (confounders, batch effects, variability between centres, ...)
- dimensionality of measurements clinical images $\left(\sim 10^{6}\right)$, transcriptome/proteome $\left(\sim 10^{5}\right)$, DNA and methylation $\left(\sim 10^{10}\right), \ldots$


## Personalized medicine: tailored treatments

Medicine of the present: one treatment fits all


Medicine of the future: more personalized diagnostics


Effect
generating 'big data' is not enough ...

- 'right drug, right dose, at the right time ...'
need predictive models $p(y \mid \boldsymbol{z})$,
$\boldsymbol{z}$ : individual's makeup (DNA, gene expr, metabolism, environment, ...)
$y$ : response to treatment
- regression:
find parameters $\boldsymbol{\theta}$ of model $p(y \mid \boldsymbol{z}, \boldsymbol{\theta})$ from historic data
curse of dimensionality ...
pre-genome medicine: $\quad N \sim 10^{3}$ data points, $\operatorname{dim}(\theta) \sim 10^{2}$
post-genome medicine: $N \sim 10^{4}$ data points, $\operatorname{dim}(\theta) \sim 10^{10}$
- simpler question: predict patient's individual risk (target aggressive treatments wisely)
cancer, heart disease, diabetes, ...:
relevant outcome is often a duration $t$,
OS (overall survival), or PFS (progression-free survival)
predictive model: $p(t \mid \boldsymbol{z}, \boldsymbol{\theta})$

Introduction
Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and p-values

Statistics in medicine: tricky business ...

"I can prove it or disprove it! What do you want me to do?"
why is statistics tricky?
Monty Hall problem
'Let's Make a Deal' (USA gameshow, 1963-1977)

standard quiz show, winner has to choose prize at the end, three doors: one with big prize, two with goats ...


- winner selects one door randomly
- Monty opens one door with a goat (two doors left ...)
- Monty gives winner the chance to change selection at last minute
would it matter?


## The main pitfalls in statistics

- accidental conditioning
(Monty Hall problem, share statistics, shop opening hours consultation, ...)
extra knowledge:
$\rightarrow$ reduces possibilities
$\rightarrow$ affects probabilities

$$
\overbrace{P(A \mid B)}^{\text {posterior }}=\frac{P(A, B)}{P(B)}=\overbrace{P(A)}^{\text {prior }} \times \frac{P(B \mid A)}{P(B)}
$$

- often counterintuitive
(Monty Hall problem, gambling,
human inability to generate random numbers, ...)
I have just thrown 10 successive sixes!
Prob $\approx 16.5 .10^{-8}$
- how many data do we need to be sure of something?

Is a genetic mutation harmless, or dangerous?
Is a given dice fair, or loaded?


typical data sets in cancer research:
$n \approx 500$ patients, at $2 \mathrm{~K} £$ each ...
what can we say after 500 throws?



- 'probability' can mean different things ...
our ignorance of
(a) something that cannot be known

(Russian roulette, we will spin the cylinder)
(b) something that is known, but not by us (Russian roulette, cylinder has already been spun)
relevant in medicine?
suppose we find survival function $S(t)=\mathrm{e}^{-t / \tau}$
explanation I: homogeneous cohort, random death times, each individual $i$ has hazard rate $1 / \tau$
explanation II: heterogeneous cohort, deterministic death times $t_{i}$, distributed according to $p(t)=\tau^{-1} \mathrm{e}^{-t / \tau}$ (potential for stratification!)
z-scores
reported in PLoS Medicine

Selective reporting
(aka cheating)


## All Trials Registered | All Results Reported

Results from half of all clinical trials are hidden.
Doctors don't have full information about the medicines we use.

## Sign the petition



Donate ;


Get involved,


Latest news ,

- missing values in data sets. ... red herrings or white sharks?

always check for informative missingness!
- correlation/association is not the same as causality!
imagine $Z$
is nr of hospital visits ...
result: positive correlation
 between $Z$ and risk!
$\beta>0$, ergo hospital visits dangerous?
or:
$Z=1,0$ : given strong chemotherapy yes/no but treatment not given if patient too weak ... result: positive association between $Z$ and health!
protective effect reported, even if chemotherapy ineffective!
two chemotherapies, $A$ and $B$, data on response rates from 880 patients
$Q$ : which treatment is better?

|  | CHEMO A | CHEMO B |
| :--- | :--- | :--- |
| response rate | $25 \%(76 / 300)$ | $28 \%(162 / 580)$ |

so treatment $B$ is better,
now we zoom in ...

|  | CHEMO A | CHEMO B |
| :--- | :--- | :--- |
| medical centre 1 | $40 \%(40 / 100)$ | $30 \%(150 / 500)$ |
| medical centre 2 | $18 \%(36 / 200)$ | $15 \%(12 / 80)$ |
| response rate | $25 \%(76 / 300)$ | $28 \%(162 / 580)$ |

still sure that $B$ is better?
Simpson's paradox

Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

## Terminology and objectives

- Data available

$$
\mathcal{D}=\left\{\left(\boldsymbol{z}_{1}, r_{1}, t_{1}\right), \ldots,\left(\boldsymbol{z}_{N}, r_{N}, t_{N}\right)\right\}
$$

$N$ samples/individuals ( $\left.\boldsymbol{z}_{i}, r_{i}, t_{i}\right)$, drawn independently from $p(t, r, \boldsymbol{z})$ (the population)
the 'covariates':

$$
\begin{aligned}
& z \in \mathbb{R}^{p}: \quad p \text { characteristics, measured at } t=0 \\
& \text { uncontrolled: e.g. gender, genome, ... } \\
& \text { controlled : } \\
& \text { e.g. medical treatment, ,.. } \\
& \text { modifiable : } \\
& \text { e.g. smoking, BMI, nutrition, ... }
\end{aligned}
$$

the 'clinical outcome':

$$
\begin{aligned}
& t \in \mathbb{R}^{+}: \begin{array}{l}
\text { failure time } \\
\text { e.g. death, onset/recurrence of disease, } \ldots
\end{array} \\
& r \in\{0,1, \ldots, R\}: \begin{array}{l}
\text { risk type that triggered failure }
\end{array} \\
& r=1 \ldots R: \text { true risks/diseases } \\
& r=0: \text { end of observation ('censoring') }
\end{aligned}
$$

| BMI |  | SELENIUM | PHYS_ACT_LEIS | PHYS_ACT_WORK | Smoking | Time | PCcens |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 22.63671875 | 105 | 2 | 1 | 0 | 33.6892539357 | 0 |
| 2 | 34.21875 | 65 | 2 | 2 | 0 | 24.810403833 | 1 |
| 3 | 20.06640625 | 72 | 2 | 3 | 2 | 23.1047227926 | 2 |
| 4 | 28.3984375 | 81 | 2 | 0 | 2 | 33.6783025325 | 0 |
| 5 | 22.94921875 | 73 | 0 | 3 | 2 | 32.8843258042 | 0 |
| 6 | 20.59765625 | 73 | 2 | 0 | 2 | 23.2936344969 | 1 |
| 7 | 26.46875 | 70 | 2 | 3 | 2 | 27.9069130732 | 2 |
| 8 | 26.38671875 | 91 | 1 | 0 | 1 | 26.6119096509 | 2 |
| 9 | 24.296875 | 73 | -10 | 0 | 0 | 26.379192334 | 1 |
| 10 | 30.01953125 | 84 | 1 | 1 | 2 | 26.8254620123 | 2 |
| 11 | 25.4296875 | 95 | 1 | 0 | 2 | 30.6557152635 | 2 |
| 13 | 23.19921875 | 67 | 3 | 1 | 0 | 33.2457221081 | 0 |
| 14 | 22.90625 | 82 | 2 | 2 | 0 | 33.6399726215 | 0 |
| 15 | 21.62890625 | 53 | 1 | 1 | 1 | 33.2457221081 | 0 |
| 16 | 23.046875 | 77 | 2 | 0 | 2 | 33.6125941136 | 0 |
| 17 | 21.70703125 | 76 | 1 | 2 | 2 | 27.8466803559 | 1 |
| 18 | 22.91796875 | 102 | 1 | 0 | 2 | 33.6125941136 | 0 |
| 19 | 24.5078125 | 57 | 2 | 0 | 2 | 25.8097193703 | 2 |
| 20 | 26.58984375 | 72 | 1 | 2 | 2 | 26.803559206 | 2 |
| 21 | 26.76953125 | 78 | 2 | 0 | 2 | 33.234770705 | 0 |
| 22 | 20.4296875 | 75 | -10 | 1 | 1 | 33.5934291581 | 0 |
| 23 | 25.0078125 | 69 | 0 | 3 | 0 | 33.6125941136 | 0 |
| 24 | 24.296875 | 73 | 2 | 3 | 2 | 21.6098562628 | 1 |
| 25 | 23.65625 | 75 | 2 | 3 | 1 | 33.2320328542 | 0 |
| 26 | 25.9296875 | 90 | 1 | 1 | 2 | 33.5359342916 | 0 |
| 27 | 23.3671875 | 58 | 1 | 1 | 2 | 33.2320328542 | 0 |
| 28 | 30.08984375 | 77 | -10 | -10 | 2 | 32.0438056126 | 2 |
| 29 | 31.08984375 | 66 | 1 | 0 | 0 | 29.4893908282 | 1 |
| 30 | 27.13671875 | 82 | 1 | -10 | 1 | 33.1526351814 | 0 |
| 31 | 19.828125 | 68 | 2 | 2 | 2 | 28.2600958248 | 2 |
| 32 | 27.30859375 | 97 | 2 | 3 | 1 | 33.5742642026 | 0 |
| 33 | 23.41796875 | 77 | 2 | 0 | 0 | 25.9219712526 | $29 / 56$ |



- Objective
find and quantify patterns that relate
covariates $\boldsymbol{z}$ to clinical outcomes ( $r, t$ ), in order to:
- predict clinical outcome for individuals
- discover disease mechanisms
- design interventions (modifiable covariates)
- Complications
- 'noise' caused by censoring
- we only know the earliest event (different risks prevent each other from happening)
- correlations between risks
- heterogeneity in patient cohorts
- overfitting danger, when $p$ is large relative to $N$

The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

## Survival probability and cause-specific hazard rates

- Joint event time statistics
imaginary situation: all events can be observed,
$t_{r}$ : time at which event $r$ occurs,
event time distributiion: $P\left(t_{0}, \ldots, t_{R}\right)$
normalisation:

$$
\int_{0}^{\infty} \ldots \int_{0}^{\infty} \mathrm{d} t_{0} \ldots \mathrm{~d} t_{R} P\left(t_{0}, \ldots, t_{R}\right)=1
$$

- Integrated event time distribution

$$
S\left(t_{0}, \ldots, t_{R}\right)=\int_{t_{0}}^{\infty} \ldots \int_{t_{R}}^{\infty} \mathrm{d} s_{0} \ldots \mathrm{~d} s_{R} P\left(s_{0}, \ldots, s_{R}\right)
$$

meaning: probability that event 0 occurs later than $t_{0}$, and event 1 occurs later than $t_{1}$, and $\ldots$

$$
S(0, \ldots, 0)=\int_{0}^{\infty} \ldots \int_{0}^{\infty} \mathrm{d} s_{0} \ldots \mathrm{~d} s_{R} P\left(s_{0}, \ldots, s_{R}\right)=1
$$

- Survival function $S(t)$
probability that all events happen later than time $t$ :

$$
S(t)=\int_{t}^{\infty} \ldots \int_{t}^{\infty} \mathrm{d} s_{0} \ldots \mathrm{~d} s_{R} P\left(s_{0}, \ldots, s_{R}\right)=S(t, t, \ldots, t)
$$

- Cause-specific hazard rates $h_{r}(t)$ how do individual risks impact on survival?

$$
h_{r}(t)=-\left[\frac{\partial}{\partial t_{r}} \log S\left(t_{0}, \ldots, t_{R}\right)\right]_{t_{k}=t \text { for all } k}
$$

work out, using $\frac{\mathrm{d}}{\mathrm{d} z} \theta(z)=\delta(z)$ :

$$
\begin{aligned}
h_{r}(t) & =\left[\frac{\frac{\partial}{\partial t_{r}} \int_{0}^{\infty} \ldots \int_{0}^{\infty} \mathrm{d} s_{0} \ldots \mathrm{~d} s_{R} P\left(s_{0}, \ldots, s_{R}\right) \prod_{k} \theta\left(s_{k}-t_{k}\right)}{S\left(t_{0}, \ldots, t_{R}\right)}\right]_{t_{k}=t \forall k} \\
& =\left[\frac{\int_{0}^{\infty} \cdots \int_{0}^{\infty} \mathrm{d} s_{0} \ldots \mathrm{~d} s_{R} P\left(s_{0}, \ldots, s_{R}\right) \delta\left(s_{r}-t_{r}\right) \prod_{k \neq r} \theta\left(s_{k}-t_{k}\right)}{S\left(t_{0}, \ldots, t_{R}\right)}\right]_{t_{k}=t \forall k} \\
& =\frac{\int_{t}^{\infty} \cdots \int_{t}^{\infty}\left(\prod_{r \neq \mu}^{R} \mathrm{~d} s_{r}\right) P\left(s_{0}, \ldots, s_{\mu-1}, t, s_{\mu+1}, \ldots, s_{R}\right)}{S(t)}
\end{aligned}
$$

$h_{r}(t) \mathrm{d} t$ : probability that event $r$ happens in time interval $[t, t+\mathrm{d} t)$, given that no event has happened yet prior to $t$

$$
h_{r}(t) \mathrm{d} t \quad=\operatorname{Prob}\left(t_{r} \in[t, t+\mathrm{d} t) \mid \text { no events yet at time } t\right) \quad(\mathrm{d} t \downarrow 0)
$$

The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and p-values

## Data likelihood in terms of cause specific hazard rates

most of the relevant quantities in survival analysis can be written in terms of the cause specific hazard rates

- Survival function

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} t} \log S(t) & =\frac{\mathrm{d}}{\mathrm{~d} t} \log S(t, t, \ldots, t) \\
& =\sum_{r=0}^{R}\left[\frac{\partial}{\partial t_{r}} \log S\left(t_{0}, \ldots, t_{R}\right)\right]_{t_{r}=t \forall r} \\
& =-\sum_{r=0}^{R} h_{r}(t)
\end{aligned}
$$

Hence, using $S(0)=1$,

$$
\log S(t)=\log S(0)-\int_{0}^{t} \mathrm{~d} s \sum_{r=0}^{R} h_{r}(s)=-\sum_{r=0}^{R} \int_{0}^{t} \mathrm{~d} s h_{r}(s)
$$

result:

$$
S(t)=\mathrm{e}^{-\sum_{r=0}^{R} \int_{0}^{t} \mathrm{ds} h_{r}(s)}
$$

- Data likelihood
$p(t, r) \mathrm{d} t$ : likelihood to observe first event being of type $r$, and occurring in time interval $[t, t+\mathrm{d} t)$ (with $\mathrm{d} t \downarrow 0$ )

To observe the above, the following statements must be true:
time of the event is in $[t, t+\mathrm{d} t)$, type of the event is $r$, no events occurred prior to $t$

$$
\theta\left(t_{r}-t\right) \theta\left(t+\mathrm{d} t-t_{r}\right) \prod_{k \neq r} \theta\left(t_{k}-t\right)=1
$$

hence

$$
\begin{aligned}
& p(t, r)= \lim _{\mathrm{d} t \downarrow 0} \frac{1}{\mathrm{~d} t} \operatorname{Prob}\left(\theta\left(t_{r}-t\right) \theta\left(t+\mathrm{d} t-t_{r}\right) \prod_{k \neq r} \theta\left(t_{k}-t\right)=1\right) \\
&= \lim _{\mathrm{d} \downarrow \downarrow 0} \frac{1}{\mathrm{~d} t} \int_{0}^{\infty} \cdots \int_{0}^{\infty} \mathrm{d} t_{0} \ldots t_{R} P\left(t_{1}, \ldots, t_{R}\right) \theta\left(t_{r}-t\right) \theta\left(t+\mathrm{d} t-t_{r}\right) \prod_{k \neq r} \theta\left(t_{k}-t\right) \\
&= \int_{0}^{\infty} \cdots \int_{0}^{\infty} \mathrm{d} t_{0} \ldots t_{R} P\left(t_{1}, \ldots, t_{R}\right) \lim _{\epsilon \downarrow 0} h_{\epsilon}\left(t_{r}-t\right) \prod_{k \neq r} \theta\left(t_{k}-t\right) \\
& \quad h_{\epsilon}(z)=\epsilon^{-1} \theta(z) \theta(\epsilon-z)= \begin{cases}\epsilon^{-1} & \text { for } z \in[0, \epsilon] \\
0 & \text { elsewhere }\end{cases}
\end{aligned}
$$

note: $\quad \lim _{\epsilon \downarrow 0} h_{\epsilon}(z)=\delta(z)$, so

$$
\begin{aligned}
p(t, r) & =\int_{0}^{\infty} \cdots \int_{0}^{\infty} \mathrm{d} t_{0} \ldots t_{R} P\left(t_{1}, \ldots, t_{R}\right) \lim _{\epsilon \downarrow 0} h_{\epsilon}\left(t_{r}-t\right) \prod_{k \neq r} \theta\left(t_{k}-t\right) \\
& =\int_{0}^{\infty} \cdots \int_{0}^{\infty} \mathrm{d} t_{0} \ldots t_{R} P\left(t_{1}, \ldots, t_{R}\right) \delta\left(t_{r}-t\right) \prod_{k \neq r} \theta\left(t_{k}-t\right) \\
& =h_{t}(t) S(t)=h_{t}(t) \mathrm{e}^{-\sum_{r^{\prime}=0}^{R} \int_{0}^{t} \mathrm{~d} s h_{r^{\prime}(s)}}
\end{aligned}
$$

- Further relation

$$
\begin{aligned}
p(t) & =\sum_{r=0}^{R} p(t, r)=\left(\sum_{r=0}^{R} h_{t}(t)\right) \mathrm{e}^{-\sum_{r^{\prime}=0}^{R} \int_{0}^{t} \mathrm{ds} h_{r^{\prime}}(s)} \\
& =-\frac{\mathrm{d}}{\mathrm{~d} t} \mathrm{e}^{-\sum_{r^{\prime}=0}^{R} J_{0}^{t} \mathrm{~d} s h_{r^{\prime}}(s)}=-\frac{\mathrm{d}}{\mathrm{~d} t} S(t)
\end{aligned}
$$

- Cause-specific hazard rates in terms of data probabilities

$$
S(t)=S(0)+\int_{0}^{t} \mathrm{~d} t^{\prime} \frac{\mathrm{d}}{\mathrm{~d} t^{\prime}} S\left(t^{\prime}\right)=1-\int_{0}^{t} \mathrm{~d} t^{\prime} p\left(t^{\prime}\right)=\int_{t}^{\infty} \mathrm{d} s p(s)
$$

substitute into formula for $p(t, r)$ :

$$
h_{r}(t)=\frac{p(t, r)}{\sum_{r^{\prime}=0}^{R} \int_{t}^{\infty} \mathrm{d} s p\left(s, r^{\prime}\right)}
$$

Post-genome medicine Statistics in medicine is tricky

The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

## Pitfalls and misconceptions

cause specific hazard rates can be tricky ...

$$
S(t)=\prod_{r} \mathrm{e}^{-\int_{0}^{t} \mathrm{~d} s h_{r}(s)}
$$

- Survival function formula factorizes over risks, does this imply that the risks are uncorrelated?
No. All risks $k \neq r$ can contribute to each $h_{r}(t)$ via the conditioning, i.e. the likelihood that nothing has happened yet prior to $t$. Risks may well interact strongly with each other, but we can no longer see this after we have calculated the rates $\left\{h_{r}(t)\right\}$ and forget about the times $\left(t_{0}, \ldots, t_{R}\right)$.
- Do we get the survival function for the situation where risk $\mu$ is disabled (e.g. a disease removed from the world) by setting $h_{\mu}(t)$ to zero?

$$
S(t) \rightarrow \mathrm{e}^{-\sum_{r \neq \mu} \int_{0}^{t} \mathrm{~d} s h_{r}(s)}
$$

No. We would have $h_{\mu}(t)=0$ for all $t$, but that is not all. Disabling risk $\mu$ can change also all hazard rates $h_{r}(t)$ with $r \neq \mu$, due to correlations among the different risks in combination with the conditioning.

Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

## Special cases

- Time-independent hazard rates

$$
h_{r}(t)=h_{r}: \quad S(t)=\mathrm{e}^{-t \sum_{r=0}^{R} h_{r}} \quad p(t, r)=h_{r} \mathrm{e}^{-t \sum_{r^{\prime}=0}^{R} h_{r^{\prime}}}
$$

- A single risk, $R=1$

One risk,
hazard rate $h(t): \quad S(t)=\mathrm{e}^{-\int_{0}^{t} \mathrm{~d} s h(s)} \quad p(t)=h(t) \mathrm{e}^{-\int_{0}^{t} \mathrm{~d} s h(s)}$

- Most probable event time distribution for $R=1$

Suppose we know only average event time $\langle t\rangle=\int_{0}^{\infty} \mathrm{d} t t p(t)$, most probable $p(t)$ : maximize Shannon entropy $H=-\int_{0}^{\infty} \mathrm{d} t p(t) \log p(t)$, subject to $\int_{0}^{\infty} \mathrm{d} t p(t)=1$ and $\int_{0}^{\infty} \mathrm{d} t t p(t)=\langle t\rangle$
Lagrange's method:

$$
\begin{aligned}
& \frac{\delta}{\delta p(t)} \int_{0}^{\infty} \mathrm{d} s p(s) \log p(s)=\frac{\delta}{\delta p(t)}\left\{\lambda_{0} \int_{0}^{\infty} \mathrm{d} s p(s)+\lambda_{1} \int_{0}^{\infty} \mathrm{d} s p(s) s\right\} \\
& 1+\log p(t)=\lambda_{0}+\lambda_{1} t \quad \text { so } p(t)=\mathrm{e}^{\lambda_{0}-1+\lambda_{1} t}
\end{aligned}
$$

use constraints:

$$
p(t)=\langle t\rangle^{-1} \mathrm{e}^{-t /\langle t\rangle}
$$

The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

Independently
distributed event times

$$
P\left(t_{0}, \ldots, t_{R}\right)=\prod_{r=0}^{R} P_{r}\left(t_{r}\right)
$$

- Integrated event time distr

$$
\begin{gathered}
S\left(t_{0}, \ldots, t_{R}\right)=\int_{t_{0}}^{\infty} \cdots \int_{t_{R}}^{\infty} \mathrm{d} s_{0} \ldots \mathrm{~d} s_{R} \prod_{r=0}^{R} P_{r}\left(t_{r}\right)=\prod_{r=0}^{R} S_{r}\left(t_{r}\right) \\
S_{r}(t)=\int_{t}^{\infty} \mathrm{d} s P_{r}(s)
\end{gathered}
$$

- Cause specific hazard rates

$$
h_{r}(t)=-\left[\frac{\partial}{\partial t_{r}} \sum_{r^{\prime}=0}^{R} \log S_{r^{\prime}}\left(t_{r^{\prime}}\right)\right]_{t_{k}=t \text { for all k}}=-\frac{\mathrm{d}}{\mathrm{~d} t} \log S_{r}(t)
$$

hence

$$
S_{r}(t)=\mathrm{e}^{-\int_{0}^{t} \mathrm{~d} s h_{r}(s)}
$$

joint event time distr now follows from the cause-specific hazard rates

$$
\begin{aligned}
& P_{r}(t)=-\frac{\mathrm{d}}{\mathrm{~d} t} S_{r}(t)=-\frac{\mathrm{d}}{\mathrm{~d} t} \mathrm{e}^{-\int_{0}^{t} \mathrm{~d} s h_{r}(s)}=h_{t}(t) \mathrm{e}^{-\int_{0}^{t} \mathrm{~d} s h_{r}(s)} \\
& P\left(t_{0}, \ldots, t_{R}\right)=\prod_{r=0}^{R}\left[h_{t}\left(t_{r}\right) \mathrm{e}^{-\int_{0}^{t_{r} \mathrm{~d} s} h_{r}(s)}\right]
\end{aligned}
$$

Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and p-values

## The identifiability problem (Tsiatis)

- Observable from data: $p(t, r)$, equivalently: $\left\{h_{0}(t), \ldots, h_{R}(t)\right\}$, since

$$
p(t, r)=h_{r}(t) \mathrm{e}^{-\sum_{r^{\prime}=0}^{R} \int_{0}^{t} \mathrm{~d} s h_{r^{\prime}(s)}}, \quad h_{r}(t)=\frac{p(t, r)}{\sum_{r^{\prime}=0}^{R} \int_{t}^{\infty} \mathrm{d} s p\left(s, r^{\prime}\right)}
$$

- For any $\left\{h_{0}(t), \ldots, h_{R}(t)\right\}$, even those corresponding to statistically dependent event times, there exists a distribution for independent event times that will give exactly the same cause-specifc hazard rates, namely

$$
P\left(t_{0}, \ldots, t_{R}\right)=\prod_{r=0}^{R} \prod_{r=0}^{R}\left[h_{t}\left(t_{r}\right) \mathrm{e}^{-\int_{0}^{t_{r}} \mathrm{ds} h_{r}(s)}\right]
$$

Hence, survival data alone do not generally permit us to identify the underlying joint distribution of event times - in particular, we cannot infer whether or not the event times of the different risks are independent.
a serious problem ...

- The Bayesian view
on the identifiability problem
- multiple hypotheses $H$ may explain our data
- but not all are equally probable ...
- calculate each $\operatorname{Prob}(H \mid \mathcal{D})$ from Bayes' formula
- Illustration
true data:

$$
p\left(t_{2}\right)=a \mathrm{e}^{-a t_{2}}, \quad \begin{cases}\text { with prob } \epsilon: & t_{1}=t_{2}+\tau \\ \text { with prob } 1-\epsilon: & \text { draw } t_{1} \text { from } p\left(t_{1}\right)=b \mathrm{e}^{-b t_{1}}\end{cases}
$$

explanation assuming
risk independence:

$$
p\left(t_{2}\right)=a \mathrm{e}^{-a t_{2}}, \quad p\left(t_{1}\right)=\underbrace{-\left(\epsilon+(1-\epsilon) \mathrm{e}^{-b t_{1}}\right) \log \left(\epsilon+(1-\epsilon) \mathrm{e}^{-b t_{1}}\right)}_{\text {with prob } \epsilon: \text { event } 1 \text { never happens }}
$$

implausible if e.g. risk 2 is cancer, risk 1 is death ...

Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

## Proportional hazards regression (Cox)

## definitions and assumptions

- Survival analysis with covariates

Add $\boldsymbol{z}$ as conditions to definitions and identities

$$
\begin{aligned}
& S(t) \rightarrow S(t \mid \boldsymbol{z}), \quad h_{r}(t) \rightarrow h_{r}(t \mid \boldsymbol{z}), \quad p(t, r) \rightarrow p(t, r \mid \boldsymbol{z}) \\
& S(t \mid \boldsymbol{z})=\int_{t}^{\infty} \cdots \int_{t}^{\infty} \mathrm{d} s_{0} \ldots \mathrm{~d} s_{R} P\left(s_{0}, \ldots, s_{R} \mid \boldsymbol{z}\right) \\
& S(t \mid \boldsymbol{z})=\mathrm{e}^{-\sum_{r=0}^{R} \int_{0}^{t} \mathrm{~d} s h_{r}(s \mid \boldsymbol{Z})}, \quad p(t, r \mid \boldsymbol{z})=h_{r}(t \mid \boldsymbol{z}) \mathrm{e}^{\left.-\sum_{r^{\prime}=0}^{R} \int_{0}^{t} \mathrm{~d} s h_{r^{\prime}(s \mid \boldsymbol{Z}}\right)}
\end{aligned}
$$

- Cox model

Parametrized form for the hazard rates:

$$
h_{r}(t \mid \boldsymbol{Z})=\lambda_{r}(t) \mathrm{e}^{\boldsymbol{\beta}^{r} \cdot \boldsymbol{Z}}, \quad \boldsymbol{\beta}^{r}=\left(\beta_{1}^{r}, \ldots, \beta_{p}^{r}\right), \quad \boldsymbol{\beta}^{r} \cdot \boldsymbol{z}=\sum_{\mu=1}^{p} \beta_{\mu}^{r} z_{\mu}
$$

$\lambda_{r}(t)$ : 'base hazard rate' of risk $r$ (covariate-independent contribution to risk)
$\beta^{r}$ : 'association parameters' of risk $r$ (impact of covariate values on risk)

- Main features of Cox's choice
- 'Proportional hazards'
due to exponential form, effect of each covariate is multiplicative:

$$
h_{r}(t)=\underbrace{\lambda_{r}(t)}_{\text {base hazard rate }{ }^{\prime} \text { 'proportional hazards' }} \times \underbrace{\mathrm{e}^{\beta_{1}^{r} z_{1}} \times \ldots \times \mathrm{e}^{\beta_{p}^{r} z_{p}}}
$$

- Effects of covariates on risk independent of time
- There exists a hyper-plane in covariate space that separates high risk individuals from low risk individuals

$$
\begin{array}{rll}
\text { 'high risk individuals' : } & \beta_{1}^{r} z_{1}+\ldots+\beta_{p}^{r} z_{p} & \text { large } \\
\text { 'Iow risk individuals' } & \beta_{1}^{r} z_{1}+\ldots+\beta_{p}^{r} z_{p} & \text { small }
\end{array}
$$

- One can quantify risk impact of each individual covariate $\mu$ in a single time-independent number: the 'hazard ratio'

$$
H R_{\mu}^{r}=\frac{\left.h_{r}(t \mid \boldsymbol{z})\right|_{z_{\mu}=1}}{\left.h_{r}(t \mid \boldsymbol{z})\right|_{z_{\mu}=0}}=\frac{\lambda_{r}(t) \mathrm{e}^{\beta_{\mu}^{r} \cdot 1+\sum_{\nu \neq \mu} \beta_{\nu}^{r} z_{\nu}}}{\lambda_{r}(t) \mathrm{e}^{\beta_{\mu}^{r} \cdot 0+\sum_{\nu \neq \mu} \beta_{\nu}^{r} z_{\nu}}}=\mathrm{e}^{\beta_{\mu}^{r}}
$$

If no impact on risk: $\quad \beta_{\mu}^{r}=0, \operatorname{HR}_{\mu}^{r}=1$.

Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

## Detour: parameter estimation from data

given data $\mathcal{D}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{N}\right\}$,
and a model $p(\boldsymbol{x} \mid \boldsymbol{\theta})$ to explain these data, what can we say about the parameters $\boldsymbol{\theta}$ ?

- Bayesian parameter inference
assume the $\left\{\boldsymbol{x}_{i}\right\}$ were indeed drawn randomly \& independently from $p(\boldsymbol{x} \mid \boldsymbol{\theta})$,

$$
p(\mathcal{D} \mid \boldsymbol{\theta})=\prod_{i=1}^{N} p\left(\boldsymbol{x}_{i} \mid \boldsymbol{\theta}\right)
$$

Bayes' identity:

$$
p(\boldsymbol{\theta} \mid \mathcal{D})=\frac{p(\mathcal{D} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta})}{p(\mathcal{D})}=\frac{p(\mathcal{D} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta})}{\int \mathrm{d} \boldsymbol{\theta}^{\prime} p\left(\mathcal{D} \mid \boldsymbol{\theta}^{\prime}\right) p\left(\boldsymbol{\theta}^{\prime}\right)}, \quad p(\boldsymbol{\theta}): \text { prior }
$$

- Simplifications
minus log-likelihood regularizer
MAP : $\quad \boldsymbol{\theta}^{\star}=\operatorname{argmax} p(\boldsymbol{\theta} \mid \mathcal{D})=\operatorname{argmin}[\overbrace{-\log p(\mathcal{D} \mid \boldsymbol{\theta})} \overbrace{-\log p(\boldsymbol{\theta})}]$
ML: $\quad \boldsymbol{\theta}^{\star}=\operatorname{argmin}[-\log p(\mathcal{D} \mid \boldsymbol{\theta})] \quad$ i.e. $p(\boldsymbol{\theta})=$ constant
- Maximimum Likelihood (ML) regression

$$
\boldsymbol{\theta}^{\star}=\operatorname{argmin}[-\log p(\mathcal{D} \mid \boldsymbol{\theta})]
$$

define empirical data distribution

$$
\hat{p}(\boldsymbol{x})=\frac{1}{N} \sum_{i=1}^{N} \delta\left(\boldsymbol{x}-\boldsymbol{x}_{i}\right)
$$

note:

$$
\begin{aligned}
-\frac{1}{N} \log p(\mathcal{D} \mid \boldsymbol{\theta}) & =-\frac{1}{N} \log \prod_{i=1}^{N} p\left(\boldsymbol{x}_{i} \mid \boldsymbol{\theta}\right)=-\frac{1}{N} \sum_{i=1}^{N} \log p\left(\boldsymbol{x}_{i} \mid \boldsymbol{\theta}\right) \\
& =-\int \mathrm{d} \boldsymbol{x} \hat{p}(\boldsymbol{x}) \log p(\boldsymbol{x} \mid \boldsymbol{\theta}) \\
& =\int \mathrm{d} \boldsymbol{x} \hat{p}(\boldsymbol{x}) \log \left[\frac{\hat{p}(\boldsymbol{x})}{p(\boldsymbol{x} \mid \boldsymbol{\theta})}\right]-\int \mathrm{d} \boldsymbol{x} \hat{p}(\boldsymbol{x}) \log \hat{p}(\boldsymbol{x}) \\
& =\underbrace{D\left(\hat{p}| | p_{\boldsymbol{\theta}}\right)}_{K L \text { distance }}+\underbrace{H[\hat{p}]}_{\text {Shannon entropy }}
\end{aligned}
$$

hence: ML finds the parameter vector $\boldsymbol{\theta}$ that minimizes the KL distance between $\hat{p}(\boldsymbol{x})$ and $p(\boldsymbol{x} \mid \boldsymbol{\theta})$

- Beyond most probable parameters: error bars
return to full posterior distribution

$$
p(\boldsymbol{\theta} \mid \mathcal{D})=\frac{\mathrm{e}^{-\Omega(\boldsymbol{\theta}, \mathcal{D})}}{\int \mathrm{d} \boldsymbol{\theta}^{\prime} \mathrm{e}^{-\Omega\left(\boldsymbol{\theta}^{\prime}, \mathcal{D}\right)}}, \quad \Omega(\boldsymbol{\theta}, \mathcal{D})=-\log p(\mathcal{D} \mid \boldsymbol{\theta})-\log p(\boldsymbol{\theta})
$$

expand $\Omega$ around minimum $\boldsymbol{\theta}^{\star}$ :

$$
\Omega(\boldsymbol{\theta}, \mathcal{D})=\Omega\left(\boldsymbol{\theta}^{\star}, \mathcal{D}\right)+\frac{1}{2}\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{\star}\right) \cdot \boldsymbol{A}\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{\star}\right)+\mathcal{O}\left(\left|\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{\star}\right)\right|^{3}\right)
$$

truncate after quadratic term:
$p(\boldsymbol{\theta} \mid \mathcal{D}) \approx\left[\frac{\operatorname{det} \boldsymbol{A}}{(2 \pi)^{N}}\right]^{\frac{1}{2}} \mathrm{e}^{-\frac{1}{2}\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{\star}\right) \cdot \boldsymbol{A}\left(\boldsymbol{\theta}-\boldsymbol{\theta}^{\star}\right)}, \quad\left\langle\left(\theta_{\mu}-\theta_{\mu}^{\star}\right)\left(\theta_{\nu}-\theta_{\nu}^{\star}\right)\right\rangle=\left(\boldsymbol{A}^{-1}\right)_{\mu \nu}$
hence, error bars for MAP/ML estimators:
$\theta_{\mu}=\theta_{\mu}^{\star} \pm\left(\boldsymbol{A}^{-1}\right)_{\mu \mu}, \quad \boldsymbol{A}_{\mu \nu}=\frac{\partial^{2}}{\partial \theta_{\mu} \partial \theta_{\nu}}[-\log p(\mathcal{D} \mid \boldsymbol{\theta})-\log p(\boldsymbol{\theta})]_{\boldsymbol{\theta}^{\star}}$

Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and p-values

ML parameters of Cox's model

$$
\begin{aligned}
& \boldsymbol{\theta}=\left\{\lambda_{r}, \boldsymbol{\beta}^{r}\right\}: \\
& p(t, r \mid \boldsymbol{z}, \boldsymbol{\theta})=h_{r}(t \mid \boldsymbol{z}, \boldsymbol{\theta}) \mathrm{e}^{-\sum_{r^{\prime}=0}^{\mathcal{R}} \int_{0}^{t} \mathrm{~d} n_{r^{\prime}(s \mid \boldsymbol{Z}, \boldsymbol{\theta})}}, \quad h_{r}(t \mid \boldsymbol{z}, \boldsymbol{\theta})=\lambda_{r}(t) \mathrm{e}^{\boldsymbol{\beta}^{r} \cdot \boldsymbol{z}}
\end{aligned}
$$

- ML inference

$$
\begin{aligned}
\boldsymbol{\theta}^{\star} & =\operatorname{argmin}_{\boldsymbol{\theta}}[-\log p(\mathcal{D} \mid \boldsymbol{\theta})] \\
& =\operatorname{argmin}_{\boldsymbol{\theta}}\left[-\sum_{i=1}^{N} \log p\left(t_{i}, r_{i} \mid \boldsymbol{z}_{i}, \boldsymbol{\theta}\right)\right] \\
& =\operatorname{argmin}_{\boldsymbol{\theta}}\left[-\sum_{i=1}^{N} \log h_{r_{i}}\left(t_{i} \mid \boldsymbol{z}_{i}, \boldsymbol{\theta}\right)+\sum_{i=1}^{N} \sum_{r=0}^{R} \int_{0}^{t_{i}} \mathrm{~d} t \lambda_{r}(t) \mathrm{e}^{\boldsymbol{\beta}^{r} \cdot \boldsymbol{z}_{i}}\right] \\
& =\operatorname{argmin}_{\boldsymbol{\theta}}\left[-\sum_{i=1}^{N} \log \lambda_{r_{i}}\left(t_{i}\right)-\sum_{i=1}^{N} \boldsymbol{\beta}^{r_{i}} \cdot \boldsymbol{z}_{i}+\sum_{i=1}^{N} \sum_{r=0}^{R} \int_{0}^{t_{i}} \mathrm{~d} t \lambda_{r}(t) \mathrm{e}^{\boldsymbol{\beta}^{r} \cdot \boldsymbol{z}_{i}}\right] \\
= & \operatorname{argmin}_{\boldsymbol{\theta}} \sum_{r=0}^{R}\left[-\sum_{i=1}^{N} \delta_{r, r_{i}} \log \lambda_{r}\left(t_{i}\right)-\sum_{i=1}^{N} \delta_{r, r_{i}} \boldsymbol{\beta}^{r} \cdot \boldsymbol{z}_{i}+\sum_{i=1}^{N} \int_{0}^{t_{i}} \mathrm{~d} t \lambda_{r}(t) \mathrm{e}^{\boldsymbol{\beta}^{r} \cdot \boldsymbol{z}_{i}}\right]
\end{aligned}
$$

To minimize

$$
\begin{aligned}
& \Psi\left[\left\{\lambda_{r}, \boldsymbol{\beta}^{r}\right\}\right]=\sum_{r=0}^{R}\left[-\int \mathrm{d} t \log \lambda_{r}(t) \sum_{i=1}^{N} \delta_{r, r_{i}} \delta\left(t-t_{i}\right)-\boldsymbol{\beta}^{r} \cdot \sum_{i=1}^{N} \delta_{r, r_{i}} \boldsymbol{z}_{i}\right. \\
&\left.+\int \mathrm{d} t \lambda_{r}(t) \sum_{i=1}^{N} \theta\left(t_{i}-t\right) \mathrm{e}^{\boldsymbol{\beta}} \cdot \boldsymbol{z}_{i}\right]
\end{aligned}
$$

- Minimize over functions $\lambda_{r}(t)$ first

$$
\begin{gathered}
\frac{\delta}{\delta \lambda_{r}(t)}\left[\sum_{i=1}^{N} \int \mathrm{~d} t \theta\left(t_{i}-t\right) \lambda_{r}(t) \mathrm{e}^{\boldsymbol{\beta}^{r} \cdot \mathbf{z}_{i}}-\int \mathrm{d} t \log \lambda_{r}(t) \sum_{i=1}^{N} \delta_{r, r_{i}} \delta\left(t-t_{i}\right)\right]=0 \\
\text { (functional differentiation) } \\
\sum_{i=1}^{N} \theta\left(t_{i}-t\right) \mathrm{e}^{\boldsymbol{\beta}^{r} \cdot \mathbf{Z}_{i}}-\frac{1}{\lambda_{r}(t)} \sum_{i=1}^{N} \delta_{r, r_{i}} \delta\left(t-t_{i}\right)=0 \\
\lambda_{r}(t)=\frac{\sum_{i=1}^{N} \delta_{r, r_{i}} \delta\left(t-t_{i}\right)}{\sum_{i=1}^{N} \theta\left(t_{i}-t\right) \mathrm{e}^{r} \cdot \boldsymbol{Z}_{i}} \quad \text { Breslow's estimator }
\end{gathered}
$$

- Insert into $\Psi$, then minimize over $\left\{\boldsymbol{\beta}^{r}\right\}$
- To minimize

$$
\begin{aligned}
\Psi\left[\left\{\boldsymbol{\beta}^{r}\right\}\right]= & \sum_{r=0}^{R}\left[-\int \mathrm{d} t \log \left(\frac{\sum_{i=1}^{N} \delta_{r, r_{i}} \delta\left(t-t_{i}\right)}{\sum_{i=1}^{N} \theta\left(t_{i}-t\right) \mathrm{e} \boldsymbol{\beta}^{r} \cdot \boldsymbol{z}_{i}}\right) \sum_{j=1}^{N} \delta_{r, r_{j}} \delta\left(t-t_{j}\right)\right. \\
& \left.-\boldsymbol{\beta}^{r} \cdot \sum_{i=1}^{N} \delta_{r, r_{i}} \boldsymbol{z}_{i}+\int \mathrm{d} t \sum_{i=1}^{N} \delta_{r, r_{i}} \delta\left(t-t_{i}\right)\right] \\
= & \sum_{r=0}^{R}\left[\int \mathrm{~d} t \log \left(\sum_{i=1}^{N} \theta\left(t_{i}-t\right) \mathrm{e}^{\boldsymbol{\beta}^{r} \cdot \boldsymbol{z}_{i}}\right) \sum_{j=1}^{N} \delta_{r, r_{j}} \delta\left(t-t_{j}\right)\right. \\
& \left.-\boldsymbol{\beta}^{r} \cdot \sum_{i=1}^{N} \delta_{r, r_{i}} \boldsymbol{z}_{i}\right]+ \text { terms independent of }\left\{\boldsymbol{\beta}^{r}\right\} \\
= & \sum_{r=0}^{R} \Psi_{r}\left(\boldsymbol{\beta}^{r}\right)+\text { terms independent of }\left\{\boldsymbol{\beta}^{r}\right\}
\end{aligned}
$$

with

$$
\Psi_{r}(\boldsymbol{\beta})=\sum_{j=1}^{N} \delta_{r, r_{j}} \log \left(\sum_{i=1}^{N} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \boldsymbol{z}_{i}}\right)-\boldsymbol{\beta} \cdot \sum_{i=1}^{N} \delta_{r, r_{i}} \boldsymbol{z}_{i}
$$

hence

$$
\boldsymbol{\beta}^{r \star}=\operatorname{argmin}_{\boldsymbol{\beta}} \Psi_{r}(\boldsymbol{\beta})
$$

- Each risk $r$, find minima of $\Psi_{r}$ :

$$
\begin{aligned}
\frac{\partial}{\partial \beta_{\mu}} \Psi_{r}(\boldsymbol{\beta}) & =\frac{\partial}{\partial \beta_{\mu}}\left[\sum_{j=1}^{N} \delta_{r, r_{j}} \log \left(\sum_{i=1}^{N} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \mathbf{z}_{i}}\right)-\boldsymbol{\beta} \cdot \sum_{j=1}^{N} \delta_{r, r_{j}} \boldsymbol{z}_{j}\right] \\
& =\sum_{j=1}^{N} \delta_{r, r_{j}} \frac{\sum_{i=1}^{N} z_{i \mu} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \boldsymbol{z}_{i}}}{\sum_{i=1}^{N} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \boldsymbol{z}_{i}}}-\sum_{j=1}^{N} \delta_{r, r_{j}} z_{j \mu}
\end{aligned}
$$

so $\boldsymbol{\beta}^{r \star}$ is solution of:

$$
\sum_{j=1}^{N} \delta_{r, r_{j}}\left[\frac{\sum_{i=1}^{N} z_{i \mu} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \mathbf{z}_{i}}}{\sum_{i=1}^{N} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \mathbf{z}_{i}}}-z_{j \mu}\right]=0
$$

$p$ coupled nonlinear equations, for each risk

Final protocol:

1. Solve $\left\{\beta^{r \star}\right\}$ from above eqns (numerically)
2. Calculate $\left\{\lambda^{r}(t)\right\}$ (from Breslow's formula)
3. Predict outcomes via $p(r, t \mid \boldsymbol{z})$ for new patients (using Cox's model, with ML parameters)

Post-genome medicine
Statistics in medicine is tricky
The formalism of survival analysis
Terminology and objectives
Survival probability and cause-specific hazard rates
Data likelihood in terms of cause specific hazard rates
Pitfalls and misconceptions
Special cases
Event time correlations and identifiability
Independently distributed event times
The identifiability problem
Proportional hazards regression
Definitions and assumptions
Parameter estimation from data
ML parameters of Cox's model
Uniqueness, error bars, and and $p$-values

Uniqueness, error bars, and and $p$-values
Curvature of $\Psi_{r}(\beta)$ :

$$
\begin{aligned}
\frac{\partial^{2} \Psi_{r}(\boldsymbol{\beta})}{\partial \beta_{\mu} \beta_{\nu}} & =\frac{\partial}{\partial \beta_{\nu}}\left[\sum_{j=1}^{N} \delta_{r, r_{j}} \frac{\sum_{i=1}^{N} z_{i \mu} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \boldsymbol{z}_{i}}}{\sum_{i=1}^{N} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \mathbf{z}_{i}}}-\sum_{j=1}^{N} \delta_{r, r_{j}} z_{j \mu}\right] \\
& =\sum_{j=1}^{N} \delta_{r, r_{j}}\left[\left\langle z_{\mu} z_{\nu}\right\rangle_{j}-\left\langle z_{\mu}\right\rangle_{j}\left\langle z_{\nu}\right\rangle_{j}\right]
\end{aligned}
$$

with

$$
\langle f(\boldsymbol{z})\rangle_{j}=\sum_{i=1}^{N} w(i \mid j) f\left(\boldsymbol{z}_{i}\right), \quad w(i \mid j)=\frac{\theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \mathbf{z}_{i}}}{\sum_{i=1}^{N} \theta\left(t_{i}-t_{j}\right) \mathrm{e} \cdot \boldsymbol{\beta} \cdot \boldsymbol{z}_{i}}
$$

properties, consequences:

- Convexity
curvature matrix is positive definite, i.e. $\Psi_{r}(\beta)$ convex, since

$$
\begin{aligned}
\forall \boldsymbol{x} \in \mathbb{R}^{p}: \quad \sum_{\mu, \nu=1}^{p} x_{\mu} x_{\nu} \frac{\partial^{2} \Psi_{r}(\boldsymbol{\beta})}{\partial \beta_{\mu} \beta_{\nu}} & =\sum_{j=1}^{N} \delta_{r, r_{j}}\left[\left\langle(\boldsymbol{x} \cdot \boldsymbol{z})^{2}\right\rangle_{j}-\langle\boldsymbol{x} \cdot \boldsymbol{z}\rangle_{j}^{2}\right] \\
& =\sum_{j=1}^{N} \delta_{r, r_{j}}\left[\left\langle\left(\boldsymbol{x} \cdot \boldsymbol{z}-\langle\boldsymbol{x} \cdot \mathbf{z}\rangle_{j}\right)^{2}\right\rangle_{j}\right] \geq 0
\end{aligned}
$$ since $\Psi_{r}(\beta)$ convex, can have only one minimum

- Error bars for association parameters

Neglect fluctuations in $\left\{\lambda_{r}(t)\right\}$, focus on $\Delta \beta_{\mu}^{r \star}$ :

$$
\begin{aligned}
A(r)_{\mu \nu} & =\sum_{j=1}^{N} \delta_{r, r_{j}}\left[\left\langle z_{\mu} z_{\nu}\right\rangle_{j}-\left\langle z_{\mu}\right\rangle_{j}\left\langle z_{\nu}\right\rangle_{j}\right] \\
\langle f(\boldsymbol{z})\rangle_{j} & =\sum_{i=1}^{N} w(i \mid j) f\left(\boldsymbol{z}_{i}\right), \quad w(i \mid j)=\frac{\theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \boldsymbol{z}_{i}}}{\sum_{i=1}^{N} \theta\left(t_{i}-t_{j}\right) \mathrm{e}^{\boldsymbol{\beta} \cdot \boldsymbol{z}_{i}}}
\end{aligned}
$$

Then

$$
\beta_{\mu}^{r}=\beta_{\mu}^{r \star} \pm \sigma_{\mu}^{r}, \quad \sigma_{\mu}^{r}=\left(\boldsymbol{A}(r)^{-1}\right)_{\mu \mu}
$$

- p-values of inferred $\beta_{\mu}^{\star}$
definition: the probability to find a value $\beta_{\mu}^{\star}$ (or one further away from zero) due to fluctuations, when the true value is zero
approx: assume above error bar is correct, disregard correlations,

$$
\begin{aligned}
p-\text { value } & =\operatorname{Prob}\left(\left|\beta_{\mu}\right| \geq\left|\beta_{\mu}^{\star}\right|\right)=1-\frac{2}{\sigma_{\mu} \sqrt{2 \pi}} \int_{0}^{\left|\beta_{\mu}^{\star}\right|} \mathrm{d} \beta \mathrm{e}^{-\frac{1}{2} \beta^{2} / \sigma_{\mu}^{2}} \\
& =1-\operatorname{Erf}\left(\left|\beta_{\mu}^{\star}\right| / \sigma_{\mu} \sqrt{2}\right), \quad \beta_{\mu}^{\star} / \sigma_{\mu}: z-\text { score }
\end{aligned}
$$

## Explanation for Simson's paradox

|  | CHEMO A | CHEMO B |
| :--- | :--- | :--- |
| medical centre 1 | $40 \%(40 / 100)$ | $30 \%(150 / 500)$ |
| medical centre 2 | $18 \%(36 / 200)$ | $15 \%(12 / 80)$ |
| response rate | $25 \%(76 / 300)$ | $28 \%(162 / 580)$ |

$P($ response $\mid$ chemo $)=\sum_{\text {centres }} P($ response $\mid$ chemo, centre $) P($ centre $\mid$ chemo $)$

$$
\begin{aligned}
& P(\text { resp } \mid A)=\frac{40}{100} \frac{100}{300}+\frac{36}{200} \frac{200}{300}=25 \% \\
& P(\text { resp } \mid B)=\frac{30}{100} \frac{500}{580}+\frac{15}{100} \frac{80}{580}=28 \%
\end{aligned}
$$

if chemo choice indep of centre:

$$
\begin{aligned}
& P(\text { resp } \mid A)=\frac{40}{100} \frac{1}{2}+\frac{36}{200} \frac{1}{2}=29 \% \\
& P(\text { resp } \mid B)=\frac{30}{100} \frac{1}{2}+\frac{15}{100} \frac{1}{2}=22.5 \%
\end{aligned}
$$

