Bridging biostatistics and health economics

Mark Clements, Xiaoyang Du and Mattias Rantalainen

2025-03-27

Introduction: Who am I?

- Short summary: a New Zealander who trained in Australia and married a Swede
- Lektor in biostatistics at KI
- My research group: Modelling and simulation with applications to cancer epidemiology and health economics
- Some of my R/C++ packages:

`rstpm2`	flexible parametric survival models
`microsimulation`	discrete event simulation, with tools for health economics
`prostata`	simulation model for prostate cancer screening (`https://github.com/mclements/prostata`)
`magree`	multi-rater agreement of ordinal outcomes (Fortran-77:)
`biostat3`	datasets and teaching material for survival analysis
`collett`	datasets and teaching material for Collett (2023)

Introduction: Evaluation of AI-assisted breast pathology for treatment assignment

- Deep learning algorithms using digital histopathology images show considerable promise for clinical risk predictions.
- Our broader research group has developed a CE-IVD approved AI algorithm (Stratipath Breast) for use in risk stratification for intermediate risk breast cancer patients.
  - The algorithm is fitted to discriminate between Nottingham Histological Grade (NHG) 1 (low risk) and NHG 3 (high risk) patients in order to further risk stratify intermediate risk NHG 2 patients.
  - Validating the AI algorithm on progression-free survival, Sharma et al. [2023] found that a median split in risk predictions lead to a hazard ratio of 1.99 (95% confidence interval 1.10, 3.60). This suggests assigning chemotherapy to NHG 2 patients with higher predicted risks.
  - COI: Mattias Rantalainen is a co-founder for the Stratipath company.
- Can we develop and apply methods for cohort-based continuous time Markov multi-state models to assess the cost-effectiveness for the Stratipath Breast algorithm?

Introduction: Outline of a health technology assessment (HTA) I

In contrast to the deep learning algorithm, the HTA models are simpler:

- Specify the strategies that are of interest
  - AI algorithm: assign chemotherapy to those at higher predicted risks
  - Reference: arbitrarily assign chemotherapy (alternatively, no-one or everyone)
- Assume a small set of health states
  - For those not assigned chemotherapy: Baseline, progression and death
  - For those assigned chemotherapy: Baseline, progression, adverse events from chemotherapy and death
  - Comment: the predicted risk groups are used to define the natural history model (yes, this is circular:)
- Model the transition intensities between those states
- Represent the utilities (values typically between 0 and 1) for each of those states

Introduction: Outline of a health technology assessment (HTA) II

- Represent the costs for each of those states
- Using a model and a cost-effectiveness threshold τ (cost per unit of utility gained), calculate long-term costs C, utilities U and net monetary benefit $\text{NMB} = U\tau - C$ for each strategy and compare the strategies.

Introduction: Outline of a health technology assessment (HTA) III

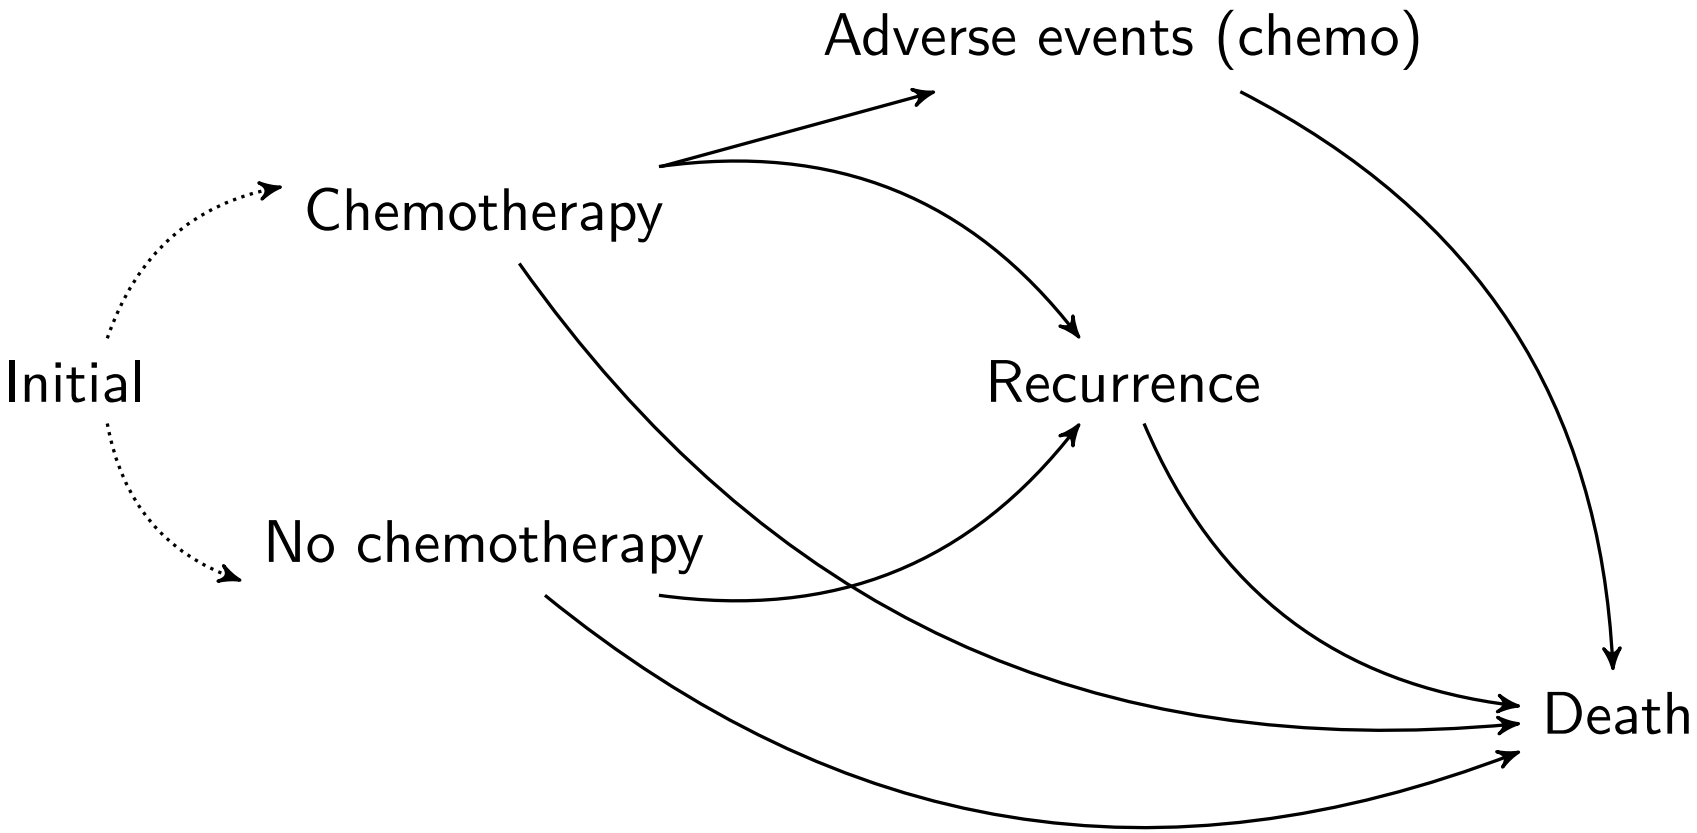

Grade for artistic merit: 2/10.

Introduction: Model classes for HTA

For HTA, it is common to use:

1. Cohort-based partitioned survival models : sequential states and use survival differences
2. Cohort-based discrete time Markov multi-state models
   - Issue: Reasoning about effects is often easier on the hazards scale (e.g. hazard ratios) rather than on the probability scale.
   - Issue: Practitioners often add probabilities, which ignores competing events and multiple events within a time unit.
3. Individual-based continuous time multi-state models, which are computationally expensive

In biostatistics, a more natural alternative would be a cohort-based continuous time Markov multi-state model.

Methods: Markov decision process

- Our cost-effectiveness analysis can be considered as a continuous-time Markov decision process with the four-tuple $(\mathcal{S}, \mathcal{A}, \boldsymbol{Q}, R)$, where:
  - $\mathcal{S}$ is a finite and discrete state space
  - $\mathcal{A}$ is a finite and discrete set of actions (or strategies)
  - $\boldsymbol{Q}(t|a)$ is a continuous-time transition intensity matrix at time t between states $s \in \mathcal{S}$ that depends on the action $a \in \mathcal{A}$
  - $R(t|a)$ is a reward at time t that depends on the action $a \in \mathcal{A}$

Methods: Kolmogorov's forward differential equations

The transition probabilities can be modelled as follows:

- Let $P_{ij}(t_0, t|a)$ be the proportion in state j at time t given in state i at time t_0 and given action a; let $\boldsymbol{P}(t_0, t|a) = (P_{ij}(t_0, t|a))$ be the transition probability matrix.
- Let the transition intensities from state i to state j $(i \neq j)$ at time t for action a be $Q_{ij}(t|a)$; let $Q_{ii}(t|a) = -\sum_{j \neq i} Q_{ij}(t|a)$; let $\boldsymbol{Q}(t|a) = (Q_{ij}(t|a))$ be the transition intensity matrix.
- Initial values are given for $\boldsymbol{P}(t_0, t_0|a) = \boldsymbol{I}$
- Kolmogorov's forward differential equations are then

$$\frac{\partial \boldsymbol{P}(t_0, t|a)}{\partial t} = \boldsymbol{P}(t_0, t|a)\boldsymbol{Q}(t|a) \tag{1}$$

- The state occupation probabilities of being in state i at time t can then be calculated from the initial values $\boldsymbol{\pi}(t_0, t_0|a)$ at t_0 and action a and then for time t from the transition intensity matrix by $\boldsymbol{\pi}(t_0, t|a) = \boldsymbol{P}(t_0, t|a)^T \boldsymbol{\pi}(t_0, t_0|a)$.

R code for state occupation probabilities for an illness-death model I

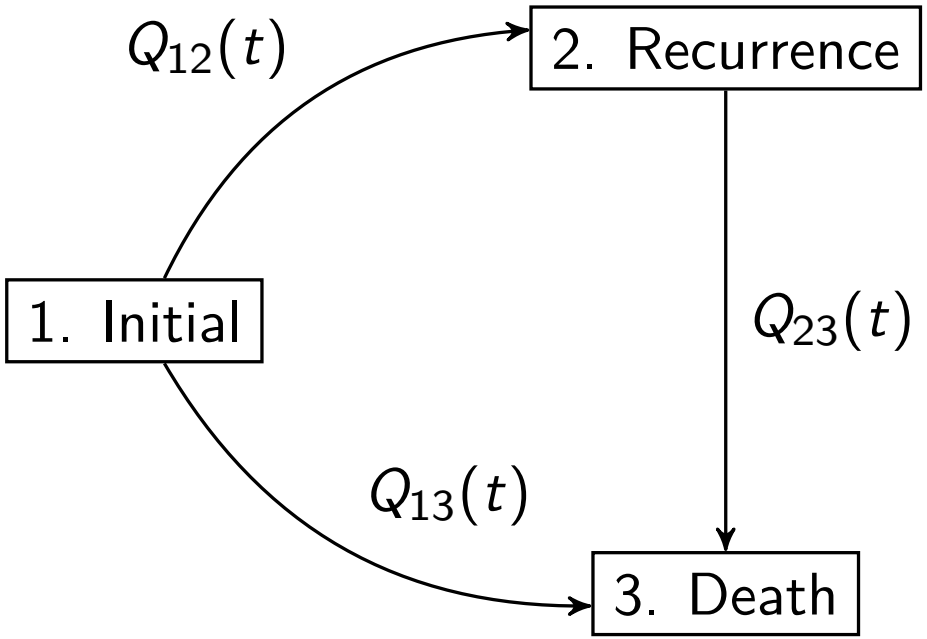

R code for state occupation probabilities for an illness-death model II

```
library(survival) # colon, Surv
library(rstpm2)   # markov_msm
nms = c("Initial","Recurrence","Death")
tmat = matrix(c(NA,Q12=1,Q13=2,
                NA,NA,Q23=3,
                NA,NA,NA),3,byrow=TRUE,
              dimnames=list(nms,nms))
recurrence = subset(survival::colon, etype==1) |>
    stpm2(Surv(time,status)~rx, data=_, df=3)
ode1 = markov_msm(list(recurrence, # Q12
                       \(time) 0.02/365, # Q13
                       \(time) 0.2/365), #  Q23
                  trans=tmat,
                  t=seq(0,10*365,length=101L),
                  newdata=data.frame(rx=c("Obs","Lev","Lev+5FU")),
                  tmvar="time")
plot(ode1, strata=~rx, lattice=TRUE, xlab="Time from study entry (days)")
```

R code for state occupation probabilities for an illness-death model III

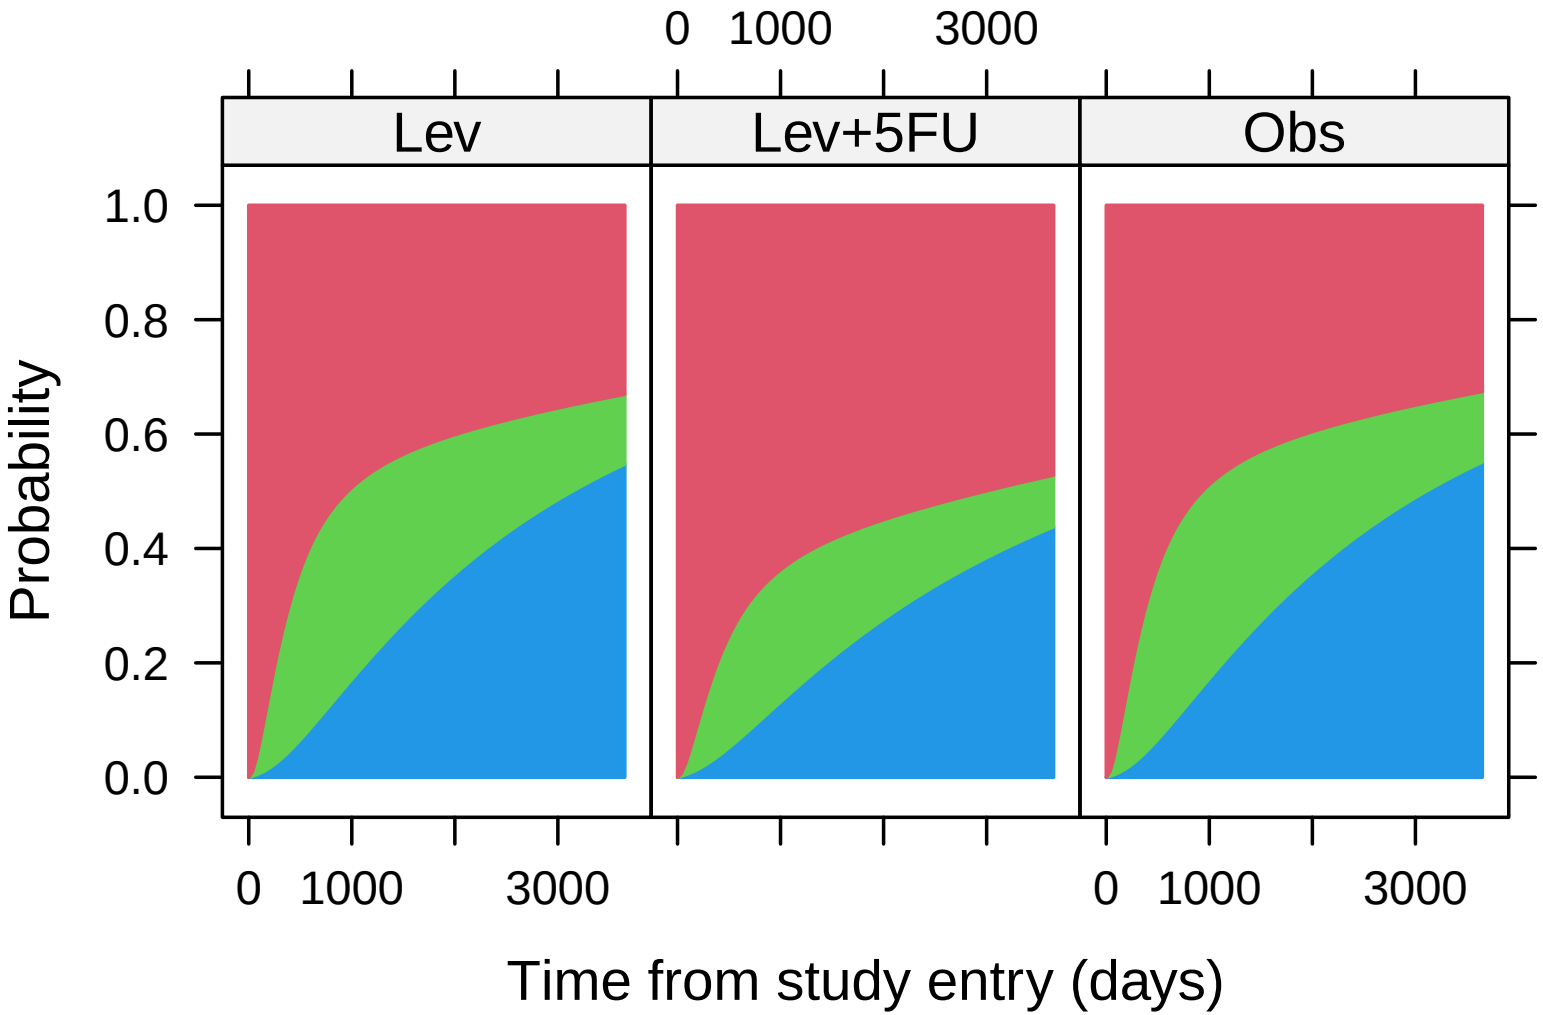

Methods: Rewards using net monetary benefit

- We propose calculating net monetary benefit by adding further ordinary differential equations.
- Let the cost-effectiveness threshold be τ (e.g. 50,000 € per quality-adjusted life-year (QALY) gained).
- For utilities, let the discount rate be λ_u (e.g. 0.03 or $\log(1 + 0.03)$), with state-specific utilities $u_i(t|a)$ for state i at time t for action a. Let the cumulative expected utilities (that is, QALYs) to time t and action a be represented by $U(t|a)$, where $U(0|a) = 0$. Then

$$\frac{\partial U(t|a)}{\partial t} = e^{-\lambda_u t} \sum_i u_i(t|a)\pi_i(t_0, t|a) \tag{2}$$

Methods: Rewards using net monetary benefit

- For costs, let the discount rate be λ_c, with costs per person-time of $c_i(t|a)$ for state i and costs per transition of $D_{ij}(t|a)$ for transitions from state i to state j at time t and action a. Let the cumulative expected costs be represented by $C(t|a)$, where $C(0|a) = 0$. Then

$$\frac{\partial C(t|a)}{\partial t} = e^{-\lambda_c t} \sum_i \pi_i(t_0, t|a) \left(c_i(t|a) + \sum_{j \neq i} D_{ij}(t|a) Q_{ij}(t|a) \right) \tag{3}$$

- At the maximum time $t_{\max}$, define the net monetary benefit with cost-effectiveness threshold τ as the reward, such that

$$R(t_{\max}|a) \equiv \text{NMB}_a(\tau) \equiv U(t_{\max}|a)\tau - C(t_{\max}|a)$$

Methods: Parameter uncertainty

- The most common approach in health economics is to bootstrap or resample for the parameters for the transmission intensity models, utilities and costs to investigate the uncertainty in the net monetary benefit.
- A Bayesian approach would be elegant. Ordinary differential equations can be implemented in Stan, JAGS and WinBUGS. TODO.
- We propose an alternative approach using the multivariate delta method. This requires gradients for NMB with respect to the uncertain parameters.
  - We can calculate those gradients using finite differences.
  - Alternatively, we can add sensitivity equations to the set of ordinary differential equations to calculate the gradients (additional slide).

Methods: Multivariate delta method and acceptability

- Let the block-diagonal variance-covariance matrix for the stacked parameters be $\boldsymbol{\Sigma}$. We are now in the position to use the multivariate delta method, such that

$$\text{var}(\boldsymbol{R}(t_{\text{max}}|\boldsymbol{a})) = \boldsymbol{R}'_m(t_{\text{max}}|\boldsymbol{a})^T \boldsymbol{\Sigma} \boldsymbol{R}'_m(t_{\text{max}}|\boldsymbol{a})$$

  for a set of actions $\boldsymbol{a}$.

- We can also calculate the acceptability probability, which is the proportion of the posterior distribution where strategy j is optimal:

$$Pr(R(t_{\text{max}}|a_j) = \max_k R(t_{\text{max}}|a_k))$$

  which can be calculated by integration.

Software implementations

- Equations 1-3 can be easily solved using the `deSolve` package in R.
- We have implemented those equations together with sensitivity equations in the `rstpm2` package in R.
- However, those implementations are comparatively slow. We have completed an implementation that extends the `hesim` package in R/C++ which will allow for bootstrapping in practical times. See `https://github.com/mclements/hesim`.

Reflections: What have I learnt from health economists?

- Clearly define which strategies are important – they have a strong policy and decision-making perspective
- Use a lifetime perspective
- Consider all downstream health events
- Cost calculations are non-trivial:)
- Calculations of health state values are non-trivial:)
- An evaluation is specific to a population – specificity vs generalisability

Reflections: Why do health economists need us (that is, biostatisticians)?

- We have a deeper understanding of:
  - Data
  - Study design
  - Estimation
  - Test characteristics
  - Uncertainty
  - Modelling and simulation
  - Bias-variance trade-offs
  - Computations

Discussion

- We have proposed and implemented a class of continuous time multi-state models using ordinary differential equations with applications to HTA (and other domains).
- We will use this model class to evaluate the cost-effectiveness of the Stratipath Breast algorithm.
- We will push changes to the `hesim` package for our implementation.
- Potential extension to semi-Markov models using Kolmogorov's forward integro-differential equations [Buchardt et al., 2015].
- To conclude: I strongly encourage you to work more closely with health economists. My experience has been very positive:)

Acknowledgements and references

- Funders: VINNOVA, Cancerfonden, Swedish Research Council.
- Contributing colleagues: Irma Fredrisksson, Keith Humphreys, Paul Lambert, Shuang Hao.

References

K. Buchardt, T. Møller, and K. B. Schmidt. Cash flows and policyholder behaviour in the semi-Markov life insurance setup. *Scandinavian Actuarial Journal*, 2015(8): 660–688, 2015.

A. Sharma, S. K. Lovgren, K. L. Eriksson, Y. Wang, S. Robertson, J. Hartman, and M. Rantalainen. Validation of an AI-based solution for breast cancer risk stratification using routine digital histopathology images. *medRxiv*, 2023. doi: 10.1101/2023.10.10.23296761. URL `https://www.medrxiv.org/content/early/2023/10/10/2023.10.10.23296761`.

Additional slide: Sensitivity equations

- For a derived parameter $\psi(\boldsymbol{\theta})$ which is a function of a vector $\boldsymbol{\theta}$ of parameters, define $\psi'_m = \dfrac{\partial \psi}{\partial \theta_m}$. We can then calculate:

$$\frac{\partial \boldsymbol{P}'_m(t_0, t|a)}{\partial t} = \boldsymbol{P}'_m(t_0, t|a)\boldsymbol{Q}(t|a) + \boldsymbol{P}(t_0, t|a)\boldsymbol{Q}'_m(t|a)$$

$$\boldsymbol{\pi}'_m(t_0, t|a) = \boldsymbol{P}'_m(t_0, t|a)^T \boldsymbol{\pi}(t_0, t_0|a)$$

$$\frac{\partial U'_m(t|a)}{\partial t} = e^{-\lambda_u t} \sum_i \left(u'_{im}(t|a)\pi_i(t_0, t|a) + u_i(t|a)\pi'_{im}(t_0, t|a) \right)$$

$$\begin{aligned}\frac{\partial C'_m(t|a)}{\partial t} = e^{-\lambda_c t} \sum_i \Bigg(& \pi'_{im}(t_0, t|a)\left(c_i(t|a) + \sum_j D_{ij}(t|a)Q_{ij}(t|a) \right) + \\ & \pi_i(t_0, t|a)\left(c'_{im}(t|a) + \sum_j (D'_{ijm}(t|a)Q_{ij}(t|a) + D_{ij}(t|a)Q'_{ijm}(t|a)) \right) + \\ & \sum_k \left(\pi'_{km}(t_0, t|a)Q_{ki}(t|a)f_i(t|a) + \pi_k(t_0, t|a)Q'_{kim}(t|a)f_i(t|a) + \pi_k(t_0, t|a)Q_{ki}(t|a)f'_{im}(t|a) \right) \Bigg)\end{aligned}$$

$$R'_m(t|a) = U'_m(t|a)\tau - C'_m(t|a)$$