Journal of Radiation Research, 2013, 54, i49–i55
doi: 10.1093/jrr/rrt040

Data-driven Markov models and their application in the evaluation of adverse events in radiotherapy

Daniel ABLER[1,2,*], Vassiliki KANELLOPOULOS[1,3], Jim DAVIES[4], Manjit DOSANJH[1], Raj JENA[5], Norman KIRKBY[3] and Ken PEACH[2]

[1]CERN, CH1211 Geneva 23, Switzerland
[2]Department of Physics, University of Oxford, Oxford, OX1 3RH, UK
[3]Faculty of Engineering and Physical Sciences, University of Surrey, Guildford, Surrey GU2 7XH, UK
[4]Department of Computer Science, University of Oxford, Oxford, OX1 3QD, UK
[5]Cambridge University Hospitals NHS Foundation Trust, Hills Road, Cambridge, CB2 0QQ, UK
*Corresponding author. CERN, 28/1-003, CH1211 Geneva 23. Tel: +41-22-76-71456; Fax: +41-22-76-68158; Email: daniel.abler@cern.ch

(Received 21 January 2013; revised 8 March 2013; accepted 21 March 2013)

Decision-making processes in medicine rely increasingly on modelling and simulation techniques; they are especially useful when combining evidence from multiple sources. Markov models are frequently used to synthesize the available evidence for such simulation studies, by describing disease and treatment progress, as well as associated factors such as the treatment's effects on a patient's life and the costs to society. When the same decision problem is investigated by multiple stakeholders, differing modelling assumptions are often applied, making synthesis and interpretation of the results difficult. This paper proposes a standardized approach towards the creation of Markov models. It introduces the notion of 'general Markov models', providing a common definition of the Markov models that underlie many similar decision problems, and develops a language for their specification. We demonstrate the application of this language by developing a general Markov model for adverse event analysis in radiotherapy and argue that the proposed method can automate the creation of Markov models from existing data. The approach has the potential to support the radiotherapy community in conducting systematic analyses involving predictive modelling of existing and upcoming radiotherapy data. We expect it to facilitate the application of modelling techniques in medical decision problems beyond the field of radiotherapy, and to improve the comparability of their results.

Keywords: decision analytic modelling; Markov model; health informatics; adverse events; radiotherapy

INTRODUCTION

Normal tissue complications play an important role in the evaluation of radiation therapy. Treatment decisions are not only based on the prognosis of tumour control but on other possible side effects, 'adverse events' (AE), on the patient: severe limitations of the function of organs or tissues in vicinity of the irradiated tumour which can have a permanent impact on the patient's quality of life.

While important for medical decision-making, realistic long-term data on such effects is difficult to obtain, mainly due to the limited duration of medical studies. In situations where relevant real-world data is not available, modelling provides an analytic framework for the synthesis of evidence and allows extrapolation of intermediate clinical endpoints to final outcomes or outcomes beyond the duration of clinical studies [1]. Investigating the 'optimal' distribution of limited health care resources often involves comparing the value ('utility') of an outcome (a 'health state') to the costs for achieving it [2]. Models for medical decision-making therefore typically use a simplified representation of the underlying disease or intervention process, together with cost and utility information.

Depending on the problem to be modelled, the most suitable approach has to be chosen from a variety of established modelling structures [3], the decision model has to be defined by experts to fit the given decision problem and the model parameters may then be determined based on available data sources, literature or expert opinion [4, 5].

However, in situations where the same decision problem is investigated by multiple stakeholders, synthesis and interpretation of the results may be difficult: e.g. the literature on cost-effectiveness of particle therapy was found to be 'non-comparable' due to differences in modelling assumptions and model parameters [6]. Homogenization of modelling assumptions may therefore increase the comparability of simulation outcomes and, indeed, similar model structures can be used to investigate a family of related questions, such as cost-effectiveness of an intervention for multiple distinct tumour entities [7].

This paper attempts to formalize that approach by introducing the notion of 'general models' as generic definitions of model structures, describing the common features of a family of models. So-called Markov models are among the most commonly used types of models for health technology assessment [8] and are particularly suitable for describing problems that involve repeating events over time and do not require interaction between individuals [3, 9].

The paper therefore focuses on 'general models' for Markov models and proposes a methodology for their derivation from existing data, based on 'Markov Model Templates' (MMT), a high level description of health states and their relations.

Figure 1 illustrates the proposed approach: the characteristics of a 'general Markov model' are expressed in computable form by the MMT language. Any such description, an 'MMT instance', can be processed by a programme (denoted 'MMT interpreter' in Fig. 1) which constructs Markov model states from data available in a data repository, following the rules specified by the MMT instance, and computes transition probabilities among the resulting Markov states. Depending on the given data, the same 'general Markov model' may result in distinct specific Markov models as indicated in Fig. 1.

This standardized approach towards the creation of Markov models may, for example, facilitate the systematic analysis of existing and future radiotherapy data by generating specific Markov models for a range of parameters, such as different tumour entities or beam qualities.

In the following, a general Markov model for the evaluation of adverse events in radiotherapy and a specification language (MMT) for general Markov models are developed. The expressiveness of the MMT language is then demonstrated using the example of the general adverse event model.

MATERIALS AND METHODS

Markov models for medical decision-making

In medical decision-making Markov models are often used to describe the disease progress of chronic diseases, represented by a finite number of discrete and mutually exclusive 'health states' that are connected by 'transitions' corresponding to clinically important events [10, 11]. Transition probabilities among them express the likelihood for a patient to change from one health state to another.

Measures for cost and outcome can be associated with states and transitions to account for resource use over the course of treatment and the gain or loss of life quality as perceived by the patient. Based on this information, cost and outcome of a particular treatment strategy can be estimated by simulating the path of a fictitious patient population through the model (detailed example in [11]). Transition probabilities are evaluated in fixed time intervals over a given number of 'cycles', and utilities and costs according to transitions and state occupation are accumulated. Depending on the evaluation type, costs, utilities and transition probabilities can be applied deterministically, according to a point estimate, or probabilistically, taking

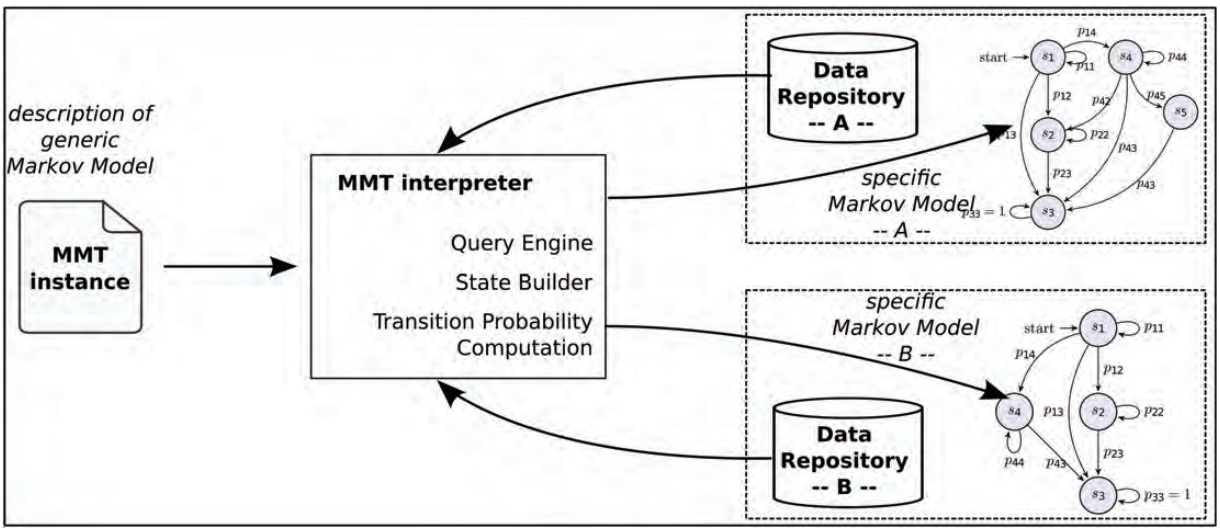

Fig. 1. Automatic creation of Markov models from distinct data sources, based on common Markov Model Template (MMT).

into account the variability of the respective measure in every evaluation.

Generalized Markov model for adverse events

We introduce a general Markov model that describes the course of treatment-related side effects after radiotherapy. The model allows the analysis of side effects across a range of tumours, which is of interest when examining global toxicity levels associated with a specific treatment technique or technology. This model combined with a cost model of treatment costs and quality of life could for example be used to assess the benefit of new radiotherapy technologies, such as volumetric modulated arc therapy (VMAT), CyberKnife and proton beam therapy.

Figure 2 shows the generic Markov model. The model is divided into two sub-models. The 'acute AE model' describes 'acute' side effects that occur during and within six weeks after treatment. The 'chronic AE model' illustrates the course of 'late' effects after treatment. Since it is envisaged that data is used from treatments comprising a variety of tumour volumes and treatment doses, side effects are grouped coarsely into 'mild' and 'severe' effects, where 'mild' covers toxicity grades less than 3 including 'no side effect', and 'severe' comprises toxicity Grades 3 and 4. All examples presented here are based on the Common Terminology Criteria for Adverse Events (CTCAE). The health states in the acute sub-model are possible combinations of acute side effects (a1; a2; a1,a2; etc.) and treatment-related death. The late-effects model includes possible combinations of late effects (c1; c2; c1,c2; etc.). Over time, patients may develop further chronic adverse events, progressing into health states with a higher number of adverse events or into the 'Death' state. When examining treatments that differ in outcome, the health states after treatment need to be paired with the characteristic of either having 'tumour control' or having 'no tumour control'.

The length of a Markov cycle is ideally chosen to be the shortest clinically meaningful time interval [8], in the case of chronic adverse events approximately one year.

Language for generalized Markov models

In order to be able to express the conditions for states and relations within a general Markov model, the MMT language was developed using the Spoofax Language Workbench [13]. Figure 3 illustrates the main concepts of this language.

'State Groups' specify groups of Markov health states that share common characteristics such as the 'severe adverse event' states or the 'death' state of Fig. 2. Transitions between distinct 'State Groups' are defined by 'State Group Transitions'.

The MMT language defines patient populations and applicable health states by queries over data sources. To ensure that a given MMT instance, specifying a general Markov model, is applicable to a wide range of data sources, only minimum assumptions on the logical representation of clinical information in those data sources have been made. We assume that the clinical information available about a patient can be factored into 'clinical contexts' (diagnosis, treatment, follow-up, etc.), 'observations' within such clinical contexts, and 'medical findings' recorded in each of those observations. Depending on the data source, the context of clinical data for each patient can be modelled explicitly or be inferred, e.g. by the temporal relation

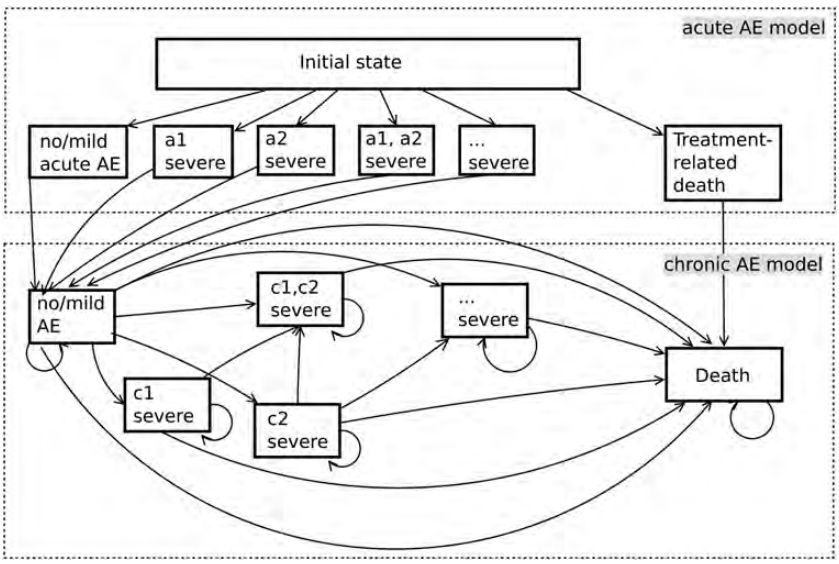

Fig. 2. State transition diagram for a general Markov model describing adverse events after radiotherapy. Generalization from [12].

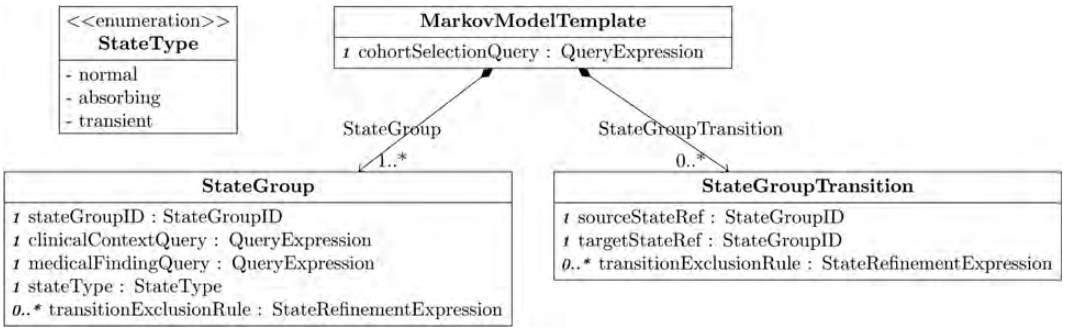

Fig. 3. Main concepts of 'Markov Model Template' (MMT) language for the description of generic Markov models.

between treatment and observation. A medical finding in this model corresponds to an attribute-value pair in a data repository and can be identified by means of a query expression.

Every Markov model reflects the disease progression of a certain patient population, with specific demographic characteristics and medical conditions, undergoing a well-defined treatment. The 'Cohort Selection Query' allows preselecting the patient population accordingly for any MMT instance.

We define Markov health states as sets of medical findings, recorded in a single observation. Using this definition, health states can be specified by query expressions over the clinical context in which an observation had been made (e.g. follow-up after radiation therapy), 'Clinical Context Query', and the medical findings (e.g. severe adverse event), 'Medical Finding Query', found in any observation in that context. Depending on the specificity of the query expressions and the granularity and completeness of the respective data source, the queries defining a 'State Group' may identify zero, one or many medical findings for every patient and observation. Since any combination of medical findings forms one distinct health state, a single 'State Group' definition may result in multiple specific Markov health states. All these Markov states are members of that 'State Group', 'parameterized' by their unique combination of medical findings.

'State Type' defines a property of all health states within a 'State Group' that restricts the possible transitions between the members of that group. While 'normal' states allow transitions to all other members and to the states themselves, 'transient' states can only be occupied for a certain time period by any individual and therefore are modelled (as a sequence of) states without transitions back into the respective state itself. Once entered, 'absorbing' states, such as the 'Death' state in Fig. 2, cannot be left and therefore only allow transitions to themselves with a probability of one. In addition to defining the 'State Type' of members of any 'State Group', further modelling assumptions can be expressed by excluding certain transitions within a 'State Group' *a priori* and irrespective of the information the data source may provide. For example, in the generic adverse event model (Fig. 2) we may want to include a modelling assumption that any chronic adverse event a patient experiences will persist over the patient's lifetime. 'Transition Exclusion Rules' allow adding 'State Refinement Expressions' to any 'State Group' for defining such constraints. These expressions apply only for 'parameterized states' and allow the specification of forbidden transitions by formulating conditions on 'Source State' and 'Target State'.

Normally, not all states required in a Markov model can be expressed by a single query, so that multiple 'State Groups' need to be defined for the generic description of the full Markov model. A 'State Group Transition' allows the definition of transitions between members of two state groups by referencing the identifier of a source and target 'State Group'. In case the referenced 'State Group' contains more than one state, it is assumed that transitions between all members of the source and target state group are possible. Again, 'Transition Exclusion Rules' provide a way to restrict the possible transitions between individual health states within source and target 'State Groups'.

RESULTS

Figures 4 and 5 show the generic Markov model for adverse event analysis (Fig. 2), expressed in terms of the MMT language. The generic model consists of two distinct MMT instances, corresponding to Markov models for acute (Fig. 4) and chronic (Fig. 5) adverse events, respectively. 'Query Expressions' in this example are defined in the Structured Query Language (SQL) and are based on a specific database schema. The schema was developed to document patient-, tumour- and treatment-related information as well as adverse event reports after radiotherapy [14].

Since both models describe distinct clinical contexts, 'during treatment' for acute adverse events and 'after treatment'/'follow-up' for chronic adverse events, their MMT instances use different 'Clinical Context Query' expressions in their 'State Group' definitions, based on the temporal relation between observation date and the end of radiotherapy treatment.

```
MarkovModelTemplate GenericAEMarkovModelAcuteSevereAEs{
        // all patients in database, no further constraints on patient cohort
        cohortSelectionQuery : "SELECT  PATIENT.ID_Patient FROM db.PATIENT;"
        // all severe acute AEs, 'noAE' state corresponds to empty return set
        StateGroup severeAcuteAEs {
                clinicalContextQuery : "db.OBSERVATION.Date_Recorded < db.RADIOTHERAPY.Treatment_End"
                medicalFindingQuery : "db.SIDEEFFECT.Score > 2  AND db.SIDEEFFECT.Score < 5"
                stateType : absorbing
        }
        // 'death' state corresponds to AE score = 5
        StateGroup treatmentRelatedDeath{
                clinicalContextQuery : "db.OBSERVATION.Date_Recorded < db.RADIOTHERAPY.Treatment_End"
                medicalFindingQuery : "db.SIDEEFFECT.Score = 5"
                stateType : absorbing
        }
        // no transitions between acute AE and death state groups (StateType 'absorbing' ) since
        // state populations only represent single observation snapshot (very short time interval).
}
```

Fig. 4. MMT representation of generic Markov model for acute severe adverse events.

```
MarkovModelTemplate GenericAEMarkovModelChronicSevereAEs{
        // all patients in database, no further constraints on patient cohort
        cohortSelectionQuery : "SELECT  PATIENT.ID_Patient FROM db.PATIENT;"
        // all severe chronic AEs, 'noAE' state corresponds to empty return set
        StateGroup severeChronicAEs {
                clinicalContextQuery : "db.OBSERVATION.Date_Recorded > db.RADIOTHERAPY.Treatment_End"
                medicalFindingQuery : "db.SIDEEFFECT.Score > 2  AND db.SIDEEFFECT.Score < 5"
                // allow transitions from every state identifed by query to itself and to any other state within stateGroup
                stateType : normal
                // exclude all transitions for which medical findings in source state are not present in target state
                transitionExclusionRule : not(sourceState SubSetOf targetState)
        }
        StateGroup death{
                clinicalContextQuery : "db.OBSERVATION.Date_Recorded > db.RADIOTHERAPY.Treatment_End"
                medicalFindingQuery : "db.SIDEEFFECT.Score = 5"
                // death is 'absorbing state'  i.e. only transition to itself  allowed
                stateType : absorbing
        }
        StateGroupTransition TransitionChronicAeToDeath{
                sourceStateRef :  severeChronicAEs
                targetStateRef :   death
                // example of  transitionExclusionRules in St ateGroupTransitions:
                // exclude transitions from AE state 'no AE' , i.e. where sourceState query returned  no  medical finding,
                // to any other state in the target StateGroup
                transitionExclusionRule : |sourceState| = 0
        }
}
```

Fig. 5. MMT representation of generic Markov model for chronic severe adverse events.

The 'State Groups' in Figs 4 and 5 are further distinguished by their 'Medical Finding Queries'. Health states corresponding to a combination of severe adverse events, such as (a1; a2; a1,a2; etc.) in Fig. 2, are defined by queries for adverse events recorded with severity Grade 3 or 4. When multiple adverse events for every patient and observation are found to match these criteria, the 'State Group' definition results in multiple Markov health states, as illustrated in Fig. 2. Health states corresponding to the death of a patient, if treatment-related, are defined by the existence of an adverse event with severity score 5.

A 'Transition Exclusion Rule' in the MMT 'State Group', representing chronic adverse event states (Fig. 5), expresses the modelling requirement that any chronic adverse event, suffered once, will persist over the lifetime of the patient.

From these queries, the relevant health states for a Markov model can be determined, describing the disease progress of the selected patient population.

A proof-of-concept implementation of the 'state builder' component in the 'MMT interpreter' in Fig. 1 successfully operates with the following selection algorithm. First, the data of all individuals fulfilling the 'Cohort Selection Query' is partitioned into distinct observations, from which the ones relevant to a particular 'State Group' are selected according to the 'Clinical Context Query'. For each of those observations, medical findings corresponding to the 'Medical Finding Query' are identified: every unique combination of medical findings forms one distinct Markov health state within that 'State Group'.

Transition probabilities can then be calculated from the identified Markov health states and the transition rules

specified in 'State Group' and 'State Group Transition', following well-established procedures such as [15, 16].

DISCUSSION

We introduced the notion of 'general Markov models', developed a language for their specification and demonstrated their application for adverse event analysis in radiotherapy. The proposed method allows the encoding of the medical understanding of a disease in computable definitions of health-states and transitions, expressed as query and constraint expressions. These definitions can be used to automatically create relevant Markov health states and transitions based on available data and to compute the transition probabilities.

Feasibility to identify model structure from information specified in MMT instances was successfully tested on a generated toxicity dataset; the module for the computation of transition probabilities between states is under development.

In agreement with modelling recommendations such as [5, 17], the structure of the resulting Markov model is directly determined by a modeller's choice (through the MMT instance), and needs to be designed according to current medical knowledge of the health condition. As recognized by [4, 5], details of the final structure of the model may be subject to the availability of data, since this directly determines the realized health states.

While Markov models generated in this way may not be of completely general validity, they provide an accurate abstraction of the disease progress observed within the dataset used for their generation. As such, they can be used for making predictions within the context in which they have been created.

We anticipate the proposed general model for adverse events to be particularly useful for the analysis of historic clinical follow-up data, where both acute and long-term toxicities may have been reported only to a low level of detail. Its use for this purpose remains to be validated.

The model can easily be adapted to reflect the granularity and type of data available. For newer and prospective studies, the model will be able to make use of traditional objective toxicity data, collected according to reporting recommendations that propose the recording of specific objective endpoints for long-term toxicities [18], as well as subjective patient reported outcome measures (PROMs). Since subjective symptom clustering from validated PROM tools has been shown to be able to act as a sensitive predictor of long-term toxicity [19], we envisage including both forms of measurement, objective measures and PROMs, into the Markov model.

Application of the proposed approach to Markov model generation in practice requires further development of information models and software tools: examples in Figs 4 and 5 assume information to be organized in a fixed data model against which queries for the definition of health states can be formulated. To facilitate the re-use of Markov model templates, however, the definition of health states must be possible regardless of the underlying data model. Query expressions may be formulated against a simplified data schema, and mappings supplied to the MMT interpreter in order to work with distinct data sources.

When a suitable query translation mechanism has been established, the same generic MMT can be applied to different (interoperable) datasets, resulting in multiple Markov models with differing parameters but adhering to the same modelling principles. Other parameters, such as costing or quality of life information associated with states and transitions need to be determined and synthesized into the generated Markov model. After the design of the initial general Markov model, the generation of specific Markov models does not require additional modelling experience, thus lowering the expertise required for creating decision analytic models.

Making the resulting Markov model more accessible for decision analytic simulation studies requires a further language or exchange format for Markov model simulations, such as [20], that specifies all the information relevant to a simulation task in computable form. Software platforms for decision analytic modelling and simulations can then be used to compute the outcome of a simulation study based on the generated Markov models.

We expect such an automated approach for the creation of Markov models to facilitate the application of modelling techniques in medical decision problems and to improve the comparability of their results.

FUNDING

This work was supported by two Marie Curie Initial Training Fellowships of the European Community's Seventh Framework Programme under contract number [PITN-GA-2008-215840-PARTNER].

REFERENCES

1. Buxton MJ, Drummond MF, Van Hout RL *et al*. Modelling in ecomomic evaluation: an unavoidable fact of life. *Health Econ* 1997;**6**:217–27.
2. Sackett DL, Torrance GW. The utility of different health states as perceived by the general public. *J Chronic Dis* 1978;**31**:697–704.
3. Brennan A, Chick SE, Davies R. A taxonomy of model structures for economic evaluation of health technologies. *Health Econ* 2006;**15**:1295–310.
4. Philips Z, Bojke L, Sculpher M *et al*. Good practice guidelines for decision-analytic modelling in health technology assessment: a review and consolidation of quality assessment. P*harmacoEconomics* 2006;**4**:355.

5. Weinstein MC, O'Brien B, Hornberger J *et al*. Principles of good practice for decision analytic modelling in health-care evaluation: report of the ISPOR task force on good research practices – modelling studies. *Value Hlth* 2003;**6**: 9–17.
6. Pijls-Johannesma M, Pommier P, Lievens Y. Cost-effectiveness of particle therapy: current evidence and future needs. *Radiother Oncol* 2008;**89**:127–34.
7. Lundkvist J, Ekman M, Rehn Ericsson S *et al*. Proton therapy of cancer: potential clinical advantages and cost-effectiveness. *Acta Oncol* 2005;**44**:850–61.
8. Stahl JE. Modelling methods for pharmacoeconomics and health technology assessment: an overview and guide. *PharmacoEconomics* 2008;**2**:131–48.
9. Barton P, Bryan S, Robinson S. Modelling in the economic evaluation of health care: selecting the appropriate approach. *J Health Serv Res Policy* 2004;**9**:110–8.
10. Sonnenberg FA, Beck JR. Markov models in medical decision making. *Med Decis Making* 1993;**13**:322–38.
11. Briggs A, Sculpher M. An introduction to Markov modelling for economic evaluation. *PharmacoEconomics* 1998;**13**:397–409.
12. Grutters JPC, Pijls-Johannesma M, De Reuysscher D *et al*. The cost-effectiveness of particle therapy in non-small cell lung cancer: exploring decision uncertainty and areas for future research. *Cancer Treat Rev* 2010;**36**:468–76.
13. The Spoofax Language Workbench. http://strategoxt.org/Spoofax (10 December 2012, date last accessed).
14. Kanellopoulos V. A data model for the follow-up and evaluation of side effects – analysis of extracted relevant clinical and biological data with the intent to test/develop models to enable prediction of outcome and indications for ion therapy. PARTNER Work Package 24 – Deliverable 4. *Technical report*. CERN, 2011.
15. Welton NJ, Ades AE. Estimation of Markov chain transition probabilities and rates from fully and partially observed data: uncertainty propagation, evidence synthesis, and model calibration. *Med Decis Making* 2005;**25**:633–45.
16. Sherlaw-Johnson C, Gallivan S, Burridge J. Estimating a Markov transition matrix from observational data. *J Oper Res Soc* 1995;**46**:405–10.
17. Nuijten C, Pronk MH, Brorens MJA *et al*. Reporting format for economic evaluation: Part II: Focus on modelling studies. *PharmacoEconomics* 1998;**14**:259–68.
18. Bentzen S, Constine L, Deasy J *et al*. Quantitative Analyses of Normal Tissue Effects in the Clinic (QUANTEC): an introduction to the scientific issues. *Int J Radiat Oncol* 2010;**76** (3 Suppl):S3–S9.
19. Gulliford S, Foo K, Morgan R *et al*. Dose-volume constraints to reduce rectal side effects from prostate radiotherapy: evidence from MRC RT01 Trial ISRCTN 47772397. *Int J Radiat Oncol* 2010;**76**:747–54.
20. Abler D, Harris S, Davies J. Meta-modelling Markov model simulations for cost effectiveness analyses. Poster presentation at ICTR-PHE 2012. *Radiother Oncol* 2012;**102**:S91.