

DR. SWEET PING NG (Orcid ID : 0000-0003-1721-0680)

DR. ADAM S GARDEN (Orcid ID : 0000-0001-7926-2641)

DR. JASON M JOHNSON (Orcid ID : 0000-0001-7724-5784)

DR. ERICH M STURGIS (Orcid ID : 0000-0002-1069-4822)

Article type : Original Article

Surveillance Imaging for Patients with Head and Neck Cancer Treated with Definitive Radiotherapy: A Partially Observed Markov Decision Process Model

Sweet Ping Ng MBBS FRANZCR^{1,2}, Temitayo Ajayi MA³, Andrew Schaefer PhD³, Courtney Pollard III MD PhD¹, Houda Bahig MD PhD^{1,4}, Adam S Garden MD¹, David I Rosenthal MD¹, G Brandon Gunn MD¹, Steven J Frank MD¹, Jack Phan MD PhD¹, William H. Morrison MD¹, Jason M Johnson MD⁵, Abdallah Mohamed MD¹, Erich M Sturgis MD MPH⁶, Clifton D Fuller MD PhD¹.

¹Department of Radiation Oncology, Division of Radiation Oncology, The University of Texas MD Anderson Cancer Center, Houston, Texas USA

²Department of Radiation Oncology, Peter MacCallum Cancer Centre, Melbourne, Victoria, Australia.

³Department of Computational and Applied Mathematics, Rice University, Houston, Texas, USA

⁴Department of Radiation Oncology, Centre Hospitalier de l'Université de Montreal, Montreal, Canada

⁵Department of Diagnostic Radiology, The University of Texas MD Anderson Cancer Center, Houston, Texas USA

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the [Version of Record](#). Please cite this article as [doi: 10.1002/CNCR.32597](https://doi.org/10.1002/CNCR.32597)

This article is protected by copyright. All rights reserved

⁶Department of Head and Neck Surgery, Division of Surgery, The University of Texas MD Anderson Cancer Center, Houston, Texas USA

Word Count: 3556 words

Corresponding author:

Dr. Sweet Ping Ng

Peter MacCallum Cancer Centre

305 Grattan Street, Melbourne, Victoria 3000 Australia

Email: sweet.ng@petermac.org

Phone: +61412636394

Preliminary results of this study were presented at the American Head and Neck Society (AHNS) 2019 Annual Meeting in May 2019.

Running head: POMDP Model for Head and Neck Cancer

Keywords: surveillance imaging; head and neck cancer; radiotherapy; markov modeling

Conflict of Interest: None

Dr. Ng is funded by the Australian Postgraduate Award, RSNA Fellow Grant, and RANZCR research grants. Dr. Fuller is a Sabin Family Foundation Fellow. Dr. Fuller receives funding and salary support from the National Institutes of Health (NIH), including: the National Institute for Dental and Craniofacial Research Establishing Outcome Measures Award (1R01DE025248/R56DE025248) and an Academic Industrial Partnership Grant (R01DE028290); a National Science Foundation (NSF), Division of Mathematical Sciences, Joint NIH/NSF Initiative on Quantitative Approaches to Biomedical Big Data (QuBBD) Grant (NSF 1557679); a National Institute of Biomedical Imaging and Bioengineering (NIBIB) Research Education Programs for Residents and Clinical Fellows Grant (R25EB025787-01); the NIH Big Data to Knowledge (BD2K) Program of the National Cancer Institute (NCI) Early Stage Development of Technologies in Biomedical Computing, Informatics, and Big Data

This article is protected by copyright. All rights reserved

Science Award (1R01CA214825); NCI Early Phase Clinical Trials in Imaging and Image-Guided Interventions Program (1R01CA218148); an NIH/NCI Cancer Center Support Grant (CCSG) Pilot Research Program Award from the UT MD Anderson CCSG Radiation Oncology and Cancer Imaging Program (P30CA016672) and an NIH/NCI Head and Neck Specialized Programs of Research Excellence (SPORE) Developmental Research Program Award (P50 CA097007). Dr. Fuller has received direct industry grant support, honoraria, and travel funding from Elekta AB.

Total number of text pages: 24

Total number of tables: 3

Total number of figures: 1

AUTHOR CONTRIBUTIONS

Sweet Ping Ng: Conceptualization, data curation, investigation, methodology, project administration, visualization, writing - original draft, and writing - review and editing. Temitayo Ajayi: formal analysis, methodology and writing - review and editing. Andrew J Schaefer: formal analysis, methodology and writing - review and editing. Courtney Pollard III: data curation, investigation, visualization and writing - review and editing. Houda Bahig: data curation, investigation and writing - review and editing. Adam S Garden: investigation, methodology, writing - review and editing. David I Rosenthal: supervision, resources, writing - review and editing. G Brandon Gunn: resources and writing - review and editing. Steven J Frank: resources and writing - review and editing. Jack Phan: resources and writing - review and editing. William H. Morrison: resources, methodology and writing - review and editing. Jason M Johnson: resources and writing - review and editing. Abdallah SR Mohamed: resources and writing - review and editing. Erich M Sturgis - resources and writing - review and editing. Clifton D Fuller: project administration, supervision, visualization, resources and writing - review and editing.

PRECIS

Our POMDP model suggests that less frequent surveillance scan policies can maintain adequate information on relapse status for patients with HNC treated with radiotherapy. This model could potentially translate to a more cost-effective surveillance program for this

This article is protected by copyright. All rights reserved

group of patients.

Abstract

Purpose: The goal of this model is to guide surveillance imaging policies after definitive radiotherapy.

Methods: A partially observed Markov decision process model was formulated to determine the optimal times to scan patients. Transition probabilities were computed using a dataset of 1508 patients with HNC who received definitive radiotherapy between years 2000 - 2010. Kernel density estimation was used to smooth the sample distributions. The reward function was derived using cost estimates from the literature. Additional model parameters were either estimated using data in the literature or clinical expertise.

Results: When considering all forms of relapse, our model showed that the optimal time between scans is longer than the time intervals used in the institutional guidelines. The optimal policy dictates that there should be less time between surveillance scans immediately following treatment compared to years after treatment. Comparable results also held when only locoregional relapses were considered as relapse events in the model. Simulation results for the inclusive relapse cases showed that less than 15% of patients experienced relapse over a simulated 36-month surveillance program.

Conclusion: This model suggests that less frequent surveillance scan policies can maintain adequate information on relapse status for patients with HNC treated with radiotherapy. This model could potentially translate to a more cost-effective surveillance program for this group of patients.

Introduction

Head and neck cancer (HNC) accounts for approximately four percent of cancer cases in the United States annually ¹. Radiation therapy is one of the curative-intent modalities for HNC. Patients who achieved complete response post-therapy are entered in a surveillance program to monitor for disease recurrence, late effects of therapy and second primary malignancies. Surveillance imaging is performed at regular intervals as dictated by

institutional and/or trial policies. This imaging is performed on the premise that if they can detect asymptomatic or clinically occult recurrences earlier, then salvage is possible, improving patient outcomes. There is limited high-level evidence to support this. Besides, there is no reliable evidence as to the optimal frequency of imaging in the surveillance period. With the increasing number of cancer survivors and rising healthcare costs ^{2,3}, a more practical and effective surveillance model is desirable.

Here, we introduce a possible surveillance model for patients with HNC treated with definitive radiotherapy using a partially observed Markov decision process (POMDPs). POMDPs are an established methodology for estimation of state-dependent alteration probabilities over time and are well suited for determining maximum thresholds/ time-points for specific intervention when incomplete knowledge precludes accurate prediction of an event or injury. In particular, POMDPs feature the stochastic evolution of a process, also found in Markov decision processes, and they also do not assume complete knowledge of the process's state at any given time. Due of these features, POMDPs have been utilized for processes wherein a decision must be made, for example, when to intervene in ischemic cardiac disease, under incomplete knowledge ^{4,5}.

The specific aims of the current study include:

1. Development of a POMDP model of locoregional and distant metastases detection on surveillance imaging and clinical examination for head and neck squamous carcinoma survivors, after definitive radiotherapy/ chemoradiotherapy.
2. Evaluation of model performance and robustness given clinical patient data (i.e., model "real world" performance), using an extant institutional dataset.
3. Optimization of institutional surveillance imaging schedule for survivorship based on the derived POMDP model.

Methods

Decision model

A partially observed Markov decision process (POMDP) enables one to model the surveillance problem mathematically. Markov decision processes (MDPs) model decision-

making over time. In an MDP, the agent (the decision-maker) has complete information of the current state when optimizing his/her choice of action ⁶. In contrast, POMDPs introduce uncertainty about the current state of the system, which is directly relevant to the surveillance problem ⁷. The knowledge of the patient's health is incomplete during most time points in a surveillance program.

Formally, a POMDP is a tuple $(S, A, T, R, h, \gamma, \Omega, O)$, where:

- S is the state space
- $A(s)$ is the set of actions in state $s \in S$
- $T: S \times S \rightarrow \mathbf{R}$ denotes the transition probabilities between states
- $R: S \times A \rightarrow \mathbf{R}$ is the reward function
- $h \in \mathbf{N} \cup \infty$ is the horizon of the model
- γ is the discount factor
- $\Omega(s)$ is the observation space in state $s \in S$
- $O: S \times A \times \Omega(s) \rightarrow \mathbf{R}$ denotes the observation probabilities

The presented models in this study have a finite state space, finite action spaces, and finite observation spaces. We consider both finite and infinite horizon models. An institutional dataset of patients with mucosal HNC treated with radiotherapy was used in this model ⁸. Patients' demographics, tumor and treatment characteristics, number of surveillance imaging and modality, and disease recurrence details were provided.

Assumptions

In order to define a tractable model, we make some assumptions. We provide justifications of these assumptions below.

- A) Patients begin the process healthy.
- B) Patients do not recover. A patient only transition to more severe health states or remains in the same health state.
- C) A single "healthy" state designation is sufficient. The health status of the patient is "healthy" if they have not relapsed. No other diseases are considered.
- D) All patients behave the same, according to the specified patient stratification.
- E) There is a maximum number of periods, M , a patient can go without a check-up.

F) Surveillance scans are “perfect,” that is, they detect health and relapse with no error.

The patients included in the dataset (IRB-approved) were all HNC patients who received definitive radiotherapy and achieved complete response post-treatment, and hence, were considered healthy at the initial time point (first post-treatment). This addresses Assumption A. Assumption B is mild as in general, patients do not recover from cancer without treatment. Assumption C is made to simplify the model and emphasize that the goal is to provide surveillance for cancer, not other diseases. We use Assumption D with the knowledge that further stratification of patient type is readily available without impacting the essence of the model. We use a maximum length of time without a scan, Assumption E, in order to ensure patients are seen with some regularity. Finally, the scan accuracy is essentially a parameter in the model, and we chose perfect scans, Assumption F for simplicity. Thus, we emphasize that the surveillance scans in the model do not necessarily represent a specific imaging technique (e.g., PET or CT) or an image of a specific region of the patient. However, this assumption does simplify the model and enables some reflection of decision-making in the surveillance planning process.

Decision Model

The unknown actual health state of the patient during surveillance motivated the use of a POMDP model. The states combined health status (healthy or relapsed) with the amount of time in months since the last scan. There were also terminal states to represent patients who were treated or died. The model was based on a third-party’s perspective in order to consider costs and benefits from all areas. The model was solved using pomdp-solve version 5.4. Both finite horizon and infinite horizon models were solved to represent short- and long-term surveillance policies. We used two different definitions of a relapse event (which corresponded to the states in the POMDP). In the first model, any form of relapse (local, regional, or distant) was considered a relapse event. The transition probabilities were estimated with this definition in mind. In the second model, only locoregional relapses were acknowledged as relapses. Therefore, if a patient experienced a non-locoregional (distant) relapse, the patient transitioned to a terminal state without penalty.

States

The state space S consists of elements s that contain health status information and the time since the last surveillance scan. We define $S = S_{\{healthy\}} \cup S_{\{relapsed\}} \cup \{death\} \cup \{treated\} \cup \{patient\ discovery\}$, where

- $S_{\{healthy\}} = \cup (healthy, t)$, for $t = 1, \dots, M$
- $S_{\{relapsed\}} = \cup (relapsed, t)$, for $t = 1, \dots, M$

A directed graphical model of states and transitions in the POMDP is given in Figure 1. With some probability, the patient may discover relapse (symptomatic or adverse clinical finding on examination), and the “Patient Discovery” state models this occurrence.

Transition Probabilities

The transition probabilities were estimated using the institutional dataset. All patients in the dataset had received definitive radiotherapy and were assumed to be completely healthy at the start of the model (see Assumption A). Patients were censored if their activity was no longer monitored before an event (relapse or death) occurred. Patient’s relapse time was estimated by interpolating between the last known period in which the patient was healthy and the first period in which relapse was discovered. The resulting probability distribution estimates were smoothed out using kernel density estimation with Gaussian kernels, and the bandwidth parameter was set automatically by the software in R version 3.4.0. The bandwidth parameter in kernel density estimation was varied to observe the effect of this parameter (see eFigure 2). However, only the transition probabilities from the default bandwidth were used in the POMDPs. The transition probabilities were then computed by first computing probabilities conditioned on the transition event occurring and then multiplying the resultant distribution by the probability that the event occurred.

Observation Probabilities

The observation probabilities were deterministic based on the state and action. If the patient was alive in the model and the action chosen was to wait, then the patient observed “no information.” In contrast, if the patient was alive and the action chosen was to scan, then the patient observed the underlying health state with perfect accuracy.

Reward Function and Actions

The reward function was determined through data available in the literature to estimate the

costs and benefits of various outcomes. The cost of a surveillance scan and cost of death (for a relapsed patient) was estimated based on data from Huntington et al. ¹⁰. The benefit of detecting cancer was estimated based on data from previous studies ^{11,12}. Details on the possible actions and construction of the reward function are described in Supplements (eMethods and eTable 1). To explore how the costs affected the study results, the model was simulated with larger costs and reward values (eTable 1).

Results

Patient dataset

The dataset comprised 1508 patients with mucosal head and neck squamous cell carcinoma treated with definitive radiotherapy between years 2000 – 2010 ⁸. The majority (82%) had oropharyngeal cancer, and 63% were never smokers or former light smokers (less than ten pack years). The median radiation dose/ fractionation delivered was 6996 cGy in 33 fractions. Human papillomavirus (HPV) or p16 data was not available as this cohort was diagnosed before routine HPV/p16 testing. The median follow-up was 99 months, with a median surveillance imaging period of 59 months. eTable 2 summarizes the number of surveillance scans recorded during this period. At the time of analysis, 190 patients have developed a disease recurrence, and 400 patients have died. Of the recurrences, 133 were within first 24 months post-treatment.

Finite horizon models

Initially, all relapses were considered equally in the determination of transition probabilities. As the majority of relapses were within the first 24 months, first we investigate if the model will suggest frequent imaging. We generated models for finite horizon length to focus on the first three years post-treatment. From our model, the optimal frequency at which to scan patients is less than the current institutional policy. For context, the current policy is presented in Table 1. Rather than frequent imaging every three months in the first year, our model has suggested imaging at 9 to 11 months after initial negative post-treatment imaging (Table 2). Following that, the model suggested imaging at approximately 10 months interval.

To test the robustness of the results, a conservative model was also constructed using the values from the third column of eTable 1. The recommended policy from the conservative model is displayed in Table 3. Even with the more conservative parameter values, the interval between follow-up visits is greater than standard practice. For brevity, the other model settings only include the neutral parameter values.

Finite horizon policy with an infinite horizon policy

Although the risk of disease recurrence decreases over time, the risk never reaches zero. Hence we generated iterations of the model to incorporate long-term considerations. The resulting policies are summarized in Table 2. Although the finite model above recommended imaging at approximately every 9 months in the first three years, the suggested interval increases up to 18 months after three years. When the model was re-simulated with much higher or exaggerated costs and reward values, the suggested interval for imaging dropped to every 7 months for initial 3 years then every 12 months thereafter.

Locoregional relapse

As patients with distant disease are generally considered incurable (except some with small volume oligometastatic disease), we generated a model which focused on detecting locoregional relapses alone. The primary aim for this is to emphasize on detecting potentially salvageable disease early. Therefore, optimal policies were computed when only locoregional relapses were considered as relapse events in the model. Compared to considering all relapses, this resulted in fewer scans recommended by the policies and less variation in the policies when the length and start date of the policy was changed. eFigure 1 provides the fitted transition probabilities over time when all relapses were considered and when relapse events were restricted to only locoregional relapses. eFigure 2 displays how the probability estimates change based on the selection of the bandwidth parameter. Table 4 displays a summary of recommended policies that combine a policy over a finite horizon with another policy over an infinite horizon.

Finite Horizon Simulation

To examine the effectiveness of the policies given by the POMDP, a simulation was built for the 36-month surveillance program where patients imaged at 9, 18 and 26 months after

radiotherapy (Table 2). Patients' survival was used as a metric of the effectiveness of a policy. Therefore, the simulation results focus on the status of each patient at the end of the simulation. eTable 3 summarizes the states of 500 simulated patients at the end of the 36-month surveillance program.

In the simulation, the majority of patients (more than 80%) did not relapse during the surveillance period. Simulation results for the inclusive relapse case showed that less than 15% of patients who developed relapse over simulated 36-month surveillance program. Within the surveillance period, approximately 84% of patients who relapsed were eligible for salvage therapy. Less than 8% of those with disease relapse (that is, less than 1.5% of the cohort) died due to undetected disease if the POMDP suggested imaging schedule is used. In addition, the average reward function value at the end of the surveillance period was higher when using the POMDP policy. Four fewer patients died under the standard policy.

Discussion

With improvements in cancer prevention, detection and treatment in HNC, the question then is if the current surveillance imaging recommendation in the post-treatment period for those who completed radiotherapy needs reviewing and renewing. Currently, there is no clear consensus with regards to the use of surveillance imaging in those who achieved a complete response after definitive radiotherapy. While the National Comprehensive Cancer Network (NCCN) recommends regular clinical assessments at follow-up with imaging to be performed if clinically indicated or routinely annually¹³, many institutional and trial protocols required patient to have surveillance imaging at follow-up at regular intervals. Most would use a schedule of every three months initially with increasing intervals between scans in subsequent years¹⁴. There is evidence that frequent surveillance imaging may not impact patients' outcomes^{15,16}. Thompson et al.¹⁷, in a study of patients with lymphoma, suggested that surveillance imaging can exacerbate patients' anxiety. Also, a recent survey of 175 patients with HNC at one-year post-treatment follow-up in Ontario, Canada indicated that only 66% of patients felt that they needed imaging¹⁸.

There have been several studies exploring positron emission tomography (PET) imaging as a

modality to detect residual HNC after treatment and to predict for future recurrences¹⁹⁻²⁸. In a meta-analysis comprising of 51 studies, Gupta et al.²⁹ reported that the negative predictive value of a negative post-treatment PET scan is high at 95%. Furthermore, the probability of a patient developing the recurrent disease in the future after achieving a negative post-treatment PET imaging is highly unlikely with a negative predictive value of above 90%^{21,22,24-27}. Therefore, PET imaging can potentially be used as a modality to stratify surveillance imaging intensity for patients further. Shah et al.³⁰ has investigated the safety of stratifying follow-up intensity according to PET response after treatment and found that those with negative post-treatment PET imaging can be safely followed up every six months rather than every three months, with no difference in overall survival or time to detection of recurrence. As the cohort in our dataset was treated between years 2000 to 2010, before PET imaging was widely used as a primary modality for post-treatment assessment, we have not factored this into the current model. With more follow-up of patients treated in the contemporary era, we expect that our model could be improved shortly to incorporate new follow-up data with PET-CT.

Using POMDP modeling, our study has shown that less frequent use of surveillance imaging for patients with HNC treated definitively with radiotherapy may be adequate. This model provides a potential framework to establish a more cost- and time-effective surveillance program. The premise for regular surveillance follow-up is to detect clinically-occult recurrences early, allowing actionable salvage, and thereby improving subsequent patient outcomes. It has been a challenge to prove if surveillance imaging does improve patient outcomes. There is increasing evidence that the majority of patients with recurrent disease had symptoms or suspicious clinical findings at follow-up^{8,16,31}. Schwartz et al.¹⁶ reported that 86% of recurrent disease was detected clinically and no survival benefit was found in those where the disease was detected via imaging alone. It may be useful to explore a more personalized surveillance schedule to tailor to patients and disease factors. For example, patients with higher risk disease (e.g., T3/T4 oral cavity with nodal involvement) may require more regular clinical follow-up than those with good-prognosis disease (e.g., HPV-positive oropharyngeal cancer). However, this can be challenging to achieve due to the increasing number of patients and it may be more practical and feasible to have a more general surveillance policy. Besides, the patient's compliance and awareness of the need to

seek medical attention if he/she developed any new symptom of concern should be taken into account. Kissun et al.³¹ found that one in 5 patients were non-compliant with follow-up, and despite the majority of patients were symptomatic from recurrent disease, less than half sought medical help in a timely manner.

This study has several limitations. Firstly, the patients were not stratified into groups sharing similar characteristics. Hence, the transition probabilities estimated for the model reflected aggregated behavior over the entire population, as opposed to more specific behavior. Secondly, the number of states in the model was limited to improve tractability; however, the model assumed that a single "healthy" indicator in the patient's state was sufficient. In reality, the transition from a patient in good health to suffering from relapse is a gradual process. Thirdly, the reward function was estimated based on data in the literature in an attempt to reflect the preferences of the parties involved. The assumption that the reward function accurately summarizes the relevant preferences must be considered when interpreting the results. The reward function does not incorporate the scarcity of resources used to conduct imaging. Lastly, our cohort was treated in the 'pre-immunotherapy' era. With the current emergence of more effective systemic therapy, such as immunotherapy and targeted therapy, it may be possible that early detection of recurrent or distant disease although asymptomatic may be treated effectively with new drugs. The recent KEYNOTE-040 and 048 data indicated that pembrolizumab treated patients with recurrent or metastatic HNC had better overall survival compared to the standard arm^{32,33}. With increasing follow up and clinical data from a contemporary cohort, our model can be updated to reflect current practice changes and needs.

This model currently relies on a single institutional database⁸. Although the data used in this model was based on a single institution dataset, it included follow-up data on more than 1500 patients with a median follow-up of 99 months – allowing the model to make assumptions as 'real-life' as possible. As with all model-based applications, further multi-institutional and/or international data with a large volume of patient follow-up is required to determine whether the extant model, which has institutional relevance, has generalizability across larger (i.e. pan-institutional) cohorts. Alternatively, the same approach can be readily integrated to derive institution-specific policies and schedules at a

local level, given a sufficient number of priors. Finally, the POMDP model derived is well suited to machine learning/reinforcement learning applications, wherein the policies and schedules are dynamically updated as increasing and newer patient data is added over time.

Ultimately, the aim is to optimize cost- and time-effectiveness of follow-up for an increasing number of survivors of HNC, while maximizing the probability of detecting salvageable disease early to optimize salvage outcomes. Nevertheless, our model suggests that less frequent surveillance imaging can provide adequate information on relapse status for patients with HNC treated definitively with radiotherapy. This model could potentially translate to a more cost- and time-effective surveillance program for this group of patients.

References

1. Cancer Facts and Figures 2017. American Cancer Society Web site, American Cancer Society, 2017
2. Mariotto AB, Yabroff KR, Shao Y, et al: Projections of the cost of cancer care in the United States: 2010-2020. *J Natl Cancer Inst* 103:117-28, 2011
3. van Agthoven M, van Ineveld BM, de Boer MF, et al: The costs of head and neck oncology: primary tumours, recurrent tumours and long-term follow-up. *Eur J Cancer* 37:2204-11, 2001
4. Hauskrecht M, Fraser H: Modeling treatment of ischemic heart disease with partially observable Markov decision processes. *Proc AMIA Symp*:538-42, 1998
5. Hauskrecht M, Fraser H: Planning treatment of ischemic heart disease with partially observable Markov decision processes. *Artif Intell Med* 18:221-44, 2000
6. Puterman ML: *Markov decision processes: discrete stochastic dynamic programming*, Wiley-Interscience, 2005
7. Lovejoy WS: A survey of algorithmic methods for partially observed Markov decision processes. *Ann Oper Res* 28:47-65, 1991
8. Ng SP, Pollard C, 3rd, Berends J, et al: Usefulness of surveillance imaging in patients with head and neck cancer who are treated with definitive radiotherapy. *Cancer*, 2019

9. Cassandra A, 2015
10. Huntington SF, Svoboda J, Doshi JA: Cost-effectiveness analysis of routine surveillance imaging of patients with diffuse large B-cell lymphoma in first remission. *J Clin Oncol* 33:1467-74, 2015
11. Brown B, Diamantopoulos A, Bernier J, et al: An economic evaluation of cetuximab combined with radiotherapy for patients with locally advanced head and neck cancer in Belgium, France, Italy, Switzerland, and the United Kingdom. *Value Health* 11:791-9, 2008
12. Smith AF, Hall PS, Hulme CT, et al: Cost-effectiveness analysis of PET-CT-guided management for locally advanced head and neck cancer. *Eur J Cancer* 85:6-14, 2017
13. Head and Neck Cancers (Version 2.2017). <http://www.nccn.org>, National Comprehensive Cancer Network
14. Ang KK, Harris J, Wheeler R, et al: Human papillomavirus and survival of patients with oropharyngeal cancer. *N Engl J Med* 363:24-35, 2010
15. Zhao X, Rao S: Surveillance imaging following treatment of head and neck cancer. *Semin Oncol* 44:323-329, 2017
16. Schwartz DL, Barker J, Jr., Chansky K, et al: Postradiotherapy surveillance practice for head and neck squamous cell carcinoma--too much for too little? *Head Neck* 25:990-9, 2003
17. Thompson CA, Charlson ME, Schenkein E, et al: Surveillance CT scans are a source of anxiety and fear of recurrence in long-term lymphoma survivors. *Ann Oncol* 21:2262-6, 2010
18. Brennan KE, Hall SF, Yoo J, et al: Routine follow-up care after curative treatment of head and neck cancer: A survey of patients' needs and preferences for healthcare services. *Eur J Cancer Care (Engl)*:e12993, 2019
19. Ho AS, Tsao GJ, Chen FW, et al: Impact of positron emission tomography/computed tomography surveillance at 12 and 24 months for detecting head and neck cancer recurrence. *Cancer* 119:1349-56, 2013
20. Gupta T, Jain S, Agarwal JP, et al: Diagnostic performance of response assessment FDG-PET/CT in patients with head and neck squamous cell carcinoma treated with high-precision definitive (chemo)radiation. *Radiother Oncol* 97:194-9, 2010
21. Koshkareva Y, Branstetter B, Gaughan JP, et al: Predictive accuracy of first

post-treatment PET/CT in HPV-related oropharyngeal squamous cell carcinoma. *Laryngoscope* 124:1843-7, 2014

22. Ng SP, Johnson JM, Gunn GB, et al: Significance of Negative Posttreatment 18-FDG PET/CT Imaging in Patients With p16/HPV-Positive Oropharyngeal Cancer. *Int J Radiat Oncol Biol Phys* 102:1029-1035, 2018

23. Sherriff JM, Ogunremi B, Colley S, et al: The role of positron emission tomography/CT imaging in head and neck cancer patients after radical chemoradiotherapy. *Br J Radiol* 85:e1120-6, 2012

24. McDermott M, Hughes M, Rath T, et al: Negative predictive value of surveillance PET/CT in head and neck squamous cell cancer. *AJNR Am J Neuroradiol* 34:1632-6, 2013

25. Yao M, Smith RB, Hoffman HT, et al: Clinical significance of postradiotherapy [18F]-fluorodeoxyglucose positron emission tomography imaging in management of head-and-neck cancer-a long-term outcome report. *Int J Radiat Oncol Biol Phys* 74:9-14, 2009

26. Vainshtein JM, Spector ME, Stenmark MH, et al: Reliability of post-chemoradiotherapy F-18-FDG PET/CT for prediction of locoregional failure in human papillomavirus-associated oropharyngeal cancer. *Oral Oncol* 50:234-9, 2014

27. Nishimura G, Yabuki K, Hata M, et al: Imaging strategy for response evaluation to chemoradiotherapy of the nodal disease in patients with head and neck squamous cell carcinoma. *Int J Clin Oncol* 21:658-667, 2016

28. Liu HY, Milne R, Lock G, et al: Utility of a repeat PET/CT scan in HPV-associated Oropharyngeal Cancer following incomplete nodal response from (chemo)radiotherapy. *Oral Oncol* 88:153-159, 2019

29. Gupta T, Master Z, Kannan S, et al: Diagnostic performance of post-treatment FDG PET or FDG PET/CT imaging in head and neck cancer: a systematic review and meta-analysis. *Eur J Nucl Med Mol Imaging* 38:2083-95, 2011

30. Shah K, Te Marvelde L, Collins M, et al: Safety and cost analysis of an (18)FDG-PET-CT response based follow-up strategy for head and neck cancers treated with primary radiation or chemoradiation. *Oral Oncol* 51:529-35, 2015

31. Kissun D, Magennis P, Lowe D, et al: Timing and presentation of recurrent oral and oropharyngeal squamous cell carcinoma and awareness in the outpatient clinic. *Br J Oral Maxillofac Surg* 44:371-6, 2006

32. Cohen EEW, Soulieres D, Le Tourneau C, et al: Pembrolizumab versus methotrexate, docetaxel, or cetuximab for recurrent or metastatic head-and-neck squamous cell carcinoma (KEYNOTE-040): a randomised, open-label, phase 3 study. *Lancet* 393:156-167, 2019

33. Burtness B, Bratland Å, Fuereder T, et al: LBA8_PRKEYNOTE-048: Phase III study of first-line pembrolizumab (P) for recurrent/metastatic head and neck squamous cell carcinoma (R/M HNSCC). *Annals of Oncology* 29, 2018

Figure

Figure 1. A directed graphical model of states and transitions in the partially observed Markov decision process (POMDP).

Some states are a combination of health status (e.g. “Healthy”) and a number indicating the time passed since the last surveillance imaging. The arrows indicate that the transition probability between the two states is non-zero in general.

Tables

	Year 1	Year 2	Year 3
Scan Interval	3 months	6 months	12 months

Table 1. The current common schedule of surveillance scans for head and neck cancer patients.

Finite Horizon Length (months)	When to scan – finite policy (months after first post-treatment scan)			Infinite Scan Interval
	Scan 1	Scan 2	Scan 3	
12	10	NA	NA	10
24	11	22	NA	18

36	9	18	26	18
----	---	----	----	----

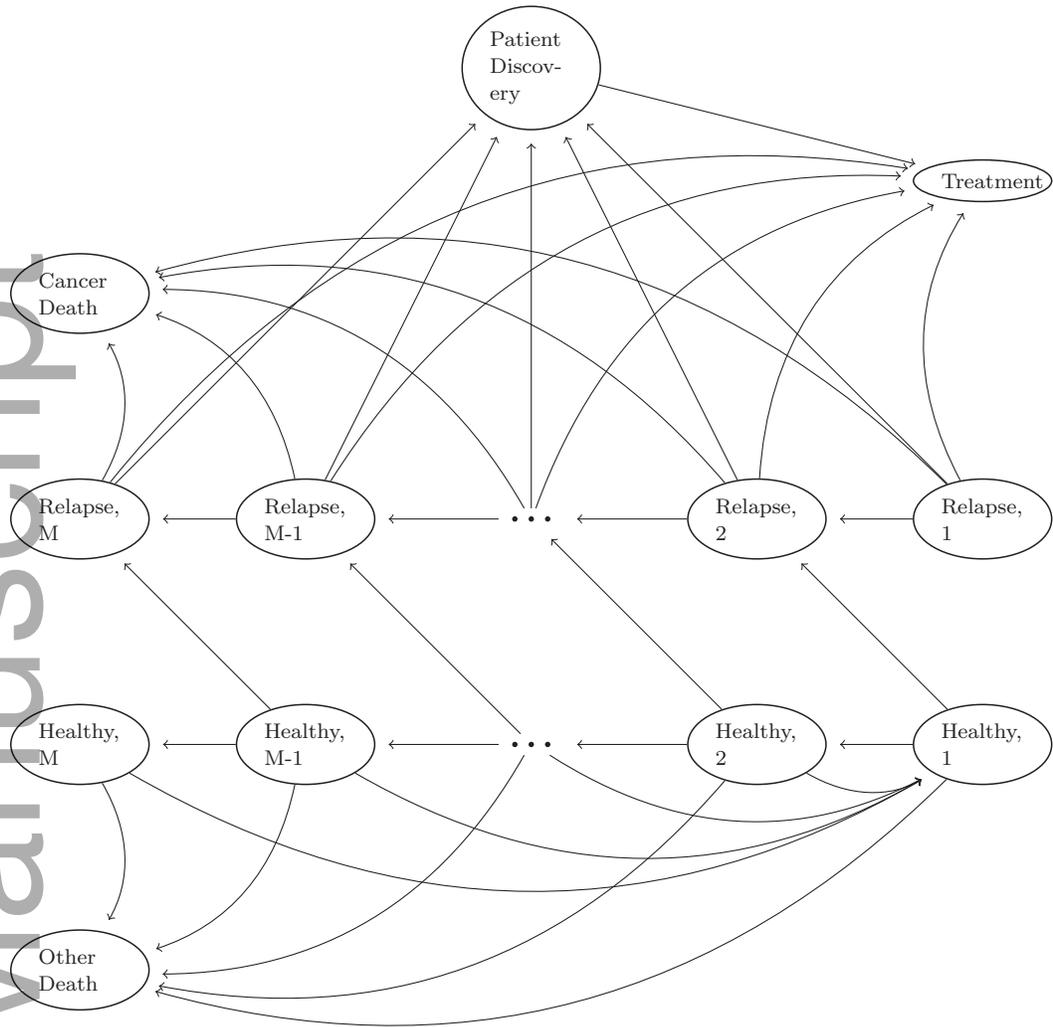
Table 2. The recommended schedule for surveillance scans for various finite horizon models and an infinite horizon policy.

Finite Horizon Length (months)	When to Scan (months after first post-treatment scan)			
	Scan 1	Scan 2	Scan 3	Scan 4
12	10	NA	NA	NA
24	7	14	22	NA
36	7	14	21	28

Table 3. The recommended schedule for surveillance scans for various finite horizon models when all relapses are considered as relapse events modeled with conservative parameters.

Finite Horizon Length (Months)	Finite Scan 1	Finite Scan 2	Finite Scan 3	Infinite Scan Interval
12	10	NA	NA	18
24	19	NA	NA	18
36	16	34	NA	18

Table 4. Recommended policies that combine a finite horizon policy with an infinite horizon policy when only locoregional relapses are considered as relapse events.





Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Ng, SP;Ajayi, T;Schaefer, AJ;Pollard, C;Bahig, H;Garden, AS;Rosenthal, DI;Gunn, GB;Frank, SJ;Phan, J;Morrison, WH;Johnson, JM;Mohamed, ASR;Sturgis, EM;Fuller, CD

Title:

Surveillance imaging for patients with head and neck cancer treated with definitive radiotherapy: A partially observed Markov decision process model

Date:

2019-11-14

Citation:

Ng, S. P., Ajayi, T., Schaefer, A. J., Pollard, C., Bahig, H., Garden, A. S., Rosenthal, D. I., Gunn, G. B., Frank, S. J., Phan, J., Morrison, W. H., Johnson, J. M., Mohamed, A. S. R., Sturgis, E. M. & Fuller, C. D. (2019). Surveillance imaging for patients with head and neck cancer treated with definitive radiotherapy: A partially observed Markov decision process model. *CANCER*, 126 (4), pp.749-756. <https://doi.org/10.1002/cncr.32597>.

Persistent Link:

<http://hdl.handle.net/11343/286618>