

COMMENT

Open Access



From algorithms to action: improving patient care requires causality

Wouter A. C. van Amsterdam^{1,2*}, Pim A. de Jong², Joost J. C. Verhoeff³, Tim Leiner^{2,4} and Rajesh Ranganath⁵

Abstract

In cancer research there is much interest in building and validating outcome prediction models to support treatment decisions. However, because most outcome prediction models are developed and validated without regard to the causal aspects of treatment decision making, many published outcome prediction models may cause harm when used for decision making, despite being found accurate in validation studies. Guidelines on prediction model validation and the checklist for risk model endorsement by the American Joint Committee on Cancer do not protect against prediction models that are accurate during development and validation but harmful when used for decision making. We explain why this is the case and how to build and validate models that are useful for decision making.

Keywords Causal inference, Oncology, Tailored treatment decision making, Prediction research, Prognosis research

Introduction

Treatment decisions in cancer care are guided by treatment effect estimates from randomized controlled trials (RCTs). RCTs estimate the *average* effect of one treatment versus another in a certain population. However, treatments may not be equally effective for every patient in a population. Knowing the effectiveness of treatments tailored to specific patient and tumor characteristics would enable individualized treatment decisions. Getting tailored treatment effects by averaging outcomes in different patient subgroups in RCTs requires an infeasible number of patients to have sufficient statistical power in all relevant subgroups for all possible treatments. Instead,

we must rely on statistical modeling, potentially using observational data from non-randomized studies to further the individualization of treatment decisions.

The American Joint Committee on Cancer (AJCC) recommends that researchers develop outcome prediction models in an effort to individualize treatment decisions [1, 2]. Outcome prediction models, sometimes called risk models or prognosis models, use patient and tumor characteristics to predict a patient outcome such as cancer recurrence or overall survival. The assumption is that the predictions are useful for treatment decisions using rules such as “prescribe chemotherapy only if the outcome prediction model predicts the patient has a high risk of recurrence”. Many outcome prediction models are published every year. Recognizing the importance of reliable predictions, the AJCC published a checklist for outcome prediction models to ensure dependable prediction accuracy in the patient population for which the outcome prediction model was designed [1]. However, accurate outcome predictions do not imply that these predictions yield good treatment decisions. In this comment, we show that outcome prediction models rely on a fixed treatment policy which implies that outcome prediction models that were found to accurately predict outcomes in validation studies can still lead to patient harm when

*Correspondence:

Wouter A. C. van Amsterdam
w.a.c.vanamsterdam-3@umcutrecht.nl

¹ Julius Center for Health Sciences and Primary Care, University Medical Center Utrecht, Utrecht, the Netherlands

² Department of Radiology, University Medical Center Utrecht, Utrecht, the Netherlands

³ Department of Radiotherapy, University Medical Center Utrecht, Utrecht, the Netherlands

⁴ Mayo Clinic, Rochester, MN, USA

⁵ Courant Institute of Mathematical Sciences, Department of Computer Science, New York University, New York City, NY, USA



© The Author(s) 2024. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

used to inform treatment decisions. We then give guidance on how to evaluate whether a model has value for decision-making and how to develop models that are useful for individualized treatment decisions.

Accurate predictions have unknown value for decision-making

Individualizing treatment decisions means changing the *treatment policy*. For example, if for a specific cancer type and stage the current treatment policy is to give the same treatment to all patients, then individualizing treatment decisions means recommending treatments tailored to a patient’s characteristics. The value of an outcome prediction model is not in how well it predicts under a certain historic treatment policy, but rather what is the effect of deploying this model on treatment decisions and patient outcomes?

Consider an outcome prediction model that uses pre-treatment tumor characteristics to predict an outcome but ignores whatever treatment the patients may have had, i.e. *treatment-naïve* models (e.g. [3–5]). Interestingly, the decision to ignore treatments in the outcome

prediction model is in line with the AJCC checklist for outcome prediction models (item 12 [1]). However, these outcome prediction models can cause more harm than good when used to support treatment decisions, even when they are accurate under the historic treatment policy. Consider for example an outcome prediction model that predicts overall survival for stage IV lung cancer patients based on the pre-treatment growth-rate of the tumor. An accurate model would predict shorter survival for patients with faster growing tumors. Applying this outcome prediction model, a clinician could decide to refrain from palliative radiotherapy in patients with faster growing tumors under the assumption that their life expectancy is too short to benefit from radiotherapy. This decision based on the outcome prediction model would be unjustified and harmful, as faster growing tumors are more susceptible to radiotherapy [6]. See Fig. 1 for an illustration of introducing an outcome prediction model for treatment decisions.

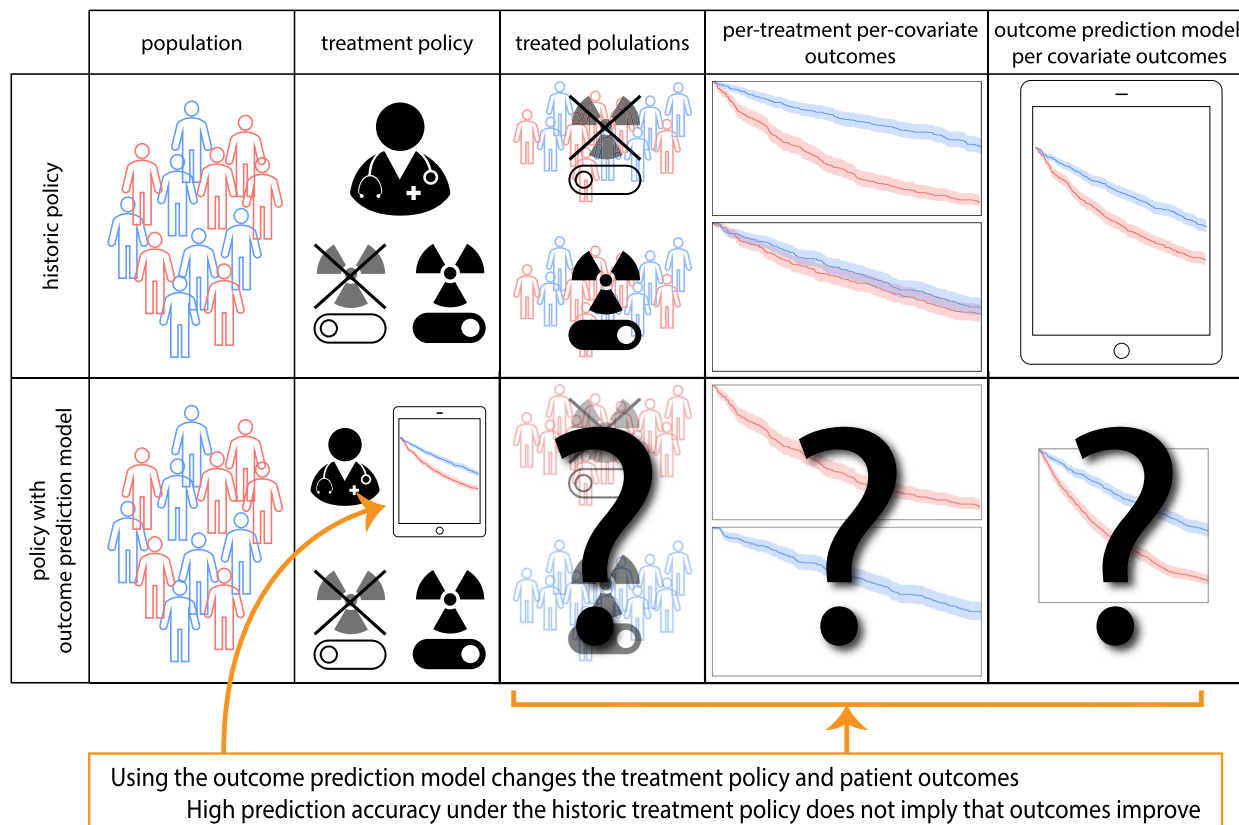


Fig. 1 Illustration of the use of outcome prediction models that ignore treatment allocations in the historical data (i.e. are *treatment naïve*) for treatment decision making. These models change the treatment decisions and thus patient outcomes but whether this change improves patient outcomes is not determined by the prediction accuracy of the outcome prediction model

Prospective validation does not test value for decision-making

The gold standard for evaluating the accuracy of an outcome prediction model is *prospective validation* [1, 7]. In a prospective validation, patient characteristics and outcomes are recorded for a new patient cohort according to a predefined protocol. Comparing the outcome prediction model's predictions with the observed outcomes provides an estimate of how accurate the outcome prediction model is outside the cohort in which the model was developed. The outcome prediction model from the lung cancer example above, if well-estimated, would be found accurate in a prospective validation that uses the historic treatment policy because the outcome prediction model was developed under the same historic policy. It would then fulfill all the AJCC checklist items but still lead to patient harm when used for treatment decisions because the differential effect of radiotherapy depending on tumor growth-rate is not accounted for in the outcome prediction model.

As an additional validation step, one may conduct a prospective validation study where the outcome prediction model is used for treatment decisions in new patients, thus changing the treatment policy. If such a validation were carried out for the lung cancer survival outcome prediction model, the patients with fast-growing tumors would be given radiotherapy less often due to the predictions of the outcome prediction model, leading to even worse survival for these patients than before

introduction of the outcome prediction model. Introducing the outcome prediction model for decision making caused harm because under the new policy treatments are withheld from those who would have benefited most (the patients with fast-growing tumors). However, in this validation study with model deployment the prediction model is still accurate as the model already predicted that patients with fast-growing tumors have a poor prognosis.

Models should improve decisions

The crux of the issue with outcome prediction models is that they answer the question "What is the chance of the outcome given these patient and tumor characteristics, *with the assumption that we will keep making the same treatment decisions as we always did?*". Similar issues exist with other kinds of outcome prediction models which make predictions using the historical treatments but without regards to the policy for how those treatments were assigned (i.e. *post-decision* models such as [8–12]). Post-decision outcome prediction models are also in line with the AJCC checklist (item 12 [1]). To improve treatment decisions however, we need models with a foreseeable positive effect on outcomes when used in decision-making.

Outcome prediction models assume treatment decisions follow the historical policy and thereby cannot inform us on the effect of a new policy derived from the outcome prediction model. This reliance on the historical treatment policy leads to a fundamental gap between

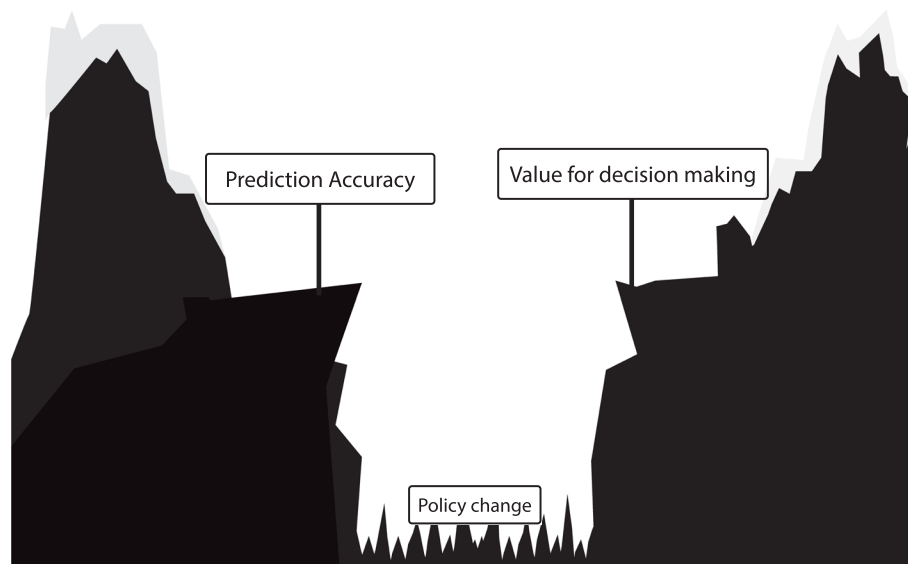


Fig. 2 Illustration of the difference between outcome prediction model accuracy and its value for treatment decision making. Validation of an outcome prediction model following the AJCC checklist leads to a reliable estimate of the outcome prediction model's accuracy if the treatment policy does not change. However, because the outcome prediction model relies on a fixed historic treatment policy, prediction accuracy does not imply value for decision making, as visualized with the gap. This gap can only be bridged with *causality*

a prediction model's accuracy and its value for treatment decision-making in clinical practice (Fig. 2). Bridging the gap from prediction accuracy to value for decision making is only possible with *causality*. Evaluating the effect of a prediction model-based treatment policy on patient outcomes requires a causal study design or causal assumptions.

How to validate models used for treatment decisions?

The ultimate test of the effect of introducing a new treatment policy for example based on an outcome prediction model is a *cluster randomized controlled trial* [7, 13]. In a cluster RCT with outcome prediction models, some groups of clinicians are randomly selected to get access to the model while others are not. This allows for the estimation of the effect of introducing the model on treatment decisions and patient outcomes. For example, the cluster RCT could demonstrate that using the model leads to fewer treatment side effects and better overall survival. However, in the context of shared decision-making, patients may weigh the value of overall survival versus treatment discomfort differently [14]. These individual preferences need to be taken into account in the cluster RCT when calculating the value of introducing a model for decision-making.

As an alternative to cluster RCTs, the expected outcomes under a treatment policy (e.g. based on a prediction model) can be evaluated in data from a standard RCT. This can be done by calculating the average outcome in the subgroup of patients for whom the randomized treatment assignment was concordant with the policy [15]. Multiple policies can be compared this way, for example comparing a policy based on a new prediction model with current clinical practice. The policy with the best outcomes is preferable. However, such an analysis does not take into account that in practice the compliance with the new treatment policy might not be perfect. Notably, the validation steps recommended in the AJCC

checklist [1] provide no information on what the effect is of deploying an outcome prediction model on treatment decisions and patient outcomes.

Building models to individualize treatment decisions

Cluster RCTs are costly and time consuming. With tools from causal inference we can improve the chance of success of models for decision making. One way to construct a good individualized treatment policy is with models that predict the outcome under hypothetical interventions, where the intervention is the decision to give a certain treatment. The optimal treatment policy selects the treatment that leads to the most beneficial expected outcome.

Estimating such models for *prediction under intervention* requires unconfoundedness, which holds when there are no unknown variables that influence both the treatment assignment and the outcome (i.e. confounders). RCTs are ideal for this as unconfoundedness holds by design because the treatment assignment is random. However, individual RCTs are generally too small to include many important patient and tumor characteristics in the modeling. Observational data from regular clinical practice on the other hand are often more readily available. If all variables that influence the treatment policy are available in a particular dataset, meaning that unconfoundedness holds, there are many approaches to prediction under intervention. These include 'conventional' statistical approaches such as regression, or machine learning approaches, for example using neural networks [16].

To express available background knowledge and judge whether unconfoundedness holds in observational data, researchers can use directed acyclic graphs (DAGs) [17]. DAGs depict variables (such as treatment, outcome and confounders) and causal dependencies between the variables as *arrows* that point from a cause variable to an effect variable, see Fig. 3 for an example. Using tools from *causal inference*, the DAG determines whether a

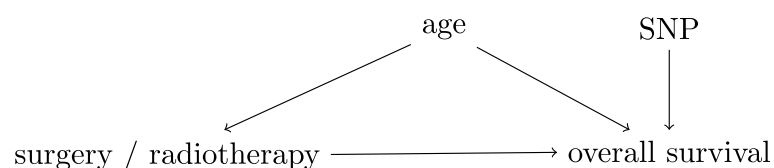


Fig. 3 Simplified Directed Acyclic Graph for the decision between surgery and radiotherapy for overall survival in lung cancer patients. As an example, consider a hypothetical study in early-stage lung cancer where researchers investigate whether the relative effectiveness of surgery versus radiotherapy for overall survival depends on a certain single-nucleotide polymorphism (SNP). The SNP assay was performed for the study only so this information did not affect the treatment decision. A DAG with four variables for this study is presented in this Figure. In this DAG, the variables *age* and *SNP* both have arrows to overall survival, but only *age* influences the treatment decision as older patients are less likely to get surgery. This DAG indicates that unconfoundedness holds when *age* is conditioned on in the analysis, as *age* is the only confounder between the treatment and the outcome [17]

prediction under intervention model can be estimated and if so what confounders need to be accounted for [17].

In some cases not all confounders are available. In this setting with unobserved confounding standard methods based on confounder adjustment cannot be used, but sometimes prediction under intervention models may be estimated using specialized methods. Two examples are methods based on proxy-variables of unmeasured confounders [18, 19] and instrumental variable methods [20] and their machine learning variants [21, 22]. These methods rely on assumptions that may not hold perfectly in reality, so figuratively speaking they might reduce the gap between model accuracy and treatment policy value, but not close the gap entirely. DAGs encode *assumptions* about the data which may not hold perfectly in practice. The effects of potential violations of these assumptions may be estimated using *sensitivity analyses* [23].

A special case for prediction under intervention is the untreated risk, which is the hypothetical outcome under no treatment (or some baseline treatment) and would be observed in the control group of an RCT. For instance, when deciding to give adjuvant therapy after breast cancer surgery, the untreated risk of recurrence is the risk of recurrence when no adjuvant therapy would be given [24]. Knowing the untreated risk is valuable when considering giving no further treatment, and as a baseline to compare other potential treatments against. Although estimating the untreated risk requires unconfoundedness, in some cases it may be estimated quite accurately even from confounded data using specialized methods [25].

Because RCTs randomly assign patients to interventions, models for prediction under intervention can be validated in RCTs with standard prediction validation approaches [7]. For shared decision-making, prediction under intervention of different treatment options allows the patient to make their own judgment on how to weigh e.g. expected overall survival with expected treatment discomfort. Whereas individual RCTs randomize a patient to a certain treatment, cluster RCTs randomize clinicians' access to a model for decision support. Thereby individual treatment decisions may still be confounded in cluster RCTs meaning that cluster RCTs cannot validate predictions from *prediction under intervention*-models directly. Both policy evaluation with cluster RCTs and prediction-under-intervention validation in standard RCTs are also possible in observational data but require unconfoundedness and thereby sensitivity analyses for potentially omitted confounders [23, 26].

Discussion

In line with American Joint Committee on Cancer recommendations [1, 2] many researchers develop outcome prediction models to individualize treatment decisions.

The AJCC checklist provides important guidelines for outcome prediction model development and validation, such as clearly defining the patient population, predictor variables and prediction time-point, in addition to validation in external datasets. These items improve the dependability of outcome prediction models for predicting outcomes in the intended patient population if there are no changes in the treatment policy [1]. However, not changing the treatment policy directly contradicts the intended purpose of these models. Outcome prediction models that satisfy all the criteria in the checklist still have unknown clinical utility because high prediction accuracy in prospective validation studies does not imply value for treatment decision-making in clinical practice [27]. Because the gap between outcome prediction model accuracy and value for decision-making is due to causal issues, it is not resolved by larger datasets, more flexible prediction algorithms (e.g. machine learning) or even by prospective validation with model deployment. In contrast, we explained how models for prediction under intervention are useful for decision-making and how to validate any model used in decision-making.

The gap between outcome prediction model accuracy and value for decision-making is due to causal issues, but it is different from the standard "correlation does not imply causation". In the standard "correlation is not causation" setting, all variables (treatment, outcome, patient/tumor characteristics) are already present in the historical data, whereas in this case, the output of the outcome prediction model cannot be a cause of the outcome. This is because the outcome prediction model is not a variable in historical data, but a shift in policy that changes the distribution of the treatment.

It was noted before that cluster RCTs are the ultimate test for the impact of a new prediction model on clinical practice due to issues related to compliance with treatment recommendations [13]. We show that because of the gap between prediction accuracy and value for treatment decision-making, many accurate outcome prediction models will fail to demonstrate value in cluster RCTs. Also, cluster RCTs measure the effect of a new policy on average outcomes but do not directly measure whether a model accurately predicts the outcome under intervening to give a certain treatment, for this individually randomized data are most valuable. For shared decision-making accurate predictions-under-intervention may be most important. Two patients with the same predicted outcomes may make a different treatment decision because each patient has their own values and preferences. In a cluster RCT, these individual values need to be accounted for when evaluating a treatment policy, for example by eliciting the values and incorporating them in the analysis when weighing for example overall survival and treatment discomfort.

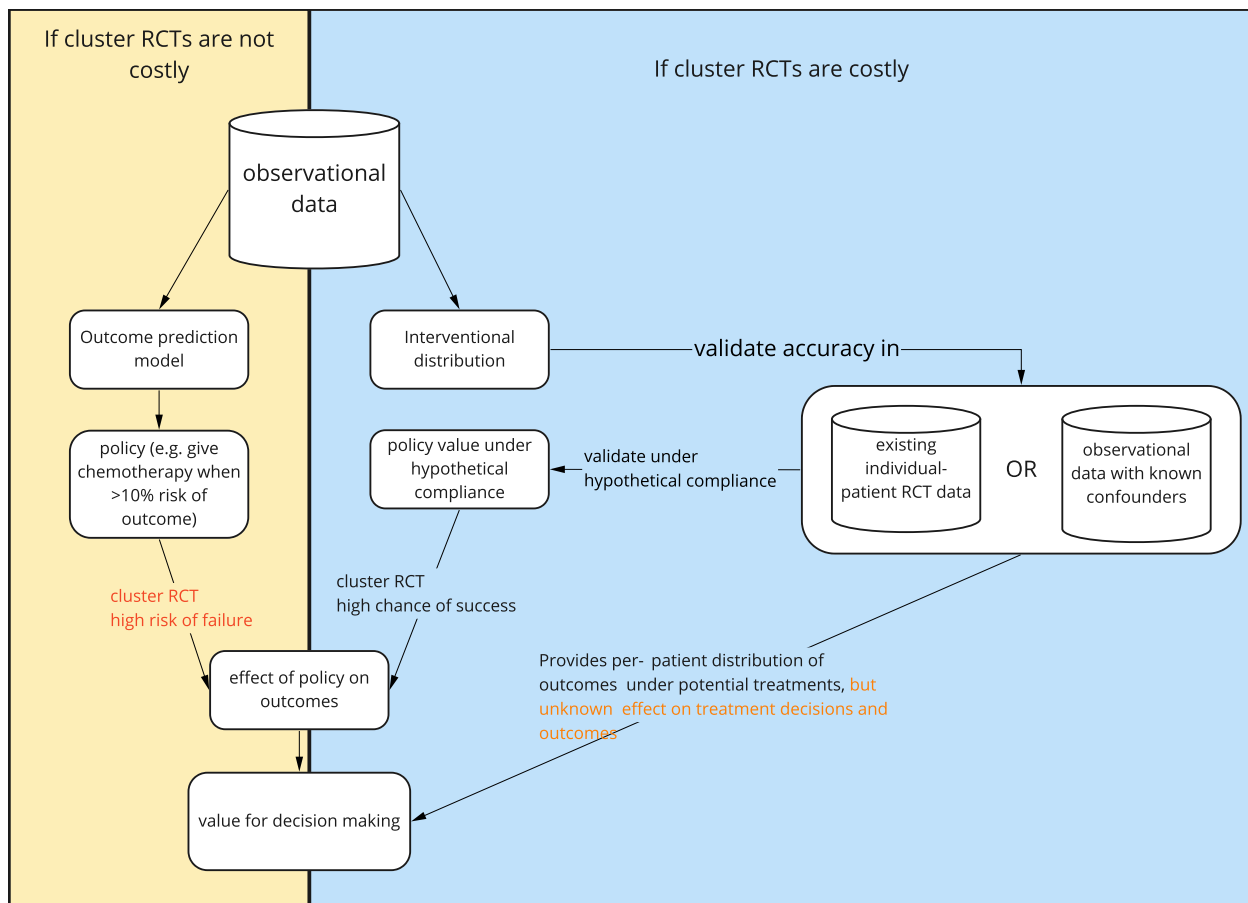


Fig. 4 Flowchart of what to do depending on the costliness of cluster randomized controlled trials. Costliness of cluster RCTs should be taken broadly, including time, money and ethical considerations

Previous work underlined the value of *prediction-under-intervention* models (sometimes referred to as *counterfactual prediction*) for supporting treatment decisions [28, 29]. Our comment highlights the potential harm of current common practice where outcome prediction models are deployed for decision making based on prediction accuracy alone, further emphasizing the relevance of prediction-under-intervention. In addition, we note how models may be validated for decision support for example with cluster RCTs.

Building models for prediction under intervention is harder than developing outcome prediction models due to the extra requirement of unconfoundedness, which involves formalizing assumptions about confounders for example with DAGs, gathering data on all confounders, often more complex statistical estimation, and sensitivity analyses. When the cost to do a cluster RCT is low, it may suffice to build outcome prediction models in line with the AJCC checklist and test them in cluster RCTs

before model deployment. As illustrated in Fig. 4, when cluster RCTs are costly, impractical or unethical, models that predict under interventions are preferable as they have foreseeable effects when used for treatment decision-making.

There is a classical distinction between treatment effect estimation and prediction that amounts to “treatment effect estimation is causal (and thus requires RCTs)” but “prediction is not causal”. When it comes to individualizing treatment decisions with prediction models, this distinction is unhelpful and confusing as the goal is to predict what would happen under different interventions. Selecting the best treatment for a patient is a causal question and requires causal answers.

Acknowledgements

Drs. Lidia Barberio, Director of “Longkanker Nederland” (the Dutch patient association for lung cancer) provided feedback on this comment. Her input broadened the scope of this work making it more relevant for patients. Specifically, we added more emphasis on the importance of including the values of the patient in treatment decision-making.

Authors' contributions

Conceptualization: WA, PJ, JV, TL, RR. Writing of draft manuscript: WA and RR. Editing and revision: all authors.

Funding

There was no specific funding for this comment.

Availability of data and materials

Not applicable.

Declarations**Ethics approval and consent to participate**

Not applicable.

Consent for publication

All authors consent with publication.

Competing interests

The authors declare no competing interests.

Received: 31 January 2023 Accepted: 15 April 2024

Published online: 26 April 2024

References

- Kattan MW, Hess KR, Amin MB, Lu Y, Moons KGM, Gershonwald JE, et al. American Joint Committee on Cancer acceptance criteria for inclusion of risk models for individualized prognosis in the practice of precision medicine. *CA Cancer J Clin*. 2016;66(5):370–4. <https://doi.org/10.3322/caac.21339>.
- Amin MB, Greene FL, Edge SB, Compton CC, Gershonwald JE, Brookland RK, et al. The Eighth Edition AJCC Cancer Staging Manual: Continuing to build a bridge from a population-based to a more “personalized” approach to cancer staging: The Eighth Edition AJCC Cancer Staging Manual. *CA Cancer J Clin*. 2017;67(2):93–99. <https://doi.org/10.3322/caac.21388>.
- Salazar R, Roepman P, Capella G, Moreno V, Simon I, Dreezen C, et al. Gene Expression Signature to Improve Prognosis Prediction of Stage II and III Colorectal Cancer. *J Clin Oncol*. 2011;29(1):17–24. <https://doi.org/10.1200/JCO.2010.30.1077>.
- Merli F, Luminari S, Tucci A, Arcari A, Rigacci L, Hawkes E, et al. Simplified Geriatric Assessment in Older Patients With Diffuse Large B-Cell Lymphoma: The Prospective Elderly Project of the Fondazione Italiana Linfomi. *J Clin Oncol Off J Am Soc Clin Oncol*. 2021;39(11):1214–22. <https://doi.org/10.1200/JCO.20.02465>.
- Courtiol P, Maussion C, Moarii M, Pronier E, Pilcer S, Sefta M, et al. Deep learning-based classification of mesothelioma improves prediction of patient outcome. *Nat Med*. 2019;25(10):1519–25. <https://doi.org/10.1038/s41591-019-0583-3>.
- Breuer K. Growth rate and radiosensitivity of human tumours—II: Radio-sensitivity of human tumours. *Eur J Cancer* (1965). 1966;2(2):173–188. [https://doi.org/10.1016/0014-2964\(66\)90009-0](https://doi.org/10.1016/0014-2964(66)90009-0).
- Moons KGM, Altman DG, Reitsma JB, Ioannidis JPA, Macaskill P, Steyerberg EW, et al. Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis (TRIPOD): Explanation and Elaboration. *Ann Intern Med*. 2015;162(1):W1. <https://doi.org/10.7326/M14-0698>.
- Ryu JS, Ryu HJ, Lee SN, Memon A, Lee SK, Nam HS, et al. Prognostic impact of minimal pleural effusion in non-small-cell lung cancer. *J Clin Oncol Off J Am Soc Clin Oncol*. 2014;32(9):960–7. <https://doi.org/10.1200/JCO.2013.50.5453>.
- Fried DV, Mawlawi O, Zhang L, Fave X, Zhou S, Ibbott G, et al. Stage III Non-Small Cell Lung Cancer: Prognostic Value of FDG PET Quantitative Imaging Features Combined with Clinical Prognostic Factors. *Radiology*. 2016;278(1):214–22. <https://doi.org/10.1148/radiol.2015142920>.
- Hippisley-Cox J, Coupland C. Development and validation of risk prediction equations to estimate survival in patients with colorectal cancer: cohort study. *BMJ (Clin Res Ed)*. 2017;357:j2497. <https://doi.org/10.1136/bmj.j2497>.
- Liu R, Rizzo S, Waliyany S, Garmhausen MR, Pal N, Huang Z, et al. Systematic pan-cancer analysis of mutation-treatment interactions using large real-world clinicogenomics data. *Nat Med*. 2022. <https://doi.org/10.1038/s41591-022-01873-5>.
- Pires da Silva I, Ahmed T, McQuade JL, Nebhan CA, Park JJ, Versluis JM, et al. Clinical Models to Define Response and Survival With Anti-PD-1 Antibodies Alone or Combined With Ipilimumab in Metastatic Melanoma. *J Clin Oncol Off J Am Soc Clin Oncol*. 2022;40(10):1068–1080. <https://doi.org/10.1200/JCO.21.01701>.
- Moons KGM, Altman DG, Vergouwe Y, Royston P. Prognosis and prognostic research: application and impact of prognostic models in clinical practice. *BMJ*. 2009;338:b606. <https://doi.org/10.1136/bmj.b606>.
- Barry MJ, Edgman-Levitan S. Shared Decision Making – The Pinnacle of Patient-Centered Care. *N Engl J Med*. 2012;366(9):780–1. <https://doi.org/10.1056/NEJMp1109283>.
- Karmali KN, Lloyd-Jones DM, van der Leeuw J, Goff Jr DC, Yusuf S, Zanchetti A, et al. Blood pressure-lowering treatment strategies based on cardiovascular risk versus blood pressure: A meta-analysis of individual participant data. *PLoS Med*. 2018;15(3):e1002538. <https://doi.org/10.1371/journal.pmed.1002538>.
- Shalit U, Johansson FD, Sontag D. Estimating individual treatment effect: generalization bounds and algorithms. 2017. [arXiv:1606.03976](https://arxiv.org/abs/1606.03976).
- Pearl J. *Causality*. Cambridge University Press; 2009.
- Miao W, Geng Z, Tchetgen Tchetgen EJ. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*. 2018;105(4):987–93. <https://doi.org/10.1093/biomet/asy038>.
- van Amsterdam WAC, Verhoeff JJC, Harlianto NI, Bartholomeus GA, Puli AM, de Jong PA, et al. Individual treatment effect estimation in the presence of unobserved confounding using proxies: a cohort study in stage III non-small cell lung cancer. *Sci Rep*. 2022;12(1):5848. <https://doi.org/10.1038/s41598-022-09775-9>.
- Wald A. The Fitting of Straight Lines if Both Variables are Subject to Error. *Ann Math Stat*. 1940;11(3):284–300. <https://doi.org/10.1214/aoms/1177731868>.
- Hartford J, Lewis G, Leyton-Brown K, Taddy M. Deep IV: a flexible approach for counterfactual prediction. In: *International Conference on Machine Learning*. PMLR; 2017. p. 1414–1423.
- Puli A, Ranganath R. General Control Functions for Causal Effect Estimation from IVs. *Adv Neural Inf Process Syst*. 2020;33:8440–51.
- Greenland S. Basic methods for sensitivity analysis of biases. *Int J Epidemiol*. 1996;25(6):1107–16. <https://doi.org/10.1093/ije/25.6.1107-a>.
- Candido dos Reis FJ, Wishart GC, Dicks EM, Greenberg D, Rashbass J, Schmidt MK, et al. An updated PREDICT breast cancer prognostication and treatment benefit prediction model with independent validation. *Breast Cancer Res*. 2017;19(1):58. <https://doi.org/10.1186/s13058-017-0852-3>.
- Van Amsterdam WAC, Ranganath R. Conditional average treatment effect estimation with marginally constrained models. *J Causal Infer*. 2023;11(1):20220027. <https://doi.org/10.1515/jci-2022-0027>.
- Keogh RH, van Geloven N. Prediction under interventions: evaluation of counterfactual performance using longitudinal observational data. 2024. <https://doi.org/10.48550/arXiv.2304.10005>.
- van Amsterdam WAC, van Geloven N, Krijthe JH, Ranganath R, Ciná G. When accurate prediction models yield harmful self-fulfilling prophecies. 2024. <https://doi.org/10.48550/arXiv.2312.01210>.
- Prosperi M, Guo Y, Sperrin M, Koopman JS, Min JS, He X, et al. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nat Mach Intel*. 2020;2(7):369–75. <https://doi.org/10.1038/s42256-020-0197-y>.
- van Geloven N, Swanson SA, Ramspek CL, Luijken K, van Diepen M, Morris TP, et al. Prediction meets causal inference: the role of treatment in clinical prediction models. *Eur J Epidemiol*. 2020;35(7):619–30. <https://doi.org/10.1007/s10654-020-00636-1>.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.