Optimized Piecewise Exponential Modeling using Dynamic Programming: An Illustration from Treatment-Naïve Advanced Melanoma

MSR 157

Oguzhan Alagoz[1] and Murat Kurt[2]

[1]University of Wisconsin-Madison, Madison, WI, USA; [2]Iovance Biotherapeutics Inc, San Carlos, CA, USA

OBJECTIVE

- **Develop a systematic and scalable dynamic programming (DP) approach to optimize the placement of knots for piecewise survival modeling.**
- **Illustrate the predictive ability and flexibility of the approach in a case-study.**

INTRODUCTION

- Extrapolation of survival outcomes beyond limited trial follow-up is a common and integral part of economic evaluations in health technology assessments.
- Extrapolation of reported survival outcomes conventionally relies on standard parametric models (SPMs) due to their ability to represent the underlying hazards in a compact, clinically intuitive and easily interpretable form. However, SPMs are limited to capture only unimodal hazard functions[1].
- Novel mechanisms of action for modern oncology treatments, such as immune checkpoint inhibitors, can lead to pseudo-progression and delayed response behaviors while generating sustained response and survival benefits in the long-run. This complex trend can manifest itself in survival plateaus and sophisticated hazard trends which may not be accurately captured by SPMs.
- Spline-based and non-parametric piecewise survival models often serve as flexible alternatives to SPMs with their ability to capture multi-modal hazard functions with multiple inflection points; however the choice for the number of knots and their locations used to partitioned the survival or hazard functions may significantly influence the accuracy of these frameworks[2].
- From a practical standpoint, while number of knots are often decided by the length of follow-up in the data and the number of modes in the hazard trends, their locations are often decided in an unsystematic fashion with limited methodological guidance.

METHODS

- The proposed DP approach aimed to approximate the smoothed, cumulative hazard function with a prespecified number of linear segments. While each linear segment represents a piece of exponential distribution for survival outcomes, adjacent segments may not be continuously connected.
- For practical purposes, without loss of generality, candidate locations for knots used to partition the smoothed, cumulative hazard function were considered on a monthly basis.
- The objective was to minimize the cumulative Bayesian Information Criteria (BIC) associated with the number of knots as well as the linear segments used to approximate the corresponding partitions of the smoothed, cumulative hazard function.
- BIC score for each candidate linear segment was indirectly estimated by minimizing sum of squared-errors associated with its fit to the corresponding partition of the cumulative hazard function
- Placement of each additional know was penalized to avoid the overfitting bias
- Locations of knots were optimized in a backward recursion using estimated BIC scores associated with adjacent candidate segments and penalties associated with each knot.
- The method consisted of four main steps:

Step 1. Reconstruction of synthetic survival data from published Kaplan-Meier (KM) curves

Step 2. Generation of cumulative smoothed hazard function using reconstructed survival data

Step 3. Formulation and solution of the following recursive equations, also referred as the "optimality equations", of the DP model for a prespecified maximum number of *N* knots

$J_k(n) = \min_{i>k}\{v(k,i) + J_i(n-1)\}$ for $k = 0,1,\dots,T-1$, and $n = 1,\dots,N$, with the boundary conditions

$J_k(0) = BIC(k,T)$, $k = 0,1,\dots,T-1$ and $J_T(n) = 0$, $n = 0,\dots,N$; where

$J_k(n)$: The optimal value function of the DP, i.e., the overall minimum BIC score associated with modeling the cumulative hazard function beyond month *k* with *n* knots.

$J_i(n-1)$: The minimum BIC score associated with modeling the cumulative hazard function beyond month *i* with *n-1* knots.

$v(k,i)$: Optimized BIC score associated with the linear segment modeling the partition of cumulative hazard function between months *k* and *i* (inclusive).

$BIC(k,T)$: The BIC score associated with the linear segment modeling the partition of cumulative hazard function between months *k* and *T* (inclusive).

T: Maximum reported time in months along the KM curve.

N: Maximum number of knots allowed.

Step 4. Identification of knot locations based on optimal solutions to the recursive equations of DP

- The DP algorithm works iteratively in a backward fashion by breaking the problem into smaller overlapping subproblems.[3] By eliminating redundant calculations and storing the solutions of subproblems, DP significantly reduces the exponential space and time complexity of knot-placement problem to a polynomial order. For example, for placement of 3 knots across 120 candidate months, there are 261,800 candidate solutions. With the DP approach, an optimal solution can be constructed without complete enumeration of these solutions and by solving only 7,260 subproblems (≥97.2% reduction in problem size) which are of smaller scale than the original problem.
- For any given month evaluated as a candidate for a knot placement, the algorithm builds a route with the lowest BIC score to all possible future months as the next candidate knot location. This procedure uses previously computed subpaths connecting the subsequent knot location to the end of the follow-up and their corresponding minimum BIC scores, and the route from the current month with minimum overall BIC score is labeled as optimal. The recursion is repeated until the time of randomization.
- The DP-approach was illustrated on published progression-free survival (PFS) and overall survival (OS) data from the CheckMate-067 trial in treatment-naive advanced melanoma.[4]
- CheckMate-067 was a global, phase III, randomized controlled trial evaluating efficacy and safety of nivolumab (NIVO) plus ipilimumab (IPI), NIVO alone, and IPI alone as first-line treatments for treatment-naïve advanced melanoma[4]. Recently, published data from this study with over 10 years of follow-up offer one of the most comprehensive long-term views on the long-term efficacy of immune checkpoint inhibitors in treatment-naïve advanced melanoma.
- Because of mature data demonstrating sustained, long-term survival benefit with clear plateaus, CheckMate-067 is often used in cost-effectiveness analyses and case studies analyzing the performance of novel survival modeling approaches for immune checkpoint inhibitors[5-11].
- Even though the DP approach can flexibly accommodate any pre-specified number of knots to partition the cumulative hazard function, to avoid overfitting bias, for each endpoint, up to 2 knots were used in the case study.
- For statistical reliability of long-term extrapolations, in both PFS and OS analysis, tail segments were required to cover ≥20% of the total number of events in the data. Sensitivity analyses explored piecewise modeling by restricting this fraction of events covered by the tail segment to ≥10%.
- The R programming language was used to implement the DP algorithm, including the *bshazard* package, which was employed to generate the smoothed hazard functions.

CONCLUSIONS

- **The DP approach offers a transparent, scalable and flexible framework to piecewise survival modeling with optimized knot locations connecting exponentially distributed segments.**
- **Unlike spline-based models, proposed approach can be augmented with additional constraints specifying distance between the locations of knots and number of events informing each segment of the piecewise model for the reliability of long-term survival projections.**
- **The proposed approach can adequately capture sophisticated time-varying hazard trends in a cohort and can be implemented on historical data to gain insights on sample size and power calculations in clinical trial design.**

RESULTS

- Across all arms, first and second knots were placed between months 21-24 and months 32-38, respectively, for OS and between months 3-9 and months 6-16, respectively, for PFS (**Table 1**).
- As a benchmark, in the NIVO arm, application of splines with 2-knots on the hazard scale placed the knots to months 3 and 7 for PFS, and to months 9 and 24 for OS.
- In the sensitivity analysis, across all arms and endpoints, optimized locations of knots were no earlier than their counterparts in the primary analysis (**Table 2**).
- For each endpoint, optimized locations of knots and fractions of events covered in each segment showed high similarity between NIVO-containing arms (**Table 1** and **Table 2**).
- In both base-case and sensitivity analyses, while the distribution of OS events across the segments in the IPI arm showed similarity to those in the NIVO-containing arms; the distribution of PFS events across the segments in the IPI arm was substantially different than those in the NIVO-containing arms.
- Compared to NIVO-containing arms, for each endpoint, optimized locations of knots were mostly earlier in time in the IPI arm in the base-case and sensitivity analyses.
- Visual inspections showed that proposed DP-based approach placed the knots around the times when the curvature of the smoothed, cumulative hazard function displays a major change (**Figure 1**).
- Variation in optimized locations of knots for both endpoints with respect to treatment and to the threshold number of events in the tail segment show the utility of the DP-approach.

Table 1. Base-case analysis results under the restriction that ≥20% of the events are allowed to occur after the last knot location

Treatment	Endpoint	Optimized Knot Locations (month)	Fraction of events: Prior to 1st knot	Fraction of events: Between the knots	Fraction of events: After 2nd knot
IPI	OS	24 & 32	68%	11%	20%
NIVO	OS	22 & 37	60%	18%	21%
NIVO+IPI	OS	21 & 38	62%	15%	23%
IPI	PFS	3 & 6	11%	62%	27%
NIVO	PFS	8 & 15	71%	9%	20%
NIVO+IPI	PFS	9 & 16	64%	14%	21%

Table 2. Sensitivity analysis results under the restriction that ≥10% of the events are allowed to occur after the last knot location

Treatment	Endpoint	Optimized Knot Locations (month)	Fraction of events: Prior to 1st knot	Fraction of events: Between the knots	Fraction of events: After 2nd knot
IPI	OS	37 & 54	83%	7%	11%
NIVO	OS	44 & 61	85%	4%	12%
NIVO+IPI	OS	50 & 73	85%	4%	11%
IPI	PFS	3 & 8	11%	72%	18%
NIVO	PFS	18 & 29	83%	7%	11%
NIVO+IPI	PFS	13 & 29	77%	12%	10%

FIGURE 1. Visual illustration of the DP approach applied to the PFS and OS data from the NIVO arm of CheckMate-067
A) Smoothed cumulative hazard function obtained from reported PFS data and corresponding optimized piecewise linear cumulative hazard function; B) Reported and modeled PFS using optimized piecewise exponential approach; C) Smoothed cumulative hazard function obtained from reported OS data and corresponding optimized piecewise linear cumulative hazard function; D) Reported and modeled OS using optimized piecewise exponential approach

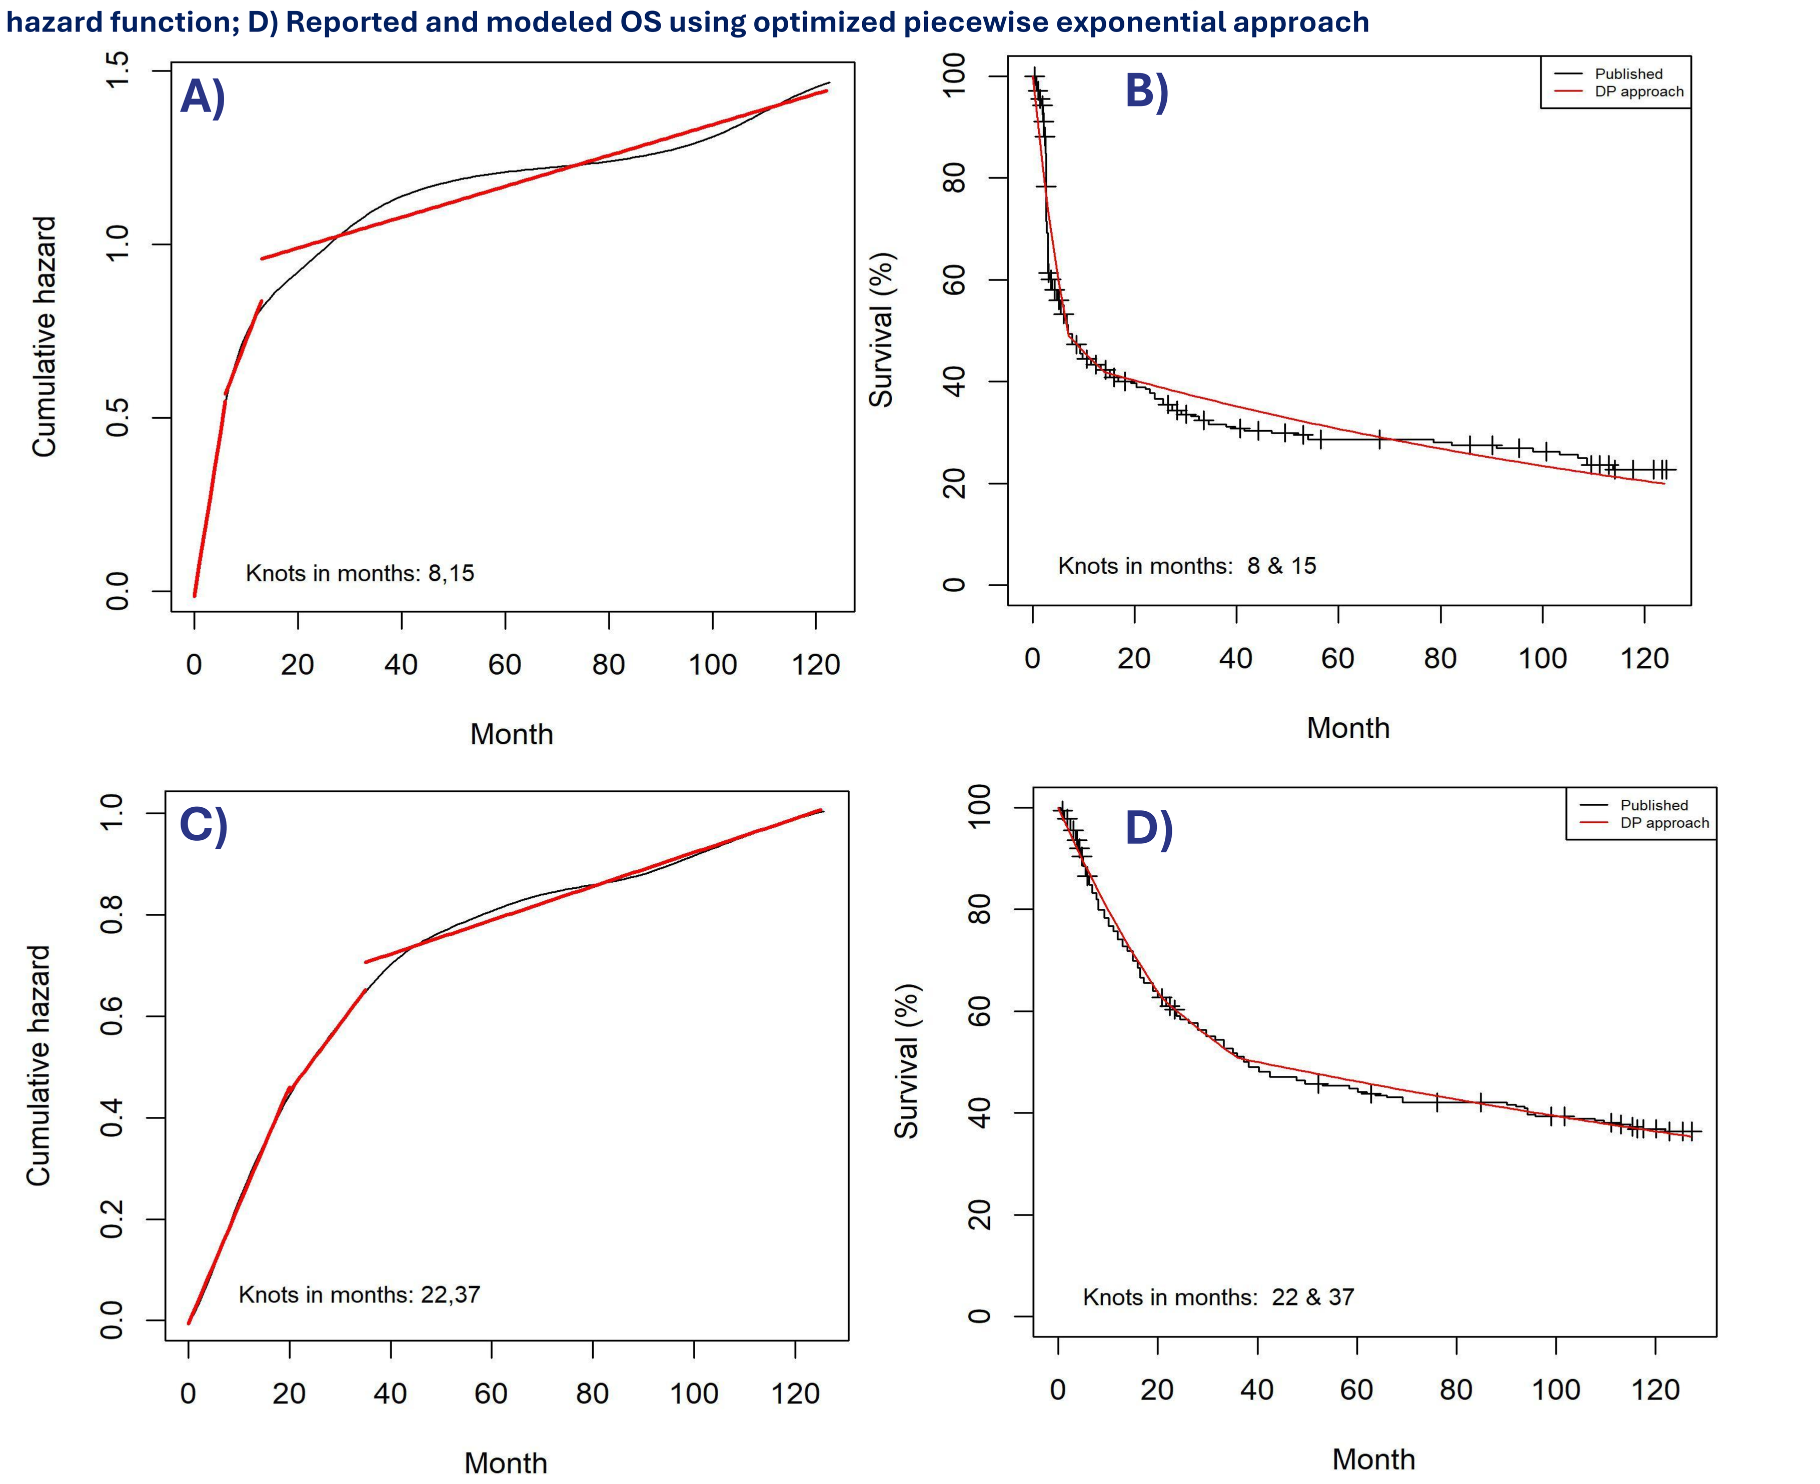

KM plots are adapted from Wolchok et al. (2024) reporting 10-year follow-up data from the Checkmate 067 study. Vertical markers represent censoring information.

LIMITATION

- The DP-approach assumes an exponential distribution for each segment of the piecewise model, however it can be generalizable to accommodate other parametric distributions which may require non-linear regression methods to calculate the BIC score associated with the segmented fits.

REFERENCES: 1. Latimer N. (2013) "NICE DSU Technical Support Document 14: Survival Analysis For Economic Evaluations Alongside Clinical Trials - Extrapolation With Patient-Level Data". Available from: https://sheffield.ac.uk/nice-dsu/tsds/full-list. Accessed on September 4, 2025.
2. Rutherford MJ, et al. (2020) "NICE DSU Technical Support Document 21: Flexible Methods for Survival Analysis" Available from https://sheffield.ac.uk/nice-dsu/tsds/full-list. Accessed on September 4, 2025.
3. Bertsekas DP (2017) "Dynamic Programming and Optimal Control", Volume 2, Fourth Edition. Athena Scientific, Belmont, MA, USA
4. Wolchok JD, et al. (2024) "Final, 10-year outcomes with nivolumab plus ipilimumab in advanced melanoma." *NEJM*, 392(1): 11-22.
5. Kurt M, et al. (2025) "Do Lifetables Overestimate Non-Cancer-Specific Survival in Oncology? A Case Study From Treatment-Naïve Advanced Melanoma". *Value in Health*, 28(6): S277-S278
6. Kurt M, et al. (2025) "Assessing the Impact of PD-L1 and BRAF Biomarkers on Long-Term Survivorship Rates Among Treatment-Naive Advanced Melanoma Patients Receiving Immune Checkpoint Inhibitors". *Value in Health*, 28(6): S298
7. Baker T, et al. (2022) "Cost-Utility of Nivolumab Plus Ipilimumab in First-Line Treatment of Advanced Melanoma in the United States: An Analysis Using Long-Term Overall Survival Data from Checkmate 067". *Pharmacoeconomics Open*, 6(5):697–710
8. Paly VF, et al. (2021) "Assessing the Impact of Modeling Non-Disease-Related Mortality on Long-Term Survivorship Rates in Previously Untreated Advanced Melanoma: A Case Study from CheckMate 067". *Value in Health*, 24(1): S55
9. Mohr P, et al. (2020) "Estimating long-term survivorship in patients with advanced melanoma treated with immune-checkpoint inhibitors: Analyses from the phase III CheckMate 067 trial". *Annals of Oncology*, 31(4): S747
10. Green N, et al. (2025) "A Bayesian Hierarchical Mixture Cure Modelling Framework to Utilize Multiple Survival Datasets for Long-Term Survivorship Estimates: A Case Study From Previously Untreated Metastatic Melanoma". *Statistics in Medicine,* 44(13-14):e70132
11. Paly VF, et al. (2022) "Heterogeneity in Survival with Immune Checkpoint Inhibitors and Its Implications for Survival Extrapolations: A Case Study in Advanced Melanoma". *MDM Policy & Practice*, 7(1): 1-15

ACKNOWLEDGMENTS: The findings and opinions expressed in this study are solely those of the authors and independent of the views of Iovance Biotherapeutics, Inc.
PRESENTED AT: International Society for Pharmacoeconomics and Outcomes Research (ISPOR) Europe 2025 Meeting | 9-12 November, 2025 | Glasgow, Scotland, UK