



HAL
open science

Survival models for dose-finding clinical trials in oncology

Anaïs Andrillon

► **To cite this version:**

Anaïs Andrillon. Survival models for dose-finding clinical trials in oncology. Cancer. Université Paris Cité, 2022. English. NNT : 2022UNIP5153 . tel-04741854

HAL Id: tel-04741854

<https://theses.hal.science/tel-04741854v1>

Submitted on 17 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

**École doctorale 393, Pierre-Louis de Santé Publique :
épidémiologie et sciences de l'information biomédicale**

Unité de recherche : INSERM, U1153 Centre de Recherche Épidémiologie
et Statistique Sorbonne Paris Cité (CRESS), équipe ECSTRRA

UNIVERSITÉ PARIS CITÉ

THÈSE DE DOCTORAT

Discipline : Santé Publique, Spécialité : Biostatistiques et biomathématiques

Dirigée par Sylvie CHEVRET et co-encadrée par Lucie BIARD

Présentée et soutenue publiquement, le 09/12/2022 par

Anaïs ANDRILLON

Survival models for dose-finding clinical trials in oncology

Composition du Jury

Xavier PAOLETTI Professeur des universités - praticien hospitalier UVSQ - Université Paris-Saclay & Institut Curie	Rapporteur
Thomas JAKI Professor, PhD Université de Cambridge & Université de Regensburg	Rapporteur
Nolan WAGES Professor, PhD Université du Commonwealth de Virginie	Examineur
Shing M. LEE Associate Professor, PhD Mailman School of Public Health, Université de Columbia	Examinatrice
Sylvie CHEVRET Professeur des universités - praticien hospitalier Université Paris Cité & Hôpital Saint-Louis	Directrice de thèse
Lucie BIARD Maître de conférences des universités Université Paris Cité & Hôpital Saint-Louis	Membre invité et co-encadrante de thèse

REMERCIEMENTS

À mon jury de thèse,

I sincerely thank Professor Thomas Jaki and Professor Xavier Paoletti, who agreed to review my thesis work. These thanks are also addressed to Professor Nolan Wages, who agreed to participate in my thesis jury. My sincere thanks also go to Professor Shing Lee, who kindly welcomed me to her research team, at Columbia University. Thank you, Shing, for your time and expertise, which was instrumental to this project. It is an honor for me to have all of you among the members of my jury. I am very grateful to be evaluated by such recognized researchers.

À mes directrices de thèse,

Lucie, un immense merci pour ton encadrement, ta bienveillance et pour avoir toujours été disponible et à l'écoute. Ton aide a été si précieuse dans la conduite de ces travaux. Merci Sylvie pour m'avoir chaleureusement proposé de travailler avec toi il y a maintenant quatre ans et surtout pour m'avoir accordé ta confiance. J'ai tellement appris à tes côtés, tu m'as fait grandir en tant que chercheur et en tant que personne. Je suis très reconnaissante pour le travail que nous avons accompli toutes les trois. Je ne vous remercierai jamais assez pour votre soutien, vos conseils et vos encouragements, même lorsque j'étais de l'autre côté de l'Atlantique. Je souhaite à tous les jeunes chercheurs de débiter leur carrière entre de si bonnes mains.

À tous les membres du SBIM,

Merci à toute l'équipe pour votre accueil chaleureux. J'ai pris beaucoup de plaisir à échanger et travailler avec vous. Merci en particulier aux doctorants (et Docteurs) Lilith, Menyssa, Luana, Jacques-Emmanuel, Valentin, et Guillaume pour tous les agréables moments partagés. Merci Anouk pour avoir été un binôme d'exception. Déborah, un grand merci pour ton soutien et tes adorables attentions depuis mon arrivée au SBIM.

À mes amis,

Anne, Morgane, Florine, Maxence, Eliette, Joséphine, Léa, Raphaël, Momo, Loïc, Amine, Isaac, et Tony. Merci à vous pour tous ces bons moments passés à vos côtés. Soyez assurés que mon amitié restera fidèle. Merci Justine pour ton soutien sans faille et cette si belle amitié qui nous unies depuis tant d'années. Merci aussi pour avoir été la meilleure partenaire de bibliothèque.

À ma famille,

Je remercie du fond du cœur ma famille, en particulier mes parents, mon frère Antoine, et ma sœur Angélique pour leur confiance et leur amour qui ont fait de moi la personne que je suis aujourd'hui. À mes parents, merci de m'avoir permis de réaliser mes études et pour votre soutien dans chacun de mes projets. Merci à tous de m'avoir permis de m'évader le temps de week-end ou vacances qui me faisaient le plus grand bien.

À mon conjoint,

Merci Joshua de croire en moi et d'être une source de bien-être, de soutien, et de réconfort. Tout au long de ces années, tu as toujours su trouver les mots pour me motiver. Tes conseils me sont et me seront toujours d'une aide précieuse. Je te remercie pour ta relecture attentive de ce manuscrit. Avec tout mon amour, merci.

Contents

Remerciements	v
Abstract	viii
Résumé	ix
Résumé long	x
List of scientific productions	xvi
Acronyms	xx
1 Introduction	1
1.1 Oncology drug development, a long, costly, and high-risk process	2
1.2 Statistical challenges in early phase oncology trials	3
1.2.1 Managing ethical concerns	3
1.2.2 Paradigm of cytotoxic chemotherapies	4
1.2.3 New classes of cancer treatment	6
1.3 Contributions	9
1.4 Manuscript organization	10
2 State of the art	12
2.1 Phase I designs	14
2.1.1 Ruled-based designs	14
2.1.1.1 ‘3 + 3’ design	15
2.1.1.2 Up-and-down designs	16
2.1.2 Model-based designs	20
2.1.2.1 CRM	22
2.1.2.2 Extensions of the CRM	24
2.1.3 Model-assisted designs	28
2.1.3.1 mTPI and Keyboard designs	28
2.1.3.2 BOIN design	30

2.1.4	Designs for a delayed toxicity outcome	33
2.1.4.1	Extensions of the CRM	34
2.1.4.2	TiTE-BOIN design	36
2.2	Phase I/II designs	37
2.2.1	Ruled-based designs	40
2.2.2	Model-based designs	41
2.2.2.1	Efficacy-toxicity trade-offs designs	41
2.2.2.2	Extensions of the CRM	44
2.2.3	Model assisted designs	49
2.2.4	Designs for delayed toxicity and efficacy outcomes	55
2.2.4.1	Extension of the CRM	55
2.2.4.2	Survival design	55
2.2.4.3	Efficacy-toxicity trade-offs designs	57
2.2.4.4	TiTE-BOIN designs	59
2.3	Nonparametric optimal Benchmark	60
3	Survival-CRM, informative Survival-CRM	63
3.1	Introduction	63
3.2	Published manuscript	67
3.3	Supplementary information	84
3.3.1	Benchmark for right-censored toxicity endpoint	84
3.3.1.1	Sensitivity analyses	87
3.3.1.2	Modified Surv-CRM	87
4	Survival-CRM-12	90
4.1	Introduction	90
4.2	Published manuscript	95
4.3	Supplementary information	110
4.3.1	Benchmark for right-censored endpoints and competing risks	110
4.3.2	Model parameters calibration	111
4.3.2.1	Dose skeletons calibration	111
4.3.2.2	Least informative variance	114
4.3.2.3	Calibration process	115
4.3.3	Sensitivity analyses	117
4.3.3.1	Simulation with varying sample size	117
4.3.3.2	Simulation of correlated time-to-events	120
5	Discussion, perspectives, and conclusion	122
	Bibliography	130

ABSTRACT

Survival models for dose-finding clinical trials in oncology

In traditional dose-finding studies, Dose-Limiting Toxicity (DLT) is determined within a fixed time observation window, where DLT is often defined as a binary outcome. In oncology dose-finding trials, for molecularly-targeted agents and immunotherapies with prolonged administration and complicated toxicity profiles, longer toxicity observation windows are required. Strategies have been proposed for phase I designs to allow a fast and continuous patient accrual and reduce trial duration but still consider the toxicity outcome as binary and fully observed at the end of the trial.

The first methodological development of the thesis concerns the Survival-Continual Reassessment Method (Surv-CRM), a phase I dose-finding design for right-censored toxicity endpoints, built on the CRM but using survival models. The Surv-CRM aims to identify the Maximum Tolerated Dose (MTD), defined as the dose with the estimated cumulative incidence of toxicity at the end of the observation window closest to a pre-specified probability of toxicity target. Moreover, in this setting of phase I trials enrolling patients in an advanced stage of diseases, the observation of a late-onset toxicity endpoint could be precluded by trial discontinuation due to death, progression, patient withdrawal, or physician discretion, defining a competing event to toxicity. To handle such treatment discontinuations, we propose the informative Survival-CRM (iSurv-CRM).

Discontinuing the trial due to disease progression could be used as efficacy information to identify the Optimal Dose (OD), resulting in phase I/II clinical trial. The second methodological development of the thesis was the Survival-CRM-12 (Surv-CRM-12), a phase I/II design to address this issue in a competing risks framework and identify dose with acceptable toxicity risk and minimum progression risk, relying on cause-specific hazards for DLT and progression without DLT.

For all the proposed methods, model parameters were estimated using Bayesian inference. The methods' performances were evaluated on simulations and compared to the performances of other designs and a nonparametric optimal benchmark. We found that the proposed designs present satisfying operating characteristics to select the correct dose at the end of the trial, and to allocate patients at the correct dose during the trial, and in terms of safety.

Keywords : Oncology; Phase I; Phase I/II; Late-onset toxicity; Treatment discontinuation; Disease progression; Competing risks; Survival data; Benchmark; Bayesian inference.

RÉSUMÉ

Modèles de survie pour les essais cliniques de recherche de dose en cancérologie

Dans les schémas d'essais cliniques précoces traditionnels, la toxicité dose limitante (TDL) est définie comme un critère binaire et mesurée sur une fenêtre d'observation fixe. Les nouvelles classes de traitement en cancérologie, comme les immunothérapies et les thérapies ciblées, ont une administration prolongée et des profils de toxicité et d'efficacité différents des chimiothérapies cytotoxiques, nécessitant des fenêtres d'observation de la toxicité prolongées. Des stratégies ont été proposées pour des essais de phase I afin de permettre un recrutement rapide et continu de patients et de réduire la durée de l'essai, mais elles considèrent toujours le critère de toxicité comme binaire et complètement observé à la fin de l'essai.

Le premier travail de thèse porte sur le développement d'un design de recherche de dose de phase I, la Surv-MRS (Méthode de Réévaluation Séquentielle basée sur un modèle de survie), une extension de la Méthode de Réévaluation Séquentielle (MRS) prenant en compte les critères censurés à droite en utilisant des modèles de survie. La Surv-MRS vise à identifier la dose maximale tolérée, définie comme la dose dont l'incidence cumulée estimée de la toxicité à la fin de la fenêtre d'observation est la plus proche d'une probabilité de toxicité cible prédéfinie. De plus, l'observation de patients à un stade avancé de la maladie sur une fenêtre prolongée favorise la survenue d'évènement en compétition, venant empêcher l'évaluation de la toxicité, comme le décès, la progression de la maladie, le retrait du patient ou une décision du médecin. Pour gérer cette censure informative due à l'interruption du traitement sur la fenêtre d'observation de la toxicité, nous avons proposé la iSurv-MRS.

L'arrêt de l'essai en raison de la progression de la maladie peut aussi être utilisé comme une information sur l'efficacité pour identifier la dose optimale, résultant en un essai clinique de phase I/II. Le deuxième travail de thèse a donc porté sur le développement de la Surv-MRS-12, un schéma de phase I/II qui inclut à la fois les informations relatives à la toxicité et à l'efficacité dans le processus d'attribution de dose. Il vise à identifier une dose optimale, définie comme la dose tolérée en termes de toxicité, optimisant le critère d'efficacité, c'est-à-dire le temps jusqu'à la progression de la maladie.

Pour l'ensemble des méthodes proposées, les paramètres du modèle ont été estimés par inférence bayésienne. Les performances des méthodes ont été évaluées par simulation et comparées aux performances d'autres méthodes et d'un parangon optimal non paramétrique (*benchmark*). Les schémas proposés présentent des caractéristiques satisfaisantes en termes d'identification de la dose correcte à la fin de l'essai, d'allocation de patients à la dose correcte au cours de l'essai, et en termes de risque de toxicité.

Mots-clés : Oncologie; Phase I; Phase I/II; Toxicité tardive; Arrêt du traitement; Risques en compétition; Progression de la maladie; Données de survie; Benchmark; Inférence bayésienne.

RÉSUMÉ LONG

Contexte : les essais cliniques de recherche de dose en cancérologie

Les essais de phase I sont les premiers essais cliniques conduits sur l'Homme et sont réalisés, en cancérologie, sur un petit échantillon de sujets malades pour des contraintes éthiques évidentes. L'objectif majeur de ces essais, encore appelés essais de recherche de dose, est de définir une dose médicamenteuse pour la suite du développement, sur des critères de tolérance. Les préoccupations éthiques dictent les principes de détermination de la dose : escalader lentement pour protéger les patients d'une sur-exposition à une dose toxique, voire létale, mais pas trop lentement pour éviter de traiter un grand nombre de patients à des doses sous-thérapeutiques. Par ailleurs, une sélection inappropriée du niveau de dose se répercute sur l'ensemble du processus de développement et peut entraîner des coûts considérables pour l'économie. Une mauvaise décision quant au niveau de dose peut soit aboutir à l'arrêt du développement d'un médicament potentiellement intéressant, soit favoriser l'évaluation d'un médicament plutôt inactif, voire dangereux. Concevoir un essai clinique de phase I avec une précision suffisante dans l'estimation des probabilités de toxicité sur la base d'échantillons de si petite taille tout en respectant des contraintes éthiques représente donc un défi statistique majeur. Pour ces raisons, les essais de recherche de dose sont séquentiels et adaptatifs, car (i) les patients sont inclus séquentiellement dans l'essai, et (ii) la dose à administrer est adaptée en fonction des informations disponibles sur les patients déjà inclus dans l'essai. Une approche communément utilisée repose sur la Méthode de Réévaluation Séquentielle (MRS) (*Continual Reassessment Method*) [O'Quigley, Pepe, and Fisher, 1990], où l'attribution de la dose est basée sur l'optimisation d'une fonction dépendant des estimations de probabilités de toxicité et de contraintes sur la tolérance.

Problématique : les nouvelles classes de traitement en cancérologie

Historiquement, les schémas d'essais cliniques précoces ont été développés pour les chimiothérapies cytotoxiques associées à des toxicités aiguës et sévères, avec l'hypothèse de stricte monotonie des relations dose-efficacité et dose-toxicité. Cela a justifié la réalisation d'essais cliniques de phase I centrés uniquement sur l'évaluation de la toxicité,

en substitution de l'efficacité, avec l'objectif d'identifier la dose dite "maximale tolérée" (DMT), indépendamment de son efficacité évaluée en phase II. En revanche, les agents anti-cancéreux tels que les immunothérapies et les thérapies ciblées, aux profils pharmacologiques d'efficacité et de toxicité différents des chimiothérapies cytotoxiques, ont remis en cause ces schémas d'essais cliniques de recherche de dose traditionnels. En premier lieu, alors que les chimiothérapies cytotoxiques sont usuellement administrées par cycles, en nombre limité, les thérapies ciblées peuvent s'administrer sur un mode continu et prolongé jusqu'à la progression de la maladie ou résistance du patient, nécessitant de ce fait des fenêtres d'observation prolongées dans les essais pour la prise en compte de toxicités tardives et des critères de jugements de substitution observables plus précocement (par exemple la survie sans progression au lieu de la survie globale). Ainsi, dans un essai de recherche de dose, en fonction du rythme d'arrivée des patients, les données de toxicité de certains patients déjà traités dans l'essai peuvent être incomplètes (ou censurées à droite) au moment d'une nouvelle inclusion. De plus, l'hypothèse de monotonie des relations dose-efficacité et dose-toxicité sur laquelle reposent la plupart des schémas de recherche de dose développés pour les chimiothérapies cytotoxiques, semble pouvoir ne pas être respectée pour ces agents anti-cancéreux. Par conséquent, la sélection de la DMT, laisse alors place à l'identification d'une "dose optimale", définie comme la dose tolérée ayant une activité pharmacologique maximale ou d'une "dose la plus souhaitable" offrant le meilleur équilibre entre toxicité et efficacité. La distinction entre phases I et II pour l'évaluation d'abord de la toxicité puis de l'efficacité est alors souvent remplacée par l'évaluation conjointe de la toxicité et de l'efficacité, résultant en un essai de phase I/II.

La plupart des schémas existants font l'hypothèse que les deux critères de toxicité et d'efficacité sont complètement observables pendant l'essai et peuvent être chacun mesurés indépendamment l'un de l'autre. Or, en pratique, dans un essai clinique, l'observation de l'un des critères peut demeurer incomplète du fait de la survenue d'un événement en lien avec l'autre critère : la survenue d'une Toxicité Dose-Limitante (TDL) ou la progression de la maladie peuvent engendrer l'arrêt du traitement et empêcher l'observation complète de l'autre critère, respectivement de l'efficacité ou la toxicité. En effet, les thérapies ciblées sont généralement administrées jusqu'à l'apparition d'une toxicité, d'une progression de

la maladie, du retrait du consentement ou à la discrétion du médecin, plutôt qu'après un nombre prédéterminé de cycles. Etant donné que la survenue d'un de ces événements engendrent l'arrêt du traitement et empêche ainsi l'observation complète de l'autre critère, la question de risques en compétition se pose.

Certaines méthodes comme la MRS dépendant du temps d'évènement (*Time-To-Event Continual Reassessment Method*) [Cheung and Chappell, 2000], ont été développées pour gérer le recrutement continu de patients pendant l'essai et permettent de réduire la durée de l'essai par rapport à la MRS. Cependant pour l'estimation de la DMT, ces modèles reposent généralement sur des critères de toxicité binaires, ne tenant pas compte du délai d'apparition de l'évènement. Pour gérer les interruptions de traitement, l'une des approches couramment utilisées consiste à considérer (à tort) que la censure est indépendante du temps de survenue de la toxicité, conduisant soit au remplacement du patient, soit à l'implémentation de stratégies pratiques pour gérer ces observations incomplètes. Cependant, la stratégie de remplacement conduit à une augmentation de la durée et du coût de l'étude, et génère un biais potentiel de sélection de patients en meilleure santé [Winther et al. 2016] et la mise en place de stratégies ad-hoc impacte les performances des schémas de recherche de dose [Biard, Cheng, et al. 2021].

Objectif

L'objectif de cette thèse a donc été de développer des schémas bayésiens, adaptatifs pour les essais de recherche de dose en cancérologie utilisant des modèles de survie adaptés aux données dépendantes des temps d'évènement et permettant de gérer trois types de censure à droite: (i) une censure administrative se produisant pendant l'essai, lorsque l'attribution de dose d'une nouvelle inclusion doit être décidée alors que le suivi des patients inclus peut être encore incomplets étant donné que seule une fraction de la fenêtre de toxicité a été observée et que le patient n'a pas encore développé de TDL (ii) une censure administrative concernant tous les patients sans TDL à la fin de la fenêtre d'observation, et (iii) une censure informative due aux sorties d'essai, susceptibles de se produire chez ces patients à un stade avancé de la maladie en raison du manque d'efficacité du traitement (décès, progression de la maladie, décision du médecin ou toute autre raison).

Contributions

Pour gérer les deux premières situations de censure, nous avons développé la méthode de réévaluation séquentielle basée sur un modèle de survie (Surv-MRS), un schéma de recherche de dose de phase I modélisant le critère de toxicité par un indicateur d'événement et le temps d'occurrence de l'événement ou le temps de censure. L'objectif de la Surv-MRS est d'identifier la dose dont l'incidence cumulée de toxicité à la fin de la fenêtre d'observation est la plus proche d'une probabilité de toxicité cible. Les performances de la Surv-MRS ont été comparées à celles de la MRS dépendant du temps d'évènement par le biais d'une étude de simulation. Bien que les modèles mathématiques soient différents, les deux schémas présentent des performances proches en termes de sélection correcte de la DMT à la fin de l'essai. En revanche, la Surv-MRS a présenté des résultats plus satisfaisants que la MRS dépendant du temps d'évènement en termes d'exposition des patients à des doses toxiques pendant l'essai.

Pour traiter la troisième situation de censure due à la survenue d'un événement en compétition dans la fenêtre d'observation de la toxicité (quelle que soit la cause) dans le cadre d'un essai de phase I, nous avons étendu la Surv-MRS à la iSurv-MRS, gérant les censures informatives. Dans le cadre de modèles de survie avec risques en compétition, deux approches permettent considérer les processus d'événements, l'une privilégiant la fonction de risque cause-spécifique, l'autre la fonction de sous-répartition. L'approche privilégiant le risque cause-spécifique revient à évaluer ce qui peut arriver au patient dans la situation hypothétique où les autres types d'événement en compétition ont été supprimés. L'approche privilégiant la fonction de sous-répartition considère un univers où tous ces événements sont possibles. Pour estimer les probabilités de toxicité dans le cadre de la iSurv-MRS, nous avons choisi de cibler l'incidence cumulée des TDL à la fin de la fenêtre d'observation [Putter et al. 2006]. Les performances de la iSurv-MRS, ont été comparées à celles de la MRS dépendant du temps d'évènement. Les résultats de l'étude de simulation montrent que les performances de la iSurv-MRS étaient nettement meilleures que celles de la MRS dépendant du temps d'évènement, d'autant plus que la censure informative était élevée.

Enfin, nous avons ensuite considéré que la sortie d'étude en raison de la progression de la maladie, en compétition avec la toxicité, pouvait être une information sur l'efficacité du traitement utilisable dans le choix de la dose correcte. En effet, la survie sans progression est de plus en plus fréquemment utilisée comme critère d'efficacité dans les essais de phase II en oncologie. Notamment dans le cadre des thérapies ciblées et l'immunothérapie où la stabilité de la maladie est souvent considérée comme un signal positif, alors que la survenue d'une progression de la maladie est traitée comme une réponse d'efficacité négative [Borcoman et al. 2019]. Nous avons donc développé la Surv-MRS-12, un schéma de recherche de dose de phase I-II dont l'objectif est d'identifier une dose optimale, définie comme une dose tolérable en terme de toxicité et optimisant le critère d'efficacité, c'est-à-dire minimisant la probabilité de progression de la maladie. Des études de simulation comparant le schéma proposé à un schéma existant de phase I/II [Takeda, Morita, and Taguri, 2020] ont montré que la Surv-MRS-12, présente des propriétés favorables en termes de sélection de la dose correcte et d'exposition des patients à des doses toxiques pendant l'essai.

Pour l'ensemble des méthodes développées, des études de sensibilités ont été réalisées pour déterminer la robustesse de la méthode à sélectionner avec précision les doses correctes face à divers scénarios de probabilité de toxicité et d'efficacité, de distributions de temps d'événements, de rythme d'entrée des patients dans l'essai, de taille d'échantillon. Un paragon (*benchmark*) non paramétrique a également été développé afin de disposer d'un outil de diagnostic pour évaluer la performance des nouveaux schémas proposés [O'Quigley, Paoletti, and Maccario, 2002]. Le paragon non paramétrique fournit une estimation de la limite supérieure du pourcentage de sélection de la dose correcte d'un schéma de recherche de dose pour un scénario donné et pour une taille d'échantillon fixe car il évalue les performances du schéma dans une situation optimale de données complètes c'est-à-dire comme si les résultats pouvaient être observés à tous les niveaux de dose pour tous les patients. Cela s'oppose aux essais cliniques réels, où les patients ne reçoivent qu'un seul niveau de dose, ce qui entraîne des informations incomplètes en raison de la non-observation des résultats à tous les autres niveaux de dose. Enfin, pour l'ensemble des méthodes, les paramètres des modèles ont été estimés par inférence bayésienne : les moyennes a posteriori des paramètres ont été calculés pour estimer les incidences cumulées

des événements à chaque niveau de dose.

Les deux premiers travaux de développement de la Surv-MRS et la iSurv-MRS ont fait l'objet d'une première publication la revue *Journal of Biopharmaceutical Statistics* [Andrillon et al. 2020] et le troisième travail méthodologique portant sur le développement de la Surv-MRS-12 a été publié dans le journal *Statistics in Medicine* [Andrillon et al. 2022]. De plus, pour faciliter l'utilisation des méthodes proposées, nous avons développé une application web disponible en libre accès (<https://shinyapps.io/SurvCRM>). Le code R est également disponible sur Github (@[anais-andrillon/SurvivalCRM](#)).

Les deux premiers travaux de développement de la Surv-MRS et la iSurv-MRS ont fait l'objet d'une première publication la revue *Journal of Biopharmaceutical Statistics* [Andrillon et al. 2020] et le troisième travail méthodologique portant sur le développement de la Surv-MRS-12 a été publié dans le journal *Statistics in Medicine* [Andrillon et al. 2022]. De plus, pour faciliter l'utilisation des méthodes proposées, nous avons développé une application web disponible en libre accès (<https://shinyapps.io/SurvCRM>). Le code R est également disponible sur Github (@[anais-andrillon/SurvivalCRM](#)).

Conclusion

Pour conclure, l'objectif de cette thèse a été de développer des schémas adaptatifs bayésiens de recherche de dose innovants pour les essais cliniques avec des critères de toxicité censurés à droite. Ces travaux permettent d'incorporer des modèles de survie en réponse à des problématiques cliniques rencontrées dans les essais cliniques de recherche de dose en oncologie. Les méthodes proposées présentent des propriétés statistiques et éthiques satisfaisantes. Nous avons également été soucieux de proposer des modèles simples, compréhensibles et faciles à mettre en œuvre, ne nécessitant ni la spécification d'un trop grand nombre de paramètres et ni des temps de calcul excessifs.

List of scientific productions

Articles

1. Andrillon A, Chevret S, Lee SM and Biard L. Dose-finding design and benchmark for a right censored endpoint. *Journal of Biopharmaceutical Statistics*. 30(6):948-963, 2020. doi:10.1080/10543406.2020.1821702
2. Andrillon A, Chevret S, Lee SM and Biard L. Surv-CRM-12: A Bayesian Phase I/II Survival CRM for right-censored toxicity endpoints with competing disease progression. *Statistics in Medicine*. 2022; 1- 14. doi:10.1002/sim.9591

Oral communications

1. ISCB - International Society for Clinical Biostatistics conference - 2020 Edition, at Krakow, Poland (online). Survival-CRM, a phase I dose-finding design for right-censored endpoint.
2. ISCB - International Society for Clinical Biostatistics conference - 2021 Edition, at Lyon, France (online). Survival-CRM-12, a phase I/II dose-finding design for right-censored toxicity endpoints with competing disease progression.
3. GSRS - Graduate Student Research Seminar - 2021, at New York, US. Calibration approach for the Survival-CRM-12, a Phase I/II dose-finding design with dependent toxicity and disease progression endpoint.
4. EPICLIN - francophone conference of CLINical EPIdemiology - 2022 Edition, at Paris, France. Survival-CRM-12, a phase I/II dose-finding design for right-censored toxicity endpoints with competing disease progression. doi:10.1016/j.respe.2022.03.095
5. CRESS - Centre of Research in Epidemiology and Statistics doctoral day - 2022 Edition, at Paris, France. Survival models for dose-finding clinical trials in oncology.
6. Seminar on complex innovative early phase trial design in onco-hematology - 2022, at Paris, France. State of the art of designs of phase I and phase I/II designs for single-agent trials in oncology.
7. 7th Early Phase Adaptive Trials Workshop - 2022 Edition, at Cambridge, UK. Survival models for dose-finding clinical trials in oncology.

Poster

-
1. *ED (École Doctorale) annual seminar - 2020 Edition, at Saint-Malo, France (online)*.
Dose-finding design and benchmark for a right censored endpoint.

Informatics tools

1. Github: [anais-andrillon/SurvivalCRM](https://github.com/anais-andrillon/SurvivalCRM)
2. RShiny App - shinyapps.io/SurvCRM

Complementary articles

1. Andrillon A, Biard L and Lee SM. Incorporating patient-reported outcomes in dose-finding clinical trials with time-to-event toxicity endpoints. Special issue for *Journal of Biopharmaceutical Statistics*. Deadline for submission: December, 2022.
2. Andrillon A, Chevret S, Lee SM and Biard L. State of the art of designs of phase I and phase I/II designs for single-agent trials in oncology. In progress.

Awards

1. *The Franco-American Fulbright Commission* - Laureate of the french national program for PhD students.
2. *EPICLIN* - Best oral communication.

Acronyms

- AEs** Adverse events. 4, 5
- AP-HP** Assistance Publique - Hôpitaux de Paris. 123
- ASCO** American Society of Clinical Oncology. 127
- ATD** Accelerated Titration Design. 16
- BCD** Biased Coin Design. 16
- bCRM** bivariate CRM. 44
- BLRM** Bayesian Logistic Regression Model. 25
- BMA-CRM** Bayesian Model Averaging CRM. 24
- BOIN** Bayesian Optimal INterval design. 28, 30, 36, 50, 52, 59
- BOIN-ET** BOIN design for Efficacy and Toxicity. 50, 51, 53, 59
- BOIN12** BOIN phase I/II design. 52, 53, 60
- cGUD** cumulative Group Up-and-Down design. 17, 18
- CR-CRM** Competing Risk-CRM. 92
- CRAN** Comprehensive R Archive Network. 128
- CRM** Continual Reassessment Method. 5, 13, 20, 22–27, 30, 34, 35, 40, 44, 46, 64, 115, 116
- CRML** CRM Likelihood. 23
- CTCAE** Common Terminology Criteria for Adverse Events. 5
- DA** Data Augmentation. 58
- DA-CRM** Data Augmentation CRM. 35
- DLT** Dose-Limiting Toxicity. 5, 7–9, 15, 18, 20, 22, 25–31, 33–36, 52, 63–65, 67, 84, 87, 88, 94, 123–126
- EM-CRM** Expectation-Maximisation CRM. 35
- EMA** European Medicines Agency. 127
- EWOC** Escalation With Overdose Control. 24, 35
- fCRM** fractional CRM. 35
- FDA** Food and Drug Administration. 5, 125, 127
- GUD** Group Up-and-Down design. 16–18
- HER2** Human Epidermal Growth Factor Receptor-2. 6

-
- iSurv-CRM** informative Survival-CRM. 9, 65, 67, 123, 124
- JUPM** Joint Unit Probability Mass. 49, 54
- KM** Kaplan-Meier. 27, 35, 63, 65
- LO-EffTox** Late-Onset EffTox design. 58
- mAbs** monoclonal Antibodies. 6
- MDD** Most Desirable Dose. 8, 38, 43, 49, 51, 52, 58, 60
- MSD** Most Successful Dose. 8, 38, 40, 44, 45, 62, 92
- MTAs** molecularly targeted agents. 6–8, 46, 91, 125
- MTD** Maximum Tolerated Dose. 5–7, 15, 17, 18, 20, 22, 24–27, 29, 30, 33, 35, 38, 51–53, 88, 112
- mTPI** modified Toxicity Probability Interval. 28–30, 36, 49
- OD** Optimal Dose. 8, 38, 45, 50, 51, 53, 56, 59, 93
- PCS** Probability of correct dose selection. 60, 61, 114–117
- PD** Pharmacodynamics. 2, 125
- PFS** progression-free survival. 91
- PK** Pharmacokinetics. 2, 125
- POS** Probability of Overdose Selection. 64
- RECIST** Response Evaluation Criteria in Solid Tumors. 91
- RP2D** Recommended Phase II Dose. 2, 123
- STEIN** Simple Toxicity and Efficacy Interval design. 49–51
- Surv-CRM** Survival-CRM. 9, 64, 65, 67, 87, 88, 123
- Surv-CRM-12** Survival-CRM-12. 9, 93–95, 114–117, 120, 123
- TEPI** Toxicity and Efficacy Probability Interval. 49
- TiTE-BOIN** Time-To-Event BOIN design. 36
- TiTE-BOIN-ET** Time-To-Event BOIN design for Efficacy and Toxicity. 59
- TiTE-BOIN12** Time-To-Event BOIN phase I/II trial design. 60
- TiTE-CRM** Time-To-Event CRM. 7, 24, 34–36, 55, 59, 64, 65, 87, 123
- TiTE-EWOC** Time-To-Event EWOC. 35
- TPI** Toxicity Probability Interval. 28
- TriCRM** Trivariate CRM. 45
- U-BOIN** Utility-based BOIN design. 51–53

UaD Up-and-Down. 14, 18

UPM Unit Probability Mass. 29

Chapter 1

Introduction

“In clinical drug development, a delicate balance needs to be achieved among clinical dose, efficacy, and toxicity to optimize the benefit/risk ratios in patients. An ideal drug candidate would have high potency and specificity to inhibit its molecular target without off-target effect, high drug exposure in disease-targeted tissues to achieve adequate efficacy at an optimal dose (ideally at low doses), and minimal drug exposure in healthy tissues to avoid toxicity at optimal doses (even at high doses).”

Acta Pharmaceutica Sinica B (APSB), Sun et al. (2022)

In the recent article “Why 90% of clinical drug development fails and how to improve it?”, Sun et al. (2022) identify possible reasons for the high failure rate of drug candidates in phase I, II, and III clinical trials and drug approval after passing the preclinical stage. They also provide some recommendations for improving drug optimization and clinical studies to ensure successful clinical drug development. As cited above, they explain how the success of clinical drug development might depend on a delicate balance among clinical dose, efficacy in disease-targeted organs, and toxicity in normal healthy organs at the time of in vitro studies, pre-clinical animal testing, and human clinical trials. The general topic of this thesis is novel clinical trial statistical methods for the early assessment of toxicity and therapeutic effect of treatments in the specific setting of oncology.

1.1 Oncology drug development, a long, costly, and high-risk process

Drug development is a complex, expensive, time-consuming, and challenging process. Around 10–15 years of continuous research and development with billions of dollars are needed for a single cancer drug to enter the pharmaceutical market [Prasad and Mailankody, 2017; Hinkson et al. 2020]. Generally, drug development proceeds through several steps, including laboratory experimentation, animal studies, human clinical trials (phase I, II, and III), and pharmacovigilance follow-up (phase IV).

First, the drug development process begins in the laboratory, *in vitro*, where the new agent is tested on cells, tissues or organs to assess its potential for therapeutic activity. Once a molecule has demonstrated some pharmacological activity, development could continue by testing the agent *in vivo*, *i.e.* on animals. During this preclinical phase whenever it exists, the new drug is administered to laboratory animals, such as mice, dogs or monkeys, to determine the effect on the animal behavior and biological level. The purpose of these experiments is to examine the safety of the new drug on animals in order to evaluate its tolerance. These preclinical studies are also intended to better understand the Pharmacokinetics (PK)¹ and Pharmacodynamics (PD)² of the drug under study. Sometimes efficacy data, such as the ability to target certain pathways, are also of major interest. These observed preclinical data may be used to predict how humans would react to the drug, as the disease progression and the PK/PD relationship in humans could be consistent, to some extent, with animal study results. The findings help to suggest a starting dose or range of doses for subsequent administration to humans, generally defined as 1/10 of the lethal dose in 10% of pre-clinically treated mice (LD10).

Once laboratory experiments and on animals have shown encouraging results, the first administration to humans may begin within the framework of a phase I clinical trial [Association of the British Pharmaceutical Industry (ABPI), 2018]. The main objective of phase I clinical trials is to evaluate the safety of the drug and identify side effects in order to establish a safe dose and then designate the Recommended Phase II Dose (RP2D).

¹Process of the uptake of drugs by the body over a period of time, including drug absorption, distribution, metabolism and elimination characteristics

²Biochemical, and physiological effects or actions of a drug's molecular on living systems (*e.g.* nadir neutrophil or platelet count, nonhematologic toxicity, molecular correlates, imaging endpoints).

They also enable investigators to assess the feasibility and the characteristics of the new treatment in humans to set the basis for further investigations of efficacy. Phase II trials are then conducted in larger groups and are designed to assess the efficacy of the drug and confirm the safety identified in phase I. Phase III studies confirm the efficacy by prospective comparison to the current standard of care ('gold standard' treatment or to a placebo). They are usually randomized, controlled, multi-center trials with large groups of patients. Finally, phase IV trials, involve the collection, detection, evaluation, monitoring and prevention of adverse reactions to the prescribed drug in real-life conditions. Pharmacovigilance detects rare or late-onset in the general population that have not previously been detected in clinical trials. Typically, each phase of the drug approval process is treated as separate clinical trials. If the drug successfully passes through the different phases, it will usually be approved by the regulatory authority for use in the general targeted population. However, over the past decades, this traditional phase I, II, III development strategy has become less common in oncology, particularly with the arrival of targeted therapies.

1.2 Statistical challenges in early phase oncology trials

1.2.1 Managing ethical concerns

For non-life-threatening diseases, phase I trials are usually conducted on healthy volunteers. However, in oncology, given the disease severity, the risk-benefit ratio of treatments is shifted. Phase I studies are therefore conducted on patients because of the aggressiveness and possible harmfulness of treatments. Although therapeutic benefit is not the primary objective of phase I trials (as opposed to phase II and III), these early trials may represent the last remaining chance for effective treatment after failure of standard therapies for severely ill patients with advanced cancer [Association of the British Pharmaceutical Industry (ABPI), 2018]. Small sample sizes and ethical considerations are the most constraining features of dose-finding trials [Gatsonis and Greenhouse, 1992]. Indeed, phase I monotherapy trials usually enroll on average 30 participants [Roberts et al. 2004]. Designing an efficient trial with sufficient accuracy in estimates to draw convincing conclusions based on such small sample sizes represents a major statistical challenge.

Moreover, inappropriate dose selection may directly impact a patient's life and result in ethical issues and costs to the economy. On the one hand, patients should not be exposed to highly toxic doses that can cause undesirable Adverse events (AEs) or even death. This has justified the use of the Fibonacci sequence for gradually incrementing dose levels in phase I clinical trials [Rogatko et al. 2007; Omura, 2003; Omura, 2006]. On the other hand, if patients are exposed to sub-therapeutic doses, this could be equally hazardous in oncology, contrary to other diseases. Thus, ethical concerns for trial participants have dictated the dose-escalation principles in dose-finding trials: go slowly to avoid a sudden jump from no observable toxicity to a toxic or even lethal dose, but not too slow so that large numbers of patients are not treated at ineffective doses, and base all these decisions on the fewest patients. Additionally, a wrong decision following a phase I trial obviously ripples through the entire development process. In other words, a misspecified dose can either stop a potentially interesting drug or promote a rather nonactive or even hazardous one to be further evaluated in phase II. It is not rare that new agents in oncology do not progress beyond early-phase clinical trials, leading to high costs and a long drug development process [Printz, 2015]. Thus, a correct dose-finding study is of the utmost importance during the drug development process. Well-performing dose-finding designs can therefore greatly impact the drug development process [Conaway and Petroni, 2019a].

1.2.2 Paradigm of cytotoxic chemotherapies

In oncology, study designs for phase I trials were originally developed for evaluating the use of cytotoxic agents in the treatment of malignancies by directly attacking all rapidly-dividing cells within the body, which include both fast-growing tumors and non-cancerous normal cells. Pharmacological considerations have led to the assumption that the highest possible dose should be administered in order to achieve greater anti-tumoral activity. The main issue in dose-finding clinical trial for such chemotherapies was actually to balance medical benefits and unwanted side effects. Consequently, the assumption of a monotone increasing dose-efficacy relationship has long been the fundamental paradigm for dose-finding (“more is better”, in terms of efficacy) though the dose increase is limited by inherent toxicity at high dose levels (“more is worse” in terms of toxicity). Toxicity is then considered as the prerequisite for efficacy.

The primary measure of toxicity is usually assumed dichotomous, and defined by the presence or absence of a Dose-Limiting Toxicity (DLT)³ within an observation window, usually, in the first cycles of treatment. At the time of trial planning, physicians have thus to identify a subset of candidate toxicities graded from the Common Terminology Criteria for Adverse Events (CTCAE)⁴ for which their occurrence is considered as a dose limitation. What constitutes a DLT varies from trial to trial. Generally, toxicity is considered as tolerable if the toxicity is acceptable, manageable, and reversible. Thus, only AEs with grade greater than or equal to 3 mostly define DLTs.

Many conventional dose-finding designs, such as the popular ‘3 + 3’ design or the Continual Reassessment Method (CRM) [O’Quigley, Pepe, and Fisher, 1990], rely on a monotonic underlying relationship between the dose and the DLTs, which guides the dose allocation during the trial. As such, the objective of phase I trials for cytotoxic agents is to identify the dose associated with a maximum ‘acceptable’ probability of toxicity, called the Maximum Tolerated Dose (MTD). The acceptable target of toxicity has long been implicit before the development of model-based designs which rely on an explicit DLT rate, pre-specified by clinicians before trial initiation, often lying between 0.20 and 0.35 [Le Tourneau, Lee, and Siu, 2009]. Then, the phase II recommended dose is usually taken as the MTD or the dose level just below the MTD [Rogatko et al. 2007; Ratain et al. 1993]. Recently, the Food and Drug Administration (FDA) also suggested that phase I trials could identify a range of doses whose efficacy could be formally evaluated through randomized trials in phase II trials.

In practice, the study begins with the inclusion of a small cohort of patients, mostly of size 3, 2, or 1. These patients are assigned to a specific dose level, usually the lowest one, then followed for a pre-specified period of time. DLTs are commonly assessed only during the first cycle(s) of treatment, typically for 3 to 4 weeks. Patients continue to be treated over several cycles (sometimes 8 cycles, depending on death, disease progression or consent withdrawal). During the trial, the DLT information is collected from patients, and then utilized in order to assign the next cohort of patients to a dose level. Therefore, phase I trials are naturally sequential and adaptive, because (i) patients enter the study in cohorts and not all at the same time, (ii) the dose of each new cohort is adaptively chosen

³In this thesis, DLT and toxicity will be used interchangeably.

⁴CTCAE, v5.0. U.S National Cancer Institute (NCI), 2018.

based on the available information from the previous patients in the trial. This allows to impose a safety constraint in case of excessive toxicity. At the end of the trial, the MTD is identified according to the objective of the trial.

1.2.3 New classes of cancer treatment

From the 1940s, cytotoxic chemotherapies with or without surgery or radiation therapy have been the reference treatments for patients with malignant tumors. Nevertheless, advances in the understanding of cancer biology have prompted the introduction of molecularly targeted agents (MTAs) and immunotherapies in the landscape of clinical trials. Since the 2000s, these new classes of anti-cancer agents have become an increasingly available and vital treatment option in several cancer indications that have transformed standard of care practice [Postel-Vinay, Arkenau, et al. 2009; Le Tourneau, Diéras, et al. 2010; Le Tourneau, Razak, et al. 2011; Le Tourneau, Gan, et al. 2012; Couzin-Frankel, 2013].

The mode of action of these new drugs differs from that of cytotoxic chemotherapies; instead of killing tumor cells, they target specific cell-signaling pathways. Non-cytotoxic compounds are a very heterogeneous group of agents. Immunotherapies aim to empower a patient's own immune system to directly fight cancer cells [Mellman et al. 2011]. Targeted therapies in oncology are usually small molecule drugs or monoclonal Antibodies (mAbs), that target a protein or enzyme carrying a mutation or another genetic alteration specific to cancer cells, absent from healthy cells. They allow altering signaling pathways while targeting tumoral cells (cell cycle regulation, metastasis, angiogenesis, apoptosis of cancer cells and tumor growth), preserving healthy cells [Han et al. 2020]. For example, trastuzumab is a monoclonal antibody targeting the Human Epidermal Growth Factor Receptor-2 (HER2) protein [Pratt et al. 2012], approved in France for the treatment of breast and advanced oeso-gastric cancers overexpressing HER2. Due to these different pharmacological characteristics and extended periods of treatment administration, novel therapies challenged standard paradigms on which clinical trial designs have long relied for the development of cytotoxic chemotherapies.

First, targeted therapies, due to their mode of action and improved tolerance, are most often administered continuously, until tumor progression or resistance, resulting in late onset immune-related toxicities [Paoletti, Postel-Vinay, et al. 2010; Postel-Vinay, Gomez-Roca, et al. 2011; June et al. 2017]. For instance, it has been reported that 57% of grade 3-4 toxicities occur after the first cycle of treatment in 36 clinical trials of molecularly targeted agents [Colin et al. 2015]. However, dose-finding trials historically use a short toxicity observation window after the onset of therapy. With the new anti-cancer agents, different DLT definitions should be considered, including longer toxicity observation windows (e.g., over multiple treatment cycles). Therefore, in a dose-finding trial, depending on the accrual rate, the DLT information of some patients already treated in the trial may be incomplete (or right-censored) at the time a dose level is to be assigned to a new cohort of patients. Some designs, such as the Time-To-Event CRM (TiTE-CRM) [Cheung and Chappell, 2000], have been adapted to handle those late-onset toxicities [Mauguen et al. 2011; Liu, Yin, and Yuan, 2013; Yuan, Lin, et al. 2018]. These designs provide practical strategies for continuous enrollment of patients when previous enrolled patients are still under observation, allowing to shorten the duration of the trial. However, such designs usually rely on binary endpoints for toxicity and does not take into account the time to the occurrence of the DLT.

Moreover, some MTAs are no longer consistent with the assumption based on the cytotoxic agents pharmacology that the more toxic, the more efficacious the drug. Indeed, a plateau of efficacy may occur when increasing the dose [Hoering et al. 2011; Postel-Vinay, Gomez-Roca, et al. 2011]. For example, when all the receptors targeted are already binded with the treatment, an increase in dose level would result in a saturation of receptors in the body, and therefore a plateau of efficacy can be expected [Postel-Vinay, Gomez-Roca, et al. 2011; Paoletti, Le Tourneau, et al. 2014]. Thus, toxicity could no longer be considered as a surrogate for efficacy, so that the target dose of phase I trials should be revised: rather than the MTD, the target dose should also consider the efficacy of the drug. Beside benefiting patients in the trial, it would also maximize the chances of success of the treatment in the later stages of clinical development [Borden and Dowlati, 2012]. Such designs that incorporate both toxicity and efficacy outcomes in dose-finding objectives are referred as seamless phase I/II designs. They aim to identify in a single study

one or several doses that either exhibit promising efficacy with acceptable levels of toxicity, defining the ‘Optimal Dose (OD)’, achieve efficacy in the absence of toxicity defining the ‘Most Successful Dose (MSD)’, or that provide the best balance between toxicity and efficacy, defining the ‘Most Desirable Dose (MDD)’. Most existing dose-finding methods assume that the outcomes of toxicity and efficacy can be each modeled independently. However, in practice, in a clinical trial, the observation of one of the outcome may remain incomplete because of the occurrence of an event related to the other. For instance, any DLT occurrence or disease progression results in discontinuation of the treatment and thus prevents the complete observation of the other endpoint, efficacy or toxicity, respectively. Indeed, as stated above, MTAs are usually administered continuously, until DLT occurrence, disease progression, consent withdrawal or at physician discretion, rather than after a predetermined number of cycles. The main cause of premature discontinuation from a phase I trial is indeed disease progression (reported in approximately 70% of the cases) [Olmos et al. 2012]. Typically, if a patient experiences either tumor progression or DLT, he/she may receive a second-line treatment off the protocol because it is highly unethical to continue treating a seriously ill patient with an ineffective or overly toxic drug. Since a first adverse event (progression or toxicity) will preclude the happening of the others within trial definitions, competing risks issues arise. Competing risks outcomes are quite common in clinical trials. Austin and Fine (2017) recently reviewed 40 clinical trials with survival outcomes and found that 31 of them (77.5%) were potentially susceptible to competing risks. However, the issue of competing risks is not always appropriately addressed and, when ignored, can lead to generate a potential bias in the estimation of the incidence rate of the event of interest. Notably, in the setting of phase I or I/II trials, the overestimation of the incidence rate of toxicity or efficacy, could lead to a rejection of a promising drug due to toxicity, or recommendation of a less effective dose, respectively. Therefore, an adaptive design addressing those competing risks outcomes in early phases of clinical evaluation appeared of prime interest.

The handling of toxicity and efficacy in phase I/II trials impacts the definition of endpoints. Existing joint evaluation designs most often use tumor response or the value of a biomarker as an efficacy marker. For the identification of the correct dose, they use either a multinomial response outcome combining toxicity and efficacy [Thall and Russell,

1998; O’Quigley, Hughes, and Fenton, 2001; Zhang, Sargent, and Mandrekar, 2006], or the dose with maximum estimated efficacy among tolerated doses [Wages and Tait, 2015; Zang et al. 2014; Yuan and Yin, 2009; Lin and Yin, 2017; Takeda, Taguri, and Morita, 2018] and/or a "trade-off" between efficacy and toxicity [Thall and Cook, 2004; Liu and Johnson, 2016]. Finally, authors have proposed schemes addressing late-onset toxicity and efficacy outcomes, which are based on survival models for toxicity and efficacy criteria [Yuan and Yin, 2009; Koopmeiners and Modiano, 2014].

1.3 Contributions

All of these considerations have led us to propose new dose-finding trial designs handling three different situations of right-censored observations: (i) administrative censoring occurring during the trial accrual, for dose assignment computations, when on-going observations may be incomplete given only a fraction of the toxicity window has been observed and the patient has not developed any dose-limiting toxicity yet, (ii) administrative censoring concerning all the patients without dose-limiting toxicity at the end of the observation window, and (iii) informative censoring due to trial discontinuations, likely to occur in those patients with advanced disease due to lack of treatment efficacy (either death, disease progression or any other reasons).

For the first two situations, we used a survival framework modeling the main toxicity outcome using both the event indicator and the event or censoring time, then estimating event probabilities by the cumulative incidence of DLT. For that, we developed a phase I dose-finding design, namely the Survival-CRM (Surv-CRM) [Andrillon et al. 2020]. To deal with the third situation and handle treatment discontinuation due to the occurrence of a competing event in the toxicity observation window (whatever the cause), we extended the Surv-CRM with the informative Survival-CRM (iSurv-CRM) [Andrillon et al. 2020].

We then considered that treatment discontinuation due to disease progression in competition with toxicity could be an efficacy information that could be used in the choice of the correct dose. In the second part of this work, we developed the Survival-CRM-12 (Surv-CRM-12), a Bayesian adaptive phase I-II design addressing late-onset competing risks outcomes [Andrillon et al. 2022]. The Surv-CRM-12 formally includes both toxicity and efficacy information through disease progression in the dose-finding objective

and the trial algorithm. It aims to identify an optimal dose, defined as the tolerable dose optimizing the efficacy outcome, that is, minimizing the probability of disease progression.

The performances of these methods have been evaluated by simulations and compared to the performances of other existing methods and a nonparametric optimal benchmark. This optimal benchmark was specifically derived for right-censored endpoints.

1.4 Manuscript organization

Chapter 2 presents a general overview of statistical methods that have been developed for the design and analysis of phase I and phase I/II clinical trials. Chapter 3 presents the first methodological development of this PhD research work, for phase I trial designs with right censored endpoints and its associated benchmark. The second development, on phase I/II design incorporating both toxicity and progression information for dose determination, is presented in Chapter 4. Both include the corresponding original manuscript of the research work that was published. Last, some discussion, perspectives and conclusion on related topics are provided in Chapter 5.

Chapter 2

State of the art of phase I and phase I/II designs for single-agent trials in oncology

2.1	Phase I designs	14
2.1.1	Ruled-based designs	14
2.1.1.1	‘3 + 3’ design	15
2.1.1.2	Up-and-down designs	16
2.1.2	Model-based designs	20
2.1.2.1	CRM	22
2.1.2.2	Extensions of the CRM	24
2.1.3	Model-assisted designs	28
2.1.3.1	mTPI and Keyboard designs	28
2.1.3.2	BOIN design	30
2.1.4	Designs for a delayed toxicity outcome	33
2.1.4.1	Extensions of the CRM	34
2.1.4.2	TiTE-BOIN design	36
2.2	Phase I/II designs	37
2.2.1	Ruled-based designs	40
2.2.2	Model-based designs	41
2.2.2.1	Efficacy-toxicity trade-offs designs	41
2.2.2.2	Extensions of the CRM	44
2.2.3	Model assisted designs	49
2.2.4	Designs for delayed toxicity and efficacy outcomes	55
2.2.4.1	Extension of the CRM	55
2.2.4.2	Survival design	55
2.2.4.3	Efficacy-toxicity trade-offs designs	57
2.2.4.4	TiTE-BOIN designs	59
2.3	Nonparametric optimal Benchmark	60

The literature proposing new dose-finding designs is very rich. There have been a large number of dose-finding methods proposed over the years, most of them trying to refine or extend the previous ones. Nevertheless, many designs share a lot of similarities and some of them have quite similar operating characteristics. Given this multitude of proposed designs, [Clertant \(2022\)](#) addressed one question “*How should we decide among the new proposals which one is the best for our propose ?*” The goal of this chapter is to provide an overview of the most popular phase I and phase I/II trials designs for single-agent trials in oncology, and attempt to highlight their advantages and weakness.

For a long time, dose escalation models have been based on predefined and simple empirical dose escalation/de-escalation rules. In such rule-based designs, the only assumption consisted of a monotonic increasing dose-toxicity relationships, related to its first dedicated use for cytotoxic drugs (section 2.1.1). They rapidly showed their limits and inflexibility, in favor of model-based designs, such as the well-know CRM [[O’Quigley, Pepe, and Fisher, 1990](#)]. Model-based designs rely on a parametric statistical model for the dose-toxicity relationship, to be estimated and which guides the dose allocation during the trial, using all available information (section 2.1.2). Model-assisted designs have then emerged for phase I clinical trials, combining the simplicity of ruled-based designs and the use of a statistical model for decision making (section 2.1.3). Extensions of all these designs have been proposed for trials with delayed toxicity endpoint (section 2.1.4) and setting of phase I/II identifying both safe and effective treatment (section 2.2).

2.1 Phase I designs

2.1.1 Ruled-based designs

Algorithm-based designs, also known as ruled-based designs, have their origins in the Up-and-Down (UaD) procedure initially developed and used in explosives testing. During World War II, at the Explosives Research Laboratory¹, experimental investigations were done to determine the height from which ammunition with an explosive charge explodes on impact.

“For example, in testing the sensitivity of explosives to shock, a common procedure is to drop a weight on specimens of the same explosive mixture from various heights. There are heights at which some specimens will explode, and others will not, and it is assumed that those which will not explode would explode were the weight dropped from a sufficiently greater height. It is supposed, therefore, that there is a critical height associated with each specimen, and that the specimen will explode when the weight is dropped from a greater height and will not explode when the weight is dropped from a lesser height.”

Journal of the American Statistical Association, [Dixon and Mood \(1948\)](#)

Based on this research, [Dixon and Mood \(1948\)](#) established the UaD method, a procedure for obtaining and analyzing the sensitivity of dosage-mortality data. With the UaD rule, given weight i has been dropped from height j , the next weight $i + 1$ is dropped at the next lower height $j - 1$ if an explosion has been observed for weight i , otherwise at the next higher height $j + 1$. This approach concentrates on testing settings near the level which elicits an explosion in 50% of experiments. Moreover, they pointed out the fact that the difficulty of estimating a critical value from a continuous variable also arises in many fields of research: in identifying a critical insecticide dose, germicides, anesthetics, and other drugs, in measuring the relation between physical threshold stimuli and sensations and perceptions, in testing point of failure of materials, etc. Forty years later [Storer and Mood \(1989\)](#) formally introduced the UaD procedure into clinical trial settings through the development of the ‘3 + 3’ design.

¹Bruceton, Pennsylvania, USA. 1941-1945

2.1.1.1 ‘3 + 3’ design

The ‘3 + 3’ design is the most well-known up-and-down design and has long been considered as the routine design for phase I clinical trials to estimate the MTD of novel drugs in oncology [Rogatko et al. 2007; Penel et al. 2009; Le Tourneau, Lee, and Siu, 2009; Le Tourneau, Gan, et al. 2012].

In the ‘3 + 3’ design, a first cohort of 3 patients is usually treated at the lowest dose level. Then, the dose escalation is empirically based : in the absence of DLT in the cohort, dose is escalated for the next cohort; if 1 DLT is observed among the 3 patients, an additional cohort of 3 patients is treated at the same level with dose escalation only if no additional DLTs; if ≥ 2 DLTs on the 3 or 6 patients, the trial is stopped and the next lower dose level is defined as the recommended dose level. MTD is decided when 6 patients are treated at a dose level with < 2 DLTs. Figure 2.1 describes a general case of the dose escalation rules of this algorithm.

In addition to its simplicity and transparency, a major reason for its popularity is that the ‘3 + 3’ design requires neither statistical grounds nor a computer program to run the trial [Le Tourneau, Lee, and Siu, 2009]. However, the inefficient operating characteristics of the ‘3 + 3’ design greatly outweigh these benefits [Harrington et al. 2013]. Several simulation studies showed a tendency for the resulting estimators of the MTD to be biased or inconsistent, especially when there are many doses and when the MTD is in the highest doses, and for a large proportion of patients to be treated at low and potentially ineffective dose levels [O’Quigley and Chevret, 1991; Faries, 1994; Ahn, 1998; Rogatko et al. 2007; Iasonos, Wilton, et al. 2008; Le Tourneau, Lee, and Siu, 2009]. Moreover, the ‘3 + 3’ design has severe restrictions in targeting a probability of toxicity. The expected toxicity rate at the estimated MTD depends on the number of doses under experimentation and ranges from 19 to 24% (rather than the anticipated 33%) [He et al. 2006]. Finally, the ‘3 + 3’ design is a memory-less design, i.e., dose escalation decisions are only based on the last one or two observed cohort(s), and the distribution of toxicities in previous cohorts is ignored.

2.1.1.2 Up-and-down designs

Other up-and-down designs have been proposed to achieve better operating characteristics than the ‘3 + 3’ method.

Biased Coin Design (BCD) – [Durham and Flournoy \(1994\)](#) proposed to modify the Up-and-Down procedure from [Dixon and Mood \(1948\)](#) to allow targeting a dose with a pre-specified probability of toxicity $\pi_{DLT} \leq 50\%$. The BCD de-escalates the dose for the next patient if the current patient has toxicity, and escalates according to a biased coin with probability $\frac{\pi_{DLT}}{(1-\pi_{DLT})}$ if there is none.

Accelerated Titration Design (ATD) – [Simon et al. \(1997\)](#) proposed a family of accelerated titration designs to reduce the number of patients assigned to subtherapeutic dose levels in the ‘3 + 3’ design, by adding an accelerated dose assignment phase in which one patient is treated per dose level, before moving to a ‘3 + 3’ design phase. ATD allows for intra-patient dose escalation, rendering the interpretation of the results difficult as cumulative or delayed toxicities may be masked.

‘A + B’ design – [Lin and Shih \(2001\)](#) generalized the ‘3 + 3’ design by introducing ‘A + B’ design, allowing to modulate the target toxicity rate according to chosen design parameters, as described by [Ivanova \(2006\)](#). The ‘A + B’ algorithm is summarized in Figure 2.1; when a new dose is introduced, a cohort of A patients are treated at a dose level, and if further observations are required on the same dose, a cohort of B further patients are then treated at the same dose level.

Group Up-and-Down design (GUD) – [Gezmu and Flournoy \(2006\)](#) developed a design in which the patients are also treated in cohorts. Parameters can be chosen in order to centre assignments around the dose with a specified toxicity rate π_{DLT} . The algorithm can be described as follows: When the last enrolled cohort of s patients was treated at dose level j and y of them experienced toxicity, the dose for the next cohort:

- is decreased to dose level $j - 1$ if $y \geq c_U$;
- is increased to dose level $j + 1$ if $y \leq c_L$;
- is repeated to dose level j if $c_L < y < c_U$

with c_L and c_U the lower and higher integers cutoffs, respectively, chosen by the physician. $0 \leq c_L < c_U \leq s$.

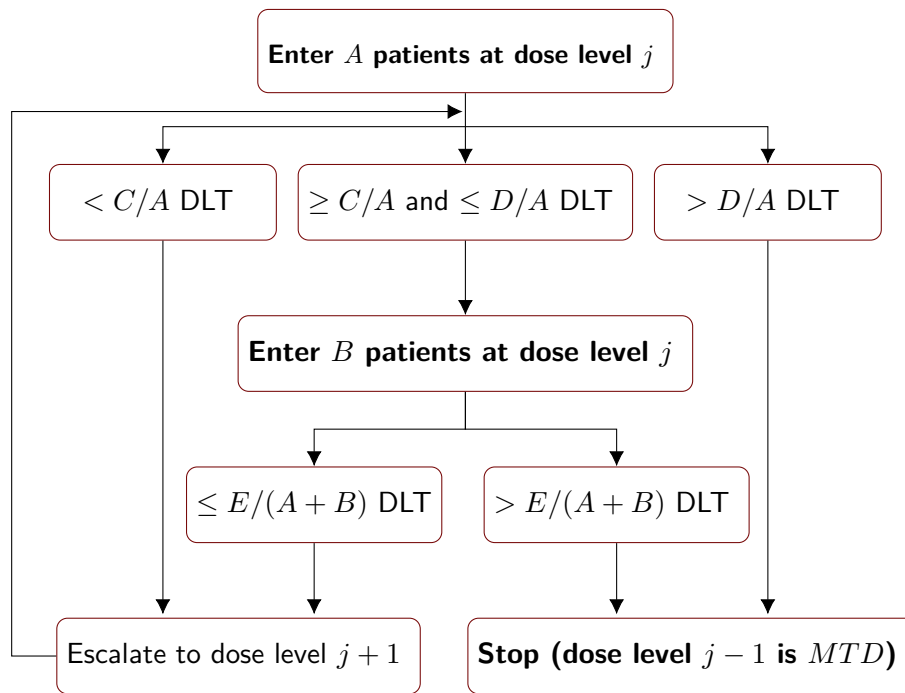


Figure 2.1: ‘A + B’ scheme without de-escalation. Of note: the ‘3 + 3’ is a special case of ‘A + B’ with $A = B = 3$; $E = C = D = 1$

cumulative Group Up-and-Down design (cGUD) – Later, [Ivanova et al. \(2007\)](#) proposed a cumulative GUD using the cumulative toxicity information at the current treating dose to make dose assignment. For dose assignment, the cGUD design uses the same principles as the GUD but considers the cumulative number of subjects treated at the current dose as the cohort size. More precisely, let s^* be the number of patients that have been assigned to the current dose level j , and assume that y^* out of the s^* patients experienced toxicity. The decision rule is then expressed in terms of estimated observed probability of toxicity ($\hat{p} = y^*/s^*$), lying within some chosen intervals:

- if $\hat{p} \geq \pi_{DLT} + \delta_U$, de-escalate to dose level $j - 1$;
- if $\hat{p} \leq \pi_{DLT} - \delta_L$, escalate to dose level $j + 1$;
- if $\pi_{DLT} - \delta_L < \hat{p} < \pi_{DLT} + \delta_U$, stay at the current dose level j

where $\delta_L = \pi_{DLT} - c_L/s^*$ and $\delta_U = c_U/s^* - \pi_{DLT}$. Similarly to the GUD design, the goal of the cGUD is to identify the MTD, which is defined as the dose associated with a pre-specified ‘target’ toxicity rate.

General remarks – Simulations comparing the operating characteristics of the presented algorithm-based designs have shown that the GUD and cGUD designs yield the best performances in selecting the MTD and assigning to the MTD during the trial. The cGUD design better protects enrolled patients during the trial [Oron et al. 2011; Liu, Cai, and Ning, 2013]. Although most of these UaD designs provide better operating characteristics than the ‘3 + 3’ design, they do not perform satisfactorily enough. It is largely recognized by statisticians that ruled-based designs, especially the standard ‘3 + 3’, are flawed and, in some sense, inadequate [Ahn, 1998; Iasonos, Wilton, et al. 2008; Le Tourneau, Lee, and Siu, 2009; Paoletti, Ezzalfani, and Le Tourneau, 2015]. Indeed, in algorithm-based designs, the dose assignment process is only based on the application of pre-defined rules to empirical counts of DLT/non-DLT responses of the last included cohort or the current treating dose. In this sense, rule-based approaches treat the MTD as being observed from a part of the data. By contrast, model-based methods presented in the next section 2.1.2 consider the MTD as a parameter to be estimated from all the data accumulated along the trial using a model, which makes sense from a statistical perspective.

Although the ‘3 + 3’ is defective, it remains one of the most widely employed methods, including for drugs whose assumptions are far from those of the 3+3, and despite many alternative model-based designs have been proposed. Indeed, methodology relating to innovative designs for cancer Phase I trials still failed to translate easily into practice. Rogatko et al. (2007) reviewed phase I cancer trials published from 1991 to 2006 and reported that “An overwhelming 98.4% of the clinical trials (1,215 of 1,235 trials) followed variations of the standard up-and-down method.” Le Tourneau, Lee, and Siu (2009) performed a similar review of 181 phase I cancer clinical trials published between 2007 and 2008 and concluded that 96.7% of dose-finding trials used rule-based algorithms. More recently, Conaway and Petroni (2019) reviewed 37 articles published in *Clinical Cancer Research* in 2018 and noted 86.5% used the ‘3 + 3’ algorithm even for trials that present innovative and complex issues. Such an important and persistent use of those designs, shown to bear many flaws, is likely related to its implementation without involving statisticians and computer programs.

A summary of previously presented phase I ruled-based designs is provided in Table 2.1.

Table 2.1: Glossary of existing ruled-based designs for phase I trials

	Advantages	Limitations	
Ruled-based designs	Easy to understand and implement	Memory-less (use only the outcome of the most recently treated patients or of the current dose)	
	Objective	Features	Available software
UaD	Median effective dose ($\pi_{DLT} = 50\%$)	1 st developed procedure to analyze the sensitivity of dosage data	None listed
'3 + 3'	Dose with a toxicity rate between 19 to 24%	Biased and inconsistent MTD estimator; patients treated at subtherapeutic doses; not flexible (target rate of toxicity, cohort size, order of doses)	R packages 'bcrm', 'escalation', 'crmPack'
BCD	Dose with a toxicity rate closest to a pre-specified rate $\pi_{DLT} \leq 50\%$	Contrary to '3 + 3', target any quantile $\pi_{DLT} \leq 0.5$	None listed
ATD	Dose with a toxicity rate $\pi_{DLT} = 25\%$	Reduce the number of undertreated patients treated in the '3+3' desig; not always suitable	Microsoft Excel macro and S-PLUS brb.nci.nih.gov
'A + B'	Dose with a DLT probability closest to a toxicity target (MTD)	More flexible than the '3 + 3'. May shorten the duration of trial	R shiny io/AplusB and R code graham-wheeler/AplusB
cGUD	Dose with a DLT probability closest to a toxicity target (MTD)	Can target any π_{DLT} ; use all information available at the current dose to make dose assignment; best-performing ruled-based design	None listed

2.1.2 Model-based designs

Over the past three decades, considerable statistical research has been conducted on early phase designs, while taking into account the ethical constraints underlying the conduct of such trials. Specifically, there has been widespread interest in model-based designs, employing statistical models to estimate the underlying relationship between the dose given to a patient and the probability of DLT [Rosenberger and Haines, 2002]. By specifying assumptions on the dose-toxicity relationship (other than a monotonically increasing relationship), strength and information is borrowed across dose levels. In the setting of small samples observed in phase I/II trials, such a parametric approach allows an interesting additional amount of information. The general idea is to base the dose-assignment on the optimization of a function depending on the current estimates and the targeted probability (typically some distance between the estimated probabilities of toxicity and the targeted probability of toxicity). Inference on the parameter(s) of the model can be either Bayesian [O’Quigley, Pepe, and Fisher, 1990; Whitehead and Brunier, 1995] or frequentist by maximum likelihood [O’Quigley and Shen, 1996]. Due to small sample sizes in phase I clinical trials, Bayesian methods are widely used, improving statistical properties such as convergence to the true MTD [Cheung, 2005]. A major advantage of these model-based methods is to be able to quantify probabilities of toxicity at each dose levels and provide the degree of confidence we may have in these estimates. Furthermore, model-based methods allow more flexibility in choosing the targeted probability of toxicity by specifically defining the MTD as a percentile of the dose-toxicity relationship. They can also be easily extended to handle more complex dose-finding setting such as delayed outcomes and/or phase I/II trials².

More specifically, model-based designs have expanded in oncology since the publication of the Continual Reassessment Method (CRM) by O’Quigley, Pepe, and Fisher (1990). The CRM was the first model-based design proposed and is still of great importance, including for the present work.

²See sections 2.1.4.1, 2.2.2, and 2.2.4.1.

TRIAL PLANIFICATION**1. Setting the clinical parameters**

- DLT definition and target toxicity probability π_{DLT}
- Dose levels tested $\mathbb{D} = \{d_1, \dots, d_J\}$ and initial dose level
- Sample size n and cohort size

2. Calibrating the model

- Functional form of the dose–toxicity model $F(x_j, \beta)$
- Skeleton $\{p_{0j}\}_{j=1}^J$ and scaled doses $\{x_j\}_{j=1}^J$ via backward substitution
- Prior distribution $\phi(\beta)$ of β

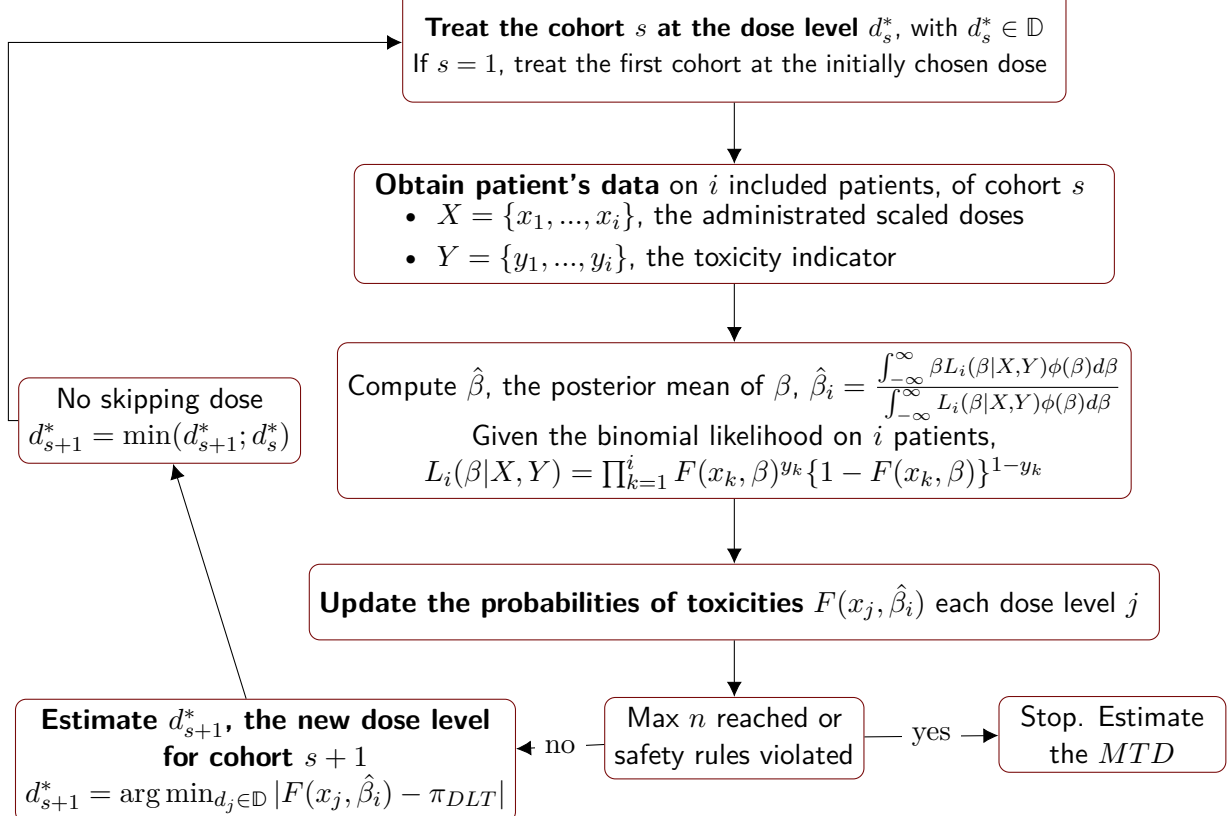
TRIAL EXECUTION

Figure 2.2: Bayesian CRM

2.1.2.1 CRM

The CRM aims to estimate the dose associated with some ‘acceptable’ targeted toxicity level, that is the dose level with estimated DLT probability closest to the pre-specified targeted probability of toxicity at trial completion. As shown in a number of simulation studies, the CRM provides faster dose escalation and improves convergence to the true MTD compared to the ‘3 + 3’ design [O’Quigley, Pepe, and Fisher, 1990; Iasonos, Wilton, et al. 2008; Le Tourneau, Lee, and Siu, 2009].

Several extensions of the CRM have been proposed, but the design was originally developed in a Bayesian framework with a one-parameter model. During the trial, patients or cohort of patients are included sequentially and, using Bayesian inference, the dose response model is updated, providing the estimated probabilities of a DLT to guide the dose assignment. The CRM allows to concentrate experimentation at the dose level for which all currently available data indicate that it is the best estimate of the correct dose level. The goal is not to perfectly estimate the dose toxicity curve but to estimate the location of the MTD as fast and efficiently as possible. The use of Bayesian inference allows a continuous learning process; as long as there are data they can be synthesized in the posterior inference. The algorithm of a trial using a CRM design is summarized in Figure 2.2 with some practical modifications compared to the first version of the CRM: no dose skipping in escalation [Goodman et al. 1995; Heyd and Carlin, 1999], the possibility of cohort accrual [Faries, 1994; Goodman et al. 1995], the use of stopping rules [O’Quigley and Reiner, 1998; Zohar and Chevret, 2001].

In the CRM design, simple models assuming monotonicity in the dose-toxicity relationship are used, e.g., a power model (empiric model) or a logistic model [O’Quigley, Pepe, and Fisher, 1990; Cheung, 2011; Chevret, 2006]. For instance, a logistic working model is defined as follows:

$$F(x_j, \beta) = \frac{\exp(\alpha_0 + \beta x_j)}{1 + \exp(\alpha_0 + \beta x_j)} \quad (2.1)$$

where $F(x_j, \beta)$ is the probability of toxicity, α_0 is a fixed constant and β is the parameter of interest to be estimated by Bayesian Inference to define the dose-toxicity curve (cf. Figure 2.3). $\{x_j\}_{j=1}^J$, are scaled dose levels, obtained by backward substitution from the

initial guesses of toxicity probabilities (skeleton) $\{p_{0j}^J\}_{j=1}^J$ of the dose-toxicity relationship to be defined before patient accrual. In the Bayesian framework, the model parameter β is assumed random and follows usually a normal prior distribution $\phi(\beta)$, that is,

$$\beta \sim N(\beta_0, \sigma_\beta^2)$$

where β_0 and σ_β^2 are, respectively, the prior mean and variance.

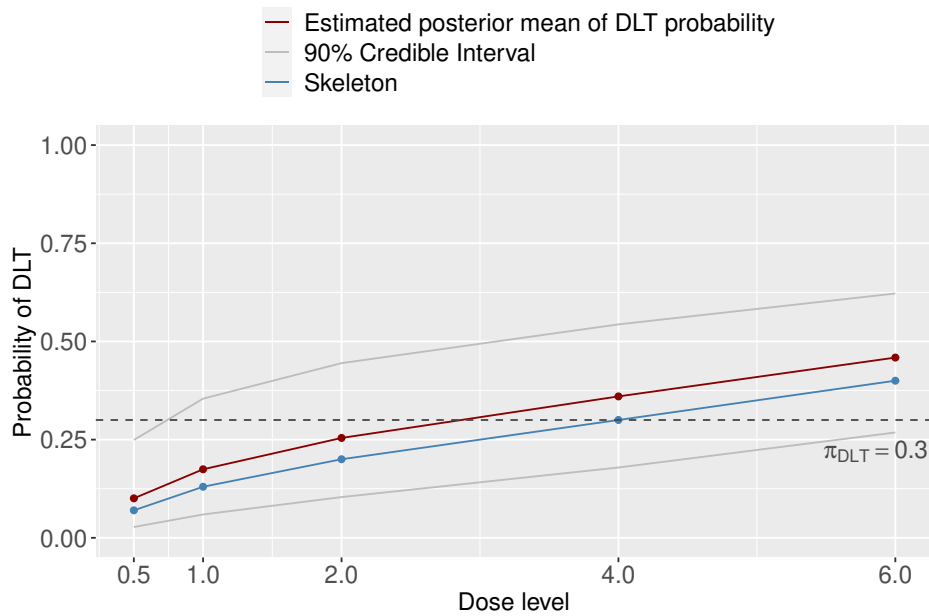


Figure 2.3: Example of final posterior mean estimates of DLT probabilities and 90% credible intervals (5th and 95th percentiles) for an hypothetical trial with the CRM, with five candidate dose levels, target toxicity probability $\pi_{DLT} = 30\%$, one-parameter logistic model as defined in equation (2.1) with intercept $\alpha_0 = 3$, and Normal prior distribution on β with standard deviation $\sigma_\beta = \sqrt{1.34}$.

General remarks – Some of the limitations of the CRM have been discussed [Shen and O’Quigley, 1996; Cheung, 2005]. First, its behavior is highly dependent on the choice of the prior distribution. In this case of phase I trials where sample size usually does not exceed 30 patients, its selection can impact the estimation of the posterior probabilities. In an attempt to address this concern, a non-Bayesian version of the CRM was also proposed, using maximum likelihood estimation to update the model parameter, called the CRM Likelihood (CRML) [O’Quigley and Shen, 1996]. The CRML requires an initial Up-

and-Down scheme before switching to the CRM. Both frequentist and Bayesian versions have similar performances regarding the recommended dose level at the end of the trial [Shen and O’Quigley, 1996; Paoletti and Kramar, 2009]. Second, asymptotic convergence can also be affected by model choices [Paoletti and Kramar, 2009]. Additionally, the specification of numerical values of the skeleton, $\{p_{0j}\}_{j=1}^J$, also impact design performances. Yin and Yuan (2009) attempted to address this issue by proposing the Bayesian Model Averaging CRM (BMA-CRM), which improves the robustness of the CRM by providing a formal basis for specifying several possible skeletons. For each adaptive decision, the BMA-CRM estimates the posterior probability for toxicity by averaging posterior probabilities estimated from the different specified skeletons. Lee and Cheung (2011) proposed a systematic calibration process to calibrate both the toxicity skeleton and the prior variance for the working model parameter, reducing the number of statistical design parameters to be specified, and making the CRM design more accessible and easier to implement. Another limitation of the CRM is that it requires a complete follow-up of the patients before carrying out a new dose determination. Cheung and Chappell (2000) proposed the TiTE-CRM to handle this challenge (see section 2.1.4.1). Finally, several extensions of the CRM for single-agent, changing the MTD definition, have also been proposed and some of them are presented in following sections 2.1.2.2 and 2.2.2.2.

2.1.2.2 Extensions of the CRM

Escalation With Overdose Control (EWOC) – The EWOC design was developed by Babb et al. (1998) to improve safety during the trial with the CRM by reducing the risk of assigning toxic doses. It is similar to the CRM, but using a two-parameters model and with an additional ethical constraint of controlling the probability of overdosing. At each new inclusion $i + 1$, the expected probability of overdosing the $i + 1^{th}$ patient with the dose d_j , given the current data, should not exceed a pre-determined value (the feasibility bound α). Thus, the EWOC uses a different definition of the MTD during and at the end of the trial, to control the risk of overdosing and assigns the $i + 1^{th}$ inclusion to the highest dose whose posterior probability of being greater than the MTD d^* is equal to or less than α , *i.e.*, $Pr(d_j > d^* | data) \leq \alpha$, with the recommended value of $\alpha = 0.25$. Of note, the EWOC approach with $\alpha = 0.5$ is equivalent to a CRM model using the posterior median

of the MTD distribution to select the next dose level [Carlin and Louis, 2009].

Bayesian Logistic Regression Model (BLRM) – Neuenschwander et al. (2009) developed a method with the objective to target a predefined toxicity interval for DLT rate instead of a single point toxicity level π_{DLT} . The BLRM used the following two-parameter logistic model to model the dose-toxicity relationship:

$$\text{logit}(p_j) = \log(\alpha) + \beta \log\left(\frac{d_j}{d^*}\right), \alpha, \beta > 0, j = 1, \dots, J \quad (2.2)$$

where p_j is the DLT rate at dose level j , and (α, β) are the unknown parameters to estimate, with a normal bivariate distribution for the prior of $\log(\alpha)$ and $\log(\beta)$. The dose levels are standardized by a reference dose level d^* , which allows for an intuitive interpretation of the model: the parameter α is the odd of a DLT for $d_j = d^*$, and the parameter β is the increase in the log odds of a DLT by a unit increase in log dose. The BLRM also uses a different definition of the MTD, defined as the dose that has the highest posterior probability of being within the target dosing interval denoted by $[\delta_1, \delta_2]$. This means that any dose level with the DLT probability within that interval $[\delta_1, \delta_2]$ can be accepted approximately as the MTD. In addition, the BLRM also imposes an overdose control rule, defined as $Pr(p_j > \delta_2 | \text{data}) \leq \alpha$, with p_j the probability of DLT at dose level j and α recommended equal to 0.25.

General remarks – Clertant (2022) and Iasonos, Wages, et al. (2016) noticed that designs with two-parameter models can generate rigidity in the dose-finding process and violation of coherence and adaptability properties. Indeed, two-parameter models may result in the method getting stuck at some level with no possibility of leaving it. Consequently, for cohort of 1 inclusions, a DLT (non-DLT) can be followed by a recommendation to escalate (de-escalate) dose level, resulting in coherence violation and a string of DLTs (non-DLTs) will not ultimately result in a recommendation to de-escalate (escalate) the level whenever such a level is available, resulting in adaptability violation.

Other extensions of the CRM, modifying the MTD definition, have also been developed. For instance, Lee and Cheung (2011) proposed to specify multiple target toxicity

thresholds for different toxicity types and grades. Typically, the MTD can be defined as the dose that satisfies both a target toxicity rate of 33% for grade 2 or higher and a target DLT rate of 20% for grade 3 or higher of toxicity. Lee, Lu, and Cheng (2020) also revised the CRM and included the patient-reported outcome in the definition of the toxicity endpoint. Instead of only capturing the DLT outcome reported by the clinicians only, the DLT outcomes from patients are also considered, resulting in two different toxicity targets for each endpoint.

A summary of presented phase I model-based designs is provided in Table 2.2. Of note, phase I model-based designs for a delayed toxicity outcome are also tabulated but will be further detailed in section 2.1.4.1.

Table 2.2: Glossary of existing model-based designs for phase I trials

	Advantages		Limitations
Model-based designs	Accurate estimation of MTD; properties of coherence and convergence; flexible (target toxicity rate, initial dose, cohort size); make full use of all information collected.		Importance of the choice of the prior and parametric model
	Objective	Features	Available software
CRM	<i>MTD*</i>	May overestimate dose for MTD	R packages ‘dfcrm’, ‘bcm’, ‘crmPack’, ‘trialr’, ‘escalation’. R interface ‘GUIP1’. R shiny trialdesign.org , uvadcos.io/crmb
EWOC	Highest dose whose posterior probability of being greater than the MTD is equal to or less than a pre-specified threshold	Keep the expected proportion of overdosed patients below a fixed boundary; could be too conservative	R interface ‘GUIP1’, R packages ‘EWOC’, ‘bcm’
BLRM	Dose with the highest posterior probability of being within a target interval	Use a two-parameter logistic model to better model the dose toxicity curve	R package ‘blrm’, ‘MoDesT’ Shiny App
TiTE-CRM	<i>MTD*</i>	Allow a fast and continuous patient accrual using weighting scheme	R package ‘dfcrm’, R interface ‘GUIP1’
TiTE-EWOC	<i>MTD*</i>	Allow a fast and continuous patient accrual using weighting scheme	R interface ‘GUIP1’
fCRM	<i>MTD*</i>	Allow a fast and continuous patient accrual by fractionizing the toxicity outcome using the KM estimator; safety deteriorated contrary to the CRM; require a preliminary ruled-based stage	R shiny clinicaltrialdesignapp , R code yangzhao98/fCRM
BMA-CRM	<i>MTD*</i>	More robust to prior misspecification by averaging on different specified skeletons	Windows software mdanderson.org , R shiny trialdesign.org
EM-CRM	<i>MTD*</i>	Allow a fast and continuous patient accrual considering pending toxicity outcome as missing data; require complex calculation; not robust with small sample size	R interface ‘GUIP1’
DA-CRM	<i>MTD*</i>	Allow a fast and continuous patient accrual considering pending toxicity outcome as missing data	Windows software mdanderson.org

* MTD is defined as the dose with a posterior mean estimate of the DLT probability closest to a toxicity target

2.1.3 Model-assisted designs

A third class of designs, known as model-assisted designs, has been developed by combining a rule-based approach and a model-based approach [Zhou, Murray, et al. 2018]. Model-assisted designs are often described as ‘transparent’ because, similarly to algorithm-based designs, dosing decisions can be tabulated before the trial begins based on the number of evaluable patients treated at current dose and the number of patients with DLT. In general, to guide dose assignment during the trial, model-assisted designs model only ‘local’ toxicity data observed at the current dose, typically using a beta-binomial model. In this respect, this approach differs from model-based designs, in which the dose–toxicity relationship is modeled across all doses levels and continuously updated during the trial.

This class of designs encompasses the toxicity probability interval design [Ji, Li, and Bekele, 2007], the modified Toxicity Probability Interval (mTPI) design [Ji, Liu, et al. 2010] and its variation mTPI-2 [Guo, Wang, et al. 2017], the Bayesian Optimal INterval design (BOIN) [Yuan, Hess, et al. 2016], and the Keyboard design [Yan, Mandrekar, and Yuan, 2017].

2.1.3.1 mTPI and Keyboard designs

mTPI design – The first interval-based design was proposed by Ji, Li, and Bekele (2007) with the Toxicity Probability Interval (TPI) design and a few years later, Ji, Liu, et al. (2010) proposed the modified TPI design (mTPI), an improved version of the TPI method, in terms of safety and number of model parameters to calibrate. The mTPI dosing decision is guided by a set of decision rules based on TPI, and uses a simple Bayesian model to infer from observed toxicity data. It requires the investigator to partition the probability of toxicity into three intervals: the underdosing interval $[0, \delta_1]$, the target dosing interval $[\delta_1, \delta_2]$, and the overdosing interval $[\delta_2, 1]$. The model assumes that p_j , the DLT probability of the current dose j , follows a priori a Beta(1, 1) distribution and a posteriori, a Beta($y_j + 1, n_j - y_j + 1$) distribution with y_j the number of patients experiencing DLT at the current dose and n_j the total number of patients treated at the current dose.

Given the observed data at the current dose j , the mTPI makes the dose assignment decision based on the Unit Probability Mass (UPM) of the three intervals. The UPM is defined as the posterior probability that the p_j is within the interval divided by the length of the interval:

- if $UPM_1 = Pr(p_j \in [0, \delta_1] | \text{data}_j) / \delta_1$ is the highest, escalate to dose level $j + 1$
- if $UPM_2 = Pr(p_j \in [\delta_1, \delta_2] | \text{data}_j) / (\delta_2 - \delta_1)$ is the highest, stay at the current dose level j
- if $UPM_3 = Pr(p_j \in [\delta_2, 1] | \text{data}_j) / (1 - \delta_2)$ is the highest, de-escalate to dose level $j - 1$

A major limitation of using UPM to guide dose assignment is that it overly minimizes UPM_3 , as the overdose interval is generally wider than the target dosing interval, resulting in a high risk of patient overdose.

Keyboard design – To address the mTPI overdosing issue, [Yan, Mandrekar, and Yuan \(2017\)](#) developed the Keyboard design, based on a series of L equal-width dosing intervals, representing possible locations of the true DLT probability for current dose level j . Of note, the mTPI-2 design [[Guo, Wang, et al. 2017](#)] was also proposed to improve the safety of mTPI in allocating patients to less toxic doses. The mTPI-2 is essentially the same as the Keyboard design. To make the dose assignment decision, the Keyboard design identifies the interval where the true DLT probability at the current dose is most likely to be located, defined as the strongest interval, I_{max} :

$$I_{max} = \operatorname{argmax}\{Pr(p_j \in I_l | \text{data}_j), l = 1, \dots, L\}$$

If I_{max} is to the left (or right) side of the target interval, it means that the current dose level j is most likely underdosing (or overdosing); therefore, the Keyboard design escalates (or de-escalates) the dose; otherwise, it stays at the current dose level j .

At the end of the trial, the dose-toxicity curve is modeled by applying an isotonic regression for instance, allowing to ensure a monotone increasing dose-toxicity relationship [[Bril et al. 1984](#)]. Then, the mTPI and Keyboard design estimate the MTD as the dose closest to a target toxicity rate.

2.1.3.2 BOIN design

Yuan, Hess, et al. (2016) developed the BOIN, another model-assisted design, which has similarities to the cumulative group up-and-down design developed by Ivanova et al. (2007), described in 2.1.1.2. Both designs evaluate whether the current dose is associated with a DLT rate sufficiently close to the target π_{DLT} or below/above the target. Unlike the mTPI and Keyboard designs, which rely on the posterior distribution of the toxicity probabilities to decide on dose assignment, the BOIN design only compares the observed toxicity rate to a dose escalation (ϕ_L) and de-escalation (ϕ_U) optimal boundaries.

Suppose j is the current dose level, and \hat{p}_j the observed DLT rate at that dose level, the BOIN design determines the next dose as follows: if $\hat{p}_j \leq \phi_L$ ($\hat{p}_j \geq \phi_U$), then escalate (de-escalate) the dose to level $j+1$ ($j-1$); otherwise, stay at the current dose level j . The final MTD estimate is the dose level associated to an observed toxicity rate, estimated by isotonic regression for example, closest to a pre-specified toxicity target.

General remarks – Among the aforementioned model-assisted designs, the main difference is the quantity compared to target intervals. The mTPI, Keyboard, and mTPI-2 designs base their decisions on $Pr(p_j \in interval|data)$ while the BOIN designs base decision on $\frac{y_j}{n_j} \in interval$. In terms of operating characteristics, the BOIN and Keyboard designs outperform the mTPI design with greater accuracy in identifying the MTD and less patient overdose [Zhou, Murray, et al. 2018]. Furthermore, the BOIN design is more a versatile method that has been extended to late-onset toxicities³ or phase I-II trials⁴. Zhou, Yuan, and L (2018) showed that the CRM, BOIN, and Keyboard designs provide comparable operating characteristics and Horton et al. (2017) concluded that the CRM outperformed the BOIN, and Keyboard methods, followed by BOIN, then mTPI in terms of percent of correct selection for the true MTD and percent of patients treated at the true MTD. However, these trends were less pronounced as the number of candidate dose levels decreased. Clertant (2022) recently summarised statistical properties that a good design should possess and concluded that intervals designs do not systematically respect all desired criteria. Indeed, mTPI, mTPI-2, BOIN and Keyboard designs could violate the

³See section 2.1.4.2

⁴See section 2.2.4.4

property of coherence, i.e., for one-by-one inclusions, a DLT (non-DLT) can be followed by a recommendation to increase (decrease) dose level, and the property of adaptability, i.e., a string of DLTs (non-DLTs) will not ultimately result in a recommendation to decrease (increase) the level whenever such a level is available. That has been also highlighted by [Wages, Iasonos, et al. \(2020\)](#) through a simulation study. This could be explained by the fact that these interval designs base their decisions on the accumulated data only at the current doses, ignoring information at adjacent dose levels above or below the current dose. The lack of any dose-toxicity model during the trial hampers the ability to share information via extrapolation and interpolation across different dose levels, thus preventing the method to use all available information on the level the design decides to move on during the trial. Elsewhere a model is used at the end of patient inclusion at the trial completion. The most common model is based on isotonic regression, although more familiar models such as the logistic could be equally well be used.

A summary of phase I model-assisted designs just presented in this current section is provided in Table 2.3. Of note, phase I model-assisted designs for a delayed toxicity outcome are also tabulated but will be detailed in the next section 2.1.4.

Table 2.3: Glossary of existing model-assisted designs for phase I trials

	Advantages	Limitations	
Model-assisted designs	Straightforward and transparent; dosing decision can be tabulated before trial initiation	Use only data from the current dose for dosing decision; problem of coherence, adaptability and convergence; require specification of several clinician-input parameters	
	Objective	Features	Available software
mTPI	Dose assignment: Maximize the unit probability mass of intervals; final estimation: <i>MTD</i> *	High risk of overdosing patients	R package ‘ escalation ’ and R code mdanderson.org
Keyboard	Dose assignment: Depend on the probability that the posterior DLT probability belongs to an interval; final estimation: <i>MTD</i> *	Simplicity of the mTPI design with better overdose control	R packages ‘ Keyboard ’ and R shiny trialdesign.org
BOIN	Dose assignment: Compare the observed toxicity rate at the current dose with two pre-specified boundaries; final estimation: <i>MTD</i> *	Simple comparison of the observed DLT rate to optimal thresholds	R packages ‘ BOIN ’, ‘ escalation ’ and R shiny trialdesign.org
TiTE-BOIN	Idem as the BOIN design	Allow a fast and continuous patient accrual considering pending toxicity outcome as missing data	R shiny trialdesign.org
TiTE-mTPI TiTE-Keyboard	Idem as the Keyboard and mTPI	Allow a fast and continuous patient accrual using weighting scheme	R shiny trialdesign.org

* *MTD* is defined as the dose with a posterior mean estimate of the DLT probability is closest to a toxicity target

2.1.4 Designs for a delayed toxicity outcome

The aforementioned designs assume toxicity outcome is binary and quickly ascertainable to make a real-time decision of dose assignment for new patients. By the time of the next dose assignment, the toxicity outcomes of patients already enrolled in the trial have to be fully observed. However, as discussed in chapter 1, late-onset toxicities are common in phase I clinical trials for targeted therapies and immunotherapies. Such late-onset toxicities can cause major logistic difficulties for the implementation of adaptive trial designs. Moreover, this problematic also arises in setting of multiple inclusion within an observation window, which is quite common in practice. Indeed, a logistical challenge arises when the DLT outcome is not captured soon enough, relative to the accrual rate.

For example, suppose a trial where the DLT is assessed over a 28-day observation window, and patients are enrolled in the trial by cohorts of three patients. If the accrual rate is two patients per week, then on average, five new patients could be accrued while waiting to evaluate the previous three patient's outcomes. Allowing accrual of those patients results in a prolonged trial duration.

A first partial strategy consists of treating the new cohort as soon as patients arrive by choosing the 'best' dose based on the most recent complete dose-outcome data. However, patients who have not experienced toxicity by the time that information is needed for the next dose assignment may experience toxicity later in the follow-up. Thus, such a design might underestimate the current toxicity rates, which might cause mis-identification of the MTD and, even more seriously, result in treating an undesirably large number of patients at too toxic doses.

A safer approach would be to suspend patient accrual and wait until every patient's outcome in the cohort has been fully observed before performing a new dose assignment so that an adaptive statistical rule can be applied. In most settings, this approach of repeatedly interrupting accrual may result in an unfeasibly long study. In addition, frequently suspending patient accrual is impractical, wastes resources, and causes tremendous administrative inconvenience.

Several phase I designs have proposed practical strategies to address these shortcomings by allowing for continuous accrual and real-time dose assignment for new patients when previous patient's outcomes are pending. Three general approaches have been proposed in the literature for phase I trials: (i) weighting the observation with incomplete follow-up time with a 'partial credit' [Cheung and Chappell, 2000; Mauguen et al. 2011; Lin and Yuan, 2020], (ii) modeling the toxicity as time-to-event endpoint [Yin, Zheng, and Xu, 2013], and (iii) regarding unobserved toxicity as missing data and handling them using missing data methods [Yuan and Yin, 2011; Liu, Yin, and Yuan, 2013]. Of note, from the perspective of interim analysis, in which a part of the existing patients have not been completely followed when a new cohort of patients comes in, the toxicity outcome could be considered as missing data. However, theoretically speaking, the late-onset outcomes are different from missing data because the complete observation will eventually be collected and toxicity are still likely to occur in the remaining follow-up time.

2.1.4.1 Extensions of the CRM

Time-To-Event CRM (TiTE-CRM) – Cheung and Chappell (2000) provided a practical solution to the common logistical problem that toxicity might not be observed immediately, resulting in longer observation windows. With the TiTE-CRM, the trial could be completed more quickly since doses can be assigned to newly enrolled patients or cohorts without waiting for complete observations on DLT. In practice, they proposed an extension of the CRM where the binomial likelihood is penalized by the follow-up ratio of each included patient.

Using the same notation as for the CRM, let $F(x, \beta)$ be the dose-toxicity working model. Given the data accrued up to the first i patients, we can estimate β using a likelihood that weights incomplete observations, defined as:

$$L_i(\beta|X, Y, W) = \prod_{k=1}^i \{w_{k,i} F(x_k, \beta)\}^{y_{k,i}} \{1 - w_{k,i} F(x_k, \beta)\}^{1-y_{k,i}}$$

where $y_{k,i}$ is the toxicity indicator for the k^{th} patient, $w_{k,i}$ its weight, and x_k its scaled administrated dose level. Cheung and Chappell (2000) evaluated through a simulation

study the impact of different weight functions, and concluded that the linear one is suitable in most cases. That is, suppose $[0, t^*]$ be the planned observation window for the DLT and $U_{k,i}$ the available follow-up time of the k^{th} patient prior to the entry of the $(i + 1)^{th}$ patient: $w_{k,i} = \min(\frac{U_{k,i}}{t^*}, 1)$, with $w_{k,i} = 1$ when patient i experienced a DLT during the observation window $[0, t^*]$.

Time-To-Event EWOC (TiTE-EWOC) – [Mauguen et al. 2011](#) also revised the CRM design with similar weighting schemes, but using the EWOC extension, and proposed the TiTE-EWOC. Similarly to the TiTE-CRM, the aim of this design is to enable continuous recruitment and to decrease the trial duration without impairing the characteristics of the EWOC design, especially its ability to control overdose.

fractional CRM (fCRM) – [Yin, Zheng, and Xu \(2013\)](#) also proposed the fCRM to address delayed toxicity issues, by ‘fractionizing’ the toxicity outcome to a value between 0 and 1 if it is not observed yet. The fractional contribution of each patient is calculated using the nonparametric Kaplan-Meier (KM) estimator and represents the conditional probability of DLT occurrence during the remaining evaluation window given that the patient has not experienced any DLT at the time of decision making. Compared to the TiTE-CRM, the fCRM design offers desirable operating characteristics in terms of MTD selection but its safety is slightly deteriorated and the procedure requires a preliminary stage until the first DLT occurs, based on decision rules on complete observations.

Expectation-Maximisation CRM (EM-CRM) and Data Augmentation CRM (DA-CRM) – Other authors have proposed to extend the CRM for delayed toxicity but using a different approach than the weighted likelihood. [Yuan and Yin \(2011\)](#) and [Liu, Yin, and Yuan \(2013\)](#) treated the pending DLT observations as missing data, and proposed the EM-CRM and DA-CRM respectively. The EM-CRM yields good operating characteristics, but requires repeated model fitting and estimation, which can be perceived as a limitation in practice. Performance results of the DA-CRM in terms of selecting the MTD are quite similar to those of the TiTE-CRM.

2.1.4.2 TiTE-BOIN design

Yuan, Lin, et al. 2018 extended the BOIN design to the Time-To-Event BOIN design (TiTE-BOIN) design to allow for continuous accrual while toxicity data for some enrolled patients are incomplete. The TiTE-BOIN relies in the imputation of the unobserved, pending DLT data, using the follow-up time of the patient whose toxicity profile is still unavailable. Let p_j denote the DLT rate at the current dose j , y_k , denote the toxicity binary outcome indicating if the k^{th} patient experienced a DLT ($y_k = 1$) or not ($y_k = 0$). Suppose that at a certain time in the trial, a total of i patients have been included, among whom only $no_{pend} < i$ patients have completed the DLT assessment. Let O denote the set of patients whose toxicity outcome is complete. The dose assignment decision is determined by comparing the estimate of p_j , \hat{p}_j , to boundaries. \hat{p}_j is defined as

$$\hat{p}_j = \frac{\sum_{k \in O} y_k + (p_{j,obs} / (1 - p_{j,obs})) (no_{pend} - STFT)}{i}$$

with $STFT = \sum_{k \in M} t_k / t^*$, the standard total follow-up times (STFT) of patients, where t_k is the follow-up time for the pending patient k at the current dose, t^* is the DLT assessment window and M denotes the set of pending individuals. The unknown value $p_{j,obs}$ is replaced with its Bayesian posterior mean estimate $\hat{p}_{j,obs}$ based on the observed data. Assuming a beta-binomial model with the prior $p_{j,obs} \sim Beta(\alpha, \beta)$ gives $\hat{p}_{j,obs} = (\sum_{k \in O} y_k + \alpha) / (no_{obs} + \alpha + \beta)$, where no_{obs} patients have not completed their DLT assessment. After this imputation, \hat{p}_j can be compared with ϕ_L and ϕ_U to determine dose escalation/de-escalation similarly to the BOIN design (see section 2.1.3.2). In comparison to the TiTE-CRM, the TiTE-BOIN yields comparable performances.

Another method was recently proposed by Lin and Yuan (2020) to deal with fast accrual that is applicable to all model-assisted designs, including the Keyboard, the mTPI, and the BOIN designs. Rather than predicting the pending DLT data as in the TiTE-BOIN, Lin and Yuan (2020) use a weighted likelihood to account for both observed and pending DLT, close to the TiTE-CRM approach.

2.2 Phase I/II designs

Several authors have proposed methods for dose-finding accounting for both toxicity and efficacy. Such seamless combination of phase I and phase II clinical trials have many advantages over conducting these two phases separately. It allows speeding up the drug development process, improving the dose-finding determination by identifying an efficient drug while controlling its toxicity, and enlarging the sample size to produce more reliable estimates of toxicity and efficacy. In phase I/II, the outcome of interest could be measured differently, depending on the objective of the trial. Several definitions have been proposed and used. The toxicity outcome can be considered as bivariate, ordinal or multiple, and the efficacy can be bivariate, ordinal or continuous [Houede et al. 2010; Bekele and Shen, 2005]. A measure of the immune response can be also of interest depending on the evaluated treatment [Liu, Guo, and Yuan, 2018]. This section overviews several phase I/II designs of a monotherapy considering only bivariate toxicity-efficacy outcomes.

Modeling the outcomes of interest – Different approaches for evaluating bivariate binary efficacy and toxicity outcomes have been proposed. They are broadly categorized into three types. In the first type of approaches, the joint distribution of efficacy and toxicity is reduced in a dichotomous variable, i.e., $Y' = 1$ if success (efficacy without toxicity), and $Y' = 0$, otherwise [Kpamegan and Flournoy, 2001]. In the second approach, the joint distribution of binary efficacy and toxicity is collapsed into an ordinal trinary variable, such as no toxicity and no efficacy, no toxicity but efficacy, and toxicity (whichever efficacy is observed or not) [Thall and Russell, 1998; O'Quigley, Hughes, and Fenton, 2001; Ivanova, 2003; Zhang, Sargent, and Mandrekar, 2006]. This relies on the situations where toxicity could not be considered as acceptable at all, mostly due to its severity that precludes any interest in the drug. Finally, in the third approach, mostly used in practice, the bivariate structure of outcomes is maintained in a joint probability distribution [Braun, 2002; Thall and Cook, 2004; Yin, Li, and Ji, 2006; Jin et al. 2014; Wages and Tait, 2015; Liu and Johnson, 2016; Li et al. 2017; Takeda, Taguri, and Morita, 2018; Lin, Zhou, et al. 2020]. This strategy of modeling the outcome of interest is mostly linked to the question of scientific interest when developing and testing a novel treatment: 'What is a 'correct' dose?'

Defining the ‘correct’ dose – There is no consensus on the definition and naming of the ‘correct’ dose in phase I/II trials, as it exists for MTD in phase I designs. However, depending on the trial objectives, general strategies for choosing the ‘correct dose’ can be identified. The first and simplest approach is to use a binary measure of ‘success’, *i.e.*, to identify the dose that maximizes the success, defined as efficacy with no toxicity [Thall and Russell, 1998; Kpamegan and Flournoy, 2001; O’Quigley, Hughes, and Fenton, 2001; Ivanova, 2003; Zhang, Sargent, and Mandrekar, 2006]. It will be referred thereafter to as the ‘most successful dose’ (MSD). However, such dichotomization results in a loss of important information. On the opposite, a second approach relies on the full use of the bivariate outcome, targeting a ‘most desirable dose’ (MDD) based on an efficacy–toxicity trade-off of desirability. To quantify such a dose desirability, designs rely either on efficacy/toxicity contours [Thall and Cook, 2004; Yin, Li, and Ji, 2006; Koopmeiners and Modiano, 2014; Jin et al. 2014] or on utility scores to balance the trade-off between efficacy and toxicity [Li et al. 2017; Lin and Yin, 2017; Zhou, Lee, and Yuan, 2019; Lin, Zhou, et al. 2020]. The other approaches compromise over the two previous ones. They aim at identifying an ‘optimal dose’ (OD) defined on some criterion of toxicity and efficacy. It could be the dose that maximizes the probability of efficacy among ‘acceptable’ doses in terms of safety, defined as the doses below some threshold of maximal toxicity rate [Wages and Tait, 2015] or below the target of the MTD [Takeda, Taguri, and Morita, 2018; Takeda, Morita, and Taguri, 2020]. The OD could also be the dose that targets toxicity and efficacy thresholds [Braun, 2002]. Finally, in case of time-to-event outcomes, the criterion for the OD could rely on the minimization of area under survival curve of the time-to-efficacy while maximizing the area under survival curve of the time-to-toxicity up to a pre-specified follow-up time [Yuan and Yin, 2009].

Table 2.4 attempts to provide a glossary of the most common designs for phase I/II according to the objective of the trial, the type of design and modeling outcome, if the design handles pending data while a continuous accrual, accounts for time-to-event information in the modeling, and non monotonic toxicity and efficacy relationships.

Table 2.4: Summary of the proposed designs for phase I/II dose-finding trials, according to their main objective, design (ruled- or model-, or assisted-based) and outcome (binary, trinary, or fully bivariate)

Design name	Author(s)	Objective	Design type	No outcome levels	Pending data	Time-to-event	Monotonic dose-toxicity relationship	Monotonic dose-efficacy relationship
Trinomial ordinal outcome	Thall and Russell (1998)	MSD	Model	3	No	No	Yes	Yes
Optimizing up-and-down	Kpamegan and Flournoy (2001)	MSD	Rules	2	No	No	Yes	Yes
Repeated Sequential Probability Ratio Test	O’Quigley, Hughes, and Fenton (2001)	MSD	Model	3	No	No	Yes	Yes
bCRM	Braun (2002)	OD	Model	4	No	No	Yes	Yes
Play-the-Winner-Like	Ivanova (2003)	MSD	Rules	3	No	No	Yes	Yes
EffTox	Thall and Cook (2004)	MDD	Model	4	No	No	Yes	No
TriCRM	Zhang, Sargent, and Mandrekar (2006)	MSD	Model	3	No	No	Yes	Yes
Odds-Ratio Trade-off	Yin, Li, and Ji (2006)	MDD	Model	4	No	No	Yes	No
Survival joint model	Yuan and Yin (2009)	OD	Model		Yes	Yes	Yes	Yes
Survival EffTox	Koopmeiners and Modiano (2014)	MDD	Model	4	Yes	Yes	Yes	No
LO-EffTox	Jin et al. (2014)	MDD	Model	4	Yes	No	Yes	No
Phase I/II CRM	Wages and Tait (2015)	OD	Model	4	No	No	Yes	No
Robust Bayesian EffTox	Liu and Johnson (2016)	MDD	Model	4	Yes	No	Yes	Yes
TEPI	Li et al. (2017)	MDD	Assisted	4	No	No	Yes	No
STEIN	Lin and Yin (2017)	MDD	Assisted	4	No	No	Yes	No
BOIN-ET	Takeda, Taguri, and Morita (2018)	OD	Assisted	4	No	No	Yes	No
U-BOIN	Zhou, Lee, and Yuan (2019)	MDD	Assisted	4	No	No	Yes	No
TiTE - phase I/II CRM	Yan, Tait, et al. (2019)	OD	Model	4	No	No	Yes	No
TiTE-BOIN-ET	Takeda, Morita, and Taguri (2020)	OD	Assisted	4	Yes	No	Yes	No
BOIN12	Lin, Zhou, et al. (2020)	MDD	Assisted	4	No	No	Yes	No
TiTE-BOIN12	Zhou, Lin, and Lee (2022)	MDD	Assisted	4	Yes	No	Yes	No

Objective: Most successful dose (MSD); Most desirable dose (MDD); Optimal dose (OD). **Design type:** Ruled-based design (Rules); Model-based design (Model); Model-assisted design (Assisted). **Number of (No) outcome levels:** Dichotomous variable (2); Ordinal trinary variable (3); Four levels variable (4).

Each design tabulated in Table 2.4 is detailed in this section 2.2, which is organized as the previous section 2.1 for phase I designs. Ruled-based designs for phase I/II trials are first introduced in section 2.2.1. Then, model-based designs are presented, including designs based on efficacy–toxicity trade-offs (section 2.2.2.1) and designs based on the CRM (section 2.2.2.2). Model-assisted for toxicity and efficacy criteria are then detailed in section 2.2.3. The section closes with designs for delayed toxicity and efficacy endpoints (section 2.2.4).

2.2.1 Ruled-based designs

Optimizing up-and-down design – Kpamegan and Flournoy (2001) proposed an algorithm-based design summarizing the pair of binary toxicity and efficacy outcomes (Y_T, Y_E) into a single binary variable Y' , with $Y' = 1$, if $(Y_T, Y_E) = (0, 1)$ (efficacy without toxicity) and $Y' = 0$ otherwise. During the trial, subjects are treated in pairs, and assigned to a new dose level according only to the outcome of the previous treated pair. The most successful dose (MSD) is then defined as the dose that maximizes $Pr(Y_T = 0, Y_E = 1)$.

Play-the-winner-like design – Ivanova (2003) also proposed an algorithm-based design, considering a setting where efficacy may be of no interest when toxicity occurs, resulting in a trinomial outcome:

$$Y' = \begin{cases} 0 & \text{if } \{Y_T = 0, Y_E = 0\} \quad (\text{No efficacy and no toxicity}) \\ 1 & \text{if } \{Y_T = 0, Y_E = 1\} \quad (\text{Success: efficacy and no toxicity}) \\ 2 & \text{if } \{Y_T = 1\} \quad (\text{Toxicity}) \end{cases} \quad (2.3)$$

with Y_E and Y_T some independent measures on a binary scale (0,1) of efficacy and toxicity, respectively. This design is based on the play-the-winner rule in that it continues to treat patients at the dose level that yields success [Zelen, 1969]. When the most recent patient is treated at dose d_j , and we observed the outcome Y'_j , for the next patient:

- de-escalate the dose to d_{j-1} if $Y'_j = 2$
- remain at d_j if $Y'_j = 1$
- escalate the dose to d_{j+1} if $Y'_j = 0$

The Play-the-Winner-Like design aims to find the MSD that maximizes the success probability $Pr(Y_T = 0, Y_E = 1)$ while meeting threshold $Pr(Y_T = 0, Y_E = 1) \geq \pi_E$.

2.2.2 Model-based designs

2.2.2.1 Efficacy-toxicity trade-offs designs

Trinomial ordinal outcome design – The first efficacy–toxicity trade–off Bayesian design was proposed by [Thall and Russell \(1998\)](#). They consider a three-category outcome model, as defined in equation (2.3) and note the probabilities of these three possible outcomes for a patient assigned to dose x , by $p_l(x) = Pr(Y' = l|x)$, $l = 0, 1, 2$. A proportional odds model for the cumulative probabilities at dose x is used, as follows:

$$\begin{aligned} \text{logit}(p_1(x, B) + p_2(x, B)) &= \mu + \alpha + \beta x \\ \text{logit}(p_2(x, B)) &= \mu + \beta x \end{aligned} \quad (2.4)$$

where $B = \{\mu, x, \alpha, \beta\}$, $\alpha > 0$ and $\beta > 0$. Bayesian inference was performed on the set of parameters B .

For dose assignment, [Thall and Russell \(1998\)](#) define a set of doses as acceptable, satisfying:

$$Pr(p_1(x, B) > \underline{\pi}_1 | data) > c_1 \quad \text{and} \quad Pr(p_2(x, B) < \overline{\pi}_2 | data) > c_2 \quad (2.5)$$

where $\underline{\pi}_1$ and $\overline{\pi}_2$ are fixed standards pre-specified by the clinician (both usually chosen lying between 0.05 and 0.20), and c_1 and c_2 are fixed probability cutoffs (usually chosen at 0.90). Dose escalation/de-escalation is decided as follows. If the current dose level is unacceptably toxic (equation (2.5)), de-escalate one dose level; if the current dose has acceptable toxicity and unacceptably low success, escalate one dose level; if the current dose is acceptable, then treat the next cohort at the acceptable dose level having the largest success probability criterion $Pr(p_1(x, B) \geq \underline{\pi}_1 | data)$.

EffTox design – The method from [Thall and Russell \(1998\)](#), considered as too restrictive, was extended by [Thall and Cook \(2004\)](#) into the Efftox design to accommodate settings where both toxicity and efficacy may occur, resulting in the following 4-levels outcome:

$$Y' = \begin{cases} 0 & \text{if } \{Y_T = 0, Y_E = 0\} \\ 1 & \text{if } \{Y_T = 0, Y_E = 1\} \\ 2 & \text{if } \{Y_T = 1, Y_E = 0\} \\ 3 & \text{if } \{Y_T = 1, Y_E = 1\} \end{cases} \quad (2.6)$$

To link the toxicity and efficacy endpoints, the probability of response and toxicity is jointly modeled using a Gumbel copula model, as follows:

$$\begin{aligned} Pr(Y_E = a, Y_T = b) &= p_E^a p_T^b (1 - p_E)^{(1-a)} (1 - p_T)^{(1-b)} \\ &+ (-1)^{(a+b)} p_E p_T (1 - p_E) (1 - p_T) \frac{(\exp(\rho) - 1)}{(\exp(\rho) + 1)} \end{aligned} \quad (2.7)$$

with $(a, b) \in \{0, 1\}$, and ρ an association parameter. The marginal probabilities of toxicity, $p_T(\cdot)$ and efficacy, $p_E(\cdot)$, are modeled using the logit models:

$$\begin{aligned} \text{logit}(p_T(x, B)) &= \beta_{T,0} + \beta_{T,1}x \\ \text{logit}(p_E(x, B)) &= \beta_{E,0} + \beta_{E,1}x + \beta_{E,2}x^2 \end{aligned} \quad (2.8)$$

where $B = \{\beta_{E,0}, \beta_{E,1}, \beta_{E,2}, \beta_{T,0}, \beta_{T,1}\}$ is the model parameter vector and x the scaled dose. For efficacy, a quadratic form is used to allow a nonmonotone relationship.

In the EffTox design, the dose assignment scheme is guided through two dimensional targets toxicity-efficacy trade-off contours, in which curves are fitted to target physician-elicited targets (π_E^*, π_T^*) . On each contour, (π_E^*, π_T^*) pairs are equally desirable and embody a trade-off between the chance of obtaining efficacy and the risk of toxicity. The next assigned dose, defined as the ‘maximum desirable dose’ by the authors, is an acceptable dose associated with the contour closest to the ideal point at the lower right corner of the efficacy/toxicity probability space, that is the dose with the best compromise between the probability of toxicity and of activity. A dose is considered to be admissible if both efficacy and the toxicity requirements are satisfied, in a similar way to what is defined in equation (2.5):

$$Pr(p_E(x, B) > \underline{\pi}_E | data) > c_E \quad \text{and} \quad Pr(p_T(x, B) < \overline{\pi}_T | data) > c_T \quad (2.9)$$

with values $\underline{\pi}_E, \overline{\pi}_T, c_E, c_T$ provided by the investigators.

General remarks – Before trial initiation, the EffTox design requires a close collaboration between statisticians and clinicians for parameterization. Firstly, it requires the prior probabilities of efficacy and toxicity at each dose to be elicited from clinicians, in order to help define the prior distribution of each model parameters $\{\beta_{E,0}, \beta_{E,1}, \beta_{E,2}, \beta_{T,0}, \beta_{T,1}, \rho\}$.

These priors should be strong enough to sensibly guide initial dose assignment decisions, but weak enough to be overridden by patient outcomes when they diverge from prior beliefs. Secondly, the investigators and statisticians should construct the trade-off contour, by fitting a curve to target values elicited from clinicians. This elicitation of the set of all (π_E^*, π_T^*) is often not obvious while the performances of the EffTox strongly depend on this correct calibration of target contours. Moreover, the curve of the contours is also important; it should be steep enough to avoid the rigidity of the design at sub-optimal doses [Yuan, Nguyen, and Thall, 2016]. Finally, in practice, the parameterization of the EffTox design is a complex process because the interaction of the contours and the admissibility criteria could be tricky to manage [Brock, 2016], and an iterative process as it requires extensive preliminary simulation studies inferring on possible dose transition pathways [Brock et al. 2017].

Odds-ratio trade-off design – Yin, Li, and Ji (2006) developed a nonparametric version of the EffTox design where the joint probabilities of efficacy and toxicity are modeled nonparametrically through a Dale odds ratio approach [Dale, 1986]. Probabilities of toxicity are constrained in increasing monotone order while leaving efficacy probabilities free of constraints. Similarly to the EffTox design presented in section 2.2.2.1, the odds-ratio trade-off design aims to identify a MDD which is an acceptable dose with highest desirability. Additionally, Bayesian criteria are used to define acceptable doses, similarly to the ones defined above in equation (2.9). This design gives similar operating characteristics and computational complexity compared to the EffTox model. One main difference with the EffTox design relies on dose desirability quantification, which is based on the odds ratio equivalence contour between toxicity and efficacy. It can be seen as less subjective than the contours of the EffTox design constructed from physician-specified values. However, it could also be viewed as a limitation because it assumes a proportional odds model, which might be unstable under specific dosage configurations where the dose levels double or triple in consecutive order, for instance.

2.2.2.2 Extensions of the CRM

Repeated Sequential Probability Ratio Test Design – O’Quigley, Hughes, and Fenton (2001) proposed a design derived from the CRM, applied to HIV studies with the aim at identifying the ‘most successful dose’ (MSD), i.e. a dose level that maximizes $Pr(Y_T = 0, Y_E = 1 | data)$ over safe doses.

This design differs from the approach proposed by Thall and Cook (2004) described in section 2.2.2.1. The first difference is the choice of dose-efficacy modeling: O’Quigley, Hughes, and Fenton (2001) consider an ordinal trinary variable, as defined in equation (2.3). Then, they propose to model the probability of efficacy at each dose level conditionally on the absence of toxicity, that is $p_{E|\bar{T}}(\cdot)$. The second difference relies on the parsimonious specification of the model. Similarly to the original CRM, one-parameter models are defined to model dose-toxicity and dose-response relationships at dose level j :

$$\begin{aligned} p_{T,j}(\alpha_j, a) &= \alpha_j^a \\ p_{E|\bar{T},j}(\beta_j, b) &= \beta_j^b \end{aligned} \tag{2.10}$$

where $0 < \alpha_1 < \dots < \alpha_J < 1$, $0 < \beta_1 < \dots < \beta_J < 1$ and $a \in \mathbb{R}^+$, $b \in \mathbb{R}^+$.

Probabilities $p_{T,j}(\cdot)$ and $p_{E|\bar{T},j}(\cdot)$ are sequentially estimated before each new inclusion, allowing estimation of success probabilities at each dose level j defined as $S_j(\cdot) = p_{E|\bar{T},j}(\cdot)(1 - p_{T,j}(\cdot))$.

The dose-finding algorithm uses two stages. The first escalating stage is based on a ‘3+3’ algorithm. The second stage has two component. First, a likelihood-based CRM is used to target a low-toxicity rate level, $\bar{\pi}_T$ ($\bar{\pi}_T = 0.10$ in the paper). Second, a sequential statistical test is applied to test $H_0: S_j(\cdot) = s_0$ versus $H_1: S_j(\cdot) = s_1$, from pre-specified $s_0 < s_1$. A conclusion in favor of H_0 leads to removal of dose level j and lower doses, and to increase the toxicity target (to at most $\bar{\pi}_T = 0.30$ in the paper). A conclusion in favor of H_1 leads to declare dose level j as the MSD.

bivariate CRM (bCRM) – Braun (2002) extended the CRM to model bivariate binary efficacy and toxicity outcomes. In the bCRM, the marginal probabilities of toxicity and efficacy at dose level j , $p_{T,j}(x_j, \beta_1)$ and $p_{E,j}(x_j, \beta_2)$, respectively, are both one-

parameter models, combined into a joint copula model with correlation parameter ρ . Based on the likelihood of the bivariate distribution of (Y_T, Y_E) and the prior distribution of the parameters $B = (\beta_1, \beta_2, \rho)$, the posterior mean of parameters are computed and marginal probabilities $p_T(x_j, \hat{B})$ and $p_E(x_j, \hat{B})$ are updated. The optimal dose is then determined as the dose that minimizes the weighted Euclidian distance between estimated efficacy and toxicity probabilities and targets (π_E^*, π_t^*) :

$$d^* = \operatorname{argmin}_{d \in \mathcal{D}} \{w(p_T(x_j, \hat{B}) - \pi_t^*)^2 + (1 - w)(p_E(x_j, \hat{B}) - \pi_E^*)^2\} \quad (2.11)$$

with π_E^* close to 1, π_t^* close to 0, and w recommended equal to 0.5.

Simulations showed that the performances of the bCRM in identifying the OD depends on the steepness around that dose. The starting dose also influences the performances of the design.

Trivariate CRM (TriCRM) – Zhang, Sargent, and Mandrekar (2006) - presented the TriCRM, a simple modified version of the design proposed by [Thall and Russell \(1998\)](#), and based on the CRM, but for a trinomial ordinal outcome (cf. equation (2.3)). The relationship between dose and toxicity/efficacy is modeled using a proportional odds model given in equation (2.4). Using the notations defined in section 2.2.2.1, the sum of probabilities of each level of the trinomial outcome is equal to 1 (i.e. $p_0(x, B) + p_1(x, B) + p_2(x, B) = 1$, for any dose x and some parameter set B). For a given estimate of B , the design uses two decision functions to determine the 'admissible doses' set:

$$\delta_1(x, B) = \mathbb{1}\{p_2(x, B) < \overline{\pi_T}\} \quad (2.12)$$

and

$$\delta_2(x, B) = p_1(x, B) - wp_2(x, B) \quad (2.13)$$

the later function representing the difference between the success (efficacy and no toxicity) probability and toxicity probability if the weight $w = 1$ and the success probability if $w = 0$. The most successful dose (MSD) is then defined as the dose maximizing $\delta_2(x, B)$ subject to $\delta_1(x, B) = 1$.

Seamless phase I/II CRM – [Wages and Tait \(2015\)](#) developed a design considering the four-levels outcome defined in equation (2.6) for MTAs. As in the CRM, the dose-toxicity relationship is modeled with a power model. However, in order to account for the uncertainty surrounding the shape of the true dose-efficacy curve, [Wages and Tait \(2015\)](#) make use of a class of working models. More specifically, they formulate a total of L different efficacy skeletons including unimodal and plateau relationships. Probability of efficacy under model m at dose d_j is then modeled as follow:

$$\mathcal{G}_m(d_j, \beta_m) = x_{jm}^{\exp(\beta_m)} \quad (2.14)$$

with β_m the model parameter to be estimated, $0 < x_{1m} < \dots < x_{Jm} < 1$ the skeleton values under working model m , and J the number of candidate doses. They rely on a selection technique based on Bayesian inference. The general idea is to select the working model m^* that best support the data, *i.e.*, the working model with the greater posterior model probability.

The authors define the 'best dose' (denominated in this thesis as the 'optimal dose') as the dose level that maximizes efficacy while ensuring safety, such that:

$$d^* = \arg \max_{d_j \in \mathcal{A}_i} p_E(d_j) \quad (2.15)$$

with $p_E(d_j) = \mathcal{G}_{m^*}(d_j, \beta_{m^*})$ the efficacy probabilities at dose level j , \mathcal{A}_i the set of acceptable doses in terms of toxicity after i entered patients, defined as $\mathcal{A}_i = \{d_j : p_T(d_j) \leq \pi_{DLT}; j = 1, \dots, J\}$. Early in the trial, when a limited amount of data is collected, an adaptive randomization phase is used to allocate patients to acceptable doses with randomization probabilities weighted according to the estimated efficacy probabilities. It prevents the rigidity of the method on a dose that was tried early in the trial. Additionally, safety and futility rules for early stopping of the trial are also proposed in the study design. The trial is terminated in the presence of undesirable toxicity, defined as:

$$p_T^-(d_1) > \pi_{DLT} \quad (2.16)$$

with $p_T^-(d_j)$ the lower bound of the exact binomial confidence interval for toxicity at the lowest dose d_1 , above which there is 95% confidence that the true toxicity for d_1 falls. Moreover, the trial is terminated for futility if:

$$p_E^+(d_j) < \pi_E \quad (2.17)$$

with $p_E^+(d_j)$ the upper bound of the exact binomial confidence interval for efficacy at the current dose d_j , below which there is 95% confidence that the true efficacy probability for d_j falls.

A summary of all phase I/II model-based designs including those previously presented in section 2.2.2 and those to be presented in section 2.2.4 is provided in Table 2.5. Of note, phase I/II model-based designs for a delayed toxicity and efficacy outcomes are also tabulated but will be detailed in the next section 2.2.4. The designs are presented in chronological order (from top to bottom).

Table 2.5: Glossary of existing parametric designs for phase I/II trials

	Objective	Features	Available software
Trinomial ordinal outcome	MSD: Dose having the highest success, among acceptable doses	Do not consider both toxicity and efficacy occurrence	None listed
Repeated Sequential Probability Ratio Test	MSD: Dose having the highest success among acceptable doses	Require a ‘3 + 3’ 1 st stage; do not consider both toxicity and efficacy occurrence	None listed
bCRM	OD: Minimizes Euclidian distance between efficacy and toxicity probabilities and targets	Influence of the starting dose on the performance	None listed
EffTox	MDD: Dose having the highest toxicity-efficacy trade-off among acceptable doses	Require close interaction with clinician to define trade-off contours	R packages ‘ trialr ’, ‘ escalation ’, Python package ‘ clintrials ’ and Windows software mdanderson.org
TriCRM	MSD: Dose having the highest success among acceptable doses	Consider a trinomial ordinal outcome	None listed
Odds-ratio trade-off	Idem as the EffTox (MDD)	Contrary to the EffTox, dose desirability quantification does not depend on physician-elicited values	None listed
Survival joint model	OD: dose with the higher AUSC value while satisfying acceptable toxicity requirement	Handle patients who will never respond using a cure rate model; several model parameters to estimate	None listed
Survival EffTox	Idem as the EffTox (MDD)	Handle patients who will never respond using a cure rate model; several model parameters to estimate.	
LO-EffTox	Idem as the EffTox (MDD)	Consider pending toxicity and efficacy outcomes as missing data	None listed
Phase I/II CRM	OD: Dose having the highest efficacy among safe doses	Account for the uncertainty in the dose-efficacy relationship with efficacy skeletons class	R package ‘ escalation ’ and R shiny uvadcos.io/wtdesign
Robust Bayesian EffTox	MDD: Dose having the highest utility function	No parametric models for the dose–toxicity/efficacy relationships	None listed
TiTE - Phase I/II CRM	OD: Dose having the highest efficacy among safe doses	Account for the uncertainty in the dose-efficacy relationship with efficacy skeletons class; suitable for delayed outcomes	None listed

2.2.3 Model assisted designs

Several extensions of model-assisted phase I designs have been also proposed for seamless phase I/II clinical trials.

Toxicity and Efficacy Probability Interval (TEPI) – Li et al. (2017) first extended the mTPI design [Ji, Liu, et al. 2010] presented in section 2.1.3.1 for both toxicity and efficacy endpoints. Similarly to the mTPI design, the unit intervals for the probability of toxicity (p_T) and the probability of efficacy (p_E) are partitioned into subintervals. Joint Unit Probability Mass (JUPM) for the toxicity and efficacy probability intervals is calculated for each interval combination. The TEPI design recommends dose escalation/de-escalation according to the combination with the largest JUPM value. At the end of the trial, the ‘most desirable dose’ (MDD) is selected based on a utility score to balance the toxicity and efficacy trade-offs. The utility function is defined as $U(p_T, p_E) = f_T(p_T)f_E(p_E)$, where f_T is a decreasing step function and f_E is an increasing step function, defined according to pre-specified cutoff values. The MDD is the dose that maximizes the estimated posterior expected utility.

Similarly to the mTPI design, dose assignment decision in the TEPI design is made on the basis of only the data from the current dose, resulting in a higher risk of missing the correct dose. Moreover, to use of unit probability mass results in a higher risk of overdosing patients.

Simple Toxicity and Efficacy Interval design (STEIN) – Lin and Yin (2017) proposed the STEIN based on optimized intervals for toxicity and efficacy. As illustrated in Figure 2.4, for dose assignment decision during the trial, pairs of observed toxicity and efficacy probabilities at the current dose level are compared to an efficacy threshold (ψ) and cutoffs defining an indifferent tolerance interval of the target probability (ϕ_L and ϕ_U). If the pair lies inside the green region, stay; if the pair lies inside the red region, de-escalate; otherwise, the next dose level is the dose with the highest posterior probability of efficacy among all doses in the local admissible set (*i.e.* $Pr(p_E > \psi | data)$). The admissible set A_j at dose level j is determined based on the observed toxicity probability at the current

dose j , $\hat{p}_{T,j}$, with respect to ϕ_L and ϕ_U :

$$A_j = \begin{cases} \{j-1, j, j+1\} & \text{if } \hat{p}_{T,j} \leq \phi_L \\ \{j-1, j\} & \text{if } \phi_L < \hat{p}_{T,j} < \phi_U \end{cases} \quad (2.18)$$

For final decision at the end of the trial, probabilities of toxicities and efficacy are generally estimated through isotonic regressions and the ‘most desirable dose’ is determined as the dose level that has the maximum utility weighted according to the importance of toxicity and efficacy, among the dose levels that have been tried during the trial.




Figure 2.4: Partitioned regions under the STEIN and the BOIN-ET designs.

Note: ψ is a threshold for efficacy and ϕ_L and ϕ_U are cutoffs for toxicity. $\hat{p}_{T,j}$ and $\hat{p}_{E,j}$ are observed probabilities of event at current dose j .

With the BOIN-ET, given current dose level j , $\phi_L < \hat{p}_{T,j} < \phi_U$ and $\hat{p}_{E,j} < \psi$, dose assignment is decided as follows: if dose level $j+1$ was never used to treat patients, escalate the dose to level $j+1$; if previous rule not applicable, choose the dose that has the maximum probability of efficacy among dose levels $\{j-1, j, j+1\}$; if the doses have the same maximum probabilities of efficacy, randomly choose one dose among them.

BOIN design for Efficacy and Toxicity (BOIN-ET) – Takeda, Taguri, and Morita (2018) extended the BOIN design [Yuan, Hess, et al. 2016] to a design considering both toxicity and efficacy, with the aim to identify an ‘optimal dose’ (OD) among J dose levels. The method used to guide dose escalation during the trial is similar to the BOIN⁵, by simply comparing toxicity probability $p_{T,j}$ and efficacy probability $p_{E,j}$ at dose level j with pre-specified lower (ϕ_L) and upper (ϕ_U) cutoffs for toxicity and threshold for efficacy (ψ) as illustrated in Figure 2.4.

The BOIN-ET design shares some similarities to the STEIN design described above. A first difference is the way the clinician-input parameters are calibrated: with the STEIN

⁵See section 2.1.3.2

approach, the set of cutoffs, $\{\phi_L, \phi_U, \psi\}$, is chosen by physicians, while in the BOIN-ET designs, optimal values for cutoffs are obtained by minimizing the posterior probability of incorrect decisions at each dose level j :

$$\begin{aligned} Pr(\text{incorrect} | \text{data}) &= Pr(H_{1j})Pr(S \text{ or } D|H_{1j}) + Pr(H_{2j})Pr(E \text{ or } D|H_{2j}) \\ &\quad + Pr(H_{3j}) \times 0 + Pr(H_{4j})Pr(E \text{ or } D|H_{4j}) \\ &\quad + Pr(H_{5j})Pr(E \text{ or } S|H_{5j})Pr(H_{6j})Pr(E \text{ or } S|H_{6j}) \end{aligned} \quad (2.19)$$

with ‘S’ for stay, ‘D’ for de-escalate, ‘E’ for escalate, and the following hypotheses:

$$\begin{aligned} H_{1j} &: (p_{T,j} = \underline{\pi}_T, p_{E,j} = \underline{\pi}_E), H_{2j} : (p_{T,j} = \underline{\pi}_T, p_{E,j} = \pi_{eff}), \\ H_{3j} &: (p_{T,j} = \pi_{DLT}, p_{E,j} = \underline{\pi}_E), H_{4j} : (p_{T,j} = \pi_{DLT}, p_{E,j} = \pi_{eff}), \\ H_{5j} &: (p_{T,j} = \overline{\pi}_T, p_{E,j} = \underline{\pi}_E), H_{6j} : (p_{T,j} = \overline{\pi}_T, p_{E,j} = \pi_{eff}) \end{aligned} \quad (2.20)$$

and pre-specified values by the investigator: π_{DLT} the target toxicity probability, $\pi_{E_{eff}}$ the target efficacy probability, $\underline{\pi}_T$ the highest toxicity probability that is deemed subtherapeutic such that dose escalation should be pursued, $\overline{\pi}_T$ the lowest toxicity probability that is deemed overly toxic such that dose de-escalation is needed, and $\underline{\pi}_E$ the highest efficacy probability that is deemed subtherapeutic so that another dose level should be pursued.

The second difference with the STEIN design is the optimal dose definition. Indeed, for final OD estimation with the BOIN-ET design, the efficacy probability is modeled through a logistic regression with the fractional polynomial with 2 degrees of freedom. The OD is then defined at the dose that maximizes the efficacy probability among the doses lower or equal to the MTD. The MTD is defined as the doses closest to the target toxicity probability π_{DLT} .

Utility-based BOIN design (U-BOIN) – Zhou, Lee, and Yuan (2019) developed the U-BOIN phase I/II design to find the ‘most desirable dose’ (MDD). The U-BOIN considers four possible levels for the categorical outcome Y' , as presented in equation (2.6). They assume that Y' follows a Dirichlet-multinomial model at dose level j , resulting in

$$(p_{0,j}, \dots, p_{3,j}) \sim \text{Dirichlet}(\alpha_0, \dots, \alpha_3) \quad (2.21)$$

where $p_{l,j} = Pr(Y' = l|d = j)$ is the posterior probability estimate corresponding to level l of outcome Y' at dose level j (cf. equation (2.6)), $\alpha_l > 0$, $l = 0, \dots, 3$, and $\sum_{l \in \{0,1,2,3\}} p_{l,j} = 1$.

The U-BOIN is a two stage phase I/II design. In stage I the BOIN design is implemented and in stage II both toxicity and efficacy data are used for dose assignment and a utility function U_j measures the dose risk-benefit trade-off at each dose j :

$$\hat{U}_j = \sum_{l \in \{0,1,2,3\}} u_l p_{l,j} \quad (2.22)$$

with u_l a utility score of each outcome of Y' , between 0 and 100, elicited by physicians.

Then, the MDD is defined as the dose with the highest utility value among acceptable doses, defined as the doses for which both of the following criteria are satisfied:

$$Pr(p_{E,j} > \underline{\pi}_E | data) > c_E \quad \text{and} \quad Pr(p_{T,j} < \overline{\pi}_T | data) > c_T \quad (2.23)$$

with $p_{T,j}$ the marginal toxicity probability, $p_{E,j}$ the marginal efficacy probability at dose level j and $\overline{\pi}_T$ the maximum tolerable DLT rate and $\underline{\pi}_E$ the lowest acceptable efficacy rate.

BOIN phase I/II design (BOIN12) – Another variation of the BOIN design have been proposed to handle toxicity and efficacy dose-finding objectives: the BOIN12 design. Under the BOIN12 design, patients are adaptively assigned to the ‘most desirable dose’ with the optimal toxicity-efficacy trade-off [Lin, Zhou, et al. 2020]. Similarly to the U-BOIN, the outcome Y' has 4 possible levels and physicians have to elicit a utility score for each level of Y' in order to calculate a desirability (or mean utility) over the four possible outcomes. The decision rules of the BOIN12 design are similar to the BOIN with an additional step of dose desirability calculation when the DLT rate at the current dose j is lower than a de-escalation boundary, ϕ_U . The final MDD is selected as the dose level that has the highest estimated utility among the doses that are not higher than the MTD. The MTD is defined here as the dose level that has the isotonicly estimated toxicity probability closest to a toxicity upper limit $\overline{\pi}_T$, with $\overline{\pi}_T$ slightly higher than π_{DLT} , the target toxicity rate used in conventional toxicity-based phase I designs.

General remarks – To conclude, U-BOIN and the BOIN12 designs share a lot of similarities; both use a utility function to measure dose risk-benefit trade-off, contrary to the BOIN-ET which directly works on the probabilities of toxicity and efficacy. However, there are two major differences between the U-BOIN and BOIN12. The first difference is that the U-BOIN uses a 2-stage approach (1st stage focusing on toxicity, and 2nd stage focusing on optimisation), which can be interesting when efficacy is long to be scored. In practice if we need to determine the optimal dose as a primary endpoint, use the BOIN12, while if we need to determine the optimal dose as a primary endpoint and the MTD as a secondary endpoint, use the U-BOIN. The second difference relies on the definition of the set of admissible doses to identify the OD: doses that do not exceed the estimated MTD with the BOIN12 versus non-toxic and non-futile doses with the U-BOIN. Of note, the U-BOIN requires larger sample sizes than BOIN12 for the same accuracy. Finally, in some sense, the U-BOIN and BOIN12 designs use the same benefit-risk trade-offs throughout the trial to determine dose transition and selection, contrary to the BOIN-ET where the method used to guide dose escalation is different from that used for final estimation.

Table 2.6 provides a summary of phase I/II model-assisted designs presented in this current section and those to be presented in section 2.2.4.4.

Table 2.6: Glossary of existing model-assisted designs for phase I/II trials

	Objective	Features	Available software
TEPI	Dose assignment: Maximize the joint unit probability mass for the toxicity and efficacy probability intervals; final estimation: Dose with the highest estimated posterior expected utility	Dose assignment decision is made on the basis of only the data of the current dose and on JUPM, leading to a decreased PCS and an increased risk of overdosing patients	R packages ‘keyboard’
STEIN	Dose assignment: Compare the observed toxicity rate at the current dose with three pre-specified boundaries; final estimation: Dose with the highest utility among the dose levels that have been tried during the trial	Suitable for different dose-efficacy relationships	None listed
BOIN-ET	Dose assignment: Compare the observed toxicity rate at the current dose with three optimal pre-specified boundaries; final estimation: Dose with the highest efficacy probability among the doses lower or equal to the MTD	Suitable for different dose-efficacy relationships	SAS code upon request
U-BOIN	Dose assignment and final estimation: Dose with the highest utility value among admissible doses (non toxic and non futile)	Suitable for different dose-efficacy relationships; require a 1 st stage focusing on toxicity	R shiny trialdesign.org
BOIN12	Dose assignment and final estimation: Dose with the highest estimated utility among dose than not exceed the estimated MTD	Suitable for different dose-efficacy relationships	R shiny trialdesign.org
TITE-BOIN-ET	Idem as BOIN-ET	Allow a fast and continuous patient accrual using weighting scheme	SAS code upon request
TITE-BOIN12	Idem as BOIN12	Allow a fast and continuous patient accrual considering pending toxicity and efficacy outcomes as missing data	R shiny trialdesign.org

2.2.4 Designs for delayed toxicity and efficacy outcomes

In this section, the designs presented are appropriate when both toxicity and efficacy binary endpoints cannot be observed in a reasonably time-frame⁶. In this setting, subjects are enrolled in the trial, even when only partial information had been acquired from the previous patient. Other designs considering only the efficacy as a time-to-event endpoint while the toxicity remains a binary endpoint exist [Lei et al. 2011; Guo and Yuan, 2015] but will not be discussed in this section where only designs handling both late-onset toxicity and efficacy are considered.

2.2.4.1 Extension of the CRM

Seamless phase I/II Time-to-event CRM – Yan, Tait, et al. 2019 generalize the TiTE-CRM to bivariate outcomes. More specifically, they proposed a time-to-event extension to the seamless phase I/II CRM propose by Wages and Tait (2015), detailed in section 2.2.2.2. They handle pending data due to partially observed subjects throughout the trial by a weighted likelihood, in a similar way to the TiTE-CRM [Cheung and Chappell, 2000]. Similarly to the method from Wages and Tait (2015), this design accommodates a wide range of dose-efficacy curves.

2.2.4.2 Survival design

Survival joint model for toxicity and efficacy – Yuan and Yin (2009) proposed a model adapted to a rapid accrual rate relatively to the toxicity and efficacy observation windows, which could delay treatment assignment. Therefore, in addition to allowing a continuous accrual of patients during the trial without the need to wait for complete follow-up of previously included patients, their proposed adaptive Bayesian design takes into account information about the time when the event occurred. The following survival function S_T for the time-to-toxicity (t_T) is defined:

$$S_T(t_T|x, \lambda_T, \alpha_T, \beta_T) = \exp(-\lambda_T t_T^{\alpha_T} \exp(\beta_T x)) \quad (2.24)$$

⁶See Section 2.1.4 for further explanation of the issues raised by delayed toxicity endpoints in phase I clinical trials.

with β_T the parameter for the dose variable. The hazard function for the time-to-toxicity is assumed to follow a Cox proportional hazards model, with a baseline hazard modeled by a Weibull distribution with λ_T and α_T the scale and shape parameters, respectively. Additionally, independent gamma prior distributions are assigned on λ_T , α_T , and β_T .

For the time-to-efficacy, they postulate a mixture cure rate model [Berkson and P, 2006] assuming that a certain fraction of patients could be cured. The survival function of the mixture cure rate model for the time-to-efficacy (t_E) is given by:

$$S_E^*(t_E|x, \lambda_E, \alpha_E, \beta_E) = (1 - c) + c S_E(t_E|x, \lambda_E, \alpha_E, \beta_E) \quad (2.25)$$

with c the fraction of the population that is cured, and S_E the survival function for the time-to-efficacy for the fraction of patients who possibly cured, modeled by a proportional hazards model with a Weibull baseline hazard:

$$S_E(t_E|x, \lambda_E, \alpha_E, \beta_E) = \exp(-\lambda_E t_E^{\alpha_E} \exp(\beta_E x)) \quad (2.26)$$

where β_E is the parameter for the dose variable and λ_E and α_E the scale and shape parameters, respectively. Independent gamma priors on each parameter is assumed. Furthermore, Yuan and Yin (2009) jointly modeled toxicity and activity using a Clayton copula model, such that

$$S(t_T, t_E|x, \rho, \theta) = \{S_T(t_T|x)^{1-\rho} + S_E(t_E|x)^{-1/\rho} - 1\}^{-\rho} \quad (2.27)$$

where $\rho > 0$ measures the correlation of the two outcomes, and $\theta = (\lambda_T, \alpha_T, \beta_T, \lambda_E, \alpha_E, \beta_E, \rho, c)$, the model parameter vector to be estimated by Bayesian inference.

For the dose-finding algorithm, the authors select the OD using areas under the survival curves (AUSC). More specifically, the OD is defined as the dose that maximizes the ratio of the AUSC of the times to toxicity and efficacy up to a pre-specified follow-up time, among doses satisfying safety conditions.

General remarks – This approach has the advantage to take into account both the toxicity and efficacy rates at the end of the follow-up time and how quickly patients ex-

perience toxicity or efficacy. However, a first limitation of their approach is related to the number of model parameters to be estimated (8 in total, *i.e.* $(\lambda_T, \alpha_T, \beta_T, \lambda_E, \alpha_E, \beta_E, \rho, c)$). It could be challenging given the small number of participants usually enrolled in phase I/II trials. Moreover, the design requires a preliminary stage before the formal Bayesian design, such as ‘3+3’-like design. It could extend study duration in cases where we expect to have a long follow-up period.

2.2.4.3 Efficacy-toxicity trade-offs designs

Survival EffTox design – Koopmeiners and Modiano (2014) modeled both toxicity and efficacy endpoints as time-to-event outcomes. Time-to-toxicity is modeled with a standard cure-rate model where $1 - p_T(d_j)$ is the fraction of the population that is cured and $p_T(d_j)$ is the probability of experiencing toxicity before the end of the safety assessment window at dose level d_j , defined as follows:

$$\text{logit}(p_T(d_j)) = \beta_{0,T} + \beta_{1,T}(d_j - 1) \quad (2.28)$$

Subtracting one from the dose d_1 results in $\text{logit}(p_T(d_1)) = \beta_{0,T}$, *i.e.* the log odds of the probability of toxicity for the first dose level, allowing to simplify interpretation and prior specification for parameter $\beta_{0,T}$. They use overall survival as an efficacy endpoint and model the time-to-death with a mixture distribution. $p_S(d_j)$ corresponds to the cumulative incidence of death within the assessment window at dose level d_j , and is defined as

$$\text{logit}(1 - p_S(d_j)) = \beta_{0,S} + \beta_{1,S}(d_j - 1) + \beta_{2,S}(d_j - 1)^2 \quad (2.29)$$

A Gumbel Copula is used to model the joint probability of time-to-toxicity and time-to-death. Similarly to the joint model of time-to-event toxicity and efficacy presented previously [Yuan and Yin, 2009], this approach requires the estimation by Bayesian inference of a notable number of model parameters (*i.e.* $(\beta_{0,T}, \beta_{1,T}, \beta_{0,S}, \beta_{1,S}, \beta_{2,S})$).

The dose assignment decision is based on a trade-off between toxicity and efficacy using the algorithm proposed by Thall and Cook (2004).

Late-Onset EffTox design (LO-EffTox) – Jin et al. (2014) proposed the LO-EffTox using Data Augmentation (DA) to facilitate the trial conduct with Delayed Outcomes in EffTox design of Thall and Cook (2004). At the time when a new dose assignment decision is to be made, any unobserved outcomes are considered as missing data. A Bayesian DA algorithm is applied to impute each missing outcome using partial follow-up times and complete outcome data. This is a similar approach to that used by Yuan and Yin (2011) in a phase I setting, that treats delayed outcomes as missing data and uses the EM algorithm under a frequentist framework to estimate toxicity probabilities (cf section 2.1.4). The LO-EffTox design also relies on predictive probabilities and imputation to obtain a completed dataset and make dose assignment decision, but under a Bayesian framework. The authors considered piecewise exponential marginal models for time-to-efficacy and time-to-toxicity (combined in a joint model via the Clayton copula), with weakly informative priors. For the dose escalation algorithm, the EffTox design is used to guide dose assignments.

Robust Bayesian EffTox design – Liu and Johnson (2016) developed a phase I/II design where the toxicity and efficacy probabilities follow a Markov structure. The toxicity (or efficacy) probability of the current dose equals the toxicity (or efficacy) probability of the previous dose plus a positive random variate generated from a scaled Beta distribution. The distribution of toxicity and efficacy outcomes are modeled jointly using a Gumbel distribution, as defined in the EffTox design (cf. equation (2.7)). The objective is to find the MDD, the dose that maximizes a utility function, defined as a trade-off between toxicity and efficacy:

$$U(p_E, p_T) = p_E - w_1 p_T - w_2 p_T \mathbb{1}(p_T > \bar{\pi}_T), \quad (2.30)$$

where $\mathbb{1}(\cdot)$ is the indicator function, w_1 and w_2 are non-negative weights, and $\bar{\pi}_T$ is an upper toxicity threshold pre-specified by clinicians. w_1 can be interpreted as the number of units of efficacy that patients are willing to trade for one unit of decrease in toxicity, and w_2 indicates how strong is the preference for choosing doses with toxicity probabilities below $\bar{\pi}_T$. The design also imposed stopping rules if there is an evidence of either unexpectedly high toxicity or unexpectedly low efficacy. Finally, the approach is close

from the [Cheung and Chappell \(2000\)](#)'s strategy to handle delayed outcomes, assigning weights to the outcomes of patients whose final outcome status has not been observed. The design is qualified as robust it does not assume a parametric model for the shape of the dose–toxicity and dose–efficacy curves. However, it impose monotonicity constraints.

2.2.4.4 TiTE-BOIN designs

Time-To-Event BOIN design for Efficacy and Toxicity (TiTE-BOIN-ET) – [Takeda, Morita, and Taguri \(2020\)](#) extended the BOIN-ET design presented in section 2.2.3, to time-to-event outcomes in order to accelerate the OD determination based on cumulative and pending data of both efficacy and toxicity, resulting in the TiTE-BOIN-ET design. They use a statistical approach close to the TiTE-CRM, also developed by [Lin and Yuan \(2020\)](#) for the TiTE-mTPI and the TiTE-Keyboard designs. Let consider the toxicity, and note $y_{T,k}$, $k = 1, \dots, i$, the binary observed toxicity data until the i^{th} patient by the time of decision making. Let $\delta_{T,k}$ indicate that the toxicity outcome $y_{T,k}$ has been ascertained ($\delta_{T,k} = 1$) or is still pending ($\delta_{T,k} = 0$) up to that time. Given the observed interim toxicity data at the current dose level j , the joint likelihood function for toxicity is given by

$$L(p_{T,j}|data) \propto \prod_{i=1}^k p_{T,j}^{\delta_{T,i}y_{T,i}} (1 - p_{T,j})^{\delta_{T,i}(1-y_{T,i})} (1 - w_{T,i}p_{T,j})^{(1-y_{T,i})} \quad (2.31)$$

where $w_{T,i}$ can be interpreted as a weight, adjusting for the fact that the toxicity outcome has not been ascertained yet. The time-to-toxicity outcome is assumed uniformly distributed over the assessment period $[0, t^*]$ [[Lin and Yuan, 2020](#)], leading to $w_{T,i} = u_{T,i}/t^*$ with $u_{T,i}$ the actual follow-up time for patient i by the decision time.

The same approach is applied to the efficacy outcome. The method used for dose assignment during the trial is similar to the BOIN and BOIN-ET designs, *i.e.* based on targeting toxicity and efficacy thresholds. At the end of the trial, the OD is the dose that maximizes the efficacy probability among the doses lower or equal to the MTD, similarly to the BOIN-ET design.

Time-To-Event BOIN phase I/II trial design (TiTE-BOIN12) – Zhou, Lin, and Lee (2022) recently proposed an extension of the BOIN12 design presented in section 2.2.3, using Bayesian data augmentation to impute Y_T and Y_E from their conditional posteriors, requiring intensive computer resources. After the imputation, the BOIN12 can be directly applied to guide dose transition and identification of MDD.

2.3 Nonparametric optimal Benchmark

In this setting of early phase clinical trials, with a small sample sizes, all dose levels cannot be explored, resulting in limited statistical models. In phase I, the main objective is not to capture the entire dose-toxicity relationships, with toxicity probability estimations that are reliable all dose levels, but to correctly estimate the toxicity probabilities locally around the MTD in order to recommend the appropriate dose level at the end of the trial. Therefore, the performances of the proposed designs are mainly evaluated in terms of the Probability of correct dose selection (PCS) of the MTD and the percentage of patients treated at the MTD during the trial.

Usually, a new proposed design is evaluated against other existing designs sharing similar dose objectives. However, such evaluations are often based on simulation studies, depending on various scenarios under which the data are generated, but that often differ across the studies. For instance, one can conclude that one design succeeds in selecting the correct dose under one true scenario of dose-toxicity curves, and fails in another. This is especially true for Bayesian models that are moreover sensitive to prior information that may favor the performance in some scenarios. This makes the selection of scenarios to evaluate a proposed design subjective and complicated. To evaluate the performances of a design by itself, *i.e.*, differently than by comparison to other existing designs, a first solution is to study the characteristics of the design on a larger sample size. However, this approach is somewhat theoretical, and far from sample sizes observed in early phase settings.

O’Quigley, Paoletti, and Maccario, 2002 developed a nonparametric optimal design

built upon the concept of “complete information”. The general idea is to generate counterfactuals [Rubin, 1974], which are the toxicity outcomes at all dose levels even those that the patient has not been administered. Indeed, in simulating a trial, each patient’s latent outcome can be simulated at all available dose levels, resulting in a collection of toxicity outcomes Y_{ij} for all of the n patients treated at each of the J dose levels:

$$\{Y_{ij}, i = 1, \dots, n, j = 1, \dots, J\}$$

with Y_{ij} the binary toxicity outcome for the i^{th} patient treated at dose level j . Contrarily, in real clinical trial, each patient receives one dose level and only the corresponding Y_{ij} is observable for the i^{th} patient. Therefore, one can only collect “incomplete” information. For instance, let consider a trial investigating five dose levels. Suppose that a patient is given dose level 3 and experiences a DLT. The monotonicity assumption implies that toxicity would necessarily be observed at upper dose levels (4 and 5). By contrast, there is no information regarding whether the patient would have suffered a toxic response on lower dose levels (1 and 2).

While not applicable in a real trial, the benchmark can be used to evaluate the performances of any dose-finding scheme given a particular simulation scenario and a given sample size. Under the benchmark, the percent of selection of the MTD is optimal in the sense it is an unbiased estimator of the true PCS and its variance achieves the Cramer-Rao lower bound; thus, it can be used as an upper bound for the performance. However, it is known that a particular method can provide a higher PCS than the original benchmark under a given scenario. For instance, parametric methods capturing the dose-response relationships since they allow sharing information across dose levels might outperform the nonparametric benchmark in complex scenarios, combined with small sample sizes [Cheung, 2011]. This is especially true with Bayesian approach and strong prior information [Paoletti, O’Quigley, and Maccario, 2004].

Several authors have revised the original benchmark proposed by O’Quigley, Paoletti, and Maccario (2002), built under the assumption of a binary toxicity endpoint monotonically increasing with dose levels and holding in single-agent trials. Zohar and O’Quigley

(2006) further developed the benchmark for estimating the MSD. Cheung (2014) developed the benchmark for phase I/II clinical trials simultaneously evaluating binary toxicity and efficacy endpoints, or for phase I trials with multiple toxicity endpoints. Mozgunov et al. generalized a benchmark for dose-finding designs to various settings with several discrete and continuous outcomes or a dual-agent combination [Mozgunov, Jaki, and Paoletti, 2020; Mozgunov, Paoletti, and Jaki, 2021].

Chapter 3

Phase I dose-finding design and benchmark for a right censored toxicity endpoint

3.1	Introduction	63
3.2	Published manuscript	67
3.3	Supplementary information	84
3.3.1	Benchmark for right-censored toxicity endpoint	84
3.3.1.1	Sensitivity analyses	87
3.3.1.2	Modified Surv-CRM	87

3.1 Introduction

In phase I clinical trials, and notably in oncology, the outcome of interest, *i.e.* the DLT, could be right-censored, when at a sequential analysis, only a fraction of the toxicity window has been observed for a patient who has not developed any DLT yet. Such censoring is likely administrative, and thus independent from the time-to-DLT. Several strategies for phase I designs have been developed to allow continuous accrual and real-time dose assignment for incoming patients when the previous patients outcomes are pending. Three general approaches have been proposed in the literature for phase I trials: (i) using a weighting scheme to reflect the partial follow-up time [Cheung and Chappell, 2000; Mauguen et al. 2011; Lin and Yuan, 2020], (ii) using a KM estimator to ‘fractionate’ the pending toxicity data [Yin, Zheng, and Xu, 2013], or more recently, (iii) using a missing data approach to impute pending toxicity data [Yuan, Lin, et al. 2018; Yuan and Yin, 2011; Liu, Yin, and Yuan, 2013]. Although some of these methods used the denomination

'time-to-event', at the end of the trial the toxicity outcome remains treated as a fully observed binary covariate ignoring when the toxicity outcome event occurred.

Additionally, for patients without DLT at the end of the toxicity observation window, the DLT outcome could be also right-censored. Again, such right-censoring could be considered as non-informative, since it only results from an administrative choice of the time window, irrespective of the patient hazard of toxicity.

Consequently, a natural approach to handle the two situations of right-censored data presented above is the use of survival models for time-to-event data. In time-to-event data, two components are recorded: an observed time and a censoring indicator. In our setting, the time component is the time interval between the date the treatment onset and the date of toxicity or the censoring date.


Therefore, the first methodological objective of the thesis was to refine the models used in the CRM to a survival framework. We thus proposed the *'Survival-CRM'* (Surv-CRM), a phase I dose-finding design for a right-censored time-to-toxicity endpoint. The dose-finding objective of the Surv-CRM is similar to that of the CRM, with the specificity to estimate the probability of toxicity at each dose level by the cumulative incidence of toxicity at the end of the observation window. Through a simulation study, we compared the performance of the Surv-CRM, to that of the TiTE-CRM and of a nonparametric benchmark. Although Surv-CRM and TiTE-CRM working models were different, the designs performed closely in terms of correct dose selection at the end of the trial. Moreover, the proposed design presented slightly better safety results compared to the TiTE-CRM as measured in terms of the Probability of Overdose Selection (POS) at the end of the trial and on the average number of patients experiencing a DLT during the trial. We showed the robustness of the proposed method against various toxicity probability scenarios, time-to-event distributions, and patient accrual rates. Figure 3.1 illustrates an hypothetical phase I clinical trial conducted with the Surv-CRM.

Furthermore, during a phase I clinical trial that enroll patients with an advanced disease, it is not rare that the observation of the toxicity endpoint is avoided due to death, disease progression, patient withdrawal, physician discretion, or any other reason, resulting in trial discontinuations and proposal of alternative therapies if available for ethical reasons. Such a right-censoring situation may rely on the dose received by the

patient, and thus, on his/her hazard of toxicity. Such an informative-censoring of DLT data may bias the estimates of the cumulative incidence function by the KM estimator. More specifically, DLT and any event resulting in trial discontinuation are mutually exclusive events within the toxicity observation window, since only the adverse event that happens first is observable. Indeed, trial discontinuation precludes complete safety assessment as required by the design, resulting in the patient to be non-evaluable and either in the patient replacement, or in several other ad-hoc practical strategies to handle those incomplete follow-ups, for instance using the TiTE-CRM design [Biard, Cheng, et al. 2021]. However, the replacement strategy leads to an increased study duration and cost, and generate a potential bias of selection of healthier patients [Winther et al. 2016]. The second strategy impacts the design performances [Biard, Cheng, et al. 2021].

Hence, to handle such potentially informative censoring, we have extended the Surv-CRM by considering a competing-risk model, referring to the ‘informative Survival-CRM’ (iSurv-CRM). We defined the cause-specific hazard for toxicity as an increasing function of the dose level and the cause-specific hazard for discontinuation as a decreasing function of the dose level. To estimate the probabilities of toxicity, we estimated the cumulative incidence of DLT at the end of the observation window, corresponding to a sub-distribution function, depending on the cause-specific hazard for discontinuation [Putter et al. 2006]. The same dose-finding objective as the Surv-CRM was used. The performances of the iSurv-CRM, were compared to those of the TiTE-CRM. The TiTE-CRM handled incomplete toxicity observations caused by competing discontinuations by assigning weights, reflecting the partial follow-up times, to patients without DLT who discontinued the trial during the observation window. The iSurv-CRM was also evaluated through a nonparametric benchmark approach refined for competing risks data. Results of the simulation study show close clearly better performances of the iSurv-CRM compared to the TiTE-CRM when there were only limited censoring, though the iSurv-CRM outperformed the Surv-CRM otherwise. The benefits of the proposed method are all the more important when informative censoring was high.

Figure 3.1: Illustration of data at the time of a sequential analysis (A) and final analysis (B) of an hypothetical phase I clinical trial using the Surv-CRM.



Let consider a phase I trial with a 0.25 targeted DLT rate, five dose levels, and a total sample size of 25 patients. The DLT assessment window is set at 42 days. The trial starts by treating the first patient at the lowest dose level and slowly escalates until dose level 4, where two patients experienced a DLT. At the arrival of the 18th patient (on day 138), the assigned dose level should be decided (Figure A). By that time (day 138), observations were complete for the first 13 patients. However, at that time, patients 14 to 17 observations are pending, as they spent only a fraction of the scheduled 42-day follow-up without experiencing DLT so far, *i.e.* they are censored at the decision-making time. It would be observed later that patient 17 eventually experienced a DLT. Following the Surv-CRM dose-finding allocation, the 18th patient is treated at dose level 4. At the end of the trial (Figure B), all patients had a complete follow-up: 5 patients experienced a DLT at dose level 4, the others were censored as no DLT was observed during their 42 days evaluation. It took about 8 months to finish the whole trial, contrary to about 30 months if we have applied a design requiring a full DLT assessment before enrolling any new patient cohort.

3.2 Published manuscript

The development of the Surv-CRM, the iSurv-CRM and its evaluation through the non-parametric benchmark are presented in this section, with the article [Andrillon et al. \(2020\)](#) published in the *Journal of Biopharmaceutical Statistics*. Moreover, to facilitate the use of the designs, we have developed graphical user interface-based software that allows users to conduct a trial using the proposed design. The software is freely available at <https://tzw6pc-anais0andrillon.shinyapps.io/SurvCRM/>.



Dose-finding design and benchmark for a right censored endpoint

Anaïs Andrillon^a, Sylvie Chevret^a, Shing M Lee^b, and Lucie Biard^a

^aINSERM U1153 Team ECSTRRA, Université De Paris, Paris, France; ^bDepartment of Biostatistics, Mailman School of Public Health, Columbia University, New York, NY, USA

ABSTRACT

Dose-finding trials aim to determine a safe dose to be tested in larger trials for efficacy. In oncology, designs generally assume conventional monotonic increasing dose-toxicity relationships, mostly with binary outcomes (e.g., dose-limiting toxicity or not), measured in the first cycle of therapy or for a fixed number of cycles. However, with new anti-cancer agents such as molecularly targeted therapies and immunotherapies, late-onset toxicities have become more frequent. Designs with prolonged observation windows and censored endpoints analyzed using survival models, appear particularly suited to these settings. Moreover, in this setting, the observation of the late-onset toxicity endpoint could be precluded by trial discontinuation due to death, progression, patient withdrawal, or physician discretion, defining a competing event to toxicity. We propose extensions of the Continual Reassessment Method (CRM) dose-finding design using survival working models for right-censored endpoints and for handling treatment discontinuation in the toxicity observation window, namely the Survival-CRM (Surv-CRM) and the informative survival-CRM (iSurv-CRM). We also developed a benchmark approach for its evaluation. In a simulation study, we compared the performance of the Surv-CRM and iSurv-CRM, to those of the Time-to-event (TITE)-CRM and the nonparametric benchmark. The performance of the proposed methods was consistent with the complexity of scenarios as assessed by the nonparametric benchmark. Without treatment discontinuations, the Surv-CRM provides proportions of correct dose selection close to those of the TITE-CRM with fewer observed toxicities and patients assigned to overtotoxic dose levels. In the presence of treatment discontinuation, the iSurv-CRM outperforms the TITE-CRM in identifying the correct dose level.

ARTICLE HISTORY

Received 1 September 2020
Accepted 2 September 2020

KEYWORDS

Dose-finding; survival data; late-onset toxicity; competing risks; treatment discontinuation; benchmark

1. Introduction

The general goal of a phase I cancer clinical trial is to define an “acceptable,” usually safe, dose level among several candidates. Toxicity is generally assessed in the first cycle of treatment, based on the observation of dose limiting toxicities (DLT), assuming a monotonic increasing dose-DLT relationship. This optimal paradigm does not work when considering the new classes of anti-cancer agents, such as molecularly-targeted therapies (MTAs) and immunotherapies drugs. Indeed, their modes of action, often targeting the immune system or cancer-cell specific pathways, differ from those of cytotoxic chemotherapies. In particular, new agents may confer late-onset toxicities, related to their specific pharmacology or to prolonged administration, contrary to cytotoxic chemotherapies, that are mainly associated with acute severe toxicities and interval administration (Mathijssen et al. 2014; Postel-Vinay et al. 2011; Wages et al. 2018). Thus, new agents require the use of specific designs with prolonged observation windows (Postel-Vinay et al. 2014) and dose optimization for drugs in later phases of development (Lee et al. 2016,

2019). Furthermore, with longer observation windows, treatment discontinuations during the trial, due to death, progression, patient withdrawal or physician discretion may occur. Toxicity assessment may be precluded by termination of the patient's exposure to the drug.

Some designs have been adapted for continuous accrual and incomplete observations during trials, such as the Time-to-Event Continual Reassessment Method (Cheung and Chappell 2000; Wheeler et al. 2018; Yuan and Yin 2009). The TITE-CRM design addresses the issue of late-onset toxicity by allowing staggered patient accrual relaxing the need for complete DLT follow-up of previously included patients. By allowing the enrollment of new patients when previous patients are still under observation, it is suitable for long observation windows and late-onset toxicities. Moreover, it can shorten the duration of a trial compared to the standard CRM. Although the TITE-CRM handles incomplete observations it relies on binary endpoints for toxicity and does not take into account the time to occurrence of toxicity. Furthermore, even if some practical strategies address treatment discontinuation, operating characteristics of the design may be impacted (Biard et al. 2020). To handle late-onset toxicities while allowing for continuous accrual, we propose new dose-finding designs, the survival-CRM (Surv-CRM) and the informative survival CRM (iSurv-CRM), that formally include the information of time to toxicity using a survival working model for right-censored endpoints, thus allowing the toxicity outcomes to be delayed or unobserved due to competing events within the trial observation window.

As for any clinical trial design, it is crucial to have diagnostic tools to evaluate the performance of a newly proposed design in term of false/correct trial conclusion. O'Quigley et al. (2002) first introduced the nonparametric optimal benchmark to provide an upper bound estimate on the performance of a dose-finding design in the setting of binary toxicity endpoints, under a given scenario, and at a fixed sample size. These benchmark values for dose selection provide meaningful information on the complexity of scenario in terms of maximum tolerated dose (MTD) identification and provide a reference for the performance of a dose-finding algorithm with complete potential information, that is, as if outcomes could be observed at all dose levels for all patients. It contrasts with real clinical trials, in which patients are allocated to receive only one dose level, thus resulting in *incomplete information* due to non observation of patient outcomes at all other dose levels. Cheung (2014) developed the benchmark for phase I/II clinical trials simultaneously evaluating binary toxicity and efficacy endpoints, or for phase I trials with multiple toxicity endpoints. Similarly, Mozgunov et al. (2020) generalized a benchmark for dose finding designs to various settings with several discrete and continuous outcomes. To our knowledge, no specific benchmark approaches for survival dose-finding models has been implemented.

In this work, we propose the survival-CRM (Surv-CRM) design building on the CRM dose-finding design, but using survival models for right-censored DLT endpoints allowing the outcomes to be delayed. To handle possible informative censoring, we also propose the informative survival-CRM (iSurv-CRM), that extends the Surv-CRM by considering a competing-risk model. In addition, we develop a nonparametric benchmark approach for evaluation of dose-finding designs with right censored time-to-event endpoints.

2. Methods

We consider dose-finding clinical trials where patients are followed-up for an observation window $[0, t^*]$ defined for DLT observation and safety assessment. By time t^* , a patient may have two different outcomes: DLT under treatment or not. In the case of novel agents, the need for prolonged observation windows leads us to consider a time-to-event working model which allows for right-censored observations.

Denote n the maximum sample size of the trial. Let $\mathbb{D} = \{d_1, d_2, \dots, d_m\}$ be the set of numerical labels for the m doses investigated in the trial. We assumed no censoring due to lost to follow-up.

Given our framework of oncology phase I dose-finding trials, where patients are usually in advanced stage of their disease, such an assumption is likely valid.

2.1. Problem formulation

2.1.1. Survival-CRM design (Surv-CRM)

We first considered censoring occurs during trial accrual only, at the time a dose assignment computation, when a fraction of the toxicity window has been observed for some patients and they have not developed any DLT yet. In this setting, censored observations are administrative, and thus independent from the time-to-toxicity endpoint, and could be considered as noninformative.

Let T , be the time to first DLT and define $\lambda(t)$ the instantaneous hazard function, and $\Lambda(t)$ the cumulative hazard function:

$$\Lambda(t) = \int_0^t \lambda(s) ds. \tag{1}$$

The survival function $S(t)$, that is the probability of being free of DLT at time t , is given by:

$$S(t) = \exp(-\Lambda(t)). \tag{2}$$

and the cumulative incidence function, $F(t) = \text{Prob}(T \leq t) = 1 - S(t)$, can be expressed as:

$$F(t) = \int_0^t f(s) ds = \int_0^t \lambda(s) S(s) ds \tag{3}$$

where $f(\cdot)$ is the density function of the time to DLT, T .

At the end of a dose-finding trial, given an observation window $[0, t^*]$, for each included patient $i, i = 1, \dots, n$, observations consist in pairs (X_i^*, Y_i^*) , where $X_i^* = \min(T_i, t^*)$ and $Y_i = \mathbb{I}(T_i \leq t^*)$, with $\mathbb{I}(\cdot)$ the indicator function ($Y_i = 1$ if the patient experienced a DLT before t^* , or $Y_i^* = 0$ otherwise).

During the trial, at the time of the dose assignment for a new patient, DLT observations of previously included patients may be right-censored, if only a fraction of the observation window has passed and the patient has not developed any DLT yet. For the MTD estimation during the trial at date τ , we therefore updated data on patient i such as: (X_i^τ, Y_i^τ) where $X_i^\tau = \min(T_i, \tau - \tau_i^0)$ and $Y_i^\tau = \mathbb{I}(T_i \leq \tau - \tau_i^0)$, where τ_i^0 is the date of arrival of patient i .

The proposed survival-CRM is an extension of the CRM with an exponential working model for the cumulative incidence of DLT. Specifically, we assumed T to be exponentially distributed and we assumed the instantaneous hazard of toxicity as an increasing function of the dose:

$$\lambda_j(\beta, d_j) = \exp(d_j \times \exp(\beta)), \text{ for } j \in \{1, \dots, m\} \tag{4}$$

where β is the unknown model parameter considered a random variable with a prior distribution $\Phi(\beta)$ assumed. Given our setting of right-censored data, the exponential distribution was chosen to model failure times. The cumulative incidence of toxicity at the end of observation window t^* for dose $d_j \in \mathbb{D}$ is then defined as follows:

$$F(t^*, \beta, d_j) = 1 - \exp(-t^* \times \exp(d_j \times \exp(\beta))) \tag{5}$$

The survival likelihood on n patients can be expressed as follows:

$$L_n(X, Y|\beta) = \prod_{i=1}^n f(X_i, \beta, d_i)^{Y_i} S(X_i, \beta, d_i)^{1-Y_i} \tag{6}$$

with $d_i \in \mathbb{D}$, the dose allocated to patient i , X_i the right-censored time-to-toxicity and Y_i the toxicity outcome.

Given the prior distribution of β , $\Phi(\beta)$, the posterior distribution for β is estimated using Bayes formula:

$$\Phi(\beta | (X, Y)) = \frac{L_n(X, Y | \beta) \Phi(\beta) d\beta}{\int_{-\infty}^{\infty} L_n(X, Y | \beta) \Phi(\beta) d\beta} \quad (7)$$

and the posterior mean $\hat{\beta}_n$ can be plugged in equation (5) to obtain estimates of the DLT cumulative incidence at each dose level:

$$\hat{\beta}_n = \frac{\int_{-\infty}^{\infty} \beta L_n(X, Y | \beta) \Phi(\beta) d\beta}{\int_{-\infty}^{\infty} L_n(X, Y | \beta) \Phi(\beta) d\beta} \quad (8)$$

2.1.2. Informative survival-CRM design (iSurv-CRM)

We then considered non negligible occurrences of treatment discontinuation in the observation window $[0, t^*]$ due to disease progression, death, withdrawal, and physician discretion, that preclude complete or reliable DLT assessment. In this setting, DLT and discontinuation of treatment are mutually exclusive events within $[0, t^*]$. We thus extended the survival-CRM (Surv-CRM) to the informative survival-CRM (iSurv-CRM) by considering a competing-risk working model in the dose-finding setting.

In a competing risks framework (Putter et al. 2006), let $k = 1$ denote the occurrence of DLT and $k = 2$ that of treatment discontinuation. Let $\lambda_{jk}(t)$ be the cause-specific hazard and $\Lambda_{jk}(t)$ the cause-specific cumulative hazard of cause k at dose d_j . The cause-specific hazard for toxicity $\lambda_{j1}(\cdot)$ was defined according to equation (4) and the cause-specific hazard of discontinuation $\lambda_{j2}(\cdot)$ as follows:

$$\lambda_{2j}(\beta_2, d_j) = \exp(-d_j \times \exp(\beta_2)), \text{ for } j \in \{1, \dots, m\} \quad (9)$$

resulting in a decreasing function of the dose. The cumulative incidence of DLT, at time t^* for dose level j , corresponding to a subdistribution function, depends on cause 2 through:

$$F_1(t^*, \lambda_{1j}, \lambda_{2j}) \int_0^{t^*} \lambda_{1j} \times S(s) ds = \frac{\lambda_{1j} \times (1 - \exp(-(\lambda_{1j} + \lambda_{2j}) \times t^*))}{\lambda_{1j} + \lambda_{2j}} \quad (10)$$

where the event-free survival at time t is defined by:

$$S(t) = \exp(-(\Lambda_1(t) + \Lambda_2(t))) \quad (11)$$

Inference on cause-specific hazards may then be performed separately for toxicity and treatment discontinuation, since the log-likelihood factors into two pieces, one involving λ_1 and the other involving λ_2 (Jeong and Fine 2006).

2.2. Dose-finding algorithm

The objective of the dose-finding trial is to identify the maximum tolerated dose (MTD), d^* , defined as the dose which has the estimated cumulative incidence of toxicity at time t^* closest to a pre-specified target π_{DLT} among the set \mathbb{D} of candidate dose levels $d_j, j = 1 \dots, m$:

$$d^* = \arg \min_{d_j \in \mathbb{D}} |\hat{F}(t^*, \beta, d_j) - \pi_{DLT}| \quad (12)$$

In the setting of competing risks, it becomes

$$d^* = \arg \min_{d_j \in \mathbb{D}} |\hat{F}_1(t^*, \lambda_{1j}, \lambda_{2j}) - \pi_{DLT}| \quad (13)$$

2.2.1. Trial design implementation

In the Surv-CRM and iSurv-CRM designs presented above, dose assignments of new patients are performed from the sequentially updated probabilities of DLT, based on available data at the time of analysis. Sequential estimates of the cumulative incidences of DLT were used to assign the estimated dose \hat{d}^* to the next patient cohort. We used a one-stage CRM design, with a Bayesian estimation of the posterior distribution of parameter β (equation (7)), with a normal prior for β , with mean 0 and standard deviation $\sqrt{1.34}$ (O'Quigley and Shen 1996). This prior allowed a reasonable compromise between flexibility and severity for the skeleton, though sensitivity analyses to this choice were also performed. Then, for patient i 's dose assignment, initiating a new cohort, we computed the posterior mean $\hat{\beta}_i$ of the model parameter (equation (7)) based on data history $H_i = (X_{i-1}, Y_{i-1})$, and used it as plug-in estimate to obtain estimated cumulative incidences of DLT for each dose level j . The new cohort, starting with patient i , was allocated to the MTD estimate, defined using (13), given H_i . The assignment continued in a sequential fashion until the prespecified sample size n was reached.

2.2.2. Benchmark implementation

For the benchmark, we assumed the whole information about a patient's toxicity outcomes could be summarized in a tolerance profile drawn from a uniform distribution $U(1, 0)$. Precisely, for patient i , with profile u_i , the quantile transformation $T_{ij} = F^1(u_i, \lambda_j)$ was applied to obtain the time to DLT outcome that this patient would have had at each dose level j , with F the assumed cumulative distribution function of time-to-toxicity, in our case an exponential distribution. Independent administrative censoring was then applied at time t^* to obtain trial observations. Basically, for a given patient i , this results in a set of outcomes at each of the m candidate dose levels, $(X_{i1}^*, Y_{i1}^*), \dots, (X_{im}^*, Y_{im}^*)$, $i = 1, \dots, n$, also called the *complete information* about patient i .

Given our setting of right-censored data and in the absence of informative censoring, the non parametric Kaplan Meier estimator was used to estimate the cumulative incidences of DLT at each dose j , $\hat{F}(t^*, \lambda_j)$. In the presence of informative censoring, cumulative incidences were estimated as described by Gray (1988). Non parametric estimators for the benchmark were chosen, similarly to the original benchmark proposal for binary endpoints (O'Quigley et al. 2002). In case of a limited sample size n , there might be no DLTs at a given dose level; in these cases, $\hat{F}(t^*, \lambda_j)$ was set to 0. Last, the MTD, d^* , was estimated from the complete dataset as the dose level associated with the estimated probability of toxicity at time t^* , $\hat{F}(t^*, \lambda_j)$ closest to π_{DLT} (equation (13)).

The algorithm presented in Table 1 (Supplementary materials) provides step-by-step guidance on how the benchmark can be constructed based on simulated patient data.

2.3. Simulation study

We simulated trials to assess the performance of our proposed Surv-CRM and iSurv-CRM designs, in comparison to the TITE-CRM and the proposed benchmark.

2.3.1. Data generation

Assuming that the set of probabilities of DLT at time t^* at each dose level is known, complete trial data were generated according to the tolerance profile procedure presented for the benchmark. Inverse transform sampling was applied on the time to DLT cumulative distribution function. We considered time-constant and time-varying toxicity hazards of DLT, using Weibull distributions. Values of 1, 0.5, and 3 were used for the shape parameter to obtain time-constant, decreasing or increasing hazards respectively, given a standardized timeframe for observation $t^* = 1$ (Figure 1), Supplementary

materials). In case of time-varying hazard, Weibull scale parameters were computed from the cumulative incidence of toxicity at time t^* for each scenario, and according to the desired shape value.

Independent administrative censoring was applied at time t^* to mimic the trial observation window and obtain $X_i = \min(T_i, t^*)$ and $Y_i = \mathbb{I}(T_i \leq t^*)$. While the trial design only uses the observed outcomes with the doses assigned following the trial algorithm, the performance of the benchmark is obtained using the *complete information* on patient outcomes at each dose level, hence providing the best performance for the dose-finding objective.

Competing risks data were simulated following Beyersmann et al. (2009). Observations were sampled from an exponential model for the time to first event T (event-free-survival), with all-cause hazard $(\lambda_1 + \lambda_2)$ and the cause of failure at the sampled time T_i determined by a random draw from a Bernoulli distribution with probability $\frac{\lambda_1}{\lambda_1 + \lambda_2}$ for toxicity (otherwise treatment discontinuation).

2.3.2. Simulation setting

Without loss of generality, we set the observation window t^* to 1 (one time unit). We considered $m = 5$ candidate dose levels, and, within the time window t^* , a target toxicity probability of DLT at $\pi_{DLT} = 0.25$. The sample size was set at $n = 25$, with cohort size of 1.

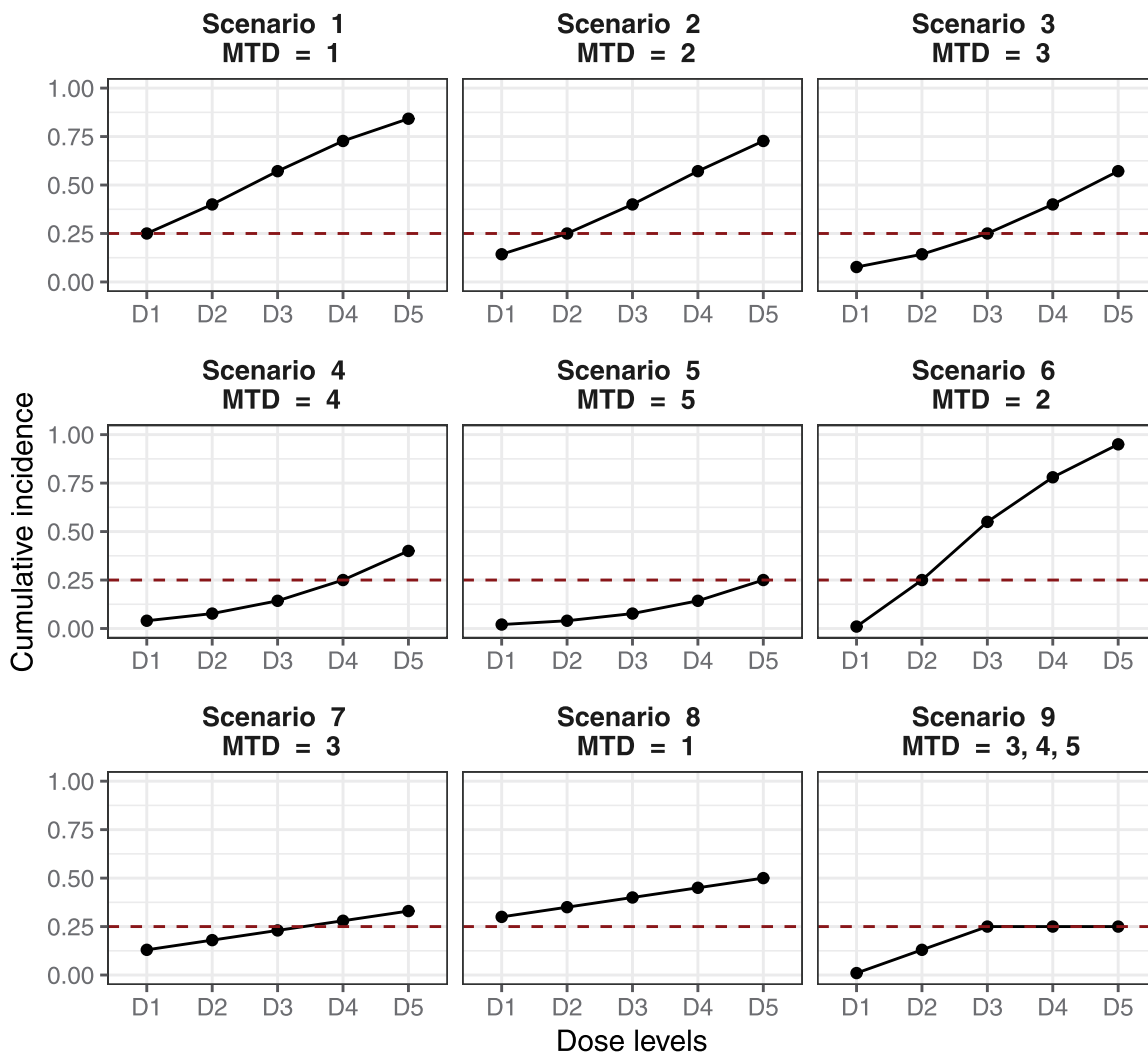


Figure 1. Simulation scenarios: cumulative incidence of DLT at time t^* by dose. Simulated trials included $m = 5$ candidate dose levels.

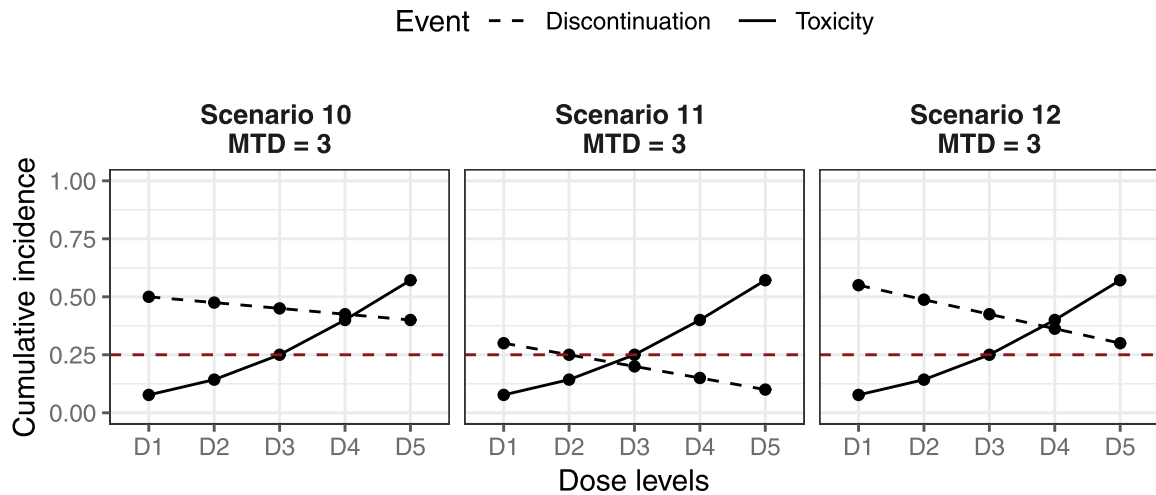


Figure 2. Simulation scenarios: cumulative incidence of DLT (solid line) and discontinuation (dashed line) at time t^* by dose. Simulated trials included $m = 5$ candidate dose levels.

Different non-parametric scenarios of dose-toxicity relationships were examined (Figure 1). They were defined based on the cumulative incidence of DLT at time t^* by dose level and they encompass different shapes of dose-DLT relationships, including monotone increasing (Sc1–Sc8, cytotoxicity scenarios) and plateau (Sc9). For the first five scenarios (Sc1–Sc5), the same interval in DLT probability was chosen between dose levels, with the MTD shifted from dose 1 to dose 5, respectively. The remaining four scenarios were sensitivity scenarios. For scenarios 6 (Sc6), a steeper dose-toxicity curve around the MTD was considered. Scenarios 7 and 8 are similar to scenarios 3 and 1, respectively, but with a less steep slope. Simulations with competing events were performed under three additional scenarios, Sc10 to 12 illustrated in Figure 2. We examined three different dose-discontinuation relationships, combined with toxicity scenario Sc3, previously presented. We considered high and moderate risks of discontinuation with cumulative incidence at t^* decreasing from 0.5 to 0.4 respectively for scenario 10, and from 0.3 to 0.1 for scenario 11. In scenario 12, the cumulative incidence ranges from 0.55 to 0.30.

For a comparison purpose, we also applied a TITE-CRM design to the simulated data. We used the Bayesian framework with an empirical dose-toxicity model and a linear weighting scheme. We applied the TITE-CRM design to all scenarios, without (Sc1–9) and with (Sc10–12) competing discontinuation. For Sc10–12, we handled incomplete toxicity observations caused by competing discontinuation by assigning and maintaining weights throughout the trial to patients without DLT who discontinued the trial during the observation window.

Designs skeletons were calibrated using the indifference interval approach (Lee and Cheung 2009; Wheeler et al. 2019), applied on the toxicity working model estimating the cumulative incidence of toxicity at the end of the observation window defined in equation (5) for Surv-CRM, and based on the empirical model for the TITE-CRM. For the Surv-CRM, we implemented the algorithm described by Lee and Cheung (2009) with target toxicity probability $\pi_{DLT} = 0.25$, $m = 5$ dose levels and dose level 3 as anticipated MTD. Based on 2,000 simulations, on a set of scenarios defined following Lee and Cheung (2009), we set the halfwidth of the indifference interval at 0.05, as providing the highest PCS corresponding to the following skeleton {0.069, 0.151, 0.250, 0.346, 0.426}. For the iSurv-CRM, we calibrated the toxicity skeleton using the best guess prior proposed by Polley (2011), i.e. {0.05, 0.10, 0.25, 0.35, 0.50} and we set the prior estimates of the treatment discontinuation probabilities to {0.50, 0.45, 0.40, 0.35, 0.35} based on clinical considerations. For the TITE-CRM, the halfwidth of the indifference interval was set at 0.07, based on the recommendations of Lee and Cheung (2009), corresponding to the following skeleton: {0.043; 0.124; 0.250; 0.398; 0.542}.

Simulated trials were initiated with the first dose level, d_1 . The cohort size was one patient and no dose skipping was allowed during dose escalation. The trial was terminated either when the maximum sample size was reached, or for safety decisions, i.e. if the posterior probability of $F(t^*|d_1)$ being above the target π_{DLT} was at least 0.95. In the latter, we considered all doses over-toxic and terminated the trial. A modified Surv-CRM with adaptive wait time between two consecutive patients was also examined following the rule proposed by Polley (2011). As a sensitivity analysis, we considered time-constant and time-varying toxicity hazards of DLT and examined different patient accrual schemes by varying the expected number of arrivals per observation window from 1 to 8. Varying the patient-to-window ratio allowed us to simulate short or long observation windows.

2.3.3. Performance measures

A total of $N = 10,000$ independent replications of each scenario were run, using either the survival-CRM or the TITE-CRM design, while benchmark performances were obtained from complete data. Based on those N replicates, for each design, we computed the probability of correct selection (PCS) of the true MTD and safety indexes such as the probability of overdose selection (POS), the average number of patients who experienced a DLT during the trial, and the overdose (OD) number, defined as the average number of patients treated at a dose above the true MTD during the trial (Cheung 2011). We also computed the percentage of early trial stopping for safety decisions, (P_{stop}).

Finally, we calculated the accuracy index, proposed by Cheung (2011), that incorporates information at all dose levels into a single summary measure:

$$\mathcal{A}_n = 1 - m \times \frac{\sum_{j=1}^m |p_j - \pi_{DLT}| \times PS_j}{\sum_{j=1}^m |p_j - \pi_{DLT}|}$$

where p_j is the true probability of DLT for dose level j and PS_j is the probability of selecting dose level j at the end of the trial. The index summarizes the distribution of the selected doses through a weighted average, accounting for the discrepancy of the scenario around the target, π_{DLT} . It ranges from $1 - m$ to 1, reached when the probability of selecting the MTD is equal to 1 and the true probability pre-defined associated to the MTD is equal to the target. This measure penalizes selecting doses distant from the true MTD: \mathcal{A}_n could be negative when the probability of selecting wrong doses is higher than or equal to $1/m$.

3. Results

Tables 1 and 2 summarize the performance of both trial designs and the benchmark, for respectively the five main scenarios (Sc1–Sc5) and the four sensitivity scenarios (Sc6–Sc9), when patient accrual is set at four patients per time window, in the absence of competing discontinuations.

First, as expected, the performance of the methods depended on the scenario. In the case of constant hazards of toxicity over time, the highest benchmark accuracy index, 0.97, was reached for the steepest scenario (Sc6), while the lowest, 0.13, was observed with the flattest one (Sc7); corresponding benchmark PCS were 96% and 24%, respectively. The performance of the proposed Surv-CRM design were in line with these findings, with accuracy index of 0.81 for Sc6 and 0.19 for Sc7. Indeed, in case of a flat dose-toxicity relationship with dose level 3 being the MTD (Sc7), the proposed design failed to identify the right dose the majority of time (PCS of 30%), in line with the benchmark (PCS of 24%). However, in scenario 7, d_3 and d_4 have DLT probability of 0.23 and 0.28 respectively, so that both are reasonable MTD candidates. Indeed, the proportion of correct selection within a 5%-acceptable region, which captures dose candidates lying closely to the true MTD, is 53% (result not shown). In the case of a true plateau relationship (Sc9), the PCS was 83% for the Surv-CRM versus 80% for the benchmark. For scenarios 2 to 4, with the MTD shifted from dose 2 to 4, PCS was 52% for Sc2,

Table 1. Simulation results for Sc1 to Sc5 of the survival-CRM, TITE-CRM and benchmark: percent of correct selection (PCS); accuracy index (A_{25}); percent of overdose selection (POS); percent of stopped trials for safety (P_{stop}); relative bias (R. bias in estimated probability of DLT at the true MTD); average overdose number (OD); number of observed DLT (No. DLT) and number of patients treated with the true MTD (No. MTD) during the trial. $N=10,000$ simulated trials with $n=25$ and $\pi_{DLT} = 25\%$. Accrual rate of four patients per observation window. (n/a: not applicable).

Method	Hazard	PCS	A_{25}	POS	P_{stop}	R. Bias	OD	No. DLT	No. MTD
S1: 0.25; 0.40; 0.57; 0.73; 0.84									
Surv-CRM	decreasing	79	0.89	21	18	0.044	7.52	7.28	17.48
	constant	79	0.89	21	12	0.032	8.57	7.78	16.43
	increasing	76	0.88	24	4	0.004	10.88	8.79	14.12
TITE-CRM	decreasing	74	0.87	26	14	-0.083	9.07	7.70	15.93
	constant	75	0.87	25	12	-0.073	9.97	8.05	15.03
	increasing	76	0.88	24	9	-0.053	11.56	8.78	13.44
Benchmark		80	0.90	20	.	0.000	.	.	.
S2: 0.14; 0.25; 0.40; 0.57; 0.73									
Surv-CRM	decreasing	50	0.71	18	2	0.073	5.81	6.29	8.95
	constant	52	0.71	19	1	0.054	6.90	6.73	9.07
	increasing	54	0.71	22	0	0.015	9.20	7.64	9.09
TITE-CRM	decreasing	57	0.73	21	1	0.018	6.85	6.71	10.06
	constant	57	0.73	21	1	0.021	7.71	7.04	9.90
	increasing	57	0.74	20	1	0.027	9.35	7.71	9.49
Benchmark		57	0.74	19	.	0.000	.	.	.
S3: 0.08; 0.14; 0.25; 0.40; 0.57									
Surv-CRM	decreasing	49	0.56	16	0	0.090	4.99	5.62	8.20
	constant	50	0.57	18	0	0.070	5.98	6.02	8.45
	increasing	53	0.59	20	0	0.031	8.06	6.80	8.59
TITE-CRM	decreasing	53	0.59	19	0	0.072	5.74	5.99	9.04
	constant	54	0.60	18	0	0.070	6.50	6.28	9.08
	increasing	55	0.61	18	0	0.067	8.03	6.84	8.91
Benchmark		57	0.63	19	.	0.000	.	.	.
S4: 0.04; 0.08; 0.14; 0.25; 0.40									
Surv-CRM	decreasing	47	0.48	15	0	0.106	3.94	4.98	7.91
	constant	49	0.50	16	0	0.084	4.76	5.29	8.26
	increasing	51	0.51	20	0	0.046	6.47	5.84	8.59
TITE-CRM	decreasing	50	0.51	17	0	0.118	4.25	5.22	8.76
	constant	51	0.52	17	0	0.113	4.90	5.44	8.88
	increasing	52	0.53	17	0	0.105	6.20	5.82	8.92
Benchmark		57	0.57	19	.	0.000	.	.	.
S5: 0.02; 0.04; 0.08; 0.14; 0.25									
Surv-CRM	decreasing	61	0.67	n/a	0	0.105	n/a	4.02	10.13
	constant	63	0.70	n/a	0	0.084	n/a	4.19	11.14
	increasing	69	0.75	n/a	0	0.049	n/a	4.47	12.89
TITE-CRM	decreasing	63	0.70	n/a	0	0.164	n/a	4.13	10.70
	constant	64	0.71	n/a	0	0.158	n/a	4.25	11.39
	increasing	65	0.71	n/a	0	0.148	n/a	4.45	12.64
Benchmark		75	0.80	n/a	.	0.000	.	.	.

50% for Sc3 and 49% for Sc4 compared to 57% with the benchmark. Of note, with flat scenarios, Sc7 and Sc9, the Surv-CRM design slightly outperformed the benchmark.

Additionally, with constant hazards, the probability of overdose selection (POS) with the survival-CRM was slightly lower to that of the TITE-CRM for most of scenarios. In particular, for Sc1 and Sc8, with MTD at the first dose level, the POS was 21% and 32% with Surv-CRM respectively, versus 25% and 43% with the TITE-CRM. In terms of patient safety during the trials, the average number of patients assigned to a dose above the MTD (OD) during the trial ranged from 4.79 to 11.29 using the survival-CRM versus 4.90 to 13.19 using the TITE-CRM. Finally, the average number of patients who experienced a DLT during the trial (No. DLT) was slightly lower with the Surv-CRM compared to the TITE-CRM for all scenarios. All selection performance measures were not markedly modified when the proportion of right-censored observations decreased, due to a slower accrual rate (Table 2, Supplementary material).

Table 2. Simulation results for Sc6 to Sc9 of the survival-CRM, TITE-CRM and benchmark: percent of correct selection (PCS); accuracy index (A_{25}); percent of overdose selection (POS); percent of stopped trials for safety (P_{stop}); relative bias (R. bias in estimated probability of DLT at the true MTD); average overdose number (OD); number of observed DLT (No. DLT) and number of patients treated with the true MTD (No. MTD) during the trial. $N=10,000$ simulated trials with $n=25$ and $\pi_{DLT} = 25\%$. Accrual rate of 4 patients per observation window. (n/a: not applicable).

Method	Hazard	PCS	A_{25}	POS	P_{stop}	R. Bias	OD	No. DLT	No. MTD
S6: 0.01; 0.25; 0.55; 0.78; 0.95									
Surv-CRM	decreasing	75	0.81	14	0	-0.003	4.72	6.22	13.27
	constant	76	0.81	13	0	0.001	5.60	6.75	12.95
	increasing	76	0.82	12	0	0.006	7.24	7.85	12.42
TITE-CRM	decreasing	82	0.86	12	0	-0.021	4.91	6.67	14.66
	constant	82	0.86	10	0	-0.001	5.51	6.99	14.12
	increasing	80	0.85	7	0	0.033	6.64	7.67	13.01
Benchmark		96	0.97	3	.	0.000	.	.	.
S7: 0.13; 0.18; 0.23; 0.28; 0.33									
Surv-CRM	decreasing	29	0.15	32	1	0.116	7.46	5.36	5.64
	constant	30	0.19	36	0	0.061	8.58	5.57	5.80
	increasing	30	0.23	44	0	-0.036	10.72	5.95	5.84
TITE-CRM	decreasing	30	0.21	37	1	0.077	8.41	5.55	5.92
	constant	31	0.22	39	1	0.042	9.26	5.71	6.03
	increasing	31	0.25	43	0	-0.021	10.90	5.99	6.03
Benchmark		24	0.13	47	.	0.001	.	.	.
S8: 0.30; 0.35; 0.40; 0.45; 0.50									
Surv-CRM	decreasing	71	0.53	29	27	-0.158	10.20	7.47	14.80
	constant	68	0.51	32	18	-0.184	11.29	7.86	13.71
	increasing	61	0.47	39	6	-0.249	13.76	8.54	11.24
TITE-CRM	decreasing	59	0.47	41	20	-0.316	12.08	7.87	12.92
	constant	57	0.46	43	16	-0.327	13.19	8.10	11.81
	increasing	53	0.43	47	10	-0.354	15.29	8.55	9.71
Benchmark		69	0.51	31	.	-0.002	.	.	.
S9: 0.01; 0.13: 0.25; 0.25; 0.25									
Surv-CRM	decreasing	81	0.67	n/a	0	0.148	n/a	4.93	16.58
	constant	83	0.71	n/a	0	0.105	n/a	5.12	17.77
	increasing	88	0.79	n/a	0	0.027	n/a	5.42	19.51
TITE-CRM	decreasing	85	0.75	n/a	0	0.231	n/a	5.12	17.64
	constant	86	0.76	n/a	0	0.199	n/a	5.25	18.49
	increasing	87	0.79	n/a	0	0.143	n/a	5.46	19.78
Benchmark		80	0.66	n/a	.	0.000	.	.	.

Compared to constant hazard, the performance of the Surv-CRM in terms of correct selection slightly improved in case of an increasing hazard of DLT (late-onset toxicities), for Sc2–Sc5 (Table 1). However, late-onset toxicities resulted in an increased overdose selection, whereas early-onset toxicities (decreasing hazard) resulted in improved (smaller) POS. In case of an increasing hazard of DLT, though the survival-CRM had a somewhat lower PCS compared to the TITE-CRM, it outperformed the later in terms of estimation (with smaller relative biases) for all scenarios. In terms of patients' safety during the trial, consistently across scenarios, late-onset toxicities increased the average number of patients assigned to toxic doses and the average number of observed DLTs and decreased the average number of patients assigned to the MTD; this was even more obvious as the accrual rate increased (Figures 3, 4 and 5). In line with these safety results, the Surv-CRM favored stopping trials for safety more often when all dose levels were considered over-toxic. In particular, in scenarios 1 and 8 with dose level 1 as MTD, and early toxicities (time-decreasing hazard of DLT), the percentage of early trial stopping for safety was 18% and 27% with Surv-CRM, respectively versus 14% and 20% with the TITE-CRM (Tables 1 and 2). Finally, in most scenarios, with time-varying hazards and varying accrual rate, Surv-CRM was slightly safer than the TITE-CRM during the trial.

The performance of all three methods improved with larger sample sizes (Figure 2, Supplementary material). The PCS was 100% with the benchmark in all scenarios but Sc7,




Figure 3. Number of patients administered with a dose above the MTD (OD) during the trial with the survival-CRM (barplots above the x-axis) and TITE-CRM (below) designs according to the patient accrual rate by observation window and the shape of hazard. $N=10,000$, $n=25$, and $\pi_{DLT} = 0.25$.

with a minimal sample size varying according to the scenario: from 100 (Sc6) to 300 (Sc4). Consistently, in Scenario 7, using the Surv-CRM and the TITE-CRM, PCS remained low (56% and 58%, respectively), even with $n = 200$.

In the presence of treatment discontinuation precluding complete observation of the toxicity outcomes, simulation results indicated a clear contrast between the informative Surv-CRM and the TITE-CRM (Table 3). The greater the risk of discontinuation, the better the performance of selection with iSurv-CRM compared to TITE-CRM. Indeed, for Sc10 and Sc11 with high and moderate risk of treatment discontinuation, PCS was 60% and 57% with the iSurv-CRM respectively versus 32% and 47% with the TITE-CRM, respectively. For Sc12, the steepest scenario, the iSurv-CRM also outperformed the TITE-CRM (PCS 59% vs. 33%). However, the probability of overdose selection ranged from 10% to 18% with the iSurv-CRM, and 3% to 11% with the TITE-CRM. Last, the iSurv-CRM allocated patients to the MTD more frequently than the TITE-CRM.

4. Discussion

In this paper, we propose an extension of the CRM that uses a survival working model to handle right-censored outcomes, the survival-CRM design, and the informative survival-CRM design for trials with a non negligible risk of competing event related to treatment discontinuation during the desired observation window, as well as a benchmark for its evaluation. This work was motivated by the need

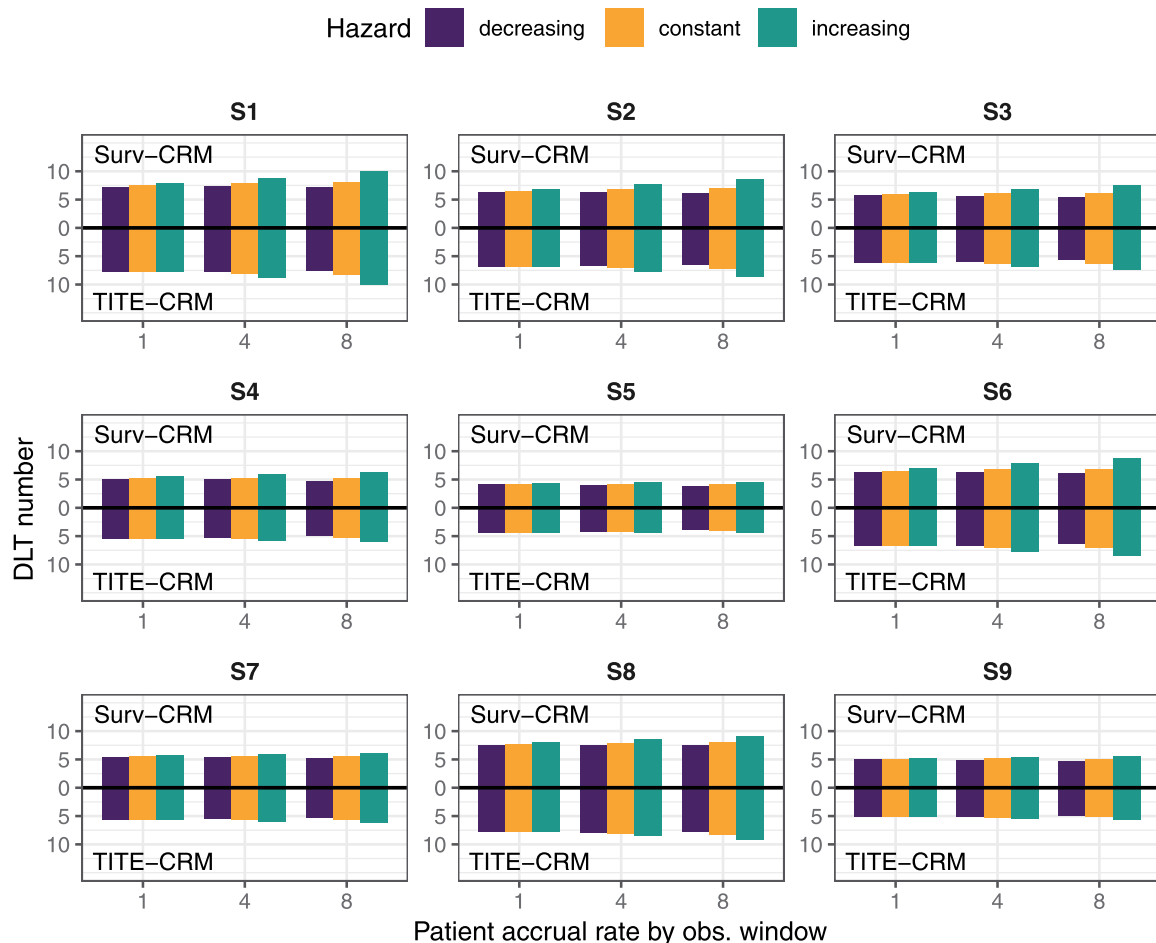


Figure 4. Number of observed DLTs during the trial with the survival-CRM (barplots above the x-axis) and TITE-CRM (below) designs according to the patient accrual rate by observation window and the shape of hazard. $N=10,000$, $n=25$, and $\pi_{DLT} = 0.25$.

for specific methods for dose finding with novel anti-cancer agents, such as targeted therapies or immunotherapies, which often require prolonged observation windows and result in some patients who do not experience a DLT but have not yet reached the end of the window. Survival inference relies of survival functions, assuming all patients will developed the event under treatment, with valid inference under independent or noninformative censoring. In our setting when the endpoint of interest is the DLT within some observation window of interest, the survival working models can provide unbiased estimates of cumulative incidence of DLT at the end of the observation window, given the choice of this window is likely independent of the toxicity process. Otherwise, when the rhythm of inclusion is short relatively to the observation window, and the observation of toxicity could be incomplete in patients already included, these data could be also considered administratively, i.e., noninformatively right-censored. Survival models have recently gained interest in dose finding due to the need for prolonged observation window with new anti-cancer therapies. Nevertheless, they have been mostly used in settings with complex endpoints, combining toxicity and efficacy endpoints or multiple response endpoints (Guo et al. 2018, 2019; Yuan and Yin 2009). For instance, Yuan and Yin (2009) proposed to jointly model toxicity and efficacy as time-to-event outcomes using a Weibull distribution for the hazard of toxicity, resulting in three model parameters to estimate, potentially less suited to our setting of Phase I clinical with small sample sizes. Thus, we proposed the Surv-CRM phase I dose-finding design using a one-parameter exponential working model for the dose-toxicity relationship, allowing for the use of those censored observations in the sequential process of MTD estimation during the trial. Nevertheless, when the observation window is long relatively to the underlying disease process, the observation of toxicity may be precluded by trial discontinuation

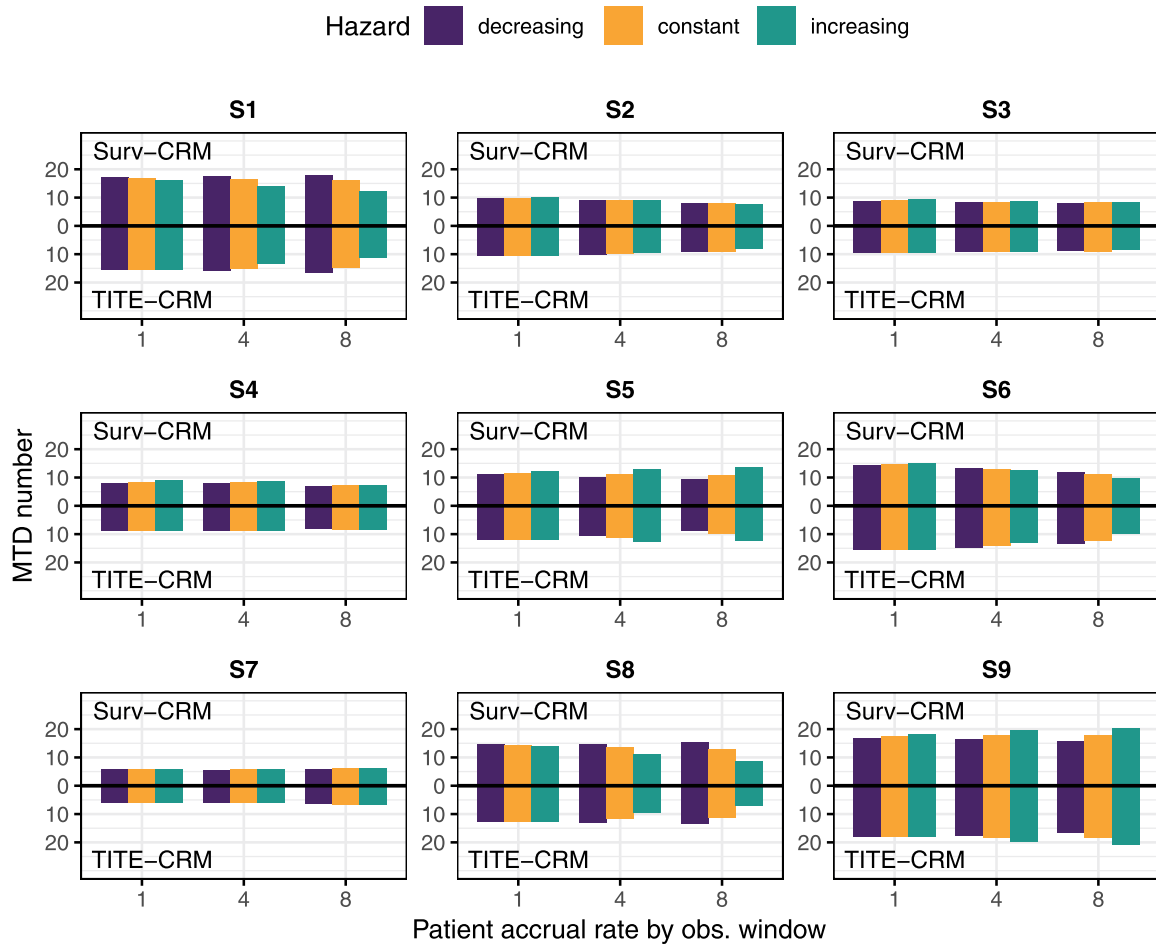


Figure 5. Number of patients treated with the true MTD (correct dose number) during the trial with the survival-CRM (barplots above the x-axis) and TITE-CRM (below) designs according to the patient arrival rate by observation window and the shape of hazard. $N=10,000$, $n=25$, and $\pi_{DLT} = 0.25$.

Table 3. Simulation results for Sc10–Sc12 of the informative survival-CRM (iSurv-CRM), TITE-CRM and benchmark: percent of correct selection (PCS); accuracy index (A_{25}); percent of overdose selection (POS); percent of stopped trials for safety (P_{stop}); relative bias (R. bias in estimated probability of DLT at the true MTD); average overdose number (OD); number of observed DLT (No. DLT) and number of patients treated with the true MTD (No. MTD) during the trial. $N=10,000$ simulated trials with $n=25$ and $\pi_{DLT} = 25\%$. Accrual rate of four patients per observation window; time-constant toxicity hazards. (n/a: not applicable).

Method	PCS	A_{25}	POS	P_{stop}	R. Bias	OD	No. DLT	No. MTD
S10: F_1 : 0.08; 0.14; 0.25 ; 0.40; 0.57								
F_2 : 0.50; 0.48; 0.45; 0.43; 0.40								
iSurv-CRM	60	0.68	10	0.28	0.133	4.14	28.54	11.51
TITE-CRM	32	0.44	3	0.48	0.447	3.51	28.17	6.80
Benchmark	57	0.62	17	.	-0.004	.	.	.
S11: F_1 : 0.08; 0.14; 0.25 ; 0.40; 0.57								
F_2 : 0.30; 0.25; 0.20; 0.15; 0.10								
iSurv-CRM	57	0.63	18	0.08	0.110	5.75	16.56	10.49
TITE-CRM	47	0.57	11	0.15	0.214	5.42	16.65	8.16
Benchmark	56	0.62	17	.	0.000	.	.	.
S12: F_1 : 0.08; 0.14; 0.25 ; 0.40; 0.57								
F_2 : 0.55; 0.49; 0.43; 0.36; 0.30								
iSurv-CRM	59	0.66	12	0.27	0.124	4.32	27.84	11.34
TITE-CRM	33	0.44	4	0.54	0.437	3.81	28.18	6.80
Benchmark	57	0.62	16	.	-0.003	.	.	.

related to lack of efficacy of the drug (due to death, progression, withdrawal, etc.). To accommodate

this particular setting, we proposed the iSurv-CRM that uses estimation of the sub-distribution function of toxicity, to define the MTD, rather than that of the survival function.

Although Surv-CRM and TITE-CRM working models are different, the designs performed closely in most simulations in terms of correct dose selection at the end of the trial. Both designs were mostly robust to the rate of right-censored observations, either due to increased patient accrual or to time-varying hazards of DLT (that is, in case of either early or late-onset toxicities). Furthermore, the proposed design outperformed the TITE-CRM in measures of patient safety during the trial; notably, it allowed a reduced number of patients treated at an overtotoxic dose level during the trial, and of those experiencing a DLT. Although the main objective of a dose-finding trial is to identify the MTD at the end of the trial, safety for patients enrolled in the trial is a matter of concern, for obvious ethical reasons, in these early phase trials. As we observed in the simulation study, inclusion of an adaptive wait time during the trial as proposed by Polley (2011) in our designs further reduces toxic doses selection without compromising the performances in correctly selecting the true MTD and keeping the trial length within a reasonable constraint. In the presence of competing discontinuations precluding complete observation of the toxicity outcomes, iSurv-CRM clearly outperformed the TITE-CRM for selecting the correct dose, expectedly when informative censoring was high. Additional performance comparisons with designs other than the TITE-CRM also assuming a time-to-event endpoint such as the PoD-TPI recently developed by Zhou et al. (2020) could also be considered in further studies.

To assess the operating characteristics of our proposed survival-CRM design, we developed a non-parametric benchmark approach (O'Quigley et al. 2002), adapted to a working model for censored data using the cumulative incidence of toxicity as the dose-finding endpoint, relying on Kaplan-Meier estimator. This provides an useful tool in assessing the design performance in various toxicity scenarios given a pre-specified sample size. The surv-CRM, as well as the TITE-CRM, outperformed the benchmark in scenarios with flat dose-toxicity relationship, illustrating how parametric methods might, in complex scenarios, combined with small samples sizes, outperform the nonparametric benchmark (Cheung 2011).

In summary, we proposed extensions of the CRM using survival working models for dose finding trials, which fit the natural censoring of data collected in these early clinical trials, and can handle an underlying competing risks framework, together with the corresponding benchmark assessment method. The Surv-CRM design, which provides the flexibility of a survival framework, showed desirable operating characteristics, close to those of the TITE-CRM, while providing a slightly safer profile for patients enrolled in these trials. When trial discontinuations within the observation window are likely to preclude the observation of DLTs in a non-negligible proportion of patients, the iSurv-CRM achieved the best performance in selecting the correct dose and allocating patients to the MTD. It should be considered when designing dose-findings trials of targeted therapies with long observation window in advanced cancer patients where treatment discontinuations are likely to occur.

Acknowledgments

The authors thank the guest editors of the special issue of the journal and two referees for valuable comments.

Declaration of interest statement

The authors declare that they have no competing interests.

Data availability statement

R code for the use of the Surv-CRM, the iSurv-CRM and the associated benchmark is freely available on GitHub (<https://github.com/SurvivalCRM>).

References

- Beyersmann, J., A. Latouche, A. Buchholz, and M. Schumacher. 2009. Simulating competing risks data in survival analysis. *Statistics in Medicine* 28 (6):0 956–971. doi:10.1002/sim.3516.
- Biard, L., B. Cheng, G. A. Manji, and S. M. Lee. 2020. A simulation study of methods for handling disease progression in dose-finding clinical trials. *Journal of Biopharmaceutical Statistics* 1–12. doi:10.1080/10543406.2020.1814796.
- Cheung, Y. K. 2011. *Dose finding by the continual reassessment method*. Chapman & Hall/CRC Biostatistics Series. New York.
- Cheung, Y. K. 2014. Simple benchmark for complex dose finding studies. *Biometrics* 70 (2):0 389–397. doi:10.1111/biom.12158.
- Cheung, Y. K., and R. Chappell. 2000. Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* 56 (4):82–1177. doi:10.1111/j.0006-341x.2000.01177.x.
- Gray, R. 1988. A class of k-sample tests for comparing the cumulative incidence of a competing risk. *Annals of Statistics* 16 (3):0 1141–1154. doi:10.1214/aos/1176350951.
- Guo, B., D. Li, and Y. Yuan. 2018. SPIRIT: A seamless phase I/II randomized design for immunotherapy trials. *Pharmaceutical Statistics* 170 (5):0 527–540. doi:10.1002/pst.1869.
- Guo, B., Y. Park, and S. Liu. 2019. A utility-based bayesian phase i–ii design for immunotherapy trials with progression-free survival end point. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 68 (2):0 411–425. doi:10.1111/rssc.12288.
- Jeong, J.-H., and J. Fine. 2006. Direct parametric inference for the cumulative incidence function. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 55 (2):0 187–200. doi:10.1111/j.1467-9876.2006.00532.x.
- Lee, S. M., D. Backenroth, Y. K. Cheung, D. L. Hershman, D. Vulih, B. Anderson, P. Ivy, and L. Minasian. 2016. Case example of dose optimization using data from bortezomib dose-finding clinical trials. *Journal of Clinical Oncology* 34 (12):0 1395–1401. doi:10.1200/JCO.2015.66.0662.
- Lee, S. M., M. Ursino, Y. K. Cheung, and S. Zohar. 2019. Dose-finding designs for cumulative toxicities using multiple constraints. *Biostatistics* 20 (1):0 17–29. doi:10.1093/biostatistics/kxx059.
- Lee, S. M., and Y. K. Cheung. 2009. Model calibration in the continual reassessment method. *Clinical Trials* 6 (3):38–227. doi:10.1177/1740774509105076.
- Mathijssen, R. H. J., A. Sparreboom, and J. Verweij. 2014. Determining the optimal dose in the development of anticancer agents. *Nature Reviews Clinical Oncology* 11 (5):0 272–281. doi:10.1038/nrclinonc.2014.40.
- Mozgunov, P., T. Jaki, and X. Paoletti. 2020. A benchmark for dose finding studies with continuous outcomes. *Biostatistics* 210 (2):0 189–201. doi:10.1093/biostatistics/kxy045.
- O’Quigley, J., and L. Z. Shen. 1996. Continual reassessment method: A likelihood approach. *Biometrics* 52 (2):0 673–84. doi:10.2307/2532905.
- O’Quigley, J., X. Paoletti, and J. Maccario. 2002. Non-parametric optimal design in dose finding studies. *Biostatistics* 3 (1):0 51–56. doi:10.1093/biostatistics/3.1.51.
- Polley, M.-Y. 2011. Practical modifications to the time-to-event continual reassessment method for phase I cancer trials with fast patient accrual and late-onset toxicities. *Statistics in Medicine* 30 (17):0 2130–2143. doi:10.1002/sim.4255.
- Postel-Vinay, S., C. Gomez-Roca, L. R. Molife, B. Anghan, A. Levy, I. Judson, J. De Bono, J.-C. Soria, S. Kaye, and X. Paoletti. 2011. Phase I trials of molecularly targeted agents: Should we pay more attention to late toxicities? *Journal of Clinical Oncology* 29 (13):01728–1735. doi:10.1200/JCO.2010.31.9236.
- Postel-Vinay, S., L. Collette, X. Paoletti, E. Rizzo, C. Massard, D. Olmos, C. Fowst, B. Levy, P. Mancini, P. Lacombe, et al. 2014. Towards new methods for the determination of dose limiting toxicities and the assessment of the recommended dose for further studies of molecularly targeted agents–dose-limiting toxicity and toxicity assessment recommendation group for early trials of targeted therapies, an European organisation for research and treatment of cancer-led study. *Eur J Cancer* 500 (12):0 2040–9. doi:10.1016/j.ejca.2014.04.031.
- Putter, H., M. Fiocco, and R. B. Geskus. 2006. Tutorial in biostatistics: Competing risks and multi-state models. *Statistics in Medicine* 260 (11):0 2389–2430. doi:10.1002/sim.2712.
- Wages, N. A., C. Chiuzan, and K. S. Panageas. 2018. Design considerations for early-phase clinical trials of immune-oncology agents. *Journal for ImmunoTherapy of Cancer* 60 (8). doi: 10.1186/s40425-018-0389-8.
- Wheeler, G. M., A. P. Mander, A. Bedding, K. Brock, V. Cornelius, A. P. Grieve, T. Jaki, S. B. Love, L. Odondi, C. J. Weir, et al. 2019. How to design a dose-finding study using the continual reassessment method. *BMC Medical Research Methodology* 190 (1):0 1–15. doi:10.1186/s12874-018-0638-z.

- Wheeler, G. M., M. J. Sweeting, and A. P. Mander. 2018. A Bayesian model-free approach to combination therapy phase I trials using censored time-to-toxicity data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 680 (2):0 309–329. doi:[10.1111/rssc.12323](https://doi.org/10.1111/rssc.12323).
- Yuan, Y., and G. Yin. 2009. Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 58 (5):0 719–736. doi:[10.1111/j.1467-9876.2009.00674.x](https://doi.org/10.1111/j.1467-9876.2009.00674.x).
- Zhou, T., W. Guo, and Y. Ji. 2020. Pod-tpi: Probability-of-decision toxicity probability interval design to accelerate phase I trials. *Statistics in Biosciences* 120 (1):0 1–22. doi:[10.1007/s12561-019-09264-0](https://doi.org/10.1007/s12561-019-09264-0).

3.3 Supplementary information

3.3.1 Benchmark for right-censored toxicity endpoint

To obtain benchmark performances for our trial given a scenario of true cumulative probabilities of DLT, the trial sample size n and the observation window t^* , we applied a non-parametric benchmark method for dose-finding trials [O’Quigley, Paoletti, and Maccario, 2002]. The approach for complex designs developed by Cheung (2014) was considered in this setting of time-to-event competing endpoints. The algorithm presented in Table S1 provides a step-by-step guidance on how the benchmark can be constructed in our setting based on simulated patient data with right-censored toxicity endpoint for a given scenario of true cumulative probability of DLT $\{p_j\}_{j=1}^m$ of m dose levels at time t^* .

Table S1: Algorithm to compute a benchmark for right-censored toxicity endpoint

- 1 Specify the observation window $[0, t^*]$ and define $F(\cdot)$, the cumulative incidence of toxicity as follows:

$$F(t^*, \lambda_j) = 1 - \exp(-\Lambda_j(t^*))$$

with the cause-specific cumulative hazard of toxicity, $\Lambda_j(t^*) = t^* \times \exp(d_j \times \exp(\beta))$.

- 2 Derive instantaneous hazards λ_j consistent with $F(\cdot)$ for each dose level according to pre-specified scenarios by finding the value of λ_j for which $p_j = F(t^*, \lambda_j)$, with $\{p_j\}_{j=1}^m$, a given scenario of true cumulative probability of DLT of m doses at time t^* . For data generation, Weibull distribution is used to generate time-constant and time-varying toxicity hazards of DLT. We define the instantaneous hazard at dose level j

$$\lambda_j(t) = \frac{a_j}{b_j} \left(\frac{t}{b_j} \right)^{a_j-1}$$

with a_j the shape parameter and b_j the scale parameter at dose level j . Values of 1, 0.5, and 3 are used for the shape parameter, a_j , to obtain time-constant, decreasing or increasing hazards respectively, as presented in Figure S1. The scale parameter b_j is then computed from p_j , the true cumulative incidence of toxicity at time t^* , resulting in

$$b_j = \exp \left(\frac{1}{a_j} \log \left(\frac{-(t^*)^{a_j}}{\log(1 - p_j)} \right) \right)$$

since $p_j = 1 - \exp(-W_j(t^*))$, with $W_j(t)$ the Weibull cumulative hazard at dose level j , defined as $W_j(t) = (\frac{t}{b_j})^{a_j}$.

- 3 Generate a sequence of n patient's toxicity profiles $\{u_i\}_{i=1}^n$ from the uniform distribution $\mathcal{U}(0, 1)$. $\{u_i\}_{i=1}^n$ summarize the tolerance related to the time-to-toxicity event T_{ij} .
- 4 Obtain time-to-toxicity using $T_{ij} = Q^{-1}(u_i, a_j, b_j)$ for each patient i and each dose level j , with $Q(\cdot)$ the cumulative distribution function of the Weibull distribution, resulting in

$$T_{ij} = b_j \times (-\log(1 - u_i))^{1/a_j}$$

- 5 Apply fixed administrative censoring at t^* to obtain trial observations corresponding to clinical trial data with pre-specified observation window t^* : $X_i = \min(T_i, t^*)$ and $Y_i = \mathbb{1}(T_i \leq t^*)$. No additional censoring process is assumed in the specific setting of early phase trials in oncology.
- 6 From all generated complete data (X_i, Y_i) , compute non-parametric estimates of cumulative incidences of toxicity by dose level j , $\hat{F}(\cdot)$, applying the [Kaplan and Meier \(1958\)](#) estimator, relying on conditional probability of not having experienced a toxicity up to t^* given not having experienced a toxicity at $t^* - \delta$:

$$\hat{F}_{KM}(t^*, d_j) = 1 - \hat{S}_{KM}(t^*, d_j) = 1 - \hat{Pr}(X > t^* | X > t^* - \delta) \hat{Pr}(X > t^* - \delta),$$

$$\text{with, } \hat{S}_{KM}(t^*, d_j) = \prod_{k, t_k < t^*} \left(1 - \frac{e_{j, t_k}}{n_{j, t_k}} \right)$$

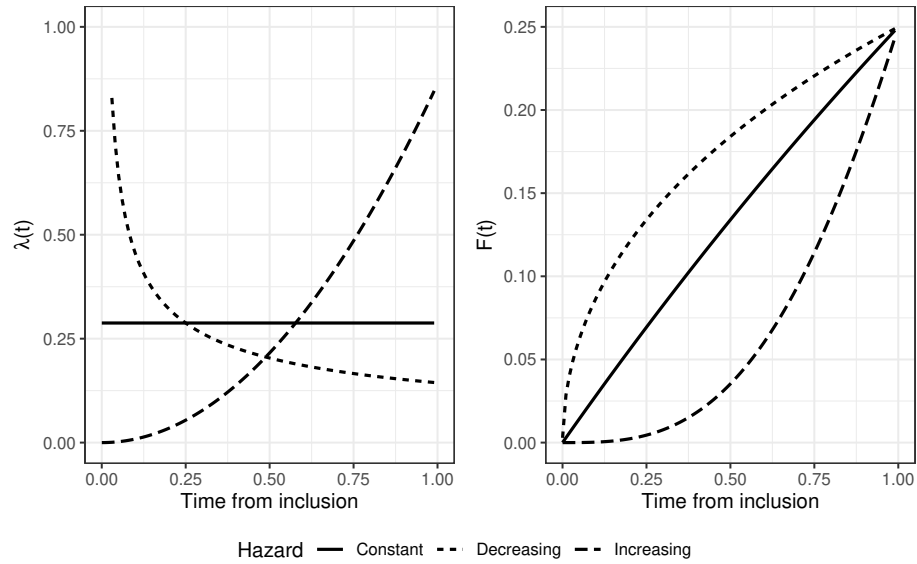
where $n_{j, t}$ is the number of at-risk patients at time t at dose level j and $e_{j, t}$ is the number of patients that undergo the event at time t at dose level j .

- 7 Estimate the MTD, applying the MTD definition to $\hat{F}(t^*, \lambda_j)$ consistently with the dose-finding objective of the evaluated design

$$d^* = \arg \min_{d_j \in \mathbb{D}} |\hat{F}_{KM}(t^*, d_j) - \pi_{DLT}|$$

- 8 Repeat steps 2-7 for $s = 1, \dots, S$ simulated trials.
- 9 Obtain a benchmark performances estimates of selecting the MTD averaging results over the S simulated trials.

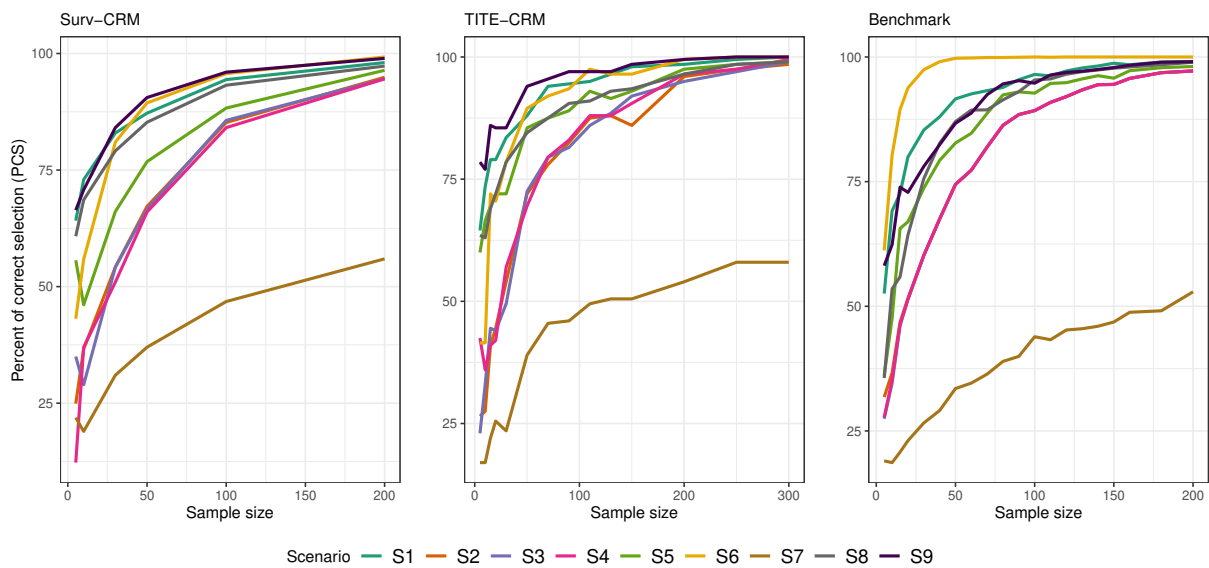
Figure S1: Hazard functions for the time-to-toxicity in the simulation study given a pre-specified cumulative incidence of toxicity at the end of the observation window $F(t^*) = 0.25$, over the observation window $t^* = 1$.



3.3.1.1 Sensitivity analyses

In sensitivity analysis, we varied the total number of patients included in a trial. Results are resumed in Figure S2.

Figure S2: Percent of Correct Selection (PCS) with the survival-CRM, TiTE-CRM and Benchmark. Constant hazard; Patient arrival rate set to 1.



3.3.1.2 Modified Surv-CRM

In a setting of continuous enrollment of patients without complete DLT follow-up of previously treated patients, dose escalation can occur too rapidly at the beginning stage of the trial in the absence of toxic outcomes and patients may be exposed to doses that are assumed to be safe early but later found to be toxic. Polley (2011) pointed out that for the TiTE-CRM and propose to add a safety lead-in, and incorporate an adaptive wait time. In the same line, we proposed a modified version of the survival-CRM incorporating an adaptive wait time between two consecutive patients. This version is referred thereafter as the ‘modified Surv-CRM’ and considered such rules: either stopping rules or adaptive waiting time. We included a wait time expressed as a decreasing linear function of the amount of information available in previous patients who have been treated at the dose administered to the current patient, increasing caution at the early stage of the trial. Using the same notation as Polley (2011), we defined parameter m , corresponding to the

maximum time the investigator is willing to wait before enrolling a new patient in the absence of any previous DLT follow-up information at the current dose and parameter c , corresponding to the threshold above which no accrual suspension is needed. We considered the following two set of values: $\{m = 0.5; c = 10\}$ and $\{m = 0.7; c = 15\}$. We observed in the results of the simulation study presented in Table S2, that the modified survival-CRM including an adaptive wait time before between two consecutive patients has comparable operating characteristics in selecting the correct dose than the unmodified Surv-CRM. As expected, with increasing values of m and c , the modified survival-CRM requires longer trials and reduces the average number of patients treated above the MTD during the trial (OD) compared to the standard Surv-CRM. For example, for Sc1, when $m = 15$ and $c = 0.7$, the average overdose number is 7.56 with the modified Surv-CRM versus 8.57 with the unmodified Surv-CRM.

Table S2: Simulation results of the survival-CRM and the modified survival-CRM with an adaptive wait time: percent of correct selection (PCS); average overdose number (OD); number of observed DLT (No. DLT) and number of patients treated with the true MTD (No. MTD) during the trial; trial length. $N=10,000$ simulated trials with $n=25$ and $\pi_{DLT} = 25\%$. Accrual rate of 4 patients per observation window; time-constant toxicity hazards. (n/a: not applicable).

	Wait time parameters	PCS	OD	No. DLT	No. MTD	Trial length
S1: 0.25; 0.40; 0.57; 0.73; 0.84						
Surv-CRM		79	8.57	7.78	16.43	6.56
Modified Surv-CRM	$c = 0.5; m = 10$	81	7.67	7.98	17.33	8.15
	$c = 0.7; m = 15$	80	7.56	7.87	17.44	10.67
S2: 0.14; 0.25; 0.40; 0.57; 0.73						
Surv-CRM		52	6.90	6.73	9.07	6.62
Modified Surv-CRM	$c = 0.5; m = 10$	51	6.69	6.71	9.10	8.68
	$c = 0.7; m = 15$	52	6.50	6.63	9.32	11.85
S3: 0.08; 0.14; 0.25; 0.40; 0.57						
Surv-CRM		50	5.98	6.02	8.45	6.66
Modified Surv-CRM	$c = 0.5; m = 10$	50	5.75	6.02	8.54	9.03
	$c = 0.7; m = 15$	51	5.64	5.96	8.64	12.54
S4: 0.04; 0.08; 0.14; 0.25; 0.40						
Surv-CRM		49	4.76	5.29	8.26	6.69
Modified Surv-CRM	$c = 0.5; m = 10$	49	5.16	5.31	7.64	9.18
	$c = 0.7; m = 15$	49	4.92	5.27	7.87	12.82
S5: 0.02; 0.04; 0.08; 0.14; 0.25						
Surv-CRM		63	n/a	4.19	11.14	6.75
Modified Surv-CRM	$c = 0.5; m = 10$	63	n/a	4.23	11.69	9.12
	$c = 0.7; m = 15$	63	n/a	4.22	11.53	12.56
S6: 0.01; 0.25; 0.55; 0.78; 0.95						
Surv-CRM		76	5.60	6.75	12.95	6.59
Modified Surv-CRM	$c = 0.5; m = 10$	76	5.42	6.69	13.25	8.66
	$c = 0.7; m = 15$	76	5.16	6.59	13.79	11.93
S7: 0.13; 0.18; 0.23; 0.28; 0.33						
Surv-CRM		30	8.58	5.57	5.80	6.69
Modified Surv-CRM	$c = 0.5; m = 10$	30	8.52	5.59	5.69	8.96
	$c = 0.7; m = 15$	30	8.37	5.56	5.66	12.35
S8: 0.30; 0.35; 0.40; 0.45; 0.50						
Surv-CRM		68	11.29	7.86	13.71	6.55
Modified Surv-CRM	$c = 0.5; m = 10$	74	9.43	8.35	15.57	8.26
	$c = 0.7; m = 15$	74	9.10	8.30	15.90	10.85
S9: 0.01; 0.13; 0.25; 0.25; 0.25						
Surv-CRM		83	n/a	5.12	17.77	6.71
Modified Surv-CRM	$c = 0.5; m = 10$	84	n/a	5.12	17.77	9.06
	$c = 0.7; m = 15$	84	n/a	5.10	17.56	12.51

Chapter 4

Phase I/II dose-finding design for a right censored toxicity endpoint with competing disease progression

4.1	Introduction	90
4.2	Published manuscript	95
4.3	Supplementary information	110
4.3.1	Benchmark for right-censored endpoints and competing risks	110
4.3.2	Model parameters calibration	111
4.3.2.1	Dose skeletons calibration	111
4.3.2.2	Least informative variance	114
4.3.2.3	Calibration process	115
4.3.3	Sensitivity analyses	117
4.3.3.1	Simulation with varying sample size	117
4.3.3.2	Simulation of correlated time-to-events	120

4.1 Introduction

Given the advanced disease of patients enrolled in oncology dose-finding trials, disease progression may occur within the toxicity observation window. Disease progression can be defined in the trial protocol as an aggravation of the patient status related to the malignancy progression. For ethical reasons, patients are taken off the study, and alternative therapies are proposed if available, preventing further assessment of DLT under the planned setting. At the same time, if patients experience a DLT, the drug is also discontin-

ued, preventing further assessment for efficacy. Dual assessment for toxicity and efficacy is only possible if the respective endpoint events do not prevent the complete observation of the other, contrarily to DLT and disease progression in our setting. Consequently, when the patient discontinues the trial, this generates an informative censoring, which defines a situation of competing risks.

For these reasons, we aimed to propose an adaptive design to address the late-onset competing risks outcomes in the early phases of clinical trials. We choose to combine toxicity and efficacy-related information in the proposed design, resulting in a phase I/II design. Using progression as efficacy information is in line with what has been proposed more and more frequently in oncology phase II trials using progression-free survival (PFS) as the efficacy endpoint. Indeed, the definition of early efficacy endpoints for some targeted agents or immunotherapy trials is often tricky. Unlike conventional chemotherapy that works by killing the cancer cell, these agents aim to stop the growth of tumors, so stable disease is often considered a positive signal, and only the occurrence of disease progression is treated as a negative efficacy response [Borcoman et al. 2019]. MTAs and immunotherapy are not necessarily expected to yield rapid tumor shrinkage to achieve a clinical response by standard response criteria (e.g., Response Evaluation Criteria in Solid Tumors (RECIST)). A certain percentage of patients who receive immunotherapy often achieve a long-term durable response [Topalian et al. 2011]. As a result, evaluating the PFS rate is a relevant possibility rather than the tumor response rate.

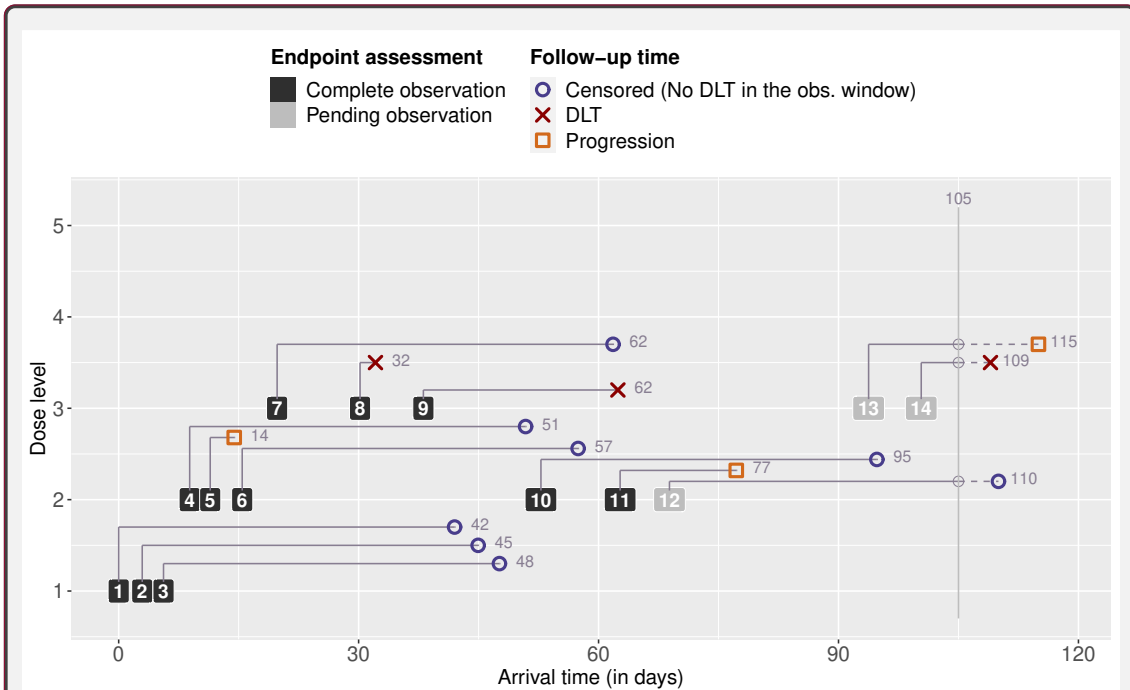
In this modeling setting of competing risks, we may have two perspectives for the same problem. We can target either the observed cumulative incidence of events or the marginal incidence of events. Targeting the observed cumulative incidence of events means considering what will happen to the patients within the observation window. In practice, it may be relevant to decision-maker on the industrial side in terms of actual effective drug requirements, and for regulatory agencies in terms of reimbursement consideration for instance. On the contrary, targeting marginal cumulative incidence of events means considering what may happen to the patients in the hypothetical situation where there is no competing event(s) [Putter et al. 2006]. The marginal cumulative incidence of toxicity could be interpreted as the “potential” probability of toxicity at the end of the observation window if there were no progression events. This approach is theoretical

because the marginal incidence of events could not be observed. However, this perspective is in line with what is done when patients who progress during a trial are replaced, *i.e.*, progression is considered independent from toxicity. To conclude, there are pros and cons to either approach for defining the dose-finding target, and this must be considered when planning as it impacts the choice of the toxicity probability target specification.

As presented in Chapter 2, many phase I/II designs have been proposed in the literature. Still, most designs do not account for competing risks and late-onset outcomes. [Biard, Lee, and Cheng \(2021\)](#) recently developed the Competing Risk-CRM (CR-CRM), a seamless phase I/II design for late-onset competing risks outcomes. [Biard, Lee, and Cheng \(2021\)](#) used a survival framework and targeted theoretical marginal incidence of events. The CR-CRM, is conducted in the three following stages: (i) rule-based initiation, (ii) toxicity-centered escalation, and (iii) toxicity and efficacy optimization. Indeed, the parameter estimation is based on the maximum likelihood, thus requiring an elaboration such as an initial dose-escalation stage. The CR-CRM aimed to identify a dose or a set of doses associated with the minimum progression marginal risk, while satisfying traditional toxicity thresholds. [Zhang, Cao, et al. \(2021\)](#) also proposed an approach to handle competing and delayed toxicity and efficacy outcomes. They used the continuation-ratio method to model the trinomial response outcome similar to [Zhang, Sargent, and Mandrekar \(2006\)](#) as described in section 2.2.2.2 and defined the following three possible levels: ('no adverse event,' 'tumor progression occurring first,' and 'DLT occurring first.'). They modeled the late-onset competing risks outcome by the cause-specific hazard rate. They used a piecewise exponential model to characterize the cause-specific hazard for the competing risks survival outcome for patients who would experience an adverse event during the total follow-up time interval, *i.e.*, the observation window $[0, t^*]$ is partitioned into several disjoint intervals. Pending toxicity and efficacy outcomes at the time of interim analysis are treated as missing data and imputed from the observations via a Bayesian data augmentation method. They used a utility function to measure the desirability of the drug at each dose level, through utility values specified by physicians for each level of the outcome. The design aims to allocate patients to the 'most successful dose' (MSD) as defined in Chapter 2, *i.e.* the dose yielding the highest posterior mean utility within an admissible set.

For this reason, we developed a Bayesian adaptive phase I/II design accounting for right-censored toxicity endpoints with competing disease progression, referring to the ‘Survival-Continual Reassessment Method-12’ (Surv-CRM-12) design. The Surv-CRM-12 is conducted in a single stage allowing for the use of both toxicity and efficacy-related information for the entire trial. It targets the cumulative probability of events within the observation window, considering competing risks. It corresponds to an approach based on the observed patient outcomes, accounting for all competing events in dose-finding and clinical decision-making. The design aims to identify the ‘optimal dose’ (OD) as the dose level(s) with minimum progression risk among acceptable doses in terms of safety. Numerical studies comparing the proposed Surv-CRM-12 to an existing phase I/II designs for time-to-event endpoints [Takeda, Morita, and Taguri, 2020] and to a nonparametric benchmark showed that the proposal yields favorable operating characteristics in dose selection, patient allocation and safety. Sensitive analysis confirmed the robustness of the proposed design. As an illustration, Figure 4.1 represents an hypothetical trial conducted using the Surv-CRM-12 design.

Figure 4.1: Illustration of data at the time of sequential analysis of an hypothetical phase I/II clinical trial using the Surv-CRM-12.



Note: Each number in the squares represents an included patient. Horizontal segment represents the follow-up, along which toxicity is indicated by a red cross, progression by a orange square and censoring by a purple circle.

Let consider a phase I/II trial with a 0.25 targeted DLT rate, five dose levels, and a DLT assessment window set at 42 days. The trial starts by treating the first patient at the lowest dose level and slowly escalates until dose level 3. When the 15th patient enters the trial (on day 105), a complete follow-up of 42 days has been observed for 11 patients, and two patients have already experienced a DLT (patients 8 and 9) at dose level 3. Patients 5 and 11 should not be considered as censored, as they already experienced disease progression. However, at that decision-making time (day 105), patients 12 to 14 are pending, as they have spent only a fraction of the scheduled 42-day follow-up without experiencing neither a DLT nor a progression yet. It would be observed later that patient 13 has finally experienced a progression and patient 14 a DLT. Following the Surv-CRM-12 dose-finding allocation, the 15th patient will be treated at dose level 2.

4.2 Published manuscript

The development of the Surv-CRM-12 has been published in the journal *Statistics in Medicine* [Andrillon et al. 2022]. The paper is presented in this section and associated works are presented in section 4.3.

RESEARCH ARTICLE

Surv-CRM-12: A Bayesian phase I/II survival CRM for right-censored toxicity endpoints with competing disease progression

Anaïs Andrillon^{1,2}  | Sylvie Chevret¹ | Shing M. Lee²  | Lucie Biard¹ 

¹ECSTRRA Team, UMR-1153, Université de Paris, INSERM, AP-HP, Hôpital Saint Louis, Paris, France

²Department of Biostatistics, Mailman School of Public Health, Columbia University, New York, New York, USA

Correspondence

Anaïs Andrillon, Hôpital Saint Louis, Service de Biostatistique et Information Médicale, 1 Avenue Claude Vellefaux, 75010 Paris, France.

Email: anaïs.andrillon@u-paris.fr

Funding information

National Center for Advancing Translational Sciences, Grant/Award Number: UL1TR001873

The growing interest in new classes of anti-cancer agents, such as molecularly-targeted therapies and immunotherapies with modes of action different from those of cytotoxic chemotherapies, has changed the dose-finding paradigm. In this setting, the observation of late-onset toxicity endpoints may be precluded by treatment and trial discontinuation due to disease progression, defining a competing event to toxicity. Trial designs where dose-finding is modeled in the framework of a survival competing risks model appear particularly well-suited. We aim to provide a phase I/II dose-finding design that allows dose-limiting toxicity (DLT) outcomes to be delayed or unobserved due to competing progression within the possibly long observation window. The proposed design named the Survival-continual reassessment method-12, uses survival models for right-censored DLT and progression endpoints. In this competing risks framework, cause-specific hazards for DLT and progression-free of DLT were considered, with model parameters estimated using Bayesian inference. It aims to identify the optimal dose (OD), by minimizing the cumulative incidence of disease progression, given an acceptable toxicity threshold. In a simulation study, design operating characteristics were evaluated and compared to the TITE-BOIN-ET design and a nonparametric benchmark approach. The performance of the proposed method was consistent with the complexity of scenarios as assessed by the nonparametric benchmark. We found that the proposed design presents satisfying operating characteristics in selecting the OD and safety.

KEYWORDS

competing risks, dose-finding, oncology, phase I/II, survival data

Abbreviations: CRM, continual reassessment method; DLT, dose limiting toxicity; MTD, maximum tolerated dose; MTA, molecularly-targeted therapies; PCS, probability of correct selection.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2022 The Authors. *Statistics in Medicine* published by John Wiley & Sons Ltd.

1 | INTRODUCTION

Historically, phase I cancer dose-finding clinical trial designs were developed for cytotoxic drugs and aimed at defining the maximal tolerated dose (MTD) of a new drug from a small and finite set of fixed dose levels; in this setting, the outcome is commonly measured by the occurrence of dose-limiting toxicities (DLT) over a short period of time. This paradigm has been challenged with the arrival of new classes of anticancer agents, such as molecularly-targeted therapies (MTAs) and immunotherapies. Their modes of action target the immune system or tumor specific pathways and differ from those of cytotoxic chemotherapies. In contrast to cytotoxic chemotherapies, which are administered at discrete time points and may cause acute severe toxicities, new agents may be responsible for late-onset toxicities, related to their pharmacological properties or prolonged and continuous administration.¹⁻³ To handle such prolonged observation windows of toxicities,⁴ a number of model-based designs have been proposed. Cheung and Chappell,⁵ introduced the time-to-event continual reassessment method (TITE-CRM), followed by Braun,⁶ Yuan and Yin,⁷ and Liu et al.⁸ Model-assisted designs have been also developed to provide solution for phase I trials with late-onset toxicity. Yuan et al⁹ proposed the Time-to-event Bayesian optimal interval design (TITE-BOIN) based on the single mean imputation method, while Lin and Yuan¹⁰ formulated a general methodology to allow model-assisted trial designs to accommodate pending DLT data.

With prolonged observation windows, competing events observed during the trial, such as disease progressions or death, may also occur, resulting in early treatment or study discontinuation. Toxicity assessment may be precluded by such early terminations of the patient exposure to the drug. Usually, the occurrence of a progression followed by the trial discontinuation within the toxicity window results either in patient replacement leading to increased study duration and cost, or implementation of practical strategies to handle those partial observations.¹¹ However, considering progression as an independent censoring event ignores this information related to efficacy, and may impact the operating characteristics of the design, with decreased performances in selecting the correct dose.¹¹ Accounting for any discontinuation event as informative censoring, using the cumulative incidence of DLT in a survival CRM setting improved the design operating characteristics compared to the binomial TITE-CRM.¹² The information provided by efficacy-related outcomes, such as disease progression, could be also used in the process of selecting a dose, resulting in phase I/II trial designs.¹³⁻²⁰ However, most of these phase I/II designs treat toxicity and efficacy as binary outcomes observed within the timeframe of the trial. Some addressed the issue of late-onset outcomes in phase I/II, assuming independent toxicity and efficacy outcomes.^{16,21,22} Nevertheless, patients who experience toxicity generally have their dose reduced or discontinued and are inconsistently followed for efficacy or progression. Therefore, the occurrences of progression and DLT compete with each other.

In this particular setting of phase I/II with delayed outcomes, Takeda et al²² proposed a Time-to-event Bayesian optimal interval design (TITE-BOIN-ET) to identify the optimal dose (OD), allowing staggered patient accrual by relaxing the need for complete follow-up of previously included patients. They applied the likelihood approach with pending toxicity data proposed by Lin and Yuan for both efficacy and toxicity outcomes.¹⁰ Allowing the continuous enrollment of new patients, when previous patients are still under observation, can shorten the duration of the trial, which is suitable for prolonged observation windows and late-onset toxicities. However, although the TITE-BOIN-ET handles incomplete observations, it relies on binary endpoints for toxicity and efficacy, as stated above. Owing to the competing risks setting, a comprehensive method that formally includes the information of time-to-toxicity and time-to-progression using a survival working model for right-censored endpoints, thus allowing the toxicity outcomes to be delayed or unobserved due to competing events within the observation window, appears fully appropriate for new anti-cancer agents. In this setting, Biard et al recently developed the CR-CRM, a seamless phase I/II design for novel anti-cancer agents with competing disease progressions,²³ using a survival likelihood framework with three stages: first rule-based initiation, second toxicity-centered escalation, and third toxicity and efficacy optimization. The CR-CRM relies on a cause-specific approach targeting the marginal latent incidences of events (toxicity and progression) for dose finding, close to a hypothetical situation where the time-to-toxicity can be observed in the absence of failure due to progression,²⁴ which may be hard to estimate in practice.

In this article, we develop a novel Bayesian Survival-continual reassessment method for phase I/II dose-finding trials, named the Survival-CRM-12 (Surv-CRM-12), using survival models for right-censored and competing DLT and progression endpoints: it allows the outcomes to be delayed or unobserved within the observation window. The proposed method aims to identify the OD, by minimizing the cumulative incidence of disease progression, given an acceptable toxicity threshold. With a single stage design, the Surv-CRM-12 design is easier to implement and allows for

the use of both toxicity and efficacy information for the entire trial. The Surv-CRM-12 also targets the cumulative probability of events within the observation window, taking competing risks into account. It corresponds to an approach based on the observed patient outcomes, accounting for all competing events in dose finding and clinical decision making.

2 | MOTIVATING EXAMPLE

The present work was retrospectively motivated by a phase I/II clinical trial setting in patients with low-risk myelodysplastic syndrome (MDS) to evaluate the tolerance and efficacy of a new targeted therapy combined with erythropoietin (EPO). More specifically, the goal of the trial was to determine the OD level of this targeted therapy in terms of both toxicity and efficacy, with constant dose of EPO. Outcomes were assessed after the first treatment cycle, up to 42 days. The DLT was defined as any grade ≥ 3 drug-related adverse event lasting more than 7 days, or any persisting cytopenias at day 42 without evidence of disease progression in bone marrow and/or peripheral blood. Treatment response was defined at day 21 by an increase in hemoglobin level of 1.5 g/dL or above. Progression free survival (PFS) was a secondary endpoint. The dose-finding was scheduled to use the TITE-BOIN-ET design,²² accounting for possible continuous recruitment during the 42-days observation window for DLT. However, we wondered whether using a survival model to make full use of the DLT and PFS, defined as right-censored competing risks outcomes, would be more informative in this setting.

3 | METHODS

3.1 | Problem formulation

We consider a dose-finding clinical trial where patients are followed up for DLT within an observation window $[0, t^*]$, with the additional issue of a non negligible risk of treatment discontinuation due to disease progression, that precludes the complete and reliable assessment of DLT. In this time-to-event framework in oncology, data are collected as dates and t^* is defined in days, resulting in $t^* > 1$. During the trial, an aggravation of the patient status may occur, related to the progression of the malignancy, leading to treatment discontinuation in order to propose alternative therapies, if available. Similarly, patients who experience toxicity are generally dose reduced or discontinued and are inconsistently followed for efficacy or progression. In this setting, DLT and progression are competing events within $[0, t^*]$, since both result in trial or treatment discontinuation. Consequently, by time t^* , patient can experience one of these three outcomes within the trial: “DLT,” “disease progression,” and “neither DLT nor progression.” Death is considered as a DLT if related to the toxicity and as progression otherwise. Note that we assume no other cause of treatment discontinuation, and no censoring due to lost to follow-up, since it is unlikely in a oncology dose-finding trial, where patients are usually closely followed and in advanced stages of their disease.

Let N , be the maximum sample size of the trial, $\mathcal{D} = \{d_1, \dots, d_J\}$ the set of J doses of the new drug investigated in the trial. The goal of our proposed design is to identify a dose that optimizes the efficacy outcome, among tolerable doses. First, we define the set of tolerable doses, \mathcal{A} , as those dose levels at or below the MTD, defined as the dose with toxicity probability, p_{1j} closest to the target π_{DLT} .

$$\mathcal{A} = \{d_j \in \mathcal{D} : d_j \leq \arg \min_{d_j \in \mathcal{D}} |p_{1j} - \pi_{\text{DLT}}|\}. \quad (1)$$

Then, we incorporate information on time-to-progression to identify the “optimal” dose (OD), d_v . We define the OD as the dose level(s) with minimum progression risk, among \mathcal{A} , the acceptable set of doses:

$$d_v = \arg \min_{d_j \in \mathcal{A}} p_{2j} \quad (2)$$

with p_{2j} , the probability of disease progression at time t^* for dose level j .

3.2 | Survival model

Given the setting described above, we use a competing risks framework^{25,26} to model the probabilities of DLT and disease progression. Let T denote the time-to-failure and k the cause of failure, with $k = 1$ for DLT and $k = 2$ for disease progression. We define $\lambda_{kj}(t|x_{kj}, \beta_k)$, the cause-specific instantaneous, and $\Lambda_{kj}(t|x_{kj}, \beta_k)$, the cumulative hazards of cause k at dose level j , with x_{kj} , the scaled dose obtained by substituting the initial guess of toxicity and progression probabilities.

We model both competing event processes with exponential hazard distributions.^{27,28} Specifically, we assume the cause-specific hazard for DLT $\lambda_{1j}(\cdot)$ is an increasing function of the dose:

$$\lambda_{1j}(t|x_{1j}, \beta_1) = \lambda_{1j} = \exp(x_{1j} \times \exp(\beta_1)), \quad \text{for } j \in \{1, \dots, m\}, \quad (3)$$

and define the cause-specific hazard for progression $\lambda_{2j}(\cdot)$ resulting in a decreasing function of the dose, as follows:

$$\lambda_{2j}(t|x_{2j}, \beta_2) = \lambda_{2j} = \exp(-x_{2j} \times \exp(\beta_2)), \quad \text{for } j \in \{1, \dots, m\}, \quad (4)$$

where β_1 and β_2 are model parameters of the dose-toxicity and dose-progression relationships, respectively.

$F_{kj}(\cdot)$, the cumulative incidence of event k , at time t for dose level j , that is, the subdistribution function depending on both DLT and progression, expresses the probability of failing from cause k at or before t with dose level j , as follows:

$$F_{kj}(t, \lambda_{1j}, \lambda_{2j}) = \int_0^t \lambda_{kj} \times S(s) ds = \frac{\lambda_{kj} \times (1 - \exp(-(\lambda_{1j} + \lambda_{2j}) \times t))}{\lambda_{1j} + \lambda_{2j}} \quad (5)$$

with $S(t) = \exp(-(\Lambda_1(t) + \Lambda_2(t)))$, the event-free survival, that is, the probability of not having experienced a failure, from any cause (DLT or disease progression), at time t .

Given $B = (\beta_1, \beta_2)$ and N included patients, the survival likelihood can be expressed as follows:

$$L_n(\mathbf{C}, \mathbf{Y}, \mathbf{D}|B) = \prod_{i=1}^N f_1(C_i, B, D_i)^{I(Y_i=1)} f_2(C_i, B, D_i)^{I(Y_i=2)} S(C_i, B, D_i)^{I(Y_i=0)} \quad (6)$$

with f_k the density function of T , the time-to-event k , \mathbf{C} the right-censored failure time and \mathbf{Y} the observed outcome defined as $\mathbf{Y} = 1$ when DLT occurs first, $\mathbf{Y} = 2$ when disease progression occurs first and $\mathbf{Y} = 0$ when neither DLT nor progression occurs within the observation window.

Inference on cause-specific hazards could be performed for DLT and disease progression separately, since the log-likelihood factors into two pieces, one involving λ_1 and the other involving λ_2 .^{27,28} Given the set of doses \mathbf{D} , the likelihood $L_n(\beta_k|\mathbf{C}, \mathbf{Y}, \mathbf{D})$ and normal prior distributions $\phi(\beta_k)$ of β_k , the posterior distributions for β_k for event k using Bayes inference is given by:

$$\phi(\beta_k)(\beta_k|\mathbf{C}, \mathbf{Y}, \mathbf{D}) \propto L_n(\mathbf{C}, \mathbf{Y}, \mathbf{D}|\beta_k)\phi(\beta_k). \quad (7)$$

The posterior mean $\hat{\beta}_k$ can be plugged in Equations (3) to (5) to compute $\hat{F}_{kj}(t, \lambda_{1j}, \lambda_{2j})$, the mean estimated cumulative incidence of event k at each dose level j :

$$\hat{\beta}_k = \frac{\int_{-\infty}^{\infty} \beta_k L_n(\mathbf{C}, \mathbf{Y}, \mathbf{D}|\beta_k)\phi(\beta_k) d\beta_k}{\int_{-\infty}^{\infty} L_n(\mathbf{C}, \mathbf{Y}, \mathbf{D}|\beta_k)\Phi(\beta_k) d\beta_k}. \quad (8)$$

The toxicity and progression cumulative incidence estimates are then used to make decision regarding dose allocation as described in Section 3.3.

3.3 | Dose-finding algorithm

Based on the aforementioned probability model, we propose a one-stage dose-finding trial using both toxicity and efficacy information in the dose-assignment process with the goal to identify a dose that optimizes the efficacy outcome, that is, the time-to-progression among a set of acceptable doses, \mathcal{A}_i , as defined below.

After the i th patient being included, the set of acceptable doses (in terms of DLT), that excludes overly toxic doses, is estimated as follows:

$$\mathcal{A}_i = \{d_j \in \mathcal{D} : d_j \leq \arg \min_{d_j \in \mathcal{D}} |\hat{F}_{1j}(t^*, \lambda_{1j}, \lambda_{2j}) - \pi_{\text{DLT}}|\} \quad (9)$$

with $\hat{F}_{1j}(t^*, \lambda_{1j}, \lambda_{2j})$, the vector of estimates of choice for $F_{1j}(t, \lambda_{1j}, \lambda_{2j})$ at each dose level, given available data.

The set of good doses, \mathcal{G}_i , is then defined as the dose levels among \mathcal{A}_i with an estimated incidence of progression at time t^* , $\hat{F}_{2j}(t^*, \lambda_{1j}, \lambda_{2j})$, lying within a range ϵ_P from the minimum estimate:

$$\mathcal{G}_i = \{d_j \in \mathcal{A}_i : \hat{F}_{2j}(t^*, \lambda_{1j}, \lambda_{2j}) - \min(\hat{F}_{2j}(t^*, \lambda_{1j}, \lambda_{2j})) \leq \epsilon_P\}, \quad (10)$$

where ϵ_P is a fixed known parameter that can be tuned based on clinical considerations related to the risk of progression under the standard of care.

To choose the dose allocated to the upcoming ($i + 1$)th patient (or cohort) among the set of good doses \mathcal{G}_i , we use an adaptive randomization algorithm,¹⁷ unless only a single good dose level is identified. Otherwise, randomization across those good doses prevents the method from getting “stuck” on a single dose that has been tried early in the trial. Based on the estimated cumulative incidence of progression, $\hat{F}_{2j}(t^*, \lambda_{1j}, \lambda_{2j})$, for doses $d_j \in \mathcal{G}_i$, randomization probability R_j is computed as follows:

$$R_j = \frac{1 - \hat{F}_{2j}(t^*, \lambda_{1j}, \lambda_{2j})}{\sum_{d_j \in \mathcal{G}_i} 1 - \hat{F}_{2j}(t^*, \lambda_{1j}, \lambda_{2j})}, \quad (11)$$

and is used to allocate the next patient or cohort to dose $d_j \in \mathcal{G}_i$ with probability R_j .

The trial ended either when the maximum sample size N has been reached, or for safety decisions, that is, when the posterior probability of $F_{11}(t^* | d_1)$ exceeding the target π_{DLT} is at least 0.95.²⁹ In the latter, we consider all doses as over-toxic and terminate the trial. If the maximum sample size N has been reached and the follow-up completed for every patient, the OD, \hat{d}_v , is then estimated as the dose with the minimum estimated cumulative incidence of progression among the acceptable set of doses, \mathcal{A}_N^* :

$$\hat{d}_v = \arg \min_{d_j \in \mathcal{A}_N^*} \hat{F}_{2j}(t^*, \lambda_{1j}, \lambda_{2j}). \quad (12)$$

An optional preliminary stage centered on the dose-toxicity relationship could also be conducted before the optimization presented above if clinical considerations justify increased caution related to patients safety at the beginning of the trial. More precisely, we propose to first use a continual reassessment method (CRM)-like approach, in order to guide sequential dose-assignments of the first subset of $r \times N$ patients of the total sample size N , with $r \in [0, 1]$ a design parameter to be pre-specified when planning the trial. Relying on the monotonicity assumption of the dose-toxicity relationship, the objective of this first stage is to identify the dose with a cumulative incidence of DLT at time t^* , $F_{1j}(\cdot)$, closest to a pre-specified target π_{DLT} . At this stage, an upcoming ($i + 1$)th cohort of patients is therefore allocated to dose d_j^{i+1} in the following way:

$$d_j^{i+1} = \arg \min_{d_j \in \mathcal{D}} |\hat{F}_{1j}(t^*, \lambda_{1j}, \lambda_{2j}) - \pi_{\text{DLT}}|. \quad (13)$$

After the rN patients have been included, we incorporate information on the dose-progression relationship in the dose-assignment process as presented above.

4 | SIMULATION STUDY

We conducted a comprehensive simulation study to explore the properties of our proposed Surv-CRM-12 design across a variety of settings. The operating characteristics of the Surv-CRM-12 design were compared with that of the TITE-BOIN-ET²² approach. Indeed, the objective of the TITE-BOIN-ET is also to determine the OD defined as the dose

that maximizes the efficacy probability satisfying a tolerability constraint. Additionally, performances of the proposed design were assessed using a nonparametric benchmark approach.^{30,31}

4.1 | Setting

In line with our motivating example, we set the observation window $t^* = 42$ days. We considered $J = 5$ candidate dose levels, and, within the time window t^* , a target toxicity probability of DLT at $\pi_{\text{DLT}} = 0.25$ and a threshold to define the good doses, $\epsilon_P = 0.10$. We simulated patient accrual using a Poisson process with an expected accrual rate of 4 patients in 42 days. The total sample size was set at $N = 45$ using a single stage design for dose assignment (ie, $rN = 0$). Simulated trials were initiated with the first dose level, d_1 . The cohort size was 1, and no dose skipping was allowed.

To assess the robustness of our proposed method, sensitivity analyses were also conducted. We first conducted a simulation study to examine the sensitivity of our design to the specification of priors. We considered normal priors with different SDs for the unknown parameters β_1 and β_2 under each scenarios. We also varied the sample size from $N = 30$ to $N = 150$, and the number of patients allocated to the first optional toxicity-centered stage, $r = 0.33$ and $r = 0.5$. Furthermore, we also evaluated our method in the case of correlated time-to-event data using the Clayton model following Yuan and Yin's approach.¹⁶

4.2 | Calibration of model parameters

To jointly calibrate the skeletons and the prior variance of the parameters β_1 and β_2 , we extended the calibration procedure using the concepts of the least informative prior variance and indifference intervals developed by Lee and Cheung.³² Specifically, two design skeletons, for toxicity and progression respectively, were calibrated using the indifference interval approach,^{33,34} using the working models of the proposed design (Equation 5).

Let θ_{DLT} be the target probability of toxicity and θ_{PROG} the expected probability of progression at the best guess or expected OD d_v , at t^* . Let $\Theta_k = [b_{k1}, b_{kJ+1}]$ be the parameter space for the working model of event k (ie, $b_k \in \Theta_k$ with $k = 1$ for toxicity and $k = 2$ for progression). We define the indifference intervals for event k , $\mathcal{H}_{k1} = [b_{k1}, b_{k2}]$, $\mathcal{H}_{kj} = (b_{kj}, b_{kj+1})$ for $j = 2, \dots, J-1$ and $\mathcal{H}_{kJ} = (b_{kJ}, b_{kJ+1}]$, where $B_j = (b_{1j}, b_{2j})$ solves

$$F_{kj}(t^*, X_{j-1}; B_j) + F_{kj}(t^*, X_j; B_j) = \begin{cases} 2\theta_{\text{DLT}} & \text{for } k = 1 \\ 2\theta_{\text{PROG}} & \text{for } k = 2 \end{cases}$$

with $X_j = (x_{1j}, x_{2j})$ the scaled doses.

By specifying a common half-width indifference interval for all dose levels, that is δ , X_1, \dots, X_J can be obtained recursively using the approach presented in Lee and Cheung,³³ as described below.

$$x_{1j-1} = \log \left(\frac{-\log \{1 - (\theta_{\text{DLT}} - \delta) - (\theta_{\text{PROG}} + \delta)\}}{\left(1 + \frac{\theta_{\text{PROG}} + \delta}{\theta_{\text{DLT}} - \delta}\right) \times t^*} \right) \times \frac{1}{\exp(b_{1j})},$$

$$x_{2j-1} = -\log \left(\frac{(\theta_{\text{PROG}} + \delta) \times \exp(\exp(b_{1j}) \times x_{1j-1})}{\theta_{\text{DLT}} - \delta} \right) \times \frac{1}{\exp(b_{2j})},$$

for $j = 2, \dots, v$,

$$x_{1j+1} = \log \left(\frac{-\log \{1 - (\theta_{\text{DLT}} + \delta) - (\theta_{\text{PROG}} - \delta)\}}{\left(1 + \frac{\theta_{\text{PROG}} - \delta}{\theta_{\text{DLT}} + \delta}\right) \times t^*} \right) \times \frac{1}{\exp(b_{1j+1})},$$

$$x_{2j+1} = -\log \left(\frac{(\theta_{\text{PROG}} - \delta) \times \exp(\exp(b_{1j+1}) \times x_{1j+1})}{\theta_{\text{DLT}} + \delta} \right) \times \frac{1}{\exp(b_{2j+1})},$$

for $j = v, \dots, J-1$.

Then, the least informative prior normal SD for $\beta_1, \sigma_{\beta_1}^{LI}$, can be obtained for a given δ using the toxicity working model with the least informative prior approach described in Lee and Cheung.³² The least informative SD is the one such that the prior probability of selecting a dose corresponds to a uniform distribution among the candidate dose levels (ie, $1/J$). We assume the prior normal SD for β_2 is equal to the SD of β_1 .

Based on this approach, we calibrated both skeletons (toxicity and progression) and the prior variance for the working models parameters, β_1 and β_2 , in our trial setting. We first evaluated the performances of the Surv-CRM-12 across values of δ based on a set of calibration scenarios where the true probabilities of DLT and progression followed plateau configurations as described by Lee and Cheung.³³ The algorithm was implemented with $t^* = 42$, a target toxicity probability $\theta_{DLT} = 0.25$ and an expected progression probability $\theta_{PROG} = 0.40$ at $d_v, J = 5$ dose levels and dose level 3 as the OD. We iterated δ from 0.01 to $0.6\theta_{DLT}$ in a discrete domain with a grid half-width of 0.01 and calculate $\sigma_{\beta_1}^{LI}$ for each δ value. Based on 2000 simulations, on a set of calibration scenarios of the plateau configuration defined following,³³ we set the half-width of the indifference interval at 0.07 and the prior SD at 0.379, as providing the highest average probability of correct selection (PCS) across the calibration scenarios, corresponding to the following skeletons for toxicity and progression respectively $\{0.055, 0.130, 0.250, 0.406, 0.571\}$ and $\{0.666, 0.541, 0.400, 0.266, 0.158\}$. As a sensitivity analysis, we also considered fixing the SD at $\sqrt{1.34}$ ³⁵ and only calibrating the two skeletons using the indifference interval approach.

4.3 | Simulated scenarios

Several non-parametric scenarios of dose-toxicity and dose-progression relationships were examined (Figure 1). They differed in terms of the cumulative incidences of DLT and disease progression at time t^* by dose level and encompassed different realistic shapes of dose-DLT and dose-progression relationships.

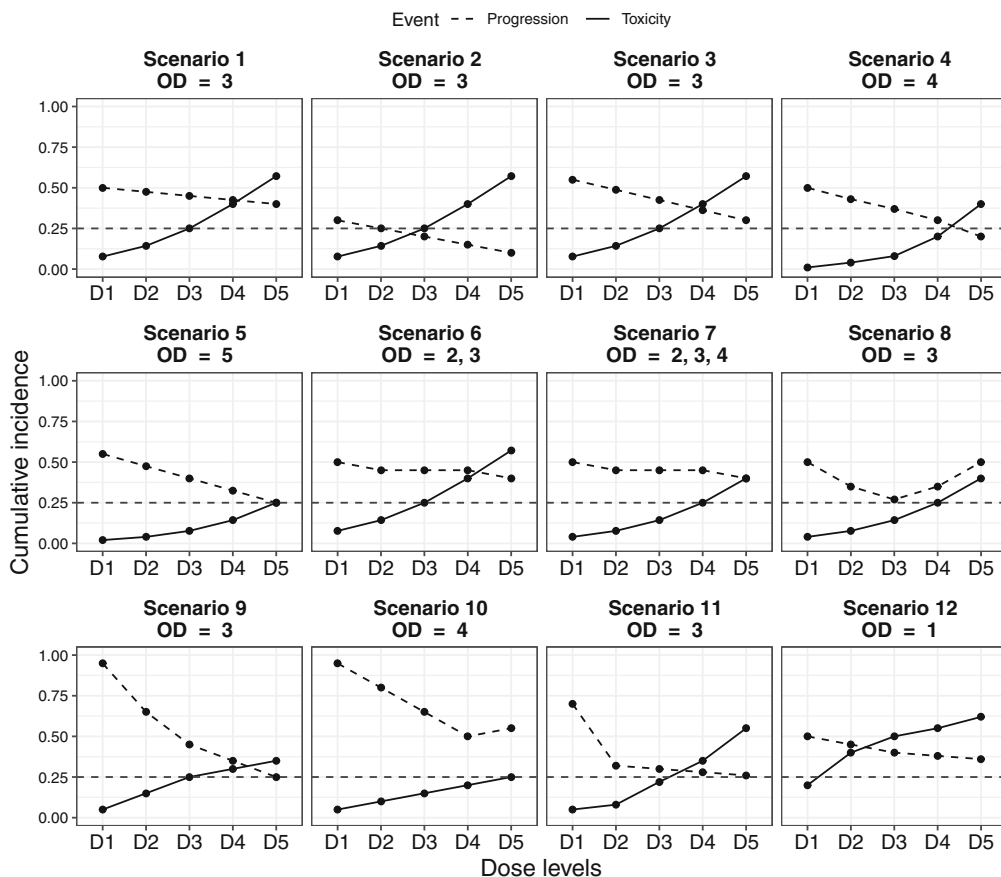


FIGURE 1 Simulation scenarios: True cumulative incidences of DLT (solid line) and progression (dashed line) at time t^* for each dose level, D1 to D5. OD, optimal dose

The first seven scenarios (Sc) had varying optimal and acceptable doses in position over the range of the five candidate dose levels: for Sc1 to Sc3, Sc6, Sc9, and Sc11, acceptable doses were d_1 to d_3 ; for Sc4, Sc7, and Sc8, d_1 to d_4 , and for Sc5 and Sc10, d_1 to d_5 . Scenarios also differed regarding the shape of dose-progression relationships including high (Sc9-Sc11), low (Sc2) and moderate risk of disease progression, steeper slope (Sc1-Sc5) and plateau (Sc6 and Sc7). We also examined operating characteristics of our proposed design in comparison with the benchmark approach and the TITE-BOIN-ET design on a U-shape dose-progression scenario (Sc8). To avoid cherry-picked scenarios that favor the proposed design, we examined four additional scenarios previously used in the study by Takeda et al²² (Sc9-Sc12). For simplicity, we defined the true cumulative incidence of progression at time t^* as the complementary of the true probability of efficacy.

Trial data were generated according to a tolerance profile procedure^{30,31} drawn from a uniform distribution $\mathbb{U}(0, 1)$ summarizing patient's toxicity and progression outcomes. Instantaneous cause-specific hazards of both events, (λ_1 and λ_2) were back computed from cumulative incidence values of event assumed at time t^* at each dose level, using exponentially distributed event times. Competing risks data were then simulated following Beyersmann et al's approach.³⁶ Observations were sampled from an exponential model for the time-to-first event T (event-free survival), with all-cause hazard ($\lambda_1 + \lambda_2$) and the cause of failure determined by a random draw from a Bernoulli distribution with probability $\frac{\lambda_1}{\lambda_1 + \lambda_2}$ for toxicity (otherwise disease progression). Last, administrative censoring was applied at time t^* to mimic the trial observation window. For trials generated under the Clayton model for time-to-event data, details are available in the Supporting Web Materials.

4.4 | Comparative methods

To evaluate performances of our proposed design, we developed a nonparametric benchmark approach providing an upper bound estimate on selecting the OD under a given scenario, and at a fixed sample size.^{12,30,31} We revised the benchmark approach proposed by Cheung³¹ for complex designs with a bivariate dose-finding objective, to account for right-censored endpoints and competing risks events. These benchmark values for dose selection provide meaningful information on the complexity of scenario in terms of OD identification and provide a reference for the performance of a dose-finding algorithm with complete potential information, that is, as if outcomes could be observed at all dose levels for all patients. In our setting of a bivariate dose-finding objective relying on the cumulative incidence of DLT and progression, we computed a benchmark based on two endpoints (toxicity and progression).³¹ In the presence of informative censoring, cumulative incidences were estimated as described by Gray.³⁷ Non parametric estimators for the benchmark were chosen, similarly to the original benchmark proposal for binary endpoints.³⁰ The OD was estimated from the complete dataset based on the trial objective defined in Equation (2). Details of the benchmark approach are provided in Supporting Web Material (Table S1).

For a comparison purpose in dose selection performances, we also applied the TITE-BOIN-ET design²² to the simulation scenarios. With the TITE-BOIN-ET design, the method used to guide dose escalation is different from the method used for final estimation. Briefly, for OD estimation at the end of the trial, regression models for each of the toxicity and efficacy outcomes are fitted. For toxicity, an isotonic regression is performed and for efficacy, a model with fractional polynomials with two degrees of freedom is applied. The OD is then determined as the dose that maximizes the efficacy probability among a set satisfying tolerability, that is, dose levels equal or below the MTD (ie, dose with toxicity probability closest to the target toxicity probability π_{DLT}). In our setting, we assumed efficacy when a patient is alive without progression. Thus, the true cumulative incidence of progression at time t^* are defined as the complementary of the true probability of efficacy. A set-up for simulations as close as possible to our simulation setting was used for a fair comparison with respect to the maximum number of patients ($N = 45$), cohort size ($=1$), starting dose (d_1), assessment windows for toxicity and efficacy (42 days) and accrual rate (one patient per 10 days). For dose assignment during the trial, the target toxicity probability and the target efficacy probability were $\pi_{DLT} = 0.25$ and $\pi_{EFF} = 0.95$, respectively. To match the Surv-CRM-12 dose-finding objective of minimizing time-to-progression, which corresponds to maximizing the efficacy probability in the TITE-BOIN-ET setting, we chose to set a target efficacy probability, π_{EFF} , close to 100%. We set the highest acceptable toxicity probability $\pi_{1 DLT} = 0.4\pi_{DLT}$ (ie, 60% deviation from the target), the lowest DLT probability that is overly toxic $\pi_{2 DLT} = 1.4\pi_{DLT}$, (ie, 40% deviation from the target), and the highest efficacy probability that is subtherapeutic $\pi_{1 EFF} = 0.6\pi_{EFF}$ (ie, 40% deviation from the target). As a result, the optimal values of the lower and upper cutoffs on toxicity and the cutoff on efficacy were estimated to be $\{0.21; 0.29; 0.80\}$. We defined safety decisions as close as possible to our setting. Skipping dose(s) was not allowed. Early termination criteria

were set as follows: when posterior probability of efficacy was less than 0.05 (ie, $1 - \pi_{\text{EFF}}$) was above 0.95, or the posterior probability of toxicity was above 0.25 (ie, π_{DLT}) was above 0.95, the dose was removed. If all doses were removed, the dose-finding trial was terminated. An early termination rule was applied based on at least 3 patients enrolled in a trial.

All the simulations were based on 10 000 replicates for the Surv-CRM-12 and TITE-BOIN-ET designs, and benchmark. For each scenario, we computed the percentage of times that each of the designs selected each dose level as the final OD, percentage of early trial stopping for safety decisions (P_{stop}), the average overdose number (No. OV), defined as the average number of patients treated at a dose above the true OD during the trial,³⁸ the average number of patients that experienced DLT (No. DLT) and disease progression (No. Prog), and the average number of patients treated at the true OD (No. OD).

4.5 | Results

Table 1 summarizes the simulation results with a total sample size of 45 patients. As expected, the performance of the methods depended on the scenario. The highest percents of correct selection (PCS) reached by the benchmark, 91% and 97%, were observed in scenarios with more than one OD due to a plateau shape of the dose-progression relationship (Sc6 and Sc7, respectively). Performances of the proposed Surv-CRM-12 design were in line with these findings, as highest PCS, 91% and 93%, were reached for Sc6 and Sc7 respectively. On the other hand, the Surv-CRM-12 selected the relatively higher dose levels as the ODs even if the risks of toxicity were increased but acceptable among the dose levels when the risks of disease progression were the same. In scenarios with moderate risk of disease progression and monotone decreasing dose-progression (Sc1 to Sc5), the Surv-CRM-12 correctly selected the OD in more than 60% of cases and outperformed the TITE-BOIN-ET method. The Surv-CRM-12 seemed robust to different dose-progression relationships given a fixed risk of toxicity. Indeed, for Sc1 to Sc3, with the same dose-toxicity relationship and different flat dose-progression relationships, the Surv-CRM-12 provided a satisfactory PCS greater than or equal to 71% and outperformed the benchmark with PCS fluctuating between 52% and 70%. In scenarios with OD at extreme dose levels (Sc5 and Sc12), the Surv-CRM-12 picked the OD in 64% and 66% of cases respectively, contrary to 76% and 90% with the benchmark. As expected, the lowest PCS, 27%, was observed in the case of a U-shape progression scenario (Sc8), while the PCS with the non-parametric TITE-BOIN-ET was 51%, and 80% with the benchmark. Finally, with huge risks of progression, 95% for the first dose level (Sc9, Sc10), the Surv-CRM-12 and the TITE-BOIN-ET hesitated between the correct dose level and the one above: PCS were 55% and 48% with the Surv-CRM-12 for Sc9 and Sc10 respectively, slightly higher than the TITE-BOIN-ET (PCS of 43% and 34% respectively), while the PCS with the benchmark were 47% and 69% respectively. To conclude in terms of picking the correct dose, the Surv-CRM-12 outperformed the TITE-BOIN-ET method over all scenarios except in the case of a U-shape dose-progression relationship.

In terms of patient exposure during the trial over the 12 scenarios, the mean number of patients allocated to overdoses (OV) was 10.19 out of a total sample size of 45 patients with the Surv-CRM-12 design, below the 14.08 mean value observed with the TITE-BOIN-ET. However, fewer patients experienced DLT during the trial using the TITE-BOIN-ET: the average number of DLT over all scenarios was 8.73 vs 10.14 with Surv-CRM-12. Finally, for 11 out of the 12 scenarios, the average number of patients treated at the true OD was higher with the Surv-CRM-12 design and the average OD over all scenarios was 24.95 against 17.88 with the TITE-BOIN-ET.

As shown by Lee and Cheung,³² having a larger prior variance changes the distribution of the prior probability of dose selection from uniform to U-shaped, leading in higher chances of extreme doses selections. Indeed, with SD equal to $\sqrt{1.34}$, PCS was higher for Sc5 with OD at level 5, equal to 75%. For Sc12, with OD in dose level 1, PCS was lower with SD equal to $\sqrt{1.34}$, equal to 59% but the percent of stopped trial for safety decision was much higher (equal to 18% contrary to 4% with the least informative SD) (Table S2 in Supporting Web Materials).

Expectedly, the performances of all three compared designs improved with increasing sample sizes (Figure S2, Supporting Web Materials). Moreover, the impact of the sample size of the optional first stage was limited (Table S3, Supporting Web Materials), so that this could be left to clinicians' preference.

Finally, sensitivity simulations showed that the performance of the Surv-CRM-12 was mostly preserved in case of correlated events. The PCS was maintained under data generated with correlated times to toxicity and progression, but the POS increased; similar results were observed with the benchmark (Figure S3, Supporting Web Materials).

TABLE 1 Simulation results for Sc1 to Sc6 of the Surv-CRM-12 design, TITE-BOIN-ET and benchmark; percent of stopped trials for safety (P_{stop}); percent of selection; number of overdose (No. OV); number of observed DLT (No. DLT); number of observed progression (No. Prog) and number of patients treated with the true OD (No. OD) during the trial

Method	P_{stop} (%)	Percent of selection by dose level (%)					No. OV	No. DLT	No. Prog	No. OD
		1	2	3	4	5				
Scenario 1										
True F_1		0.08	0.14	0.25	0.40	0.57				
True F_2		0.50	0.48	0.45	0.43	0.40				
Surv-CRM-12	0	0	18	73	9	0	7.0	11.0	20.4	25.5
TITE-BOIN-ET	1	15	30	41	9	4	13.1	10.4	21.5	15.3
Benchmark	n/a	12	34	52	1	0	n/a	n/a	n/a	n/a
Scenario 2										
True F_1		0.08	0.14	0.25	0.40	0.57				
True F_2		0.30	0.25	0.20	0.15	0.10				
Surv-CRM-12	0	0	18	71	10	0	6.4	10.5	9.5	24.2
TITE-BOIN-ET	1	14	30	47	8	1	6.3	9.1	10	15.2
Benchmark	n/a	3	28	64	5	0	n/a	n/a	n/a	n/a
Scenario 3										
True F_1		0.08	0.14	0.25	0.40	0.57				
True F_2		0.55	0.49	0.43	0.36	0.30				
Surv-CRM-12	0	0	17	73	10	0	7.5	11.0	19.5	25.3
TITE-BOIN-ET	1	12	25	47	12	3	13.1	10.7	20.3	15.8
Benchmark	n/a	2	25	70	3	0	n/a	n/a	n/a	n/a
Scenario 4										
True F_1		0.01	0.04	0.08	0.20	0.40				
True F_2		0.50	0.43	0.37	0.30	0.20				
Surv-CRM-12	0	0	0	4	78	19	7.2	8.9	13.8	26.5
TITE-BOIN-ET	0	4	5	16	56	21	14	7.9	16.6	16
Benchmark	n/a	0	0	9	76	15	n/a	n/a	n/a	n/a
Scenario 5										
True F_1		0.02	0.04	0.08	0.14	0.25				
True F_2		0.55	0.48	0.40	0.32	0.25				
Surv-CRM-12	0	0	0	1	35	64	n/a	7.5	14.3	17.1
TITE-BOIN-ET	5	2	3	11	25	54	n/a	6.2	19.4	26.9
Benchmark	n/a	0	0	1	23	76	n/a	n/a	n/a	n/a
Scenario 6										
True F_1		0.08	0.14	0.25	0.40	0.57				
True F_2		0.50	0.45	0.45	0.45	0.40				
Surv-CRM-12	0	0	19	72	9	0	6.9	10.9	20.6	36.5
TITE-BOIN-ET	1	15	33	39	8	4	13	10.3	21.4	28
Benchmark	n/a	8	46	45	1	0	n/a	n/a	n/a	n/a

TABLE 1 (Continued)

Method	P_{stop} (%)	Percent of selection by dose level (%)					No. OV	No. DLT	No. Prog	No. OD
		1	2	3	4	5				
Scenario 7										
True F_1		0.04	0.08	0.14	0.25	0.40				
True F_2		0.50	0.45	0.45	0.45	0.40				
Surv-CRM-12	0	0	0	26	67	6	3.7	9.4	20.1	40.1
TITE-BOIN-ET	0	8	15	30	30	17	14.1	7.5	24	28
Benchmark	n/a	2	27	36	34	1	n/a	n/a	n/a	n/a
Scenario 8										
True F_1		0.04	0.08	0.14	0.25	0.40				
True F_2		0.50	0.35	0.27	0.35	0.50				
Surv-CRM-12	0	0	0	27	67	5	24.2	9.1	15.1	16.3
TITE-BOIN-ET	0	4	14	50	21	11	21.4	6.8	19.8	14.2
Benchmark	n/a	0	5	80	15	0	n/a	n/a	n/a	n/a
Scenario 9										
True F_1		0.05	0.15	0.25	0.30	0.35				
True F_2		0.95	0.65	0.45	0.35	0.25				
Surv-CRM-12	0	0	14	55	26	5	14.1	11.0	21.0	20.8
TITE-BOIN-ET	1	1	18	43	25	13	18.1	9.5	23.6	13.9
Benchmark	n/a	0	10	47	31	11	n/a	n/a	n/a	n/a
Scenario 10										
True F_1		0.05	0.10	0.15	0.20	0.25				
True F_2		0.95	0.80	0.65	0.50	0.55				
Surv-CRM-12	0	0	1	17	48	35	10.9	8.5	26.2	18.3
TITE-BOIN-ET	0	1	4	21	34	40	22.3	4.7	34	8.8
Benchmark	n/a	0	1	11	69	19	n/a	n/a	n/a	n/a
Scenario 11										
True F_1		0.05	0.08	0.22	0.35	0.55				
True F_2		0.70	0.32	0.30	0.28	0.26				
Surv-CRM-12	0	0	4	71	25	0	12.3	10.6	13.8	25.8
TITE-BOIN-ET	0	1	23	55	18	3	14.2	9.6	15.1	16.9
Benchmark	n/a	0	29	61	10	0	n/a	n/a	n/a	n/a
Scenario 12										
True F_1		0.20	0.40	0.50	0.55	0.62				
True F_2		0.50	0.45	0.40	0.38	0.36				
Surv-CRM-12	4	66	29	0	0	0	22.1	13.7	21.2	22.9
TITE-BOIN-ET	5	75	18	2	0	1	21.7	13.1	19.8	20.3
Benchmark	n/a	90	10	0	0	0	n/a	n/a	n/a	n/a

Note: 10 000 simulated trials (Surv-CRM-12 design and benchmark) with $N=45$, $\pi_{\text{DLT}} = 0.25$ and $\epsilon_P = 0.1$. 10 000 simulated trials (TITE-BOIN-ET) with $\pi_{\text{DLT}} = 0.25$, $\pi_{\text{EFF}} = 0.95$, $\pi_{1\text{ DLT}} = 0.4$, $\pi_{\text{DLT}} = 1.4\pi_{\text{DLT}}$ and $\pi_{1\text{ EFF}} = 0.6\pi_{\text{EFF}}$. Correct selection results based on ODs are given in boldface. Abbreviation: n/a, not applicable.

5 | DISCUSSION

In this article, we proposed a novel one-stage Bayesian adaptive phase I/II design, named the Survival-CRM-12, for novel anti-cancer drugs such as immunotherapies and MTAs, with the purpose to identify the OD defined as the dose with acceptable risk of toxicity and a minimum risk of disease progression. Such a targeted dose, sometimes called the “best” dose¹⁷ appeared suitable for such early phase I/II trials, with the aim of maximizing the likelihood of efficacy while only ensuring safety constrains. Our proposed design allows continuous enrollment of patients, and handles competing risks, including the information of time-to-event outcomes using survival working models for right-censored endpoints. It allows the toxicity outcome to be delayed or unobserved due to competing progression within the observation window. Such an issue of late-onset outcome is common and important in the field of novel anticancer agents. Indeed, it has been reported that 57% of the grades 3-4 toxicities occur after the treatment cycle 1 in 36 clinical trials of molecularly targeted agents.³⁹ Moreover, disease progression is one of the main cause of premature discontinuation from a Phase I trial (approximately 70% of cases),⁴⁰ and should be managed other than by ad-hoc strategies of replacing patients. As demonstrated in the simulation study across a variety of realistic scenarios, and assessed by the benchmark, our proposed method appears an interesting approach in this setting, exhibiting desirable operating characteristics. Indeed, our design performs well in selecting the OD, as well as treating patient at safe doses during the clinical trial.

Studies have examined the impact of model specification on the performance of the dose-finding design and have shown the importance of selecting appropriate model parameters via simulation studies.^{41,42} Instead of arbitrarily setting initial guesses of the probabilities of DLT and disease progression, and prior variance for model parameters, we proposed a systematic calibration process based on previous literature.^{32,33} The calibration method which used jointly the half-width of the indifference interval both for the toxicity and progression working models, and the least informative prior variance had satisfactory operating characteristics. Jointly calibrating the initial guesses of the probabilities of toxicity and progression at each dose and the prior variance, simplified the calibration process as it reduced the number of parameters to be specified. This approach makes the Bayesian Surv-CRM-12 model-based design more accessible and easier to implement.

In this phase I/II setting, we chose to model disease progression instead of efficacy, highlighting its competition with the observation of toxicities. Indeed, unlike conventional chemotherapy that works by shrinking the tumor, immunotherapeutic agents often delay cancer progression and prolong survival without achieving rapid tumor shrinkage, and a certain percentage of patients who receive immunotherapy often achieve a long-term durable response.⁴³ Therefore, for some early phase immunotherapeutic trials, it is preferable to use as efficacy criterion the disease progression, rather than the conventional treatment response. Recently, Zhang et al⁴⁴ developed a dose finding phase I/II design accounting for the late-onset competing risk endpoint, relying on piecewise cause-specific hazard to characterize the competing risk outcomes. However, it treats late-onset outcomes as missing data while we incorporated them in our model, and requires the definition of a utility function for identifying the OD.

This study has some limitations. First, we modeled the cause-specific hazard for progression as a decreasing function of dose. Although a model-based approach allows the borrowing of information across dose levels, a potential limitation related to this choice of working model is that the Surv-CRM-12 tends to select higher doses among safe doses in cases of plateau scenarios for dose-progression relationship and it does not perform well for U-shaped dose-progression relationships. Nonparametric approaches which do not rely on any assumption on the dose-toxicity and dose-efficacy curves, may be useful in those situations,⁴⁵ as the TITE-BOIN-ET approach proposed. This approach is part of the model-assisted designs as its decision rule can be predetermined, but requires specification during the trial planning of several cutoffs by the clinician. A nonparametric approach was also considered with the benchmark method using the complete profiles of outcomes under all dose levels instead of the time-to-event outcome under the assigned dose in the Surv-CRM-12 or TITE-BOIN-ET designs. Nevertheless, the Surv-CRM-12 model-based design outperformed the benchmark in some scenarios, which is not observed with the TITE-BOIN-ET model-assisted design. In sparse settings such as phase I/II clinical, parametric working models may outperform the non-parametric benchmark in capturing the dose-response relationships since they allow sharing information across dose levels.³⁸ Second, the “optimal” dose level could have been defined as the dose with both the highest probability of efficacy and the lowest probability of toxicity. This could be the target in particularly frail populations, for example, in neonates and children. However, this assumes equal weighting of toxicity and efficacy, and at the time, the growing accelerated approvals of drugs, notably in oncology, suggests that the target of interest is mostly efficacy.⁴⁶ Nevertheless, we agree that some consensus regarding the target dose in such trials, should be reached in the future. This would allow a more relevant comparison of all the proposed designs.

Finally, with the emergence of anti-cancer molecularly-targeted agents and immunotherapy agents changing the landscape of dose-finding in oncology, it is essential to incorporate both late-onset toxicities and progression information

which are often collected as mutually exclusive events into the identification of OD. The proposed phase I/II Surv-CRM-12 design achieved the best performance in selecting the OD and allocating patients to the OD when trial discontinuation after disease progression within the observation window is likely to preclude the observation of DLTs in a non-negligible proportion of patients. It should be considered when designing dose-finding trials of targeted therapies with prolonged observation windows in advanced cancer patients where treatment discontinuations due to disease progression are likely to occur.

ACKNOWLEDGEMENTS

The authors thank the referees for their constructive comments and insightful suggestions on the manuscript. Shing M. Lee was supported by the National Center for Advancing Translational Sciences, National Institutes of Health, through Grant Number UL1TR001873.

CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

DATA AVAILABILITY STATEMENT

Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.

ORCID

Anaïs Andrillon  <https://orcid.org/0000-0003-4804-4611>

Shing M. Lee  <https://orcid.org/0000-0001-8413-6869>

Lucie Biard  <https://orcid.org/0000-0003-0660-1264>

REFERENCES

1. Postel-Vinay S, Gomez-Roca C, Molife LR, et al. Phase I trials of molecularly targeted agents: should we pay more attention to late toxicities? *J Clin Oncol*. 2011;29(13):1728-1735.
2. Mathijssen RHJ, Sparreboom A, Verweij J. Determining the optimal dose in the development of anticancer agents. *Nat Rev Clin Oncol*. 2014;11(5):272-281.
3. Wages NA, Chiuzan C, Panageas KS. Design considerations for early-phase clinical trials of immune-oncology agents. *J Immunother Cancer*. 2018;6(1):81.
4. Postel-Vinay S, Collette L, Paoletti X, et al. Towards new methods for the determination of dose limiting toxicities and the assessment of the recommended dose for further studies of molecularly targeted agents—Dose-Limiting Toxicity and toxicity assessment recommendation group for early trials of targeted therapies, an European organisation for research and treatment of cancer-led study. *Eur J Cancer*. 2014;50(12):2040-2049.
5. Cheung YK, Chappell R. Sequential designs for Phase I clinical trials with late-Onset toxicities. *Biometrics*. 2000;56(4):1177-1182.
6. Braun TM. Generalizing the TITE-CRM to adapt for early and late-onset toxicities. *Stat Med*. 2006;25(12):2071-2083.
7. Yuan Y, Yin G. Robust EM continual reassessment method in oncology dose finding. *J Am Stat Assoc*. 2011;106(495):818-831.
8. Liu S, Yin G, Yuan Y. Bayesian data augmentation dose finding continual reassessment method and delayed toxicity. *Ann Appl Stat*. 2013;7(4):1837-2457.
9. Yuan Y, Lin R, Li D, Nie L, Warren KE. Time-to-event Bayesian optimal interval design to accelerate Phase I trials. *Clin Cancer Res*. 2018;24(20):4921-4930.
10. Lin R, Yuan Y. Time-to-event model-assisted designs for dose-finding trials with delayed toxicity. *Biostatistics*. 2020;21(4):807-824.
11. Biard L, Cheng B, Manji GA, Lee SM. A simulation study of methods for handling disease progression in dose-finding clinical trials. *J Biopharm Stat*. 2020;31(2):156-167.
12. Andrillon A, Chevret S, Lee SM, Biard L. Dose-finding design and benchmark for a right censored endpoint. *J Biopharm Stat*. 2020;30(6):948-963.
13. Thall PF, Russell KE. A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in Phase I/II clinical trials. *Biometrics*. 1998;54(1):251-264.
14. O'Quigley J, Hughes MD, Fenton T. Dose-finding designs for HIV studies. *Biometrics*. 2001;57(4):1018-1029.
15. Braun TM. The bivariate continual reassessment method. Extending the CRM to Phase I trials of two competing outcomes. *Control Clin Trials*. 2002;23(3):240-256.
16. Yuan Y, Yin G. Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *J R Stat Soc Ser C Appl Stat*. 2009;58(5):719-736.
17. Wages NA, Tait C. Seamless Phase I/II adaptive design for oncology trials of molecularly targeted agents. *J Biopharm Stat*. 2015;25(5):903-920.
18. Lin R, Yin G. STEIN: a simple toxicity and efficacy interval design for seamless Phase I/II clinical trials. *Stat Med*. 2017;36(26):4106-4120.

19. Takeda K, Taguri M, Morita S. BOIN-ET: Bayesian optimal interval design for dose finding based on both efficacy and toxicity outcomes. *Pharm Stat.* 2018;17(4):383-395.
20. Lin R, Zhou Y, Yan F, Li D, Yuan Y. BOIN12: Bayesian optimal interval Phase I/II trial design for utility-based dose finding in immunotherapy and targeted therapies. *JCO Precis Oncol.* 2020;4:1393-1402.
21. Jin I, Liu S, Thall P, Yuan Y. Using data augmentation to facilitate conduct of Phase I-II clinical trials with delayed outcomes. *J Am Stat Assoc.* 2014;109(506):525-536.
22. Takeda K, Morita S, Taguri M. TITE-BOIN-ET: time-to-event Bayesian optimal interval design to accelerate dose-finding based on both efficacy and toxicity outcomes. *Pharm Stat.* 2020;19(3):335-349.
23. Biard L, Lee SM, Cheng B. Seamless Phase I/II design for novel anticancer agents with competing disease progression. *Stat Med.* 2021;40(21):4568-4581.
24. Tsiatis A. A nonidentifiability aspect of the problem of competing risks. *Proc Natl Acad Sci.* 1975;72(1):20-22.
25. Andersen PK, Abildstrom SZ, Rosthøj S. Competing risks as a multi-state model. *Stat Methods Med Res.* 2002;11(2):203-215.
26. Putter H, Fiocco M, Geskus RB. Tutorial in biostatistics: competing risks and multi-state models. *Stat Med.* 2006;26(11):2389-2430.
27. Jeong JH, Fine J. Direct parametric inference for the cumulative incidence function. *J R Stat Soc Ser C Appl Stat.* 2006;55(2):187-200.
28. Benichou J, Gail MH. Estimates of absolute cause-specific risk in cohort studies. *Biometrics.* 1990;46(3):813-826.
29. Ivanova A, Wang Y, Foster MC. The rapid enrollment design for Phase I clinical trials. *Stat Med.* 2016;35(15):2516-2124.
30. O'Quigley J, Paoletti X, Maccario J. Non-parametric optimal design in dose finding studies. *Biostatistics.* 2002;3(1):51-56.
31. Cheung YK. Simple benchmark for complex dose finding studies. *Biometrics.* 2014;70(2):389-397.
32. Lee SM, Cheung YK. Calibration of prior variance in the Bayesian continual reassessment method. *Stat Med.* 2011;30(17):2081-2089.
33. Lee SM, Cheung YK. Model calibration in the continual reassessment method. *Clin Trials.* 2009;6(3):227-238.
34. Wheeler GM, Mander AP, Bedding A, et al. How to design a dose-finding study using the continual reassessment method. *BMC Med Res Methodol.* 2019;19(1):1-15.
35. O'Quigley J, Shen LZ. Continual reassessment method: a likelihood approach. *Biometrics.* 1996;52(2):673-684.
36. Beyersmann J, Latouche A, Buchholz A, Schumacher M. Simulating competing risks data in survival analysis. *Stat Med.* 2009;28(6):956-971.
37. Gray RJ. A class of K-sample tests for comparing the cumulative incidence of a competing risk. *Ann Stat.* 1988;16(3):1141-1154.
38. Cheung YK. *Dose Finding by the Continual Reassessment Method.* Biostatistics Series. New York: Chapman & Hall/CRC Press; 2011.
39. Colin P, Micallef S, Delattre M, Mancini P, Parent E. Towards using a full spectrum of early clinical trial data: a retrospective analysis to compare potential longitudinal categorical models for molecular targeted therapies in oncology. *Stat Med.* 2015;34(22):2999-3016.
40. Olmos D, A'hern R, Marsoni S, et al. Patient selection for oncology Phase I trials: a multi-institutional study of prognostic factors. *J Clin Oncol.* 2012;30(9):996-1004.
41. Chevret S. The continual reassessment method in cancer Phase I clinical trials: a simulation study. *Stat Med.* 1993;12(12):1093-1108.
42. Paoletti X, Kramar A. A comparison of model choices for the continual reassessment method in Phase I cancer trials. *Stat Med.* 2009;28(24):3012-3028.
43. Topalian SL, Weiner GJ, Pardoll DM. Cancer immunotherapy comes of age. *J Clin Oncol.* 2011;29(36):4828-4836.
44. Zhang Y, Cao S, Zhang C, Jin I, Zang Y. A Bayesian adaptive Phase I/II clinical trial design with late-onset competing risk outcomes. *Biometrics.* 2020;77(3):796-808.
45. Cheung YK. Sequential implementation of stepwise procedures for identifying the maximum tolerated dose. *J Am Stat Assoc.* 2007;102(480):1448-1461.
46. Sachs RE, Donohue JM, Dusetzina SB. Accelerated approval - taking the FDA's concerns seriously. *N Engl J Med.* 2022;387(3):199-201.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Andrillon A, Chevret S, Lee SM, Biard L. Surv-CRM-12: A Bayesian phase I/II survival CRM for right-censored toxicity endpoints with competing disease progression. *Statistics in Medicine.* 2022;1-14. doi: 10.1002/sim.9591

4.3 Supplementary information

4.3.1 Benchmark for right-censored endpoints and competing risks

The algorithm used for the benchmark construction within a competing risks setting is presented in Table S1, based on S simulated trial outcomes for a given scenario of true cumulative probability of DLT of J doses, $\{p_{1j}\}_{j=1}^J$, and true cumulative probability of progression of J doses, $\{p_{2j}\}_{j=1}^J$, at time t^* . Complete time-to-event data with competing toxicity and progression in the setting of a dose-finding clinical trial was generated following Beyersmann et al. (2009)'s approach for each patient and each dose level.

Table S1: Algorithm to compute a benchmark with competing risks

- 1 Define $F_1(\cdot)$ and $F_2(\cdot)$, the cumulative distribution function (CDF) of time-to-toxicity and time-to-progression, respectively, and specify the observation window $[0, t^*]$. In this competing risks setting, the cumulative incidence of event k is defined as

$$F_k(t^*, \lambda_{1j}, \lambda_{2j}) = \frac{\lambda_{kj} \times (1 - \exp(-(\lambda_{1j} + \lambda_{2j}) \times t^*))}{\lambda_{1j} + \lambda_{2j}}$$

with the cause-specific hazard of failure from toxicity models as, $\lambda_{1j}(d_j, \beta_1) = \exp(d_j \times \exp(\beta_1))$ and the cause-specific hazard of failure from progression models as, $\lambda_{2j}(d_j, \beta_2) = \exp(-d_j \times \exp(\beta_2))$.

- 2 Derive instantaneous hazards λ_{kj} consistent with $F_k(\cdot)$ at each dose level according to pre-specified scenarios of toxicity and progression by computing the values of $(\lambda_{1j}, \lambda_{2j})$ for which
$$\begin{cases} p_{1j} = F_1(t^*, \lambda_{1j}, \lambda_{2j}) \\ p_{2j} = F_2(t^*, \lambda_{1j}, \lambda_{2j}) \end{cases}$$
, resulting in

$$\begin{cases} \lambda_{1j} = \frac{-\log(1-p_{1j}-p_{2j}) \times t^*}{1+p_{2j}/p_{1j}} \\ \lambda_{2j} = \lambda_{1j}/p_{1j} * p_{2j} \end{cases}$$

- 3 Generate a sequence of n patient profiles $\{u_i\}_{i=1}^n$ from the uniform distribution $\mathcal{U}(0, 1)$.
- 4 Obtain all-cause event times using $T_{ij} = Q^{-1}(u_i, \lambda_{1j} + \lambda_{2j})$ for each patient i and each dose level j , with $Q(\cdot)$, the CDF of the exponential distribution.
- 5 Given time T_{ij} , determine the type to event by a random draw from a Bernoulli distribution with probability $\frac{\lambda_{1j}}{\lambda_{1j} + \lambda_{2j}}$ for toxicity ($k = 1$; and otherwise progression, $k = 2$).

- 6** Apply administrative censoring if $T_i > t^*$, to obtain clinical trial data with pre-specified observation window t^* . $X_i = \min(T_i, t^*)$ and

$$Y_i = \begin{cases} 1 & \text{if } T_i \leq t^* \text{ and patient } i \text{ experienced toxicity} \\ 2 & \text{if } T_i \leq t^* \text{ and patient } i \text{ experienced treatment progression} \\ 0 & \text{otherwise} \end{cases}$$

- 7** From all (X_i, Y_i) , compute non parametric estimates of the cumulative incidences $\hat{F}_{1j}(t^*, \lambda_{1j}, \lambda_{2j})$ and $\hat{F}_{2j}(t^*, \lambda_{1j}, \lambda_{2j})$ applying the [Gray \(1988\)](#) estimator for each dose level j .
- 8** Estimate the correct dose, consistently with the dose-finding objective of the evaluated design.
- 9** Repeat steps 2-7 for $s = 1, \dots, S$ simulated trials.
- 10** Obtain a benchmark estimate of the probability of correct selection averaging results over the S simulated trials.

4.3.2 Model parameters calibration

4.3.2.1 Dose skeletons calibration

This section discusses the calibration of the initial guesses of toxicity and progression probabilities associated with the tested doses, so-called 'dose skeletons'. Scaled doses $\{x_{1j}\}_{j=1}^J$ and $\{x_{2j}\}_{j=1}^J$ are obtained via backward substitution by solving

$$\begin{cases} p_{10j} = F_1(t^*, x_{1j}, x_{2j}, \beta_{10}, \beta_{20}) \\ p_{20j} = F_2(t^*, x_{1j}, x_{2j}, \beta_{10}, \beta_{20}) \end{cases}$$

where (β_{10}, β_{20}) are the prior mean of the model parameters (β_1, β_2) and (p_{10j}, p_{20j}) are the initial guesses of the toxicity and progression probabilities associated with dose level j . To obtain those dose levels, we used a calibration procedure relied on both the toxicity and progression working models.

Design skeletons for toxicity and progression were calibrated using the indifference interval approach [[Lee and Cheung, 2009](#); [Cheung, 2011](#)], applied to the toxicity and progression working models estimating the cumulative incidences of toxicity and progression

at the end of the observation window t^* . Given θ_{DLT} and θ_{PROG} , respectively the target probability of toxicity and the expected probability of progression at the optimal dose ν at t^* , let $\mathcal{B}_k = [b_{k1}, b_{kJ+1}]$ be the parameter space for the working model of event k (i.e. $\beta_k \in \mathcal{B}_k$ with $k = 1$ for toxicity and $k = 2$ for progression). We defined the intervals for event k as follows:

$$\begin{cases} \mathcal{H}_{k1} = [b_{k1}, b_{k2}) \\ \mathcal{H}_{kj} = (b_{kj}, b_{kj+1}) & \text{for } j = 2, \dots, J-1 \\ \mathcal{H}_{kJ} = (b_{kJ}, b_{kJ+1}] \end{cases}$$

where $B_j = (b_{1j}, b_{2j})$ is solution for

$$F_k(t^*, X_{j-1}; B_j) + F_k(t^*, X_j; B_j) = \begin{cases} 2\theta_{DLT} & \text{for } k = 1 \\ 2\theta_{PROG} & \text{for } k = 2 \end{cases}$$

with $X_j = (x_{1j}, x_{2j})$ the dose labels. It means that with this parameter B_j , the estimated mean of probability of toxicity (progression) of doses level j and $j-1$ is equal to θ_{DLT} (θ_{PROG}). Following [Cheung and Chappell \(2002\)](#), for calibration purposes, we relied on targeting MTD and we defined indifference intervals for a given correct dose level ν as the interval of event probabilities associated with the neighbouring doses of the true MTD as these doses may be selected instead of the true MTD (ν) under large samples. The indifference for ν is denoted by:

$$\begin{cases} [NA, F_k(t^*, X_{\nu+1}; B_{\nu+1})] & \text{for } \nu = 1 \\ [F_k(t^*, X_{\nu-1}; B_\nu), F_k(t^*, X_{\nu+1}; B_{\nu+1})] & \text{for } \nu = 2, \dots, J-1 \\ [F_k(t^*, X_{\nu-1}; B_\nu), NA] & \text{for } \nu = J \end{cases}$$

To calculate initial guesses of probability of toxicity and progression from indifference intervals, the idea is first to obtain recursively scaled doses, $\{X_1, \dots, X_J\}$, given a pre-specified half-width indifference interval, that is δ . Indeed, we first calculated $X_\nu = (x_{1\nu}, x_{2\nu})$, the dose labels associated to the prior ν , via backward substitution since

$$\begin{cases} F_1(t^*, X_\nu; B_0) = \theta_{DLT} \\ F_2(t^*, X_\nu; B_0) = \theta_{PROG} \end{cases} \quad (4.1)$$

where $B_0 = (\beta_{10}, \beta_{20})$ the prior mean of $B = (\beta_1, \beta_2)$. If an indifference interval of length 2δ includes the desired dose levels for toxicity and progression, the remaining dose levels $X_{\nu-1} = (x_{1\nu-1}, x_{2\nu-1})$ and $X_{\nu+1} = (x_{1\nu+1}, x_{2\nu+1})$ can be obtained by solving the following equations:

$$F_k(t^*, X_{\nu-1}; B_\nu) + F_k(t^*, X_\nu; B_\nu) = \begin{cases} 2\theta_{DLT} & \text{for } k = 1 \\ 2\theta_{PROG} & \text{for } k = 2 \end{cases},$$

$$F_k(t^*, X_\nu; B_{\nu+1}) + F_k(t^*, X_{\nu+1}; B_{\nu+1}) = \begin{cases} 2\theta_{DLT} & \text{for } k = 1 \\ 2\theta_{PROG} & \text{for } k = 2 \end{cases}$$

which is based on the definition of \mathcal{B}_k , and

$$F_k(t^*, X_{\nu-1}; B_\nu) = \begin{cases} \theta_{DLT} - \delta & \text{for } k = 1 \\ \theta_{PROG} + \delta & \text{for } k = 2 \end{cases},$$

$$F_k(t^*, X_{\nu+1}; B_{\nu+1}) = \begin{cases} \theta_{DLT} + \delta & \text{for } k = 1 \\ \theta_{PROG} - \delta & \text{for } k = 2 \end{cases}$$

as $[F_k(t^*, X_{\nu-1}; B_\nu), F_k(t^*, X_{\nu+1}; B_{\nu+1})]$ is the indifference interval for ν of length 2δ .

To calculate the initial guesses of the probabilities of toxicity and progression, we used the following equations:

$$x_{1j-1} = \log \left(\frac{-\log \{1 - (\theta_{DLT} - \delta) - (\theta_{PROG} + \delta)\}}{(1 + \frac{\theta_{PROG} + \delta}{\theta_{DLT} - \delta}) \times t^*} \right) \times \frac{1}{\exp(b_{1j})},$$

$$x_{2j-1} = -\log \left(\frac{(\theta_{PROG} + \delta) \times \exp(\exp(b_{1j}) \times x_{1j-1})}{\theta_{DLT} - \delta} \right) \times \frac{1}{\exp(b_{2j})}$$

for $j = 2, \dots, \nu$,

$$x_{1j+1} = \log \left(\frac{-\log \{1 - (\theta_{DLT} + \delta) - (\theta_{PROG} - \delta)\}}{(1 + \frac{\theta_{PROG} - \delta}{\theta_{DLT} + \delta}) \times t^*} \right) \times \frac{1}{\exp(b_{1j+1})},$$

$$x_{2j+1} = -\log \left(\frac{(\theta_{PROG} - \delta) \times \exp(\exp(b_{1j+1}) \times x_{1j+1})}{\theta_{DLT} + \delta} \right) \times \frac{1}{\exp(b_{2j+1})}$$

for $j = \nu, \dots, J - 1$.

Selecting the initial guesses based on δ simplified the calibration process as it reduced the number of parameters from $2 \times J$ to only two parameters. This skeleton guarantees that the target probabilities of the DLT and progression lie in the specified indifference intervals.

We also explored another approach considering two different widths of indifference intervals, one related to toxicity and another related to progression contrary to the approach presented above considering a unique half-width of indifference interval δ . The contribution of using two parameters was limited in terms of PCS improving, and very consuming in terms of simulation. Indeed, both half-widths of indifference interval should be iterated from 0.01 to $0.6\theta_{DLT}$ in a discrete domain with a grid half-width of 0.01, resulting in a significant number of combinations to be explored.

4.3.2.2 Least informative variance

In this section, we discuss the specification of the prior standard deviation σ_{β_1} , for the model parameter β_1 in a Surv-CRM-12. Of note, for simplicity, we assumed the prior normal standard deviation for β_2 is equal to the standard deviation of β_1 . In conventional Bayesian inference, a large value of σ_{β_1} corresponds to a vague prior, whereas a small σ_{β_1} conveys strong prior beliefs. In the context of dose-finding, however, the notion of a non-informative prior is to be defined with reference to our belief about the correct, rather than that about the model parameter β_1 .

As presented bellow in Table SA, we computed $Pr(\beta_1 \in B_{1j})$, the probability that the dose j is the correct dose under the prior distribution for a given half-width indifference interval δ , θ_{DLT} , θ_{PROG} , number of candidate doses ($J = 5$) and expected optimal dose level (ν). $Pr(\beta_1 \in B_{1j}) = Pr(\nu_{\beta_1} = j)$ where ν_{β_1} is the model-based OD when β_1 is the true parameter value, so that $\nu_{\beta_1} = j$ if and only if β_1 belongs to the home set B_{1j} . Then, $Pr(\beta_1 \in B_{1j})$ could be calculated, as we assumed a normal formulation for the prior distribution of the model parameter β_1 , resulting in

$$Pr(\beta_1 \in B_{1j}) = \Phi(b_{1,j+1}/\sigma_{\beta_1}) - \Phi(b_{1,j}/\sigma_{\beta_1})$$

where $B_{1j} = (b_{1,j}, b_{1,j+1})$ as defined in previous section 4.3.2.1, and Φ is the cumulative distribution function of a standard normal. As illustrated in Table SA, and as shown by Lee and Cheung (2011), depending on the value of σ_1 , the distribution of ν_{β_1} can be unimodal with peak at dose level 3 when $\sigma_1 = 0.5$ for instance, uniform when $\sigma_1 = 0.8$ and gradually turns into a U-shaped as σ_1 increases (*e.g.* $\sigma_1 = 1.16$ or 1.5). For example, when

$\sigma_1 = 1.16$ (which corresponds to $\sigma_1 = \sqrt{1.34}$, *i.e.*, the conventional standard deviation used when conducting a trial with the CRM), the prior implies that either dose 1 or dose 5 is the correct dose with over 50% probability. Thus, a non-informative prior in terms of ν_{β_1} corresponds to the uniform distribution, that is, $Pr(\beta_1 \in B_{1j}) = 1/J$ for all j and J the total number of candidate doses. In summary, the vagueness of a prior in dose-finding setting can be determined by how close the probability of selection of each dose is to a uniform distribution.

Table SA: The distribution of ν_{β_1} with $\delta = 0.05$, for a trial with $\theta_{DLT} = 0.25$, $\theta_{PROG} = 0.4$, $J = 5$ dose levels and dose level 3 as the correct dose ($\nu = 3$).

j	$(bj, bj + 1)$	$Pr(\beta_1 \in B_{1j})$ with $\sigma_1 =$			
		0.5	0.8	1.16	1.5
1	$(-\infty, -0.67)$	0.09	0.20	0.28	0.33
2	$(-0.67, -0.22)$	0.24	0.19	0.14	0.11
3	$(-0.22, 0.22)$	0.34	0.22	0.15	0.12
4	$(-0.22, 0.66)$	0.24	0.19	0.14	0.11
5	$(0.66, \infty)$	0.09	0.20	0.28	0.33

4.3.2.3 Calibration process

In practice to calibrate δ and $\sigma_{\beta_1}^{LI}$, we performed a simulation study. We calculated $\sigma_{\beta_1}^{LI}$ for each value of δ and selected the $(\delta, \sigma_{\beta_1}^{LI})$ that maximized the average PCS across the calibration set. More specifically, we performed simulations using the Surv-CRM-12 under a set of calibration scenarios, such that the true probabilities of DLT and progression follow the plateau configuration, where $\mu_{1j} = F_{1L}$ and $\mu_{2j} = F_{2U}$ for $j < l$, $\mu_{1j} = F_{1U}$ and $\mu_{2j} = F_{2L}$ for $j > l$ and $\mu_{1j} = \theta_{DLT}$ and $\mu_{2j} = \theta_{PROG}$ for $j = l$ where $l = 1, \dots, J$, $F_{1L} = \theta_{DLT}/(2 - \theta_{DLT})$, $F_{1U} = 2\theta_{DLT}/(1 + \theta_{DLT})$, $F_{2U} = 2\theta_{PROG}/(1 + \theta_{PROG})$ and $F_{2L} = \theta_{PROG}/(2 - \theta_{PROG})$.

As a sensitivity analysis, calibration was also performed to obtain the dose skeleton for the simulation study only using the indifference intervals approach [Lee and Cheung, 2009] and given a normal prior with standard deviation fixed at $\sqrt{1.34}$, *i.e.*, the value commonly

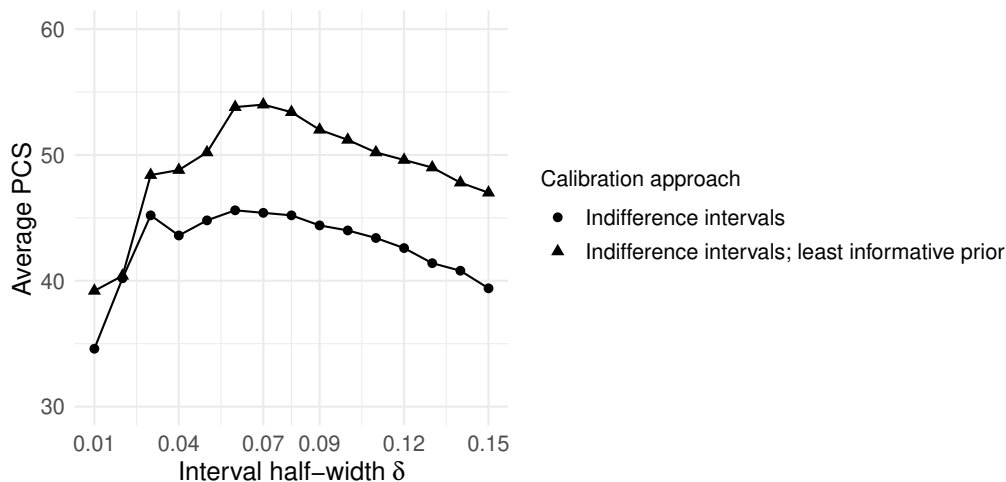


Figure S1: Average Probability of Correct Selection (PCS) of the optimal dose across the set of calibration scenarios, by calibration interval half-width according to the calibration approach, with $N = 45$.

used for the CRM [O’Quigley and Shen, 1996]. For the latest calibration approach, based on 2,000 independent replications, we set δ at 0.06, as providing the highest average PCS across the calibration scenarios. As an illustration, Figure S1 reported the average PCS by interval half-width, across the calibration scenarios, over 2,000 simulations according the calibration approach. Table S2 presents simulation results using the Surv-CRM-12 with dose skeleton calibration but without performing any calibration on the prior variance.

Table S2: Simulation results for Sc1 to Sc12 of the Surv-CRM-12 design; percent of discontinued trials for safety (P_{stop}); percent of selection (PS), number of overdose (No. OV); number of observed DLT (No. DLT); number of observed progressions (No. Prog) and number of patients treated with the true OD (No. OD) during the trial. 10,000 simulated trials with $N = 45$, $\pi_{DLT} = 0.25$, $\epsilon_P = 0.1$, and $\sigma_{\beta_1} = \sigma_{\beta_2} = \sqrt{1.34}$. Correct selection results based on ODs are given in boldface. (n/a: not applicable).

Scenario	P_{stop} (%)	PS by dose level (%)					No. OV	No. DLT	No. Prog	No. OD
		1	2	3	4	5				
Sc1	1	0	22	68	10	0	8.49	10.90	20.45	20.95
Sc2	0	0	21	67	12	0	7.61	10.30	9.74	18.92
Sc3	1	0	21	67	11	0	9.31	11.12	19.56	20.77
Sc4	0	1	1	3	67	28	13.39	10.02	13.45	20.13
Sc5	0	0	0	1	24	75	0.00	8.34	13.78	24.03
Sc6	1	0	21	73	5	0	6.60	10.45	20.39	34.13
Sc7	0	0	1	29	65	5	6.38	9.70	20.08	36.42
Sc8	0	0	1	25	66	8	25.98	9.74	16.05	12.85
Sc9	4	0	11	41	33	12	21.04	11.46	20.01	14.54
Sc10	3	0	1	10	34	53	19.65	8.99	26.28	13.94
Sc11	0	0	4	68	27	0	14.49	11.01	14.12	21.68
Sc12	18	59	22	1	0	0	18.27	12.70	21.52	26.73

4.3.3 Sensitivity analyses

4.3.3.1 Simulation with varying sample size

Figure S2 presents the PCS for the three methods according to the total sample size of the trial.

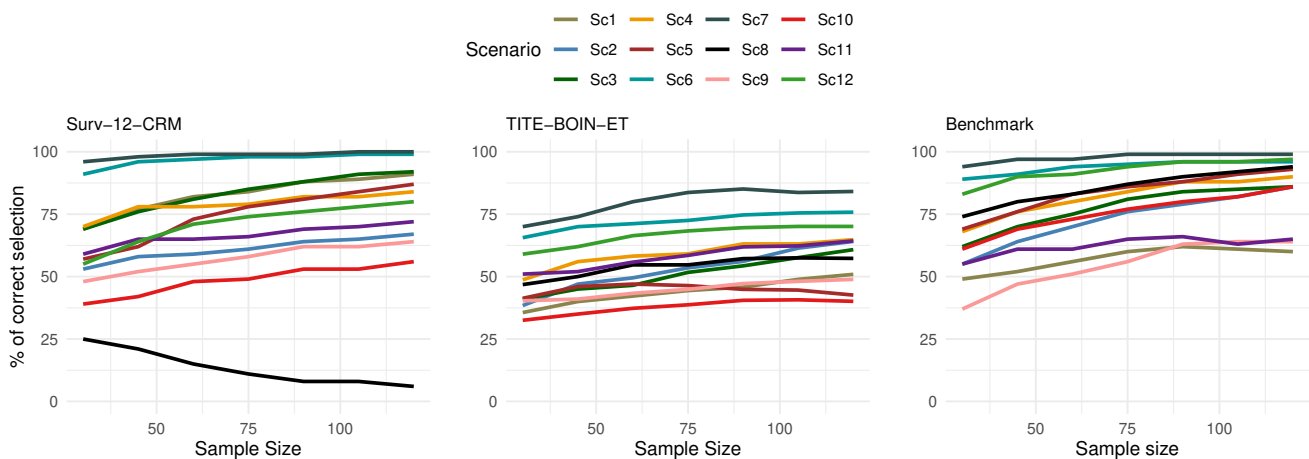


Figure S2: Percent of correct selection (PCS) with the Surv-CRM-12, the TiTE-BOIN-ET designs and the benchmark according to the total sample size N .

Table S3 reports simulation results with a total sample size of $N = 45$ and $N = 90$ depending on the number of patients in the optional toxicity centered stage, with the proposed Surv-CRM-12 design.

Table S3: Simulation results for Sc1 to Sc12 of the Surv-CRM-12 design according the sample size N , and the number of patients allocated to the first optional stage, rN ; percent of stopped trials for safety (P_{stop}); percent of selection (PS), number of overdose (No. OV); number of observed DLT (No. DLT); number of observed progression (No. Prog) and number of patients treated with the true OD (No. OD) during the trial. 10,000 simulated trials, $\pi_{DLT} = 0.25$ and $\epsilon_P = 0.1$. (n/a: not applicable).

	P_{stop} (%)	PS by dose level (%)					No. OV	No. DLT	No. Prog	No. OD
		1	2	3	4	5				
Scenario 1										
True F_1		0.08	0.14	0.25	0.40	0.57				
True F_2		0.50	0.48	0.45	0.43	0.40				
$N = 45; rN = 15$	0	0	18	73	9	0	7.09	10.97	20.35	25.40
$N = 45; rN = 23$	0	0	18	73	9	0	7.06	10.97	20.35	25.48
$N = 90; rN = 20$	0	0	11	85	4	0	10.32	21.99	40.68	61.21

$N = 90; rN = 45$	0	0	11	85	4	0	10.19	21.99	40.65	61.29
-------------------	---	---	----	----	---	---	-------	-------	-------	-------

Scenario 2True F_1 0.08 0.14 **0.25** 0.40, 0.57True F_2 0.30 0.25 **0.20** 0.15 0.10

$N = 45; rN = 15$	0	0	17	72	10	0	7.67	11.02	9.31	25.01
-------------------	---	---	----	----	----	---	------	-------	------	-------

$N = 45; rN = 23$	0	0	17	72	10	0	7.70	11.02	9.30	24.97
-------------------	---	---	----	----	----	---	------	-------	------	-------

$N = 90; rN = 20$	0	0	10	85	5	0	10.97	22.08	18.48	60.65
-------------------	---	---	----	----	---	---	-------	-------	-------	-------

$N = 90; rN = 45$	0	0	10	86	4	0	10.98	22.08	18.48	60.66
-------------------	---	---	----	----	---	---	-------	-------	-------	-------

Scenario 3True F_1 0.08 0.14 **0.25** 0.40 0.57True F_2 0.55 0.49 **0.43** 0.36 0.30

$N = 45; rN = 15$	0	0	17	73	10	0	7.56	11.07	19.44	25.28
-------------------	---	---	----	----	----	---	------	-------	-------	-------

$N = 45; rN = 23$	0	0	17	73	10	0	7.54	11.05	19.46	25.35
-------------------	---	---	----	----	----	---	------	-------	-------	-------

$N = 90; rN = 20$	0	0	10	85	4	0	10.98	22.15	38.68	61.17
-------------------	---	---	----	----	---	---	-------	-------	-------	-------

$N = 90; rN = 45$	0	0	10	85	4	0	10.99	22.15	38.69	61.17
-------------------	---	---	----	----	---	---	-------	-------	-------	-------

Scenario 4True F_1 0.05 0.10 0.15 **0.20** 0.25True F_2 0.95 0.75 0.50 **0.40** 0.30

$N = 45; rN = 15$	0	0	0	4	79	18	7.85	9.20	13.66	26.71
-------------------	---	---	---	---	----	----	------	------	-------	-------

$N = 45; rN = 23$	0	0	0	4	78	18	7.89	9.20	13.67	26.65
-------------------	---	---	---	---	----	----	------	------	-------	-------

$N = 90; rN = 20$	0	0	0	0	87	13	14.90	19.52	26.47	63.79
-------------------	---	---	---	---	----	----	-------	-------	-------	-------

$N = 90; rN = 45$	0	0	0	0	87	13	14.97	19.52	26.48	63.72
-------------------	---	---	---	---	----	----	-------	-------	-------	-------

Scenario 5True F_1 0.02 0.04 0.08 0.14 **0.25**True F_2 0.55 0.48 0.40 0.32 **0.25**

$N = 45; rN = 15$	0	0	0	1	34	65	0.00	7.72	14.17	18.71
-------------------	---	---	---	---	----	----	------	------	-------	-------

$N = 45; rN = 23$	0	0	0	1	34	65	0.00	7.72	14.18	18.68
-------------------	---	---	---	---	----	----	------	------	-------	-------

$N = 90; rN = 20$	0	0	0	0	21	79	0.00	17.63	26.34	50.95
-------------------	---	---	---	---	----	----	------	-------	-------	-------

$N = 90; rN = 45$	0	0	0	0	21	79	0.00	17.62	26.35	50.97
-------------------	---	---	---	---	----	----	------	-------	-------	-------

Scenario 6True F_1 0.08 **0.14** **0.25** 0.40 0.57True F_2 0.50 **0.45** **0.45** 0.45 0.40

$N = 45; rN = 15$	0	0	19	72	9	0	6.97	10.93	20.26	36.36
-------------------	---	---	----	----	---	---	------	-------	-------	-------

$N = 45; rN = 23$	0	0	19	73	8	0	6.87	10.94	20.24	36.46
-------------------	---	---	----	----	---	---	------	-------	-------	-------

$N = 90; rN = 20$	0	0	11	85	4	0	9.99	21.91	40.51	78.35
-------------------	---	---	----	----	---	---	------	-------	-------	-------

$N = 90; rN = 45$	0	0	11	85	4	0	9.99	21.92	40.51	78.35
-------------------	---	---	----	----	---	---	------	-------	-------	-------

Scenario 7True F_1 0.04 **0.08** **0.14** **0.25** 0.40True F_2 0.50 **0.45** **0.45** **0.45** 0.40

$N = 45; rN = 15$	0	0	0	26	67	7	3.80	9.40	20.14	40.07
$N = 45; rN = 23$	0	0	0	26	68	7	3.81	9.40	20.14	40.06
$N = 90; rN = 20$	0	0	0	15	82	3	6.09	19.96	40.23	82.78
$N = 90; rN = 45$	0	0	0	15	82	3	6.20	19.97	40.22	82.67

Scenario 8

True F_1	0.04	0.08	0.14	0.25	0.40
True F_2	0.50	0.35	0.27	0.35	0.50

$N = 45; rN = 15$	0	0	0	26	68	6	25.34	9.34	15.15	15.64
$N = 45; rN = 23$	0	0	0	26	68	5	25.36	9.33	15.16	15.61
$N = 90; rN = 20$	0	0	0	16	82	2	60.71	19.77	30.30	25.14
$N = 90; rN = 45$	0	0	0	16	82	2	60.85	19.77	30.31	25.00

Scenario 9

True F_1	0.05	0.15	0.25	0.30	0.35
True F_2	0.95	0.65	0.45	0.35	0.25

$N = 45; rN = 15$	0	0	14	55	26	5	14.06	10.94	20.99	20.81
$N = 45; rN = 23$	0	0	14	55	26	5	14.07	10.95	20.99	20.75
$N = 90; rN = 20$	0	0	8	64	25	2	27.66	22.43	40.61	47.55
$N = 90; rN = 45$	0	0	8	64	26	2	27.63	22.43	40.62	47.63

Scenario 10

True F_1	0.05	0.10	0.15	0.20	0.25
True F_2	0.95	0.80	0.65	0.50	0.55

$N = 45; rN = 15$	0	0	1	17	48	35	10.93	8.51	26.19	18.28
$N = 45; rN = 23$	0	0	1	17	48	35	10.91	8.50	26.20	18.33
$N = 90; rN = 20$	0	0	0	8	50	42	28.72	18.08	50.34	40.39
$N = 90; rN = 45$	0	0	0	8	50	42	28.77	18.08	50.33	40.38

Scenario 11

True F_1	0.05	0.08	0.22	0.35	0.55
True F_2	0.70	0.32	0.30	0.28	0.26

$N = 45; rN = 15$	0	0	4	71	25	0	12.64	10.75	13.78	25.74
$N = 45; rN = 23$	0	0	4	71	25	0	12.68	10.73	13.79	25.68
$N = 90; rN = 20$	0	0	1	80	19	0	23.01	21.83	27.10	59.53
$N = 90; rN = 45$	0	0	1	81	19	0	22.97	21.83	27.09	59.57

Scenario 12

True F_1	0.20	0.40	0.50	0.55	0.62
True F_2	0.50	0.45	0.40	0.38	0.36

$N = 45; rN = 15$	4	66	29	0	0	0	22.18	13.72	21.22	22.82
$N = 45; rN = 23$	4	66	30	0	0	0	22.13	13.71	21.22	22.87
$N = 90; rN = 20$	6	77	17	0	0	0	33.19	24.80	43.13	56.81
$N = 90; rN = 45$	6	77	17	0	0	0	33.28	24.80	43.15	56.72

4.3.3.2 Simulation of correlated time-to-events

To take into account the correlation between dose-limiting toxicity (DLT) and efficacy, we modeled times to DLT and progression using the Clayton model, as proposed by [Yuan and Yin \(2009\)](#). We thus generated correlated pairs of time-to-toxicity and time-to-progression (t_{DLT}, t_{PROG}) , for each patient. The joint density of the survival times of both outcomes, $f(t_{DLT}, t_{PROG})$, was defined as:



$$f(t_{DLT}, t_{PROG}) = \frac{c+1}{c} \{S_{DLT}(t_{DLT})^{-1/c} + S_{PROG}(t_{PROG})^{-1/c} - 1\}^{-c-2} \\ f_{DLT}(t_{DLT})f_{PROG}(t_{PROG})\{S_{DLT}(t_{DLT})S_{PROG}(t_{PROG})\}^{-1/c-1}$$

with $f_{DLT}(t_{DLT})$ and $f_{PROG}(t_{PROG})$ the marginal density functions for (t_{DLT}, t_{PROG}) , $S_{DLT}(t_{DLT})$ and $S_{PROG}(t_{PROG})$ the survival functions for t_{DLT} and t_{PROG} respectively, and c the correlation between times to toxicity and progression.


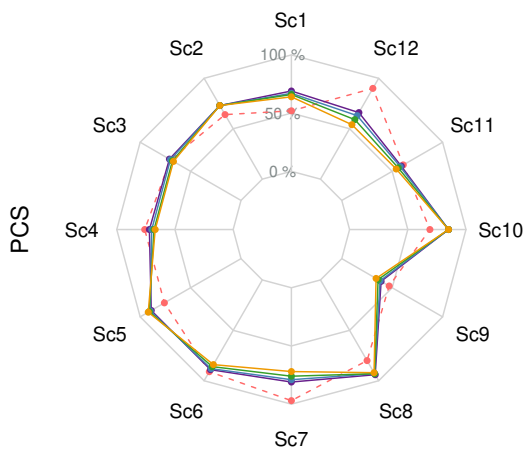
The bivariate random variable (t_{DLT}, t_{PROG}) was generated by first simulating t_{DLT} from its marginal distribution function $f_{DLT}(t_{DLT})$, then generating t_{PROG} from its conditional distribution $f_{PROG|DLT}(t_{PROG}|t_{DLT})$. In particular, we generated two independent random profiles (u_1, u_2) from the uniform distribution $\mathcal{U}(0, 1)$. By inverse transform sampling, we then obtained $t_{DLT} = S_{DLT}^{-1}(1 - u_1)$ and $t_{PROG} = S_{PROG|DLT}^{-1}(1 - u_2)$. We assumed exponential marginal survival functions, S_{DLT} and S_{PROG} .

Figure S3 reports the percents of correct and overdose selection depending on the correlation between time-to-toxicity and time-to-progression, with the proposed Surv-CRM-12 and the corresponding benchmark.

Surv-CRM-12



Benchmark



—●— $c = 1$
 —●— $c = 1.2$
 —●— $c = 1.5$
 —●— $c = 1.8$
 - - ● - - Not correlated

Figure S3: Percent of correct selection (PCS) and Percent of overdose selection (POS) with the Surv-CRM-12 and the benchmark using Clayton model for data generation according to the correlation between time-to-toxicity and time-to-progression. A small value of c represents a high correlation. When $c \rightarrow 0$, the correlation approaches 1 and, when $c \rightarrow \infty$, the correlation converges to 0.

Chapter 5

Discussion, perspectives, and conclusion

This thesis work was organized around several statistical dose-finding methods assessing adverse and therapeutic effects of targeted anti-cancer agents which have different administration schedules and modes of action than cytotoxic chemotherapies. The first part of the thesis consisted in a state-of-the-art, highlighting the main methods proposed for cancer dose-finding phase I and phase I/II trials. We attempted to emphasize the most relevant features of existing designs, which could help methodologists and clinicians in choosing the most appropriate approach to the clinical question investigated when planning such trials. Combining phases I and II of clinical development for conducting seamless designs is an attractive approach [Cuffe et al. 2014] and an active area of research as presented in chapter 2. It allows for simultaneously monitoring the toxicity and efficacy outcomes in a single trial to improve the dose level recommendation, reducing the costs, and speeding up the development process. Note that the state-of-the-art on phase I/II designs for a single agent in oncology is currently the object of a paper in progress.

At the start of this thesis work, no dose-finding design enabled to account for both administrative and informative right-censoring of the toxicity endpoints. Thus, the main objective of this thesis was to develop innovative adaptive dose-finding designs for clinical trials handling the right-censoring of the toxicity endpoints. More specifically, we aimed to develop designs satisfying safety requirements, with good operational characteristics whichever the possible locations of the correct dose(s), and outperforming the existing designs. We were also concerned with proposing simple, understandable, and easy-to-implement designs, not requiring too many tuning parameters to be specified, and no excessive computation times.

First, we developed the 'Survival-CRM' (Surv-CRM), a Bayesian adaptive phase I design that aims to identify the MTD, allowing administrative censoring that may occur either due to incomplete observation during the trial accrual or at the end of the trial. In this setting of administrative censoring, the benefit of the Surv-CRM compared to the TiTE-CRM was limited. However, in many settings, such as metastatic diseases, the length of the DLT assessment window results in an increased likelihood that patients may discontinue the trial prior to completing the full observation window because of death, disease progression, physician discretion or consent withdrawal. In this setting, *i.e.* when a substantial number of patients were expected to discontinue the trial within the toxicity window, the use of the TiTE-CRM was no longer appropriate. Consequently, accounting for the occurrence of such competing events was required. To answer those issues, we developed the 'informative Survival-CRM' (iSurv-CRM). Numerical studies showed that the proposed iSurv-CRM design improves the operating characteristics (notably in dose selection, patient allocation and safety) compared to the TiTE-CRM. Therefore, the iSurv-CRM should be considered when the expected level of trial discontinuations is high. We then proposed the 'Survival-CRM-12' (Surv-CRM-12) to tackle the issue of late-onset outcomes that compete with the observation of toxicity in phase I/II clinical trials. In this design, the occurrence of either the disease progression or the DLT prevents the observation of the other event, within the observation window. Numerical simulation studies confirmed the desirable performances of the proposed Surv-CRM-12 design.

A phase I/II clinical trial, the SALMA study, sponsored by the *Assistance Publique - Hôpitaux de Paris (AP-HP)* has been designed using the iSurv-CRM, and will be open to recruitment in autumn 2022 (Principal Investigator: R. ITZYKSON. NCT05580861). This clinical trial aims at assessing repurposed sulfasalazine (SSZ) in combination with standard-of-care induction therapy (idarubicin (IDA) and cytarabine (AraC)) in patients age above 60 years with newly diagnosed unfavorable Acute Myeloid Leukemias (AML). The main objective is to identify the MTD and to recommend a RP2D of SSZ out the seven following candidate dose levels¹. The trial was designed with a sample size of 20

¹Candidate Dose Levels (DL) of the SALMA study, administered orally:

- DL-2: 0.5 g, twice a day hence 1.0 g/d, days 1-8;

patients for dose-finding, followed by a dose expansion cohort at the estimated MTD. DLT was defined as any grade 3 or higher neutropenia, thrombocytopenia or non-hematological adverse events occurring within the 42 days (end of the induction cycle). The probability of DLT should not exceed 33%. About 50% of progressive diseases were expected by the end of the induction cycle for the first dose level and about 30% at the MTD. Thus, the dose-finding design relied on the iSurv-CRM² in this setting of high patient accrual rate relatively to the DLT observation window and accounting for right-censored toxicity in sequential dose assignment process. Moreover, this design allows accounting for observed progressive diseases during the safety assessment window.

In this thesis, we placed ourselves in a survival setting with only two competing risks (DLT and disease progression), while patients who experienced neither disease progression nor DLT at the end of the observation window were administratively censored. However, an essential assumption in the survival setting is that all patients in the study will eventually experience one of the events of interest. Therefore, patients are always considered at risk of having the event outside the observation window, even if they completed the trial follow-up. However, in cancer research, the event of interest may not occur for some individuals. Patients could be cured after therapy for many cancer types, even after a very long period of follow-up time. We propose that such a framework could be accommodated by using a cure model, a special case of survival models where a portion of subjects in the population never experience the event of interest [Berkson and P, 2006]. Such an approach has been used by Yuan and Yin (2009) to model time-to-efficacy, as described in section 2.2.4.2, chapter 2. More specifically, they considered a mixture cure rate model, considering two subpopulations: a susceptible subpopulation that possibly may respond and a non-susceptible subpopulation that will never respond to the treatment. Thus, modeling

-
- DL-1: 0.5 g, three times a day hence 1.5 g/d, days 1-8;
 - DL1: 0.5 g, three times a day hence 1.5 g/d, days 1-15;
 - DL2: 1 g, three times a day hence 3.0 g/d, days 1-15;
 - DL3: 1.5 g, three times a day hence 4.5 g/d, days 1-15;
 - DL4: 2.0 g, three times a day hence 6.0 g/d, days 1-15.
 - DL5: 2.0 g, four times a day hence 8.0 g/d, days 1-15.

²Additional safety rules were implemented, notably escalation to a new dose level j is allowed only if at least 3 patients have been treated at dose level $j - 1$ among all included patients so far.

the time-to-event data with a cured fraction appears to be an interesting way of modeling such dose-activity data and a promising way of increasing the amount of information on the treatment effects for targeted therapies.

In phase I/II clinical trials, we are concerned with finding a therapeutic dose of a new drug that maximizes the efficacy as well as controls the toxicity. In this work, we defined the optimal dose as the dose level(s) with the minimal progression risk and an acceptable risk of toxicity. Indeed, the recent and growing accelerated approvals of drugs, notably in oncology [Lythgoe et al. 2022], suggested that the target of interest was mostly efficacy. Thus, we did not impose any additional constraints on tolerance. However, given the dose-decreasing working model for progression, our method tends to select high doses among the set of safe doses in the case of plateau scenarios for the dose-progression relationship. By contrast, the objective could be to target the lowest safe dose that achieves the highest efficacy. Thus, in case of a plateau dose-efficacy relationship and an increasing dose-toxicity relationship, only the doses with the lowest probability of toxicity would be considered as optimal doses. This is in line with the recent project of the FDA (‘Optimus’) that recommends considering the full information in determining the optimal dose, including non-clinical data, PK/PD data, and early efficacy and safety data, and suggests selecting a safer (lower) dose as the optimal dose if the dose have sufficient efficacy. Then the ‘optimal’ dose level could be defined as the dose level with both the highest probability of efficacy and the lowest probability of toxicity. This is an appealing approach to define the correct dose. In all cases, some standardization in defining the targeted dose, and its naming, appears mandatory; it would improve the comparison of designs and the understanding of all those designs for practitioners.

We decided to model the dose-toxicity relationship as an increasing function of the dose as it appears reasonable to assume that the DLT rate increases with the dose level. By contrast, modeling the dose-progression relationship as a decreasing function of the dose could be a limitation of our proposals. Indeed, in the setting of immunotherapies and MTAs, due to the biological mechanism of the MTAs [Ellis, 2003; Morgan et al. 2003], the occurrence of disease progression usually decreases up to some dose where a plateau

in the risk of progression is likely observed. Under the competing risks framework, we could consider it fair to assume a monotonic decreasing cumulative incidence of disease progression. Indeed, in a situation of plateau dose-progression relationship, the observed estimated cumulative incidence of progression decreases when the observed estimated cumulative incidence of DLT increases. However, a future work perspective would consist in proposing an approach to capture a possible plateau dose-progression relationship as proposed by Yan, Tait, et al. (2019). Furthermore, when a U-shaped dose-progression relationship is expected, another functional form for the dose-progression relationship including a quadratic term could be of interest, as well as nonparametric approaches such that from Takeda, Morita, and Taguri (2020).

I would like to conclude this work by commenting on the challenges I have faced during my three years of clinical research at AP-HP. Namely, the promotion of innovative statistical methodologies, the enhancement of their accessibility, and the need for close collaboration between clinicians and statisticians.

A correct dose-finding study is the bedrock of medical advancement and the development of new drugs [Bhatt, 2010]. Phase I clinical trials are conducted on small samples, and the conclusions drawn are then applied to the entire target population. Therefore, statistical inference is a pivotal component of clinical trial development and has become increasingly recognized by regulatory agencies [Berry, 2006]. Model-based designs for phase I have been proposed more than 30 years ago [O’Quigley, Pepe, and Fisher, 1990], but, even though accessible in practice and statistically efficient, standard rule-based designs such as the ‘3 + 3’ design [Storer and Mood, 1989], well-known to be associated with poor statistical and operating characteristics, are still often used. This can be easily exemplified by a *PubMed* search. Using (‘3 + 3’ design) AND (‘phase 1’ OR ‘phase I’) as the search terms in September 2022, achieved 1,392 results, while switching to (‘CRM’ design OR ‘continual reassessment method’) AND (‘phase 1’ OR ‘phase I’) only retrieved 341. The transfer of statistical innovations into the clinical practice and medical literature, is a primary well-known issue [Altman and N, 1994]. If the institutional review board and the medical community have easily accepted rule-based designs, this is mostly because they

are very easy to understand and implement by practitioners. By contrast, there is a large consensus within the statistical community on the superiority of Bayesian methods as this facilitates the use of informative priors and innovative dose-finding clinical trial designs by directly incorporating available information in decision-making. This can be illustrated by a large supportive literature that has been accruing for over ten years: “*Bayesian clinical trials: no more excuses*” was the title of an editorial of *Clinical Trials* [Gönen, 2009], “*Bayesian adaptive clinical trials: a dream for statisticians only?*” asked Chevret (2012), a number of review have been published [O’Quigley and Zohar, 2006; Berry, 2006; Jaki et al. 2013; Love et al. 2017; Zhou, Murray, et al. 2018; Clertant, 2022], and several reporting guidance [Gaydos et al. 2006; LoRusso et al. 2010; Nie et al. 2016; Austin and Fine, 2017]. The American Society of Clinical Oncology (ASCO), the US FDA, and the European Medicines Agency (EMA) also recognized that novel adaptive designs using statistical models are of great importance and encourage sensible usage of them in phase I trials [US Food and Drug Administration, 2019; European Medicines Agency, 2020]. Therefore, all actors involved in clinical research should keep on promoting the use of efficient statistical methods.

To improve the diffusion of adequate and innovative model-based designs in early phase dose-finding trials, the implementation must be simple, with few parameters to be tuned. A number of complex modeling and non parsimonious frameworks have been proposed recently in this area of research. Developing methods requiring the specification of a large number of design parameters could be a real obstacle to their implementation in practice. It is particularly critical given the small sample sizes of early phase trials and the high sensitivity of the dose-finding methods to the specification of these parameters. Consequently, selecting appropriate parameters is a crucial process for the success of a design application that benefit both trial participants and future patients. For these reasons, we were concerned about providing a systematic approach to calibrate the main design parameters for each developed method. More specifically, we proposed guidelines and recommendations for the specification of the initial guesses of toxicity/progression probabilities and the variance of the prior distribution of the parameter(s) model(s). Proposing such calibration methods allowed both to obtain ‘good’ choices for the statistical compo-

nents and reduce the number of clinician-input parameters. Moreover, the development of complex dose-finding methods should not stop at the conception stage but designs should be ready for use. Making such designs accessible is not an easy task, but a real effort has been undertaken recently. Implementation barriers have been largely addressed thanks to the development of several user-friendly software through the Comprehensive R Archive Network (CRAN) or open-source software and web app and [MD Anderson Cancer Center, 2017; Pallmann et al. 2020]. The publication of method implementation tutorials and example-based guides also helps dissemination [Paoletti, Baron, et al. 2006; Wheeler et al. 2019; Sabanés Bové et al. 2019; Burnett et al. 2020; Lee, Wages, et al. 2021]. Of note, we developed a R Shiny application to promote the implementation of our proposed methods³. To conclude, proposing understandable, easy-to-implement and user-friendly methods that respond to contemporary clinical issues is of primary interest.

Finally, to facilitate the outreach of an innovative method, a change in trial design and logistics is required, with cooperative work between clinicians and statisticians from the planning phase up to the analysis of the data. First, methodologists should ensure that the selected dose-finding design meets the desired clinical objectives, experimental setup, and practical constraints of the trial. Indeed, a successful clinical trial is based not only on a positive clinical conclusion but also on the most appropriate methodology to answer the research question. At this planning stage, interdisciplinary collaboration is critical since the research questions need to be considered from multiple viewpoints. Furthermore, after the trial has been carried out, the joint publication of papers with clinicians and statisticians introducing such examples of trials using adequate dose-finding designs is also an essential step to improve the dissemination of a new method. Indeed, investigators tend to implement methods used in published literature in their disease area. As stated above, the large number of published phase I oncology trials based on the ‘3 + 3’ design is likely to encourage other investigators to use the ‘3 + 3’ design. A collaborative effort is thus mandatory to overcome the lack of understanding of statistical methods, concerns about obtaining regulatory approval, or even reluctance to break from traditional methods, in order to increase the use of Bayesian adaptive designs and conduct well-designed trials.

³The application is freely available at <https://tzw6pc-anais0andrillon.shinyapps.io/SurvCRM/> and will be finalized by the thesis defense.

In summary, the main objective of this thesis was to develop innovative adaptive Bayesian dose-finding designs for clinical trials handling right-censored toxicity endpoints. Thus, this thesis was an attempt to develop comprehensive survival models and competing risks framework to address complex issues in the context of dose-finding clinical trials for new targeted cancer drugs. We proposed well-performing methods that achieve desirable performances, are accessible, and are easy to implement in practice.

Bibliography

- Ahn, C. (1998). “An evaluation of phase I cancer clinical trial designs”. *Stat Med*; 17 (14), pp. 1537–1549 (cit. on pp. 15, 18).
- Altman, D. G. and G. S. N (1994). “Transfer of technology from statistical journals to the biomedical literature. Past trends and future predictions”. *JAMA*; 272 (2), pp. 129–132 (cit. on p. 126).
- Andrillon, A., S. Chevret, S. M. Lee, and L. Biard (2020). “Dose-finding design and benchmark for a right censored endpoint”. *J Biopharm Stat*; 23, pp. 1–16 (cit. on pp. xv, 9, 67).
- (2022). “Surv-CRM-12: A Bayesian phase I/II survival CRM for right-censored toxicity endpoints with competing disease progression”. *Stat Med*; n/a, pp. 1–14 (cit. on pp. xv, 9, 95).
- Association of the British Pharmaceutical Industry (ABPI) (2018). *Guidelines for Phase I clinical trials 2018 edition*. URL: <https://www.abpi.org.uk/publications/guidelines-for-phase-i-clinical-trials-2018-edition/> (cit. on pp. 2, 3).
- Austin, P. C. and J. P. Fine (2017). “Accounting for competing risks in randomized controlled trials: a review and recommendations for improvement”. *Stat Med*; 36 (8), pp. 1203–1209 (cit. on pp. 8, 127).
- Babb, J., A. Rogatko, and S. Zacks (1998). “Cancer phase I clinical trials: efficient dose escalation with overdose control”. *Stat Med*; 17 (10), pp. 1103–1120 (cit. on p. 24).
- Bekele, B. N. and Y. Shen (2005). “A Bayesian approach to jointly modeling toxicity and biomarker expression in a phase I/II dose-finding trial”. *Biometrics*; 61 (2), pp. 344–354 (cit. on p. 37).
- Berkson, J. and G. R. P (2006). “Survival Curve for Cancer Patients Following Treatment”. *J Am Stat Assoc*; 47 (259), pp. 501–15 (cit. on pp. 56, 124).

-
- Berry, D. A. (2006). “Bayesian clinical trials”. *Nat Rev Drug Discov*; 5 (1), pp. 27–36 (cit. on pp. 126, 127).
- Beyersmann, J., A. Latouche, A. Buchholz, and M. Schumacher (2009). “Simulating competing risks data in survival analysis”. *Stat Med*; 28 (6), pp. 956–971 (cit. on p. 110).
- Bhatt, A. (2010). “Evolution of clinical research: a history before and beyond james lind”. *Perspect Clin Res*; 1 (1), pp. 6–10 (cit. on p. 126).
- Biard, L., B. Cheng, G. A. Manji, and S. M. Lee (2021). “A simulation study of methods for handling disease progression in dose-finding clinical trials”. *J Biopharm Stat*; 31 (2), pp. 156–167 (cit. on pp. xii, 65).
- Biard, L., S. M. Lee, and B. Cheng (2021). “Seamless phase I/II design for novel anticancer agents with competing disease progression”. *Stat Med*; 40 (21), pp. 4568–4581 (cit. on p. 92).
- Borcoman, E. et al. (2019). “Novel patterns of response under immunotherapy”. *Ann Oncol*; 30 (3), pp. 385–396 (cit. on pp. xiv, 91).
- Borden, E. C. and A. Dowlati (2012). “Phase I trials of targeted anticancer drugs: a need to refocus”. *Nat Rev Drug Discov*; 11 (12), pp. 889–890 (cit. on p. 7).
- Braun, T. M. (2002). “The bivariate continual reassessment method. extending the CRM to phase I trials of two competing outcomes”. *Control Clin Trials*; 23 (3), pp. 240–256 (cit. on pp. 37–39, 44).
- Bril, G., R. Dykstra, C. Pillers, and T. Robertson (1984). “Algorithm AS 206: isotonic regression in two independent variables”. *J R Stat Soc Ser C Appl Stat*; 33 (3), pp. 352–357 (cit. on p. 29).
- Brock, K. et al. (2017). “Implementing the EffTox dose-finding design in the Matchpoint trial”. *BMC Med Res Methodol*; 17 (1), p. 112 (cit. on p. 43).
- Brock, K. (2016). *Snapshot of clintrials package used in 'Implementing the EffTox Dose-Finding Design in the Matchpoint Trial' manuscript* (cit. on p. 43).
- Burnett, T. et al. (2020). “Adding flexibility to clinical trial designs: an example-based guide to the practical use of adaptive designs”. *BMC Med*; 18 (1) (cit. on p. 128).
- Carlin, B. P. and T. A. Louis (2009). *Bayesian Methods for Data Analysis. Third Edition*. Chapman & Hall/CRC, Boca Raton, FL (cit. on p. 25).

- Cheung, Y. K. (2005). “Coherence principles in dose-finding studies”. *Biometrika*; 92 (4), pp. 863–873 (cit. on pp. 20, 23).
- (2011). *Dose Finding by the Continual Reassessment Method*. New-York: Chapman & Hall/CRC Biostatistics Series. Chap. 1 (cit. on pp. 22, 61, 111).
- (2014). “Simple benchmark for complex dose finding studies”. *Biometrics*; 70 (2), pp. 389–397 (cit. on pp. 62, 84).
- Cheung, Y. K. and R. Chappell (2000). “Sequential Designs for Phase I Clinical Trials with Late-Onset Toxicities”. *Biometrics*; 56 (4), pp. 1177–1182 (cit. on pp. xii, 7, 24, 34, 55, 59, 63).
- (2002). “A simple technique to evaluate model sensitivity in the continual reassessment method”. *Biometrics*; 58 (3), pp. 671–674 (cit. on p. 112).
- Chevret, S. (2006). “Statistical methods for dose-finding experiments” (cit. on p. 22).
- (2012). “Bayesian adaptive clinical trials: a dream for statisticians only?” *Stat Med*; 31 (11-12), pp. 1002–1013 (cit. on p. 127).
- Clertant, M. (2022). “Early-Phase Oncology Trials: Why So Many Designs?” *J Clin Oncol*; Aug, JCO.21.02493 (cit. on pp. 13, 25, 30, 127).
- Colin, P. et al. (2015). “Towards using a full spectrum of early clinical trial data: a retrospective analysis to compare potential longitudinal categorical models for molecular targeted therapies in oncology”. *Stat Med*; 34 (22), pp. 2999–3016 (cit. on p. 7).
- Conaway, M. R. and G. R. Petroni (2019a). “The impact of early-phase trial design in the drug development process”. *Clin Cancer Res*; 25 (2), pp. 819–827 (cit. on p. 4).
- (2019b). “The Role of Early-Phase Design-Response”. *Clin Cancer Res*; 25 (10), p. 3191 (cit. on p. 18).
- Couzin-Frankel, J. (2013). “Cancer immunotherapy”. *Science*; 342 (6165) (cit. on p. 6).
- Cuffe, R., D. Lawrence, A. Stone, and M. Vandemeulebroecke (2014). “When is a seamless study desirable? case studies from different pharmaceutical sponsors”. *Pharm Stat*; 13 (4), pp. 229–37 (cit. on p. 122).
- Dale, J. R. (1986). “Global cross-ratio models for bivariate, discrete, ordered responses”. *Biometrics*; 42 (4), pp. 909–917 (cit. on p. 43).

- Dixon, W. J. and A. M. Mood (1948). “A Method for Obtaining and Analyzing Sensitivity Data”. *Journal of the American Statistical Association*; 43 (241), pp. 109–126 (cit. on pp. 14, 16).
- Durham, S. D. and N. Flournoy (1994). “Random walks for quantile estimation”. *Statistical Decision Theory and Related Topics V*, pp. 467–476 (cit. on p. 16).
- Ellis, L. M. (2003). “Antiangiogenic therapy: more promise and, yet again, more questions”. *J Clin Oncol*; 21 (21) (cit. on p. 125).
- European Medicines Agency (2020). *Guideline on the clinical evaluation of anticancer medicinal products in man*. <https://www.ema.europa.eu/en/evaluation-anticancer-medicinal-products-man> (cit. on p. 127).
- Faries, D. (1994). “Practical modifications of the continual reassessment method for phase I cancer clinical trials”. *J Biopharm Stat*; 4 (2), pp. 147–164 (cit. on pp. 15, 22).
- Gatsonis, C. and J. B. Greenhouse (1992). “Bayesian methods for phase I clinical trials”. *Stat Med*; 11 (10), pp. 1377–1389 (cit. on p. 3).
- Gaydos, B. et al. (2006). “Adaptive Dose-Response Studies”. *Drug Information Journal*; 40 (4), pp. 451–461 (cit. on p. 127).
- Gezmu, M. and N. Flournoy (2006). “Group Up-and-Down designs for dose-finding”. *Journal of Statistical Planning and Inference*; 36 (1), pp. 1749–1764 (cit. on p. 16).
- Gönen, M. (2009). “Bayesian clinical trials: no more excuses”. *Clin Trials*; 6 (3), pp. 203–204 (cit. on p. 127).
- Goodman, S. N., M. L. Zahurak, and S. Piantadosi (1995). “Some practical improvements in the continual reassessment method for phase I studies”. *Stat Med*; 14 (11), pp. 1149–1161 (cit. on p. 22).
- Gray, R. J. (1988). “A class of K-sample tests for comparing the cumulative incidence of a competing risk”. *Annals of Statistics*; 16 (3), pp. 1141–1154 (cit. on p. 111).
- Guo, B. and Y. Yuan (2015). “A Bayesian dose-finding design for phase I/II clinical trials with nonignorable dropouts”. *Stat Med*; 34 (10), pp. 1721–1732 (cit. on p. 55).
- Guo, W., S. J. Wang, et al. (2017). “A Bayesian interval dose-finding design addressing Ockham’s razor: mTPI-2”. *Contemp Clin Trials*; 58, pp. 23–33 (cit. on pp. 28, 29).
- Han, Y., D. Liu, and L. Li (2020). “PD-1/PD-L1 pathway: current researches in cancer”. *Am J Cancer Res*; 10 (3), pp. 727–742 (cit. on p. 6).

- Harrington, J. A. et al. (2013). “Adaptive designs for dual-agent phase I dose-escalation studies”. *Nature Reviews Clinical Oncology*; 10 (5), pp. 277–288 (cit. on p. 15).
- He, W., J. Liu, B. Binkowitz, and H. Quan (2006). “A model-based approach in the estimation of the maximum tolerated dose in phase I cancer clinical trials”. *Biostatistics*; 25 (12), pp. 2027–2042 (cit. on p. 15).
- Heyd, J. M. and B. P. Carlin (1999). “Adaptive design improvements in the continual reassessment method for phase I studies”. *Stat Med*; 18 (11), pp. 1307–1321 (cit. on p. 22).
- Hinkson, I. V., B. Madej, and E. A. Stahlberg (2020). “Accelerating therapeutics for opportunities in medicine: a paradigm shift in drug discovery”. *Frontiers in Pharmacology*; 11 (cit. on p. 2).
- Hoering, A., M. LeBlanc, and J. Crowley (2011). “Seamless phase I-II trial design for assessing toxicity and efficacy for targeted agents”. *Clin Cancer Res*; 17 (4), pp. 640–646 (cit. on p. 7).
- Horton, B. J., N. A. Wages, and M. R. Conaway (2017). “Performance of toxicity probability interval based designs in contrast to the continual reassessment method”. *Stat Med*; 36 (2), pp. 291–300 (cit. on p. 30).
- Houede, N. et al. (2010). “Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials”. *Biometrics*; 66 (2), pp. 532–540 (cit. on p. 37).
- Iasonos, A., N. A. Wages, et al. (2016). “Dimension of model parameter space and operating characteristics in adaptive dose-finding studies”. *Stat Med*; 35 (21), pp. 3760–3775 (cit. on p. 25).
- Iasonos, A., A. S. Wilton, et al. (2008). “A comprehensive comparison of the continual reassessment method to the standard 3 + 3 dose escalation scheme in Phase I dose-finding studies”. *Clin Trials*; 5 (5), pp. 465–477 (cit. on pp. 15, 18, 22).
- Ivanova, A. (2003). “A play-the-winner-type urn design with reduced variability”. *Metrika*; 58, pp. 1–13 (cit. on pp. 37–40).
- (2006). “Escalation, group and A + B designs for dose-finding trials”. *Stat Med*; 25 (21), pp. 3668–3678 (cit. on p. 16).

-
- Ivanova, A., N. Flournoy, and Y. Chung (2007). “Cumulative cohort design for dose finding”. *Journal of Statistical Planning and Inference*; 137, pp. 2316–2327 (cit. on pp. 17, 30).
- Jaki, T., S. Clive, and C. J. Weir (2013). “Principles of dose finding studies in cancer: a comparison of trial designs”. *Cancer Chemother Pharmacol*; 71 (5), pp. 1107–1114 (cit. on p. 127).
- Ji, Y., Y. Li, and B. Bekele (2007). “Dose-finding in phase I clinical trials based on toxicity probability intervals”. *Clin Trials*; 4 (3), pp. 235–244 (cit. on p. 28).
- Ji, Y., P. Liu, Y. Li, and B. N. Bekele (2010). “A modified toxicity probability interval method for dose-finding trials”. *Contemp Clin Trials*; 7 (6), pp. 653–663 (cit. on pp. 28, 49).
- Jin, I. H., S. Liu, P. F. Thall, and Y. Yuan (2014). “Using data augmentation to facilitate conduct of phase I-II clinical trials with delayed outcomes”. *J Am Stat Assoc*; 109 (506), pp. 525–536 (cit. on pp. 37–39, 58).
- June, C. H., J. T. Warshauer, and J. A. Bluestone (2017). “Is autoimmunity the Achilles’ heel of cancer immunotherapy?” *Nat Med*; 23 (5) (cit. on p. 7).
- Kaplan, E. L. and P. Meier (1958). “Nonparametric Estimation from Incomplete Observations”. *Journal of the American Statistical Association*; 53 (282), pp. 457–481 (cit. on p. 85).
- Koopmeiners, J. S. and J. Modiano (2014). “A Bayesian adaptive Phase I-II clinical trial for evaluating efficacy and toxicity with delayed outcomes”. *Clin Trials*; 11 (1), pp. 38–48 (cit. on pp. 9, 38, 39, 57).
- Kpamegan, E. E. and N. Flournoy (2001). “An Optimizing Up-and-Down Design”. In: *Optimum Design 2000*. Ed. by A. Atkinson, B. Bogacka, and A. Zhigljavsky. Boston, MA: Springer US, pp. 211–224 (cit. on pp. 37–40).
- Le Tourneau, C., V. Diéras, et al. (2010). “Current challenges for the early clinical development of anticancer drugs in the era of molecularly targeted agents”. *Target Oncol*; 5 (1) (cit. on p. 6).
- Le Tourneau, C., H. K. Gan, A. R. Razak, and X. Paoletti (2012). “Efficiency of new dose escalation designs in dose-finding phase I trials of molecularly targeted agents”. *PLoS One*; 7 (12), e510396 (cit. on pp. 6, 15).

- Le Tourneau, C., J. J. Lee, and L. L. Siu (2009). “Dose escalation methods in phase I cancer clinical trials”. *J Natl Cancer Inst*; 101 (10), pp. 708–720 (cit. on pp. 5, 15, 18, 22).
- Le Tourneau, C., A. R. Razak, et al. (2011). “Heterogeneity in the definition of dose-limiting toxicity in phase I cancer clinical trials of molecularly targeted agents: a review of the literature”. *Eur J Cancer*; 47 (10) (cit. on p. 6).
- Lee, S. M. and Y. K. Cheung (2009). “Model Calibration in the Continual Reassessment Method”. *Clin Trials*; 6 (3), pp. 227–38 (cit. on pp. 111, 115).
- (2011a). “Calibration of prior variance in the Bayesian continual reassessment method”. *Stat Med*; 30 (17), pp. 2081–2089 (cit. on pp. 24, 114).
- (2011b). “Continual reassessment method with multiple toxicity constraints”. *Biostatistics*; 12 (2), pp. 386–398 (cit. on p. 25).
- Lee, S. M., X. Lu, and B. Cheng (2020). “Incorporating patient-reported outcomes in dose-finding clinical trials”. *Stat Med*; 39 (3), pp. 310–325 (cit. on p. 26).
- Lee, S. M., N. A. Wages, K. A. Goodman, and A. C. Lockhart (2021). “Designing Dose-Finding Phase I Clinical Trials: Top 10 Questions That Should Be Discussed With Your Statistician”. *JCO Precision Oncology*; 5, pp. 317–324 (cit. on p. 128).
- Lei, X., Y. Yuan, and G. Yin (2011). “Bayesian phase II adaptive randomization by jointly modeling time-to-event efficacy and binary toxicity”. *Lifetime Data Anal*; 17 (1), pp. 156–174 (cit. on p. 55).
- Li, D. H., J. B. Whitmore, W. Guo, and Y. Ji (2017). “Toxicity and efficacy probability interval design for phase I adoptive cell therapy dose-finding clinical trials”. *Clin. Cancer Res*; 23 (1), pp. 13–20 (cit. on pp. 37–39, 49).
- Lin, R. and G. Yin (2017). “STEIN: A simple toxicity and efficacy interval design for seamless phase I/II clinical trials”. *Stat Med*; 36 (26), pp. 4106–4120 (cit. on pp. 9, 38, 39, 49).
- Lin, R. and Y. Yuan (2020). “Time-to-event model-assisted designs for dose-finding trials with delayed toxicity”. *Biostatistics*; 21 (4) (cit. on pp. 34, 36, 59, 63).
- Lin, R., Y. Zhou, et al. (2020). “BOIN12: Bayesian Optimal Interval Phase I/II Trial Design for Utility-Based Dose Finding in Immunotherapy and Targeted Therapies”. *JCO Precis Oncol*; 4, PO.20.00257 (cit. on pp. 37–39, 52).

- Lin, Y. and W. J. Shih (2001). “Statistical properties of the traditional algorithm-based designs for phase I cancer clinical trials”. *Biostatistics*; 2 (2), pp. 203–215 (cit. on p. 16).
- Liu, S., C. Cai, and J. Ning (2013). “Up-and-Down designs for phase I clinical trials”. *Contemp Clin Trials*; 36 (1), pp. 218–227 (cit. on p. 18).
- Liu, S., B. Guo, and Y. Yuan (2018). “A Bayesian Phase I/II Trial Design for Immunotherapy”. *J Am Stat Assoc*; 113 (523), pp. 1016–1027 (cit. on p. 37).
- Liu, S. and V. E. Johnson (2016). “A robust Bayesian dose-finding design for phase I/II clinical trials”. *Biostatistics*; 17 (2), pp. 249–263 (cit. on pp. 9, 37, 39, 58).
- Liu, S., G. Yin, and Y. Yuan (2013). “Bayesian data augmentation dose finding with continual reassessment method and delayed toxicity”. *The Annals of Applied Statistics*; 7 (4), pp. 2138–2215 (cit. on pp. 7, 34, 35, 63).
- LoRusso, P., S. Boerner, and L. Seymour (2010). “An overview of the optimal planning, design, and conduct of phase I studies of new therapeutics”. *Clin Cancer Res*; 16 (6), pp. 1710–1718 (cit. on p. 127).
- Love, S. B. et al. (2017). “Embracing model-based designs for dose-finding trials”. *Br J Cancer*; 117 (3), pp. 332–339 (cit. on p. 127).
- Lythgoe, M. P. et al. (2022). “Cancer Therapy Approval Timings, Review Speed, and Publication of Pivotal Registration Trials in the US and Europe, 2010-2019”. *JAMA Netw Open*; 5 (6), e2216183 (cit. on p. 125).
- Mauguen, A., M. C. Le Deley, and S. Zohar (2011). “Dose-finding approach for dose escalation with overdose control considering incomplete observations”. *Stat Med*; 30 (13), pp. 1584–1594 (cit. on pp. 7, 34, 35, 63).
- MD Anderson Cancer Center (2017). *Biostatistics Software*. <https://biostatistics.mdanderson.org/SoftwareDownload> (cit. on p. 128).
- Mellman, I., G. Coukos, and G. Dranoff (2011). “Cancer immunotherapy comes of age”. *Nature*; 480 (7378), pp. 480–489 (cit. on p. 6).
- Morgan, B. et al. (2003). “Antiangiogenic therapy: more promise and, yet again, more questions”. *J Clin Oncol*; 21 (21) (cit. on p. 125).
- Mozgunov, P., T. Jaki, and X. Paoletti (2020). “A benchmark for dose finding studies with continuous outcomes”. *Biostatistics*; 21 (2), pp. 189–201 (cit. on p. 62).

- Mozgunov, P., X. Paoletti, and T. Jaki (2021). “A benchmark for dose-finding studies with unknown ordering”. *Biostatistics*; 01 (cit. on p. 62).
- Neuenschwander, B., M. Branson, and T. Gsponer (2009). “Critical aspects of the Bayesian approach to phase I cancer trials”. *Statistics in medicine*; 27 (13), pp. 2420–2439 (cit. on p. 25).
- Nie, L. et al. (2016). “Rendering the 3+ 3 design to rest: more efficient approaches to oncology dose-finding trials in the era of targeted therapy”. *Clin Cancer Res*; 22 (11), pp. 2623–2629 (cit. on p. 127).
- O’Quigley, J. and S. Chevret (1991). “Methods for dose finding studies in cancer clinical trials: A review and results of a Monte Carlo study”. *Stat Med*; 10 (11), pp. 1647–1664 (cit. on p. 15).
- O’Quigley, J., X. Paoletti, and J. Maccario (2002). “Non-parametric optimal design in dose finding studies”. *Biostatistics*; 3 (1), pp. 51–56 (cit. on pp. xiv, 60, 61, 84).
- O’Quigley, J. and L. Z. Shen (1996). “Continual reassessment method: a likelihood approach”. *Biometrics*; 52 (2), pp. 673–684 (cit. on pp. 20, 23, 116).
- O’Quigley, J. and S. Zohar (2006). “Experimental designs for phase I and phase I/II dose-finding studies”. *Br J Cancer*; 94 (5), pp. 609–613 (cit. on p. 127).
- O’Quigley, J., M. D. Hughes, and T. Fenton (2001). “Dose-Finding Designs for HIV Studies”. *Biometrics*; 57 (4) (cit. on pp. 9, 37–39, 44).
- O’Quigley, J., L. Pepe, and L. Fisher (1990). “Continual reassessment method : a practical design for phase I clinical trials in cancer”. *Biometrics*; 46 (1), pp. 33–48 (cit. on pp. x, 5, 13, 20, 22, 126).
- O’Quigley, J. and E. Reiner (1998). “ A stopping rule for the continual reassessment method”. *Stat Med*; 85 (3), pp. 741–748 (cit. on p. 22).
- Olmos, D. et al. (2012). “ Patient selection for oncology phase I trials: a multi-institutional study of prognostic factors”. *J Clin Oncol*; 30 (9) (cit. on p. 8).
- Omura, G. A. (2003). “Modified Fibonacci search”. *J Clin Oncol*; 21 (16), p. 3177 (cit. on p. 4).
- (2006). “Phase 1 dose-finding trials and fibonacci”. *Clin Cancer Res*; 12 (1), p. 321 (cit. on p. 4).

-
- Oron, A. P., D. Azriel, and P. Hoff (2011). “Dose-finding designs: The role of convergence properties”. *Int. J. Biostat*; 7 (cit. on p. 18).
- Pallmann, P. et al. (2020). “Designing and evaluating dose-escalation studies made easy: The MoDEsT web app”. *Clin Trials*; 17 (2) (cit. on p. 128).
- Paoletti, X., B. Baron, et al. (2006). “Using the continual reassessment method: lessons learned from an EORTC phase I dose finding study”. *Eur J Cancer*; 42 (10), pp. 1362–1368 (cit. on p. 128).
- Paoletti, X., M. Ezzalfani, and C. Le Tourneau (2015). “Statistical controversies in clinical research: Requiem for the 3 + 3 design for phase I trials”. *Ann Oncol*; 26 (9), pp. 1808–1812 (cit. on p. 18).
- Paoletti, X. and A. Kramar (2009). “A comparison of model choices for the Continual Reassessment Method in phase I cancer trials”. *Stat Med*; 28 (24), pp. 3012–3028 (cit. on p. 24).
- Paoletti, X., C. Le Tourneau, et al. (2014). “Defining dose-limiting toxicity for phase 1 trials of molecularly targeted agents: results of a DLT-TARGETT international survey.” *Eur J Cancer*; 50 (12), pp. 2050–2056 (cit. on p. 7).
- Paoletti, X., J. O’Quigley, and J. Maccario (2004). “Design efficiency in dose finding studies”. *Computational Statistics & Data Analysis*; 45 (2), pp. 197–214 (cit. on p. 61).
- Paoletti, X., S. Postel-Vinay, et al. (2010). “Méthodes pour la recherche de dose de traitements ciblés : nouvelles pistes. [Dose finding methods for targeted agents: new perspectives]”. *Bull Cancer*; 97 (12) (cit. on p. 7).
- Penel, N. et al. (2009). ““Classical 3 + 3 design” versus “accelerated titration designs”: analysis of 270 phase 1 trials investigating anti-cancer agents”. *PLoS One*; 27 (6), pp. 552–556 (cit. on p. 15).
- Polley, M. C. (2011). “Practical modifications to the Time-to-Event Continual Reassessment Method for phase I cancer trials with fast patient accrual and late-onset toxicities”. *Stat Med*; 30 (17), pp. 2130–2143 (cit. on p. 87).
- Postel-Vinay, S., H. T. Arkenau, et al. (2009). “Clinical benefit in Phase-I trials of novel molecularly targeted agents: does dose matter? ” *Br J Cancer*; 100 (9) (cit. on p. 6).

- Postel-Vinay, S., C. Gomez-Roca, et al. (2011). “Phase I trials of molecularly targeted agents: should we pay more attention to late toxicities?” *J Clin Oncol*; 29 (13) (cit. on p. 7).
- Prasad, V. and S. Mailankody (2017). “Research and Development Spending to Bring a Single Cancer Drug to Market and Revenues After Approval”. *JAMA Intern Med*; 177 (11), pp. 1569–1575 (cit. on p. 2).
- Pratt, V. M. et al. (2012). “Medical Genetics Summaries”. *Bethesda (MD): National Center for Biotechnology Information (US)* (cit. on p. 6).
- Printz, C. (2015). “Failure rate: Why many cancer drugs don’t receive FDA approval, and what can be done about it”. *Cancer*; 121 (10), pp. 1529–1530 (cit. on p. 4).
- Putter, H., M. Fiocco, and R. B. Geskus (2006). “Tutorial in biostatistics: competing risks and multi-state models”. *Stat Med*; 26 (11), pp. 2389–2430 (cit. on pp. xiii, 65, 91).
- Ratain, M. J., R. Mick, R. L. Schilsky, and M. Siegler (1993). “Statistical and ethical issues in the design and conduct of phase I and II clinical trials of new anticancer agents”. *J Natl Cancer Inst*; 85 (20), pp. 1637–1643 (cit. on p. 5).
- Roberts, T. G. et al. (2004). “Trends in the risks and benefits to patients with cancer participating in phase 1 clinical trial”. *JAMA*; 292 (17) (cit. on p. 3).
- Rogatko, A. et al. (2007). “Translation of innovative designs into phase I trials”. *J Clin Oncol*; 25 (31), pp. 4982–4986 (cit. on pp. 4, 5, 15, 18).
- Rosenberger, W. F. and L. M. Haines (2002). “Competing designs for phase I clinical trials: A review”. *Stat Med*; 21 (18), pp. 2757–2770 (cit. on p. 20).
- Rubin, D. B. (1974). “Estimating causal effects of treatments in randomized and non-randomized studies”. *Journal of Educational Psychology*; 66 (5), pp. 688–701 (cit. on p. 61).
- Sabanés Bové, D., W. Y. Yeung, G. Palermo, and T. Jaki (2019). “Model-Based Dose Escalation Designs in R with crmPack”. *Journal of Statistical Software*; 89 (10), pp. 1–22 (cit. on p. 128).
- Shen, L. Z. and J. O’Quigley (1996). “Consistency of continual reassessment method under model misspecification”. *Biometrika*; 83, pp. 395–405 (cit. on pp. 23, 24).
- Simon, R. et al. (1997). “Accelerated titration designs for phase I clinical trials in oncology”. *Journal of the National Cancer Institute*; 2 (2), pp. 203–215 (cit. on p. 16).

-
- Storer, B. E. and A. M. Mood (1989). “Design and analysis of phase I clinical trials”. *Biometrics*; 45 (3), pp. 925–937 (cit. on pp. 14, 126).
- Sun, D., W. Gao, H. Hu, and S. Zhou (2022). “Why 90% of clinical drug development fails and how to improve it?” *Acta Pharmaceutica Sinica B* (cit. on p. 1).
- Takeda, K., S. Morita, and M. Taguri (2020). “TITE-BOIN-ET: Time-to-event Bayesian optimal interval design to accelerate dose-finding based on both efficacy and toxicity outcomes”. *Pharm Stat*; 19 (3), pp. 335–349 (cit. on pp. xiv, 38, 39, 59, 93, 126).
- Takeda, K., M. Taguri, and S. Morita (2018). “BOIN-ET: Bayesian optimal interval design for dose finding based on both efficacy and toxicity outcomes”. *Pharmaceutical Statistics*; 17 (4), pp. 383–395 (cit. on pp. 9, 37–39, 50).
- Thall, P. F. and J. D. Cook (2004). “Dose-finding based on efficacy-toxicity trade-offs”. *Biometrics*; 60 (3) (cit. on pp. 9, 37–39, 41, 44, 57, 58).
- Thall, P. F. and K. E. Russell (1998). “A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials”. *Biometrics*; 54 (1) (cit. on pp. 8, 37–39, 41, 45).
- Topalian, S. L., G. J. Weiner, and D. M. Pardoll (2011). “Cancer immunotherapy comes of age”. *J Clin Oncol*; 29 (36), pp. 4828–4836 (cit. on p. 91).
- US Food and Drug Administration (2019). *Guidance for Industry – Adaptive design clinical trials for drugs and biologics*. <https://www.fda.gov/media/78495/download> (cit. on p. 127).
- Wages, N. A., A. Iasonos, J. O’Quigley, and M. R. Conaway (2020). “Coherence principles in interval-based dose finding”. *Pharm. Stat.*; 19 (2), pp. 137–144 (cit. on p. 31).
- Wages, N. A. and C. Tait (2015). “Seamless Phase I/II Adaptive Design for Oncology Trials of Molecularly Targeted Agents”. *Journal of Biopharmaceutical Statistics*; 25 (5), pp. 903–920 (cit. on pp. 9, 37–39, 46, 55).
- Wheeler, G. M. et al. (2019). “How to design a dose-finding study using the continual reassessment method”. *BMC Med Res Methodol*; 19 (1), p. 18 (cit. on p. 128).
- Whitehead, J. and H. Brunier (1995). “Bayesian decision procedures for dose determining experiments”. *Stat Med*; 14 (9-10), pp. 885–893 (cit. on p. 20).

- Winther, S. B., T. L. Jørgensen, P. Pfeiffer, and C. Qvortrup (2016). “Can we predict toxicity and efficacy in older patients with cancer? Older patients with colorectal cancer as an example”. *ESMO open*; 1 (13), e000021 (cit. on pp. xii, 65).
- Yan, D., C. Tait, et al. (2019). “Generalization of the time-to-event continual reassessment method to bivariate outcomes”. *J Biopharm Stat*; 29 (4), pp. 635–647 (cit. on pp. 39, 55, 126).
- Yan, F., S. J. Mandrekar, and Y. Yuan (2017). “Keyboard: a novel Bayesian toxicity probability interval Design for Phase I Clinical Trials”. *Clin Cancer Res*; 23 (15), pp. 3994–4003 (cit. on pp. 28, 29).
- Yin, G., Y. Li, and Y. Ji (2006). “Bayesian dose-finding in phase I/II clinical trials using toxicity and efficacy odds ratios”. *Biometrics*; 62 (3), pp. 777–784 (cit. on pp. 37–39, 43).
- Yin, G. and Y. Yuan (2009). “Bayesian Model Averaging Continual Reassessment Method in Phase I Clinical Trials”. *Journal of the American Statistical Association*; 104 (487), pp. 954–968 (cit. on p. 24).
- Yin, G., S. Zheng, and J. Xu (2013). “Fractional Dose-Finding Methods with Late-Onset Toxicity in Phase I Clinical Trials”. *J Biopharm Stat*; 23 (4), pp. 856–870 (cit. on pp. 34, 35, 63).
- Yuan, Y., K. R. Hess, S. G. Hilsenbeck, and M. R. Gilbert (2016). “Bayesian optimal interval design: a simple and well-performing design for phase I oncology trials”. *Clin Cancer Res*; 7 (6), pp. 653–663 (cit. on pp. 28, 30, 50).
- Yuan, Y., R. Lin, D. Li, and L. Nie (2018). “Time-to-event Bayesian Optimal Interval Design to Accelerate Phase I Trials”. *Clinical Cancer Research*; 24 (20) (cit. on pp. 7, 36, 63).
- Yuan, Y., H. Q. Nguyen, and P. F. Thall (2016). *Bayesian Designs for Phase I–II Clinical Trials*. Ed. by 1. edn. Boca Raton: CRC Press. Chapman & Hall/CRC Biostatistics Series (cit. on p. 43).
- Yuan, Y. and G. Yin (2009). “Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes”. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*; 58 (5), pp. 719–736 (cit. on pp. 9, 38, 39, 55–57, 120, 124).

-
- (2011). “Robust EM Continual Reassessment Method in Oncology Dose Finding”. *J Am Stat Assoc*; 106 (495), pp. 818–831 (cit. on pp. 34, 35, 58, 63).
- Zang, Y., J. J. Lee, and Y. Yuan (2014). “Adaptive designs for identifying optimal biological dose for molecularly targeted agents”. *Clin Trials*; 11 (3), pp. 319–327 (cit. on p. 9).
- Zelen, M. (1969). “Play the winner rule and the controlled clinical trial”. *J Am Stat Assoc*; 64 (325), pp. 131–146 (cit. on p. 40).
- Zhang, W., D. J. Sargent, and S. Mandrekar (2006). “An adaptive dose-finding design incorporating both toxicity and efficacy”. *Stat Med*; 25 (14) (cit. on pp. 9, 37–39, 45, 92).
- Zhang, Y., S. Cao, et al. (2021). “A Bayesian adaptive phase I/II clinical trial design with late-onset competing risk outcomes”. *Biometrics*; 77 (3), pp. 796–808 (cit. on p. 92).
- Zhou, H., T. A. Murray, H. Pan, and Y. Yuan (2018). “Comparative review of novel model-assisted designs for phase I clinical trials”. *Stat Med*; 37 (14), pp. 2208–2222 (cit. on pp. 28, 30, 127).
- Zhou, H., Y. Yuan, and N. L (2018). “Accuracy, safety, and reliability of novel phase I trial designs”. *Clin Cancer Res*; 24 (18), pp. 4357–4364 (cit. on p. 30).
- Zhou, Y., J. J. Lee, and Y. Yuan (2019). “A utility-based Bayesian optimal interval (U-BOIN) phase I/II design to identify the optimal biological dose for targeted and immune therapies”. *Stat Med*; 38 (28), pp. 5299–5316 (cit. on pp. 38, 39, 51).
- Zhou, Y., R. Lin, and J. Lee (2022). “TITE-BOIN12: A Bayesian phase I/II trial design to find the optimal biological dose with late-onset toxicity and efficacy”. *Stat Med*; 41 (11), pp. 1918–1931 (cit. on pp. 39, 60).
- Zohar, S. and S. Chevret (2001). “The continual reassessment method : Comparison of bayesian stopping rules for dose-ranging studies”. *Stat Med*; 20 (19), pp. 2827–2843 (cit. on p. 22).
- Zohar, S. and J. O’Quigley (2006). “Optimal designs for estimating the most successful dose”. *Stat Med*; 25 (24) (cit. on pp. 61, 62).