Statistics and Its Interface Volume 3 (2010) 177–183

Maximum-relevance weighted likelihood estimator: Application to the continual reassessment method

Matthieu Resche-Rigon*, Sarah Zohar and Sylvie Chevret

Typical phase I dose-finding clinical trials, notably in cancer, are characterized by a small number of patients (less than 40), a relatively high number of dose levels (4 to 6) and sequential dose allocation rules. In this setting, the Continual Reassessment Method (CRM) has been recommended as a dose allocation rule that provides a consistent method to converge to the maximal tolerated dose (MTD), possibly based on likelihood (CRML). In this adaptive design setting, we derived a Relevance Weighted Likelihood to propose a robust estimation of the MTD. The main idea is to weight the individual contributions to likelihood using a decreasing function of rank. We compare this method to the CRML throughout simulations.

Keywords and phrases: Relevance weighted likelihood, Phase I, Dose-finding clinical trials, Continual reassessment method.

1. INTRODUCTION

The primary objective of phase I dose-finding cancer trials is to estimate the maximum tolerated dose (MTD) of a new drug among a low number of dose levels. Traditional cancer phase I designs are rule-based designs that treat the MTD as a sample statistic. More recently, model-guided designs have been developed, where the MTD is considered as the dose that corresponds to a prespecified probability of toxicity in the patient population, i.e., some percentile of interest. In most of these designs, an inference process on a parametric framework is used to guide dose escalation [13, 5, 22, 1, 20].

The Continual Reassessment Method (CRM) was proposed by O'Quigley et al. to estimate the MTD [13]. While originally based on a Bayes inference, a likelihood version of the original CRM (CRML) was developed thereafter [15]. This further requires a set of heterogeneous responses such that the first stage is a rule-based design using three patient cohorts that end upon observation of the first toxicity.

Whether rule- or model-based, CRM(L) is an adaptive design in which future design points are selected on the basis of previous responses at earlier design points. Indeed, the dose administered to any patient is selected on the basis of previous design points, namely patient doses and responses. Because dose assignments depend on previous data collection, the observations are dependent. Therefore, information drawn from any observation is used throughout the estimation process. Thus, the influence of each observation is expected to be related to its rank, with the first observations having the highest influence. This influence could be reinforced by the small sample size of the dose-finding trials. Moreover, as a binary regression, there is no response symmetry. The greater the number of estimated probabilities that are far from 0.5, the more the models are sensitive to a small number of observations [3]. Indeed, when the estimated probabilities are lower than 0.5, most information is dependent on the very few patients (who had a response). It is well-known that the Robbins-Monro procedure does not perform well in the estimation of extreme quantiles [18]. In a phase I dose-finding cancer trial setting, the target probability is usually below 50% and is commonly between 20% and 30% [13]. In phase II dose-finding trials that focus on the probability of failure, the targets are likely to be at 10% or even 5%, rather than 30% [17]. For situations involving such potentially extreme quantiles, the poor robustness of CRM(L) as a model-guided design has been duly noted [17].

Actually, the hidden assumption of CRML statistical modeling – i.e., that the probability distribution of dose-response is homogeneous – can be violated [9]. In the setting of adaptive designs with time heterogeneity, relevance weighted likelihood (ReWL) methods have been proposed by Hu and Rosenberger and more recently by Duan and Hu for doubly adaptive biased coin designs [9, 10, 4]. They consist of weighting the individual contributions to likelihood according to their relevance to decrease the influence of first observations on global conclusions [11, 7]. In the context of individual rank-related influence, we propose to develop a robust method for CRML by weighting the individual contributions to likelihood according to their rank using ReWL. This could be easily applied to the Bayesian CRM.

The paper is organized as follows. First, we present the weighted estimator of ReWL. Then Section 3 provides a simulation study to assess its relative performance compared to standard likelihood CRML. The results are presented in Section 4. Finally, a discussion with practical considerations is provided in Section 5.

*Corresponding author.

2. THE ROBUST LIKELIHOOD CONTINUAL REASSESSMENT METHOD

The CRM is based on a sequential estimation of the MTD from a finite set of dose levels and a fixed sample size n. Let d_i $(i = i, \ldots, k)$ denote the dose levels of the drug to be tested and p denote the target probability of response. The relationship between dose and response is modeled using a one-parameter model $\psi(x_i, \theta)$, where θ is the parameter to be estimated, x_i a function of doses given by $\psi^{-1}(p_i, \theta_0)$, p_i is the initial guess of toxicity probability associated with the dose level d_i and θ_0, the initial guessed value of the parameter θ. In this paper, we used the power model $\psi(x_i, \theta) = x_i^{\exp(\theta)}$ with $\theta_0 = 0$ as described by O'Quigley and Shen and recently recommended by Paoletti and Kramar [15, 16]. Since $\theta_0 = 0$ is chosen, the ψ function then reduces to $\psi(x_i, \theta) = p_i^{\exp(\theta)}$.

Let $\{(x(r), y_r);\ r = 1, \ldots, j\}$ be the accumulated data after the inclusion of the j^{th} patient, with $j \leq n$, $x(r)$ is the administered dose to the r^{th} patient, and y_r his(her) binary outcome.

The likelihood function $L_j(\theta)$ after j patients is defined by:

$$(1) \qquad L_j(\theta) = \prod_{r=1}^{j} \psi(x(r), \theta)^{y_r} (1 - \psi(x(r), \theta))^{(1-y_r)}$$

Attribution of doses is iteratively performed after each observation by the selection of the dose level $x(j+1)$, which minimizes $(\psi(x_i, \hat{\theta}_j) - p)^2; i = 1, \ldots, k$ where $\hat{\theta}_j$ is the updated model parameter through maximum likelihood estimation [15]. This will be referred as the CRML below.

To reduce the impact of first observations, we proposed to adapt weighted likelihood estimators such as those proposed by Hu and Rosenberger [9, 10] to the CRML. Each individual component of the likelihood (1) is thus weighted differently, so that the likelihood after j patients becomes:

$$(2) \qquad L_j^w(\theta) = \prod_{r=1}^{j} \psi(x(r), \theta)^{y_r w_r} (1 - \psi(x(r), \theta))^{(1-y_r) w_r}$$

where w_r is the weight of the r^{th} patient, out of a total of j patients. To slowly increase weights over ranks, the weight w_r of the r^{th} included patient was defined as follows:

$$(3) \qquad w_r = \log(\log(r + 2))^{\gamma}$$

with $\gamma \in [0, 4.5]$. After j included patients, $\hat{\theta}_j$ and $\hat{\gamma}_j$ are estimated using maximization of the weighted likelihood. The administered dose level to the next patient is that dose level associated with the estimated probability of response closest to the target. This allocation procedure will be further denoted ReWL CRM.

3. SIMULATION STUDY

We simulated phase I cancer dose-finding trials aiming at estimating the 10^{th} percentile of the dose-toxicity relationship. Six dose levels were considered, with initial guesses of toxic probabilities (the so-called working model) fixed at 0.01, 0.05, 0.1, 0.2, 0.3 and 0.5.

Six scenarios of actual toxic probabilities, $S_l(x) = P(Y = 1|x)$ $(l = 1, \ldots, 6)$ were examined (Figure 1). In scenario 1, the actual probabilities are equal to the working model. In scenario 2, the first dose level is noticeably nontoxic (10 fold lower than in scenario 1). Scenario 3 is similar to scenario 2, although the rate of toxicity above the MTD is greater than in scenario 2. Scenario 4 and 5 are close to scenario 3, though with increased toxicity from the first dose level. Finally, in scenario 6, toxicity is noticeably excessive for all doses. Note that, in scenario 3, the differential in toxicity between doses around the MTD is higher than in scenarios 1–2 and 4–5, so that the MTD is simpler to identify.

The trial sample size was fixed at $n = 24$. The first dose level was administered to the first patient. The dose allocation scheme and inference used the standard CRML and ReWL CRML, unless there was no heterogeneity in responses when the standard '3+3' scheme was first used, similarly to the CRML design of O'Quigley and Shen [15]. No skipping was allowed.

To better assess the robustness of the method — that is, to highlight the rank influence responsible for the dose-response heterogeneity — we evaluate the impact of outliers on results. One approach is to establish an outlier-generating model that allows a small number of observations from a random sample to come from a distribution differing from the targeted distribution. The observations from the outlier-generating model are called contaminants. To obtain unexpected or rare observations, we simulated a contaminated population, and assessed their influence according to their rank in the recruited sample. Highly toxic contaminants were generated from $S_l(x)^{\exp(\beta)}$ where $\beta = -2$, except for scenario 6, where $\beta = 2$, and the contaminants were concentrated within each quarter of the sample. The overall proportion of contaminants was fixed at 0.10; the expected number of contaminants was 2.4, all observed within each quarter of the sample — that is, within each subset of six patients.

To confirm the increased individual influence in the case of low target levels, we reran simulated trials using a 0.05 target from the first three scenarios, where the MTD was the second dose level. Finally, to assess the performances of the method when dealing with higher target levels, we reran the analyses using 0.30 as the target.

In each situation, operating characteristics were computed and compared from 20,000 simulated trials, namely, the percentage of dose correct selection (PCS) and the estimated response probabilities with mean bias and mean squared error (MSE), and the overall toxicities observed in the trial.

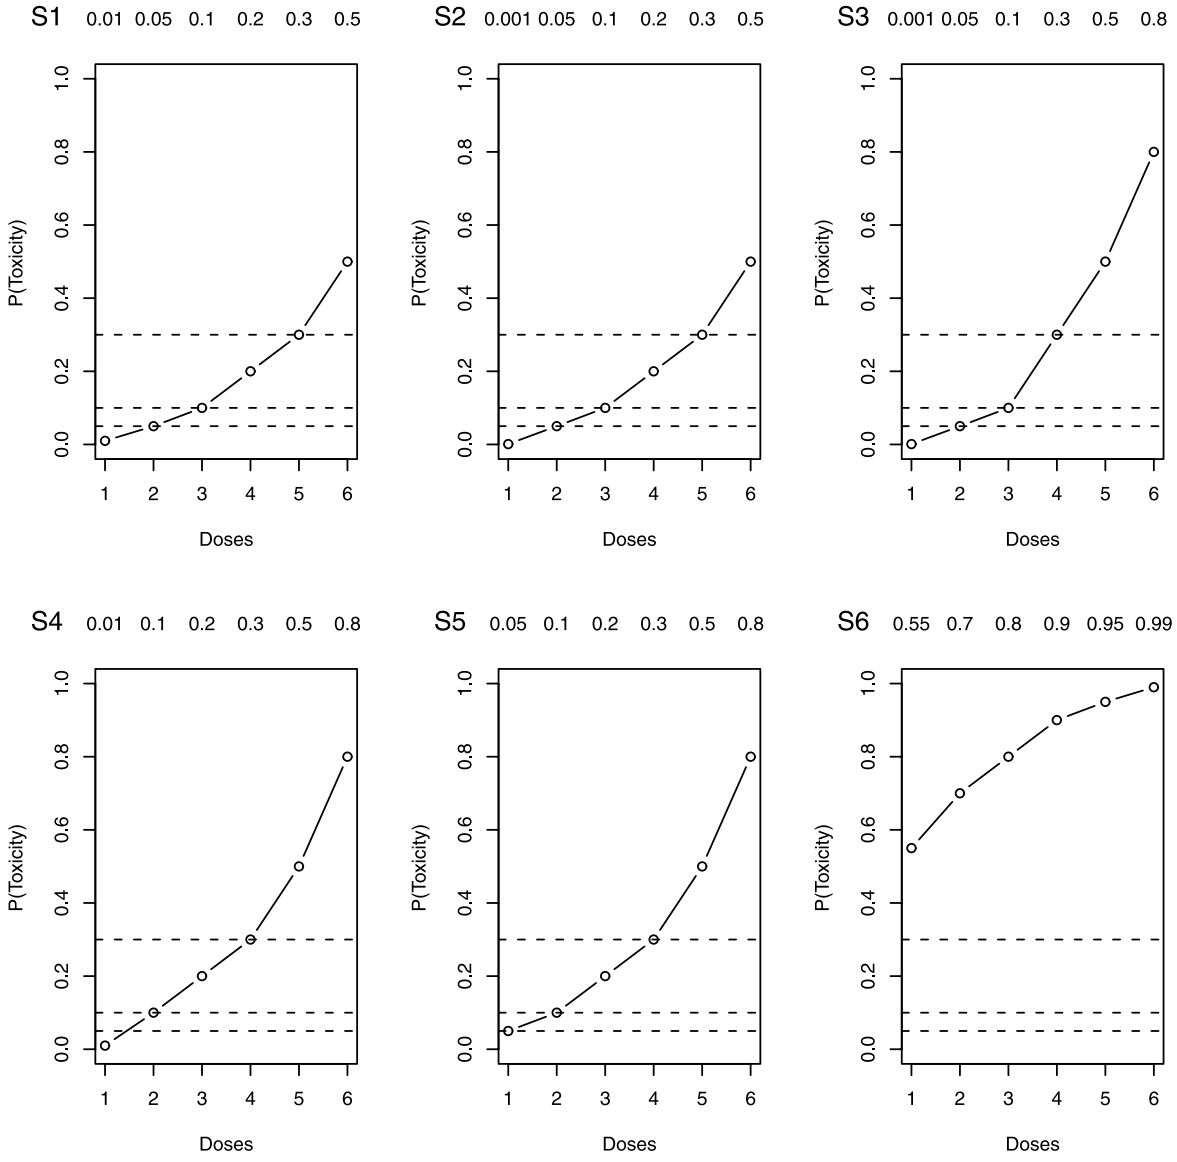

Figure 1. Dose-toxicity curves: Scenarios 1 to 6.

Simulations were carried out in the S language (R-cran 2.8 software). The code is available upon request to the first author.

4. RESULTS

4.1 Comparison of CRML and ReWL-CRM according to the target

Table 1 reports simulation results when dealing with target probabilities of 0.30, 0.10 and 0.05. When dealing with a targeted 30^{th} percentile of the dose-toxicity relationship, performances of CRML and ReWL CRM were close, either in terms of PCS, bias and MSE, or in terms of observed toxicities. However, when the target was 0.10, PCS was improved, whereas biases and MSE decreased when using ReWL CRM, compared to the use of CRML. Observed proportions of toxicities were close. This was further observed when dealing with a 0.05 target, where the gain in PCS achieved by the use of ReWL CRM was approximatively about 10% in the three scenarios. These findings illustrate how, as stated above, the individual influence of CRM is related to the rank, with greater influence in the case of low-targeted percentiles.

4.2 Increased heterogeneity in the dose-response

To better exemplify how the ReWL CRM outperforms CRML by erasing the influence of first individuals, we further simulated a contaminated population within each quar-

Table 1. Comparison of the Standard CRML or the ReWL CRM when estimating the MTD defined as the 5th, 10th or 30th percentiles of the dose-toxicity relationship, based on 24-patient cohorts and 20,000 replicated trials from the six different scenarios (S)

Target	S.	Method	Dose recommendation (%)			Bias	MSE	% toxicities
			Too low	MTD	Too high			
0.30	1	CRML	27.9	52.2	19.9	−0.010	0.098	19.8
		ReWL	25.5	53.8	20.7	−0.016	0.098	20.0
	2	CRML	26.6	53.0	20.4	−0.013	0.098	20.0
		ReWL	25.1	54.0	20.9	−0.018	0.098	19.9
	3	CRML	10.6	60.0	29.4	−0.019	0.068	24.9
		ReWL	8.9	59.7	31.5	−0.024	0.069	24.1
	4	CRML	24.8	53.1	22.0	0.009	0.080	24.5
		ReWL	23.1	53.4	23.5	0.005	0.080	24.1
	5	CRML	29.1	51.1	19.8	0.018	0.085	23.9
		ReWL	25.4	52.2	22.5	0.010	0.084	24.4
	6	CRML		99.9	0.1	0.001	0.102	55.4
		ReWL		99.8	0.2	0.003	0.113	55.8
0.10	1	CRML	30.1	41.3	28.6	0.007	0.065	9.5
		ReWL	27.5	43.8	28.7	−0.000	0.061	9.4
	2	CRML	27.7	42.9	29.4	0.002	0.060	9.6
		ReWL	26.8	45.0	28.2	−0.002	0.057	9.4
	3	CRML	36.8	49.2	14.1	0.059	0.179	13.1
		ReWL	30.0	55.6	14.4	0.010	0.052	9.7
	4	CRML	19.5	46.2	34.3	−0.002	0.049	10.7
		ReWL	18.6	46.4	35.0	−0.006	0.050	10.6
	5	CRML	34.8	38.5	26.7	0.025	0.076	11.2
		ReWL	29.2	39.4	31.4	0.013	0.071	11.3
	6	CRML		100.0		0.001	0.102	55.2
		ReWL		100.0		0.003	0.113	55.2
0.05	1	CRML	36.4	36.4	27.2	0.011	0.046	7.0
		ReWL	26.3	45.7	28.0	0.007	0.043	6.8
	2	CRML	33.1	38.9	27.9	0.004	0.035	6.8
		ReWL	21.7	49.5	28.7	0.002	0.034	6.6
	3	CRML	36.6	40.4	23.0	0.008	0.035	6.7
		ReWL	24.9	50.9	24.2	0.006	0.034	6.5

ter of the sample. The results are displayed in Figure 2 with a 0.10 target of toxic probability. Actually, the earlier the contaminants, the higher the difference in PCS and bias from the two methods, with improved performances of the ReWL CRM. This confirms that the ReWL CRM erases the influence of first observations, which is obvious in the CRM.

5. DISCUSSION

It has been established that the CRM is consistent under model misspecifications but not generally. This paper pointed out the rank influence in the CRM(L) when estimating the MTD, which was assessed throughout this simulation study. We also wondered whether the robustness of the CRM could be improved by downweighting the influence of first observations. Indeed, because observations are made sequentially by the dose-finding design, the probability distribution of the responses has been reported potentially time heterogeneous [9]. In such a setting, the potential for time trends could bias the results from standard likelihood analyses, and the weighted likelihood methodology was selected to account for this time trend. The results of our simulation study showed the superiority of the ReWL CRM over CRML with respect to both the correct estimation of the MTD and accuracy, especially in cases where the MTD was defined as a low percentile of the dose-toxicity relationship. Indeed, when the percentile of interest was low, the ReWL CRM was less sensitive than CRML in terms of how close the converged recommendation was to the target. Moreover, both the bias and the mean square error were reduced, in agreement with previous reports from other settings [9]. Finally, the ReWL estimator depends on the relevance weights that express the statistician's perceived relationship within the studied population and are usually chosen on intuitive grounds [7]. In such situations, as demonstrated by Hu in

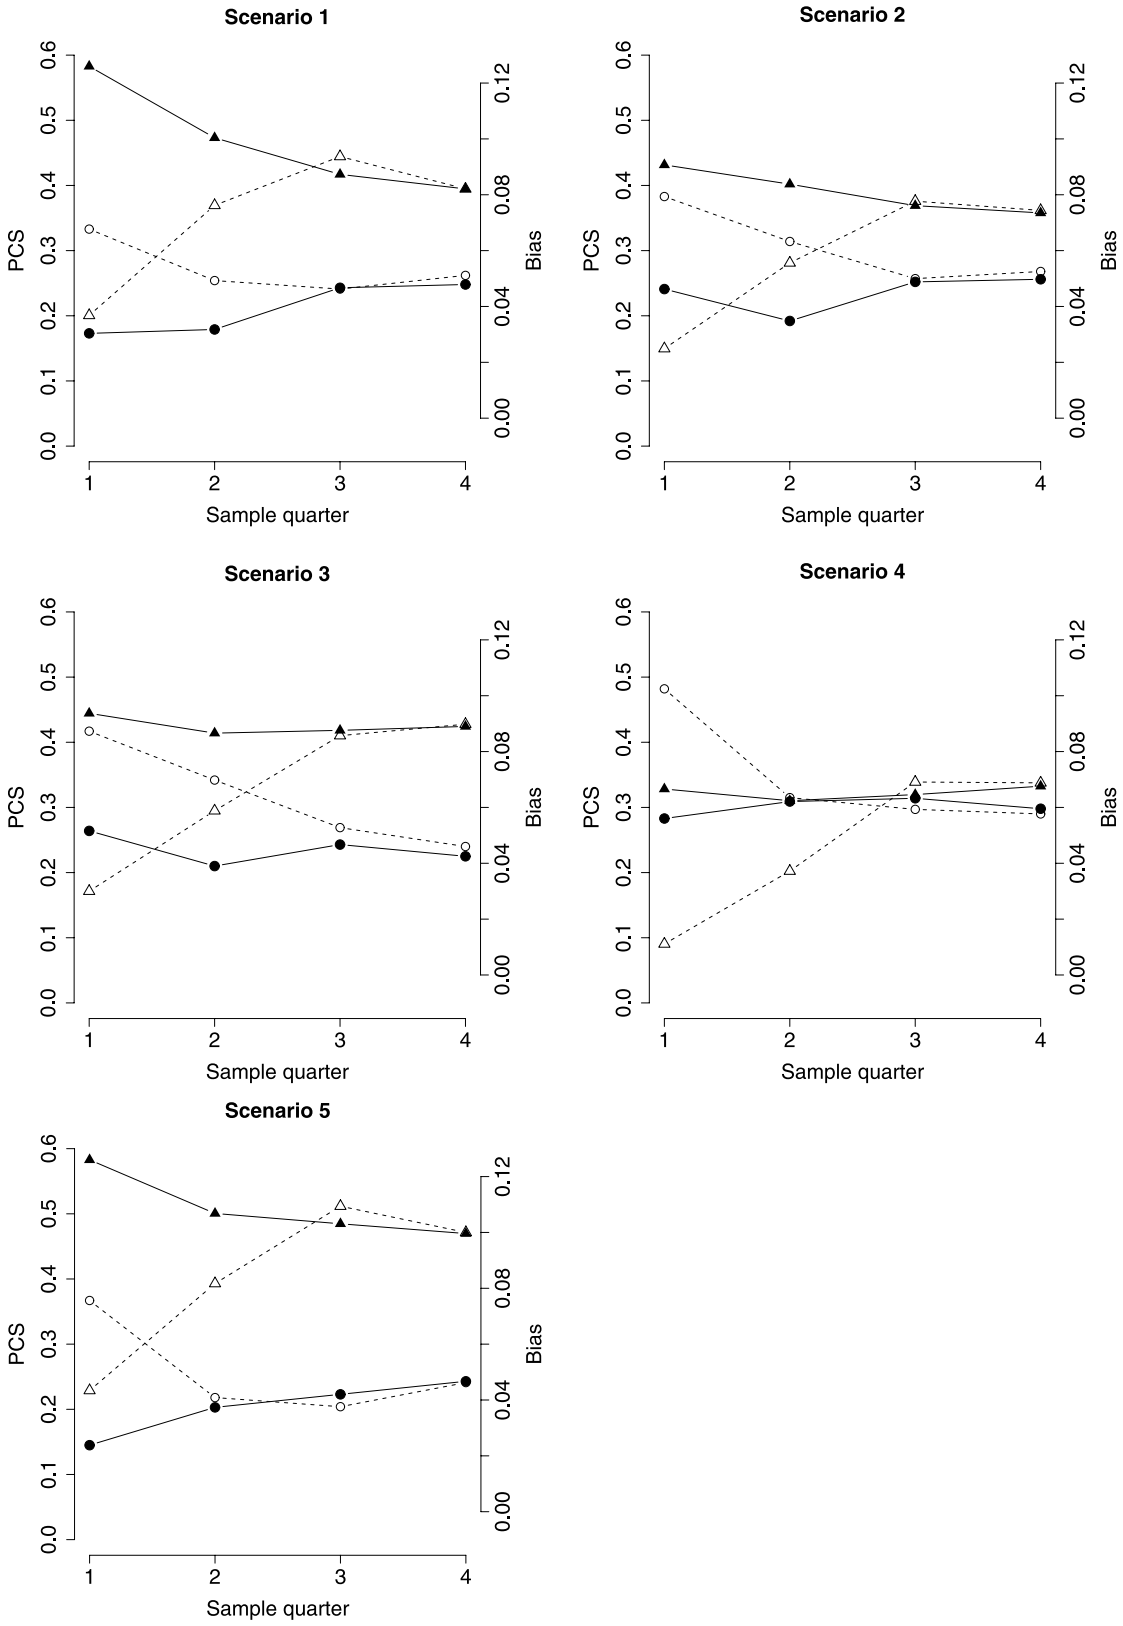

Figure 2. Comparison of CRML and ReWL CRM with regards to the percentage of correct dose selection (PCS) (● and ○ respectively) and estimated bias in toxicity probability at the recommended dose level after 24 inclusions (▲ and △ respectively), when an average 2.4 observations within each sample quarter are drawn from a contaminant population.

1997 and by Hu, Rosenberger and Zidek in 2000 for dependent data, relevance weighted likelihood estimator is consistent [7, 8]. For a large homogeneous sample, results would be similar for weighted or unweighted likelihood estimator. As an illustration, with a target of 0.3, ReWL CRM results are close to CRM results. Nevertheless the theoretical proof of the asymptotic properties of MWLE under CRM is an open question and should be investigated in future studies.

Several robust methods have been proposed for linear regression, that have been modified for the logistic model. These methods, such as the M-estimator [6, 21] and the E-estimator [19], consist in downweighting observations with large residuals at the time of analysis. When analysis is performed sequentially, this is questionable. Notably, detecting "large" residuals in comparison to the others is an open issue. Thus, we retained the broadest concept of weighted likelihood that best handles the sequential nature of the CRML. The theory of weighted likelihood has been used in a diverse group of applications. Actually, it has been already used in the setting of CRM [12, 2]. O'Quigley developed the so-called "retrospective CRM" to re-analyze dose-findings trials through the CRM by weighting observations with the frequency of previous dose allocation [12]. Cheung proposed the time-to-event CRM (TITE-CRM), which allows patients to be entered in a staggered fashion, with weights depending on the time-to-analysis [2]. We used a decreasing function of the rank over a bounded interval to insure a potentially heavy, decreased influence of first observations but close weights for the last ones [17]. Of note, when $\gamma = 0$, weights were all equal to one, so that the weighted likelihood (2) is equivalent to the standard likelihood (1).

Finally, we recommend the use of ReWL CRM when targets are below 10%. Although a simulation study cannot represent a universally valid truth in a mathematical sense, it allows learning about the properties of the design in various situations. Moreover, we choose to compare ReWL CRM to CRM on scenarios under which we know that CRM performs well. Most scenarios actually included the true MTD, but we also considered an extreme scenario, that investigated a dose range completely located over the true MTD. In all our simulations, ReWL CRM performs similarly or even better than the CRML in terms of PCS, and it was as efficient as the CRML from an ethical viewpoint. Notably, the overall percentage of observed toxicities was not increased despite the slight shift to the right of the underlying administered dose distribution. Nevertheless, this should not preclude the classical rules of prudence for conducting dose-finding trials: treating patients one-by-one (even with cohort sizes greater than one); including patients once a previous patient's response has been observed; and sequentially computing stopping rules based on toxicity to avoid false conclusion [14, 23].

Received 16 December 2009

REFERENCES

[1] Babb, J., Rogatko, A., and Zacks, S. (1998). Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103–1120.

[2] Cheung, Y. K. and Chappell, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56**, 1177–82. MR1815616

[3] Copas, J. (1988). Binary regression models for contaminated data. *Journal of the Royal Statistical Society, Series B* **50**, 225–265. MR0964178

[4] Duan, L. and Hu, F. (2009). Doubly adaptive biased coin designs with heterogeneous responses. *Journal of Statistical Planning and Inference* **139**(9), 3220–3230. MR2535195

[5] Gatsonis, C. and Greenhouse, J. B. (1992). Bayesian methods for phase I clinical trials. *Statistics in Medicine* **11**, 1377–1389.

[6] Hampel, F., Ronchetti, E., Rousseeuw, P., and Stahel, W. (1986). *Robust Statistics*. Wiley, New York. MR0829458

[7] Hu, F. (1997). The asymptotic properties of the maximum relevance weighted likelihood estimators. *The Canadian Journal of Statistics* **25**, 45–49. MR1451672

[8] Hu, F., Rosenberger, W., and Zidek, J. (2000). Relevance weighted likelihood for dependent data. *Metrika* **51**(3), 223–243. MR1795371

[9] Hu, F. and Rosenberger, W. F. (2000). Analysis of time trends in adaptive designs with application to a neurolophysiology experiment. *Statistics in Medicine* **19**, 2067–2042.

[10] Hu, F. and Rosenberger, W. F. (2006). *The Theory of Response-Adaptive Randomization in Clinical Trials*. Wiley series in probability and statistics. Wiley-Interscience, Hoboken, New Jersey, Chapter 7, 107–119. MR2245329

[11] Hu, F. and Zidek, J. (2002). The weighted likelihood. *The Canadian Journal of Statistics* **30**, 347–371. MR1944367

[12] O'Quigley, J. (2005). Retrospective analysis of sequential dose-finding designs. *Biometrics* **61**, 749–756. MR2196163

[13] O'Quigley, J., Pepe, M., and Fisher, L. (1990). Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics* **46**, 33–48. MR1059105

[14] O'Quigley, J. and Reiner, E. (1998). A stopping rule for the continual reassessment method. *Biometrika* **85**, 741–748. MR1665850

[15] O'Quigley, J. and Shen, L. Z. (1996). Continual reassessment method: A likelihood approach. *Biometrics* **52**, 673–684.

[16] Paoletti, X. and Kramar, A. (2009). A comparison of model choices for the continual reassessment method in phase I cancer trials. *Statistics in Medicine* **28**(24), 3012–3028.

[17] Resche-Rigon, M., Zohar, C., and Chevret, S. (2008). Adaptive designs for dose-finding in non-cancer phase II trials: Influence of early unexpected outcomes. *Clinical Trials* **5**, 595–606.

[18] Roshan Joseph, V. (2004). Efficient Robbins-Monro procedure for binary data. *Biometrika* **91**, 461–470. MR2081313

[19] Ruckstuhl, A. and Welsh, A. (2001). Robust fitting of the binomial model. *The Annals of Statistics* **29**, 1117–1136. MR1869243

[20] Thall, P. F. and Russell, K. E. (1998). A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* **54**, 251–264.

[21] Victoria-Feser, M. (2002). Robust inference with binary data. *Psychometrika* **67**, 21–32. MR1960827

[22] Whitehead, J. and Brunier, H. (1995). Bayesian decision procedures for dose determining experiments. *Statistics in Medicine* **14**, 885–893; discussion 895–899.

[23] Zohar, S. and Chevret, S. (2001). The continual reassessment method: Comparison of bayesian stopping rules for dose-ranging studies. *Statistics in Medicine* **20**, 2827–2843.

Matthieu Resche-Rigon
Université Paris Diderot - Paris 7, Paris, France
Département de Biostatistique et Informatique Médicale
Inserm, UMRS 717, Paris, France
Hôpital Saint-Louis, AP-HP
1 avenue Claude Vellefaux
75010 Paris, France
E-mail address:
matthieu.resche-rigon@paris7.jussieu.fr

Sarah Zohar
Inserm, UMRS 717, Paris, France
Département de Biostatistique et Informatique Médicale
Hôpital Saint-Louis, AP-HP
1 avenue Claude Vellefaux
75010 Paris, France
E-mail address: sarah.zohar@paris7.jussieu.fr

Sylvie Chevret
Université Paris Diderot - Paris 7, Paris, France
Département de Biostatistique et Informatique Médicale
Inserm, UMRS 717, Paris, France
Hôpital Saint-Louis, AP-HP
1 avenue Claude Vellefaux
75010 Paris, France
E-mail address: sylvie.chevret@paris7.jussieu.fr