

dfcomb: An R-package for phase I/II trials of drug combinations

Marie-Karelle Riviere, Jacques-Henri Jourdan, Sarah Zohar

▶ **To cite this version:**

Marie-Karelle Riviere, Jacques-Henri Jourdan, Sarah Zohar. dfcomb: An R-package for phase I/II trials of drug combinations. Computer Methods and Programs in Biomedicine, 2016, 125, pp.117-133. 10.1016/j.cmpb.2015.10.018 . hal-01297367

HAL Id: hal-01297367

https://hal.sorbonne-universite.fr/hal-01297367

Submitted on 4 Apr 2016

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.

dfcomb: an R-package for Phase I/II Trials of Drug Combinations

Marie-Karelle Riviere[(a)], Jacques-Henri Jourdan[(b)], and Sarah Zohar[(a)]

(a) INSERM, U1138, Equipe 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6, Paris, France
(b) INRIA Paris-Rocquencourt

Abstract

In this paper, we present the *dfcomb* R package for the implementation of a single prospective clinical trial or simulation studies of phase I combination trials in oncology. The aim is to present the features of the package and to illustrate how to use it in practice though different examples. The use of combination clinical trials is growing, but the implementation of existing model-based methods is complex, so this package should promote the use of innovative adaptive designs for early phases combination trials.

Keywords: Dose-finding; Phase I; Combination; Adaptive design; R package

1 Introduction

Most phase I cytotoxic dose-finding studies seek to establish a dose high enough to be able to observe potential efficacy while maintaining the toxicity rate within certain pre-defined acceptable limits. Phase I studies in oncology focus on determining the maximum tolerated dose (MTD) that will be used in further experiments [1, 2, 3]. However, nowadays, it is rare to find new agents that perform much better than existing therapeutic strategies. When combining two or more agents, there may be a synergistic effect in terms of efficacy. That is why investigators wish to increase overall anti-tumor action and survival by combining several agents: cytotoxics or targeted molecules or both. As a result, it is difficult to suppose that each agent will act independently in terms of toxicity.

When combining two agents, either a cytotoxic with a cytotoxic or a cytotoxic with a molecularly targeted agent (MTA), the underlying modeling assumptions diverge. Indeed, these two types of combinations differ in the shape of the dose-efficacy curves. For cytotoxic agents, a higher dose yields a greater response, but also leads to higher toxicity. However, the dose-efficacy relationship of the MTA may not follow the same monotonic pattern. The efficacy of some MTAs often increases at low dose levels and then plateaus (or approximately plateaus) at higher dose levels once a saturation level has been reached [4, 5, 6, 7]. Recently, we proposed two novel statistical designs for each combination type [8, 9]. In practice, these methods are based on Bayesian approaches that require high statistical and programming skills.

Many of the newly developed statistical methods for combination designs do not include software for their implementation. Some exceptions include POCRM [10] and PIPE [11]. As a result, many combination studies are still conducted with methods used for single agents [12].

In this article, we present a new R [13] package *dfcomb*, which provides functions for the conduct of a prospective single trial and simulation studies of the combination of two cytotoxics or a cytotoxic and an MTA in phase I dose-finding trials [8, 9]. The primary

features of *dfcomb* are as follows:

- Generates simulation replicates of phase I/II clinical trials for combination studies of two cytotoxic agents or a cytotoxic with an MTA.
- Determines the next or recommended combination in a phase I/II combination clinical trial.

The primary objective of this paper is to present the features of each method and to describe how to use the R *dfcomb* package in practice. Section 2 describes the models and computational methods of the package. Section 3 presents the R functions with simulations and examples of two dataset clinical trials, and in section 4 we conclude with a discussion.

2 Computational methods

2.1 Combination of two agents where toxicity and efficacy increase with both agents (`CombIncrease_sim` and `CombIncrease_next`)

Riviere et al [8] proposed a Bayesian phase I dose-finding design for clinical trials combining two agents where the toxicity of both agents is assumed to increase with dose levels. As in classic phase I design, the design only accounted for toxicity because the efficacy was assumed to be correlated with toxicity and to increase with dose levels. Therefore, the aim was to determine the maximum tolerated dose (MTD) in the context of the two-drug combination, defined as the combination associated with a toxicity closest to a maximum pre-defined dose-limiting toxicity (DLT) rate.

The combination toxicity probability at dose (j, k), $\pi_{j,k}$, composed of dose level j of agent 1 ($j = 1, \ldots, J$) and dose level k of agent 2 ($k = 1, \ldots, K$), was modeled using a 4-parameter logistic model:

$$\text{logit}(\pi_{\text{j,k}}) = \beta_0 + \beta_1 \text{u}_\text{j} + \beta_2 \text{v}_\text{k} + \beta_3 \text{u}_\text{j} \text{v}_\text{k},$$

where u_j and v_k denote the standardized doses of agents 1 and 2, β_0, β_1, β_2 and β_3 are unknown parameters such that $\beta_1 > 0$, $\beta_2 > 0$, $\forall k, \beta_1 + \beta_3 v_k > 0$ and $\forall j, \beta_2 + \beta_3 u_j > 0$, ensuring that the toxicity probability increases with the dose levels of both agents, and intercept $\beta_0 \in \mathbb{R}$. Gibbs sampling was used to sample in posterior distributions and estimate all posterior probabilities and parameters.

During the trial, determination of dose escalation and de-escalation to adjacent combinations was based on the uncertainty in $P(\pi_{j,k} < \theta|\text{data})$, that is on the probability that the current combination is safe enough. At the end of the trial, the MTD was selected as the tested combination with highest posterior probability, $P(\pi_{j,k} \in [\theta - \delta; \theta + \delta])$, of being in a pre-defined interval around the targeted toxicity θ. We implemented an algorithm-based start-up phase to gather enough information at the beginning of the trial in order to have reliable estimates of posterior probabilities. Moreover, we added stopping rules in cases where all dose levels were too toxic or all dose levels were far from the targeted toxicity (and therefore potentially ineffective). The details of the method can be found in [8].

2.2 Combination of two agents where the efficacy of one of the agents can plateau (`CombPlateau_sim` and `CombPlateau_next`)

Riviere et al [9] proposed a Bayesian phase I/II dose-finding design for clinical trials combining two agents where the toxicity and efficacy of both agents are assumed to increase with dose levels, but the efficacy of one agent can plateau (e.g. molecularly targeted agent).

The toxicity is assumed to be a binary outcome (1 if DLT, 0 otherwise), and the toxicity probability at combination (j, k) was modeled using a logistic model:

$$\text{logit}(p_{jk}) = \beta_0 + \beta_1 u_j + \beta_2 v_k$$

where and u_j and v_k are standardized doses ascribed to the jth dose level of agent 1 and the kth dose level of agent 2, β_0, β_1, and β_2 are unknown parameters such that $\beta_1 > 0$ and

$\beta_2 > 0$, so that toxicity monotonically increases with the dose levels of both agents 1 and 2.

In contrast, the efficacy response may require a longer follow-up time to be scored and was therefore considered as a time-to-event outcome. When combining an MTA with a cytotoxic agent, efficacy is expected to increase monotonically with the dose of the cytotoxic agent, but initially increases and then plateaus with the dose of the MTA, after the MTA reaches a level of saturation. We modeled the time to efficacy for the combination (j, k) using a proportional hazard model, $\lambda_{jk}(t)$, augmented with a plateau parameter τ:

$$\lambda_{jk}(t) = \lambda_0(t)\exp\{\gamma_1 w_j + \gamma_2(z_k \mathbb{1}(k < \tau) + z_\tau \mathbb{1}(k \geq \tau))\},$$

where $\lambda_0(t)$ is the baseline hazard, $\mathbb{1}(.)$ denotes the indicator function, and w_j and z_k are standardized doses. We assume that $\gamma_1 > 0$ and $\gamma_2 > 0$, and therefore efficacy increases monotonically with the dose of the cytotoxic agent 1. The plateau parameter τ is an integer between 1 and K and indicates at which dose level of agent 2 efficacy reaches a plateau. We assume an exponential distribution for the time to efficacy with a constant baseline hazard, i.e., $\lambda_0(t) = \lambda_0$, resulting in the following survival function for the time to efficacy:

$$S_{jk}(t) = \exp[-\lambda_0 t \exp\{\gamma_1 w_j + \gamma_2(z_k \mathbb{1}(k < \tau) + z_\tau \mathbb{1}(k \geq \tau))\}].$$

Then, the response rate, q_{jk}, at the end of the full follow-up time T for patients treated at the combination (j, k), is given by $q_{jk} = 1 - S_{jk}(T)$.

Gibbs sampling was used to sample in toxicity and efficacy posterior distributions and to estimate all posterior probabilities and parameters. The next optimal combination to be administered was selected as the admissible (safe: $P(p_{jk} > \theta) < C_{\mathrm{T}}$, and effective: $P(S_{jk}(T) > \xi) \geq C_{\mathrm{E}}$; where C_{T} and C_{E} are probability thresholds that need to be calibrated through a simulation study) combination with the highest estimate of efficacy, along with the lowest estimate of toxicity (if plateau reached).

We implemented an algorithm-based start-up phase to gather enough information at the

beginning of the trial in order to have reliable estimates of posterior probabilities. If during the trial, all combinations were inadmissible, the trial was stopped without a recommendation of combination. The details of the method can be found in [9]

3 R-functions

In this section, for each of the two methods described above, we present code examples and the corresponding outputs for: (1) simulation replicates of clinical trials, and (2) application to data of an ongoing clinical trial.

3.1 Combination of two agents where toxicity and efficacy increase with both agents (`CombIncrease_sim` and `CombIncrease_next`)

3.1.1 Simulation

We will illustrate the features of the `CombIncrease_sim` function. We considered 5 dose levels for the first agent (`ndose_a1=5`) and 3 dose levels for the second agent (`ndose_a2=3`). Then we defined our scenario representing the true toxicity probabilities of the combination. True toxicity probabilities should be entered in a matrix with agent 1 in rows and agent 2 in columns, with increasing toxicity probabilities with both row and column numbers, as for example:

```
p_tox <- matrix(c(0.05,0.10,0.15,0.30,0.45,
                  0.10,0.15,0.30,0.45,0.55,
                  0.15,0.30,0.45,0.50,0.60),nrow=5,ncol=3)
```

We chose a targeted toxicity probability (`target`) of 0.30, which is common in oncology due to the aggressiveness of the disease, and a targeted interval ([`target_min`;`target_max`]) of $[0.20; 0.40]$. We then provided initial guesses for toxicity probabilities for the dose levels

of agent 1, and agent 2. Typically, if information from phase I clinical trials for these agents individually is available, it can be incorporated to set up the prior toxicity probabilities.

```
prior_tox_a1 <- c(0.12, 0.2, 0.3, 0.4, 0.5)
prior_tox_a2 <- c(0.2, 0.3, 0.4)
```

In this example, if there is no early termination of the trial, a maximum total number of 20 cohorts (`n_cohort`) of 3 patients (`cohort`) should be included (that is 60 patients). The toxicity outcome was binary (`tite=FALSE`), 1 if a DLT is observed, 0 otherwise. We performed 2000 replicates of phase I clinical trials (`nsim`).

In order to determine the combination of drugs for the next cohort of patients, probability thresholds c_e and c_d have to be defined. They are critical for the performance of the design as they control the dose escalation and de-escalation. When setting up a new clinical trial, the values of c_e and c_d should be carefully calibrated through simulation, to ensure good operating characteristics of the design. In practice, this can be performed as follows. First, define a set of representative dose–toxicity scenarios that may be encountered in the trial, and then perform simulations using different values of c_e and c_d to evaluate the performance of the design. This is a trial-and-error process and may involve repeatedly tuning the values of c_e and c_d based on the simulation results. The goal is to find the values of c_e and c_d that yield good overall performance across different scenarios (e.g., the percentage of correct selection of the MTD, the number of patients exposed to over-toxic combinations or under-toxic combinations). We fixed these settings to default values, which seems to give good performance in a large variety of scenarios [8].

Finally, investigators may wish to stop a trial early when all combinations are estimated to be unacceptable in terms of toxicity. Therefore, we propose a stopping rule when the lowest combination is too toxic. In our example, to stop the trial early at least 6 patients should have been included in the lowest combination (`n_min`) and the probability that the estimated toxicity probability at that combination is above the target (0.30) should be greater than 0.95, meaning that the probability of over-dosing is high. Again, for a simpler use, these arguments

were set by default to values giving, in general, good performances for inexperienced users. Then users can choose to calibrate them by simulation studies.

This results in the following call to the function:

```
sim <- CombIncrease_sim(ndose_a1=5, ndose_a2=3, p_tox=p_tox, target=0.30,
target_min=0.20, target_max=0.40, prior_tox_a1=prior_tox_a1,
prior_tox_a2=prior_tox_a2, n_cohort=20, cohort=3, tite=FALSE, nsim=2000,
c_e=0.85, c_d=0.45, c_stop=0.95, n_min=6, seed = 14061991)
```

A progress bar appears that indicates the progress of the number of simulations chosen. The simulation can be long, and it is normal to wait some time before the bar changes. Calculation time will clearly depend on the computer resources the user has at his or her disposal, but this example takes about 3 hours 45 minutes to run on a single portable computer with an Intel Core i7. If the call to the function is stopped before its end, the results for the current number of simulations ended are displayed. The generated output is as follows:

```
True toxicities:
       Agent 1
Agent 2    1    2    3    4    5
      3 0.15 0.30 0.45 0.50 0.60
      2 0.10 0.15 0.30 0.45 0.55
      1 0.05 0.10 0.15 0.30 0.45
```

True toxicity probabilities entered are recalled, and then the results show at each combination (1) the percentage of selection, (2) the mean number of patients included, and (3) the mean number of observed toxicities. Therefore, the percentage of correct selection of a true MTD (combination associated with the toxicity target of 30%) is $28.30 + 40.65 + 6.90 = 75.85\%$. The mean number of patients and toxicities of each combination can be used to evaluate the number of patients exposed to over-dosing or under-dosing.

```
Percentage of Selection:

        Agent 1
Agent 2    1     2     3    4    5
       3 6.0 28.30  6.40 0.15 0.00
       2 0.2  4.45 40.65 3.75 0.05
       1 0.0  0.05  3.05 6.90 0.05


Number of patients:

       Agent 1
Agent 2    1    2     3    4    5
      3 4.24 9.83  5.59 0.78 0.09
      2 0.96 9.28 13.44 2.10 0.07
      1 3.81 1.17  5.48 2.98 0.19


Number of toxicities:

       Agent 1
Agent 2    1    2    3    4    5
      3 0.61 2.93 2.57 0.40 0.06
      2 0.09 1.42 4.05 0.95 0.04
      1 0.19 0.12 0.80 0.93 0.09
```

The trials did not stop early and included the total sample size of 60 patients. Other entries are also recalled as follows:

```
Percentage of inconclusive trials:      0
The minimum number of cohorts to stop the trial is:      2


Number of simulations:   2000
Cohort size:     3
```

```
Number of cohort planned:        20
Total patients accrued:  60
Toxicity target:         0.3
Targeted toxicity interval:      [ 0.2 , 0.4 ]
Prior toxicity probabilities for agent 1:
[1] 0.12 0.20 0.30 0.40 0.50
Prior toxicity probabilities for agent 2:
[1] 0.2 0.3 0.4
Escalation threshold:    0.85
Deescalation threshold:  0.45
Stopping threshold:      0.95
Toxicity is not a time-to-event but binary
```

3.1.2 Application

In a real clinical trial, the `CombIncrease_next` function determines the next combination to administer. As with `CombIncrease_sim`, the number of dose levels for each agent, targeted toxicities, initial guesses for toxicity probabilities, and minimum number of patients for early termination should be specified.

```
prior_tox_a1 <- c(0.12, 0.2, 0.3, 0.4, 0.5)
prior_tox_a2 <- c(0.2, 0.3, 0.4)
```

In this example, we considered that the trial is still ongoing (`final=FALSE`), and currently 18 patients have been enrolled (`pat_incl`). For each of these 18 patients, two vectors corresponding to the dose levels of each agent administered to the patient should be given, as for example:

```
dose1 <- c(1,1,1,2,2,2,3,3,3,3,3,3,3,3,4,4,4)
dose2 <- c(1,1,1,2,2,2,3,3,3,2,2,2,1,1,1,1,1)
```

In the following example, the last three patients have received dose 4 of agent 1 and dose 1 of agent 2. The toxicity outcome is binary (`tite=FALSE`), so the user have to input a vector of toxicity containing for each patient: 1 if a dose-limiting toxicity is observed and 0 otherwise.

```
toxicity <- c(0,0,0,0,0,0,0,0,1,0,1,0,0,0,0,0,0,1)
```

If the toxicity is considered as a time-to-event (`tite=TRUE`), then two vectors containing the time-to-toxicity and the follow-up time for each patient as well as the value corresponding to the full follow-up time window should be given. The detailed example results in the following call to the function:

```
next_dose <- CombIncrease_next(ndose_a1=5, ndose_a2=3, target=0.3,
target_min=0.20, target_max=0.40, prior_tox_a1=prior_tox_a1,
prior_tox_a2=prior_tox_a2, final=FALSE, pat_incl=18, dose_adm1=dose1,
dose_adm2=dose2, tite=FALSE, toxicity=toxicity, n_min=6)
```

Another possible option include `in_startup` (set at TRUE by default). If FALSE the start-up phase is forced to be ended and thus estimation can be performed even if the clinical trial did not follow the start-up phase recommended by the original design. This code takes less than one second to run. The generated output is composed as the recalled number of patients and toxicities at each combination:

```
Number of patients:
       Agent 1
Agent 2 1 2 3 4 5
      3 0 0 3 0 0
      2 0 3 3 0 0
      1 3 0 3 3 0
```

```
Number of toxicities:
       Agent 1
Agent 2 1 2 3 4 5
      3 0 0 1 0 0
      2 0 0 1 0 0
      1 0 0 0 1 0
```

The estimated posterior probabilities for toxicity, under/over-dosing are reported. They are used to evaluate the safety of combinations and lead to the recommendation of the next most appropriate combination:

```
Toxicity prob:
       Agent 1
Agent 2    1    2    3    4    5
      3 0.10 0.21 0.39 0.57 0.70
      2 0.03 0.06 0.18 0.42 0.65
      1 0.01 0.02 0.07 0.25 0.58


P(toxicity prob < target):
       Agent 1
Agent 2    1    2    3    4    5
      3 0.95 0.77 0.32 0.10 0.05
      2 1.00 0.99 0.87 0.28 0.08
      1 1.00 1.00 0.99 0.68 0.17


Prob underdosing:
       Agent 1
Agent 2    1    2    3    4    5
      3 0.85 0.54 0.13 0.03 0.01
```

```
      2 0.99 0.96 0.63 0.11 0.03
      1 1.00 0.99 0.95 0.42 0.08


Prob targeted interval:
       Agent 1
Agent 2    1    2    3    4    5
      3 0.13 0.36 0.42 0.18 0.09
      2 0.01 0.04 0.34 0.38 0.14
      1 0.00 0.00 0.05 0.43 0.19


Prob overdosing:
       Agent 1
Agent 2    1   2    3    4    5
      3 0.02 0.1 0.45 0.78 0.90
      2 0.00 0.0 0.03 0.51 0.83
      1 0.00 0.0 0.00 0.15 0.73
```

The next recommended combination is $(4, 1)$ (dose level 4 of agent 1, and dose level 1 of agent 2). Entries as well as accumulated information are reported.

```
Start-up phase ended:    YES
NEXT RECOMMENDED COMBINATION:    (4,1)


Number of patients included:     18
Toxicity target:         0.3
Targeted toxicity interval:      [ 0.2 , 0.4 ]
Prior toxicity probabilities for agent 1:
[1] 0.12 0.20 0.30 0.40 0.50
Prior toxicity probabilities for agent 2:
```

```
[1] 0.2 0.3 0.4
The minimum number of patients to stop the trial is:     6
Escalation threshold:    0.85
Deescalation threshold:  0.55
Stopping threshold:      0.95
Toxicity is not a time-to-event but binary
```

3.2 Combination of two agents where the efficacy of one of the agents can plateau (`CombPlateau_sim` and `CombPlateau_next`)

3.2.1 Simulation

We will illustrate the features of the `CombPlateau_sim` function. We considered 3 dose levels for the cytotoxic agent (`ndose_a1`) and 4 dose levels for the molecularly targeted agent (`ndose_a2`). We defined the true toxicity and efficacy probabilities of the combination. They should be entered in rows for the cytotoxic agent and in columns for the MTA, with increasing toxicity/efficacy probabilities with both row and column numbers:

```
p_tox <- t(matrix(c(0.10,0.15,0.30,0.45,
                    0.15,0.30,0.45,0.50,
                    0.30,0.45,0.55,0.65),nrow=4,ncol=3))
p_eff <- t(matrix(c(0.25,0.25,0.26,0.27,
                    0.40,0.41,0.41,0.42,
                    0.55,0.55,0.56,0.56),nrow=4,ncol=3))
```

We chose a maximum targeted toxicity probability of 0.30 (`tox_max`) and a minimum efficacy probability of 0.20 (`eff_min`). We then provided initial guesses for toxicity/efficacy probabilities for the dose levels of both agents:

```
prior_tox_a1 <- c(0.2, 0.3, 0.4)
```

```
prior_eff_a1 <- c(0.3, 0.4, 0.5)
prior_tox_a2 <- c(0.12, 0.2, 0.3, 0.4)
prior_eff_a2 <- c(0.3, 0.4, 0.5, 0.59)
```

Again, if information from phase I clinical trials for these agents used as single-agents is available, it can be incorporated to set up the prior toxicity probabilities. In this example, if there is no early termination of the trial, a maximum sample size of 75 enrolled patients is needed (`n`), with a cohort of 3 patients throughout the trial (`cohort_start`, `cohort`). As efficacy is considered as a time-to-event outcome (time-to-progression), the full follow-up time to assess efficacy should be specified (`time_full`). In this example we chose a follow-up time of 7 weeks. The patient accrual is simulated using a Poisson process. We chose a rate of 0.28 (`poisson_rate`) meaning that on average 1 patient arrives every 3.5 weeks. Even if toxicity should be quickly ascertainable, it is possible to consider that before dose determination the cohort should be followed up for a minimum time to assess toxicity. Usually in oncology, 3 weeks are needed to evaluate toxicity. In the following example, we did not constrain a minimum follow-up time before dose determination and set `cycle` to 0. We performed 2000 replicates of phase I/II clinical trials (`nsim`).

In order to determine the combination of drugs for the next cohort of patients, toxicity and efficacy thresholds c_{tox} and c_{eff} have to be defined. These probabilities are used to define combinations that are admissible in terms of toxicity and efficacy. A smaller c_{tox} will be more stringent and require a safer combination in terms of toxicity, and a higher c_{eff} will require a more effective combination. When setting up a new clinical trial, the values of c_{tox} and c_{eff} should be calibrated through simulation depending on the desired operating characteristics of the design. The following call to the function was performed:

```
sim <- CombPlateau_sim(ndose_a1=3, ndose_a2=4, p_tox=p_tox, p_eff=p_eff,
tox_max=0.30, eff_min=0.20, prior_tox_a1=prior_tox_a1,
prior_tox_a2=prior_tox_a2, prior_eff_a1=prior_eff_a1,
prior_eff_a2=prior_eff_a2, n=75, cohort_start=3, cohort=3, time_full=7,
```

```
poisson_rate=0.28, cycle=0, nsim=2000, c_tox=0.85, c_eff=0.10,
seed = 2174892, threads=0)
```

The seed for simulations can be chosen. Finally, as these simulations are computationally intensive, the run time can be long. The argument `threads` is used to specify the number of threads of the computer to use; by default 0 uses all threads. As this design deals with both toxicity and efficacy, simulations in general take longer to execute than for the previous design. We run this code in a little over 2 days (52 hours).

The generated output is first composed of the recalled toxicity and efficacy probabilities:

```
True toxicities:
       Agent 1
Agent 2    1    2    3
      4 0.45 0.50 0.65
      3 0.30 0.45 0.55
      2 0.15 0.30 0.45
      1 0.10 0.15 0.30


True efficacies:
       Agent 1
Agent 2    1    2    3
      4 0.27 0.42 0.56
      3 0.26 0.41 0.56
      2 0.25 0.41 0.55
      1 0.25 0.40 0.55

Percentage of Selection:
        Agent 1
Agent 2     1   2     3
```

```
        4 0.35 0.2  0.00
        3 1.05 0.4  0.20
        2 0.50 7.0  5.55
        1 0.05 8.5 73.25

Number of patients:
        Agent 1
Agent 2    1    2     3
      4 1.92 0.26  0.29
      3 3.71 1.27  1.29
      2 4.00 5.45  8.65
      1 3.88 7.07 35.13

Number of toxicities:
        Agent 1
Agent 2    1    2     3
      4 0.85 0.12  0.19
      3 1.11 0.58  0.71
      2 0.61 1.64  3.90
      1 0.41 1.04 10.51

Number of efficacies:
        Agent 1
Agent 2    1    2     3
      4 0.50 0.11  0.16
      3 0.95 0.53  0.71
      2 1.02 2.25  4.80
```

```
      1 0.98 2.83 19.28

Percentage of inconclusive trials:      2.95
```

The results show that the percentage of correct selection of the optimal combination $(3, 1)$ is 73.25%. The trial was stopped at 2.95% because no combination was found that was acceptable in terms of both toxicity and efficacy.

```
Number of simulations:   2000
Cohort size for the start-up phase:      3
Cohort size for the model-based phase:   3
Total sample size:       75
Total patients accrued:  72.9
Toxicity upper bound:    0.3
Efficacy lower bound:    0.2
Prior toxicity probabilities for agent 1:
[1] 0.2 0.3 0.4
Prior toxicity probabilities for agent 2:
[1] 0.12 0.20 0.30 0.40
Prior efficacy probabilities for agent 1:
[1] 0.3 0.4 0.5
Prior efficacy probabilities for agent 2:
[1] 0.30 0.40 0.50 0.59
Toxicity threshold:      0.85
Efficacy threshold:      0.1
Toxicity is binary
Efficacy is a time-to-event
Full follow-up time:     7
Minimum waiting time between two dose cohorts is of one toxicity cycle of: 0
```

```
Patient arrival is modeled as a Poisson process with rate: 0.28
Trial mean duration: 268.3669
```

The mean trial duration was 268 weeks, that is about 5 years, if only one center is recruiting, with on average one patient enrolled every 3.5 weeks. Entries and accumulated data are reported.

3.2.2 Application

When used for a real clinical trial, the `CombPlateau_next` function determines the next combination to administer. As with `CombPlateau_sim`, the number of dose levels for each agent, maximum toxicity, minimum efficacy, and initial guesses for toxicity/efficacy probabilities should be specified:

```
prior_tox_a1 <- c(0.2, 0.3, 0.4)
prior_eff_a1 <- c(0.3, 0.4, 0.5)
prior_tox_a2 <- c(0.12, 0.2, 0.3, 0.4)
prior_eff_a2 <- c(0.3, 0.4, 0.5, 0.59)
```

`stage` is an integer with value 0 if less than half of the total sample size has been included, 1 if more than half of the total sample size has been included but the trial is still ongoing, and 2 if the trial is over and the dose recommendation should be made. We considered that the trial is ongoing with less than half of the total sample size included, setting `stage` to 0, and the start-up phase is finished (`in_startup=FALSE`). After the start-up phase, patients were included by cohorts of 3 patients (`cohort`). Currently, 18 patients have been enrolled (`pat_incl`). For each patient, two vectors corresponding to the dose levels of each agent administered to the patient should be given:

```
dose_adm1 <- c(1,1,1,1,1,1,2,2,2,3,3,3,3,3,3,2,2,2)
dose_adm2 <- c(1,1,1,2,2,2,1,1,1,1,1,1,1,1,1,2,2,2)
```

A vector containing the binary toxicity outcome for each patient has to be defined.

```
toxicity <- c(0,0,0,0,0,1,0,0,0,0,0,1,0,0,1,0,0,0)
```

The full follow-up time for efficacy assessment was set to 7 weeks. As efficacy is assessed as a time-to-progression, for each patient, the follow-up time (maximum equal to the full follow-up time) as well as the time-to-progression are required. If no progression is observed for the patient, then the value should be set to +Inf.

```
time_prog <- c(1.6,4.2,3.5,5.1,2.4,4.8,2.8,4.4,+Inf,3.9,+Inf,4.6,1.8,+Inf,
0.5, 5.4,2.8,+Inf)
time_follow <- c(rep(7,15), 4.9, 3.1, 1.3)
```

The example detailed above results in the following call to the function:

```
next_dose <- CombPlateau_next(ndose_a1=3, ndose_a2=4, tox_max=0.30,
eff_min=0.20, prior_tox_a1, prior_tox_a2, prior_eff_a1, prior_eff_a2,
stage=0, in_startup=FALSE, cohort=3, pat_incl=18, dose_adm1=dose_adm1,
dose_adm2=dose_adm2, toxicity=toxicity, time_full=7, time_prog=time_prog,
time_follow=time_follow)
```

This code takes two seconds to run. The output reports the number of patients included, observed toxicities and progressions at each combination according to the entries:

```
Number of patients:
       Agent 1
Agent 2 1 2 3
      4 0 0 0
      3 0 0 0
      2 3 3 0
      1 3 3 6


Number of toxicities:
```

```
       Agent 1
Agent 2 1 2 3
      4 0 0 0
      3 0 0 0
      2 1 0 0
      1 0 0 2


Number of progressions:
       Agent 1
Agent 2 1 2 3
      4 0 0 0
      3 0 0 0
      2 3 1 0
      1 3 2 4
```

Then estimated toxicities and efficacies leading to recommendation of combinations are given:

```
Toxicity prob:
       Agent 1
Agent 2    1    2    3
      4 0.21 0.31 0.40
      3 0.17 0.25 0.34
      2 0.13 0.19 0.26
      1 0.08 0.13 0.18


P(toxicity prob > tox_max):
       Agent 1
Agent 2    1    2    3
```

```
      4 0.35 0.51 0.67
      3 0.23 0.39 0.58
      2 0.08 0.19 0.40
      1 0.01 0.04 0.18

Efficacy prob:
       Agent 1
Agent 2    1    2    3
      4 0.35 0.51 0.66
      3 0.32 0.48 0.63
      2 0.24 0.39 0.54
      1 0.13 0.26 0.40

P(efficacy prob < eff_min):
       Agent 1
Agent 2    1    2    3
      4 0.75 0.97 0.99
      3 0.72 0.96 0.99
      2 0.61 0.95 0.99
      1 0.24 0.73 0.93

Prob plateau: Agent 2
   4    3    2    1
0.38 0.27 0.19 0.15

Start-up phase ended:    YES
NEXT RECOMMENDED COMBINATION:    (3,2)
```

The next recommended combination is $(3, 2)$. With the current data, the plateau location for the MTA is estimated at dose level 4 (highest posterior probability), which means that with the dose involved, the efficacy plateau should not be reached. Other entries are also recalled:

```
Number of patients included:     18
Maximum toxicity:        0.3
Minimum efficacy:        0.2
Prior toxicity probabilities for agent 1:
[1] 0.2 0.3 0.4
Prior toxicity probabilities for agent 2:
[1] 0.12 0.20 0.30 0.40
Prior efficacy probabilities for agent 1:
[1] 0.3 0.4 0.5
Prior efficacy probabilities for agent 2:
[1] 0.30 0.40 0.50 0.59
Toxicity threshold:      0.85
Efficacy threshold:      0.1
Toxicity is binary
Efficacy is a time-to-event
Full follow-up time:     7
Minimum waiting time between two dose cohorts is of one toxicity cycle of: 0
```

4 Conclusions

The *dfcomb* package implements two methods for phase I clinical trials in oncology for the combination of two agents depending on the underlying assumptions: either both toxicity and efficacy increase with the dose of each agent involved (e.g. combination of two cytotoxic

agents), or the efficacy can plateau with the dose of one agent (e.g. combination of a cytotoxic agent with a molecularly targeted agent). For each method, the package implements a single prospective clinical trial or simulation studies. In this package, several flexible inputs can be used, such as: (1) prior information on each agent gathered before the trial, (2) the cohort size of patients allocated at each combination level as well as the total sample size, (3) the targeted toxicity, (4) binary or censured data among other possible entries. In the last decade some innovative methods have been proposed for combination studies. However, the implementation of such complex approaches is difficult as most commercial software programs do not provide functions or macros for these approaches. Indeed, many early phase clinical trial methods are not used in practice as programs or scripts are not available. Our R package enables statisticians and physicians to implement adaptive model-based dose-finding designs when two anti-cancer agents are studied.

References

[1] D. Faries. Practical modifications of the continual reassessment method for phase I cancer clinical trials. *J Biopharm Stat*, 4:147–64, 1994.

[2] E. L. Korn, D. Midthune, T. T. Chen, L. V. Rubinstein, M. C. Christian, and R. M. Simon. A comparison of two phase I trial designs. *Stat Med*, 13:1799–806, 1994.

[3] J. O'Quigley and S. Zohar. Experimental designs for phase I and phase I/II dose-finding studies. *Br. J. Cancer*, 94(5):609–613, Mar 2006.

[4] B. J. Druker. Perspectives on the development of a molecularly targeted agent. *Cancer Cell*, 1(1):31–36, Feb 2002.

[5] A. Gafter-Gvili, A. Leader, R. Gurion, L. Vidal, R. Ram, A. Shacham-Abulafia, I. Ben-Bassat, M. Lishner, O. Shpilberg, and P. Raanani. High-dose imatinib for newly di-

agnosed chronic phase chronic myeloid leukemia patients–systematic review and meta-analysis. *Am. J. Hematol.*, 86(8):657–662, Aug 2011.

[6] L. M. Ellis. Antiangiogenic therapy: more promise and, yet again, more questions. *J. Clin. Oncol.*, 21(21):3897–3899, Nov 2003.

[7] B. Morgan, A. L. Thomas, J. Drevs, J. Hennig, M. Buchert, A. Jivan, M. A. Horsfield, K. Mross, H. A. Ball, L. Lee, W. Mietlowski, S. Fuxuis, C. Unger, K. O'Byrne, A. Henry, G. R. Cherryman, D. Laurent, M. Dugan, D. Marme, and W. P. Steward. Dynamic contrast-enhanced magnetic resonance imaging as a biomarker for the pharmacological response of PTK787/ZK 222584, an inhibitor of the vascular endothelial growth factor receptor tyrosine kinases, in patients with advanced colorectal cancer and liver metastases: results from two phase I studies. *J. Clin. Oncol.*, 21(21):3955–3964, Nov 2003.

[8] M. K. Riviere, Y. Yuan, F. Dubois, and S. Zohar. A Bayesian dose-finding design for drug combination clinical trials based on the logistic model. *Pharm Stat*, 13(4):247–257, 2014.

[9] M. K. Riviere, Y. Yuan, F. Dubois, and S. Zohar. A Bayesian dose finding design for clinical trials combining a cytotoxic agent with a molecularly targeted agent. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 64(1):215–229, 2015.

[10] N. A. Wages and N. Varhegyi. pocrm: an R-package for phase I trials of combinations of agents. *Comput Methods Programs Biomed*, 112(1):211–218, Oct 2013.

[11] A. P. Mander and M. J. Sweeting. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Stat Med*, 34(8):1261–1276, Apr 2015.

[12] M. K. Riviere, C. Le Tourneau, X. Paoletti, F. Dubois, and S. Zohar. Designs of drug-combination phase I trials in oncology: a systematic review of the literature. *Ann. Oncol.*, Nov 2014.

[13] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014.