

OPTIMAL STOPPING FOR RESPONSE-GUIDED DOSING

JAKOB KOTAS

Department of Mathematics
University of Portland
5000 N Willamette Blvd
Portland, OR 97203, USA

ABSTRACT. In response-guided dosing (RGD), the goal is to make optimal dosing decisions based on the stochastic evolution of a patient’s disease condition. Typically, RGD is formulated as a finite-horizon problem with decision-making occurring over a predetermined time frame. In this paper we relax the latter assumption to allow for the possibility of ending treatment early. This could occur due to remission of the disease or a finding of futility in treatment of the disease. Our framework is formulated as a stochastic dynamic program (DP) where a stop/do-not-stop decision is made in discrete sessions, and if stopping is not chosen, an optimal dose is determined for that session. Numerical simulations for rheumatoid arthritis are presented, and monotonicity of the stop/do-not-stop threshold with respect to time is proven.

1. Introduction. Optimal stopping of stochastic dynamic programs (DPs) (also known as Markov decision processes) has been an area of interest in operations research for decades [14, 5, 16]. This paper attempts to apply the theory of optimal stopping to the problem of response-guided dosing, where patients receive dosing specific to their individual disease progression over time.

Treatment paradigms for various diseases allow for stopping due to adverse events, and in some cases guidelines have been constructed for when to stop treatment. For some diseases, a recommendation to stop treatment is made typically at the end of a gradual tapering-down of dose for patients who respond well to treatment and are considered to be in remission. For others, patients are given a standard dose and the treatment decision at each time step is of the stop/do-not-stop type. In addition, stopping treatment may occur for patients in poor disease states due to a finding of futility or a desire to switch to a different drug or type of treatment.

2. Literature review. Discontinuation of pharmacological therapy has been studied in a number of diseases. For rheumatoid arthritis (RA), a protocol for discontinuing the biologic agent infliximab has been developed by Maas et al.: patients whose 28-joint disease activity score (DAS28) is below 3.2, and have received stable dose for at least 6 months, have their doses tapered down by 25% of the original dose every 8-12 weeks until discontinuation of treatment is achieved or the patient

2010 *Mathematics Subject Classification.* Primary: 60G40, 62L15; Secondary: 92C50.

Key words and phrases. Optimal stopping, data-driven decision-making, Markov decision process, medical dosing, rheumatoid arthritis.

The first author was funded in part by the National Science Foundation through grant #CMMI 1536717 when the author was a doctoral student at the University of Washington.

experiences a flare-up [23]. Another study on adapting dose of the biologic agent infliximab based on patient response ended up stopping treatment for 7 of 76 patients due to adverse events [4]. One meta-analysis compared gradual lowering of dose (also called “down-titration”) and discontinuation versus continuation of the drugs adalimumab and etanercept in RA patients with low DAS28 scores with mixed results, finding that stopping treatment produces benefits in some, but not all patients [24].

Infliximab is also used to treat other inflammatory bowel diseases (IBD) including Crohn’s disease and ulcerative colitis (UC). Other studies have focused mainly on patient outcomes after the decision to stop infliximab treatment. Several studies have been conducted on the risk of IBD disease relapse after a decision to interrupt treatment of infliximab [11, 12, 19]. A prevalence study found that an “important proportion” of RA patients in remission were directed to down-titrate or discontinue treatment the drug, indicating that the stopping decision is not uncommon in practice, though a patient-specific numerical framework does not exist [10]. One study found that 62% of patients who stopped a second-line drug in combination therapy for RA did not experience a flare within one year; yet patients who continued the second-line drug had a lower chance of flare [22]. A meta-analysis of flare rates for RA patients with low DAS28 scores or in remission found that “more than one-third of patients” may down-titrate or stop disease-modifying anti-rheumatic drugs (DMARD) without risk of a flare for one year [9].

Some clinical trials for hepatitis have included the possibility of stopping treatment within a response-guided framework. A response-guided clinical trial using telaprevir for hepatitis C directed patients with an HCV RNA level greater than 1000 IU per mL at week 4, or who had virologic failure at week 12 or between weeks 24 and 36 of the study, to stop treatment [15]. Jacobson et al. developed stopping rules for patients destined to fail boceprevir-based combination therapy for hepatitis C based on phase 3 trial databases; the rules were then applied retroactively to determine how many patients could have stopped treatment early to minimize drug toxicity, resistance, and costs [6]. Along the same vein, Davis, et al. performed a retrospective analysis of patients taking pegylated interferon alfa-2b and ribavirin for hepatitis C to identify a rule that would have stopped treatment early for some patients; they found that patients who did not achieve an early virologic response of at least 2 logs in the first 12 weeks compared with baseline was predictive of ultimate futility of the therapy, and thus could have been stopped early [3]. Response-guided dosing of peginterferon in hepatitis B studies have established a rule to stop treatment if there is no decline of serum HBsAg levels from baseline to weeks 12 or 24 [20, 21].

While these studies have considered the decision of when to stop treatment, the rules developed are ad hoc, specific to individual drugs and diseases. Stopping treatment is typically considered only at one of a few pre-specified time-points during treatment, and is not considered as an alternative to dosing in each individual session. In addition, stopping is generally considered only for cases of drug futility and not disease remission. Furthermore, the stopping criteria that have been developed for specific drugs and diseases have not been built using a mathematically rigorous optimal stopping approach.

In the operations research literature, several authors have considered optimal stopping rules for clinical trials by creating a stochastic dynamic programming approach. Papers by Berry and Müller et al. considered Bayesian approaches

to phase II clinical trials [1, 13]. Their work, like ours, considers a sequential decision problem where patients in the trial are dosed over discrete sessions. At each time point, three arms are considered: continuation of pharmaceutical treatment; stopping for futility; or stopping for efficacy, with direction to enroll in a phase III clinical trial. Along the way, dose-response parameters are learned. Our work differs from theirs in that we do not consider dosing in the context of a clinical trial; rather, we look at drugs that have already been brought to market and thus information about the dose-response parameter is assumed to be known.

In this paper, we extend the previous stochastic DP model of Kotas and Ghatge [7] to allow for stopping treatment as an alternative to administering dose in any session. That paper modeled the disease progression of an individual patient as a finite-horizon, fixed-length Markov decision process. The optimal solution balances improving the patient’s disease state as much as possible at the end of treatment with the costs incurred due to adverse effects in each treatment session. This paper’s additional contribution is to allow stopping, which in essence adds an additional option to the decision-space, so that in any session a dose may still be administered, or a decision to stop may be made. If the decision-maker stops treatment, then no dose may be administered in future sessions, and as a result no future per-session costs are incurred.

3. Model without stopping. Our model with stopping is an extension of the stochastic DP model for RGD by Kotas and Ghatge [7]. For completeness, we give an overview of that model here.

Let T denote the number of treatment sessions, indexed by $t = 1, 2, \dots, T$, in a treatment course. The time-interval between two consecutive treatment sessions may be hours, days, weeks, or months depending on the disease. For simplicity of notation, we assume that these intervals are equal. At the beginning of each treatment session, the physician observes a numerical score of the patient’s disease condition, and chooses a dose for that session. These numerical scores belong to a convex set $X \subseteq \mathbb{R}$. Smaller real numbers in this set represent less severe disease. The disease condition at the beginning of treatment session t is denoted by $x_t \in X$. The dose level chosen by the physician for this session after observing x_t is denoted by d_t . Possible dose levels d_t belong to the interval $D \triangleq [0, \bar{d}] \subset \mathbb{R}$, where \bar{d} is a finite upper bound on permissible dose levels.

For $t = 1, 2, \dots, T$, disease conditions evolve according to dynamics

$$x_{t+1} = x_t + f(d_t; \theta), \text{ for } x_t, x_{t+1} \in X, \text{ and } d_t \in D, \quad (1)$$

where θ are independent and identically distributed (iid) random variables that take values from a finite set $\Omega \triangleq \{v_1, v_2, \dots, v_k\}$ with probabilities p_1, p_2, \dots, p_k respectively (therefore, θ are Categorically distributed). We assume $f(\cdot; \theta)$ is continuous and non-increasing over D for each $\theta \in \Omega$, indicating an improvement in disease state with increasing dose. Note that the assumption that θ is Categorically distributed is not restrictive as any continuous distribution with finite support can be arbitrarily well approximated by a Categorical distribution with sufficient number of bins k . The distribution is assumed to be known, perhaps as a result of completing a trial on a cohort of patients, as described by [8]. The assumption that the state transition function is additive in state is made for algebraic simplicity and is appropriate for a wide range of dose-response functions used in practice, including exponential, exponential linear-quadratic, logistic, Michaelis-Menten, Hill’s, Emax, power law, Gompertz, and beta-Poisson (in some cases after taking logarithms to

turn functions that are multiplicative in state to additive, and working in log-units) [17].

Aversion to dose is modeled using a continuous cost function $c : D \rightarrow \mathbb{R}_+$. Since D is compact, continuity of $c(\cdot)$ implies that it is bounded. Examples include linear, quadratic, and exponential functions. Aversion to disease conditions x_{T+1} at the end of the treatment course is modeled using a continuous and bounded cost function $h : X \rightarrow \mathbb{R}_+$. Examples include linear, quadratic, exponential, and ramp. The cost in the ramp function is zero up to a threshold and then increases with disease score; this can be used to model the treat-to-target approach.

Let $J_t(x_t)$ denote the minimum total expected cost not yet incurred, given that the disease condition at the beginning of the t th session is x_t . These optimal cost-to-go functions $J_t(\cdot)$ are unique solutions of Bellman's equations

$$J_t(x_t) = \min_{d_t \in D} \left(c(d_t) + \sum_{j=1}^k J_{t+1}(x_t + f(d_t; v_j)) p_j \right),$$

with $J_{T+1}(x_{T+1}) = h(x_{T+1})$. (2)

Problem (2) involves optimizing a continuous function over the nonempty compact set D and hence it has an optimal solution. The set of doses that attain the minimum in (2) define an optimal policy: for state $x_t \in X$ in session t , it is denoted by $A_t^*(x_t) \subseteq D$.

Bellman's equations (2) can be solved approximately easily using backward induction through discretization of X (along with truncation if needed) and of D [2]. Because we do not assume particular properties for the underlying cost functions such as convexity, the problem does not have a specific structure that can be taken advantage of for numerical calculation. Nevertheless, for practical problems of interest and even on a relatively fine grid spacing for a single patient, the problem is computationally tractable through enumeration of the cost-to-go function for all possible doses d_t and then selecting the smallest value in the array.

4. Model with stopping. The aforementioned model took the number of equally-spaced treatment sessions T to be known a priori, with sessions indexed by $t = 1, 2, \dots, T$. Now, we also take T to be known, but allow for the possibility of ending treatment early. Thus, T can be thought of as a loose upper bound on the number of treatment sessions. In session t , we can choose to give a dose, which is denoted d_t , where $d_t \in D \triangleq [0, \bar{d}] \forall t$ where $\bar{d} < \infty$ is the maximum permissible dose in one session; alternatively a decision to stop treatment can be made (for all $t < T$). Note that a dose of zero is not equivalent to a decision to stop; a dose of zero still implies that a patient will return for treatment in the following session, although they receive no dose now. Stopping, on the other hand, implies that the patient will not return for treatment after the current session. For the purposes of graphical illustration, we denote the decision to stop treatment by $d_t = -1$. This value is not physically meaningful, but as it lies outside the permitted dose range $D \triangleq [0, \bar{d}]$, it allows for simple visualization.

If treatment is terminated when the state is $x_t \in X$, where higher values correspond to more severe disease states, the patient derives a terminal cost $h : X \rightarrow \mathbb{R}_+$. If treatment is continued, then a dose level $d_t \in D$ for that session must be chosen with associated cost $c : D \rightarrow \mathbb{R}_+$. The state then stochastically evolves according

to the state transition function

$$x_{t+1} = x_t + f(d_t; \theta), \quad (3)$$

as in the model without stopping described in section 3. The Bellman's equations then become:

$$J_t(x_t) = \min \left\{ \min_{d_t \in D} \left(c(d_t) + \sum_{j=1}^k J_{t+1}(x_t + f(d_t; v_j)) p_j \right), h(x_t) \right\},$$

with $J_{T+1}(x_{T+1}) = h(x_{T+1})$. (4)

where the outer minimization problem chooses the more optimal of continuing treatment or stopping now. The Bellman's equations can be solved numerically using backward induction with discretization of the state-action space.

5. Stopping for rheumatoid arthritis. We reconsider the rheumatoid arthritis problem based on OPTION trial data which was discussed in the Kotas and Ghate model [7, 18]. We begin by reviewing that model.

We seek to determine an optimal dose of the drug tocilizumab in combination therapy with a fixed dose of methotrexate for rheumatoid arthritis. The patient state x_t in session t is taken to be the natural logarithm of the DAS28_t score, a widely-used measure of disease progression in RA, which takes positive values. Doses are administered over a maximum of $T = 7$ sessions, equally spaced at 4 weeks. The dose to be administered in session t is designated as d_t , measured in units of mg tocilizumab per kg body weight, take on values in $[0, 10]$. The dose-response function is taken to be a modified log-Michaelis Menten with additive noise:

$$x_{t+1} = x_t + \ln \kappa_2 - \ln(\kappa_1 + \kappa_2 + d_t) + \theta \quad (5)$$

where $\kappa_1, \kappa_2 > 0$ are parameters indicating the effectiveness of methotrexate and inverse effectiveness of tocilizumab, respectively, and θ are iid Categorical distributed random variables, specifically a discretization of the Normal(0, 0.05²) distribution truncated to within $\pm 3\sigma$, renormalized and discretized to bins of width 0.01σ . Based on data from the OPTION clinical trial, we estimated values for the parameters: $\kappa_1 = 4.5295$ mg/kg and $\kappa_2 = 124.1593$ mg/kg [18]. The initial state is taken to be $x_1 = \ln(6.8)$.

The terminal state cost function is taken to be exponential: $h(x_{T+1}) = \exp(x_{T+1}) = \text{DAS28}_{T+1}$. The per-session cost function was taken to be linear: $c(d_t) = \gamma d_t$, with γ constant, and a value of γ was estimated to be $\gamma = 0.028557$ (mg/kg)⁻¹ using data from the OPTION trial [18]. The determination of these parameter values is described fully in [7]; we omit that here for brevity.

For our stopping variation, all parameters were set to the same values as [7], except we consider a slightly different cost function $c(\cdot)$. We generalize our cost function to the two-parameter linear-affine form $c(d) = \gamma d + b$. Here, γ is still the cost per unit of dose, but we also allow for a fixed cost per session b , which can model administrative overhead, or the cost associated with storing drugs even if they are not administered to a patient, or the cost associated with the patient traveling to a clinic only to receive no dose. With $b = 0$ we recover the linear cost function of [7].

Bellman's equations (2) were solved approximately using backward induction with a discretization of 0.01 for the state x_t for all t , and a discretization of 0.01 for the dose d_t for all t .

Results of numerical simulation are given in Figure 1. Three subfigures are shown corresponding to the values $b = 0$, $b = 0.05$, and $b = 0.1$ which represent increasing levels of fixed cost using the cost function $c(d) = \gamma d + b$. On each plot, the optimal policy is plotted. On week 0, the patient's state is known to be x_1 ; thus just one point of the optimal dose is given. In subsequent sessions, due to the uncertainty of patient response, a range of state values are possible. These are plotted as curves of different colors for each session.

The value of the parameter b represents a fixed per-session cost that is penalized the same amount as the cost associated with a dose of b/γ in one session. For $b = 0.05$, this works out to a dose of approximately 1.75 mg/kg, and for $b = 0.1$, approximately 3.5 mg/kg. These values are significantly below the maximum allowable dose in one session of 10 mg/kg and thus both model the reasonable situation where our primary cost of concern is still the cost associated with dose and not the fixed costs. At the same time, the values of the fixed costs are not so low as to be negligible.

In Figure 1, we make the following observations. With $b > 0$ we see that stopping indeed becomes optimal for some of the lowest disease states. This is intuitive, as for a very low disease state, the certain fixed per-session cost in upcoming sessions outweighs the possibility that the disease state will rise to a point where dose would be given, so a decision to stop is optimal. For intermediate state values, we see a wait-and-see policy, where no dose is given but neither is treatment stopped, as the possibility of a flare-up to the point that positive dose is optimal is higher. For still higher disease states, we observe that a positive dose is given. Finally, we note that as b increases, the threshold state below which stopping is optimal increases. Again this is reasonable, as the decision to stop becomes more appealing the higher the per-session cost becomes.

We also note that when $b = 0$, stopping is never optimal for the state values we consider. In essence the only options become either giving a positive dose, or being in the wait-and-see region. This is also intuitive, since if there is no cost associated with storing the drug, or administrative overhead, or having the patient travel for the appointment, etc., then it is always most beneficial for the patient to return, on the off-chance (no matter how small) that their state has deteriorated to the point where a positive dose would be prescribed. In reality, we would expect a value of b that is always positive, even if small, because these fixed costs do exist; thus the situation where b is strictly zero is unlikely to occur in practice.

6. Monotonicity of stopping threshold state with respect to time. In our numerical experiments, we have observed that if stopping is ever optimal, it is optimal below a threshold state. Let us define this threshold state in session t as x_t^* . One question that naturally arises in the multi-period problem is whether x_t^* is a function of t , and if so, if x_t^* is monotone in one direction or the other.

A zoom-in of Subfigure 1c is given in Figure 2, centered on x_t^* , and with a finer discretization. Figure 2 suggests that x_t^* appears to monotonically increase with respect to session number. That is, the less treatment time remaining, the higher the threshold state between stopping and not stopping.

This result is not only numerically observed, but provable. In fact a proof for a general DP problem with stopping is found in section 4.4, volume 1 of [2]. For convenience we provide a counterpart of this proof using our notation.

By the Bellman's equations (4), it is optimal to stop at time t for all states x_t in the set

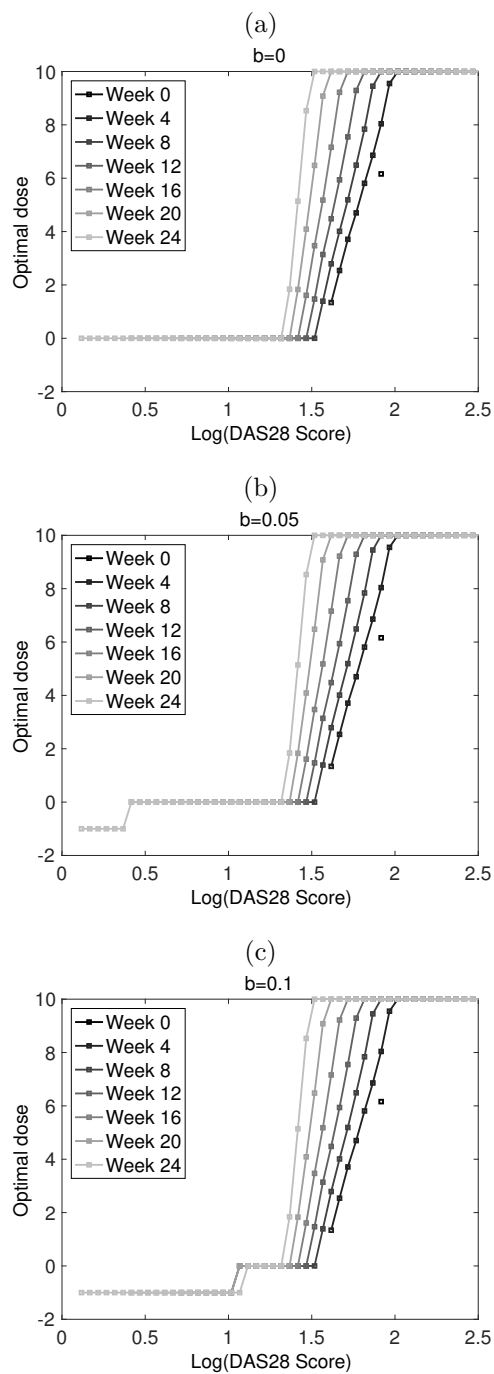


FIGURE 1. Optimal policy of tocilizumab dosing for rheumatoid arthritis. The cost functions are $c(d) = 0.028557d + b$, with fixed cost parameter b varying among subfigures. A dose of -1 indicates a decision to stop treatment.

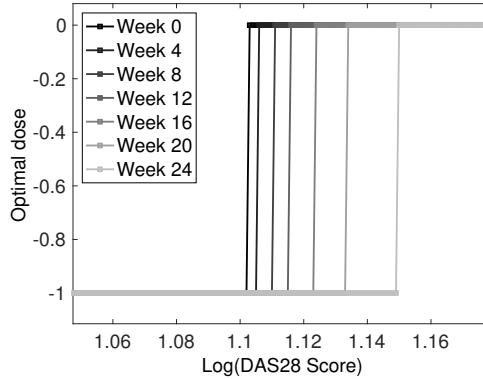


FIGURE 2. Optimal policy with the cost function $c(d) = 0.028557d + 0.1$. The plot is a zoomed-in version of Figure 1c around the threshold area between stopping and not stopping, with refined discretization. A dose of -1 indicates a decision to stop.

$$T_t = \left\{ x_t \mid h(x_t) \leq \min_{d_t \in D} c(d_t) + \sum_{j=1}^k J_{t+1}(x_t + f(d_t; v_j)) p_j \right\} \quad (6)$$

Equation (4) along with the boundary condition of the DP, $J_T(x_T) = h(x_T)$, implies that

$$J_{T-1}(x) \leq J_T(x) \quad \forall x \quad (7)$$

Using equation 4 along with the stationarity property of our problem and the monotonicity property of DP, we obtain via induction

$$J_t(x) \leq J_{t+1}(x) \quad \forall x, t. \quad (8)$$

Using this fact, we see

$$T_1 \subset T_2 \subset \dots \subset T_{N-1}. \quad (9)$$

In our numerical simulations we have observed that $T_t = (-\infty, x_t^*)$ for all t . This combined with equation (9) gives that the upper limit of the stopping set, x_t^* , increases monotonically with t .

7. Conclusions. We have presented an extension of the stochastic DP model of [7] where the decision-maker can decide to stop treatment at any treatment session. If a decision to stop is made in the current period, all future per-session costs are avoided, and the patient's final disease state is taken to be the current state. Intuitively, we expect the decision to stop will be optimal, if ever, at low disease states. At these states, the future per-session cost outweighs the benefit of lowering the disease state through treatment. In some cases, we may also find the existence of a wait-and-see region, where zero dose is given in a particular session, but the decision to stop is not made— this can incur a per-session cost, but the possibility of giving a dose later to lower the disease state outweighs that per-session cost, so we continue. At the highest disease states, positive doses are given as the benefit of reducing disease state wins out over the per-session costs.

We reconsidered the rheumatoid arthritis example of [7] again, but this time allowing for stopping. For the original problem, stopping was never optimal over

the states considered. However, by adding a fixed per-session cost $b > 0$ to the cost function $c(d) = \gamma d + b$, we found that stopping is optimal for some of the lowest disease states. This indicates that stopping is optimal when the future fixed per-session costs outweigh the potential benefit of giving dose later.

In the literature, stopping is mentioned not only for patients in very low disease states (remission,) but also sometimes for very high disease states, as this indicates a failure of the drug to have an effect on the patient. In practice, this would often indicate the need to switch to a different drug or treatment scheme. As we were only considering the dose of a single drug in our framework, this situation did not arise for us, but could be another interesting direction for future work.

REFERENCES

- [1] D. A. Berry and P. Müller and A. P. Grieve and M. Smith and T. Parke and R. Blazek and N. Mitchard and M. Krams, [Adaptive Bayesian designs for dose-ranging drug trials](#), *Case Studies in Bayesian Statistics*, **5** (2002), 99–181.
- [2] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 4th edition, Athena Scientific, Cambridge, MA, 2012.
- [3] G. L. Davis, J. B. Wong, J. G. McHutchison, M. P. Manns, J. Harvey and J. Albrecht, [Early virologic response to treatment with peginterferon alfa-2b plus ribavirin in patients with chronic hepatitis C](#), *Hepatology*, **38** (2003), 645–652.
- [4] M. Flendrie, M. C. W. Creemers and P. L. C. M. van Riel, [Titration of infliximab treatment in rheumatoid arthritis patients based on response patterns](#), *Rheumatology*, **46** (2007), 146–149.
- [5] B. I. Grigelionis and A. N. Shiryaev, On Stefan’s problem and optimal stopping rules for Markov processes, *Theory of Probability and its Applications*, **4** (1966), 541–558.
- [6] I. M. Jacobson, P. Marcellin, S. Zeuzem, M. S. Sulkowski, R. Esteban, F. Poordad, S. Bruno, M. H. Burroughs, L. D. Pedicone, N. Boparai, W. Deng, M. J. DiNubile, K. M. Gottesdiener, C. A. Brass, J. K. Albrecht and J.-P. Bronowicki, [Refinement of stopping rules during treatment of hepatitis C genotype 1 infection with boceprevir and peginterferon/ribavirin](#), *Hepatology*, **56** (2012), 567–575.
- [7] J. Kotas and A. Ghate, Response-guided dosing for rheumatoid arthritis, *IIE Transactions on Healthcare Systems Engineering*, **6** (2016), 1–21.
- [8] J. Kotas and A. Ghate, [Bayesian learning of dose-response parameters from a cohort under response-guided dosing](#), *European Journal of Operational Research*, **265** (2018), 328–343.
- [9] T. M. Kuijper, F. B. G. Lamers-Karnebeek, J. W. G. Jacobs, J. M. W. Hazes and J.J. Luime, [Flare rate in patients with rheumatoid arthritis in low disease activity or remission when tapering or stopping synthetic or biologic DMARD: a systematic review](#), *Journal of Rheumatology*, **42** (2015), 2012–2022.
- [10] L. Lojo, G. Bonilla, D. Peiteado, A. Villalba, C. Plasencia, L. Nuno, A. Balsa and E. Martin-Mola, [Down-titration and discontinuation of infliximab, adalimumab and etanercept in established rheumatoid arthritis](#), *Annals of the Rheumatic Diseases*, **72** (2013), 237.
- [11] E. Louis, G. Vernier-Massouille, J.-C. Grimaud, Y. Bouhnik, D. Laharie, J.-L. Dupas, H. Pillant, L. Picon, M. Veyrac, M. Flamant, G. Savoye, R. Jian, M. De Vos, G. Paintaud, E. Piver, J.-F. Colombel, J.-Y. Mary and M. Lemann, [Infliximab discontinuation in Crohn’s disease patients in stable remission on combined therapy with immunosuppressors: a prospective ongoing cohort study](#), *Gastroenterology*, **136** (2009), A146.
- [12] M. Marino, E. Zucchi, M. Fabbro, I. Lodolo, R. Maieron, S. Vadalà and M. Zilli, [Outcome of infliximab discontinuation in IBD patients and therapy rechallenging in relapsers: Single centre preliminary data](#), *Journal of Crohn’s and Colitis*, **8** (2014), S232–S233.
- [13] P. Müller, D. A. Berry, A. P. Grieve and M. Krams, A Bayesian decision-theoretic dose-finding trial, *Decision Analysis*, **3** (2006), 197–207.
- [14] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, Inc., New York, 1994.
- [15] K. E. Sherman, S. L. Flamm, N. H. Afdhal, D. R. Nelson, M. S. Sulkowski, G. T. Everson, M. W. Fried, M. Adler, H. W. Reesink, M. Martin, A. J. Sankoh, N. Adda, R. S. Kauffman, S. George, C. I. Wright, F. Poordad for the ILLUMINATE study team, [Response-guided](#)

- telaprevir combination treatment for hepatitis C virus infection, *The New England Journal of Medicine*, **365** (2011), 1014–1024.
- [16] A. N. Shiryaev, *Optimal Stopping Rules*, 1st edition, Springer-Verlag, Berlin Heidelberg, 2008.
- [17] W. Slob, Dose-response modeling of continuous endpoints, *Toxicological Sciences*, **66** (2002), 298–312.
- [18] J. Smolen, A. Beaulieu, A. Rubbert-Roth, C. Ramos-Remus, J. Rovensky, E. Alecock, T. Woodworth, R. Alten and OPTION trial investigators, Effect of interleukin-6 receptor inhibition with tocilizumab in patients with rheumatoid arthritis (OPTION study): a double-blind, placebo-controlled, randomised trial, *The Lancet*, **371** (2008), 987–997.
- [19] A. Sofi, A. Ali, S. A. Khuder and A. Nawras, Meta-analysis– maintenance of remission following discontinuation of infliximab in patients with Crohn’s disease, *Gastroenterology*, **144** (2013), S637.
- [20] M. J. Sonneveld, B. E. Hansen, T. Piratvisuth, J.-D. Jia, S. Zeuzem, E. Gane, Y.-F. Law, Q. Xie, E. J. Heathcote, H.L.-Y. Chan and H.L.A. Janssen, Response-guided peginterferon therapy in hepatitis B e antigen-positive chronic hepatitis B using serum hepatitis B surface antigen levels, *Hepatology*, **58** (2013), 872–880.
- [21] M. J. Sonneveld, V. Rijckborst, C. A. Boucher, B. E. Hansen and H.L. Janssen, Prediction of sustained response to peginterferon alfa-2b for hepatitis B e antigen-positive chronic hepatitis B using on-treatment hepatitis B surface antigen decline, *Hepatology*, **52** (2012), 1251–1257.
- [22] S. ten Wolde, F. C. Breedveld, B. A. C. Dijkmans, J. Hermans, J. P. Vandenbroucke, M. A. F. J. van de Laar, H. M. Markusse, M. Janssen and H. R. van den Brink, Randomised placebo-controlled study of stopping second-line drugs in rheumatoid arthritis, *The Lancet*, **347** (1996), 347–352.
- [23] A. van der Maas, W. Kievit, B. J. F. van den Bemt, F. H. J. van den Hoogen, P. L. C. M. van Riel and A. A. den Broeder, Down-titration and discontinuation of infliximab in rheumatoid arthritis patients with stable low disease activity and stable treatment: an observational cohort study, *Annals of Rheumatoid Disease*, **71** (2012) 1849–1854.
- [24] N. van Herwaarden, A. A. den Broeder, W. Jacobs, A. van der Maas, J. W. Bijlsma, R. F. van Vollenhoven and B. J. van den Bemt, Down-titration and discontinuation strategies of tumor necrosis factor-blocking agents for rheumatoid arthritis patients with low disease activity, *Cochrane Database of Systematic Reviews*, **9** (2013).

Received for publication April 2018.

E-mail address: kotas@up.edu