Check for updates

RESEARCH ARTICLE

UPDATE Dose Titration Algorithm Tuning (DTAT) should supersede 'the' Maximum Tolerated Dose (MTD) in oncology dose-finding trials [version 3; referees: 3 approved, 1 approved with reservations]

David C. Norris ID

Precision Methodologies, LLC, Seattle, USA

Abstract

Background. Absent adaptive, individualized dose-finding in early-phase oncology trials, subsequent 'confirmatory' Phase III trials risk suboptimal dosing, with resulting loss of statistical power and reduced probability of technical success for the investigational therapy. While progress has been made toward explicitly adaptive dose-finding and quantitative modeling of dose-response relationships, most such work continues to be organized around a concept of 'the' maximum tolerated dose (MTD). The purpose of this paper is to demonstrate concretely how the aim of early-phase trials might be conceived, not as 'dose-finding', but as *dose titration algorithm (DTA)*-finding. **Methods.** A Phase I dosing study is simulated, for a notional cytotoxic chemotherapy drug, with neutropenia constituting the critical dose-limiting toxicity. The drug's population pharmacokinetics and myelosuppression dynamics are simulated using published parameter estimates for docetaxel. The amenability of this model to linearization is explored empirically. The properties of a simple DTA targeting neutrophil nadir of 500 cells/mm^3 using a Newton-Raphson heuristic are explored through simulation in 25 simulated study subjects. **Results.** Individual-level myelosuppression dynamics in the simulation model approximately linearize under simple transformations of neutrophil concentration and drug dose. The simulated dose titration exhibits largely satisfactory convergence, with great variance in individualized optimal dosing. Some titration courses exhibit overshooting. **Conclusions.** The large inter-individual variability in simulated optimal dosing underscores the need to replace 'the' MTD with an individualized concept of MTD_i . To illustrate this principle, the simplest possible DTA capable of realizing such a concept is demonstrated. Qualitative phenomena observed in this demonstration support discussion of the notion of *tuning* such algorithms. Although here illustrated specifically in relation to cytotoxic chemotherapy, the DTAT principle appears similarly applicable to Phase I studies of cancer immunotherapy and molecularly targeted agents.

Open Peer Review

Referee Status: ✔ ? ✔ ✔

	Invited Referees			
	1	2	3	4
UPDATE version 3 published 17 Jul 2017				
REVISED version 2 published 30 Mar 2017			✔ report	✔ report
version 1 published 07 Feb 2017	✔ report	? report		

1 **Matthew E. Nielsen**, University of North Carolina at Chapel Hill, USA

2 **Natalja Strelkowa**, Boehringer Ingelheim Pharma GmbH & Co. KG, Germany

3 **Felipe Maglietti** ID , University of Buenos Aires, Argentina

4 **Yuan Ji**, University of Chicago, USA

Discuss this article

Comments (2)

Corresponding author: David C. Norris (david@precisionmethods.guru)

How to cite this article: Norris DC. **Dose Titration Algorithm Tuning (DTAT) should supersede 'the' Maximum Tolerated Dose (MTD) in oncology dose-finding trials [version 3; referees: 3 approved, 1 approved with reservations]** *F1000Research* 2017, **6**:112 (doi: 10.12688/f1000research.10624.3)

First published: 07 Feb 2017, **6**:112 (doi: 10.12688/f1000research.10624.1)

UPDATE   Updates from Version 2

This paper has been updated to remove incorrect copyright licences, which appeared in the previous version 2.

Amendments from Version 1

Changes in v2, stimulated largely by Dr. Natalja Strelkowa's reviewer comments, follow closely my posted reply. Chiefly, I have sought in v2 to separate the merely incidental aspects of the concrete simulation here presented from the essential content of the DTAT principle. In particular, (a) I now directly contrast DTAT's dose titration algorithm abstraction against the fallacious abstraction of 'the' MTD, introducing the separate acronym 'DTA' specifically for added emphasis, and indicating a connection with Karl Popper's concept of objective knowledge; (b) I stress that recursive filtering methods (such as the Kalman filter) have contributed heuristically to DTAT, but in no way delimit or define DTAT; (c) I discuss at some length (adding 3 new citations) how the DTAT principle applies beyond cytotoxic chemotherapy as simulated here – specifically, to cancer immunotherapies and molecularly targeted agents; (d) I intensify my invective against 'the' MTD, directly condemning it as a fallacy of misplaced concreteness, and employing the scare quotes now even in my title. Additionally, Dr. Strelkowa's very important point about the need for fail-safe upper bounds in DTAs receives stress both in the main text and in the caption to Figure 5. Finally, I strike the suggestions made in v1 that a linear Kalman filter might improve on the Newton-Raphson method used here, since I have realized that the transformations which 'linearize' Figure 4 merely simplify, but do not actually linearize, the nadir dynamics. (As I indicate in the text, full-information methods now appear in any case the most promising general line of attack.)

See referee reports

Introduction

Despite advances in Bayesian adaptive designs[1,2] and model-based dose-finding[3], oncology dose-finding studies remain conceptually in the thrall of 'the' maximum tolerated dose (MTD). This fallacious concept stands opposed to the long-recognized heterogeneity of cancer patients' pharmacokinetics and pharmacodynamics (PK/PD), and to the diversity of their individual values and goals of care. Under this conceptual yoke, these dose-finding studies constitute a significant choke-point in drug development, where a severe discount may be applied to the potential value in new molecules through the hobbling of subsequent 'efficacy' trials by inadequate individual-level dosing[4].

Strangely, Bayesian innovation in dose-finding studies has proceeded apace without issuing a meaningful challenge to the inherently frequentist conception of an MTD as determined by whole-cohort frequencies of dose-limiting toxicities (DLTs). Thus, even as Bayesianism has made progress toward the ethical imperative of *efficient use of data*[5] in such studies, it has neglected to confront the distinct ethical dimension of *individualism*[6]. This seems a great irony, as the dynamic learning model of Bayesianism is equally suited, and indeed equally essential, to solving the latter problem.

This paper demonstrates individualized dose-finding in a simulated Phase I study of a cytotoxic chemotherapy drug for which neutropenia constitutes the critical dose-limiting toxicity. Importantly, myelosuppression is interpreted also as a monotone

index of therapeutic efficacy, without the added complication of a dose-response 'plateau'[7] such as postulated for molecularly targeted agents (MTAs). This creates a problem setting where simple heuristics apply, simplifying the demonstration undertaken here. The aim of this exercise is to elaborate a concrete setting in which 'dose-finding study' may readily be seen as a misnomer. Under the view advanced here, early-phase studies of this kind should be conceived as *dose titration algorithm tuning* (DTAT) studies.

The idea that 'dose finding studies' should yield dose titration algorithms (DTAs) is not new. More than a quarter-century ago, Sheiner and colleagues[8] advocated a learn-as-you-go concept for "dose-ranging studies", addressing concerns about "parallel-dose designs" that are not far removed from the motivations for the present paper. As in the advocacy of Sheiner *et al.*, parametric models play an important role in this paper, although in keeping with a spirit of pragmatism I relax this dependence to some extent by means of a semiparametric dynamic on-line learning heuristic.

Methods

A hypothetical cytotoxic drug is considered, modeled notionally after docetaxel, to be infused in multiple 3-week cycles. The pharmacokinetics are taken to follow a 2-compartment model with parameters as estimated for docetaxel in a recent population pharmacokinetic study[9]. Chemotherapy-induced neutropenia (CIN) is taken to follow a myelosuppression model due to Friberg *et al.*[10]. Together, these models form a population pharmacokinetic/ pharmacodynamic (PK/PD) model within which DTAs may be simulated and tuned for optimality. For simulation purposes, and anticipating the future value of ready access to a variety of inference procedures in follow-on work, this PK/PD model is implemented in R package **pomp**[11]. R version 3.3.2 was used[12].

Basic behaviors of the models are illustrated by simulation graphics generated for 25 individuals randomly generated from the population PK/PD model. Properties with specific relevance to absolute neutrophil count nadir (ANC_{nadir})-targeted dose titration are then investigated, with an eye to demonstrating the predictability of nadir timing. In particular, an approximate linearization of neutrophil nadir level and timing is demonstrated, achieved through suitable power-law transformation of infusion doses and logarithmic transformation of neutrophil concentration. Within this transformed parameter space, a simple recursive DTA is defined on the basic heuristic of the Newton-Raphson method for root-finding. For simplicity, monitoring of CIN is not modeled endogenously to this algorithm, but is treated as exogenous such that nadir timing and level are known precisely. A 'DTAT' study is simulated and visualized for 25 patients, with the *tuning parameters* of the recursive titration algorithm held fixed. The visualization supports a discussion of how these parameters might be *tuned* over the course of a Phase I study. All simulations and figures in this paper were generated by a single R script, archived on OSF[13].

Pharmacokinetic model

We take the population pharmacokinetics of our cytotoxic drug to obey a 2-compartment model, with parameters drawn notionally from estimates published for docetaxel [9, Table 2]; see Table 1.

Figure 1 shows illustrative pharmacokinetic profiles for 25 randomly-generated individuals from this population, administered a 100 mg dose.

Myelosuppression model

Chemotherapy-induced neutropenia is simulated using the 5-compartment model of Friberg *et al.* [10, Table 4], in which myelocytes (here, neutrophils) arise from progenitor cells in a proliferative compartment, mature through a series of 3 transitional states, and emerge into the systemic circulation; see Figure 2. Transit between successive compartments in this model is a Poisson process with time constant k_{tr}, total mean transit time being therefore given by $MTT = 4/k_{tr}$. See Table 2.

Figure 3 shows illustrative myelosuppression profiles for 25 randomly-generated individuals from this population, administered a 100 mg dose.

Linearizing CIN dynamics by dose rescaling

When parametrized by $\sqrt[4]{dose}$, individuals' trajectories in $(\log(ANC_{nadir}) \times time_{nadir})$-space may be approximately linearized, as shown in Figure 4. This linearization recommends the fourth-root and logarithmic transformations employed hereafter for drug dose and ANC.

Dose titration

Recursive nonlinear filtering, as implemented in the extended Kalman filter (EKF) or its more modern adaptations[14], constitutes a powerful conceptual framework for approaching model-based dose

Table 1. Two-compartment pharmacokinetics of docetaxel, from Onoue *et al.* (2016) [9, Table 2]. *CL*: clearance; *Q*: intercompartmental clearance; V_c: volume of central compartment; V_p volume of peripheral compartment; *CV*: coefficient of variation. (*) A CV for V_c was unavailable in 9 and has been set arbitrarily to 0.1.

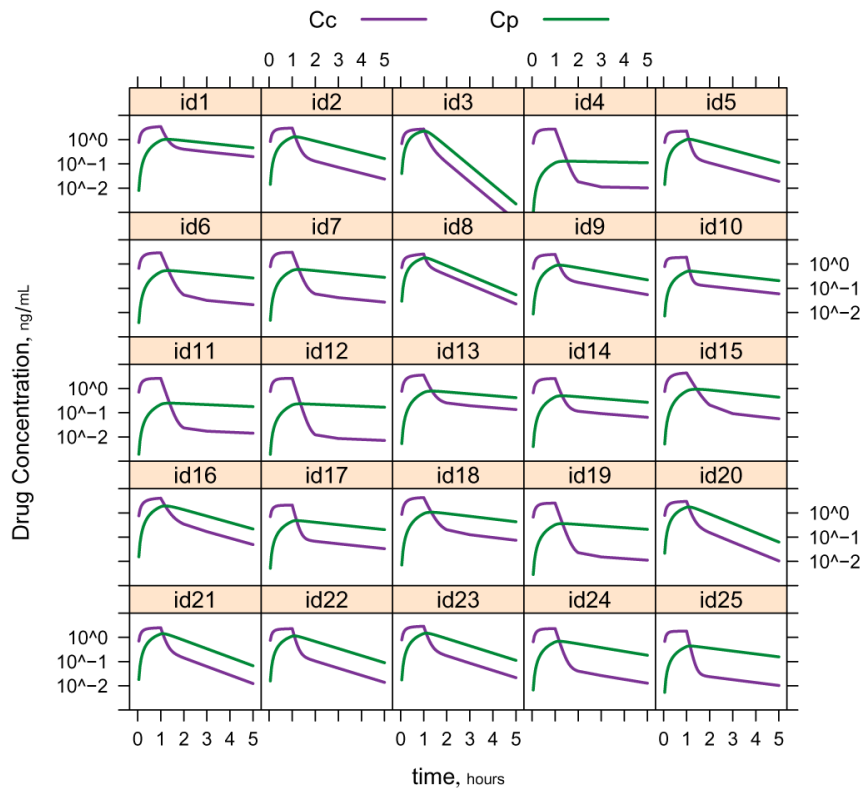| Param	Units	Mean	CV
CL	L/hour	32.6	0.295
Q	L/hour	5.34	0.551
V_c	L	5.77	0.1*
V_p	L	11.0	0.598

Figure 1. Two-compartment pharmacokinetics of a 1-hour infusion of 100 mg of the modeled drug, for 25 randomly generated individuals in our population pharmacokinetic model. C_c and C_p are drug concentrations in the central and peripheral compartments, respectively.

$$k_{prol} = k_{tr} \left(\frac{Circ}{Circ_0}\right)^{-\gamma}$$

$Prol$

k_{tr}

$Transit_1$

k_{tr}

$Transit_2$

k_{tr}

$Transit_3$
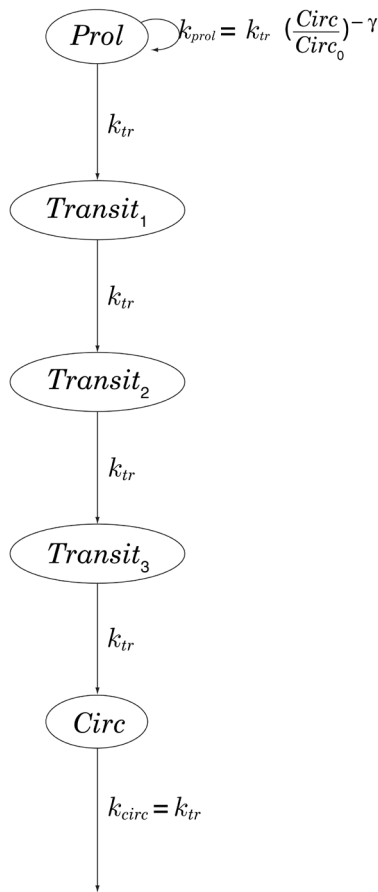
k_{tr}

$Circ$

$k_{circ} = k_{tr}$

Figure 2. Chemotherapy-induced myelosuppression model of Friberg et al. [10, Table 4]. *Prol*: proliferative compartment; *Transit$_n$*: maturation compartments; *Circ*: systemic circulation; k_{tr}: transition rate; k_{prol}: rate of proliferation of progenitor cells, regulated by a negative-feedback loop parametrized by $\gamma > 0$.

Table 2. Parameters of the chemotherapy-induced myelosuppression model of Friberg et al. [10, Table 4]. *Circ$_0$*: baseline neutrophil concentration; *MTT*: mean transit time between the 5 model compartments; γ: exponent of feedback loop; EC_{50}, E_{max}: parameters of a model (of the standard E$_{max}$ type) governing docetaxel-induced depletion in the proliferative compartment.

Param	Units	Mean	CV
$Circ_0$	cells/mm^3	5050	0.42
MTT	hours	89.3	0.16
γ	-	0.163	0.039
E_{max}	μM	83.9	0.33
EC_{50}	μM	7.17	0.50

titration[15]. Indeed, the 'tuning' in 'DTAT' was itself suggested by the practice of tuning a Kalman filter[16] for optimal performance.

For present purposes, however, it suffices to implement a model-free recursive titration algorithm built on the Newton-Raphson method, with a numerically-estimated derivative based on most recent infusion doses and their corresponding ANC nadirs. In this algorithm, a relaxation factor $\omega = 0.75$ is applied to any proposed dose *increase*, with safety in mind. Whereas the slope of $\log(ANC_{nadir})$ with respect to $\sqrt[3]{dose}$ is expected to be strictly negative *at steady state*, hysteresis effects arising during initial steps of dose titration do sometimes yield positive numerical estimates for this slope; so the slope estimates are constrained to be ≤ 0. The infusion dose for cycle 1 is 50 mg, and the cycle-2 dose is calculated conservatively using a slope -2.0, which is larger (in absolute terms) than for any of our simulated patients except `id1` and `id13`; see Figure 4. For reference, these starting values for the *tuning parameters* of the titration algorithm are collected in Table 3.

With the illustrative purpose of this article again in mind, we treat *neutropenia monitoring* as an exogenous process yielding precise nadir timing and levels. This enables a demonstration of the main point without the encumbrance of additional modeling infrastructure peripheral to the main point.

On 'tuning'

If one considers Figure 5 as a *sequence* of titration outcomes emerging in serially enrolled study subjects, it becomes clear that even quite early in the study it will seem desirable to 'retune' the titration algorithm. For example, provided that course-1 CIN monitoring is implemented with sufficient intensity to deliver advance warning of an impending severely neutropenic nadir, so that timely colony-stimulating factor may be administered prophylactically[17], then upon review of the titration courses in the first 10 subjects it may well appear desirable to increase $dose_1$ from 50mg to 100mg. Likewise, given the third-dose overshooting that occurs in 4 of the first 10 subjects, it may seem desirable to adjust the relaxation factor ω downward. Of note, at any given time any such proposed retuning may readily be subjected to a 'dry run' using retrospective data from all convergent titration courses theretofore collected. (Hysteresis effects would however be inaccessible to a strictly data-driven dry run absent formal modeling that captures such effects.) Furthermore, the 'tuning' idea readily generalizes to the fundamental modification or even wholesale replacement of a DTA. For example, the overshooting seen for subjects `id8`, `id10`, `id12` and `id23`, inspires further thought about refining (or replacing) the admittedly very naive Newton-Raphson method employed herein. (At the very least, DTAs deployed in actual DTAT studies must incorporate fail-safe upper bounds both on *absolute dose* and on *proportional dose* increases.)

A further dimension of 'tuning' that must be discussed is the potential for driving the tuning parameters using statistical models
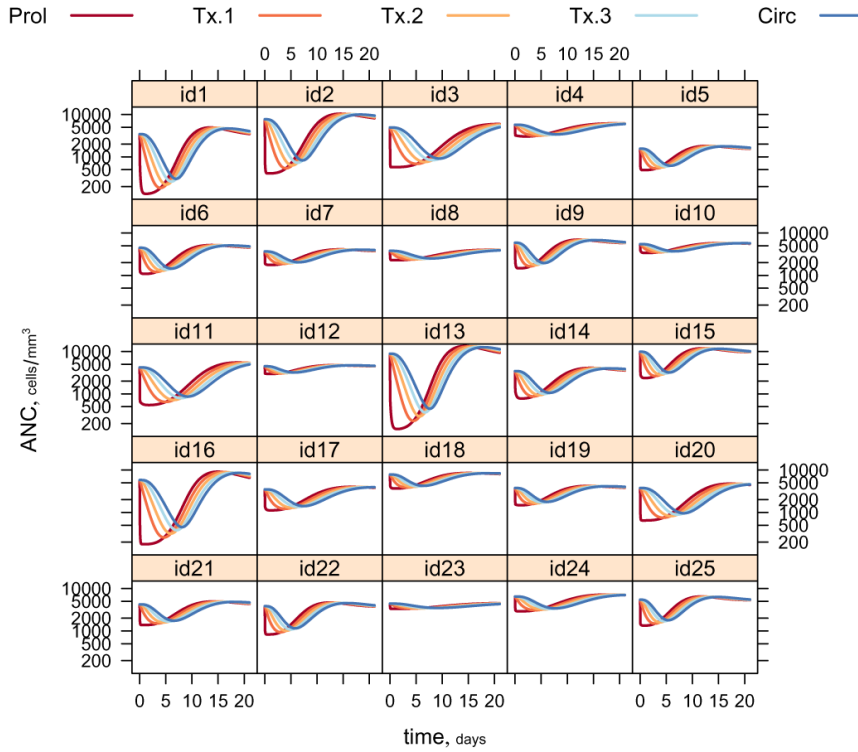
Figure 3. Myelosuppression profiles for the same 25 randomly generated individuals as in Figure 1. Note how a chemotherapeutic 'shock' to the proliferative compartment *Prol* propagates through the maturation compartments $Tx_{1,2,3}$ and thence to the systemic circulation *Circ*. (*ANC*: absolute neutrophil count.)

Figure 4. Trajectories of ANC nadirs during dose escalation in 25 randomly-generated individuals. The 10 doses plotted are evenly spaced on a fourth-root scale. Not only are the trajectories themselves nearly linear in $(\log(ANC) \times time)$-space, but each one is traversed at roughly 'constant velocity' with respect to $\sqrt[4]{dose}$.

Table 3. **Values of the *tuning parameters* of the dose titration algorithm simulated in Figure 5.**
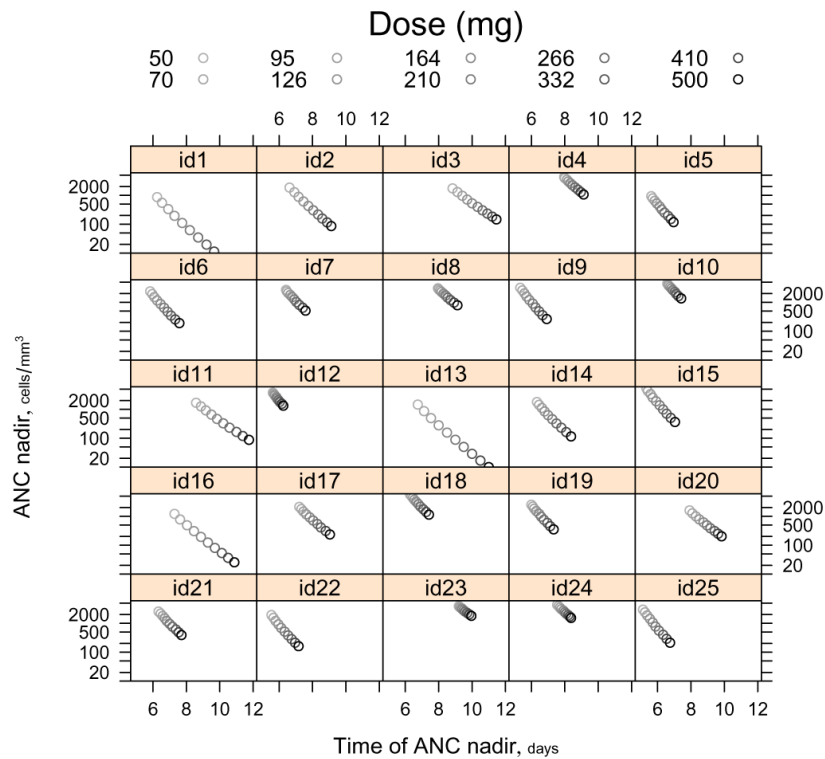
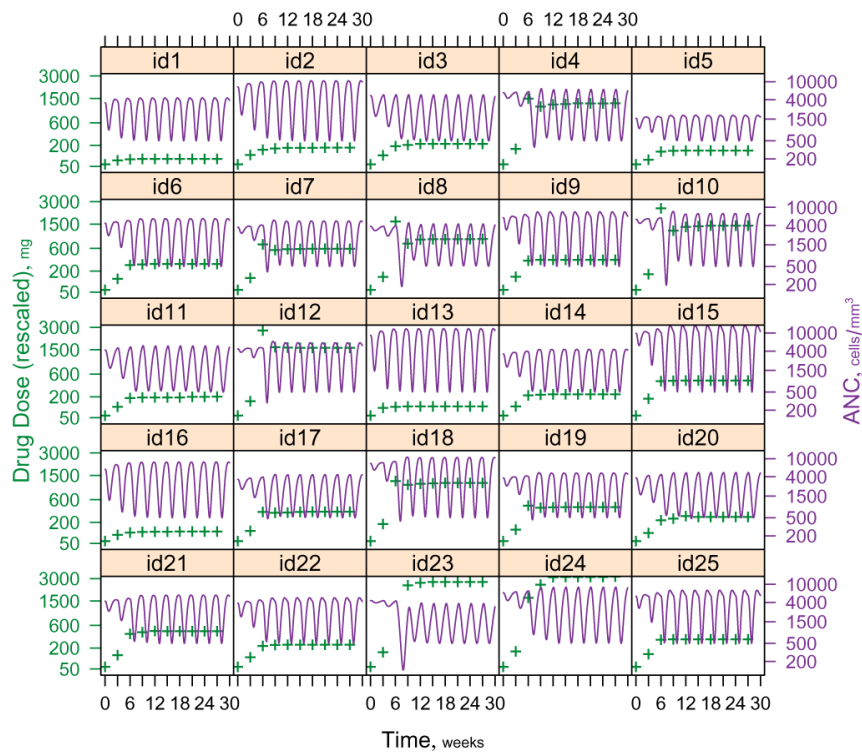| Param	Description	Value
ω	Relaxation factor	0.75
$slope_1$	Slope for cycle-2 dosing	−2.0
$dose_1$	Initial (cycle-1) dose	50 mg

Figure 5. Titration profiles in 25 simulated patients over ten 3-week cycles of chemotherapy. Note particularly the overshooting that occurs in subjects `id8`, `id10`, `id12` and `id23`. This underscores the importance, in actual DTAT applications, of imposing fail-safe upper bounds both on *absolute dose* and on *relative dose increases*. It also bears noting that the 25 MTD_i's evident in this figure span 1.5 orders of magnitude.

built on baseline covariates. Surely, to the extent that the great heterogeneity in final dosing evident in Figure 5 could be predicted based on *age*, *sex*, *weight* or indeed on pharmacogenomic testing, then $dose_1$ should be made a function of these covariates. The *recalibration* of such models as data accumulate from successive study subjects is very much a part of the full concept of 'tuning' I wish to advance.

Finally, whereas I have discussed 'tuning' here largely in terms of reflective, organic decision-making such as occurs in the creative refinement of algorithms or in data-driven statistical model development, I do not mean to exclude more formal approaches to algorithm tuning. A decision-theoretic framing of the tuning

problem should enable formal algorithm tuning to be specified and carried out meaningfully. Such framing would also have the salutary effect of bringing into view *objectively* the important matter of patients' heterogeneity with respect to values and goals of care. It seems quite likely that the balance of benefits from *aggressive titration* versus harms of *toxicities* will generally differ from one patient to another. Dose titration algorithms should most emphatically be tuned to these factors as well. For example, if a patient with more advanced disease and short expected survival nevertheless decided to enroll in a Phase I DTAT study to pursue the possibility of therapeutic benefit, then this patient's decision would indicate a subjective weighting of benefits vs. harms favoring a *higher starting dose* and *more aggressive titration*.

Discussion

It is where pharmacometrics meets the field of *optimal control* that the current literature seems to make its closest point of contact with the DTAT concept I am advancing here. In optimal-control investigations of chemotherapy[18–23], as in DTAT, relatively large decision spaces are explored. Indeed, the infinite-dimensional spaces of *control functions* posited for exploration in optimal control applications dwarf the finite-dimensional spaces of *tuning parameters* in DTAT as dramatically as the latter dwarf the finite sets of *discrete doses* trialed in now-standard Phase I studies. This intermediate 'cardinality' of DTAT reflects an important advantage in an era when, to almost universal chagrin, the detested 3+3 dose-finding design retains its hegemony due partly to widespread resistance to modeling[24]. In such an era, optimal control applications that involve detailed mathematical modeling of *tumor biology and dynamics* sadly seem consigned to the fringes of practice. Acceptance of such ambitious problem formulations, expressing as they do the spirit of a future age, must await deep cultural changes in the medical sciences and clinical practice.

As easy as it is, however, to disparage 'resistance to modeling' as some kind of antediluvian attitude, this resistance does rightfully assert the importance of unmodeled complexities that necessitate application of organic forms of clinical judgment[25]. It should be clear from the above discussion of 'tuning' that DTAT readily accommodates and veritably *invites* scrutiny, supervision and modification by clinical judgment. For example, if during the course of our DTAT study adverse effects other than neutropenia were to emerge as occasional dose-limiting toxicities, then the full concept of 'tuning' advanced above would invite dynamic, 'learn-as-you-go' modifications of the titration algorithm. Such modifications could begin with decreasing the relaxation factor ω, but might also involve efforts to classify and predict these new DLTs, and to incorporate such new understanding explicitly into the DTA yielded by the study. Indeed, whatever philosophical challenge DTAT embodies is likely to take the form of requiring an *intensified* commitment to clinical judgment, in a learn-as-you-go world where the always-provisional nature of medical knowledge must frankly be acknowledged[6,26].

I have presented the *DTAT principle* here embedded in the context of a specific simulation study. This creates the need explicitly to demarcate what I wish to advance as essential in DTAT, from what is merely incidental to the illustration offered here. DTAT makes its fundamental contribution in putting forward a **new abstraction** (the *DTA* with its *tuning parameters*) capable of *embodying knowledge objectively*[27], to supersede a **fallacious abstraction** ('the' MTD) that almost completely lacks this capability. I use the term *fallacious* advisedly, meaning to identify *'the' MTD* specifically with what Whitehead called the "*fallacy of misplaced concreteness* [which] consists in neglecting the degree of abstraction involved when an actual entity is considered merely so far as it exemplifies certain categories of thought[28]." Indeed, the 1.5 orders of magnitude spanned by the MTD$_i$'s of Figure 5 show the degree of abstraction involved in *'the' MTD* to be so egregious as to render this concept plainly useless for embodying what we need to learn from Phase I oncology studies.

Several aspects of the illustration offered here should be clearly understood as *not* essential to the DTAT principle. Firstly, notwithstanding the important heuristic role it has played in the development (and even the naming) of DTAT, *recursive filtering* in no way delimits DTAT. In fact, I now rather suspect that full-information methods will push recursive filtering to the sidelines in practical DTAT applications, and that whatever utility recursive filtering retains will derive from its use as a vehicle for illustrating DTAT to clinicians, perhaps in nomogram forms[15]. Secondly, although neutropenia-targeted dosing of a cytotoxic chemotherapy drug has provided a most propitious context for the present simulation, DTAT need not be thought limited to such drugs. In the important area of immuno-oncology, common dose-limiting toxicities (DLTs) admit monitoring on time-scales comparable to the chemotherapy induced neutropenia (CIN) simulated here. For example, the *cytokine release syndrome* (CRS) that accompanies chimeric antigen receptor (CAR-)T cell therapies typically arises within 1 week of administration (even earlier with concomitant high-dose IL-2) and constitutes a clinical *syndrome* that admits multivariate monitoring on numerous quantitative clinical and laboratory measures[29]. Even molecularly targeted agents (MTAs), for which *late* toxicities have attracted the lion's share of attention[30], remain accessible to the DTAT principle—with DTA learning occurring on a longer time scale. Of course, a DTA that reacts to diverse, lower-grade MTA toxicities[31] that patients *experience and evaluate subjectively* may resemble a process of ongoing *shared decision making* (with the oncology care team) more closely than it resembles the impersonal calculations we typically think of as 'algorithmic'. But with a suitably broadened understanding of 'algorithm'—one that accommodates what might typically be termed *protocols*—the DTAT (or perhaps, DT*PT*) principle continues to apply. In such applications, *supervision and modification by clinical judgment* as mentioned above clearly comes to the fore. But even then, the development and application of scoring systems for patient-reported clinical symptoms and quality of life would enable dose titration protocols (DT*P*s) to be described objectively in quite 'algorithmic' terms that would preserve the applicability of a 'tuning' concept.

Conclusions

I have advanced a concept of *dose titration algorithm tuning* (DTAT), drawing illustrative and orienting connections with *recursive filtering* and *optimal control*. I have concretely illustrated key elements of DTAT by simulating neutrophil-nadir-targeted titration of a hypothetical cytotoxic chemotherapy drug with pharmacokinetics and myelosuppressive dynamics patterned on previously estimated population models for docetaxel. I have also discussed the applicability of DTAT to other types of anti-cancer therapy. I believe DTAT presents a *prima facie* case for discarding the outmoded concept of 'the' maximum tolerated dose (MTD) of cancer therapeutics. This argument should be of interest to a wide range of stakeholders, from cancer patients with a stake in receiving optimal individualized 'MTD$_i$' dosing, to shareholders in pharmaceutical innovation with a stake in efficient dose-finding before Phase III trials.

Data availability

Open Science Framework: Code and Figures for v1 of F1000Research submission: Dose Titration Algorithm Tuning (DTAT) should supersede the Maximum Tolerated Dose (MTD) concept in oncology dose-finding trials, doi 10.17605/osf.io/vwnqz[13]

Endorsement

Daniela Conrado (Associate Director, Quantitative Medicine at Critical Path Institute) confirms that the author has an appropriate level of expertise to conduct this research, and confirms that the submission is of an acceptable scientific standard. Daniela Conrado declares she has no competing interests.

Competing interests

The author operates a scientific and statistical consultancy focused on precision-medicine methodologies such as those advanced in this article. **Author Endorsement**: Daniela Conrado (Associate Director, Quantitative Medicine at Critical Path Institute) confirms that the author has an appropriate level of expertise to conduct this research, and confirms that the submission is of an acceptable scientific standard. Daniela Conrado declares she has no competing interests.

Grant information

The author(s) declared that no grants were involved in supporting this work.

References

1. Lunn D, Best N, Spiegelhalter D, *et al.*: **Combining MCMC with 'sequential' PKPD modelling.** *J Pharmacokinet Pharmacodyn.* 2009; **36**(1): 19–38.
   PubMed Abstract | **Publisher Full Text**

2. Bailey S, Neuenschwander B, Laird G, *et al.*: **A Bayesian case study in oncology Phase I combination dose-finding using logistic regression with covariates.** *J Biopharm Stat.* 2009; **19**(3): 469–484.
   PubMed Abstract | **Publisher Full Text**

3. Pinheiro J, Bornkamp B, Glimm E, *et al.*: **Model-based dose finding under model uncertainty using general parametric models.** *Stat Med.* 2014; **33**(10): 1646–1661.
   PubMed Abstract | **Publisher Full Text**

4. Lisovskaja V, Burman CF: **On the choice of doses for phase III clinical trials.** *Stat Med.* 2013; **32**(10): 1661–1676.
   PubMed Abstract | **Publisher Full Text**

5. Berry DA: **Bayesian Statistics and the Efficiency and Ethics of Clinical Trials.** *Statist Sci.* 2004; **19**(1): 175–187.
   Publisher Full Text

6. Palmer CR: **Ethics, data-dependent designs, and the strategy of clinical trials: time to start learning-as-we-go?** *Stat Methods Med Res.* 2002; **11**(5): 381–402.
   PubMed Abstract | **Publisher Full Text**

7. Zang Y, Lee JJ, Yuan Y: **Adaptive designs for identifying optimal biological dose for molecularly targeted agents.** *Clin Trials.* 2014; **11**(3): 319–327.
   PubMed Abstract | **Publisher Full Text** | **Free Full Text**

8. Sheiner LB, Beal SL, Sambol NC: **Study designs for dose-ranging.** *Clin Pharmacol Ther.* 1989; **46**(1): 63–77.
   PubMed Abstract | **Publisher Full Text**

9. Onoue H, Yano I, Tanaka A, *et al.*: **Significant effect of age on docetaxel pharmacokinetics in Japanese female breast cancer patients by using the population modeling approach.** *Eur J Clin Pharmacol.* 2016; **72**(6): 703–710.
   PubMed Abstract | **Publisher Full Text**

10. Friberg LE, Henningsson A, Maas H, *et al.*: **Model of chemotherapy-induced myelosuppression with parameter consistency across drugs.** *J Clin Oncol.* 2002; **20**(24): 4713–4721.
    PubMed Abstract | **Publisher Full Text**

11. King MD, Grech-Sollars M: **A Bayesian spatial random effects model characterisation of tumour heterogeneity implemented using Markov chain Monte Carlo (MCMC) simulation [version 1; referees: 1 approved].** *F1000Res.* 2016; **5**: 2082.
    Publisher Full Text

12. R Core Team: **R: A Language and Environment for Statistical Computing.** R Foundation for Statistical Computing, Vienna, Austria, 2016.
    Reference Source

13. Norris DC: **Code and Figures for v1 of F1000Research submission: Dose Titration Algorithm Tuning (DTAT) should supersede the Maximum Tolerated Dose (MTD) concept in oncology dose-finding trials.** *Open Science Foundation.* 2017.
    Data Source

14. Julier SJ, Uhlmann JK: **New extension of the Kalman filter to nonlinear systems.** 1997; **3068**: 182–193.
    Publisher Full Text

15. Norris DC, Gohh RY, Akhlaghi F, *et al.*: **Kalman filtering for tacrolimus dose titration in the early hospital course after kidney transplant.** *F1000Res.* 2017; **6**.
    Publisher Full Text

16. Kalman RE: **A New Approach to Linear Filtering and Prediction Problems.** *J Basic Eng.* 1960; **82**(1): 35–45.
    Publisher Full Text

17. Weycker D, Li X, Edelsberg J, *et al.*: **Risk and Consequences of Chemotherapy-Induced Febrile Neutropenia in Patients With Metastatic Solid Tumors.** *J Oncol Pract.* 2015; **11**(1): 47–54.
    PubMed Abstract | **Publisher Full Text**

18. De Pillis LG, Fister KR, Gu W, *et al.*: **Optimal control of mixed immunotherapy and chemotherapy of tumors.** *J Biol Syst.* 2008; **16**(01): 51–80.
    Publisher Full Text

19. Dua P, Dua V, Pistikopoulos EN: **Optimal delivery of chemotherapeutic agents in cancer.** *Comput Chem Eng.* 2008; **32**(1–2): 99–107.
    Publisher Full Text

20. d'Onofrio A, Ledzewicz U, Maurer H, *et al.*: **On optimal delivery of combination therapy for tumors.** *Math Biosci.* 2009; **222**(1): 13–26.
    PubMed Abstract | **Publisher Full Text**

21. Krabs W, Pickl S: **An optimal control problem in cancer chemotherapy.** *Appl Math Comput.* 2010; **217**(3): 1117–1124.
    Publisher Full Text

22. Engelhart M, Lebiedz D, Sager S: **Optimal control for selected cancer chemotherapy ODE models: a view on the potential of optimal schedules and choice of objective function.** *Math Biosci.* 2011; **229**(1): 123–134.
    PubMed Abstract | **Publisher Full Text**

23. Ledzewicz U, Schättler H, Gahrooi MR, *et al.*: **On the MTD paradigm and optimal control for multi-drug cancer chemotherapy.** *Math Biosci Eng.* 2013; **10**(3): 803–819.
    PubMed Abstract | **Publisher Full Text**

24. Petroni GR, Wages NA, Paux G, *et al.*: **Implementation of adaptive methods in early-phase clinical trials.** *Stat Med.* 2017: **36**(2): 215–224.
    PubMed Abstract | **Publisher Full Text** | **Free Full Text**

25. Feinstein AR: **"*Clinical Judgment*" revisited: the distraction of quantitative models.** *Ann Intern Med.* 1994; **120**(9): 799–805.
    PubMed Abstract | **Publisher Full Text**

26. Norris DC: **Casting a realist's eye on the real world of medicine: Against Anjum's ontological relativism.** *J Eval Clin Pract.* 2017.
    PubMed Abstract | **Publisher Full Text**

27. Popper KR: **Objective Knowledge: An Evolutionary Approach.** Rev ed. Oxford [Eng.]: New York: Clarendon Press; Oxford University Press; 1979.
    Reference Source

28. Whitehead AN: **Process and Reality.** New York: Free Press; (The quotation is from Chapter I section III.). 1969.

29. Weber JS, Yang JC, Atkins MB, *et al.*: **Toxicities of Immunotherapy for the Practitioner.** *J Clin Oncol.* 2015; **33**(18): 2092–2099.
    PubMed Abstract | **Publisher Full Text** | **Free Full Text**

30. Molife LR, Alam S, Olmos D, *et al.*: **Defining the risk of toxicity in phase I oncology trials of novel molecularly targeted agents: a single centre experience.** *Ann Oncol.* 2012; **23**(8): 1968–1973.
    PubMed Abstract | **Publisher Full Text**

31. 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.
    PubMed Abstract | **Publisher Full Text**

Open Peer Review

Current Referee Status: ✔ ? ✔ ✔

Version 2

Referee Report 14 June 2017

✔ **Yuan Ji**

Department of Public Health Sciences, University of Chicago, Chicago, IL, USA

I agree fully that MTD should not be the only goal of dose finding. The paper nicely pointed that out. The author seems to suggest using an optimization approach (i.e., Newton-Raphson) rather than a model-based approach for achieving individualized dosing. There is no concrete design being proposed, or at least it is unclear to me how to design a trial based on the contents in the paper.

Also, I think the paper can be more concise and needs some more clear flow on the statistics and mathematical modelling. It is expected that a probability model (or an algorithm) is presented with an inference procedure. There are a few models proposed in this paper, but their relationship to the inference is a bit unclear. I would recommend a subsection of "Algorithm" that would connect different parts of the paper into a coherence piece. Otherwise, the paper clearly discusses an important issue in dose finding.

Is the work clearly and accurately presented and does it cite the current literature?
Yes

Is the study design appropriate and is the work technically sound?
Partly

Are sufficient details of methods and analysis provided to allow replication by others?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
Partly

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Yes

Competing Interests: No competing interests were disclosed.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Author Response 21 Jun 2017

David C. Norris, Precision Methodologies, LLC, USA

I thank Dr. Ji for his thoughtful and encouraging review. I acknowledge that this paper by itself does not provide a recipe for designing a complete clinical trial. My aim here has been rather to offer a philosophical and conceptual framework within which such designs may be conceived and constructed. Thus, as the saying goes, "Some User Assembly Required." **I do wish to encourage ambitious research teams in academia and industry to take up this challenge.** I trust that its formidable nature will attract, rather than deter, the necessary sort of interest.

A great deal more will be required to design a DTAT study than has hitherto been required to realize current designs organized around 'the' MTD. By its nature, 'the' MTD tends to abstract away all details of PK/PD, and thus to enable biostatistician trialists to proceed unimpeded by entanglements with pharmacometrics. 'The' MTD also lends itself rather easily to the construction of simple probability models, and to statistical inference over a scant few parameters. No doubt, 'the' MTD owes its astonishing longevity (despite its *prima facie* invalidity) in part to its having enabled such functional compartmentalization and easy statistical modeling. Research groups that successfully develop meaningful DTAT designs will surely comprise interdisciplinary collaborations spanning *pharmacometrics*, *biostatistics* and the *clinical* realm. Such collaborations, although perhaps difficult (Kowalski 2015, Senn 2010, Sheiner 1991), are ripe with riches for those who will pull them off (Norris 2017c).

From the perspective of *Biostatistics*, a formal probability model and associated statistical inference over its parameters are conspicuously absent from this paper. While I anticipate that formal statistical inference will usefully illuminate the DTAT principle and its applications, I doubt that it is strictly necessary to them. Formal inference in a DTAT application will likely find itself subordinated to "organic forms of clinical judgement" such as I reference in the Discussion section. Indeed, any "dynamic, 'learn-as-you-go' modifications" of a DTA would effectively pull the rug out from under a statistical inference paradigm that requires *ex ante* commitments to immutable probability models. Thus, statistical model development might potentially extend *into and throughout* the conduct of a DTAT study, so as to give a DTAT study partly the character of what is commonly called 'data-driven model development'. Performing statistical inference over the tuning parameters of a DTA intrinsically has a quite different character and level of difficulty from inference over the probability of a DLT (Rogatko *et al* 2015). The most meaningful applications of the DTAT principle will employ models incorporating diverse mechanistic biological insights informed by disparate sorts of data (Senn 2010). Not all such 'informing' will be feasibly accomplished entirely under the duress of statistical formality.

Nevertheless, I do hope in forthcoming publications to demonstrate explicitly how formal statistical inference may play a useful role in a DTAT study; as I mention in the Methods section, I have selected the **pomp** package for implementing this early work precisely to avail myself of a flexible platform for later incorporating a variety of inference procedures. **I heartily welcome expressed interest from potential collaborators in industry, academia and the regulatory sphere.**

References

Kowalski KG. My Career as a Pharmacometrician and Commentary on the Overlap Between Statistics and Pharmacometrics in Drug Development. *Statistics in Biopharmaceutical Research*. 2015;7(2):148-159. doi:10.1080/19466315.2015.1008645.

Senn S. Statisticians and pharmacokineticists: what they can still learn from each other. *Clin Pharmacol Ther*. 2010;88(3):328-334. doi:10.1038/clpt.2010.128.

Sheiner LB. The intellectual health of clinical drug evaluation. *Clin Pharmacol Ther*. 1991;50(1):4-9.

Norris DC. Costing 'the' MTD. *bioRxiv*. June 2017. doi:10.1101/150821.

Rogatko A, Cook-Wiens G, Tighiouart M, Piantadosi S. Escalation with Overdose Control is More Efficient and Safer than Accelerated Titration for Dose Finding. *Entropy (Basel)*. 2015;17(8):5288-5303. doi:10.3390/e17085288.

Competing Interests: No competing interests were disclosed.

Referee Report 18 May 2017

doi:10.5256/f1000research.12185.r22676

Felipe Maglietti (iD)

Faculty of Exact and Natural Sciences, University of Buenos Aires, Buenos Aires, Argentina

It is a very interesting article. We should start moving forward to a more personalized medicine. Dosing, specially in oncology is very complex since in many cases the short life expectancy can be ruined by having serious side effects. This approach could be very useful for determining the optimal dose in different patients.

Is the work clearly and accurately presented and does it cite the current literature?
Yes

Is the study design appropriate and is the work technically sound?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?
Yes

If applicable, is the statistical analysis and its interpretation appropriate?
I cannot comment. A qualified statistician is required.

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Yes

Competing Interests: No competing interests were disclosed.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Version 1

Referee Report 22 March 2017

doi:10.5256/f1000research.11448.r21177

? **Natalja Strelkowa**

Global Biometry and Clinical Applications, Boehringer Ingelheim Pharma GmbH & Co. KG, Biberach an der Riss, Germany

I support the main idea of the article to promote individualized dose finding for toxic drugs. The methodological framework of reinforcement learning is an interesting and promising alternative to the current practice of population based maximum tolerated dose approach (MTD).

The following points need to be stressed and/or addressed in the paper:

- The immediate toxic response must be reliably measurable for the reinforcement learning to work; long term toxicity is as far as I have understood is not included in the current framework. This needs to be addressed, in particular in light of new immune-oncology and targeted therapies where the toxic response is not immediately observable.

- The time to achieve optimal dosing scheme according to the dose finding algorithm is also of a key importance and must be much shorter than expected survival of the patient.

- For the practical implementation of the approach, hart boundaries on overshooting must be implemented. This will likely influence the time to optimal dosing for each patient.

The title and the conclusions are too strong in my view and need revision. This approach is very promising and I myself am also very interested in its application, but in the abstract and discussion section it needs to be clear that key practical questions are still to be addressed. Strong statements like: 'DTAT should supersede MTD' should be removed from the title and abstract and replaced by 'DTAT is an alternative to current practice' or 'Could be a better alternative to current practice for drugs with narrow therapeutic window' etc.

Competing Interests: No competing interests were disclosed.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

> Author Response 23 Mar 2017
>
> **David C. Norris**, Precision Methodologies, LLC, USA
>
> I thank Dr. Strelkowa greatly for her supportive and critical comments, which invite extended discussion on several points that v1 of this paper has omitted **to its detriment**. I use this reply to

indicate changes I propose to make in v2, pursuant to Dr. Strelkowa's comments. Unless otherwise indicated, these changes would seem to belong in the **Discussion** section:

- Dr. Strelkowa rightly points to general conditions that constrain the time-scale on which DTAT-based learning may operate. Specifically, DTAT cannot learn faster than the lag-time at which the targeted toxic response(s) develop. In the important area of immuno-oncology that Dr. Strelkowa highlights, common dose-limiting toxicities (DLTs) do admit monitoring on time scales comparable to the chemotherapy induced neutropenia (CIN) simulated in this paper. For example, the *cytokine release syndrome* (CRS) that accompanies chimeric antigen receptor (CAR-)T cell therapies typically arises within 1 week of administration (even earlier with concomitant high-dose IL-2) and constitutes a clinical *syndrome* that admits multivariate monitoring on numerous quantitative clinical and laboratory measures.[1] Regarding however the molecularly targeted agents (MTAs) also mentioned by Dr. Strelkowa, *late* toxicities indeed have tended to attract the lion's share of attention.[2] One reason for this is that the early toxicities of MTAs tend to be relatively milder than those of cytotoxic and immunologic therapies.[3] Nevertheless, a DTAT principle continues to apply here—just on a longer time scale. Of course, a 'dose titration *algorithm'* (DTA) that responds to MTA toxicities that patients mainly *experience and evaluate subjectively* may resemble a process of ongoing *shared decision making* (with the oncology care team) more than it resembles the impersonal calculations we typically think of as 'algorithmic'. But with a suitably broadened understanding of 'algorithm'—one that accommodates what might typically be termed **protocols**—the DTAT (or perhaps, DT**P**T) principle continues to apply. In such applications, "supervision and modification by clinical judgment" as mentioned in the **Discussion** clearly comes to the fore. But even in such contexts, the development and application of scoring systems for patient-reported clinical symptoms and quality of life would enable dose titration **p**rotocols (DT**P**s) to be described objectively in quite 'algorithmic' terms that would preserve the applicability of a 'tuning' concept.

- The issue of *competing risks* Dr. Strelkowa raises is truly important, and brings into play the inter-individual heterogeneity in "values and goals of care" that I discussed in opening the **Introduction** and (especially) toward the end of the **On 'tuning'** section, where I explicitly highlight values/goals as factors that should "most emphatically" inform DTA tuning. (In the PDF, this utterly essential point tragically breaks across pages 7-9 due to an intervening page 8 of figures, a mishap I will aim to avoid in v2.) To address Dr. Strelkowa's comment here, I plan to expand that latter discussion explicitly to incorporate *prognosis*. For example, if a patient with more advanced disease and short expected survival has (in consultation with the oncologist and oncology care team) nevertheless decided to enroll in a Phase I DTAT study to pursue the possibility of therapeutic benefit, then this patient's decision indicates a subjective weighting of benefits vs harms favoring a *higher starting dose* and *more aggressive titration*.

- I agree that any sound dose titration algorithm will incorporate fail-safe limits on dose escalation. To support a richer discussion in the **On 'tuning'** section, I purposely chose 'wrong' tuning parameters (and a naive Newton-Raphson DTA) so that Figure 5 would illustrate several problems that one would never purposely 'design in' to an actual study: (a) a too-low starting dose and (b) the potential for overshooting, such as occurred in id8, id10, id12 and id23—the last of whom received in fact an off-scale dose. For v2, in both the Figure 5 caption and the main text of **On 'tuning'**, I propose to note the important role that

fail-safe upper bounds on both *absolute dose* and *dose multipliers* will play in reducing the risk of overshooting in a practical trial.

A larger point about DTAT requires clarification in v2. The technical connections with *recursive filtering* and *optimal control*, as drawn in the paper, serve much the same **notional** purpose as choosing docetaxel as the basis for simulation. The above discussion shows the DTAT principle applicable to oncology therapeutics beyond the cytotoxics. Likewise, notwithstanding the essential heuristic role *recursive filtering* has played in DTAT's development, it should not be thought to define DTAT. (In fact, I have at this point in my further work already abandoned the linear Kalman filter approach indicated in the last sentence of v1 **Abstract**, in favor of full-information methods.) What **does** define DTAT's essential contribution to Phase I oncology study design is that it yields a **new abstraction** (the *DTA* with its *tuning parameters*) capable of *embodying knowledge objectively* [4], to supersede a **fallacious abstraction** ('the' MTD) that almost completely lacks this capability. I must make this point explicit in v2, emphasizing that the role of the technical connections I've drawn is *to illustrate* the **(DTA+tuning)** abstraction which constitutes DTAT's essential contribution. I will also try to modify the title somehow to underscore this point. I will however retain a withering treatment of 'the' MTD, an anti-precision idea whose time has passed. In support of that view, I will in the v2 **Discussion** or **Introduction** briefly discuss 'the' MTD specifically as a *fallacy of misplaced concreteness* [5]. The purpose of my strong title is **to provoke long-overdue critical thought and discourse.** Should any clinical trial methodologist leap now to the defense of 'the' MTD, I will most heartily welcome that challenge.

References

1. Weber JS, Yang JC, Atkins MB, Disis ML. Toxicities of Immunotherapy for the Practitioner. *JCO*. 2015;33(18):2092-2099. doi:10.1200/JCO.2014.60.0379.
2. 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. doi:10.1200/JCO.2010.31.9236.
3. Molife LR, Alam S, Olmos D, et al. Defining the risk of toxicity in phase I oncology trials of novel molecularly targeted agents: a single centre experience. *Ann Oncol*. 2012;23(8):1968-1973. doi:10.1093/annonc/mds030.
4. Popper KR. *Objective Knowledge: An Evolutionary Approach*. Rev. ed. Oxford [Eng.]: New York: Clarendon Press; Oxford University Press; 1979.
5. Whitehead AN. *Science and the Modern World: Lowell Lectures, 1925*. New York: The Free Press; 1997.

Competing Interests: No new competing interests have arisen since v1 was published.

Referee Report 14 March 2017

Matthew E. Nielsen
Departments of Urology, Epidemiology, and Health Policy & Management, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

The author presents a provocative modeling-based demonstration of an innovative alternative to the traditional one-size-fits-all maximally tolerated dose concept. The concept of individualized pharmacokinetic-based precision dosing is intuitively appealing and the analyses presented herein support potential utility from the development and application of such methods.

Competing Interests: No competing interests were disclosed.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Discuss this Article

Version 2

Author Response 08 Jun 2017

David C. Norris, Precision Methodologies, LLC, USA

I thank Dr. Mukherjee for his supportive comment, and for his encouragement to provide a more accessible presentation of DTAT. I have just uploaded annotated slides from my May 12, 2017 podium presentation at a Stanford Symposium on Dose Selection for Cancer Treatment Drugs. The pharmacoeconomic perspective developed in those slides—and also in the poster I presented at Life Science Innovation Northwest (LSINW 2017)—may help build a wider context within which the importance and impact of DTAT can be appreciated.

Furthermore, I have also provided a lay summary of this paper at growkudos.com. This is directed toward the broadest possible audience, including laypersons with cancer, their families and lay advocates.

Competing Interests: No competing interests were disclosed.

Reader Comment 01 May 2017

Neil Mukherjee,

A very interesting paper.
The author has presented a very interesting alternative to clinical drug trial process. The tolerated dose concept needs to change to be more tailorable to individuals. I would ask the author to explain the "Dose Titration Algorithm Tuning (DTAT)" the process and needs early on in the introduction for it to be palatable to a regular around the corner clinician.

Competing Interests: None