



Integrated testing strategies can be optimal for chemical risk classification[☆]



Marko Raseta^{a,*}, Jon Pitchford^b, James Cussens^c, John Doe^d

^a Department of Mathematics, University of York, UK

^b Department of Biology, University of York, UK

^c Department of Computer Science, University of York, UK

^d Parker Doe Partnership LLP, UK

ARTICLE INFO

Article history:

Received 1 July 2016

Revised 21 April 2017

Accepted 16 May 2017

Available online 23 May 2017

Keywords:

Optimal integrated testing strategies

Dynamic programming

ABSTRACT

There is an urgent need to refine strategies for testing the safety of chemical compounds. This need arises both from the financial and ethical costs of animal tests, but also from the opportunities presented by new in-vitro and in-silico alternatives. Here we explore the mathematical theory underpinning the formulation of optimal testing strategies in toxicology. We show how the costs and imprecisions of the various tests, and the variability in exposures and responses of individuals, can be assembled rationally to form a Markov Decision Problem. We compute the corresponding optimal policies using well developed theory based on Dynamic Programming, thereby identifying and overcoming some methodological and logical inconsistencies which may exist in the current toxicological testing. By illustrating our methods for two simple but readily generalisable examples we show how so-called integrated testing strategies, where information of different precisions from different sources is combined and where different initial test outcomes lead to different sets of future tests, can arise naturally as optimal policies.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

Notions of what society deems to be an acceptable testing regime for new chemicals are in a constant state of flux. Until 1999 it was acceptable in the EU to perform tests on guinea pigs in order to determine whether certain cosmetic products were hazardous for human skin [20]. After 1999 this was replaced by the mouse LLNA (Local Lymph Node Assay), another animal-based method. More recently the EU ethical climate has changed again; by 2018 no new chemical to be used in the cosmetics industry can be tested on animals. Instead, chemicals need to be classified reliably using information from emerging in-vitro and in-silico assays, supplemented where possible by mathematical models. These new methods are likely to be less accurate than in-vivo tests, but are generally cheaper and less ethically problematic to implement. This presents a problem common across toxicology in general: can we make good predictions about the risks associated with new chemicals without using animals at all? In other words, how best can we

assemble uncertain information based on non-animal assays, so as to arrive at optimal ethical testing regimes?

Many important papers have emerged on this topic [9–14,17].

Indeed [10] develops a theory that determines the optimal exposure level of any particular member of the population to the chemical and uses this theory to solve a decision problem of how to pick which chemical to test for hazard first from some finite set of possible chemicals. [9] develops a framework allowing one to compute the optimal battery of tests to assess a generic toxicological endpoint by means of a cost effectiveness analysis (CEA). By contrast, [11] develops a framework in which adaptive cost sensitive Integrated Testing Strategies can be derived by means of a Value of Information technique (VOI). The authors there distinguish between decision problems for competitive businesses and regulators. Furthermore, [13] begins by improving and generalising previous work [12,14] by developing more accurate potency class predictions of skin sensitisation potential of chemicals via theory of Bayesian Networks and then uses these results together with VOI framework to derive Optimal Integrated Testing Strategies for the assessment of chemical hazard of chemicals. Finally, similar to [9,11,17] uses CEA in the context of performing a cost effectiveness analysis in the special case of acute oral toxicity.

However, none of these explicitly accounts for the individual differences between humans both in the exposure (i.e. environ-

[☆] This work was supported by the NC3Rs [grant number NC K001264 1]

* Corresponding author.

E-mail addresses: marko.raseta@york.ac.uk (M. Raseta), jon.pitchford@york.ac.uk (J. Pitchford), james.cussens@york.ac.uk (J. Cussens), john.doe@parkerdoe.com (J. Doe).

mental variability) and in the toxicity corresponding caused by that exposure (i.e. individual variability). To be more specific, only [10] introduces a concept of toxicity formally but treats it as constant for all members of the population. Moreover, none of these papers combines these with the financial costs of chemical risk classification in a mathematically rigorous fashion.

Any new testing strategy must be able to deal rationally with contradictory evidence. For example, one in-vitro assay may predict that certain chemical is a skin sensitiser, while another in-silico assay may predict that the same chemical is actually safe. The classification part of the argument in [24] deals with this problem using a combination of majority voting and Bayesian Statistics [14]. proposes assembling a Bayesian Network and combines this with the Weight-of-Evidence approach to overcome this issue [23]. proposes a strategy of “averaging probabilities”, using empirical estimates of precision of each assay and then averaging these out in one “meta-assay”. Each of these solutions may be pragmatic and defensible within the authors’ given problem, but an over-arching logical framework would be a helpful step in confirming the value and risk associated with the removal of animal tests.

In what follows we shall propose a mathematical framework which seeks to simultaneously overcome the shortcomings mentioned above. The issues of imprecision, and of environment- and individual-level variability, fall naturally within theories developed for evolutionary ecology [6]. The efficient assimilation of evidence can then be handled by well-developed theories of Markov Decision Processes [4].

We finish this chapter by surveying previous work in Toxicology and Medicine that is based on this mathematical theory. To the best of our knowledge very little work in Toxicology uses Markov Decision Processes, to be more precise, these are the works of [5,15] is a rich summary of techniques used in contemporary in house pharmacological research and decision making illustrated with numerous examples.

Among a vast range of mathematical and computational techniques used are the Markov Decision Processes as applied to the optimal decision making of a pharmaceutical business on whether to proceed from earlier (Phase 1 clinical trial) to later (Phases 2 and 3 of clinical trials) stages. Our work generalises this work in a number of directions. Firstly, the models in [5] are developed for the sake of a making a commercial business more profitable and do not take into account the regulatory aspect of Toxicology, i.e. the fact that the company may actually incur fines from regulatory bodies and lawsuits from individual consumers in case they exhibit adverse outcomes as a consequence of using the drug. The Markov Decision Process model in this paper takes this into account via the mechanisms of misclassification costs: in case the company declares an unsafe chemical as safe there will be serious consequences. Equally, if the company actually declares a safe chemical as unsafe it will lose money by not selling the safe product in the market for which it possibly had an advantage over its competitors. Thus our work bridges the two worlds: it allows the company to maximise its profits while simultaneously acts in the best interest of the general public. Secondly, [5] does not take neither variable exposure to chemical among different members of the target population nor precision of measurements used to test safety in the account. Instead, it relies on toxicity levels observed in recently tested chemicals to draw conclusions on the new chemical of interest while the transition probabilities of the Markov Decision Model model are estimated from historical data and power calculations which, by nature, cannot guarantee precisions of estimates in advance in the case of unknown moments. Another problem with this approach is that there is no guarantee the new chemical will share toxicity thresholds with the previously used chemicals.

These issues motivate the truly novel part of the methodological work presented in this paper which is applicable in a gen-

eral setting not necessarily restricted to Toxicology and Medicine. Namely, the transition probabilities between states of the Markov Decision Process in this paper are based on the idea of integrating evidence from different measurements of the same quantity in a non-contradictory way by using the information on the precision from the instruments/assays used in the process. Knowing the value obtained in a less accurate measurement and its precision we can get a probability distribution on the values more accurate measurement of the same quantity can possibly take. This simple observation has far reaching consequences; namely since the states of the Markov Decision Process in our model are the collection of measurements the above allows for a logically sound way of defining transition probabilities of the model by first performing cheapest possible tests for each parameter of interest. This procedure is justified by the existence of devices of varying precision in many fields of human work. As far as Toxicology is concerned, this corresponds to in-vivo, in-vitro and in-silico tests. By choosing to start our analysis with a cheap in-silico test we remedy the problem outlined above: indeed we get the transition probabilities of the model without having to resort to further unknown characteristics of the chemical therefore bypassing a potentially circular argument.

Another interesting work involving Markov Decision Process in Toxicology is a theoretical paper of [15]. The authors approach the problem of model checking for a Markov Decision Process from the Computer Science point of view using the language of Mathematical Logic; actually as it turns out, the authors do not work with a conventional probabilistic definition of a Markov Decision Process but define their own in another set-up. Although motivated by an example of Insulin compartment model the paper soon drifts into proving results in Mathematical Logic and holds little practical value.

When it comes to applications in Medicine the literature is much larger. We discuss in detail a variety of different applications [1,2,8,16,18,21,22]. As mentioned in the above, the main novelty of this paper, logically consistent aggregation of evidence from different measurements in a non-contradictory way and without a need to resort to empirical estimates still stands.

[8] develops a general framework for learning efficient approaches to medical diagnosis. It resembles this manuscript and [3] in the sense that the states are cumulative history of observations, which in turn, guarantees the Markovian nature of the process but it resorts to the empirical estimates of transition probabilities between states based on observed frequencies [16]. develops a discrete time, but unlike this paper, an infinite-horizon Markov Decision Process to maximise the patient’s quality-adjusted life years prior to them having either a stroke or developing a Coronary Heart Disease. The infinite horizon is justified by the large number of visits to the doctor by Type 2 Diabetes patients. The model resembles the one of ours in that it has a terminal state which in turn insures convergence. Finally, transition probabilities are computed as a combination of equations based on medical knowledge and empirical observations.

[21] develops a Markov Decision Process that aims to maximise the expected lifetime or quality-adjusted life years. Similarly to [16] this model contains an absorbing state which can be reached from any other state and is an infinite-horizon problem for the same reason as [16]. Furthermore, exactly as [16], the rewards in the problem are not measured only in monetary units as in our case but instead in the units of Health Economics, namely quality-adjusted life years, while the transition probabilities of the model are estimated empirically from data [1]. develops a Markov Decision Process for optimal choice of when to go for a liver transplant and then, furthermore, should one accept the part of a liver of a living-donor or the entire cadaveric liver.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات