Proceedings of the
2nd Workshop on Artificial Intelligence and Formal Verification, Logics, Automata and Synthesis (OVERLAY),
September 25, 2020

AI-Guided Synthesis of Personalised Pharmacological Treatments via In Silico Clinical Trials

S. Sinisi[1], V. Alimguzhin[1], T. Mancini[1], E. Tronci[1], F. Mari[2], and B. Leeners[3]

[1]Computer Science Dept., Sapienza University of Rome, Italy
[2]Dept. Movement, Human & Health Sciences, University of Rome Foro Italico, Italy
[3]Division of Reproductive Endocrinology, University Hospital Zurich, Switzerland

Abstract

A key topic in *precision medicine* is to develop pharmacological treatments *optimised* for any given individual (*personalised* treatments). Model-based approaches (aka In Silico Clinical Trials, ISCT) aim at achieving this goal, by enabling the automatic synthesis and verification of personalised treatments via *simulation*, before they are actually administered to patients.

In this short paper, we introduce the area of ISCT, and review our approach to the *in silico* automatic synthesis of optimal personalised treatments, where numerical simulation of quantitative models of the human physiology and reactions to drugs is driven by Artificial Intelligence global search.

1 Introduction

The overall objective of In Silico Clinical Trials (ISCT) is to perform, *in silico*, the typical activities carried out to assess safety and efficacy of pharmacological treatments, biomedical devices, or other therapeutic procedures (see, *e.g.*, [2, 49, 61, 17]). Being entirely *model-based*, ISCT have the potential to be a key-enabler of *precision* medicine, where *personalised* treatments optimised for a given patient can be designed (and verified *in silico* via simulation) before being actually administered. ISCT build on the availability of computational models (Virtual Physiological Human, VPH, models) for the human physiology, patho-physiology, and drugs Pharmacokinetics/Pharmacodynamics (PKPD), which also define the possible physiological differences between different individuals (*i.e.*, the possible *phenotypes*). Such models range from networks of biological pathways (*qualitative* knowledge, *e.g.*, [27, 18]), to mechanistic *quantitative* physiology models at different levels of scale (*e.g.*, molecules, [55], cells [22, 3], organs [66, 48], body compartments [13, 59], tumour growth [52, 26]) up to the whole human body homeostatis (*e.g.*, [25, 46]), as well as *quantitative* models of the PKPD of medicinal drugs (*e.g.*, [32]).

2 Formal framework

VPH models. Large VPH models are often defined as *hybrid systems* whose dynamics is governed by Ordinary Differential Equations (ODEs) (*e.g.*, [58, 1, 4]). Model *inputs* represent drug administrations and other exogenous actions, while *parameters* encode *inter* as well as *intra* subject variabilities (*e.g.*,

[60, 37, 11]). Model *outputs* define the values of a set of model *observables*, *i.e.*, the biological quantities of interest that are measurable in a human patient via clinical assays. Given an assignment λ to the parameters of a VPH model $\mathcal{S}$ and an input function (defining a time-sequence of drug administrations), the equations in $\mathcal{S}$ define the time evolution of the observables (aka the *trajectory* of $\mathcal{S}$), when the system is fed with that input function starting from its initial state, and with its parameters set to λ. Of course, as we are interested in *causal* VPH models, the time evolution of $\mathcal{S}$ up to any time point only depends on the *restriction* of the input function up to that time point (see, *e.g.*, [58, 42]).

Complete populations of Virtual Phenotypes. VPH models typically take into account inter-subject variabilities (*i.e.*, the physiological differences among different individuals) by including *parameters* in their equations. Different parameter assignments yield different model evolutions (also wrt. reactions to drug administrations), thus define different Virtual Phenotypes (VPs). Intuitively, each VP represents a class of *indistinguishable* (as long as the VPH model is concerned) individuals. Computing a complete set of parameter assignments for the model at hand which entail *physiologically meaningful* evolutions for the observables is the starting point to obtain a *representative* population of VPs (hence, ideally showing *all* possible phenotypes). Searching for a complete population of (physiologically meaningful) VPs is all but easy in complex VPH models. In fact, as VPH models are often non-identifiable, over-parameterised, and with unknown inter-dependency constraints among parameters, *most* model parameterisations entail evolutions of the observables which *clearly violate* the laws of biology. In our previous work [60, 37, 11] we showed that a statistically complete (with respect to the set of possible model behaviours) population of (physiologically meaningful) VPs for large non-identifiable ODE-based models can be effectively computed by exploiting methods based on Artificial Intelligence (AI) global search and statistical model checking, starting from background domain knowledge and available clinical data.

Human Patient Digital Twins. A Digital Twin (DT) is a *digital representation* of the physiology of interest of a *given* patient, and of their reactions to relevant drugs (PKPD). The availability of a DT of a patient would open-up a plethora of new clinical opportunities in the area of precision medicine, starting from the automatic *in silico* synthesis of optimal personalised treatments. In order to build a DT for a patient, we exploit the clinical data available for that patient (in the form of a clinical record $\mathcal{C}$), as well as a complete population $\mathcal{P}$ of VPs for a VPH model $\mathcal{S}$ defining the physiology and PKPD of interest for our envisioned treatments. A clinical record $\mathcal{C}$ is defined as a pair $(\mathbf{u}^{(\mathcal{C})}, \mathbf{o}^{(\mathcal{C})})$, where $\mathbf{u}^{(\mathcal{C})}$ is a time-series of the (past) clinical actions (*e.g.*, drug administrations) performed on the patient, and $\mathbf{o}^{(\mathcal{C})} = \{\mathbf{o}_i^{(\mathcal{C})}\}_1^p$ is a set of p time series of clinical measurements performed on that patient, regarding the p biological quantities corresponding to the observables of $\mathcal{S}$. DTs are defined in terms of a function $\eta(\mathcal{C}, \lambda)$ which measures the *mismatch* between clinical measurements in $\mathcal{C}$ and the trajectory of VP (*i.e.*, the model parameterised with) $\lambda \in \mathcal{P}$, when fed with the very same input function $\mathbf{u}^{(\mathcal{C})}$ as the one reported in $\mathcal{C}$. In [57], $\eta(\mathcal{C}, \lambda)$ is the average, over the p available time series of measurements in $\mathcal{C}$, Mean Absolute Percentage Error (MAPE) between $\mathbf{o}_i^{(\mathcal{C})}$ and the trajectory of the corresponding observable of $\mathcal{S}$ parameterised with λ and fed with $\mathbf{u}^{(\mathcal{C})}$. The *digital twin* $\mathcal{P}(\mathcal{C})$ associated to clinical record $\mathcal{C}$ is then the set of *all* VPs λ in our (complete) population $\mathcal{P}$ having a mismatch $\eta(\mathcal{C}, \lambda)$ up to a given threshold $\delta \in \mathbf{R}_{0+}$: $\mathcal{P}(\mathcal{C}) = \{\lambda \in \mathcal{P} \mid \eta(\mathcal{C}, \lambda) \leq \delta\}$. In other words, the digital twin of a human patient is a *digital* representation of the patient physiology and reaction to drugs, in the form of *all* VPs entailing model behaviours that are consistent (up to an average MAPE δ) with the patient clinical measurements available in $\mathcal{C}$.

3 Computing optimal personalised treatments

Below, we review our recent work [43, 57], where we perform the automatic *in silico* synthesis, via intelligent search, of personalised treatments optimised for a given patient, using the patient DT computed from an available clinical record.

Treatment optimisation is performed with respect to a user-defined criterion. In Section 4, we show experimental results in applying our algorithm to seek for an effective personalised treatment that uses a *minimum quantity* of drug (*i.e.*, a *lightest still effective treatment*). Although this is just an example, such

optimality criterion is a relevant proxy in the attempt to minimise the probability and severity of adverse effects of the employed drugs.

Let $\mathcal{S}$ be a VPH model defining the physiology and PKPD of interest of our envisioned treatment, and $\mathcal{C}$ be the clinical record of the patient for whom an optimal personalised treatment is sought. Our approach to optimal treatment computation takes as input a finite set of possible *clinical actions* (defining *e.g.*, the possible drugs to be administered with their portfolio of dosage patterns), as well as treatment *invariants* and *goals*. The latter inputs define conditions that must be, respectively, *always* and *eventually* satisfied by a successful treatment (aka *safety* and *liveness* properties). Our algorithm searches, in the space of possible sequences of clinical actions, a treatment which, *if* administered to *all* VPs of the digital twin $\mathcal{P}(\mathcal{C})$ associated to the human patient at hand (represented by clinical record $\mathcal{C}$), *always* satisfies treatment invariants (safety) and *ends-up* to satisfy the treatment goals (effectiveness). Our algorithm is based on intelligent backtracking guided by heuristics to efficiently explore the space of possible treatments of bounded length (our search space). At any step, our algorithm exploits a black-box simulator (*e.g.*, [33]) to compute the effects of the envisioned treatment prefix on all the VPs of $\mathcal{P}(\mathcal{C})$.

4 A case study

In [43, 57] we used our algorithm to perform an ISCT aimed at computing optimal personalised treatments for the down-regulation phase of an assisted reproduction treatment on 21 human patients, for whom we had retrospective clinical records (data courtesy of Pfizer Inc. and Hannover Medical School) totalling around 800 measurements. Assisted reproduction treatments are complex and challenging, with low average success rates (around 30%) even in the top clinics, and with many factors that, to date, can be hardly kept under full control [28, 24, 29].

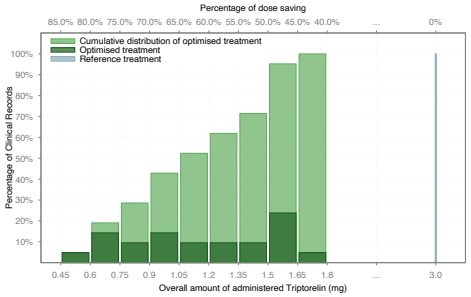

Figure 1: ISCT results.

To perform our ISCT, we used the state-of-the-art non-identifiable ODE-based VPH model described in [53] of the human female HPG axis and of the PKPD of downregulation drugs, and exploited the availability of a statistically complete population $\mathcal{P}$ of VPs for such a model, as computed in [60, 37].

For each arm of our ISCT, the associated human patient was digitally represented by a DT defined by the set of VPs in $\mathcal{P}$ with a mismatch threshold $\delta = 35\%$. Our 21 DTs contained 11 VPs each on average. Hence, 21 instances of our optimal personalised treatment computation algorithm were launched embarrassingly in parallel for 15 days on a HPC infrastructure in search for effective lightest treatments personalised for our 21 patients. Figure 1 compares the overall amount of drug used by our personalised treatments against the amount of drug used by the reference treatment currently in use at University Hospital Zurich. The overall drug saving is always at least 40% ($\sim 60\%$ on average).

5 Related work

Data-driven approaches to design individualised treatments have been investigated (*e.g.*, [50, 19]). However, when clinical data for the patient at hand is scarce, such approaches cannot be applied. For example, in our case study hormones blood concentrations are not measured every day, since those measurements are costly and invasive. In these cases, model-based approaches, exploiting PharmacoKinetics (PK), as, *e.g.*, [65], are used instead to build populations of VPs, which are leveraged to optimise and individualise drug doses [63]. In our setting, we have to face with complex non-identifiable VPH models, *e.g.*, [25, 46, 53] defining the underlying biological mechanisms. Such models are hybrid systems usually defined by highly non-linear ODEs (*e.g.*, [4]) whose inputs are discrete event sequences (*e.g.*, [38, 40, 41]). To find an optimal treatment means to find an optimal plan in hybrid domains. In the literature, there are techniques to model planning problems in hybrid domains, *e.g.*, PDDL+ [20, 62]. However, most of PDDL+ planners can deal only with linear dynamics (*e.g.*, [12]), although some attempts to handle non-linear dynamics do exist (*e.g.*, [10]). Model-checking techniques are also used to find plans (*e.g.*, [6, 7, 14, 15]) and optimal controls (*e.g.*, [16, 47, 56]). However, the typical complexity of the ODEs defining VPH models relevant

for clinical practice makes such models out of reach for symbolic approaches, and appoints numerical integration as the only viable means to compute (black-box) the model evolutions under a given input function. In particular, even considering that clinical actions have constant and equal duration (as, *e.g.*, in [30]), *no reasoning or inference* can be made on action effects in a black-box setting as ours, because the only way to interact with the models is through numerical simulation. As a consequence, classical approaches to compute, through inference, a dynamic preference order among the candidate actions to be tried during search (as those exploited in, *e.g.*, classical planning, planning for white-box hybrid systems or (Q)CSP, SAT or local search solvers, see *e.g.*, [54, 5, 21, 23, 34, 8, 9, 35, 44]) cannot be directly applied.

The automated synthesis of rational decisions and plans in black-box environments is common in several other application domains of high industrial relevance, like smart grids (*e.g.*, [36, 39, 45]), games (*e.g.*, [31]) and real-time manoeuvring of unmanned aerial vehicles (*e.g.*, [51]).

6 Conclusions

In this paper we reviewed our recent work on the *in silico* automatic synthesis of optimal personalised treatments in healthcare. Our algorithms are based on the numerical simulation of models of the human physiology and drugs PKPD guided by AI global search. We showed the application of our approach by computing personalised treatments for the downregulation phase of an assisted reproduction protocol, which are effective on the patient at hand, but minimise the overall amount of drug used.

The possibility to optimise *in silico* a complex treatment for a given human patient before its actual administration shows the potential of model-based approaches in precision medicine. This however calls for *qualified* VPH models (see, *e.g.*, [64]), which is currently one of the major obstacles for the uptake of *in silico* methods in clinical practice. Indeed, the results of our ISCT (conducted using a state-of-the-art experimentally validated VPH model) are very promising, but must be taken with care. In particular, the actual effectiveness of the personalised treatments generated by our algorithm needs to be assessed *in vivo*, in order to verify the accuracy of the model in predicting the patient reactions to *non-standard drug dosing patterns*, as those computed by our algorithm.

Acknowledgements This work was partially supported by: Italian Ministry of University and Research under grant "Dipartimenti di eccellenza 2018–2022" of the Department of Computer Science of Sapienza University of Rome; EC FP7 project PAEON (600773); INdAM "GNCS Project 2019"; Sapienza University 2018 project RG11816436BD4F21.

References

[1] R. Alur, *et al.* Hybrid modeling and simulation of biomolecular networks. In *HSCC 2001*, *LNCS*, vol. 2034. Springer, 2003.

[2] Avicenna Project. *In silico* clinical trials: How computer simulation will transform the biomedical industry, 2016.

[3] M. Bächler, *et al.* Species-specific differences in follicular antral sizes result from diffusion-based limitations on the thickness of the granulosa cell layer. *Mol Hum Reprod*, 20(3), 2014.

[4] E. Bartocci and P. Lió. Computational modeling, formal analysis, and tools for systems biology. *PLoS Comput Biol*, 12(1), 2016.

[5] A. Biere, *et al.*, eds. *Handbook of Satisfiability*, vol. 185. IOS, 2009.

[6] S. Bogomolov, *et al.* Planning as model checking in hybrid domains. In *AAAI 2014*.

[7] S. Bogomolov, *et al.* PDDL+ planning with hybrid automata: Foundations of translating must behavior. In *ICAPS 2015*. AAAI, 2015.

[8] L. Bordeaux, *et al.* A unifying framework for structural properties of CSPs: Definitions, complexity, tractability. *JAIR*, 32, 2008.

[9] L. Bordeaux, *et al.* Generalizing consistency and other constraint properties to quantified constraints. *ACM TOCL*, 10(3), 2009.

[10] D. Bryce, *et al.* Smt-based nonlinear PDDL+ planning. In *AAAI 2015*.

[11] A. Calabrese, *et al.* Generating T1DM virtual patients for in silico clinical trials via AI-guided statistical model checking. In *RCRA 2019*, *CEUR W.P.*, vol. 2538. CEUR, 2019.

[12] A. Coles and A. Coles. PDDL+ planning with events and linear processes. In *ICAPS 2014*.

[13] C. Dalla Man, *et al.* The UVA/Padova type 1 diabetes simulator: New features. *JDST*, 8, 2014.

[14] G. Della Penna, *et al.* Upmurphi: A tool for universal planning on PDDL+ problems. In *ICAPS 2009*.

[15] G. Della Penna, *et al.* Planning for autonomous planetary vehicles. In *ICAS 2010*. IEEE, 2010.

[16] G. Della Penna, *et al.* Cgmurphi: Automatic synthesis of numerical controllers for nonlinear hybrid systems. *Eur J Control*, 19(1), 2013.

[17] European Medicines Agency. Reporting of physiologically based pharmacokinetic (PBPK) modelling and simulation, 2019. EMA/CHMP/458101/2016.

[18] A. Fabregat, *et al.* The Reactome pathway knowledgebase. *Nucl Ac Res*, 46(D1), 2018.

[19] A. Fogliata, *et al.* On the pre-clinical validation of a commercial model-based optimisation engine: Application to volumetric modulated arc therapy for patients with lung or prostate cancer. *Radiother Oncol*, 113(3), 2014.

[20] M. Fox and D. Long. Modelling mixed discrete-continuous domains for planning. *JAIR*, 27, 2006.

[21] F. Glover and G. Kochenberger, eds. *Handbook of Metaheuristics*, vol. 57. Springer, 2006.

[22] H. Gong, *et al.* Analysis and verification of the HMGB1 signaling pathway. *BMC Bioinf*, 11(S-7), 2010.

[23] G. Gottlob, *et al.* Conditional constraint satisfaction: Logical foundations and complexity. In *IJCAI 2007*.

[24] M. Hengartner, *et al.* Negative affect is unrelated to fluctuations in hormone levels across the menstrual cycle: Evidence from a multisite observational study across two successive cycles. *J Psycho Res*, 99, 2017.

[25] R. Hester, *et al.* Hummod: a modeling environment for the simulation of integrative human physiology. *Front Physiol*, 2, 2011.

[26] P. Jackson, *et al.* Patient-specific mathematical neuro-oncology: Using a simple proliferation and invasion tumor model to inform clinical practice. *Bull Math Biol*, 77(5), 2015.

[27] M. Kanehisa, *et al.* Kegg: New perspectives on genomes, pathways, diseases and drugs. *Nucl Ac Res*, 45(D1), 2017.

[28] B. Leeners, *et al.* Lack of associations between female hormone levels and visuospatial working memory, divided attention and cognitive bias across two consecutive menstrual cycles. *Front Behav Neuro*, 11, 2017.

[29] B. Leeners, *et al.* Associations between natural physiological and supraphysiological estradiol levels and stress perception. *Front Psycol*, 10, 2019.

[30] H. Li and B. Williams. Generative planning for hybrid systems based on flow tubes. In *ICAPS 2008*.

[31] N. Lipovetzky, *et al.* Classical planning with simulators: Results on the atari video games. In *IJCAI 2015*.

[32] J. Lippert, *et al.* Open systems pharmacology community–an open access, open source, open science approach to modeling and simulation in pharmaceutical sciences. *CPT: PSP*, 8(12), 2019.

[33] F. Maggioli, *et al.* SBML2Modelica: Integrating biochemical models within open-standard simulation ecosystems. *Bioinformatics*, 36(7), 2020.

[34] T. Mancini, *et al.* Evaluating ASP and commercial solvers on the CSPLib. *Constraints*, 13(4), 2008.

[35] T. Mancini, *et al.* Combinatorial problem solving over relational databases: View synthesis through constraint-based local search. In *SAC 2012*. ACM, 2012.

[36] T. Mancini, *et al.* Demand-aware price policy synthesis and verification services for smart grids. In *SmartGridComm 2014*. IEEE, 2014.

[37] T. Mancini, *et al.* Computing biological model parameters by parallel statistical model checking. In *IWBBIO 2015*, *LNCS*, vol. 9044. Springer, 2015.

[38] T. Mancini, *et al.* SyLVaaS: System level formal verification as a service. In *PDP 2015*. IEEE.

[39] T. Mancini, *et al.* User flexibility aware price policy synthesis for smart grids. In *DSD 2015*. IEEE.

[40] T. Mancini, *et al.* Anytime system level verification via parallel random exhaustive hardware in the loop simulation. *Microprocessors and Microsystems*, 41, 2016.

[41] T. Mancini, *et al.* SyLVaaS: System level formal verification as a service. *Fundam Inform*, 1–2, 2016.

[42] T. Mancini, *et al.* On minimising the maximum expected verification time. *Inf Proc Lett*, 122, 2017.

[43] T. Mancini, *et al.* Computing personalised treatments through in silico clinical trials. A case study on downregulation in assisted reproduction. In *RCRA 2018*, *CEUR W.P.*, vol. 2271. CEUR, 2018.

[44] T. Mancini, *et al.* An efficient algorithm for network vulnerability analysis under malicious attacks. In *ISMIS 2018*. Springer, 2018.

[45] T. Mancini, *et al.* Parallel statistical model checking for safety verification in smart grids. In *SmartGridComm 2018*. IEEE.

[46] M. Mateják and J. Kofránek. Physiomodel – An integrative physiology in Modelica. In *EMBC 2015*.

[47] M. Mazo, *et al.* PESSOA: A tool for embedded controller synthesis. In *CAV 2010*.

[48] L. O. Müller and E. F. Toro. A global multiscale mathematical model for the human circulation with emphasis on the venous system. *Int J Num Meth Biom Eng*, 30(7), 2014.

[49] F. Pappalardo, *et al. In silico* clinical trials: concepts and early adoptions. *Brief Bioinform*, 2018.

[50] W. Raghupathi and V. Raghupathi. Big data analytics in healthcare: promise and potential. *Health Inf Sci Syst*, 2(1), 2014.

[51] M. Ramirez, *et al.* Real–time UAV maneuvering via automated planning in simulations. In *IJCAI 2017*.

[52] B. Ribba, *et al.* A tumor growth inhibition model for low-grade glioma treated with chemotherapy or radiotherapy. *Clin Cancer Res*, 18(18), 2012.

[53] S. Röblitz, *et al.* A mathematical model of the human menstrual cycle for the administration of GnRH analogues. *J Theor Biol*, 321, 2013.

[54] F. Rossi, *et al.*, eds. *Handbook of Constraint Programming*. Elsevier, 2006.

[55] P. Roy and K. Roy. Molecular docking and QSAR studies of aromatase inhibitor androstenedione derivatives. *J Pharm Pharmacol*, 62(12), 2010.

[56] M. Rungger and M. Zamani. SCOTS: A tool for the synthesis of symbolic controllers. In *HSCC 2016*.

[57] S. Sinisi, *et al.* Optimal personalised treatment computation through in silico clinical trials on patient digital twins. *Fundam Inform*, 2020.

[58] E. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*, 1998.

[59] N. Trayanova, *et al.* Personalized imaging and modeling strategies for arrhythmia prevention and therapy. *Curr Opin Biomed Eng*, 5, 2018.

[60] E. Tronci, *et al.* Patient-specific models from inter-patient biological models and clinical records. In *FMCAD 2014*. IEEE, 2014.

[61] U.S.A. Food and Drug Administration. Physiologically based pharmacokinetic analyses – format and content guidance for industry, 2018. FDA-2016-D-3969.

[62] M. Vallati, *et al.* Efficient macroscopic urban traffic models for reducing congestion: A PDDL+ planning approach. In *AAAI 2016*.

[63] S. van Dijkman, *et al.* Individualized dosing algorithms and therapeutic monitoring for antiepileptic drugs. *Clin Pharmacol & Ther*, 103(4), 2018.

[64] M. Viceconti, *et al.* Credibility of in silico trial technologies – a theoretical framing. *IEEE J Biom Health Inf*, 24(1), 2020.

[65] S. Willmann, *et al.* PK-Sim®: a physiologically based pharmacokinetic 'whole-body' model. *BIOSILICO*, 1(4), 2003.

[66] F. Yalcinkaya, *et al.* Mathematical modelling of human heart as a hydroelectromechanical system. In *ELECO 2013*.