PhD Reseach Proposal

Identification of Dynamic Parameters for Gene Networks using Bio-inspired computation and Reinforcement Learning

Background : Master 2 in Computer science, or bio-informatics

Laboratory : Laboratoire I3S, UMR 6070 UNSA-CNRS, Algorithmes-Euclide-B, 2000 route des Lucioles, B.P. 121, 06903 Sophia Antipolis CEDEX (between Nice and Cannes, French Riviera)

PhD Supervision :

- Jean-Paul Comet, full professeur in computer science and bio-informatics, lab. I3S, `jean-paul.comet@univ-cotedazur.fr`, `http://www.i3s.unice.fr/~comet/`
- Denis Pallez, professor in computer science, lab I3S, `denis.pallez@univ-cotedazur.fr`, `http://denispallez.i3s.unice.fr/`

Keywords : Modelling of gene networks, AI, Reinforcement Learning, bio-inspired methods.

Scientific Description

Objectives : In molecular biology, a large range of interactions between genes, proteins, and molecular entities has been widely studied, and large maps have been built. Nevertheless, the knowledge of interactions is not sufficient to explain the dynamics of the system and to hope being able to control it : the dynamics is also governed by parameters explaining the relative strengths between regulators. The goal of this PhD is to develop methodologies able to automate the identification of parameter values associated to a gene regulatory network, from the knowledge of gene regulatory graphs and from some behavioural properties.
The research will focus on formal methods and artificial intelligence (bio-inspired methods and reinforcement learning) in order to automate the identification of parameters of formal models that can best explain the observations. Current collaborations with biologists will give rise to fruitful applications by, for example, designing a temporised model of the coupling between circadian clock, cellular cycle and metabolism regulation.

Scientific Context : The group « Bioinfo formelle » of I3S laboratory has proposed different formal frameworks to study the behaviour of gene regulatory networks. A regulatory network, modelled as a graph, defines the static interactions of a biological system : in the gene regulatory networks modelling setting, each interaction abstracts the individual influence of a gene x on the expression of another gene y. The dynamic of the network is governed by numerous unknown parameters that we want to identify. René Thomas has proposed a discrete modelling framework that allows approximations of sigmoid functions (that usually represent the evolution of target gene expression level according to the level of the regulatory entity) by step functions. Using this modelling framework, the exploration of the qualitative dynamics of the system can be described as path exploration in a finite state space and formal methods from computer science are useful [1].

In order to take into account temporal information, which plays crucial role in a wide range of biological systems, we also have developed, for several years, a hybrid modelling framework [2,3], which measures the time to go from a state to another one. This framework consists in a particular class of hybrid automata, but the crucial problem remains in the determination of accurate values for all numerous parameters. In order to address this question, we developed a "weakest precondition calculus" inspired by Hoare's Logic (initially dedicated to imperative programs) which leads to constraints on parameters which has to be fulfilled in order to make the model's dynamics to be consistent with the observations [4, 5, 6]. However, the exploitation of the constraints generated by the Hoare Logics is not so easy : classical constraints solvers are not able to extract solutions.

Methodology : In this PhD, we envisage to use Artificial Intelligence (AI) techniques to overcome this problem whose difficulty is reinforced by the enormous number of parameters piloting the currently studied gene networks. In a first approach, we envisage to consider this problem as a large-scale optimization problem by introducing a similarity metric between the execution of the parametrized model and the observations. As this metric is continuous but not differentiable, bio-inspired meta-heuristics will be considered [7, 8, 9] as they are very competitive.

An alternative to previous approach consists in considering the problem as a sequential decision problem, which parameter to test next, and address it with reinforcement learning [10]. This technique focuses on how some software agents ought to take actions in an environment in order to maximize some notion of cumulative reward. In order to act near optimally, the agent must reason about the long term consequences of its actions (i.e., maximize future income), although the immediate reward associated with this might be negative. Thus, reinforcement learning is particularly well-suited to problems that include a long-term versus short-term reward tradeoff. Here, the rewards have to be deduced from the similarity measure between the temporal traces in the current model and the observed traces coming from experiments.

One of the most well-known technique in reinforcement learning is Monte Carlo Tree Search [11], used for beating the world champion in Go. This technique uses a policy for balancing exploration and exploitation when selecting the most 'interesting' state in the search space. Upper Confidence Bound is one of such policy but many other variants exist (AMAF, RAVE, UCB1-Tuned...). Choosing one of them generates a different search tree leading to choose a sub-optimal path. Current tendency is to optimize on-the-fly, during simulation, the policy based on bio-inspired computation [12] as it is done in [13].

Work plan : This PhD proposal is ambitious because it mixes different sub-fields of computer science, bio-informatics and AI techniques. PhD' work could be split into the following phases :

1. getting started and understanding the hybrid formalism for gene networks, the associated hybrid Weakest Precondition calculus and its implementation in OCaml (Holmes BioNet),
2. investigating the use of large scale bio-inspired computation techniques to determine some solutions of the constraints problem, and compare them to ones obtained by classical constraints solver,
3. investigating the use of Reinforcement Learning (MCTS) as an alternative for same problem,
4. investigate a hybrid approach where the policy is obtained thanks to bio-inspired methods.

Scientific Impact : A numerous number of biological phenomena is temporized and in medicine, more and more responses of human organism seem to be linked to the circadian clock [14]. Thus, chronotherapy becomes a field of research full of hope, and it becomes crucial to develop some frameworks able to handle both qualitative dynamics and associated temporal aspects. The goal of this PhD is to develop a method which will be able to automate the construction of hybrid models of biological systems including the key-point of parameters' identification, making possible the study of large gene regulatory networks. The team has outlined for several years the potential of formal methods for studying different models (circadian clock, regulation of metabolism, and cell cycle) in the context of qualitative modelling. The integration of these three models in the continuous-time setting, would be of great help in deciphering the main causalities that could explain the impact of daily drug administration time on effectiveness.

National and international collaborations : The team has long established close links with several biology teams on the territory of the Côte d'Azur University and elsewhere. A new qualitative model of the circadian cycle has been developed with the iBV (Franck Delaunay), a first qualitative model of the regulation of metabolism has been developed with the Metabologic team from Mauritius (Radjeev Khoodeeram) and with the laboratory BIRL from SupBioTech (Jean-Yves Trosset), and the adaptation of this model to metabolism of pancreatic cancer cells is underway with CRCM-Inserm in Marseilles (Sophie Vasseur).

Bio-inspired meta-heuristics have been used for identifying active module in a regulatory network in collaboration with iBV (O. Soriani) and Claude Pasquier. The work has been accepted in CSBio 2019 conference. In another domain, we used Particle Swarm Optimisation for tuning a Image Signal Processor for an industrial project.

References :

[1] G. Bernot, J.-P. Comet, A. Richard, J. Guespin, *Application of Formal Methods to Biological Regulatory Networks : Extending Thomas' Asynchronous Logical Approach with Temporal Logic.* J.T.B.. 229(3) :339-347, 2004.

[2] J. Behaegel, J.-P. Comet, G. Bernot, E. Cornillon, F. Delaunay, *A hybrid model of cell*

cycle in mammals, 6th International Conference on Computational Systems-Biology and Bioinformatics (CSBio'2015), Bangkok, Thailand. November 22-25, 2015.

[3] J. Behaegel, J.-P. Comet, G. Bernot, E. Cornillon and F. Delaunay, *A hybrid model of cell cycle in mammals.* J. of Bioinformatics and Computational Biology. 14(1) :1640001 [17 pp.], 2016.

[4] G. Bernot ; J.-P. Comet ; Z. Khalis ; A. Richard ; O.F. Roux, *A Genetically Modified Hoare Logic.* Theoretical Computer Science. 765 :145-157, 2019. DOI 10.1016/j.tcs.2018.02.003.

[5] J. Behaegel, J.-P. Comet and M. Folschette, *Constraint Identification Using Modified Hoare Logic on Hybrid Models of Gene Networks*, International Symposium on Temporal Representation and Reasoning (TIME), Mons, Belgium. 16-18 October 2017. pp. 5 :1-5 :21

[6] J. Behaegel, J.-P. Comet and M. Pelleau. *Identification of Dynamic Parameters for Gene Networks*, the 30th IEEE International Conference on Tools with Artificial Intelligence (ICTAI), Volos, Greece. November 5-7, 2018.

[7] Deb, K., Myburgh, C., *Breaking the Billion-Variable Barrier in Real-World Optimization Using a Customized Evolutionary Algorithm.* GECCO, 2016, https ://doi.org/10.1145/2908812.2908952

[8]Mahdavi, S., Shiri, M.E., Rahnamayan, S., *Metaheuristics in large-scale global continues optimization : A survey*, 2015, https ://doi.org/10.1016/j.ins.2014.10.042

[9] IEEE *Task Force on Large-Scale Global Optimization*, http ://www.tflsgo.org

[10] H. van Hasselt, *Reinforcement Learning in Continuous State and Action Spaces*, in "Reinforcement Learning : State-of-the-Art", Springer Berlin Heidelberg, 2012.

[11] C. B. Browne et al., *A Survey of Monte Carlo Tree Search Methods*, in IEEE Transactions on Computational Intelligence and AI in Games, vol. 4, no. 1, pp. 1-43, March 2012., 10.1109/TCIAIG.2012.2186810

[12] J. Del Ser, E. Osaba, D. Molina, X. Yang, S. Salcedo-Sanz, D. Camacho, S. Das, P. N. Suganthan, C. A. Coello Coello, F. Herrera, *Bio-inspired computation : Where we stand and what's next, Swarm and Evolutionary Computation*, Vol. 48, 2019, https ://doi.org/10.1016/j.swevo.2019.04.008.

[13] X. Ma, K. Driggs-Campbell, Z. Zhang, M. J. Kochenderfer, *Monte Carlo Tree Search for Policy Optimization*, IJCAI 2019, https ://doi.org/10.24963/ijcai.2019/432

[14] R. Zhang, N. Lahens, H. Ballance, M. Hughes, and J. Hogenesch, *A circadian gene expression atlas in mammals : implications for biology and medicine.* Proceedings of the National Academy of Sciences, 111(45) :16219–16224, 2014.