
IDENTIFYING DYNAMIC REGULATION WITH ADVERSARIAL SURROGATES

Ron Teichner^{1,2}, Naama Brenner^{3,2}, and Ron Meir^{1,2}

¹Viterbi Dept. of Electrical & Computer Eng., Technion, Israel Institute of Technology

²Network Biology Research Lab, Technion, Israel Institute of Technology

³Dept. of Chemical Eng., Technion, Israel Institute of Technology

ABSTRACT

Homeostasis, the ability to maintain a stable internal environment in the face of perturbations, is essential for the functioning of living systems. Given observations of a system, or even a detailed model of one, it is both valuable and extremely challenging to extract the control objectives of the homeostatic mechanisms. Lacking a clear separation between plant and controller, frameworks such as inverse optimal control and inverse reinforcement learning are unable to identify the homeostatic mechanisms. A recently developed data-driven algorithm, Identifying Regulation with Adversarial Surrogates (IRAS), detects highly regulated or conserved quantities as the solution of a min-max optimization scheme that automates classical surrogate data methods. Yet, the definition of homeostasis as regulation within narrow limits is too strict for biological systems which show sustained oscillations such as circadian rhythms. In this work, we introduce Identifying Dynamic Regulation with Adversarial Surrogates (IDRAS), a generalization of the IRAS algorithm, capable of identifying control objectives that are regulated with respect to a dynamical reference value. We test the algorithm on simulation data from realistic biological models and benchmark physical systems, demonstrating excellent empirical results.

Keywords model discovery · control · adversarial · machine learning

1 Introduction

Living systems maintain stability against internal and external perturbations, a phenomenon known as homeostasis [7, 28, 31]. This is a ubiquitous central pillar across all scales of biological organization, such as molecular circuits, physiological functions, and population dynamics. Failure of homeostatic control is associated with diseases including diabetes, autoimmunity, and obesity [31]. It is therefore vital to identify the regulated variables that the system aims to maintain at a stable setpoint. Unlike simple human-made systems, where known pre-selected variables are under control, biological systems are characterized by multiple coupled control loops as well as other dynamic structures [7]. A fundamental difference is that they are not divided to separate ‘plant’ and ‘controller’ entities, as is commonly assumed in control theory, but rather make up a complex network of interactions. In such a network regulated variables may be maintained at a stable setpoint, as in the classic example of the regulation of blood glucose concentration, a tightly regulated biological process in which the rates of glycolysis and gluconeogenesis are flexible variables [31]. A particularly interesting aspect is that, under certain conditions, biological homeostatic controllers may become oscillatory [48]. Indeed, while the occurrence of oscillations is generally avoided in control engineering, oscillatory behavior is ubiquitously found in natural systems [32, 17, 35]. A well documented example is the Baroreflex control of the cardiovascular system where slow blood-pressure and heart-rate oscillations are observed [15]. The biological complexity makes it challenging to identify the regulated variables that the system actively maintains in the vicinity of an attractor and, as we further explain, dedicated algorithmic tools are required.

Biological systems are commonly modeled by a set of dynamical equations where internal variables and control variables are not clearly separated. Therefore, the controlled objective and the control signal in a biological system model are implicit. Although control theory typically assumes the existence of a separate plant and controller (see

Chapter 1.2 in [6]), many theoretical results and analysis tools do not require such a separation [12]. Identifying Regulation with Adversarial Surrogates (IRAS), a recently published algorithm directly addresses the plant-controller coupling, and is thus dedicated to control identification in biological systems. IRAS identifies variables that are regulated within narrow limits and is thus incompatible with biological systems which show sustained oscillations such as circadian rhythms [38]. In this work we introduce Identifying Dynamic Regulation with Adversarial Surrogates (IDRAS), a fundamental extension of the IRAS algorithm, capable of identifying control objectives that are regulated with respect to a dynamic reference value. We test the algorithm on synthetic data of kinetic protein interactions, bacterial growth and division and on purely physical systems demonstrating excellent empirical results. See Appendix Section C for a review of related work.

The IRAS algorithm assumes that it is in general a *combination* of the observed variables which is internally controlled. Therefore, given a set of measurements over time, $z(t) \in \mathbb{R}^m$, IRAS finds a combination $g(\cdot)$ of the measured variables, that is maintained around a setpoint,

$$g(z(t)) \approx c_{\text{set}}, \quad (1)$$

where we can assume $c_{\text{set}} = 0$ without loss of generality. In this work, we consider systems that are regulating some unknown combination of the observables over an unknown dynamically changing reference value,

$$g(z(t)) \approx c_e(t). \quad (2)$$

By adding the subscript e we emphasize that the behavior of $c_e(t)$ can be environment dependent. For example, circadian rhythms are regulated on molecular and physiological levels, but control parameters are entrained to the daylight cycle [39, 16]. Here we do not assume any specific form neither for the combination $g(\cdot)$ nor for the process $c_e(\cdot)$. For illustration, a general model that might yield observations that satisfy $g(z(t)) = c_e(t)$ is,

$$du_e = f_{u_e}(u_e)dt + dW_{u_e}, \quad dx = f_x(x, u_e)dt + dW_x, \quad z_k = h(x(t_k)) + \eta_{z_k}, \quad (3)$$

where $x \in \mathbb{R}^n$ is the state of the observed system, $u_e \in \mathbb{R}^{n_u}$ is an input from the environment, $z_k \in \mathbb{R}^m$ is the observation sampled at time t_k , dW_{u_e} and dW_x are Wiener processes and η_{z_k} is a measurement noise.

In general, any system identification procedure that does not decouple $g(\cdot)$ from $c_e(\cdot)$ will yield results that do not generalize to other environments where $c_{e'}(t)$ has a different behavior. It is therefore vital to separately identify the regulated combination from the dynamics of the reference. In the next Section we present IDRAS, a purely data-driven algorithm that simultaneously learns a combination $g(\cdot)$ and the dynamic process that it follows, $c_e(\cdot)$.

2 Algorithm development

We are interested in identifying empirically, from a set of measurements, a variable combination which tightly follows a dynamically changing reference value, where both are unknown. The combination could represent an internal quantity of high importance to the system, and the reference value could reflect temporal trends in the environment. Given observations $z_k \in \mathbb{R}^m$ at discrete times $k \in \{1, 2, \dots, N\}$, we search for $g_\theta : \mathbb{R}^m \rightarrow \mathbb{R}$, a function parameterized by θ such that,

$$c_k = g_\theta(z_k), \quad (4)$$

where the resulting time-series c_k is the (learned) reference. We assume that this time-series follows some dynamics and it is therefore possible to learn a filter¹, namely a predictor of the value of c_k at time k based on its previous values, formally given by

$$\hat{c}_k = F_{\phi, \omega, \delta}(c_{k-T:k-1}). \quad (5)$$

The structure of the learned filter F , and the meaning of its parameters, is detailed at the end of the present Section, and $c_{k-T:k-1} = [c_{k-T}, c_{k-T+1}, \dots, c_{k-1}]$ with T a hyper-parameter. The filtering error is,

$$e_k = c_k - \hat{c}_k \triangleq \mathcal{E}_\Omega(z_{k-T:k}), \quad (6)$$




Figure 1: IDRAS algorithm outline. The observed time-series z is permuted to create the unconstrained series z^* . The shuffle player, only exposed to the $1D$ projections under $\mathcal{E}_\Omega(\cdot)$, sets the distributions under the projection are identical, (8). Then the combination player, given z and \tilde{z} , updates the parameters Ω towards minimizing the CR, (7). The block \mathcal{E}_Ω , based on (6) replaces the function $g(\cdot)$ in Fig. 3 in [45].

¹Formally (5) is a 1-step predictor, but, following [5], we refer to it as a filter.

and is a function of the parameters $\Omega = [\theta, \phi, \omega, \delta]$. Our goal is to learn both $g_\theta(\cdot)$ and $F_{\phi, \omega, \delta}(\cdot)$ such that the error is small, namely $e_k \approx 0$ for all k . We note that a straightforward optimization yields the trivial pair, $g(\cdot) \equiv 0, F(\cdot) \equiv 0$.

The problem formulated is in fact a generalization of problem (1). It is immediate to verify that by setting $F_{\phi, \omega, \delta}(\cdot) \equiv 0$, we obtain $\mathcal{E}_\Omega(z_{k-T:k}) = g_\theta(z_k)$, $e_k = c_k$ and thus we optimize to find a combination $g_\theta(\cdot)$ such that $c_k \approx 0$, namely regulated to a narrow range around a fixed value.

This simpler problem was addressed in [46] by utilizing comparisons between original time-series and random (though possibly constrained) shuffled time-series. A regulated combination is presumably composed of components that co-vary to compensate and buffer perturbations. Therefore, shuffling the temporal order of each component separately would ruin these co-variations and greatly increase the variance of the combination over time. To quantify this notion, a ‘‘coefficient of regulation’’ (CR) is defined as the ratio between the standard deviations of the combination in the original data c_k and in the shuffled data \tilde{c}_k . In [46] it was found that straightforward optimization of this measure is insufficient to escape trivial solutions and artefacts. Rather, a two-player algorithm was constructed which alternates between minimizing the CR and modifying the shuffled time-series. This algorithm terminates when the CR can no longer be minimized, while at the same time reproducing general geometric features of the data in parameter space. Relying on similar concepts, the next Section introduces the generalization of this algorithm to time-varying regulation set-points.

Identifying Dynamic Regulation with Adversarial Surrogates

IDRAS is an iterative algorithm consisting of two competing players, a generalization of the IRAS algorithm. The first player aims to minimize the CR, a measure of invariance (see [46] Section 1), while the second player aims to render the task of the first player more difficult by forcing it to extract information about the temporal structure of the data, which is absent from time-shuffled ‘‘surrogate’’ data.

Combination player At iteration i , the first player sets the parameters Ω to minimize the CR,

$$\Omega^{(i)} = \operatorname{argmin}_{\Omega} \sigma(e_{T+1:N}) / \sigma(\tilde{e}_{T+1:N}^{(i)}), \quad (7)$$

where $\sigma(e_{T+1:N})$ and $\sigma(\tilde{e}_{T+1:N}^{(i)})$ are the standard-deviations of the time-series e_k and $\tilde{e}_k^{(i)}$ respectively. The time-series e_k is the filtering error on the original data and the time-series $\tilde{e}_k^{(i)}$ is the filtering error on surrogate data, the outcome of a resampling procedure using a resampling function $\zeta^{(i-1)}(\cdot)$ that was set by the second player in the previous iteration. Formally, $e_k = \mathcal{E}_\Omega(z_{k-T:k})$, $\tilde{e}_k^{(i)} = \mathcal{E}_\Omega(z_{k-T:k-1} \tilde{z}_k^{(i)})$, and $\tilde{z}_k^{(i)} \sim P_{\tilde{z}}^{\zeta^{(i-1)}}(\tilde{z}_k^{(i)} | z_{k-T:k-1}) = P_z(\tilde{z}_k^{(i)}) \zeta^{(i-1)}(\tilde{e}_k^{(i)})$, where $P_z(z)$ is the probability of observing a measurement z at some random time.

Optimizing (7), the combination player searches for a combination $g_\theta(\cdot)$, whose filtering error (which is the 1-step prediction error) $\sigma(e_{T+1:N})$ is small w.r.t. the error in predicting a sample from a random point in time, $\sigma(\tilde{e}_{T+1:N}^{(i)})$. This encourages the combination player to find a useful (g, F) pair, representing a meaningful underlying quantity in the observed system. We note that for $\zeta(\cdot) \equiv 1$ the time-series \tilde{z}_k is a naïve random permutation of the time-series z_k . Optimizing (7) w.r.t. such unconstrained permutation leads to artifacts, see Appendix Section A for elaboration. To avoid these artifacts the second player constrains the time-series \tilde{z}_k by setting the resampling function $\zeta(\cdot)$.

Shuffle player The second player uses the current solution of the combination player, $\Omega^{(i)}$, to create a new shuffled time-series $\tilde{z}_k^{(i+1)}$, which resembles the statistical structure of the data under the 1D projection $\mathcal{E}_{\Omega^{(i)}}$. Formally, this corresponds to the selection of a resampling function $\zeta^{(i)}(\cdot)$ that minimizes the distributional distance,

$$\zeta^{(i)} = \operatorname{argmin}_{\zeta} \mathcal{D}(e_{T+1:N}, \tilde{e}_{T+1:N}), \quad (8)$$

where $e_k = \mathcal{E}_{\Omega^{(i)}}(z_{k-T:k})$, $\tilde{e}_k = \mathcal{E}_{\Omega^{(i)}}(z_{k-T:k-1}, \tilde{z}_k)$, and $\tilde{z}_k \sim P_{\tilde{z}}^{\zeta}(\tilde{z}_k | z_{k-T:k-1}) = P_z(\tilde{z}_k) \zeta(\tilde{e}_k)$.

Lemma 1. (Shuffle player’s optimal solution.) *The shuffle player, who only has access to the error time-series, can solve optimization (8) and obtain $\mathcal{D} = 0$ by choosing $\zeta(\mathcal{E}_{\Omega^{(i)}}(\cdot)) = P_{z_{(T)}}(\cdot) / P_{z_{(T)}^*}(\cdot)$, where $P_{z_{(T)}}(z_{k-T:k})$ is the probability of observing the sequence $z_{k-T:k}$ (at a random time k) and $P_{z_{(T)}^*}(z_{k-T:k-1}, z) = P_{z_{(T-1)}}(z_{k-T:k-1}) P_z(z)$ is the probability of observing the sequence $z_{k-T:k-1}$ followed by a time-random observation, sampled from the naïvely permuted time-series.*

The two players inform each other of their current step results, and the process continues iteratively until the combination player can no longer decrease the CR in (7). We refer to this algorithm as IDRAS and depict its outline in Figure 1. The filter block (5), the proof of Lemma 1 and a performance assessment metric are given in the Appendix Section A.

3 Validation

After presenting the construction of IDRAS, we seek to validate it on datasets with a known control objective, so that the quality of the results can be assessed. We chose two validation examples of biological models: a kinetic model of protein interactions and a model of bacterial life cycle. Additionally, to demonstrate the efficiency of our algorithm in studying physical systems, we validate the algorithm on a dataset that serves for benchmarking machine-learning algorithms. We briefly depict here the first examples and leave a detailed explanation of all examples to the Appendix Section B. Code reproducing all examples is available at <https://github.com/RonTeichner/IRAS>.

A kinetic model of interactions

A kinetic model describes regulatory interactions in the production of two proteins incorporating a feedback loop. In the considered model (inspired by [22]), the total amount of two proteins P and S , namely $P + S$, is controlled by M , the mRNA molecule that is produced through a process of constitutive transcription at a rate K and degraded with first order kinetics at a rate γ_M . The model is under constant perturbations to the protein expression rate K . The model is described by the differential equations

$$\begin{aligned} K(t) &= K_0(1 + 0.5 \cos(2\pi \frac{1}{\tau_K} t + \phi_K)), & dM &= (K(t) - f(P + S) - \gamma_M M)dt, \\ dP &= (k_P M - \gamma_P P)dt + \eta_P dW_P, & dS &= (k_S M - \gamma_S S)dt + \eta_S dW_S, \end{aligned} \quad (9)$$

where the mRNA M and the two proteins P, S are linked in a feedback loop. Small changes in S or P , modeled by the increments of the Wiener processes dW_P and dW_S , induce swift and sharp changes in the transcription of M and maintain $P + S$ around a reference level

$$c^*(t) \triangleq (P + S)(t) = \left(f + \frac{\gamma_M \gamma_P \gamma_S}{k_P \gamma_S + k_S \gamma_P} \right)^{-1} K(t), \quad (10)$$

which is reflected in a high negative correlation between S and P . In (9) we identify the components of the general model (3): $u_e = [K, M]$, $x = [P, S]$. Our observations contain the levels of the two proteins P and S (an observation taken in labs by measuring fluorescence intensity of constitutively expressed proteins [11]), $z_k = [P(t = k/f_s), S(t = k/f_s)]$, where $k \in [1, \dots, N]$ and f_s is the sampling rate. Figure 4a illustrates the kinetic interactions model, parameter values are listed in the Appendix.

We ran both IRAS and IDRAS algorithms in search of the control objective. Figure 4b depicts the output (dashed red) of $g_\theta(P, S)$ trained by the IRAS algorithm. Due to the lack of a regulated constant combination, IRAS did not converge. When running IDRAS a combination that tightly follows a dynamic reference was found despite its oscillating nature, as depicted in Figure 4c. See the Appendix for performance assessment metric scores.

4 Discussion

Biological systems are often “reversed engineered” to build mathematical models based on observed data [34, 41, 13, 8, 42, 25, 10]. Detecting regulatory processes is a technically challenging problem with many potential applications [1, 36, 2, 18, 20, 19, 21]. Recently an algorithm named IRAS was introduced [46] that receives as input raw dynamic measurements and provides combinations of the observables that are maximally conserved across time. Here we presented IDRAS, a generalization of IRAS, capable of identifying control objectives that are regulated w.r.t. a dynamical reference value. This algorithm allows the identification of control objectives in biological systems which show sustained oscillations, such as circadian rhythms. We provide validation in three distinct realistic examples demonstrating excellent empirical results.

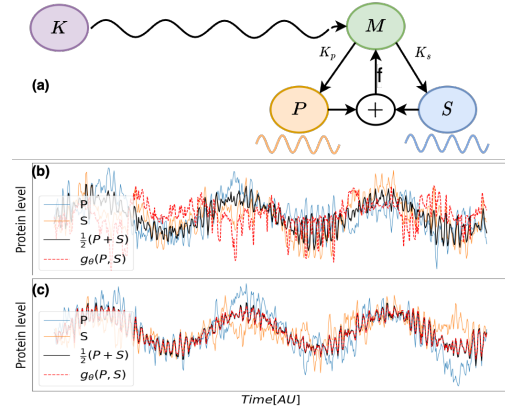


Figure 2: (a) Illustration of the closed-loop model (9). The mRNA molecule M induces the production of proteins S and P and receives a negative feedback of their sum. (b) The IRAS algorithm does not converge, due to the lack of a regulated constant combination. The combination $P + S$ (black) along with the output of IRAS (dashed red). (c) IDRAS algorithm captures the control objective and its oscillating trend.

References

- [1] Abdullahi Rashid Adem and Chaudry Masood Khalique. Symmetry reductions, exact solutions and conservation laws of a new coupled kdv system. *Communications in Nonlinear Science and Numerical Simulation*, 17(9):3465–3475, 2012.
- [2] Abdullahi Rashid Adem and Ben Muatjetjeja. Conservation laws and exact solutions for a 2d zakharov–kuznetsov equation. *Applied Mathematics Letters*, 48:109–117, 2015.
- [3] Ferran Alet, Dylan Doblar, Allan Zhou, Josh Tenenbaum, Kenji Kawaguchi, and Chelsea Finn. Noether networks: meta-learning useful conserved quantities. *Advances in Neural Information Processing Systems*, 34:16384–16397, 2021.
- [4] Ariel Amir. Cell size regulation in bacteria. *Physical review letters*, 112(20):208102, 2014.
- [5] Brian DO Anderson and John B Moore. *Optimal filtering*. Courier Corporation, 2012.
- [6] Karl Johan Åström and Richard M Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2021.
- [7] George E Billman. Homeostasis: the underappreciated and far too often ignored central organizing principle of physiology. *Frontiers in Physiology*, page 200, 2020.
- [8] Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.
- [9] Michael B Chang, Tomer Ullman, Antonio Torralba, and Joshua B Tenenbaum. A compositional object-based approach to learning physical dynamics. *arXiv preprint arXiv:1612.00341*, 2016.
- [10] Boyuan Chen, Kuang Huang, Sunand Raghupathi, Ishaan Chandratreya, Qiang Du, and Hod Lipson. Automated discovery of fundamental variables hidden in experimental data. *Nature Computational Science*, 2(7):433–442, 2022.
- [11] Yong-Joon Choi and Kazuaki Sawada. Physical sensors: Fluorescence sensors. In Roger Narayan, editor, *Encyclopedia of Sensors and Biosensors (First Edition)*, pages 1–19. Elsevier, Oxford, first edition edition, 2023.
- [12] Carlo Cosentino and Declan Bates. *Feedback control in systems biology*. Crc Press, 2011.
- [13] Bryan C Daniels and Ilya Nemenman. Automated adaptive inference of phenomenological dynamical models. *Nature communications*, 6(1):1–8, 2015.
- [14] Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. End-to-end differentiable physics for learning and control. *Advances in neural information processing systems*, 31, 2018.
- [15] M Di Rienzo, G Parati, A Radaelli, and P Castiglioni. Baroreflex contribution to blood pressure and heart rate oscillations: time scales, time-variant characteristics and nonlinearities. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1892):1301–1318, 2009.
- [16] Jeanne F Duffy and Kenneth P Wright Jr. Entrainment of the human circadian system by light. *Journal of biological rhythms*, 20(4):326–338, 2005.
- [17] Jay C Dunlap, Jennifer J Loros, and Patricia J DeCoursey. *Chronobiology: biological timekeeping*. Sinauer Associates, 2004.
- [18] OH El-Kalaawy. Variational principle, conservation laws and exact solutions for dust ion acoustic shock waves modeling modified burger equation. *Computers & Mathematics with Applications*, 72(4):1031–1041, 2016.
- [19] OH El-Kalaawy. New: Variational principle–exact solutions and conservation laws for modified ion-acoustic shock waves and double layers with electron degenerate in plasma. *Physics of Plasmas*, 24(3):032308, 2017.
- [20] OH El-Kalaawy. Modulational instability: conservation laws and bright soliton solution of ion-acoustic waves in electron-positron-ion-dust plasmas. *The European Physical Journal Plus*, 133(2):1–12, 2018.
- [21] OH El-Kalaawy, SM Moawad, and Shrouk Wael. Stability: Conservation laws, painlevé analysis and exact solutions for s-kp equation in coupled dusty plasma. *Results in physics*, 7:934–946, 2017.
- [22] Hana El-Samad. Biological feedback control—respect the loops. *Cell Systems*, 12(6):477–487, 2021.
- [23] Richard P Feynman, Robert B Leighton, and Matthew Sands. The feynman lectures on physics; vol. i. *American Journal of Physics*, 33(9):750–752, 1965.
- [24] Jacopo Grilli, Matteo Osella, Andrew S Kennard, and Marco Cosentino Lagomarsino. Relevant parameters in models of cell division control. *Physical Review E*, 95(3):032411, 2017.

- [25] Adam Haber and Elad Schneidman. Learning the architectural features that predict functional similarity of neural networks. *Physical Review X*, 12(2):021051, 2022.
- [26] Jessica B Hamrick, Kelsey R Allen, Victor Bapst, Tina Zhu, Kevin R McKee, Joshua B Tenenbaum, and Peter W Battaglia. Relational inductive bias for physical construction in humans and machines. *arXiv preprint arXiv:1806.01203*, 2018.
- [27] Po-Yi Ho, Jie Lin, and Ariel Amir. Modeling cell size regulation: From single-cell-level statistics to molecular mechanisms and population-level effects. *Annual review of biophysics*, 47:251–271, 2018.
- [28] Victoria Hsiao, Anandh Swaminathan, and Richard M Murray. Control theory for synthetic biology: recent advances in system characterization, control design, and controller implementation for synthetic biology. *IEEE Control Systems Magazine*, 38(3):32–62, 2018.
- [29] Suckjoon Jun and Sattar Taheri-Araghi. Cell-size maintenance: universal strategy revealed. *Trends in microbiology*, 23(1):4–6, 2015.
- [30] Patrick Kidger, James Foster, Xuechen Li, Harald Oberhauser, and Terry Lyons. Neural SDEs as Infinite-Dimensional GANs. *International Conference on Machine Learning*, 2021.
- [31] Maya E Kotas and Ruslan Medzhitov. Homeostasis, inflammation, and disease susceptibility. *Cell*, 160(5):816–827, 2015.
- [32] Dorothy T Krieger. The clocks that time us:(physiology of the circadian timing system). *Psychosomatic Medicine*, 44(6):559–560, 1982.
- [33] Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. *International Conference on Artificial Intelligence and Statistics*, 2020.
- [34] Lennart Ljung. System identification. In *Signal analysis and prediction*, pages 163–173. Springer, 1998.
- [35] David Lloyd and Ernest L Rossi. *Ultradian rhythms in life processes: An inquiry into fundamental principles of chronobiology and psychobiology*. Springer Science & Business Media, 2012.
- [36] Stanislav Yu Lukashchuk. Conservation laws for time-fractional subdiffusion and diffusion-wave equations. *Nonlinear Dynamics*, 80(1):791–802, 2015.
- [37] Liang Luo, Yang Bai, and Xiongfei Fu. Stochastic threshold in cell size control. *Physical Review Research*, 5(1):013173, 2023.
- [38] Martin C Moore-Ede. Physiology of the circadian timing system: predictive versus reactive homeostasis. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 250(5):R737–R752, 1986.
- [39] Colin S Pittendrigh and Dorothea H Minis. The entrainment of circadian oscillations by light and their role as photoperiodic clocks. *The American Naturalist*, 98(902):261–294, 1964.
- [40] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. *Advances in neural information processing systems*, 30, 2017.
- [41] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *science*, 324(5923):81–85, 2009.
- [42] Jingxiang Shen, Feng Liu, Yuhai Tu, and Chao Tang. Finding gene network topologies for given biological function with recurrent neural network. *Nature communications*, 12(1):1–10, 2021.
- [43] Kate Sidaway-Lee, Maria J Costa, David A Rand, Bärbel Finkenstadt, and Steven Penfield. Direct measurement of transcription rates reveals multiple mechanisms for configuration of the arabidopsis ambient temperature response. *Genome biology*, 15:1–10, 2014.
- [44] Lee Susman, Maryam Kohram, Harsh Vashistha, Jeffrey T Nechleba, Hanna Salman, and Naama Brenner. Individuality and slow dynamics in bacterial growth homeostasis. *Proceedings of the National Academy of Sciences*, 115(25):E5679–E5687, 2018.
- [45] Ron Teichner, Ron Meir, and Danny Eitan. Enhancing causal estimation through unlabeled offline data. *arXiv preprint arXiv:2202.07895*, 2022.
- [46] Ron Teichner, Aseel Shomar, Omri Barak, Naama Brenner, Shimon Marom, Ron Meir, and Danny Eytan. Identifying regulation with adversarial surrogates. *Proceedings of the National Academy of Sciences*, 120(12):e2216805120, 2023.
- [47] Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.

- [48] Kristian Thorsen, Oleg Agafonov, Christina H Selstø, Ingunn W Jolma, Xiao Y Ni, Tormod Drengstig, and Peter Ruoff. Robust concentration and frequency control in oscillatory homeostats. *PLoS One*, 9(9):e107766, 2014.
- [49] Silviu-Marian Udrescu and Max Tegmark. Ai feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16):eaay2631, 2020.
- [50] Nicholas Watters, Daniel Zoran, Theophane Weber, Peter Battaglia, Razvan Pascanu, and Andrea Tacchetti. Visual interaction networks: Learning a physics simulator from video. *Advances in neural information processing systems*, 30, 2017.
- [51] Hong Wei Yang, Min Guo, and Hailun He. Conservation laws of space-time fractional mzk equation for rossby solitary waves with complete coriolis force. *International Journal of Nonlinear Sciences and Numerical Simulation*, 20(1):17–32, 2019.
- [52] Aston Zhang, Zachary C. Lipton, Mu Li, and Alexander J. Smola. *Dive into Deep Learning*. 2020. <https://d21.ai>.

A Appendix A: Algorithm development

Disabling the shuffle player in the algorithm, that is, setting $\zeta(\cdot) \equiv 1$, the time-series \tilde{z}_k is a naïve random permutation of the time-series z_k . Optimizing (7) w.r.t such unconstrained permutation lead to artifacts. To illustrate consider that one of the observables is merely a Wiener process. The variance of a 1-step prediction error is proportional to dt while the variance of predicting a random sample increases with N , the length of the time-series. Therefore, without constraining the permutation, the CR of a trivial combination that outputs the Wiener process approaches zero as N increases. This leads to the identification of the Wiener process as the control objective, clearly an artifact. We refer the reader to Section 1.B in [46] for a proof and a detailed explanation.

Proof of Lemma 1. In what follows we prove that the resampling function

$$\zeta(\mathcal{E}_{\Omega^{(i)}}(\cdot)) = \frac{P_{z_{(T)}}(\cdot)}{P_{z_{(T)}^*}(\cdot)}, \quad (11)$$

defined in 1 yields $\mathcal{D} = 0$ in (8). Recall that,

$$e_k = \mathcal{E}_{\Omega^{(i)}}(z_{k-T:k}), \tilde{e}_k = \mathcal{E}_{\Omega^{(i)}}(z_{k-T:k-1}, \tilde{z}_k), \quad \tilde{z}_k \sim P_{\tilde{z}}^{\zeta}(\tilde{z}_k | z_{k-T:k-1}) = P_z(\tilde{z}_k)\zeta(\tilde{e}_k). \quad (12)$$

Based on (12) define

$$P_{\tilde{z}_{(T)}}(z_{k-T:k-1}, \tilde{z}_k) = P_{z_{(T-1)}}(z_{k-T:k-1})P_{\tilde{z}}^{\zeta}(\tilde{z}_k | z_{k-T:k-1}),$$

the probability that the shuffle player will concatenate \tilde{z}_k to the observed sequence $z_{k-T:k-1}$. In (8)-(12), the time-series e_k inherits its distribution via the 1D projection from the distribution $P_{z_{(T)}}$ and the time-series \tilde{e}_k inherits its distribution from $P_{\tilde{z}_{(T)}}$. We will show that for the choice of the resampling function in (11), the two 1D projections coincide. Note that

$$\begin{aligned} P_{\tilde{z}_{(T)}}(z_{k-T:k-1}, \tilde{z}_k) &= P_{z_{(T-1)}}(z_{k-T:k-1})P_{\tilde{z}}^{\zeta}(\tilde{z}_k | z_{k-T:k-1}) \\ &\stackrel{(i)}{=} P_{z_{(T-1)}}(z_{k-T:k-1})P_z(\tilde{z}_k)\zeta(\tilde{e}_k) \\ &\stackrel{(ii)}{=} P_{z_{(T)}^*}(z_{k-T:k-1}, \tilde{z}_k)\zeta(\tilde{e}_k), \end{aligned} \quad (13)$$

where (i) is by the definition of $P_{\tilde{z}}^{\zeta}$ in (12) and (ii) by the definition of $P_{z_{(T)}^*}$ in Lemma 1. The derivation in (13) implies that

$$P_{\tilde{z}_{(T)}}(\tilde{e}_k = \mathcal{E}_{\Omega^{(i)}}(z_{k-T:k-1}, \tilde{z}_k)) = P_{z_{(T)}^*}(\tilde{e}_k)\zeta(\tilde{e}_k),$$

from which by substituting (11), $\zeta(\tilde{e}_k) = \frac{P_{z_{(T)}}(\tilde{e}_k)}{P_{z_{(T)}^*}(\tilde{e}_k)}$, it directly follows that $P_{\tilde{z}_{(T)}}(\tilde{e}_k) = P_{z_{(T)}}(\tilde{e}_k)$ \square .

A.1 Filter

The architecture of the filter within the block $\mathcal{E}_{\Omega}(\cdot)$ in Fig. 1 has many degrees-of-freedom and can be chosen by the user according to prior knowledge regarding the nature of the dynamic reference. To impose few constrains on the filter $F(\cdot)$ it can be implemented by a fully connected deep neural-network. Dealing with biological systems, we assume

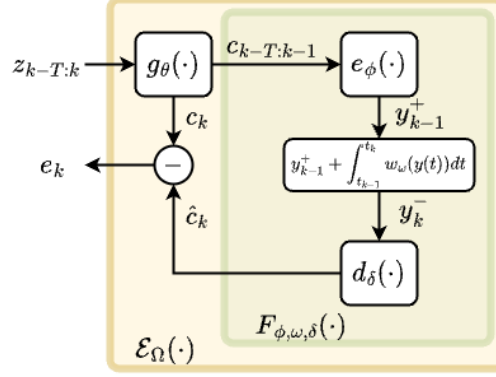


Figure 3: Architecture of the $\mathcal{E}_\Omega(\cdot)$ block (6) from Fig. 1. The filtering error e_k is the difference between the learned reference value c_k and its 1-step prediction \hat{c}_k . The filter $F_{\phi,\omega,\delta}(\cdot)$, (14), infers a latent-state y_{k-1}^+ and time-advances it to the state y_k^- from which the estimation \hat{c}_k is decoded.

that the dynamics of the reference can be modeled by a continuous-time, time-invariant latent model (see [33, 30] for details on integrating differential equations using neural-networks). The filter-block contains three parts: (i) An encoder $e_\phi : \mathbb{R}^T \rightarrow \mathbb{R}^{n_y}$, that given T consecutive values of the reference, $c_{k-T:k-1}$, infers a latent state $y_{k-1}^+ \in \mathbb{R}^{n_y}$ (with n_y a user-defined hyper-parameter), (ii) A drift function $w_\omega : \mathbb{R}^{n_y} \rightarrow \mathbb{R}^{n_y}$ describing the deterministic term in the dynamics of the latent state that serves to time-advance the latent state y_{k-1}^+ to y_k^- , (iii) An emission function $d_\delta : \mathbb{R}^{n_y} \rightarrow \mathbb{R}$ that decodes the 1-step predicted value \hat{c}_k from the latent state y_k^- . The following set of equations describe the filter-block (5) marked $\mathcal{E}_\Omega(\cdot)$ in 1,

$$y_{k-1}^+ = e_\phi(c_{k-T:k-1}), \quad y_k^- = y_{k-1}^+ + \int_{t_{k-1}}^{t_k} w_\omega(y(t))dt, \quad \hat{c}_k = d_\delta(y_k^-). \quad (14)$$

The block is depicted in Figure 3.

A.2 Performance assessment metric

IDRAS, an unsupervised learning algorithm, yields a time-series of a regulated quantity, c_k , and the corresponding time-series of 1-step predictions, \hat{c}_k . To assess whether a quantity that follows a dynamic reference value was indeed found, we calculate the normalized prediction error energy $\rho(c, \hat{c})$.

Definition A.1. The normalized error energy for two time-series $x_{1:N}$ and $\hat{x}_{1:N}$ is defined,

$$\rho(x_{1:N}, \hat{x}_{1:N}) = \frac{\sum_{k=1}^N (\bar{x}_k - \hat{\bar{x}}_k)^2}{\sum_{k=1}^N (\bar{x}_k)^2}, \quad \bar{x}_k = \frac{x_k - E[x]}{\sigma(x)},$$

where $E[x] = \frac{1}{N} \sum_{k=1}^N x_k$ and $\sigma^2(x) = \frac{1}{N} \sum_{k=1}^N (x_k - E[x])^2$, and similar for \hat{x} .

The normalized error energy $\rho(c, \hat{c})$ allows for an immediate assessment of the significance of the identified combination. A low ρ value ($\rho < 0.5$) indicates that a combination that follows a dynamic reference was found. We remind the reader that the shuffle player in IDRAS guarantees that the algorithm won't converge to trivial combinations such as $c_k \equiv \text{const}$.

B Appendix B: Validation

In all examples, to assess the performance, we calculate the normalized error energy between the known control objective c_k^* and the output of IDRAS c_k and consider $\rho(c, c^*) < 0.1$ as an excellent empirical result. We start by introducing the architecture of the networks used throughout the examples. Code reproducing all examples is available at <https://github.com/RonTeichner/IRAS>.

B.1 Architecture

All the parameterized functions used in Section 3 are feed-forward artificial neural networks with an input layer of 32 neurons, a hidden layer of 16 neurons and an output layer of appropriate dimension. The activation function of

all neurons is Leaky-ReLU except for the output neuron whose activation function is the Sigmoid function. The hyper-parameters are $T = 10$ and $n_y = 2$. We train the networks as described in Algorithm 1 for a pre-defined fixed number of epochs (150) using the Stochastic-Gradient-descent optimizer with a momentum value of 0.9 and a learning rate of 0.01. See [52] for a detailed explanation of feed-forward neural networks, activation functions and optimizers.

B.2 A kinetic model of interactions

We first validate IDRAS on simulated data generated from a kinetic model that describes regulatory interactions in the production of two proteins incorporating a feedback loop. In the considered model (inspired by [22]), the total amount of two proteins P and S , namely $P + S$, is controlled by M , the mRNA molecule that is produced through a process of constitutive transcription at a rate K and degraded with first order kinetics at a rate γ_M . The model is under constant perturbations to the protein expression rate K . These might be caused by ambient temperature dynamics that were found to have large passive effects on both mRNA synthesis and decay rates [43]. The model is described by the differential equations

$$\begin{aligned} K(t) &= K_0(1 + 0.5 \cos(2\pi \frac{1}{\tau_K} t + \phi_K)), \\ dM &= (K(t) - f(P + S) - \gamma_M M)dt, \\ dP &= (k_P M - \gamma_P P)dt + \eta_P dW_P, \\ dS &= (k_S M - \gamma_S S)dt + \eta_S dW_S, \end{aligned} \quad (15)$$

where the mRNA M and the two proteins P, S are linked in a feedback loop. Both P and S are positively affected by M , with their steady-state values proportional to it. The concentration M in turn, is negatively affected by the sum $P + S$, with the strength of this negative feedback given by the rate constant f . We note that (15) is a model of dimension 5, comprising 3 state variables and an oscillating input K that results from an underlying state-space model of dimension at least 2 (a 1D model cannot produce oscillations).

Small changes in S or P , modeled by the an increments of the Wiener processes dW_P and dW_S , induce swift and sharp changes in the transcription of M and maintain $P+S$ around a reference level

$$c^*(t) \triangleq (P + S)(t) = \frac{1}{f + \frac{\gamma_M \gamma_P \gamma_S}{k_P \gamma_S + k_S \gamma_P}} K(t), \quad (16)$$

which is reflected in a high negative correlation between S and P [43]. In (15) we identify the components of the general model (3): $u_e = [K, M]$, $x = [P, S]$. Our observations contain the levels of the two proteins P and S (an observation taken in labs by measuring fluorescence intensity of constitutively expressed proteins [11]), $z_k = [P(t = k/f_s), S(t = k/f_s)]$, where $k \in [1, \dots, N]$ and f_s is the sampling rate. Figure 4a illustrates the kinetic interactions model and SI Appendix section 1.2 lists the parameter values.

We ran both IRAS and IDRAS algorithms in search of the control objective. Figure 4b depicts the output (dashed red) of $g_\theta(P, S)$ trained by the IRAS algorithm. Due to the lack of a regulated constant combination, IRAS did not converge and scored $\rho_{\text{IRAS}}(c, c^*) = 1.82$ (see Definition A.1).

When running IDRAS, we have a simple measure to assess whether a quantity that follows a dynamic reference value was found, the measure of the normalized prediction error energy defined in Section A.2. A score of $\rho_{\text{IDRAS}}(c, \hat{c}) = 0.148$ was obtained, indicating that a combination that tightly follows a dynamic reference was found. The validation score is $\rho_{\text{IDRAS}}(c, c^*) = 0.012$, indicating that the combination $P + S$ was precisely found, despite its oscillating nature, as depicted in Figure 4c.

Following are the exact details of the parameters used.

The timescale of the feedback loop $\tau_{(P+S)} = 1/f = 0.0005$ ($F = 2000$) is much shorter than the timescale of environmental influenced oscillations in K , $\tau_K = 0.2$. We simulated a dataset of 100 observed systems sampled at rate $f_s = 1$ with $K_s = K_p = 150$, $\gamma_S = \gamma_P = 70$, $\gamma_M = 80$, $f = 2000$, $K_0 = 300$ and $dW_P, dW_S \sim \mathcal{N}(0, (0.5dt)^2)$. For each system $\phi_K \sim \mathcal{U}[0, \pi]$ and $M(0), P(0), S(0) \sim \mathcal{U}[0.02, 0.1]$.

B.3 Bacterial life cycle

The next example we consider is a realistic biological model of bacterial life cycle. We shall focus here on bacterial growth homeostasis, where growth and division proceed for many generations with significant variability and statistical

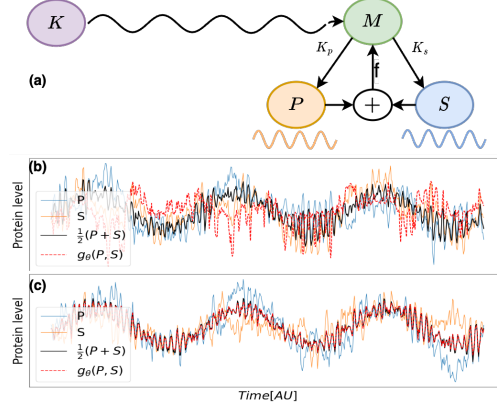


Figure 4: (a) Illustration of the closed-loop model (15). The mRNA molecule M induces the production of proteins S and P and receives a negative feedback of their sum. (b) The IRAS algorithm does not converge, due to the lack of a regulated constant combination. The combination $P + S$ (black) along with the output of IRAS (dashed red). (c) IDRAS algorithm captures the control objective and its oscillating trend.

stability. This is a problem with a long history and on which a large body of data is available. We apply our algorithm to simulation data, which mimic experimental measurements but where the regulation is known, and show that IDRAS detects the correct mode of regulation.

Most bacteria grow smoothly and divide abruptly, consistent with a threshold crossing by some division indicator at the single-cycle timescale [4, 24, 44]; the threshold itself can have dynamics over multiple cycles. Three types of division indicators corresponding to different regulation modes have been proposed: cell size (“sizer” control mechanism), added size (“adder” mechanism) and elapsed time (“timer”) [29, 27]. In a recent paper, Luo et al. [37] demonstrate that commonly used heuristic tools uncover the correct mode of regulation only under very restricted conditions. We reproduce these results and show that in contrast to the commonly used tools, IDRAS accurately identifies the cell-division mechanism.

The inter-generation threshold $u(t)$ is modeled by a stochastic Ornstein-Uhlenbeck process,

$$du = \frac{\mu_u - u}{\tau_u} dt + \sqrt{2 \frac{\sigma_u^2}{\tau_u}} dW \quad (17)$$

where the term dW is an increment of a Wiener process [37]. On the single-cycle scale, the k^{th} cell grows exponentially,

$$\begin{aligned} x(t) &= x_{k,b} e^{\alpha_k(t-t_{k,b})}, & t_{k,b} \leq t \leq t_{k,d} \\ x_{k,b} &= x(t_{k-1,d}) f_k, & t_{k,b} = t_{k-1,d} \\ t_{k,d} &= \operatorname{argmin}_t \{t \geq t_{k,b} \mid x_{k,b} e^{\alpha(t-t_{k,b})} = u(t)\}, \end{aligned} \quad (18)$$

where $x_{k,b} \triangleq x(t_{k,b})$ is the birth size, $x_{k,d} \triangleq x(t_{k,d})$ is the size at division, α_k is the exponential growth rate and f_k is the division fraction. Here we simulate a *sizer* mechanism such that the cell divides when its size crosses the threshold $u(t)$. Figure 5a depicts a simulated lineage over time - the cell-size, $x(t)$ (dashed-black), and the stochastic threshold $u(t)$ (blue). SI Appendix section 1.3 lists the parameter values.

In (17) and (18) we identify the components of the general model, (3), with $u_e = u$ and $x = x$. Our observations, derived from $x(t)$, contain the initial size, growth rate and the cycle duration, $z_k = [x_{k,b}, \alpha_k, T_k = t_{k,d} - t_{k,b}]$ where $\alpha_k = \frac{1}{T_k} \log(\frac{x_d}{x_b})$. We note that our choice of the feature vector z_k renders IDRAS’s task harder as now to correctly detect the sizer mechanism the network has to learn the combination $g() = x_b e^{\alpha T}$ and not just $g() = x_d$, in case x_d was an entry of z_k .

We ran both IRAS and IDRAS algorithms in search of the division mechanism. We chose a realistic parameter set for which heuristic identification methods based on data correlations fail to detect the control mechanism correctly [37]. Figure 5b depicts the output (dashed red) of $g_\theta(x_b, \alpha, T)$ trained by the IRAS algorithm vs the ground-truth output, the sizer division mechanism $x_b e^{\alpha T}$. Here IRAS, optimizing objective (1), outputs a combination that results from fusing together the division mechanism and the dynamic threshold $u(t)$ (with $\rho_{\text{IRAS}}(c, c^*) = 0.66$), as we now explain. In our model, the timescale of the threshold is significantly slower than the time scale of single-cells, $\tau_u \gg T$, such that the




Figure 5: (a) Bacteria life-cycles. The blue curve is the stochastic threshold process $u(t)$, (17), and the black curve is the cell-size over time, (18), such that the cell divides once it reached the dynamic threshold. (b) IRAS does not decouple the cell-size mechanism (black-curve) from the threshold trend $c(t)$ and yields a combination (dashed-red) that represents their mixture. (c) IDRAS captures the division-mechanism.

threshold is approximately constant along a single-cell cycle, $u(t) \approx u_k$ for $t_{k,b} \leq t \leq t_{k,d}$. The birth size, about half the size at division of the previous cell is thus $x_{k,b} \approx 0.5u_k$, and the size at division is (by definition) $x_{k,d} \approx u_k$. The single-cell cycle regulated combination is therefore,

$$x_{k,b}e^{\alpha_k T_k} - u_k \approx x_{k,b}e^{\alpha_k T_k} - 2x_{k,b} \approx 0, \quad (19)$$

a mixture of two terms. The first, $x_{k,b}e^{\alpha_k T_k}$, represents the sizer mechanism while the second, $-2x_{k,b}$, is an influence of the dynamic threshold process.

IDRAS, optimizing objective (2), decouples the two by separately learning the threshold dynamics and a combination which is regulated w.r.t. this dynamic reference value. Figure 5c depicts the precise identification of the sizer mechanism with a validation score $\rho_{\text{IDRAS}}(c, c^*) = 0.046$. The normalized prediction error energy is $\rho_{\text{IDRAS}}(c, \hat{c}) = 0.23$, fitting the expected value derived from the noise term in (17) up to 8% ($2\bar{T}/\tau_c$, where \bar{T} is the mean cell-cycle, about 25 minutes).

Following are the exact details of the parameters used.

The simulated dataset contains 100 lineages each consisting of 100 generations. In (18), $\mu_u = 1 [\mu\text{m}]$, $\sigma_u = 0.1 [\mu\text{m}]$ and $\tau_u = 200 [\text{min}]$; Initial conditions are $u(0) \sim N(\mu_u, \sigma_u^2)$ for each lineage. In Eq. 18, the growth rate is sampled from a Gamma-distribution, $\alpha_n \sim \Gamma(25, 9.4 \cdot 10^{-4})$ and the division fraction is distributed $f_n \sim \mathcal{N}(0.5, 0.05^2)$. The initial conditions for the first cell in each lineage are $x_{0,b} \sim \mathcal{N}(0.5, 0.05^2)$, $t_{0,b} = 0$.

B.4 Identifying Complex Physical Equations

To further challenge the IDRAS algorithm, we evaluate it on a broad range of physics problems, taken from the seminal Feynman Lectures on Physics [23], also used in the recently published ‘‘Feynman Symbolic Regression Database’’ (FSReD) [49]. We simulate physical equations incorporating oscillating terms and expect IDRAS to decouple the oscillating terms from the quantities they follow.

One such equation is equation II.6.15b, taken from Volume 2 in [23], which describes the electric field induced by a dipole. The dipole, with a dipole moment P_d (the product of the charges and their separation) induces, at distance r and angle ωt , an electric field whose transverse component is E_f ,

$$E_f = \frac{3}{4\pi\epsilon} \frac{P_d}{r^3} \cos(\omega t) \sin(\omega t).$$

We simulate 100 time-series of length 100 in which the dipole spins with angular velocity ω . For each sample, the observables (E_f , ϵ , P_d and r) are sampled uniformly in $[1, 5]$ (although not physically valid, it is the way the FSReD dataset is synthesized [49]). We expect IDRAS to identify the combination

$$g_{\theta^*}(E_f, \epsilon, P_d, r) = E_f \frac{4\pi\epsilon r^3}{3P_d} = \cos(\omega t) \sin(\omega t) = c^*(t).$$

To quantify the performance of IDRAS, we compute $\rho(c, c^*)$ for each of three chosen examples from the fields of electric-field forces (I.12.11), non-linear responses (I.50.26), and electric dipoles II.6.15b. As shown in table 1, the agreement is excellent.

Table 1: IDRAS captures physical relations

	Equation	$\rho(c, c^*)$
I.12.11	$F = q * (E_f + Bv \sin(\omega t))$	0.007
I.50.26	$x = x_1(\cos(\omega t) + \alpha \cos^2(\omega t))$	0.022
II.6.15b	$E_f = \frac{3}{4\pi\epsilon} \frac{P_d}{r^3} \cos(\omega t) \sin(\omega t)$	0.059

C Appendix C: Related work

Some dynamical systems obey conservation laws, where a combination of variables is constant along time trajectories of the system [51], as in (1). Identifying conservation laws from observed data is an active research field, [50, 40, 26, 14, 9, 47], and it is important to note that, even given the differential equations of the system, identifying a conservation law analytically - or even proving its existence - is a difficult mathematical problem [1, 36, 2, 18, 20, 19, 21]. Methods for deriving the differential equations of a system [8] leave the question of identifying the conservation law unsolved. This problem is a sub-space of the one considered here, since the concept of conservation laws usually refers to well-defined dynamical systems with fixed parameter. Identifying biological regulation includes also the more general case where parameters are varying in time but still some combination is regulated around a set-point in the face of these perturbations. Optimizing for conservation alone can lead to trivial quantities, such as predicting a constant $g(z) = c$ independent of z . In a recent paper, [3] refer to a non-trivial $g(\cdot)$ by the term *useful conservation law*. To obtain a non-trivial solution, [46] define a measure of invariance, the "Coefficient of Regulation" (CR) and an optimization algorithm that obtains meaningful invariants in (1). Inferring meaningful dynamic regulatory processes, (2), renders the task more difficult. Here we generalize IRAS to allow inferring meaningful dynamic regulatory process.