

*The Third International TRANSIT Workshop
on Cross-disciplinary Research*

Dynamics on and of Networks



25–26 March 2020

Ron Cooke Hub, University of York, UK

Programme and Abstracts

Wed 25 March 2020

09:30	<i>registration / coffee</i>
10:00	keynote 1 -- Jean-Louis Giavitto Man-Machine Synchronization
11:00	<i>coffee</i>
11:30	Daniel Franks Going Viral on the Social Network: Disease spread in killer whale social networks
12:15	Daphne Ezer and Mahiar Mahjoub Rapid gene expression dynamics of photoperiod genes are governed by the early morning gene regulatory network
13:00	<i>lunch</i>
14:00	Alison McKay Engineering supply chains
14:45	Penelope Faulkner Rainford, Yasmin Merali and Susan Stepney Simulating situated agents in a complex environment modelled as a dynamic network
15:30	Jonathan Ward State-space reduction of dynamics on networks using symmetry: a study of real-world networks
16:15	<i>tea</i>
16:45	keynote 2 -- Anje Neutel Soil webs, sea mats and scaling the turtle: Making sense of the multitude of feedbacks in flow networks
17:45	panel session
18:45	<i>close</i>

Thur 26 March 2020

08:45	<i>registration / coffee</i>
09:00	keynote 3 -- James Marshall 'Optimal' decisions on and of graphs
10:00	Fakher Raza and Dimitar Kazakov Dynamics on Pipeline Network
10:45	<i>coffee</i>
11:30	Marcin Kupilas Clumps, Filamentary Networks and Star Formation
12:15	Antoine Spicher From MGS to Global Transformations
13:00	<i>lunch</i>
14:00	Adam Katona, Danial Franks and James Walker Learning to predict the evolvability of artificial gene regulatory networks from network dynamics
14:45	Richard Watson What can adaptive networks learn?
15:30	panel session plus discussion of topic for TWCR 2021
16:30	<i>tea / close</i>

Man-Machine Synchronization

Jean-Louis Giavitto

IRCAM, France

Ensemble music is the result of a choreography of events and expectations in time. And the capacity for real-time synchronization and coordination is a common ability among trained musicians performing together a music score. But is a computer able to understand the dynamic of playing together?

Parts of the answer are provided by Antescofo, a system that attempts to enable computer-human musical interactions in the context of mixed music, i.e. when humans and computers are performing together. Antescofo provides an abstract programmer's model for an artificial musician in an ensemble with musical real-time constraints. Antescofo ability to interact musically relies on a dedicated heterogenous model of time, that encompasses event-driven and time-driven specifications, absolute and relative time, and subjective and social time. This model of time makes possible the common understanding of the coordination required to take part in musical interactions, a first step in the development of a musical companionship and a practical experience in the possible sharing of time between man and computers.

Going Viral on the Social Network: Disease spread in killer whale social networks

Dan Franks

University of York

Infectious disease outbreaks can hamper the recovery of endangered population, and often create a serious extinction risk. The southern resident killer whale (SRKW) community is a small, endangered population of highly gregarious cetaceans that faces numerous threats from toxins, noise pollution, and declines in prey abundance, with pathogens recently identified as a putative fourth threat. Here, we use a stochastic individual-based model, informed by 5 years of detailed observations of close dynamic social network contacts, to predict the severity of disease outbreaks through the contact network, with a focus on cetacean morbillivirus (CeMV). We find that the SRKW contact network is highly vulnerable to disease outbreaks, with CeMV likely to produce outbreaks in which over half the population becomes infected.

Rapid gene expression dynamics of photoperiod genes are governed by the early morning gene regulatory network

Daphne Ezer, Mahair Mahjoub

University of York, University of Sydney

In a gene regulatory network, each node represents a gene and each directed edge represents a regulatory relationship between a pair of genes. Each gene is associated with a scalar—gene expression—whose value is determined by the gene expression of the genes that have edges pointing towards it, usually subject to a time delay. Experiments measure gene expression over time (i.e. the dynamics of values on the network). However, the structure of the network is often unknown or has many missing edges. Furthermore, many parameters associated with the dynamical system underlying the network (including the time delay term) are unknown.

The early morning is a critical time period for plants, because this is when the internal clock is reset by light and temperature signals. However, early morning genes have very rapid bursts of expression that are difficult to capture experimentally. We collected gene expression data at an extremely high temporal resolution and used this to infer the time-delayed gene regulatory network that controls waves of expression that occur in the early morning. Approximately a third of the edges in this network have been confirmed with DNA binding experiments. We searched for enriched associations between functional annotations within the network topology, developing a new R package called PAFway for performing this kind of statistical analysis. This analysis showed that photoperiodism—the process of measuring and responding to day length—is enriched downstream of both circadian rhythm and light sensing processes. A small sub-network, driven by red and blue light sensors, is sufficient for explaining the pattern of expression dynamics of many key photoperiod genes. We were able to verify the structure of this sub-network by experimentally removing certain edges in the network and monitoring the changes in gene expression dynamics over time. This project provides a case study of how to utilise large experimental data sets to infer the structure of networks that govern dynamical systems.

Engineering supply chains

Alison McKay, University of Leeds

Engineering supply chains are the networks of organisations through which products are delivered to markets and, increasingly, designed and supported through their lives. With reference to TRANSIT terminology, the state of, or dynamics *on*, a given supply chain is a critical factor in the effectiveness and efficiency of the engineering process that it underpins. When supply chains malfunction, the consequences in terms of reputational damage, and cost and time overruns can be vast. For example, the UK's Crossrail project had cost overruns of circa £600 million¹ and lessons learnt reports include recommendations for the management of quality in [manufacturing and construction] supply chains². Root causes of malfunctions can often be attributed to [product] system-level design decisions, made very early in the product development process, that fix or heavily constrain: (a) key characteristics and requirements of the major sub-systems, and (b) for each of these sub-systems, whether it is to be designed and made in-house or by suppliers. With reference to TRANSIT terminology, these factors relate to the structure, or dynamics *of*, the supply chain.

To allow engineering designers, who typically have limited supply chain knowledge, to visualise implications of early design decisions on supply chain performance, we have built an Excel-based proof-of-concept software prototype³ interface for an engineering simulation package (WITNESS⁴). The interface^{5,6} takes as input an indented parts list and associated make/buy scenario, and uses them to generate a simulation model that creates risk profiles of design supply chain operations of alternative [product] system architectures and make/buy scenarios. With reference to TRANSIT terminology, we use a product architecture and associated make/buy scenario to generate a simulation model of the meta-dynamics of the supply network; and then use simulation software to explore risks in the operation of the resulting design and make supply network. A demonstration of this prototype will be included in the presentation associated with this abstract.

A distinguishing feature of engineering supply networks lies in the close relationship between the structure of the supply network and that of the product lifecycle it supports. Current product development practice tends to focus on product performance and regard the supply network structure, and so meta-dynamics, as a consequence of a design rather than something that can itself be controlled as part of the product development process. However, as recognition of societal challenges such as climate change grows, so does the importance of considering wider societal impacts of engineering processes and associated supply chains. This, in turn, is creating a need for new forms of engineering design tool that support consideration of supply chain meta-dynamics, including climate and environmental impacts, as part of the product development process. The main reason for putting forward this abstract is to create opportunities to learn from other disciplines and so inform the early design of this new generation of engineering design system.

¹ <https://www.railway-technology.com/news/uks-crossrail-project-costs-overshoots-nearly-600m/>

² <https://learninglegacy.crossrail.co.uk/documents/supply-chain-quality-requirements/>

³ An interactive demonstration of the proof of concept software will be available on request.

⁴ <https://www.lanner.com/en-gb/technology/witness-simulation-software.html>

⁵ McKay, A.; de Pennington, A.; Chittenden, R.J.; Hazlehurst, T.; Baker, R.B.; Waller, T. (2019) "Visualising the impact of early design decisions on engineering supply chains". 22nd International Conference on Engineering Design, ICED'19. Delft, Netherlands. 5-8 August 2019. www.iced19.org

⁶ McKay, A.; de Pennington, A.; Chittenden, R.J.; Hazlehurst, T.; Baker, R.B.; Waller, T. (2019, in press) "The systematic derivation & evaluation of extended enterprise processes from system architectures". Annual Systems Engineering Conference. 19–20 November 2019. The Royal Armouries Museum, Armouries Drive, Leeds. <https://incoseonline.org.uk/ASEC2019/>

Simulating situated agents in a complex environment modelled as a dynamic network

Penny Faulkner Rainford, Yasmin Merali, Susan Stepney

January 16, 2020

We are interested in modelling systems of complex agents situated in a complex environment. For example, consider modelling the elderly who have declining mobility. The people can be modelled as agents making decisions; we may use agents' individual parameters to dictate how many of them use different resources, and use this to model each agent's mobility decline over time. The environment comprises hospitals, GPs and other resources, which can be modelled in terms of the flow of patients around these resources. Each of these component models can answer some questions, but combined they can include mutual feedback. Patients change their minds and behaviours in response to the environment; the environment is modified by and responds to the patients. We wish to build a model that can capture this feedback using agents situated on a dynamic network.

Our goal is an agent-based system that supports individuality and decision making capability of agents situated in an environment. We use a network to represent the environment, with agents located at nodes. We enable feedback by allowing the network to limit agent choice, and to modify agents as they move, while also allowing agent actions to modify the network.

There are different kinds of dynamics possible in this model. There are individual agent-node interactions: a node might take a payment from an agent in exchange for a resource located there. An agent's behaviour might change over time because of the routes it takes through the network and how this changes its state. The network topology itself might change in reaction to the broader pattern of agent motion: it might add new resources or capabilities to improve agent flow through the system, or remove under-utilised edges or nodes; it might change its structure based on outside influences as the system reacts to an imposed change.

Our goal is to develop an approach that allows us to build models and simulations of many diverse systems following this structure. To enable this, we are developing a 'socio-physical systems' meta-model, and an accompanying software platform for implementation. We intend this modelling to help reveal emergent structures, unintended consequences of change and transition, created in dynamic systems over time as they find equilibrium, or as they engage in extended transient behaviour.

State-space reduction of dynamics on networks using symmetry: a study of real-world networks

Jonathan Ward

University of Leeds

Network symmetries, or graph automorphisms, occur when a permutation of node labels leaves the network's edge set unchanged. Surprisingly, such symmetries occur in real-world networks, but are limited to a few relatively simple types known as "basic symmetric motifs" [MacArthur, Sanchez-Garcia, Anderson, *Discrete Applied Mathematics* 156, 2008]. Network symmetries can be used to reduce the size of the state-space of Markov chain dynamics on networks, such as epidemic or voter models, via a process called "lumping" [Simon, Taylor, Kiss, *Journal of Mathematical Biology* 62, 2011]. We will show how the symmetries in real-world networks can be used to lump the state-space of a broad class of network dynamics and provide a formula for the size of the state-space reduction.

We discuss the feasibility of using graph automorphism lumping on real-world networks by looking at the symmetries present in over 1000 small real-world networks collected from the website networkrepository.com. We find that more than 80% of the networks analysed have non-trivial symmetry and that in 94% of these cases, the symmetry is entirely due to basic symmetric motifs. Crucially, we find 62 networks with significant lumping, i.e. cases where the size of the full state-space is not computationally feasible to analyse, but the size of the lumped state-space is feasible. This computational feasibility naturally depends on the particular computer being used, but our study suggests that regardless of this, there will always be networks with significant lumping. We also highlight the types of symmetry that give rise to significant lumping and the kinds of complex, non-BSM, symmetry that occur.

Soil webs, sea mats and scaling the turtle: Making sense of the multitude of feedbacks in flow networks

Anje-Margriet Neutel

British Antarctic Survey

It is easy to get overwhelmed by the complexity of the dynamic systems we are surrounded by and are part of. So many interactions, generating many more feedback loops, making it impossible to predict what will happen next.

Theoretical ecology has a long tradition of studying interaction networks. It has identified simple mathematical relations between their structure and stability. But when real data are used to parameterise the interaction strengths in classic theoretical models, these relations no longer hold. For example, the complexity of a network (defined as the product of system size and its connectance, the proportion of possible interactions that are realised) is traditionally linked with instability, but observations show that it has no relation with stability at all. The theoretical relation with instability is an artefact of the assumption of random interaction strengths. To explain the stability of naturally organised systems, we need to quantify the interactions, go beyond the pairs, and quantify longer feedback loops.

With examples from soil food webs, above-and-belowground Antarctic ecosystems and bryozoan competition networks on the sea floor, I hope to show that with this quantitative loop analysis we can take a step towards a general stability theory grounded in observation. I will address questions like:

- How can we compare systems encompassing many different time scales?
- How far do we need to go beyond pairwise interactions?
- Why do natural systems become more complex over time, or How can we marry stability with development?

‘Optimal’ decisions on and of graphs

James Marshall

University of Sheffield

Optimality theory for decision-making exists at two levels, optimal decisions by individuals, and optimal decisions by groups. Although these levels can be closely related, they are not often studied in conjunction. In this talk I will show how dynamical models of networks can implement decisions as if they were optimal agents, then show how groups of individually-optimal agents can reach consensus decisions (or not) when distributed over networks.

Dynamics on Pipeline Network

Fakher Raza, Dimitar Kazakov

University of York

Pipelines are everywhere. As we walk down the street it is probable that a high-speed pressure wave is travelling underground in a water pipeline. Sometimes these pressure waves can turn into water bursts causing damage to life and the environment.

The pipeline network operation is highly dynamic since the speed at which pressure waves can travel along a pipeline system typically vary between 400 m/s and 1300 m/s. The transient behaviour of the pipeline network is inherently associated with the change in operation or occurrence of an anomaly over time.

Pump trip, valve closure or station start-up is an example of an operational transient where a pressure wave travels back and forth from a fixed location and that eventually dissipates through the network until a steady state is achieved.

A pipeline burst is an example of an anomaly, resulting from human error or equipment failure. The pressure wave, in this case, initiates from the burst location and travels towards all hydraulically connected pipes pressurising and depressurising all pipes and fittings.

Pipeline transients are usually captured by the sensors installed on the network; however, it is not feasible to monitor an entire network with sensors. A pipeline simulator is required to configure multiple virtual sensors which can simulate an anomaly and calculate the intensity of the pressure wave throughout the network. The output from the simulator helps monitor the operation and various alerts can be generated. Pipeline simulators can highlight wave travel patterns at specified locations which can be stored and used for the pattern learning. This simulated data can mimic real-life trends that can be used to support anomaly detection through pattern recognition.

Pipeline network dynamics generates identifiable patterns for intelligent pattern recognition techniques and the ML models are the most suitable candidates for anomaly detection techniques.

A machine learning model will need transient data input for all possible burst locations or at least a limited number of locations scattered evenly across the zone. VariSim™ is a powerful pipeline simulator capable of modelling network hydraulic behaviour for generating the required data sets. If the burst scenarios within the hydraulic model are simulated at several different locations, the pattern data can be saved and used for training. After an accurate ML model is established, the prediction of burst locations can then be triggered every time a burst alert is raised by the sensors.

Clumps, Filamentary Networks and Star Formation

Marcin Kupilas

University of Leeds

Stars form from cold and dense cores, which are themselves found in substructures of giant molecular clouds (GMCs). Recent observations by the Herschel Space Observatory have revealed the presence of a complex filamentary network in every GMC and it is now understood that the formation of this filamentary structure is a vital step in the star formation process where cores are found at the intersections of filaments and in regions where the filament has undergone fragmentation. Numerical simulations have not yet managed to fully elucidate the physics responsible for GMC formation, as it is difficult to create initial conditions that are realistic and not artificial in nature, with the ability to reproduce observed macroscopic properties of the GMC and its dense substructures. It is important that the initial conditions are carefully considered, as the physics that forms GMCs is then partly responsible for the effect that subsequently formed stars have on their parent cloud. This work will present preliminary results of a numerical study of a shock interacting with an inhomogeneous cloud formed due to the effects of thermal instability and gravity. The implications of this work will be discussed.

From MGS to Global Transformations

Antoine Spicher, Univ Paris Est Creteil, LACL, 94000, Creteil, France

Context

The MGS project has been created to deal with the notion of *Dynamical System with a Dynamical Structure* [2], (DS)² for short. By dynamical structure, we mean that the set of state variables is changing over time together with their relative relationships. Such systems are ubiquitous. For example, think of an amorphous computation distributed on the dynamical network of mobile devices, or of the morphogenesis of a living organism. Moreover, the evolution of the structure is often intimately coupled with the evolution of a patterning on the structure.

As glimpsed by A. Turing in his work about *The Chemical Basis of Morphogenesis*, computer science can help the specification and the simulation of models for the (DS)² class of systems. MGS proposes to pursue this program for systems that can be described as a structured collection of many interacting elementary entities. In such a case, the intended behavior is described through a set of *local transformation rules* specifying how certain subgroups of entities evolve in parallel.

This approach has been fruitful and numerous examples of (DS)² have been efficiently expressed [4]. In the same work, it has been shown that the interaction-based paradigm of computation provided by MGS captures a large majority of unconventional models: data are organized following some structure depending on the computing model (from the crystalline space of cellular automata to the total lack of structure of chemical models), the dynamics being defined through a set of interaction laws specifying how a sub-structure of the data behaves locally.

Towards Global Transformations

In MGS, an elementary step of computation consists in (1) identifying evolving sub-parts, (2) locally updating these parts independently, and (3) gathering together the updated parts to get back a global object. The genericity of MGS comes from a powerful language of rules allowing the expression of steps (1) and (2) independently of the underlying data structure. However, step (3) is not specified within the rules but is left to the choice of the data structure coming with an *ad hoc* reconstruction procedure.

A direct consequence is that the sole knowledge of the rules is not enough to fully understand the dynamics. Particularly the behavior of overlapping evolving subgroups is left unspecified. Thus a *mutual exclusion* property, the so-called *maximal-parallel rule application strategy*, is considered to forbid parallel applications of intersecting rules.

This restriction of MGS – shared for the same reason by all the captured unconventional models – makes a large class of systems left out. A simple such example is mesh-refinement whose local rules consist of the replacement of coarse parts of a mesh by finer ones. In any step, all rule applications overlap with some others, and all of them need to be applied altogether to make any sense. The rules are designed to work together from the very beginning.

Global Transformations (GT) [3] is an original rule-

based model of computation. It puts the emphasis on correctly handling the fact that all rules must be applied whenever possible so that they imperatively need to be compatible on overlaps. This is realized with an efficient formalization of what we call the *mutual agreement* between the rules. Formally, it is captured by a generalization of the concept of monotony: when there is an inclusion between the left members of two rules, there must also be an inclusion between the associated right members, in a global coherent way. Additional rules are considered to specify the fate of overlaps (*e.g.*, the common edge of two neighbor triangles for mesh-refinement), and thus *in which sense* overlapping rules mutually agree.

The aforementioned monotony is to be considered with respect to the order induced by the inclusion. In this context, the implementation of a GT is only a matter of extension of the partial function induced by the rules to a total monotonic function working on all possible entries ordered by inclusion. We advocate that the inclusion order formally captures the notion of *locality* which was so far only informally addressed. Moreover while the property of continuity was only evoked for MGS programs, GTs, as monotonic functions, are by definition *continuous*.

First Results and Future Works

To our knowledge, three models of computation achieve the same properties as GTs: deterministic Lindenmayer systems, cellular automata, and causal graph dynamics. It has been shown that the former is captured by GTs [1], and we are currently working on clarifying the relationship between GTs and the latter. In fact, GTs – like MGS – are a generic tool able to capture existing computing models and to propose original ones, simply by varying the nature of the objects to be transformed. We are currently developing an implementation of GTs to investigate these models. We particularly focus our developments on modeling morphogenesis, a paradigmatic example of (DS)².

GTs are mathematically grounded so that theoretical studies are also intended in order to ease the reasoning on the specified models. A first goal is to further the topological aspects of GTs to get a better understanding of locality and causality, for example by lifting well-known results from the theory of cellular automata. Finally GTs are so far restricted to the case of deterministic models; we are currently working on their extension to non-determinism.

References

- [1] A. Fernandez, L. Maignan, and A. Spicher. Lindenmayer systems and global transformations. In *UCNC*, pages 65–78. Springer, 2019.
- [2] J.-L. Giavitto, O. Michel, J. Cohen, and A. Spicher. Computations in space and space in computations. In *Unconventional Programming Paradigms*, pages 137–152. Springer, 2004.
- [3] L. Maignan and A. Spicher. Global graph transformations. In *GCM@ ICGT*, pages 34–49, 2015.
- [4] A. Spicher and J.-L. Giavitto. Interaction-based programming in MGS. In *Advances in Unconventional Computing*, pages 305–342. Springer, 2017.

Learning to predict the evolvability of artificial gene regulatory networks from network dynamics

Adam Katona
ak1774@york.ac.uk

Daniel W. Franks
daniel.franks@york.ac.uk

James Alfred Walker
james.walker@york.ac.uk

A well designed indirect encoding can turn a difficult optimization problem to an easy one. But designing good encodings is a challenging problem. A process to automatically learn good performing indirect encodings is desirable. In biological evolution, this process is called the evolution of evolvability.

One way to define evolvability is, it is the ability to generate adaptive phenotypic variation [2]. The source of evolvability in nature is the ability of the developmental system to turn random genotypic variation into a distribution of phenotypic variation [3]. By choosing the right developmental program, evolution can increase its evolvability. Evolution have the possibility to do this, because there are patterns in developmental system, environment interaction that are deep enough to be able to generalize to future environments [3]. If an aspect of the developmental system was useful for generating adaptive variation in many past environments, it might be useful for future environments.

We propose a novel meta learning approach which learns and optimizes for patterns of the developmental system which predict evolvability. These patterns can be learned using supervised learning using a dataset of developmental systems and their evolvability. There are several challenging aspect of building such a system. There is a need to describe the developmental system with some kind of developmental features. There is also a need to measure evolvability in a way that allows for comparison of two measurements.

We attempt to apply this approach for learning patterns which predict evolvability of an artificial gene regulatory network, similar to the network used by Banzhaf et al.[1]. We use the trajectories of these networks in time or frequency domain as our features to describe the developmental system.

References

- [1] Wolfgang Banzhaf. “On the dynamics of an artificial regulatory network”. In: *European Conference on Artificial Life*. Springer. 2003, pp. 217–227.
- [2] Massimo Pigliucci. “Is evolvability evolvable?” In: *Nature Reviews Genetics* 9.1 (2008), p. 75.
- [3] Richard A Watson and Eörs Szathmáry. “How can evolution learn?” In: *Trends in ecology & evolution* 31.2 (2016), pp. 147–157.

What can adaptive networks learn?

Richard Watson
University of Southampton
(raw1@soton.ac.uk)

Brains and computers are not the only kinds of things that can learn. The ability to adapt with experience, as familiar in learning neural networks (natural and artificial), can also be exhibited by gene networks⁷, protein networks, ecological networks⁶, economic and financial networks, institutional and commercial infrastructure, and social networks^{11,15} – if they have suitable properties. These systems are adaptive networks: i.e. the behaviours of components are influenced by connections and, reflexively, connections are altered by behaviours of the components^{4,6,10}. When such systems are ‘pulsed’ – stretched by a stress or perturbation, then released and stretched again, repeatedly – the connections slowly ‘relax’ to accommodate the tensions created in each stretch^{10,12,13,7}. Over time the differential relaxation of connections causes the structure of the system to become organised to ‘mirror’ the structure of the stresses it has experienced¹⁴, with the consequence that the system more easily accommodates familiar stresses or environments in future (increased resilience), and anticipates efficient responses to new perturbations^{2,14,3,7,9,13}. In neural networks we call this learning, specifically, Hebbian learning (i.e. positive feedback on correlations) which changes connections according to past experience in a way that enables generalisation to novel instances (if the inductive bias is appropriate)^{4,5}. Such positive feedback on correlations does not require specific neural mechanisms (natural or artificial) and will occur in any complex network with imperfectly elastic interactions (i.e. with structural connections that are slightly plastic and thus deform slowly under stress)^{4,5}. Recognising this expands our understanding of intelligence and prompts exciting new questions. e.g., What can a financial market, a smart energy network or a social network learn? When does the knowledge in such a system generalise well to novel conditions? and What conditions can be applied to enhance its learning capabilities? This talk will discuss how various aspects of machine learning theory can be employed to characterise these behaviours and address these questions.

1. Rago, A., Kouvaris, K., Uller, T., & Watson, R. A. (2019). How adaptive plasticity evolves when selected against. to appear *PLoS Comp. Bio.* (BioRxiv, 339622.)
2. Kouvaris, K., Clune, J., Kounios, L., Brede, M., & Watson, R. A. (2017). How evolution learns to generalise: Using the principles of learning theory *PLoS Comp. Bio.*, 13(4), e1005358.
3. Watson, R. A., Mills, R., Buckley, C. L., Kouvaris, K., Jackson, A., Powers, S. T., ... & Power, D. (2016). Evolutionary connectionism: algorithmic principles underlying the evolution of biological organisation in evo-devo, evo-eco and evolutionary transitions. *Evolutionary biology*, 43(4), 553-581.
4. Watson, R.A. & Szathmáry, E. (2016) “How Can Evolution Learn”, *Trends in Ecology & Evolution*
5. Power, D. A., Watson, R. A., Szathmáry, E., ... Czapp, B. (2015). What can ecosystems learn? Expanding evolutionary ecology with learning theory. *Biology Direct*, 10(1), 69.
6. Watson, R.A., Wagner, G. P., Pavlicev, M., Weinreich, D. M., & Mills, R. (2014). The evolution of phenotypic correlations and “developmental memory”. *Evolution*. 68(4) 1124-1138.
7. Watson, R.A., Mills, R.M., Buckley, C.L. (2011) Global Adaptation in Networks of Selfish Components. *Artificial Life*. 17(3):147-66.
8. Davies, A. P., Watson, R.A., Mills, R., Buckley, C. L., & Noble, J. (2011). How Individual Habituation of Agent Interactions Improves Global Utility. *Artificial Life*, 17(3), 167-181.
9. Watson, R.A., Mills, R.M., Buckley, C.L. (2011) Transformations in the scale of behavior and the global optimization of constraints in adaptive networks, *Adaptive Behaviour* 19(4): 227-249.
10. Watson, R.A., Buckley, C.L., Mills, R.M. (2010) Optimisation in ‘Self-modelling’ Complex Adaptive Systems. *Complexity*. 16(5):17-26.
11. Kounios, L., Clune, J., Kouvaris, K., Wagner, G. P., Pavlicev, M., Weinreich, D. M., & Watson, R. A. (2016). Resolving the paradox of evolvability with learning theory: How evolution learns to improve evolvability on rugged fitness landscapes. *arXiv preprint arXiv:1612.05955*.
12. McCabe, C., Watson, R. A., Prichard, J., & Hall, W. (2011, June). The web as an adaptive network: coevolution of web behavior and web structure. In *Procs of the 3rd Int. Web Sci. Conf.* p22-28. ACM.
13. Caldwell, J. R., Watson, R. A., Thies, C., & Knowles, J. D. (2018). Deep Optimisation: Solving Combinatorial Optimisation Problems using Deep Neural Networks. *arXiv preprint arXiv:1811.00784*.