

The Evolutionary Origins of Hierarchy

Henok Mengistu, Joost Huizinga, Jean-Baptiste Mouret, Jeff Clune

▶ To cite this version:

Henok Mengistu, Joost Huizinga, Jean-Baptiste Mouret, Jeff Clune. The Evolutionary Origins of Hierarchy. PLoS Computational Biology, 2016, 12 (6), pp.23. 10.1371/journal.pcbi.1004829.s021 . hal-01330681

HAL Id: hal-01330681 https://inria.hal.science/hal-01330681

Submitted on 12 Jun2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



OPEN ACCESS

Citation: Mengistu H, Huizinga J, Mouret J-B, Clune J (2016) The Evolutionary Origins of Hierarchy. PLoS Comput Biol 12(6): e1004829. doi:10.1371/journal. pcbi.1004829

Editor: Olaf Sporns, Indiana University, UNITED STATES

Received: December 29, 2015

Accepted: February 25, 2016

Published: June 9, 2016

Copyright: © 2016 Mengistu et al. This is an open access article distributed under the terms of the <u>Creative Commons Attribution License</u>, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Data can be found at: https://datadryad.org/resource/doi:10.5061/dryad. 60hp6.

Funding: JC was supported by an NSF CAREER award (CAREER: 1453549). JBM was supported by the ANR Creadapt project (ANR-12-JS03-0009) and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement number 637972). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. **RESEARCH ARTICLE**

The Evolutionary Origins of Hierarchy

Henok Mengistu¹, Joost Huizinga¹, Jean-Baptiste Mouret^{2,3,4}, Jeff Clune¹*

1 Evolving AI Lab, Department of Computer Science, University of Wyoming, Laramie, Wyoming, United States of America, 2 LARSEN, Inria, Villers-Iès-Nancy, France, 3 UMR 7503 (LORIA), CNRS, Vandœuvre-Iès-Nancy, France, 4 LORIA (UMR 7503), Université de Lorraine, Vandœuvre-Iès-Nancy, France

* jeffclune@uwyo.edu

Abstract

Hierarchical organization-the recursive composition of sub-modules-is ubiquitous in biological networks, including neural, metabolic, ecological, and genetic regulatory networks, and in human-made systems, such as large organizations and the Internet. To date, most research on hierarchy in networks has been limited to quantifying this property. However, an open, important question in evolutionary biology is why hierarchical organization evolves in the first place. It has recently been shown that modularity evolves because of the presence of a cost for network connections. Here we investigate whether such connection costs also tend to cause a hierarchical organization of such modules. In computational simulations, we find that networks without a connection cost do not evolve to be hierarchical, even when the task has a hierarchical structure. However, with a connection cost, networks evolve to be both modular and hierarchical, and these networks exhibit higher overall performance and evolvability (i.e. faster adaptation to new environments). Additional analyses confirm that hierarchy independently improves adaptability after controlling for modularity. Overall, our results suggest that the same force-the cost of connections-promotes the evolution of both hierarchy and modularity, and that these properties are important drivers of network performance and adaptability. In addition to shedding light on the emergence of hierarchy across the many domains in which it appears, these findings will also accelerate future research into evolving more complex, intelligent computational brains in the fields of artificial intelligence and robotics.

Author Summary

Hierarchy is a ubiquitous organizing principle in biology, and a key reason evolution produces complex, evolvable organisms, yet its origins are poorly understood. Here we demonstrate for the first time that hierarchy evolves as a result of the costs of network connections. We confirm a previous finding that connection costs drive the evolution of modularity, and show that they also cause the evolution of hierarchy. We further confirm that hierarchy promotes evolvability in addition to evolvability caused by modularity. Because many biological and human-made phenomena can be represented as networks, and because hierarchy is a critical network property, this finding is immediately relevant **Competing Interests:** The authors have declared that no competing interests exist.

to a wide array of fields, from biology, sociology, and medical research to harnessing evolution for engineering.

Introduction

Hierarchy is an important organizational property in many biological and man-made systems, ranging from neural [1, 2], ecological [3], metabolic [4], and genetic regulatory networks [5], to the organization of companies [6], cities [7], societies [8], and the Internet [9, 10]. There are many types of hierarchy [11-13], but the one most relevant for biological networks [14, 15], especially neural networks [1, 2, 16], refers to a recursive organization of modules [10, 13]. Modules are defined as highly connected clusters of entities that are only sparsely connected to entities in other clusters [17, 19, 20]. Such hierarchy has long been recognized as a ubiquitous and beneficial design principle of both natural and man-made systems [14]. For example, in complex biological systems, the hierarchical composition of modules is thought to confer greater robustness and adaptability [1, 2, 16, 21], whereas in engineered designs, a hierarchical organization of simple structures accelerates the design, production, and redesign of artifacts [19, 22, 23].

While most studies of hierarchy focus on producing methods to quantify it [4, 9, 11, 14, 24 -28], a few have instead examined why hierarchy emerges in various systems. In some domains, the emergence of hierarchy is well understood; e.g., in complex systems, such as social networks, ecosystems, and road networks, the emergence of hierarchy can be explained by resource constraints or by local decisions and interactions [3, 15, 29-31]. But, in biological systems, where the evolution of hierarchy is shaped by natural selection, why hierarchy evolves, and whether its evolution is due to direct or indirect selection, is an open and interesting question [3, 15, 32]. Non-adaptive theories state that the hierarchy in some, but not all, types of biological networks may emerge as a by-product of random processes [29]. Most adaptive explanations claim that hierarchy is directly selected for because it confers evolvability [33, 34], which is the ability of populations to quickly adapt to novel environments [35]. Yet in computational experiments that simulate natural evolution, hierarchy rarely, if ever, evolves on its own [36-38], suggesting that alternate explanations are required to explain the evolutionary origins of hierarchy. Moreover, even if hierarchy, once present, is directly selected for because of the evolvability it confers, explanations are still required for how that hierarchy emerges in the first place.

In this paper we investigate one such hypothesis: the existence of costs for network connections creates indirect selection for the evolution of hierarchy. This hypothesis is based on two lines of reasoning. The first is that hierarchy requires a recursive composition of modules [10], and the second is that hierarchy includes sparsity. A recent study demonstrated that both modularity and sparsity evolve because of the presence of a cost for network connections [17]. Connection costs may therefore promote both modularity and sparsity, and thus may also promote the evolution of hierarchy.

It is realistic to incorporate connection costs into biological network models because it is known that there are costs to create connections, maintain them, and transmit information along them [39-41]. Additionally, evidence supports the existence of a selection pressure in biological networks to reduce the net cost of connections. While it remains an open question how strong such selection is [41], multiple studies have shown that biological neural networks, which are hierarchical [1, 2], have been organized to reduce their amount of wiring by having fewer long connections and by locating neurons optimally to reduce the wiring between them [40, 42-44].



Fig 1. The main hypothesis. Evolution with selection for performance only results in non-hierarchical and non-modular networks, which take longer to adapt to new environments. Evolving networks with a connection cost, however, creates hierarchical and functionally modular networks that can solve the overall problem by recursively solving its sub-problems. These networks also adapt to new environments faster.

A relationship between hierarchy and connection costs can also be observed in a variety of different man-made systems. For example, very large scale integrated circuits (VLSIs), which are designed to minimize wiring, are hierarchically organized [16]. In organizations such as militaries and companies, a hierarchical communication model has been shown to be an ideal configuration when there is a cost for communication links between organization members [45]. That connection costs promote hierarchy in human-made systems suggests that the same might be true in evolved systems. Here we test that hypothesis in computational simulations of evolution and our experiments confirm that hierarchy does indeed evolve when there is a cost for network connections (Fig 1).

We also investigate the hypothesis that hierarchy confers evolvability, which has long been argued [1, 2, 16, 46], but has not previously been extensively tested [16]. Our experiments confirm that hierarchical networks, evolved in response to connection costs, exhibit an enhanced ability to adapt.

Experimentally investigating the evolution of hierarchy in biological networks is impractical, because natural evolution is slow and it is not currently possible to vary the cost of biological connections. Therefore, we conduct experiments in computational simulations of evolving networks. Computational simulations of evolution have shed substantial light on open, important questions in evolutionary biology [47-49], including the evolution of modularity [17, 18, 20, 50-52], a structural property closely related to hierarchy. In such simulations, randomly generated individuals recombine, mutate, and reproduce based on a fitness function that evaluates each individual according to how well it performs a task. The task can be analogous to efficiently metabolizing resources or performing a required behavior. This process of evolution cycles for a predetermined number of generations.

We evolved computational abstractions of animal brains called artificial neural networks (ANNs) [53, 54] to solve hierarchical Boolean logic problems (Fig 2A). In addition to abstracting animal brains, ANNs have also been used as abstractions of gene regulatory networks [55]. They abstract both because they sense their environment through inputs and produce outputs, which can either be interpreted as regulating genes or moving muscles (Methods). In our experiments, we evolve the ANNs with or without a cost for network connections. Specifically, the experimental treatment selects for maximizing performance and minimizing connection



Fig 2. A cost for network connections produces networks that are significantly more hierarchical, modular, high-performing, and likely to functionally decompose a problem. The algorithms for quantifying hierarchy and modularity are described in Methods. The bars below plots indicate at

which generation a significant difference exists between the two treatments. (A) The hierarchical AND-XOR-AND problem (the default for our experiments). The top eight nodes are inputs to the problem and the bottom node is an output. (B) P&CC networks are significantly more hierarchical than PA networks. *p*-values are from the Mann-Whitney-Wilcoxon rank-sum test, which is the default statistical test throughout the paper unless otherwise stated. (C) P&CC networks are also significantly more modular than PA networks, confirming a previous finding [17, 38]. (D) P&CC networks evolve a solution to the problem significantly faster. (E) Evolved networks from the 16 highest-performing replicates in the PA treatment. The networks are non-hierarchical, non-modular, and do not tend to decompose the problem. Each network panel reports fitness/performance (F), hierarchy (H), and modularity (M). Nodes are colored if they solve one of the logic sub-functions in (A). <u>S1 Fig</u> shows networks from all 30 replicates for both treatments. (F) Evolved networks from the 16 highest-performing replicates in the P&CC treatment. The networks from the 16 highest-performing replicates in the P&CC treatment. The networks from the 16 highest-performing replicates for both treatments. (F) Evolved networks from the 16 highest-performing replicates in the P&CC treatment. The networks from the 16 highest-performing replicates in the P&CC treatment. The networks are hierarchical, modular, and decompose the problem. (G) A comparison of P&CC and PA networks from the final generation. P&CC networks are significantly more hierarchical, modular, and solve significantly more sub-problems.

doi:10.1371/journal.pcbi.1004829.g002

costs (*performance and connection cost*, P&CC), whereas the control treatment selects for performance only (*performance alone*, PA).

Following previous work on the evolution of modularity [17, 51], in our default treatments the evolving networks are layered, feed-forward (i.e. connections are allowed only between neurons in one layer and neurons in the next layer), and have eight inputs and a single output (Methods). During evaluation, each network is tested on all possible (256) input patterns of zeros and ones, and the network's output is checked against a hierarchical Boolean logic function provided with the same input (Fig 2A and Table 1). An ANN output ≥ 0 is considered True and an output <0 is considered False. A network's performance (fitness) is its percent of correct answers over all input patterns.

Results

On the main experimental problem (Fig 2A), the addition of a connection cost leads to the evolution of significantly more hierarchical networks (Fig 2B and 2G). Confirming previous findings on different problems [17, 38], the addition of a connection cost also significantly increases modularity (Fig 2C and 2G) and reduces the number of generations required to evolve a solution (Fig 2D).

Importantly, while final performance levels for the performance and connection cost (P&CC) treatment are similar to those of the performance alone (PA) treatment, there is a qualitative difference in how the networks solve the problem. P&CC networks exhibit *functional hierarchy* in that they solve the overall problem by recursively combining solutions to sub-problems (Fig 2F), whereas the PA networks tend to combine all input information in the first layer and then process it in a monolithic fashion (Fig 2E). Such functional hierarchy can be quantified as the percent of *sub-problems* a network solves (e.g. the AND and XOR gates in Fig 2A). A sub-problem is considered solved if, for all possible network inputs, there exists a threshold such that a neuron in a network exists that outputs an above-threshold value whenever the answer to the sub-problem is True, and a sub-threshold value when the answer is False, or vice-versa (Methods). This measure reveals that evolved P&CC networks solve significantly more sub-problems than their PA counterparts (Fig 2G, $p < 2.6 \times 10^{-16}$ via Fisher's exact test).

Table 1. The main problem (pictured in Fig 2A). Networks receive 8-bit vectors as inputs. As shown, a successful network could AND adjacent input pairs, XOR the resulting pairs, and AND the result. Performance is a function only of the final output, and thus does not depend on how the network solves the problem; other, non-hierarchical solutions also exist.

	Values													
Input pattern	0	0		1		1		0		1		1		1
AND gate	0				1				0				1	
XOR gate			1								1			
AND gate							1							

doi:10.1371/journal.pcbi.1004829.t001



Fig 3. Solving sub-problems is correlated with both hierarchy (left) and modularity (right). The shape sizes and enclosed numbers indicate the number of networks at that coordinate (an empty shape indicates only one network is present). The Pearson's correlation coefficient is 0.96 for hierarchy and 0.87 for modularity, indicating strong, linear, positive relationships. Both correlations are significant (*p* < 0.00001) according to a t-test with a correlation of zero as the null hypothesis.

To further investigate how the ability to solve sub-problems is related to hierarchy and modularity, we plotted the percent of sub-problems solved vs. both hierarchy and modularity (<u>Fig 3</u>). The plots show a significant, strong, positive correlation between the ability to solve sub-problems and both hierarchy and modularity.

It has been hypothesized that one advantage of network hierarchy is that it confers evolvablity [1, 2, 16]. We test this hypothesis by first evolving networks to solve one problem (the *base environment*) and then evolving those networks to solve a different problem (the *target environment*). To isolate evolvability, we keep initial performance equal by taking the first 30 runs of each treatment (PA and P&CC) that evolve a perfectly-performing network for the base environment [17]. Each of these 30 networks then seeds 30 runs in the target environment (for 900 total replicates per treatment). The base and target problems are both hierarchical and share some, but not all, sub-problems (Fig 4). Evolution in the target environment continues until the new problem is solved or 25000 generations elapse. We quantify evolvability as the number of generations required to solve the target problems. In all experiments, P&CC networks take significantly fewer generations to adapt to the new environment than PA networks. They also solve significantly more of the target problem's sub-problems (Fig 4).

One possible reason for the fast adaptation of P&CC networks is that their modular structure allows solutions to sub-problems to be re-used in different contexts [17]. A hierarchical structure may also be beneficial if both problems are hierarchical, even if the computation at points in the structure is different. For example, modules that solve XOR gates can quickly be rewired to solve EQU gates (Fig 4A). Another reason for faster P&CC adaptation could be that these networks are sparser, meaning that fewer connections need to be optimized.

To further understand why connection costs increase hierarchy, we generated 20000 random, valid networks for each number of connections a network could have (<u>Methods</u>). A



Fig 4. P&CC networks adapt significantly faster and solve significantly more sub-problems in new environments. In these experiments, networks first evolve to solve a problem perfectly in a base environment (left) and are then placed in a target environment (right) where they continue evolving to solve a different problem. The evolvability of PA and P&CC networks is quantified as the number of generations they take to solve the new problem perfectly. A pair of evolved networks is shown for both treatments. The left one shows the network with median hierarchy (here and elsewhere, rounding up) of 30 replicates in the base environment; the right one shows the median hierarchy network of the 30 runs in the target environment started with the network on the left. <u>S5-S10</u> Figs show all network pairs.

network is valid if it has a path from each of the input nodes to the output node. The networks were neither evolved nor evaluated for performance. Of these networks, those that are low-cost tend to have high hierarchy, and those with a high cost have low hierarchy ($\underline{Fig 5}$, left). This inherent association between low connection costs and high hierarchy suggests why selecting for low-cost networks promotes the evolution of hierarchical networks. It also suggests why networks evolve to be non-hierarchical without a pressure to minimize connection costs. Indeed, most evolved PA networks reside in the high-cost, low-hierarchy region, whereas all P&CC networks occupy the low-cost, high-hierarchy region ($\underline{Fig 5}$ left).

Similarly, there is also an inverse relationship between cost and modularity for these random networks (Fig 5, right), as was shown in [17]. All evolved P&CC networks are found in the low-cost, high-modularity region; PA networks are spread over the low-modularity, highcost region (Fig 5, right).

As Figs 3 and 5 suggest, network modularity and hierarchy are highly correlated (S12 Fig, Pearson's correlation coefficient = 0.92, p < 0.00001 based on a t-test with a correlation of zero as the null hypothesis). It is thus unclear whether hierarchy evolves as a direct consequence of connection costs, or if its evolution is a by-product of evolved modularity. To address this issue we ran an additional experiment where evolutionary fitness was a function of networks being *low-cost* (as in the P&CC treatment), *high-performing* (as in all treatments), and *non-modular* (achieved by selecting for low modularity scores). We call this treatment P&CC-NonMod. The results reveal that P&CC-NonMod networks have the same low level of modularity as PA networks (Fig 6A and 6B, p = 0.23), but have significantly higher hierarchy (Fig 6A and 6C,



Fig 5. Lower cost networks are more hierarchical and modular. The hierarchy (left) and modularity (right) of randomly generated (i.e. non-functional) networks is shown for each cost after being normalized per cost value and then smoothed by a Gaussian kernel density estimation function. Colors indicate the probability of a network being generated at that location (heat map). Networks evolved in either the P&CC or PA treatment are overlaid as green circles or blue triangles, respectively. Circle or triangle size and the enclosed number indicate the number of networks at that coordinate (no number means 1). All evolved P&CC networks are in the high-hierarchy, low-cost region. Most evolved PA networks are in the high-cost, low-hierarchy region.

p < 0.00001) and solve significantly more sub-problems than PA networks (Fig 6D). These results reveal that, independent of modularity, a connection cost promotes the evolution of hierarchy. Additionally, P&CC-NonMod networks are significantly more evolvable than PA networks (Fig 6E-6G), revealing that hierarchy promotes evolvability independently of the known evolvability improvements caused by modularity [17]. To gain better insight into the relationship between modularity, hierarchy, and performance, we searched for the highest-performing networks at all possible levels of modularity and hierarchy. We performed this search with the multi-dimensional archive of phenotypic elites (MAP-Elites) algorithm [56, 57]. The results show that networks with the same modularity can have a wide range of different levels of hierarchy, and vice-versa, which indicates that these network properties can vary independently (Fig 7). Additionally, the high-hierarchy, high-modularity region, in which evolved P&CC networks reside, contains more high-performing solutions than the low-hierarchy, lowmodularity region where PA networks reside (Fig 7), suggesting an explanation for why P&CC networks find high-performing solutions faster (Fig 2D). To test the generality of our hypothesis, we ran three additional experiments. First, we repeated the main experiment with two different Boolean-logic problems that each have different logic gates, but have the same number of inputs and outputs. For all problem variants, P&CC networks are significantly more hierarchical, modular, and solve more sub-problems than PA networks (Fig 8). The P&CC treatment also evolved high-performing networks significantly faster on all of these additional problems (S13 Fig).

It is known that biological networks, such as neural systems, have long-range connections [58]. The equivalent in our model are connections that can skip any number of layers. Because such connections could affect the evolution of network modularity and hierarchy, we conducted a second experiment that investigates whether the results reported above hold when connections between any layers are allowed (Methods). The results are qualitatively unchanged: a connection cost still promoted the evolution of modularity and hierarchy (S14 Fig and S15 Fig).





Fig 6. Evolving low-cost, high-performing networks that are non-modular reveals that independent of modularity, a connection cost promotes the evolution of hierarchy. (A) Networks from the 16 highest-performing P&CC-NonMod replicates (S4 Fig shows networks from all 30 trials). The networks are hierarchal, but not highly modular. (B) There is no significant difference in modularity between P&CC-NonMod and PA networks, but P&CC-NonMod networks are significantly more hierarchical (C) and solve significantly more sub-problems (D) than PA networks. (E-G) P&CC-NonMod networks also adapt significantly faster to a new environment than PA networks, suggesting that hierarchy promotes evolvability independently of modularity. (E) The base and target problem for this evolvability experiment. (F) A perfect-performing network evolved for the base problem (left) and a descendant network evolved on the target problem (right). The example networks are those with median hierarchy: S11 Fig shows all pairs. (G) P&CC-NonMod networks adapt significantly faster to the new problem.

We also conducted a third experiment similar to the main experiment, but with a problem that has eight inputs and two outputs (see <u>S17A Fig</u>). This experiment investigates whether hierarchy and modularity evolve only in the presence of problems with many inputs and a single output (as in the main experimental problem: <u>Fig 2A</u>). Again the results are qualitatively unchanged (<u>S16 Fig and S17 Fig</u>).

Additionally, our results are qualitatively unchanged when measuring hierarchy via a different, independently created hierarchy metric, specifically that of Czégel and Palla [83] (Methods). As with the default metric developed by Mones et al. [11], under the alternate metric, final-generation P&CC networks are significantly more hierarchical than final-generation PA networks in the main treatment (S18A Fig, p < 0.001), the treatment that allows connections to skip layers (S18B Fig, p < 0.001), and the two-output treatment (S18C Fig, p < 0.001). This new metric also enabled us to verify that P&CC networks are significantly more hierarchical than PA networks on a problem with three outputs (S18D Fig, p < 0.001, problem shown in S19 Fig); the Mones et al. hierarchy metric returned pathological results for the three-output case.

Overall, the results from all of these additional experiments suggest that the hypothesis that connection costs promote the evolution of modularity and hierarchy is robust to varying key assumptions of the default model.



Fig 7. Network modularity and hierarchy can independently vary, and high-performing networks exist with a wide range of modularity and hierarchy scores. The highest-performing networks evolution discovered (with the MAP-Elites algorithm) for each combination of modularity and hierarchy. A few example networks are shown, along with their fitness (F), hierarchy (H), and modularity (M). The best network from each of the PA and P&CC treatments are also overlaid as blue triangles and green circles, respectively. The size of the circles or triangles and the enclosed number indicate the number of networks at that coordinate (no number means 1).

Discussion

The evolution of hierarchy is likely caused by multiple factors. These results suggest that one of those factors is indirect selection to reduce the net cost of network connections. Adding a cost for connections has previously been shown to evolve modularity [17, 38]; the results in this paper confirm that finding and further show that a cost for connections also leads to the evolution of hierarchy. Moreover, the hierarchy that evolves is functional, in that it involves solving a problem by recursively combining solutions to sub-problems. It is likely that there are other forces that encourage the evolution of hierarchy, and that this connection cost force operates in conjunction with them; identifying other drivers of the evolution of hierarchy and their relative contributions is an interesting area for future research.

These results also reveal that, like modularity [<u>17</u>, <u>51</u>], hierarchy improves evolvability. While modularity and hierarchy are correlated, the experiments we conducted where we explicitly select for networks that are hierarchical, but non-modular, reveal that hierarchy improves evolvability even when modularity is discouraged.

An additional factor that is present in modular and hierarchical networks is *sparsity*, a term meaning that only a few connections exist of the total that could. It is possible that this property



Fig 8. Our results are qualitatively unchanged on different problems. The P&CC networks are significantly more modular, hierarchical, and solve more sub-problems than PA networks on different, hierarchical Boolean-logic problems. For each problem, an example evolved network (specifically, the one with median hierarchy) from each treatment is shown. S2 Fig and S3 Fig show the final, evolved network from each replicate for both treatments on both problems. Note that for the problem on the right, an extra layer of hidden nodes was added due to the complexity of the problem (Methods).

explains some or even all of the evolvability benefits of modular, hierarchical networks. Future work is needed to address the difficult challenge of experimentally teasing apart these related properties.

Our study suggests a biologically plausible selection cost that leads to the evolution of sparse, hierarchical networks. Our finding is thus consistent with a previous study that found that sparse networks tend to be more hierarchical [29]. However, that previous work did not investigate whether sparsity tends to evolve on its own or which forces can encourage its evolution, nor did it investigate whether sparsity leads to hierarchy in an evolutionary context. Thus, our work both reaffirms the relationship between sparsity and hierarchy, and provides insight into their evolutionary origins.

As has been pointed out for the evolution of modularity [17], even if hierarchy, once present, is directly selected for because it increases evolvability, that does not explain its evolutionary origins, because enough hierarchy has to be present in the first place before those evolvability gains can be selected for. This paper offers one explanation for how sufficient hierarchy can emerge in the first place to then provide evolvability (or other) benefits that can be selected for.

This paper shows the effect of a connection cost on the evolution of hierarchy via experiments on many variants of one class of problem (hierarchical logic problems with many inputs and one output). In future work it will be interesting to test the generality of these results across different classes of problems, including non-hierarchical problems. The data in this paper suggest that a connection cost will always make it more likely for hierarchy to evolve, but it remains an open, interesting question how wide a range of problems hierarchy will evolve on even with a connection cost.

In addition to shedding light on why biological networks evolve to be hierarchical, this work also lends additional support to the hypothesis that a connection cost may also drive the emergence of hierarchy in human-constructed networks, such as company organizations [45], road systems [59], and the Internet [9]. Furthermore, knowing how to evolve hierarchical networks can improve medical research [60], which benefits from more biologically realistic, and thus hierarchical, network models [61, 62].

The ability to evolve hierarchy will also aid fields that harness evolution to automatically generate solutions to challenging engineering problems [63, 64], as it is known that hierarchy is a beneficial property in engineered designs [19, 22]. In fact, artificial intelligence researchers have long sought to evolve computational neural models that have the properties of modularity, regularity, and hierarchy [19, 65–68], which are key enablers of intelligence in animal brains [19, 69–71]. It has recently been shown that combining techniques known to produce regularity [72] with a connection cost produces networks that are both modular and regular [38]. This work suggests that doing so can produce networks that have all three properties, a hypothesis we will confirm in future work. Being able to create networks with all three properties will both improve our efforts to study the evolution of natural intelligence and accelerate our ability to recreate it artificially.

Methods

Experimental setup

There were 30 trials per treatment. Each trial is an independent evolutionary process that is initiated with a different random seed, meaning that the sequence of stochastic events that drive evolution (e.g. mutation, selection) are different. Each trial lasted 25000 generations and had a population size of 1000. Unless otherwise stated, all analyses and visualizations are based on the highest-performing network per trial (ties are broken randomly).

Statistics

The test of statistical significance is the Mann-Whitney-Wilcoxon rank-sum test, unless otherwise stated. We report and plot medians \pm 95% bootstrapped confidence intervals of the medians, which are calculated by re-sampling the data 5000 times. For visual clarity we reduce the re-sampling noise inherent in bootstrapping by smoothing confidence intervals with a median filter (window size of 101).

Evolutionary algorithm

Networks evolve via a multi-objective evolutionary algorithm called the Non-dominated Sorting Genetic Algorithm version II (NSGA-II), which was first introduced in [73]. While the original NSGA-II weights all objectives equally, to explore the consequence of having the performance objective be more important than the connection cost objective, Clune et al. [17] created a stochastic version of NSGA-II (called probabilistic NSGA-II, or PNSGA, implemented via the Sferes_{v2} framework [84]) where each objective is considered for selection with a certain probability (a detailed explanation can be found in [17] and Supp. S21A Fig). Specifically, performance factors into selection 100% of the time and the connection cost objective factors in *p* percent of the time. Preliminary experiments for this paper demonstrated that values of *p* of 0.25, 0.5, and 1.0 led to qualitatively similar results. However, for simplicity and because the largest differences between P&CC and PA treatments resulted from p = 1, we chose that value as the default for this paper. Note that when p = 1, NSGA-II and PNSGA are identical.

Behavioral diversity

Evolutionary algorithms notoriously get stuck in local optima (locally, but not globally, highfitness areas), in part because limited computational resources require smaller population sizes than are found in nature [64]. To make computational evolution more representative of natural evolutionary populations, which exhibit more diversity, we adopt the common technique of promoting behavioral diversity [17, 64, 74–77, 79] by adding another independent objective that rewards individuals for behaving differently than other individuals in the population. This diversity objective factored into selection 100% of the time. Preliminary experiments confirmed that this diversity-promoting technique is necessary. Without it, evolution does not reliably produce functional networks in either treatment, as has been previously shown when evolving networks with properties such as modularity [17, 38].

To calculate the behavioral diversity of a network, for each input we store a network's output (response) in a binary vector; output values > 0 are stored as 1 and 0 otherwise. How different a network's behavior is from the population is calculated by computing the Hamming distance between that network's output vector and all the output vectors of all other networks (and normalizing the result to get a behavioral diversity measure between 0 and 1).

Connection cost

Following [17], the connection cost of a network is computed after finding an optimal node placement for internal (hidden) nodes (input and output nodes have fixed locations) that minimizes the overall network connection cost. These locations can be computed exactly [78]. Optimizing the location of internal nodes is biologically motivated; there is evidence that the location of neurons in animal nervous systems are optimized to minimize the summed length of connections between them [39, 40, 78]. Network visualizations show these optimal neural placements. The overall network connection cost *cc* is then computed as the summed squared length of all connections:

$$cc = \sum_{ij} (A_{ij} dist(i,j)^2)$$
(1)

Where A_{ij} is 1 if node *i* and node *j* are connected and 0 otherwise, and dist(i, j) is the euclidean distance between node *i* and *j* after moving them to their optimal positions. Network nodes are located in a two-dimensional Cartesian space (*x*, *y*). The locations of input and output nodes are fixed, and the locations of hidden nodes vary according to their optimal location as just described. For all problems, the inputs have *x* coordinates of {-3.5, -2.5, . . ., 2.5, 3.5}, a *y* coordinate of (0), and the output is located at (0,4), except for the problem in Fig.8, right, which has an output located at (0,5) because of the extra layer of hidden neurons.

Network model and its biological relevance

The default network model is multi-layered and feed-forward, meaning a node at layer n receives incoming connections only from nodes at layer n - 1 and has outgoing connections only to nodes at layer n + 1. This network model is common for investigating questions in systems biology [52, 53, 80], including studies into the evolution of modularity [17, 51]. In the experiments that allow long-range connections that can skip layers, a node at layer n can receive incoming connections from nodes at any layer < n, and it can have outgoing connections to nodes at any layer > n.

While the layered and feed-forward nature of networks may contribute to elevated hierarchy, this network architecture is the same across all treatments and we are interested in the *differences* between levels of hierarchy that occur with and without a connection cost.

For the main problem, AND-XOR-AND, networks are of the form $8 \setminus 4 \setminus 2 \setminus 1$, which means that there are 8 input nodes, 3 hidden layers each having 4, 4 and 2 nodes respectively, and 1 output node. The integers $\{-2, -1, 1, 2\}$ are the possible values for connection weights, whereas the possible values for the biases are $\{-2, -1, 0, 1, 2\}$.

Information flows through the network in discrete time steps one layer at a time. The output y_j of node j is the result of the function:

$$y_j = \tanh\left(\lambda\left(\sum_{i\in I_j}\omega_{ij}y_i + b_j\right)\right)$$

where I_j is the set of nodes connected to node j, w_{ij} is the connection strength between node iand node j, y_i is the output value of node i, and b_j is a bias. The bias determines at which input value the output changes from negative to positive. The function, tanh(x), is an activation function that guarantees that the output of a node is in the range of [-1, 1]. The slope of the transition between the two extreme output values is determined by λ , which is here set to 20 (S21B and S21C Fig).

To be consistent with [17], in all treatments, the initial number of connections in networks is randomly chosen between 20 and 100. If that chosen number of initial connections is more than the maximum number of connections that a network can have (which is 58), the network is fully connected. When random, valid networks are generated, the initial number of connections ranges from the minimum number needed, which is 11, to the maximum number possible, which is 58. Each connection is assigned a random weight selected from the possible values for connection weights and is placed between randomly chosen, unconnected neurons residing on two consecutive layers. Our results are qualitatively unchanged when the initial number of connections is smaller than the default, i.e. when evolution starts with sparse networks (S20 Fig).

Mutations

To create offspring, a parent network is copied and then randomly mutated. Each network has a 20% chance of having a single connection added. In the default experiments, candidates for new connections are pairs of unconnected nodes that reside on two consecutive layers. For experiments that allow connections to skip layers, two different layers are randomly selected and then a node from within each layer is selected randomly and those nodes are joined by a new connection. Each network also has a 20% chance of a randomly chosen connection being removed. Each node in the network has 0.067% chance of its bias being incremented or decremented with both options equally probable; five values are available $\{-2, -1, 0, 1, 2\}$ and mutations that produce values higher or lower than these values are ignored. Each connection in the network has $\frac{2.0}{n}$ chance of its weight being incremented or decremented, where *n* is the total number of connections in the network. Because weights must be in the set of possible values $\{-2, -1, 1, 2\}$, mutations that produce values higher or lower than these network and the set of possible values are ignored.

Modularity

Network modularity is calculated by finding a partition of the network into modules that maximizes the commonly used Q modularity metric for directed networks [81], with the resulting Q score being the best estimate of the true modularity of the network. Because this metric has been extensively described before [81], here we only describe it briefly. The Q metric defines network modularity, for a particular division of the network, as the number of within-module edges minus the expected number of these edges in an equivalent network, where edges are placed randomly [82]. The formula for this metric is:

$$Q = rac{1}{m} \sum_{ij} \left[A_{ij} - rac{k_i^{in} k_j^{out}}{m}
ight] \sigma_{\scriptscriptstyle ci,cj}$$

Where k_i^{in} and k_j^{out} are the in- and out-degree of node *i* and *j*, respectively, *m* is the total number of edges in the network, A_{ij} is an entry in the connectivity matrix, which is 1 if there is an edge

from node *i* to *j*, and 0 otherwise, and $\sigma_{ci, cj}$ is a function whose value is 1 if nodes *i* and *j* belong to the same module, and 0 otherwise. Because iterating over all possible network partitions is prohibitively expensive, even for the small networks presented here, we adopt the widely used, computationally efficient spectral method for approximating the true Q score, which is described in [81].

Default hierarchy metric

Our default hierarchy measure comes from [11]. It ranks nodes based on their influence. A node's influence on a network equals the portion of a network that is reachable from that node (respecting the fact that edges are directed). Based on this metric, the larger the proportion of a network a node can reach, via its outgoing edges, the more influential it is. For example, a root node has more influence because a path can be traced from it to every node in the network, whereas leaf nodes have no influence. The metric calculates network hierarchy by computing the *heterogeneity* of the influence values of all nodes in the network. Intuitively, node-influence heterogeneity is high in hierarchical networks (where some nodes have a great deal of influence and others none), and low in non-hierarchical networks (e.g. in a fully connected network the influence of nodes is perfectly homogeneous).

Because of the non-linear function that maps a node's inputs to its output, even a small change in the input to a node can change whether it fires. For that reason, it is difficult to determine the influence one node has on another based on the strength of the connection between them. We thus calculate hierarchy scores by looking only at the presence of connections between nodes, ignoring the strength of those connections. The score for a weighted directed network is calculated by:

$$\frac{\sum_{i \in V} [C_R^{max} - C_R(i)]}{N-1}$$

 C_R^{max} is the highest influence value and V represents a set of all nodes in the network. N is the number of nodes in the network. Each node in the network is represented by *i*. $C_R(i)$, the influence value of node *i*, is given by:

$$C_{\scriptscriptstyle R}(i) = rac{1}{N-1} \sum_{j: 0 < d^{
m out}(i,j) < \infty} \left[rac{1}{d^{
m out}(i,j)}
ight]$$

Here, $d^{out}(i, j)$ is the length of the path that goes from node *i* to node *j*, meaning it is the number of outgoing connections along the path.

Alternate hierarchy metric

This section gives a brief description on the second, alternate hierarchy metric that we employed to verify our main results: a complete description of the method can be found in Czégel and Palla [83]. The metric ranks network nodes based on their influence on other nodes in the network. Such influence can be measured by, for each step of the algorithm, starting random walkers at each node in the network and allowing them to randomly traverse a single network edge in reverse order based on the relative importance of each edge. The distribution of which nodes are visited is normalized after each step, and aggregated across many steps. In the limit, a unique, stable, distribution is reached that reflects the influence of each node in the network on the other nodes. Hierarchy is then quantified as the relative standard deviation of this distribution; more heterogenous distributions reveal more hierarchical structure, because they have a few nodes with high influence on other nodes whereas most have little influence. Conversely, a highly homogenous distribution indicates a non-hierarchical structure in which the influence of each node on the others is similar.

Functional hierarchy

As a proxy for quantifying functional hierarchy, we measure the percent of logic sub-problems of an overall problem that are solved by part of a network. Note that an overall logic problem can be solved without solving any specific sub-problem (e.g. in the extreme, if the entire problem is computed in one step). We determine whether a logic sub-problem is solved as follows: for all possible inputs to a network, a neuron solves a logic gate (a sub-problem) if, for its outputs, there exists any threshold value that correctly separates all the True answers from the False answers for the logic gate in question. We also consider a sub-problem as solved if it is solved by a group of neurons on the same layer. To check for this case, we consider all possible groupings of neurons in a layer (groups of all sizes are checked). We sum the outputs of the neurons in a group and see if there is a threshold that correctly separates True and False for the logic sub-problem on all possible network inputs. To prevent counting solutions multiple times, each sub-problem is considered only once: i.e. the algorithm stops searching when a sub-problem is found to be solved.

Supporting Information

S1 Fig. The addition of a connection cost leads to the evolution of hierarchical, modular, and functionally hierarchical networks. In this visualization, networks are first sorted by fitness (F), then by hierarchy (H), and finally by modularity (M). Network nodes are colored if they solve a logic subproblem of the overall problem (Methods). (A) The main experimental problem in the paper, AND-XOR-AND. (B-C) The highest-performing networks at the end of each trial of the performance alone (PA) treatment (left) are less hierarchical, modular, and functionally hierarchical than networks from the performance and connection cost (P&CC) treatment (right).

(EPS)

S2 Fig. The results from the main experiment are qualitatively the same on a second, different, hierarchical problem: AND-EQU-AND (A). The highest-performing networks at the end of each trial of the performance alone (PA) treatment (**B**) are less hierarchical, modular, and functionally hierarchal than networks from the performance and connection cost (P&CC) treatment (**C**). Networks are first sorted by fitness (F), then by hierarchy (H), and finally by modularity (M).

(EPS)

S3 Fig. The results from the main experiment are qualitatively the same on a third, different, hierarchical problem: OR-XOR/EQU-EQU. See the previous caption for a lengthier explanation. These networks have an extra layer of hidden nodes vs. the default network model owing to the extra complexity of the last logic gate, EQU. (EPS)

S4 Fig. Evolving networks with a connection cost, but an additional explicit pressure to be non-modular, produces networks that are hierarchical, but non-modular. These results show that a connection cost promotes hierarchy independent of the modularity-inducing effects of a connection cost. **(A)** The problem for this experiment, which was the default experiment for the paper (AND-XOR-AND). **(B)** Almost all of the end-of-run networks from this P&CC-NonMod treatment are hierarchical, yet have low modularity. (EPS)

S5 Fig. (part 1 of 2) The networks from the PA treatment for the first evolvability experiment, in which networks are first evolved to perfect fitness on the AND-XOR-AND **problem and then are transferred (black arrow) to the AND-EQU-AND problem (A)**. The highest-performing network from each replicate in the base environment seeds 30 independent runs in the target environment, leading to a total of 900 replicates per treatment in the target environment. (B) In this visualization the best-performing networks from the original environment are on the left side of each arrow and on the right side is an example descendant network from the target environment (specifically, the network with median hierarchy). (EPS)

S6 Fig. (part 2 of 2) The networks from the P&CC treatment for the first evolvability experiment. See Supp. <u>S5 Fig</u> for a more detailed explanation. (EPS)

S7 Fig. (part 1 of 2) The second, AND-XOR-AND to OR-XOR-AND, evolvability experiment (A) and the networks from the PA treatment for this experiment (B). Except for a different target environment, this experiment has the same setup as the evolvability experiment in Supp. <u>S5 Fig</u>. (EPS)

S8 Fig. (part 2 of 2) The P&CC treatment networks from the second, AND-XOR-AND to OR-XOR-AND, evolvability experiment (pictured in Supp. <u>S7 Fig</u>). Except for a different target environment, this experiment has the same setup as the evolvability experiment shown in Supp. <u>S5 Fig</u>.



S9 Fig. (part 1 of 2) The third evolvability experiment, OR-XOR/EQU-EQU to AND-E-QU-AND (A), and the networks from the PA treatment for this experiment (B). Except for a different base environment, this experiment has the same setup as the evolvability experiment shown in Supp. <u>S5 Fig</u>. (EPS)

S10 Fig. (part 2 of 2) The networks from the P&CC treatment for the third, OR-XOR/ EQU-EQU to AND-EQU-AND, evolvability experiment (pictured in Supp. <u>S9 Fig</u>). Except for a different base environment, this experiment has the same setup as the evolvability experiment in Supp. <u>S5 Fig</u>.

(EPS)

S11 Fig. Evolvability is improved even in networks that are hierarchical, but non-modular, demonstrating that the property of hierarchy conveys evolvability independent of modularity. (A) The base problem that networks originally evolved on (left) and the new, target problem that networks are transferred to and further evolved on (right). (B) In each pair, on the left is a perfect-performing network evolved for the base problem and on the right is an example descendant network that evolved on the target problem (specifically, the descendant network with median hierarchy). Except for being the P&CC-NonMod treatment, this evolvability experiment has the same setup as the evolvability experiment in Supp. <u>S5 Fig</u>. (EPS)

S12 Fig. There is a strong, linear, and positive correlation between network hierarchy and modularity. The Pearson's correlation coefficient is 0.92. The correlation is significant (p < 0.00001), as calculated by a t-test with a correlation of zero as the null hypothesis. Larger circles or triangles indicate the presence of more than one network at that location (the number describes how many).

(EPS)

S13 Fig. In addition to the main experimental problem, the P&CC treatment also evolved high-performing networks faster than the PA treatment on two different problems (AND-EQU-AND, left, and OR-XOR/EQU-EQU, right; both are pictured in Fig.8 in the main text). The bar below each plot indicates when a significant difference exists between the two treatments.

(EPS)

S14 Fig. The addition of a cost for network connections also promotes the evolution of modularity and hierarchy when the default model is modified to allow connections to skip layers by removing the traditional constraint that connections are only allowed between adjacent layers (Methods). (A) The problem for this experiment is the main experimental problem (AND-XOR-AND). The evolved PA networks (B) are significantly less hierarchal and less modular (S15 Fig) than the evolved P&CC networks (C). (EPS)

S15 Fig. The results from the final generation of the experiment that allows connections to skip layers by removing the traditional constraint that connections are only allowed between adjacent layers (<u>Methods</u>). This change to the model does not change the result that P&CC networks are significantly more hierarchical, modular, and solve more sub-problems than PA networks.

(EPS)

S16 Fig. Adding a second output does not change the result that P&CC networks evolve to be significantly more (A) hierarchical, (B) modular, and (C) solve a higher percent of subproblems. <u>S17 Fig</u> shows the evolved networks. (EPS)

S17 Fig. Adding a second output does not change the result that P&CC networks evolve to be significantly more hierarchical, modular, and solve a higher percent of sub-problems (see <u>S16 Fig for p-values</u>). (A) The eight-input, two-output AND-XOR-AND/OR problem for this experiment, which is similar to the main experimental problem (<u>Fig 2A</u>, main paper), except that it has more outputs. The final, evolved networks from the performance alone (PA) treatment (B) are less hierarchical, modular, and functionally hierarchical than networks from the performance and connection cost (P&CC) treatment (**C**). (EPS)

S18 Fig. The main findings of this paper are qualitatively unchanged when a different hierarchy measure is employed. With the new metric [83] (Methods), the hierarchy of final-generation P&CC networks is significantly higher than that of PA networks for (A) the main experimental problem, (B) the experiment where connections are allowed to skip layers, (C) the two-output experiment, and (D) a three-output experiment. (EPS)

S19 Fig. The three-output problem.

(EPS)

S20 Fig. Our results are qualitatively unchanged when initializing networks with sparsely connected networks. In this experiment, the minimum and maximum number of initial connections that networks start with in generation 0 are 11 and 20, respectively. Due to the fact that at least 11 connections are needed to solve the experimental problem, networks that have an initial number of connections within this range are considered sparse (note: the default range for initial number of connections is [20, 100], <u>Methods</u>). The hierarchy (A), modularity

(B), and percent of sub-problems solved (C) are significantly higher for end-of-run P&CC networks, indicating that, regardless of the initial connectivity of networks, a connection cost promotes the evolution of these traits.

(EPS)

S21 Fig. Details of the evolutionary algorithm (figure adapted from [17]). (A) A graphical depiction of the the multi-objective evolutionary algorithm in our study, which is called the Non-dominated Sorting Genetic Algorithm version II (NSGA-II) [73]. In NSGA-II, evolution starts with a population of N randomly generated networks. N offspring are generated by randomly mutating the best of these individuals (as determined by tournament selection, wherein the best organism of 2 randomly selected organisms is chosen to produce offspring asexually). The combined pool of offspring and the current population are ranked based on Pareto dominance, and the best N networks are selected to form the next generation. This process continues for a fixed number of generations or until networks with the desired performance or properties evolve. (B) An example network model. Networks are typically used by researchers to abstract the activities of many biological networks, such as gene regulatory networks and neural networks [17, 51, 52, 55, 64]. Nodes (analogous to neurons or genes) represent processing units that receive inputs from neighbors or external sources and process them to compute an output signal that is propagated to other nodes. For example, nodes at the input layer are activated by environmental stimuli and their output is passed to internal nodes. In this figure, arrows indicate a connection between two nodes, and thus illustrate the pathways through which information flows. Each connection has a weight, which is a number that controls the strength of interaction between the two nodes. Information flows through the network, ultimately determining the firing pattern of output nodes. The firing patterns of output nodes can be considered as commands that activate genes in a gene regulatory network or that move muscles in an animal body. The output value of each node, y, is a function of its weighted inputs and bias. In this paper, the specific activation function is tanh(20x), where $x = \sum_i (w_i I_i + b)$, and where I_i is the ith input, w_i the associated synaptic weight, and b a bias that, like the weight vector, is evolved. The specific function is depicted in (C). Multiplying the input by 20 makes the function more like a step function. The output range is [-1,1]. (EPS)

Author Contributions

Conceived and designed the experiments: HM JH JBM JC. Performed the experiments: HM. Analyzed the data: HM JH JBM JC. Contributed reagents/materials/analysis tools: HM JH JBM. Wrote the paper: HM JH JBM JC.

References

- 1. Meunier D, Lambiotte R, Bullmore ET. Modular and hierarchically modular organization of brain networks. Frontiers in neuroscience. 2010 Dec 8; 4:200. doi: 10.3389/fnins.2010.00200 PMID: 21151783
- 2. Meunier D, Lambiotte R, Fornito A, Ersche KD, Bullmore ET. Hierarchical modularity in human brain functional networks. Frontiers in neuroinformatics, 3. 2009. doi: 10.3389/neuro.11.037.2009
- **3.** Miller W III. The hierarchical structure of ecosystems: connections to evolution. Evolution: Education and Outreach. 2008 Jan 1; 1(1):16–24.
- Ravasz E, Somera AL, Mongru DA, Oltvai ZN, Barabási AL. Hierarchical organization of modularity in metabolic networks. Science. 2002 Aug 30; 297(5586):1551–5. doi: <u>10.1126/science.1073374</u> PMID: <u>12202830</u>
- Yu H, Gerstein M. Genomic analysis of the hierarchical structure of regulatory networks. Proceedings of the National Academy of Sciences. 2006 Oct 3; 103(40):14724–31. doi: <u>10.1073/pnas.0508637103</u>

- Rowe R, Creamer G, Hershkop S, Stolfo SJ. Automated social hierarchy detection through email network analysis. In Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis 2007 Aug 12 (pp. 109–117). ACM.
- Krugman P. Confronting the mystery of urban hierarchy. Journal of the Japanese and International economies. 1996 Dec 31; 10(4):399–418. doi: <u>10.1006/jjie.1996.0023</u>
- Guimera R, Danon L, Diaz-Guilera A, Giralt F, Arenas A. Self-similar community structure in a network of human interactions. Physical review E. 2003 Dec 17; 68(6):065103. doi: <u>10.1103/PhysRevE.68.</u> 065103
- Vázquez A, Pastor-Satorras R, Vespignani A. Large-scale topological and dynamical properties of the Internet. Physical Review E. 2002 Jun 28; 65(6):066130. doi: <u>10.1103/PhysRevE.65.066130</u>
- Ravasz E, Barabási AL. Hierarchical organization in complex networks. Physical Review E. 2003 Feb 14; 67(2):026112. doi: <u>10.1103/PhysRevE.67.026112</u>
- Mones E, Vicsek L, Vicsek T. Hierarchy measure for complex networks. PloS one. 2012 Mar 28; 7(3): e33799. doi: 10.1371/journal.pone.0033799 PMID: 22470477
- Pumain D. Hierarchy in natural and social sciences. 3rd ed. Springer; 2006 doi: <u>10.1007/1-4020-4127-6</u>
- Lane D. Hierarchy, complexity, society. In Hierarchy in natural and social sciences 2006 (pp. 81–119). Springer Netherlands.
- Sales-Pardo M, Guimera R, Moreira AA, Amaral LA. Extracting the hierarchical organization of complex systems. Proceedings of the National Academy of Sciences. 2007 Sep 25; 104(39):15224–9. doi: <u>10.</u> <u>1073/pnas.0703740104</u>
- Lorenz DM, Jeng A, Deem MW. The emergence of modularity in biological systems. Physics of life reviews. 2011 Jun 30; 8(2):129–60. doi: 10.1016/j.plrev.2011.02.003 PMID: 21353651
- Bassett DS, Greenfield DL, Meyer-Lindenberg A, Weinberger DR, Moore SW, Bullmore ET. Efficient physical embedding of topologically complex information processing networks in brains and computer circuits. PLoS Comput Biol. 2010 Apr 22; 6(4):e1000748. doi: <u>10.1371/journal.pcbi.1000748</u> PMID: <u>20421990</u>
- Clune J, Mouret J-B, Lipson H. The evolutionary origins of modularity. Proceedings of the Royal Society of London B: Biological Sciences. 2013 Mar 22; 280(1755):20122863. doi: <u>10.1098/rspb.2012.2863</u>
- Verbancsics, P, Stanley, KO. Constraining connectivity to encourage modularity in HyperNEAT. In Proceedings of the 13th annual conference on Genetic and evolutionary computation 2011 Jul 12 (pp. 1483–1490). ACM.
- **19.** Lipson H. Principles of modularity, regularity, and hierarchy for scalable systems. Journal of Biological Physics and Chemistry. 2007 Dec; 7(4):125. doi: <u>10.4024/40701.jbpc.07.04</u>
- Wagner GP, Pavlicev M, Cheverud JM. The road to modularity. Nature Reviews Genetics. 2007 Dec 1; 8(12):921–31. doi: 10.1038/nrg2267 PMID: 18007649
- Kaiser M, Hilgetag CC, Kötter R. Hierarchy and dynamics of neural networks. Frontiers in neuroinformatics. 2010 Aug 23; 4(112).
- Suh NP. The principles of design. New York: Oxford University Press; Feb 1990.
- 23. Ozaktas HM. Paradigms of connectivity for computer circuits and networks. Optical Engineering. 1992 Jul 1; 31(7):1563–7. doi: 10.1117/12.57685
- Trusina A, Maslov S, Minnhagen P, Sneppen K. Hierarchy measures in complex networks. Physical review letters. 2004 Apr 29; 92(17):178702. doi: 10.1103/PhysRevLett.92.178702 PMID: 15169201
- Corominas-Murtra B, Rodríguez-Caso C, Goñi J, Solé R. Measuring the hierarchy of feedforward networks. Chaos: An Interdisciplinary Journal of Nonlinear Science. 2011 Mar 1; 21(1):016108. doi: <u>10.</u> <u>1063/1.3562548</u>
- Dehmer M, Borgert S, Emmert-Streib F. Entropy bounds for hierarchical molecular networks. PLoS One. 2008 Aug 28; 3(8):e3079. doi: <u>10.1371/journal.pone.0003079</u> PMID: <u>18769487</u>
- Song C, Havlin S, Makse HA. Origins of fractality in the growth of complex networks. Nature Physics. 2006 Apr 1; 2(4):275–81. doi: 10.1038/nphys266
- Ryazanov AI. Dynamics of hierarchical systems. Physics-Uspekhi, 31(3), 286–287;1988. doi: <u>10.1070/</u> PU1988v031n03ABEH005744
- Corominas-Murtra B, Goñi J, Solé RV, Rodríguez-Caso C. On the origins of hierarchy in complex networks. Proceedings of the National Academy of Sciences. 2013 Aug 13; 110(33):13316–21. doi: <u>10.1073/pnas.1300832110</u>
- 30. O'Neill RV. A hierarchical concept of ecosystems. Princeton University Press; 1986.

- Wu J, David JL. A spatially explicit hierarchical approach to modeling complex ecological systems: theory and applications. Ecological Modelling. 2002 Jul 15; 153(1):7–26. doi: <u>10.1016/S0304-3800(01)</u> 00499-9
- Flack JC, Erwin D, Elliot T, Krakauer DC. Timescales, symmetry, and uncertainty reduction in the origins of hierarchy in biological systems. Evolution cooperation and complexity. 2013 Feb 22:45–74.
- Salthe SN. Evolving hierarchical systems: their structure and representation. Columbia University Press; 2013 Aug 13.
- Sun J, Deem MW. Spontaneous emergence of modularity in a model of evolving individuals. Physical review letters. 2007 Nov 30; 99(22):228107. doi: 10.1103/PhysRevLett.99.228107 PMID: 18233336
- Pigliucci M. Is evolvability evolvable?. Nature Reviews Genetics. 2008 Jan 1; 9(1):75–82. doi: <u>10.1038/</u> nrg2278 PMID: <u>18059367</u>
- Clune, J, Beckmann, BE, McKinley, PK, Ofria, C. Investigating whether HyperNEAT produces modular neural networks. In Proceedings of the 12th annual conference on Genetic and evolutionary computation 2010 Jul 7 (pp. 635–642). ACM.
- Paine RW, Tani J. How hierarchical control self-organizes in artificial adaptive systems. Adaptive Behavior. 2005 Sep 1; 13(3):211–25. doi: 10.1177/105971230501300303
- Huizinga J, Mouret J-B, Clune J. Evolving Neural Networks That Are Both Modular and Regular: Hyper-Neat Plus the Connection Cost Technique. Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation; 2014.
- Cherniak C, Mokhtarzada Z, Rodriguez-Esteban R, Changizi K. Global optimization of cerebral cortex layout. Proceedings of the National Academy of Sciences, 101(4), 1081–1086;2004. doi: <u>10.1073/</u> pnas.0305212101
- Chen BL, Hall DH, Chklovskii DB. Wiring optimization can relate neuronal structure and function. Proceedings of the National Academy of Sciences. 2006 Mar 21; 103(12):4723–8. doi: <u>10.1073/pnas.</u>0506806103
- 41. Sporns O. Networks of the Brain. MIT press (2011).
- **42.** Raj A, Chen YH. The wiring economy principle: connectivity determines anatomy in the human brain. PloS one. 2011 Sep 7; 6(9):e14832. doi: <u>10.1371/journal.pone.0014832</u> PMID: <u>21915250</u>
- Ahn YY, Jeong H, Kim BJ. Wiring cost in the organization of a biological neuronal network. Physica A: Statistical Mechanics and its Applications. 2006 Jul 15; 367:531–7. doi: 10.1016/j.physa.2005.12.013
- Laughlin SB, Sejnowski TJ. Communication in neuronal networks. Science. 2003 Sep 26; 301 (5641):1870–4. doi: <u>10.1126/science.1089662</u> PMID: <u>14512617</u>
- 45. Guimera R, Arenas A, Diaz-Guilera A. Communication and optimal hierarchical networks. Physica A: Statistical Mechanics and its Applications. 2001 Oct 1; 299(1):247–52. doi: <u>10.1016/S0378-4371(01)</u> 00302-8
- 46. Simon HA. The architecture of complexity. Proceedings of the American Philosophical Society. 1962
- Lenski RE, Ofria C, Collier TC, Adami C. Genome complexity, robustness and genetic interactions in digital organisms. Nature. 1999 Aug 12; 400(6745):661–4. doi: <u>10.1038/23245</u> PMID: <u>10458160</u>
- Lenski RE, Ofria C, Pennock RT, Adami C. The evolutionary origin of complex features. Nature. 2003 May 8; 423(6936):139–44. doi: <u>10.1038/nature01568</u>
- 49. Wilke CO, Wang JL, Ofria C, Lenski RE, Adami C. Evolution of digital organisms at high mutation rates leads to survival of the flattest. Nature. 2001 Jul 19; 412(6844):331–3. doi: <u>10.1038/35085569</u> PMID: <u>11460163</u>
- Espinosa-Soto C, Wagner A. Specialization can drive the evolution of modularity. PLoS Comput Biol. 2010 Mar 26; 6(3):e1000719. doi: <u>10.1371/journal.pcbi.1000719</u> PMID: <u>20360969</u>
- Kashtan N, Alon U. Spontaneous evolution of modularity and network motifs. Proceedings of the National Academy of Sciences. 2005 Sep 27; 102(39):13773–8. doi: <u>10.1073/pnas.0503610102</u>
- Kashtan N, Noor E, Alon U. Varying environments can speed up evolution. Proceedings of the National Academy of Sciences. 2007 Aug 21; 104(34):13711–6. doi: <u>10.1073/pnas.0611630104</u>
- Alon U. An introduction to systems biology: design principles of biological circuits. CRC press; 2006 Jul 7.
- 54. Trappenberg T. Fundamentals of computational neuroscience. Oxford University Press; 2009.
- Geard N, Wiles J. A gene network model for developing cell lineages. Artificial Life. 2005; 11(3):249– 67. doi: 10.1162/1064546054407202 PMID: 16053570
- Mouret, J-B, Clune J. Illuminating search spaces by mapping elites. arXiv preprint arXiv:1504.04909; 2015.

- Cully A, Clune J, Tarapore D, Mouret J-B. Robots that can adapt like animals. Nature 521, 503–507; 2015. doi: <u>10.1038/nature14422</u> PMID: <u>26017452</u>
- Kaiser M, Hilgetag CC. Nonoptimal component placement, but short processing paths, due to long-distance projections in neural systems. PLoS Comput Biol. 2006 Jul 21; 2(7):e95. doi: <u>10.1371/journal.</u> <u>pcbi.0020095</u> PMID: <u>16848638</u>
- Louf R, Jensen P, Barthelemy M. Emergence of hierarchy in cost-driven growth of spatial networks. Proceedings of the National Academy of Sciences. 2013 May 28; 110(22):8824–9. doi: <u>10.1073/pnas.</u> <u>1222441110</u>
- Zhang Y, Huang Z, Zhu Z, Liu J, Zheng X, Zhang Y. Network analysis of ChIP-Seq data reveals key genes in prostate cancer. European journal of medical research. 2014 Sep 3; 19(1):47. doi: <u>10.1186/</u> <u>s40001-014-0047-7</u> PMID: <u>25183411</u>
- Shmulevich I, Dougherty ER, Kim S, Zhang W. Probabilistic Boolean networks: a rule-based uncertainty model for gene regulatory networks. Bioinformatics. 2002 Feb 1; 18(2):261–74. doi: <u>10.1093/</u> <u>bioinformatics/18.2.261</u> PMID: <u>11847074</u>
- 62. Albert R. Scale-free networks in cell biology. Journal of cell science. 2005 Nov 1; 118(21):4947–57. doi: 10.1242/jcs.02714 PMID: 16254242
- 63. Koza JR, Keane MA, Streeter MJ, Mydlowec W, Yu J, Lanza G. Genetic programming IV: Routine human-competitive machine intelligence. Springer Science & Business Media; 2006 Mar 4.
- Floreano D, Mattiussi C. Bio-inspired artificial intelligence: theories, methods, and technologies. MIT press; 2008 Aug 22.
- Stanley KO, Miikkulainen R. A taxonomy for artificial embryogeny. Artificial Life. 2003; 9(2):93–130. doi: 10.1162/106454603322221487 PMID: 12906725
- Hornby, GS. Measuring, enabling and comparing modularity, regularity and hierarchy in evolutionary design. In Proceedings of the 7th annual conference on Genetic and evolutionary computation 2005 Jun 25 (pp. 1729–1736). ACM.
- Clune J, Stanley KO, Pennock RT, Ofria C. On the performance of indirect encoding across the continuum of regularity. Evolutionary Computation, IEEE Transactions on. 2011 Jun; 15(3):346–67. doi: <u>10.</u> 1109/TEVC.2010.2104157
- Gruau F. Automatic definition of modular neural networks. Adaptive behavior. 1994 Sep 1; 3(2):151– 83. doi: 10.1177/105971239400300202
- **69.** Nolfi S, Floreano D. Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines. MIT press; 2000.
- 70. Striedter GF. Principles of brain evolution. Sinauer Associates;2005.
- Wagner GP, Altenberg L. Perspective: complex adaptations and the evolution of evolvability. Evolution. 1996 Jun 1:967–76. doi: <u>10.2307/2410639</u>
- Stanley KO, D'Ambrosio DB, Gauci J. A hypercube-based encoding for evolving large-scale neural networks. Artificial life. 2009 Mar 5; 15(2):185–212. doi: 10.1162/artl.2009.15.2.15202 PMID: 19199382
- 73. Deb K. (2001). Multi-objective optimization using evolutionary algorithms (Vol. 16). John Wiley & Sons.
- Mouret J-B, Doncieux S. Overcoming the bootstrap problem in evolutionary robotics using behavioral diversity. In Evolutionary Computation, 2009. CEC'09. IEEE Congress on 2009 May 18 (pp. 1161– 1168). IEEE.
- Doncieux, S, Mouret J-B. Behavioral diversity measures for evolutionary robotics. In Evolutionary Computation (CEC), 2010 IEEE Congress on 2010 Jul 18 (pp. 1–8). IEEE.
- 76. Mouret J-B, Doncieux S. Encouraging behavioral diversity in evolutionary robotics: An empirical study. Evolutionary computation. 2012 Mar; 20(1):91–133. doi: <u>10.1162/EVCO_a_00048</u> PMID: <u>21838553</u>
- Risi, S, Vanderbleek, SD, Hughes, CE, Stanley, KO. How novelty search escapes the deceptive trap of learning to learn. In Proceedings of the 11th Annual conference on Genetic and evolutionary computation 2009 Jul 8 (pp. 153–160). ACM.
- Chklovskii DB. Exact solution for the optimal neuronal layout problem. Neural computation. 2004 Oct; 16(10):2067–78. doi: 10.1162/0899766041732422 PMID: 15333207
- Lehman J, Stanley KO. Abandoning objectives: Evolution through the search for novelty alone. Evolutionary computation. 2011 May 4; 19(2):189–223. doi: <u>10.1162/EVCO_a_00025</u> PMID: <u>20868264</u>
- Karlebach G, Shamir R. Modelling and analysis of gene regulatory networks. Nature Reviews Molecular Cell Biology. 2008 Oct 1; 9(10):770–80. doi: 10.1038/nrm2503 PMID: 18797474
- Leicht EA, Newman ME. Community structure in directed networks. Physical review letters. 2008 Mar 21; 100(11):118703. doi: <u>10.1103/PhysRevLett.100.118703</u> PMID: <u>18517839</u>

- Newman ME. Modularity and community structure in networks. Proceedings of the national academy of sciences. 2006 Jun 6; 103(23):8577–82. doi: <u>10.1073/pnas.0601602103</u>
- **83.** Czégel D, Palla G. Random walk hierarchy measure: What is more hierarchical, a chain, a tree or a star?. Scientific reports. 2015; 5.
- **84.** Mouret J-B, Doncieux S. Sferes v2: Evolvin'in the multi-core world. In Evolutionary Computation (CEC), 2010 IEEE Congress on (pp.1–8). IEEE.