

Generalization as the great leap in evolvability: insights from machine learning

Steven A. Frank 

Department of Ecology and Evolutionary Biology, University of California, Irvine, United States

Corresponding author: Department of Ecology and Evolutionary Biology, University of California, Irvine, CA 92697-2525, United States. Email: saf@stevefrank.org

Abstract

Natural selection encodes learned information in the genome. Learned solutions may be tuned specifically to past challenges, failing in altered environments. Or solutions can be general, capturing the essential structure of the challenge and performing well across variations within the abstract class. For example, a neural system might recognize the exact outlines of a rattlesnake but not other snakes, or it might recognize the essence of snakesness. The problem of how a system generalizes is a fundamental aspect of evolvability, the ability of a system to learn broad solutions to novel challenges. In recent years, machine learning has significantly advanced our understanding of when systems generalize their learned solutions and how they accomplish such generalization. One surprising discovery overturned conventional wisdom about learning. Large systems, with more adjustable parameters than the dimensions of the incoming data, do not merely memorize the data patterns in the way suggested by traditional theory. Instead, systems with more parameters generalize better than smaller systems. Because natural selection is a learning algorithm, the new theory of generalization applies to biological evolution. Specifically, increasing regulatory complexity and parameterization associates with increasing evolvability for the discovery of general solutions. This link between genomic complexity and generalization may have been a primary driving force in evolutionary history.

Keywords: genetic regulatory networks, fitness landscape, double descent, overparameterization

Introduction

The distinction between memorization and generalization is the central problem of learning. Natural selection is a basic learning algorithm, inevitably facing this same problem (Valiant, 2013; Watson & Szathmáry, 2016).

In the past, statistics and machine learning predicted a tradeoff. A system that is too simple cannot capture general patterns, it underfits. A system that is too complex memorizes every idiosyncratic fluctuation in the data, overfitting and missing any underlying generality. There should be an ideal spot of intermediate complexity at which the system is just large enough to capture the general nature of the problem without memorizing the noise.

Recent advances in machine learning have overturned that seemingly fundamental tradeoff. Vastly overparameterized systems generalize remarkably well compared with smaller systems. The great breakthroughs in artificial intelligence come from gigantic systems, with far greater numbers of parameters than the dimensionality of the data used for learning (Belkin et al., 2019; Nakkiran et al., 2021).

These observations have led to new theory about how systems learn to generalize. This article translates that theory into biology, where natural selection faces the same fundamental problem of learning and generalization (Kouvaris et al., 2017; Parter et al., 2008; Watson & Szathmáry, 2016). Three questions outline the problem.

How do biological systems generalize their solutions to challenges? For example, an immune system might memorize specific danger signals, or it might have the general abil-

ity to infer danger by combining the weight of evidence from multiple indicators.

What characteristics of genomic regulatory structure or neural wiring associate with memorization versus generalization? The new theory suggests that very large parameter spaces generalize better, which means that complexly overwired gene regulatory networks or neural networks should generalize more broadly.

What comparative predictions arise over the history of life with regard to the tendency for greater generalization? Transitions from prokaryotes to single-cell eukaryotes to multicellular eukaryotes associate with increasing regulatory dimensionality, predicting an increasing historical trend in generalization.

Evolvability and generalization synergize in an interesting way. Evolvability is the capacity to reach adaptive solutions to future challenges. Generalization arises when an adaptive solution works well in situations beyond the initial challenge that shaped it. When the two combine, each adaptation is no longer a direct fix to a particular challenge but a competence that extends across whole classes of future challenges. The lineage stops collecting narrow solutions one at a time and starts accumulating broad ones. That shift in the character of adaptive accumulation is what I call a great leap in evolvability.

This new perspective changes the way that we think about fitness landscapes. A fitness landscape describes how well genotypes perform in a particular environment. Machine learning calls that landscape a training-error surface,

measuring the distance between a genotype's performance and an optimal solution to the challenge.

The new learning theory has shown that training error alone is insufficient for understanding generalization (Belkin et al., 2019). In highly overparameterized systems, many different solutions fit the training data equally well. The excess dimensions create large subspaces of equivalent training performance but varying generalization ability. Within these subspaces, learning dynamics preferentially reach simple, smooth solutions that capture the general structure of the problem rather than its idiosyncratic details (Wilson, 2025).

High dimensionality also protects against a system's tendency to confuse true structure with noise. Systems with excess dimensions can avoid this confusion because they have enough degrees of freedom to separate noise from true pattern. Overparameterized systems do not just have more ways to fit. They have more ways to fit well once they find the right solution region, and the dynamics of learning guide them to the right place (Transtrum et al., 2025).

These principles predict that as regulatory circuits become more complex and highly parameterized, they will tend to shift from particular to general solutions. Thus, we can order life's tendency for generality. Prokaryotes, with compact genomes, likely encode relatively specific responses to specific signals. Single-celled eukaryotes, with more complex regulatory control, tend toward more general solutions. Multicellular organisms should have the greatest generalization.

Theory of generalization

Classical statistical theory predicts that a system with more parameters than data points will overfit. It will memorize every detail of the training set, including noise. An overfit system predicts outputs poorly when confronted with new inputs that differ from the original data.

Overfitting arises as one approaches the interpolation threshold, at which the number of parameters matches the complexity of the data. At that threshold, there is essentially one solution that closely fits, a purely memorized description of the data.

The surprise from recent studies in machine learning is what happens beyond the interpolation threshold. As the number of parameters grows far past the number needed to fit the data, the system's ability to predict outputs for novel data improves (Belkin et al., 2019; Nakkiran et al., 2021).

The overall error pattern can be measured with respect to a generalized ability to predict outputs for new inputs. That error pattern as a function of the number of system parameters traces a double descent (Figure 1).

At first, the error declines as the number of parameters increases from zero, because each new parameter provides some power to match the data. But then the system begins to memorize too much, and the error against new data rises toward the interpolation threshold. After that threshold, with more parameters, the error descends again, the surprising second descent.

A similar double descent pattern arises with respect to the temporal dynamics of learning (Nakkiran et al., 2021). Dynamics describe how a highly overparameterized system learns in response to a stream of training data that include

both inputs and outputs. At first, the system learns something about the early data stream, improving performance for a while when the system is tested by its ability to predict outputs for new inputs. This initial decline in the system's error is the first descent.

However, after a continuation of the training data stream, the system becomes very brittle, because it has memorized much of the detail in the observations, including noise and idiosyncrasies of particular observations that do not apply in general. Thus, the error starts to rise toward a peak at the interpolation threshold of memorization.

However, in a highly overparameterized system, continued learning often causes the error to decline again as the system passes the interpolation regime and moves into a regime of generalization. That second generalization phase is the second descent in the error curve when tested against novel inputs.

In response to these observations about the nature of learning, recent studies have developed three complementary lines of theory about how system dimensionality influences generalization.

The first theory concerns the ability of learning dynamics to reach particular solutions. When many solutions fit equally well, gradient-based learning preferentially finds simple, smooth ones (Wilson, 2025). Simple solutions occupy larger regions of parameter space, and learning dynamics are more likely to find big regions than small ones. Put another way, overparameterization creates an abundance of solutions. Among those solutions, the dynamics select for simple and general models of the world.

The second theory concerns aliasing, the confusion between true structure and noise (Transtrum et al., 2025). At the interpolation threshold, aliasing is maximal because the system has no spare capacity to distinguish signal from noise. Every degree of freedom is committed to fitting exactly. Excess dimensions past that rigidly constrained fit provide the slack needed to separate pattern from accident.

Biologically, different regulatory mappings may produce the same phenotype in the environments encountered so far but diverge in a new environment. Extra dimensions help the circuit separate the shared structure that is generally favored from incidental correlations that arose because of the particular environmental history and dimensional constraints.

The third theory concerns the geometry of the interpolation spike (Schaeffer et al., 2023). That spike requires that the data have directions of greater and lesser variation and that novel inputs used for testing must differ in their directions of variation. Near the threshold, errors concentrate along low-variance directions because fitting concentrates along high-variance directions. Adding dimensions spreads the errors widely, so that a mismatch only picks up a few small contributions to overall error rather than a big patch of errors linked along a particular constrained dimension.

Biologically, if past environments sample mostly one direction of variation, selection tunes that direction while leaving others poorly tuned. A novel environment along one of those weakly sampled directions exposes concentrated error. With more dimensions, error is spread over many small mismatches rather than concentrated along one brittle direction, making the system less sensitive to novel environments.




Figure 1. Double descent of generalization error in learning. The solid curve shows generalization error, the performance of a learned solution against novel inputs not encountered during learning. The dashed curve shows training error, the performance against the data used for learning. In the classical regime, increasing system complexity first reduces generalization error. But then the generalization error increases as the system begins to memorize idiosyncratic details of the training data. At the interpolation threshold, the system has just enough parameters to describe the training data exactly, and generalization error peaks. Beyond the threshold, in the overparameterized regime, generalization error declines even though training error remains at zero. The x-axis can be read as the number of adjustable parameters in the system or as learning time for a system with many parameters. After figure 1 of Elad et al. (2020).

These three results are complementary. Wilson (2025) explains why, among the vast number of solutions available to an overparameterized system, the ones actually found tend to be general rather than idiosyncratic. Transtrum et al. (2025) explains the spike as maximal aliasing and extra dimensions as the remedy. Schaeffer et al. (2023) explains when the interpolation spike occurs and why extra dimensions dissolve it.

Classical statistics penalizes complexity to avoid overfitting, effectively smoothing out the interpolation spike. By contrast, biology tends not to penalize complexity as strongly. Evolutionary dynamics is therefore likely to experience the full consequences of the double descent learning curve. That fact leads to new implications for understanding how regulatory systems form specialized or generalized solutions, and for understanding how fitness landscapes truly influence evolutionary process.

Fitness landscapes

The new theory of generalization changes how we should think about fitness landscapes. The classical landscape problem concerns how performance increases in relation to the fitness gradient, and how peaks and valleys favor or impede trajectories to better performance.

Traditionally, the trajectories describe how evolutionary process updates the system in response to performance in each realized environment. By contrast, generalization concerns performance in relation to the potential set of environments that might be encountered.

Now that we are thinking of generalization rather than realized performance, we can ask some new questions. How does the dimensionality of a solution—of a phenotype—alter generalization? How does the nature of the environmental challenge influence generalization?

The first step may be consideration of phenotypes. Biological fitness landscapes focus on the map of genotypes to phenotypes to fitness. The genotype-to-phenotype part of the mapping is widely considered to be a crucial and difficult component (Payne & Wagner, 2019; Wagner & Altenberg, 1996).

I favor thinking of phenotypes as biochemical or neural circuits that take inputs and compute outputs (Frank, 2025a, 2025b). That encompasses simple metric traits and also traits that are solutions to complex environmental challenges. A computational circuit perspective also provides a direct link to machine learning, which often proceeds by the improving performance trajectory of a computational circuit to solve a specifically defined problem.

In particular, the circuit framing translates evolutionary adaptation into a common type of machine learning challenge. Environmental and internal state signals are the inputs. Regulatory connections, binding strengths, thresholds, and interaction terms are the circuit parameters. Phenotypic responses are the outputs. The fitness landscape mirrors the training-error surface, ranking parameters by their performance in encountered environments.

Generalization asks different questions. Among the parameters that perform about the same over those historical environments, which does selection actually find? How do those parameters perform in novel environments? The new

machine learning theory says that generalized performance in novel environments depends on the dimensionality of the parameter space defined by the regulatory architecture.

Neurobiology is of course based on computational circuits, and gene regulatory networks have been studied extensively. But few studies have asked how and why regulatory networks generalize to novel inputs. Parter et al. (2008) showed that circuits evolved under varying environments generalize to novel environments. Xue et al. (2019) showed that the dimensionality of internal representation determines the range of available strategies. Kouvaris et al. (2017) looked at network modularity and connectivity in relation to generalization.

The study of such circuits fits well within the traditional analysis of fitness landscapes. In that case, the circuits improve in response to a sequence of realized environments that favor particular parameter combinations over others, a typical learning trajectory over a performance or fitness landscape.

The next step will be a deeper consideration of how dimensionality alters the learned solutions' ability to generalize to novel inputs. In light of the recent insights from machine learning, how can we understand the fitness landscape geometry in terms of the evolutionary dynamics of generalization? That question provides a way to extend prior work, in which Gavrillets showed that high-dimensional landscapes have extensive nearly neutral networks connecting distant genotypes (Gavrillets, 2004). The theory of generalization extends that insight by shifting attention to how learning dynamics select among solutions within those connected regions.

Discussion

The theory of generalization abstractly describes the geometry of learning. That abstract theory raises a broad prediction about biology.

Greater dimensionality of phenotypic circuits should correspond to a greater tendency for generalization. For example, single-celled eukaryotes have greater genetic regulatory capacity and dimensionality compared with most prokaryotes. Thus, yeast should generalize more than bacteria. And multicellular eukaryotes should generalize more than single cells.

Similar ideas about evolutionary flexibility have been discussed. Davidson (2006) showed that more complex organisms have greater hierarchical depth in their regulatory architecture, which he argued enables evolutionary flexibility. Wagner (2005) showed that redundant networks have broad neutral spaces that potentially connect phenotypic innovations. Kirschner & Gerhart (2005) identified organizational features, such as weak linkage between compartmentalized components, that bias variation toward functionality. But none explained why more wiring connectivity should produce generalization.

The theory here suggests that generalization arises from the sheer dimensionality of the regulatory space. The specific circuit features identified in prior work may be important for evolvability but not as the primary cause of generalizing capacity. Instead, processes that increase dimensionality may matter more. For example, neutrality and robustness have been raised as primary causes of increasing regulatory dimensionality (Frank, 2023).

Secondarily, older circuits are more likely to generalize. In the temporal version of the double descent pattern discussed earlier, a learning system first improves, then passes through a brittle memorizing phase, then reaches generalization. Applied to biology, recently evolved regulatory circuits should be more likely to show brittle, context-specific correlations among components. Older circuits should be more likely to generalize across a range of conditions.

In summary, the new framework for generalization arises from the fact that natural selection is a basic learning algorithm. Thus, the theory of generalization in learning, originally developed to explain why overparameterized machine learning systems abstract rather than memorize, also applies to biology. Regulatory connections are parameters, selective history is training data, selection is the learning optimizer, and novel environments are test data.

This perspective adds an aspect to evolutionary theory that has been missing. The fitness landscape remains useful for local questions about epistasis, adaptive walks, and similar evolutionary problems. But for the deeper question of how lineages discover general solutions to environmental challenges, the relevant factor is not a position on a fitness surface. It is how the dimensionality of the regulatory architecture alters the geometry of evolutionary trajectories within the fitness landscape.

Data availability

This article generated no new data or computer code.

Funding

The Donald Bren Foundation and US National Science Foundation grant DEB-2325755 support my research.

Conflict of interest

The author declares no conflict of interests.

References

- Belkin, M., Hsu, D., Ma, S. Y., & Mandal, S. (2019). Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32), 15849–15854. <https://doi.org/10.1073/pnas.1903070116>
- Davidson, E. H. (2006). *The regulatory genome: Gene regulatory networks in development and evolution*. Academic.
- Elad, M., Simon, D., & Aberdam, A. (2020). Another step toward demystifying deep neural networks. *Proceedings of the National Academy of Sciences*, 117, 27070–27072. <https://doi.org/10.1073/pnas.2018957117>
- Frank, S. A. (2023). Robustness and complexity. *Cell Systems*, 14(12), 1015–1020. <https://doi.org/10.1016/j.cels.2023.11.003>
- Frank, S. A. (2025a). Circuit design in biology and machine learning. I. Random networks and dimensional reduction. *Evolution; International Journal of Organic Evolution*, 79, 1403–1418. <https://doi.org/10.1093/evolut/qpaf065>
- Frank, S. A. (2025b). Circuit design in biology and machine learning. II. Anomaly detection. *Entropy*, 27(9), 896. <https://doi.org/10.3390/e27090896>
- Gavrillets, S. (2004). *Fitness landscapes and the origin of species*. Princeton University Press. <https://doi.org/10.1515/9780691187051>
- Kirschner, M., & Gerhart, J. (2005). *The plausibility of life: resolving darwin's dilemma*. Yale University Press.

- Kouvaris, K., Clune, J., Kounios, L., Brede, M., & Watson, R. A. (2017). How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation. *PLOS Computational Biology*, 13(4), e1005358. <https://doi.org/10.1371/journal.pcbi.1005358>
- Nakkiran, P., Kaplun, G., Bansal, Y., Yang, T. S., Barak, B., & Sutskever, I. (2021). Deep double descent: where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(12), 124003. <https://doi.org/10.1088/1742-5468/ac3a74>
- Parter, M., Kashtan, N., & Alon, U. (2008). Facilitated variation: how evolution learns from past environments to generalize to new environments. *PloS Computational Biology*, 4(11), e1000206. <https://doi.org/10.1371/journal.pcbi.1000206>
- Payne, J. L., & Wagner, A. (2019). The causes of evolvability and their evolution. *Nature Reviews Genetics*, 20(1), 24–38. <https://doi.org/10.1038/s41576-018-0069-z>
- Schaeffer, R., Khona, M., Robertson, Z., Boopathy, A., Pistunova, K., Rocks, J. W., Fiete, I. R., & Koyejo, O. (2023). Double descent demystified: identifying, interpreting & ablating the sources of a deep learning puzzle. *arXiv:2303.14151*.
- Transtrum, M. K., Hart, G. L. W., Jarvis, T. J., & Whitehead, J. P. (2025). Generalized aliasing explains double descent and informs model design. *Physical Review Research*, 7(4), 043268. <https://doi.org/10.1103/physrevresearch.7.043268>
- Valiant, L. (2013). *Probably approximately correct: Nature's algorithms for learning and prospering in a complex world*. Basic Books.
- Wagner, A. (2005). *Robustness and evolvability in living systems*. Princeton University Press.
- Wagner, G. P., & Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution; International Journal of Organic Evolution*, 50, 967–976. <https://doi.org/10.1111/j.1558-5646.1996.tb02339.x>
- Watson, R. A., & Szathmáry, E. (2016). How can evolution learn? *Trends in Ecology & Evolution*, 31(2), 147–157. <https://doi.org/10.1016/j.tree.2015.11.009>
- Wilson, A. G. (2025). *Deep learning is not so mysterious or different*. PMLR.
- Xue, B. K., Sartori, P., & Leibler, S. (2019). Environment-to-phenotype mapping and adaptation strategies in varying environments. *Proceedings of the National Academy of Sciences*, 116(28), 13847–13855. <https://doi.org/10.1073/pnas.1903232116>

Received April 11, 2026; revisions received May 18, 2026; accepted June 4, 2026

Associate Editor: Stephen Proulx; Handling Editor: Jason Wolf

© The Author(s) 2026. Published by Oxford University Press on behalf of The Society for the Study of Evolution (SSE).

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.