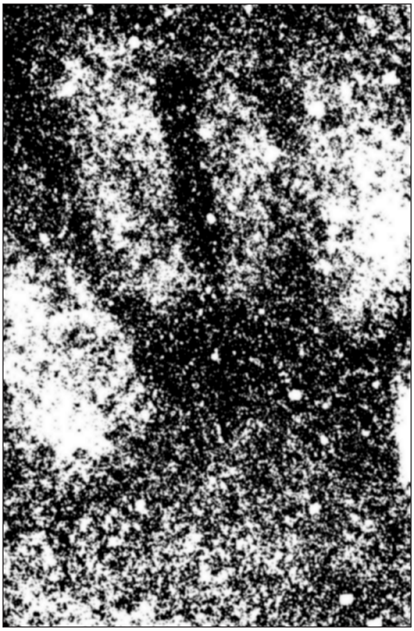

EMERGENT ENGINEERING: REFRAMING THE GRAND CHALLENGE FOR THE 21ST CENTURY

David C. Krakauer, SFI
Santa Fe Institute, 2019

How can human society have reached the moon, harnessed the unobservable mechanics of the atom, continued to build computers that become exponentially faster and cheaper each year, and yet have operated so poorly in establishing stable economies, reducing the incidence of conflict and disease, and discovering and manufacturing effective biomedical drugs? It is certainly not through lack of interest, resources, effort, and intelligence.

The war on cancer, the pursuit of greater economic equality and financial stability, the creation of online safety and security, and the invention of new nontoxic and effective pharmacological drugs have absorbed astronomical sums of money into both research and development—and yet in so many cases they have foundered and failed through the misapplication of previously highly successful ideas of engineering and design to complex systems.

There is an urgent need for novel concepts directed at achieving an evolutionary and emergent engineering, and it is our contention that they are likely to come from the domains of biological and social life—not from the deterministic world of designed mechanical artifacts.

A History of Success and Failure: The Siren Song of the Grand Challenge

There are a handful of technology projects of such sheer audacity and scale that they have become bywords for human ambition, ingenuity, and impact. Included among these projects should certainly be numbered the Apollo Program, the Manhattan Project, CERN and the LHC, the Human Genome Project, the Panama Canal, and the Great Wall of China. Comparably impressive in scale and cost is the continued application of some version of Moore's law to integrated circuit design.

Over the last several decades, new discoveries and theories have emerged to study complex systems—networks of adaptive agents—that promise to be better suited to addressing economics, disease, cybersecurity, and medical treatment than erstwhile approaches founded on very successful (where appropriate) classical engineering axioms.

The Apollo project, launched by President Kennedy in 1961 as an "urgent national need," expended $25.4 billion (around $165 billion, adjusted) and employed over four hundred thousand workers between 1959 and 1970. For the sake of comparison, the vastly less edifying financial bailout of 2008 (an effort to correct the technical deficiencies of designed market instruments) cost just short of $750 billion.

The Manhattan project was launched to produce nuclear weapons in response to the global threat from Nazi Germany and its allies. From 1939 to 1946, around 130 thousand people were employed at a cost of $2 billion ($22 billion, adjusted), spread across thirty sites, to produce a bomb. The so-called war on cancer (borrowing the martial metaphor to instill a disciplined sense of urgency and determination, and to justify huge expenditures) has been waged since it was first declared by Nixon in 1971. Research costs have been on the order of $100 billion, with the greatest gains in longevity coming through modified behavior and improved screening, whereas biomedical engineering has produced rather modest increases (a few percentage points) in life extension, according to Decennial Life Table data.

Moore's law describes the doubling of transistor density every two years from 1970 until the near present. This exponential trend has been accompanied by similar trends in reduced cost and increased memory capacity. The economic burden maintaining these trends has been considerable, including the cost of fabrication—a modern fabrication plant costs around $2 billion to build—and the percentage of revenue consumed by research and development increased by 40 percent over the decade from 1999 to 2009. Compare this to the deliciously impudent Eroom's law (Moore's law reversed), which describes how drug discovery since the 1980s has become more expensive, slowed down delivery of drugs to those in need, and produced ineffective and doubtful remedies—and this despite the giddy hyperbole surrounding drug design, machine learning, and high-throughput technologies.

Over the last several decades, new discoveries and theories have emerged to study complex systems—networks of adaptive agents—that promise to be better suited to addressing economics, disease, cybersecurity, and medical treatment than erstwhile approaches founded on very successful (where appropriate) classical engineering axioms.

The Classical Engineering Axioms

1. Design according to well-understood scientific principles that hold for all components in isolation and in aggregate.

2. Design systems with as near as possible fault-free components to very high levels of combined precision.
3. Minimize error and accept only the smallest system failure rates by eliminating uncertainty and reducing degrees of freedom of components.
4. Design systems into linear ranges of operation where collective dynamics are predictable and controllable.
5. Reduce noise and adaptability of components to prevent unexpected emergent behaviors.

To calibrate the difficulty of applying this kind of framework to complex problems, consider an appropriately complementary list of properties found in most complex systems of the kind that govern the societies in which we live, from the immune system to the nervous system, as well as ecosystems, economies, and political institutions.

The Properties of Complex Systems

1. We have few general design principles for adaptive components (cells, organisms, nations) in isolation or in the aggregate where new unforeseen properties emerge.
2. Components typically have high failure rates in all tasks and accomplish their objectives through statistical averaging and approximation across multiple scales and levels.
3. There is significant uncertainty and lack of information at both the component and aggregate level, and components have large—and often poorly understood—repertoires of behavior.
4. Most evolved complex systems operate in nonlinear and often near-critical regimes (close to thresholds and tipping points).

5. Adaptability of components is the rule—not the exception—and learning and adaptation are ongoing and irrepressible.

Nearly every assumption listed in the engineering axioms is violated by complex systems. So, what leads us to believe that we can use the insights of classical engineering to predict and control these systems? The recent historical record makes the case rather clear: we have not succeeded in this approach.

This does not mean that there is not enormous value in engineering design, or that we should give up. There is an alternative approach that respects in its own axioms the properties of complex systems and builds on the insights of classical engineering—we call this approach emergent engineering.

The Objectives of an Emergent Engineering

1. Seek to modify the reward or selective context in which semi-autonomous agents operate and design toward better incentives.
2. Accept significant component error rates and focus on mechanisms that can average and aggregate these effects to acceptable levels in the collective output.
3. Design with an eye towards distributions of outcomes and not towards deltas (single optimal outcomes), pursuing average properties throughout.
4. Develop mechanisms for controlling nonlinear dynamics and predicting and influencing critical transitions.
5. Harness adaptation to allow for continued exploration and exploitation rather than coercing systems into single states that require endless iterations of costly novel production.

This is obviously a very ambitious list, and yet there are both established and nascent areas of research, several coming from within engineering, where ideas with these characteristics have been—and are being—developed. These include nonlinear control theory, evolutionary dynamics, genetic algorithms and programming, pattern formation, agent-based modeling, mechanism design, collective computation, programmable and adaptive matter, and distributed token systems such as the blockchain.

There are three very successful areas where these and related ideas have been applied by recruiting and not obstructing the properties of complex adaptive systems: vaccine development, auction design, and neural prosthetics.

Co-opting Complexity for Emergent Engineering

A vaccine is a chemical agent that provokes a sustained reaction from the immune system of a host directed at pathogens with features similar to the vaccine. Vaccines achieve emergent engineering by exploiting the noisy, distributed, nonlinear adaptive properties of immune systems. Vaccines are effective precisely because the contribution to immunity engineered into a vaccine is negligible in comparison to the host mechanisms that it engenders and that it exploits.

An auction is a mechanism for meeting the market needs of heterogeneous buyers and sellers without prior knowledge of pricing. The auction mechanism exploits uncertainty and the adaptive and strategic self-interests of buyers and sellers to achieve a desired distribution. Auction mechanisms are effective because they are simple and shift the burden of information processing and decision making to populations of agents and away from overly complicated, over-regulated, centralized valuation mechanisms.

Neuroprosthetics are electrical brain implants that are able to compensate for the loss of essential cognitive functions to include perception and movement. The effectiveness of neuroprosthetics relies on the brain's ability to learn the inputs and outputs of the implant. An effective implant is one that exploits the noisy adaptive nature of neural tissue such that the restoration of function is largely a matter of ongoing brain rewiring and recoding and more modestly reliant on the prewiring/pre-encoding of the prosthetic.

Each of these examples illustrates engineered solutions working cooperatively with a complex system in order to predict and control a behavior rather than shoehorn complexity into regimes where the classical engineering axioms hold sway. Success depends on a deep understanding and application of the properties of complex systems in a hybrid agent–artifact setting.

Next Steps

Engineering-inspired paradigms from genomics to synthetic biology and circuit-based neuroscience, upwards into market regulation and cybersecurity, need to continue to evolve by merging in part with the adaptive sciences to embrace complex reality and communicate the practices and implications of this new approach to the critical challenges of the modern world. The disruption of the disciplines and the growth of respect fostered across the spurious borders of the institutionalized sciences—departments and schools—would be a welcome and necessary corollary.

Perhaps the greatest "phase transition" in our thinking that such an approach could engender is the maturation in our willingness to live with relatively high levels of uncertainty in the domains of complex phenomena—and thus give up on ideas like complete "cures," the elimination of "risk," the design of perfect "stability," and achieving total "security." We replace these ideals of a deterministic age with an understanding of the ever-evolving nature of adaptive processes, seeking to discover new methods for the specification of incentives, rewards, constraints, and communication, together capable of moving outcomes into a space of desirable, albeit never optimal, performance.