Pervasive Flexibility in Living Technologies through Degeneracy-Based Design

James Whitacre*,**
Vrije Universiteit Brussel

Axel Bender†
Defence Science and Technology Organisation
University of New South Wales

Abstract The capacity to adapt can greatly influence the success of systems that need to compensate for damaged parts, learn how to achieve robust performance in new environments, or exploit novel opportunities that originate from new technological interfaces or emerging markets. Many of the conditions in which technology is required to adapt cannot be anticipated during its design stage, thus creating a challenge for the designer. Inspired by the study of a range of biological systems, we propose that *degeneracy*—the realization of multiple, functionally versatile components with contextually overlapping functional redundancy—will support adaptation in technologies, because it effects pervasive flexibility, evolutionary innovation, and homeostatic robustness. We provide examples of degeneracy in a number of rudimentary living technologies, from military sociotechnical systems to swarm robotics, and we present design principles—including shared protocols, loose regulatory coupling, and functional versatility—that allow degeneracy to arise in both biological and man-made systems.

Keywords
Pervasive flexibility, adaptation, degeneracy, living technologies, distributed robustness

A version of this paper with color figures is available online at http://dx.doi.org/10.1162/artl_a_00116. Subscription required.

1 Introduction

Unanticipated requirements can arise throughout a technology's life and are a notoriously difficult engineering problem and a challenging research topic, because past routines and contingency plans will be of limited utility. Dealing with new challenges requires exploration, diversity, and bet-hedging: principles that are common to any discipline in which responses to novelty determine competitive success.

However, these conceptualizations of adaptive behavior provide only approximate descriptive accounts of how complex technologies can be designed to achieve robustness and adaptation in novel circumstances. Importantly, it is still poorly understood how adaptive options are generated and exploited within a systems context without sacrificing other objectives related to efficiency and effectiveness. Recent theories have proposed that a biological property known as degeneracy plays an important role in establishing synergistic relationships between biosystem complexity, adaptive potential, robustness, and efficiency [48]. In this article, we will explore how the development of degeneracy in living technologies [5] may help these systems acquire desirable adaptive features. We will propose a

* Contact author.
** Center Leo Apostel, Vrije Universiteit Brussel, Brussels, Belgium. E-mail: jwhitacre79@gmail.com
† Land Operations Division, Defence Science and Technology Organisation, Edinburgh, SA, Australia; School of Engineering and Information Technology, University of New South Wales, Canberra, ACT, Australia. E-mail: axel.bender@dsto.defence.gov.au

precise form of diversity, using degeneracy-based design principles, that can support distributed robustness in a number of social and technological multi-agent systems. While not a complete recipe for realizing the living technology vision, the principles emphasized here have been chosen because they are ubiquitous in evolvable biosystems, are mostly absent in fragile and non-adaptive designed systems, and are notably missing in all but a few discussions that address the requisite conditions of artificial life [10], artificial intelligence [51], and living technologies [23, 52].

2 Design and Development of Flexible Systems

2.1 Summary

If flexibility only exists in places where we perceive future need, then our resilience will be limited by our foresight. Because novel requirements are not predictable, flexible responses to novelty cannot be entirely prespecified. Instead, flexibility must be a pervasive property that emerges on demand without explicit planning or foresight.

2.2 Introduction

Researchers and engineers from largely isolated disciplines have uncovered similar principles that contribute to the design or evolution of adaptive systems and appear to be widely applicable within ecosystems, biochemical networks, systems engineering, and human organizations [12, 13, 30, 42]. Primary factors that contribute to resilience of food webs, canalization of multicellular development, physiological homeostasis, and robust control of automated manufacturing processes include intuitive engineering concepts such as functional redundancy, bet-hedging, saturation effects, and fail-safe principles.

Feedback control concepts have been particularly successful in explaining robustness in a wide range of systems, including biological networks [4, 9, 33, 53]. These findings have received considerable publicity in biology, in part because they support the view of a top-down deconstruction of biological complexity that could reveal intuitive insights into the exceptional robustness and adaptive potential of living systems [14].

Less attention has been given to evidence that the majority of robustness in biological systems cannot be attributed to simple feedback loops or component redundancy [12, 13, 19, 34, 43, 46, 47, 50]. Even though control theoretic principles are relevant to biological robustness, the sets of components displaying single-reference-point, closed-loop control principles only do so within tightly restricted (microenvironment) conditions. It is more often the case that numerous interdependent actions among diverse biological elements are integrated over space and time, with robustness emerging in a distributed fashion. In biology, this phenomenon is referred to as distributed robustness [47, 50] or emergent flexibility [52].

Biological and ecological research on distributed robustness has uncovered statistical patterns of regulatory (activation/inhibition) and mass-action interactions that are positively correlated with robustness, including nested feedback loops, bow-tie architectures, and long-tail distributions of regulatory interactions [15, 43, 45]. For instance, in the immune system [29] and metabolism [15], distributed robustness is facilitated in part by a multi-scaled bow-tie architecture: At many scales of the system there exist multiple pathways to achieving a given function or effect. The result of these multiple pathways is that the system is endowed with exceptional flexibility when operating under stressed conditions. Interestingly, these pathways are compensatory but not entirely redundant: In many circumstances they contribute to entirely different functions. As a simple example, the metabolism of glucose can take place through two distinct pathways: glycolysis and the pentose phosphate pathway. Although these pathways can substitute for each other when necessary, the entirety of their metabolic effects is not identical.

Distributed robustness can emerge in similar ways within human organizational contexts such as military forces. For instance, in traditional warfare, the tactical advance of an adversary can be suppressed by employing artillery, cavalry, air strike assets, or other force elements. A deployed military force in need of suppressing an adversary's advance therefore has multiple courses of action (COAs) available. These COAs represent different utilizations and orchestrations of the same deployed force elements. While different COAs may generate the same effect, for example, during the suppression of an adversary's advance, they differ in other tasks, such as the own force's tactical mobility. When circumstances change in an unanticipated manner, a deployed force can quickly readjust the utilizations and orchestrations of force elements and shift from the COA executed so far to another that promises success under the changed new conditions. In short, military adaptive capabilities arise from force elements that compensate and complement each other. This allows for a changeable organizational form that emerges in response to deployment contexts of large (and dangerous) uncertainty.

2.3 Need for Pervasive Flexibility

Within sociotechnical systems, it is not controversial to assert that adequate responses toward novel internal and external stresses generally require flexibility in what/when/where actions are taken by a combination of human, hardware, and electronic assets. However, because the what/when/where of novel requirements is fundamentally unpredictable, the flexibility needed to respond to this novelty cannot be prespecified on the basis of the anticipation of future conditions. Instead, the adequate provision of flexibility requires it to be a pervasive system property that can emerge without explicit planning or foresight. Importantly, if flexibility only arises in the places where we perceive future need, then resilience will be limited by our foresight, that is, our ability to predict plausible future scenarios [52].

A rich history of engineering and planning experience suggests that pervasive flexibility will be prohibitively costly and impractical, due to the inefficiency of idle redundant resources (robustness-efficiency tradeoff). The alternative—a situation where flexibility emerges on demand without requiring idle/backup resources—is hard to imagine in a technological context. In other circumstances, flexibility requires diversity in options, not redundancy. However, there are management overheads arising from diversity that add to system complexity and should slow down adaptation in large systems (complexity-adaptation tradeoff). Drawing from these perspectives, pervasive flexibility is assumed to be possible in nature only because natural selection permits inefficiency and a slow pace of adaptation. The problem with these assumptions is that they are inaccurate and misleading: Biological systems evolve within highly competitive and resource-constrained environments, and rapid evolutionary change is common in even the most complex species [39]. Biological systems are able to bypass or partly resolve the conflicts between robustness, efficiency, complexity, and evolvability that limit technological capabilities [48].

Although survival and fecundity are not perfect analogs to market-based forces, some researchers believe that the similarities are sufficient in some circumstances to warrant research into the nature-inspired design of artificial systems. Recently we proposed a theory to explain how pervasive forms of biological flexibility are achieved at high levels of efficiency through a property known as degeneracy [50]. Degeneracy refers to a unique situation where groups of agents exhibit functional redundancy in some contexts but functional diversity in others.

3 Contributing Factors for Adaptation in Biological, Social, and Engineered Systems

To understand how degeneracy can achieve the proposed effects, it is helpful to first discuss the typical options that are available to a system that is responding to novel conditions. As listed in Figure 1 and Table 1, response options typically involve one or a combination of the following:

1. reducing exposure to unwanted conditions by manipulating the environment;
2. reducing exposure to unwanted conditions through mobility within the environment;

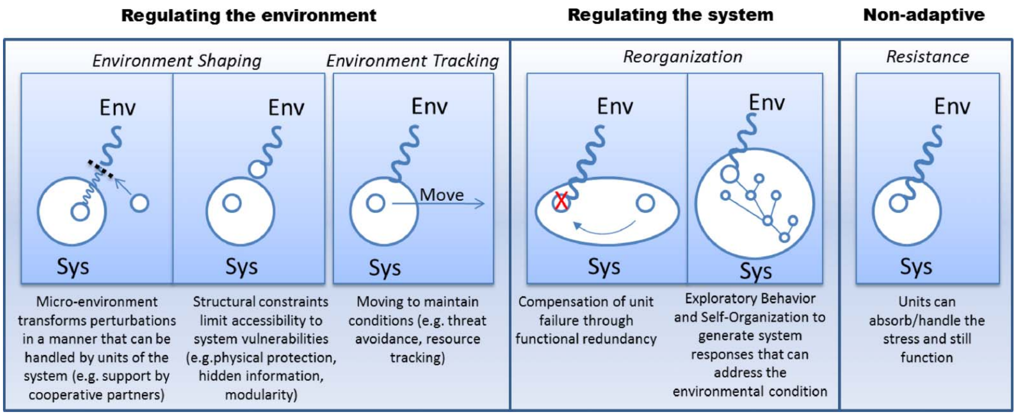

Figure 1. Four pathways for adapting to environmental stress: environmental shaping (regulate the environment so fragilities are not accessed/revealed); environmental tracking (move to new environments where performance can be maintained); reorganization (adapt system response in order to sustain or improve upon system objectives); resistance (design high-quality components that resist damage from perturbations).

3. adapting system behavior in response to environmental change in order to improve, maintain, or repair a capability output;
4. designing high-quality components that are individually robust toward stress and damage. In this case, no adaptive response is required.

Each of these options may invoke prototypical examples in the reader's mind; however, each is widely applicable across many systems contexts. Response options are a prerequisite for adaptation and involve one or a combination of the following actions:

1. *Changes in how much, when, and where resources are needed:* This form of adaptation requires options to quickly change the quantity of a particular functional output at a particular place and time. Excess backup resources can support this type of adaptation; however, idle resources reduce average efficiency and thus can be costly.
2. *Changes in task specifications:* Unexpected conditions sometimes require a function to be executed in a manner that deviates from the norm. Maintaining diversity in the options for executing a task, with each option displaying unique vulnerabilities, can provide reliability in the face of novel requirements. Option diversity is typically not random and instead reflects an accumulated knowledge of expected disturbances. For instance, bet-hedging strategies drawn from portfolio theory are used in several disciplines to reduce the likelihood of large systemic risks against known uncertainties [22, 41, 44].
3. *Functional novelty (exaptation):* New environments can reveal opportunities to utilize existing components in novel ways: a class of adaptation that is known to biologists as exaptation [25]. Maintaining diversity in options, assets, or agents can improve the likelihood of discovering exaptation opportunities. Successful R&D departments (e.g., at 3M, Apple, Google) are often in the habit of exploring potential exaptations for existing products or organizational competencies.

In summary, adaptation involves several distinct requirements, in some cases demanding quantitative changes in functional outputs through redundancy, while in other cases demanding qualitatively diverse options to sustain functional capabilities in novel environments or to explore new capability opportunities. In systems engineering and organization science, these requirements for adaptation have been addressed through largely separate programs involving redundancy, bet-hedging, exploration, and related strategies (see Figure 1 and Table 1).

Table 1. Factors that contribute to robust system performance under volatile conditions.

Description of mechanisms that enhance robustness	Biological examples	Engineering and management science examples	Military examples
Local environment shaping/regulation: Instead of achieving robustness by responding to environmental stress, it is sometimes possible to shape the environment in ways that allow a system to avoid exposure to damaging stress.	Niche construction and environment simplification alter the type and frequency of perturbations encountered. Heat shock proteins (e.g., Hsp90) assist other proteins to fold and refold into functional conformations, confer conformational robustness toward thermal fluctuations, and canalize a broad range of morphological traits. Localization of harmful pathogens through tissue inflammation or through ingestion by macrophages.	Monitoring and controlling subsystem operating environments can reduce exposure to damaging perturbations (cooling fans in computers, dust-free microchip production facilities, refrigerators). Fail-safe principles can dynamically encapsulate subsystems and prevent failures from propagating into expensive devices and systemritical operations (fuses, emergency power supplies, e.g., in hospitals)	Armor, secure/safe zones, bunkers provide protection to otherwise vulnerable assets. Smoke grenades obscure an adversary's line of sight and lower the likelihood of being acquired by the enemy's target prosecution assets. Engagement with local communities reduces their willingness to support insurgents and improves counterinsurgency through enhanced intelligence on insurgents' organization, stronghold, planned actions, etc.
Mobility: Having the ability to move or be moved into environments can enable functions to be achieved when conditions demand them or to be relocated when hostile conditions develop.	Predator avoidance, adaptive foraging, migration, and seed dispersal all provide options for populations to seek out and track suitable habitats.	Having mobile parts can help trace favorable environments (e.g., actuators that rotate solar panels toward the sun to maximize energy production) or reduce exposure to unfavorable conditions.	A considerable amount of military hardware has the express purpose of providing mobility in various circumstances. For instance, combat and assault bridges enable the rapid crossing of natural and man-made obstacles, thus enhancing operational mobility.

Table 1. (*continued*)

Description of mechanisms that enhance robustness	Biological examples	Engineering and management science examples	Military examples
		Service-oriented businesses with multimodal delivery methods for their services. Business models that provide access to multiple markets.	
Reliability through functional and pathway redundancy: Distinct components/ pathways that are functionally similar provide robustness against the loss of components.	Gene regulation, protein functionality, metabolic pathways, signaling pathways, and neural anatomy can be highly degenerate, meaning that distinct components/processes/ pathways can sometimes elicit compensatory effects, functional redundancy, and compensatory repair/buffering of system characters.	Empirically driven placement of backup devices as well as storage/maintenance/ preservation facilities can buffer against fluctuating operating conditions or component failure. Feedback diagnositics for preventative maintenance to replace components before they fail.	Excess equipment, personnel, supplies, etc., are deployed into areas of operation in order to provide redundancy when harsh operating conditions or enemy actions cause losses. Military forces are task-organized and can adjust the utilization and orchestration of force elements such that shifts from one COA to another can be executed smoothly and rapidly.
Resistance: Robustness of component toward variable conditions removes need for any system-level response.	Many types of threshold effects in biology appear as subsystems with innate (albeit bounded) resistance to change (e.g., genetic switches, TCR-mediated activation of T cells, neural activation).	High cost ultraquality components with lower rates of failure can provide reliability in circumstances where replacement is impractical.	Military forces are designed to be resistant, e.g., to hold a position when under attack or to survive in the direct fire zone. Most military equipment is ruggedized to withstand rough usage and harsh conditions.

4 Degeneracy

The distinct forms of adaptation discussed in the previous sections are each partly supported in biological systems through degeneracy. Degeneracy is a property seen in repertoires of multifunctional agents when some of the agents are functionally interoperable for certain types of tasks but uniquely functionally qualified for others (Figure 2).

Degeneracy is not restricted to biological systems and can be easily seen in many complex adaptive systems. Conceptual illustrations of degeneracy are given in Figure 3 for small and large defense systems, including the design of land mobility capability, army operations, and multinational alliances. Additional biological and human organizational examples are listed in Table 2.

5 Logical and Intuitive Benefits from Degeneracy

Some of the desirable properties that arise from degeneracy can be related to simple concepts such as functional redundancy, bet-hedging, and exploratory behavior. One simple way degenerate components contribute to adaptation is through component multifunctionality. By being able to contribute to a variety of tasks, multifunctional components can change what they do and contribute to system responses involving quantitative changes in functional outputs. Simply stated, multifunctional components can engage in one particular function if more resources for that function are needed, or be reassigned to one of its other functions if fewer of those resources are needed. Thus, multifunctional agents can support adaptive responses to changing task requirements (Figure 4a,b).

Components that are degenerate can also support bet-hedging and exploratory adaptive responses. First, degenerate components are functionally interchangeable for certain tasks but achieve these tasks in different ways. These differences can lead to differences in agent performance when tasks are executed under novel or rare conditions. Because of this context-revealed *response diversity* (Figure 4d), performing a particular task under unexpected novel conditions is more likely to be possible by a repertoire of degenerate than of redundant components. In a similar vein, degenerate components can harbor somewhat distinct vulnerabilities, thus increasing the likelihood that at least one component will not fail when confronted with a novel or rare disturbance, and thus providing a basic form of bet-hedging (Figure 4d).

Novel environments sometimes reveal opportunities for a component to be coopted to perform a new beneficial function. With behavioral differences among degenerate components, each component harbors a contextually unique potential for cooptation. A group of degenerate elements thus provides greater opportunities for exploring innovative capabilities and for responding to novel functional requirements (Figure 4e).

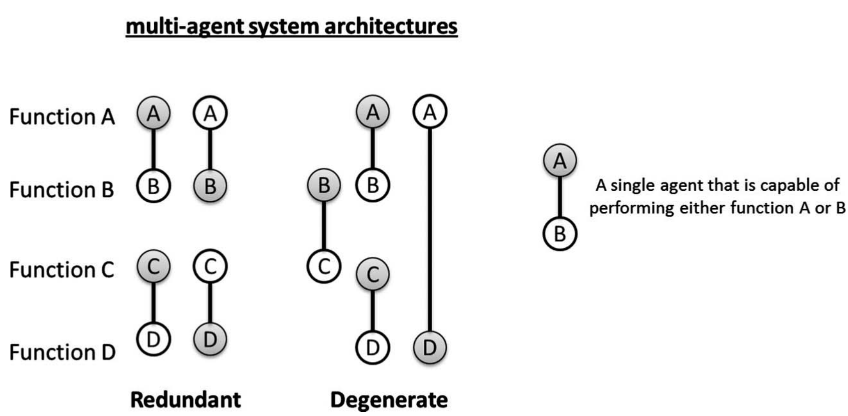

Figure 2. Bifunctional agents that are either (purely) redundant (i.e., perfectly identical in functional capabilities) or degenerate (i.e., diverse in their bifunctionality while also having overlap in one of their functions, i.e., partial redundancy).

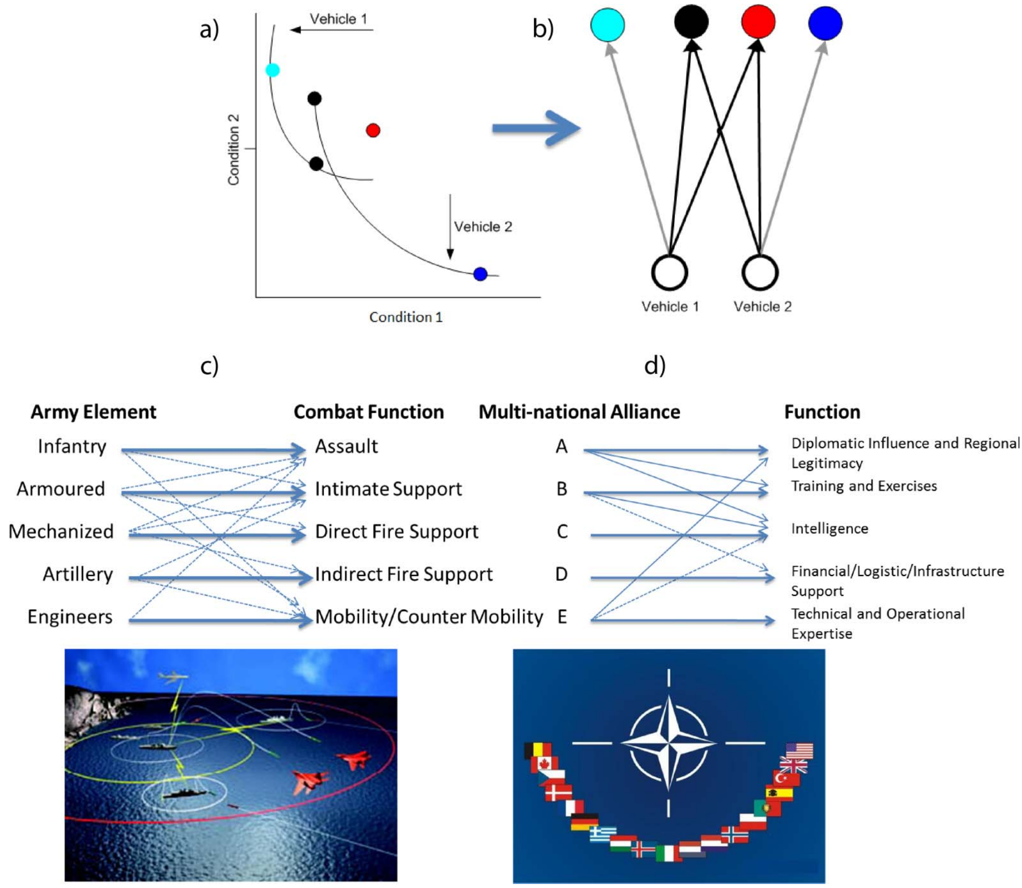

Figure 3. Examples of degeneracy at different levels of organization in defense. (a) Different vehicle types have unique tradeoffs that define the conditions under which they can operate. The continuous tradeoff surface (arcs) conceptually illustrates a vehicle's operational range. In the online version, the red node typifies a set of conditions under which either vehicle type is suitable. The black nodes represent a task where condition 1 is a hard constraint and condition 2 is a soft constraint. Either vehicle can perform the task, but at different performance levels. Blue nodes are tasks where each vehicle is uniquely qualified. (b) Vehicles 1 and 2 from (a) are degenerate, that is, under certain conditions the vehicles are interoperable, while in others the vehicles are uniquely qualified. (c) Degeneracy in the relationship between army elements and combat functions. Elements have prime roles (thick arrows) but can also support other combat functions (dotted arrows). Every combat function can thus be performed by more than one element, resulting in both resistance to enemy actions and the ability to reorganize and execute a diverse spectrum of COA. (d) Degeneracy in the relationship between nations and strategic capabilities. In multinational alliances, individual nations (A, B, …) typically lead the effort in a strategic domain (thick arrows) but also support other efforts (dotted arrows). This results in a network of assignments that enhances the alliance's resistance (e.g., to changes in membership) and capacity to innovate (e.g., when "proven" methods do not generate the desired outcomes).

5.1 Emergent Benefits from Degeneracy

Degeneracy also supports emergent forms of distributed flexibility that could be relevant to the adaptive capabilities of living technologies. First, when multifunctional components are interoperable in a subset of their functions, fluctuating task requirements can cause interoperability options to become synergistically linked and result in a basic form of distributed robustness. An example of this synergistic linkage is illustrated in Figure 5a. In the figure, agent B can perform either task 2 or 3. If agent B has no task assigned to it, then it is available to take over tasks (of type 2 or 3) assigned to agent a or c. This allows agents A and C to be available for tasks (of types 1 and 4) that agent B could not carry out. In other words, resources of agent B do not only support adaptation toward variable demands in tasks 2 and 3, they can also indirectly enable new resources to be available for tasks 1 and 4, tasks that are

Table 2. System classes where agents are multifunctional and have functions that can partially overlap with other agents. Degeneracy is observed in each case through the conditional similarity of functional capabilities.

Agent	System	Environment	Control	Agent tasks
Vehicle type	Transportation fleet	Transportation network	Centralized command and control	Transporting goods, pax
Force element	Defense force structure	Future scenarios	Strategic planning	Missions
Person	Organization	Marketplace	Management	Job roles
Deme	Ecosystem	Physical environment	Self-organized	Resource usage and creation
Gene product	Interactome	Cell	Self-organized and evolved	Energetic and steric interactions
Antigen	Immune system	Antibodies and host proteins	Immune learning	Recognizing foreign proteins

unrelated to agent B. More generally, partial interoperability allows excess resources related to a particular task to support fluctuations in unrelated tasks. As a result, small amounts of local excess resources can be utilized in a versatile manner, thereby increasing the variety of task fluctuations to which a system can respond. We have shown in [50] that this simple effect can fundamentally alter the tradeoff between robustness and efficiency.

In particular, we have discovered that some architectures allow systems to respond to enormous task fluctuations, a phenomenon that we captured in the *networked buffering* (NB) hypothesis [50]. According to NB, if partial interoperability relationships form a connected network (Figure 5c), then

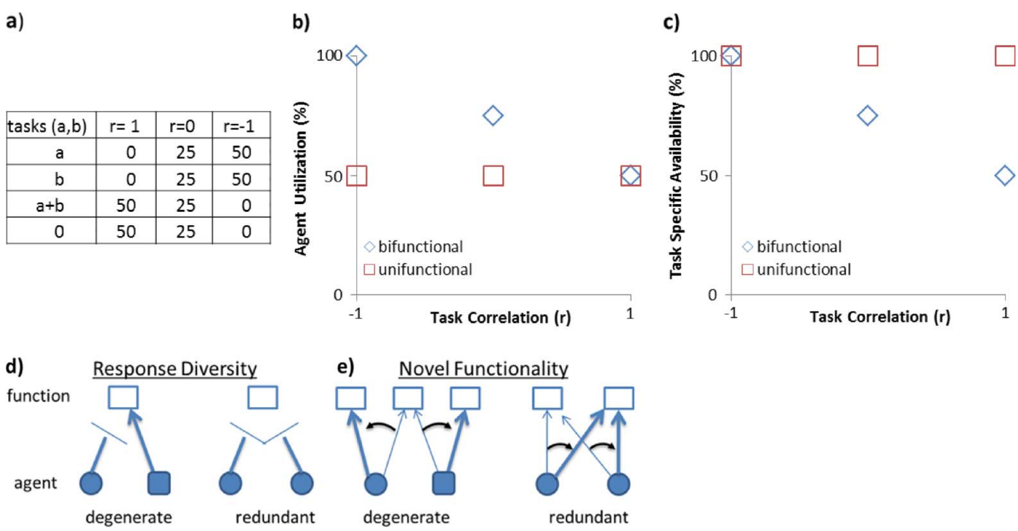

Figure 4. Multifunctional components are less likely to sit idle under fluctuating task requirements. (a) Probability of events where each task (*a*, *b*) has a 50% probability of occurring and tasks are 100% positively correlated ($r = 1$), uncorrelated ($r = 0$), and 100% negatively correlated ($r = -1$). (b) Expected utilization rate for bifunctional ($a + b$) and unifunctional (a or b) agents. (c) Expected availability for a specific task. (d) Structural differences can enhance reliability in executing a function. (e) A greater variety of novel functions can be revealed when degenerate agents are placed in new environments.

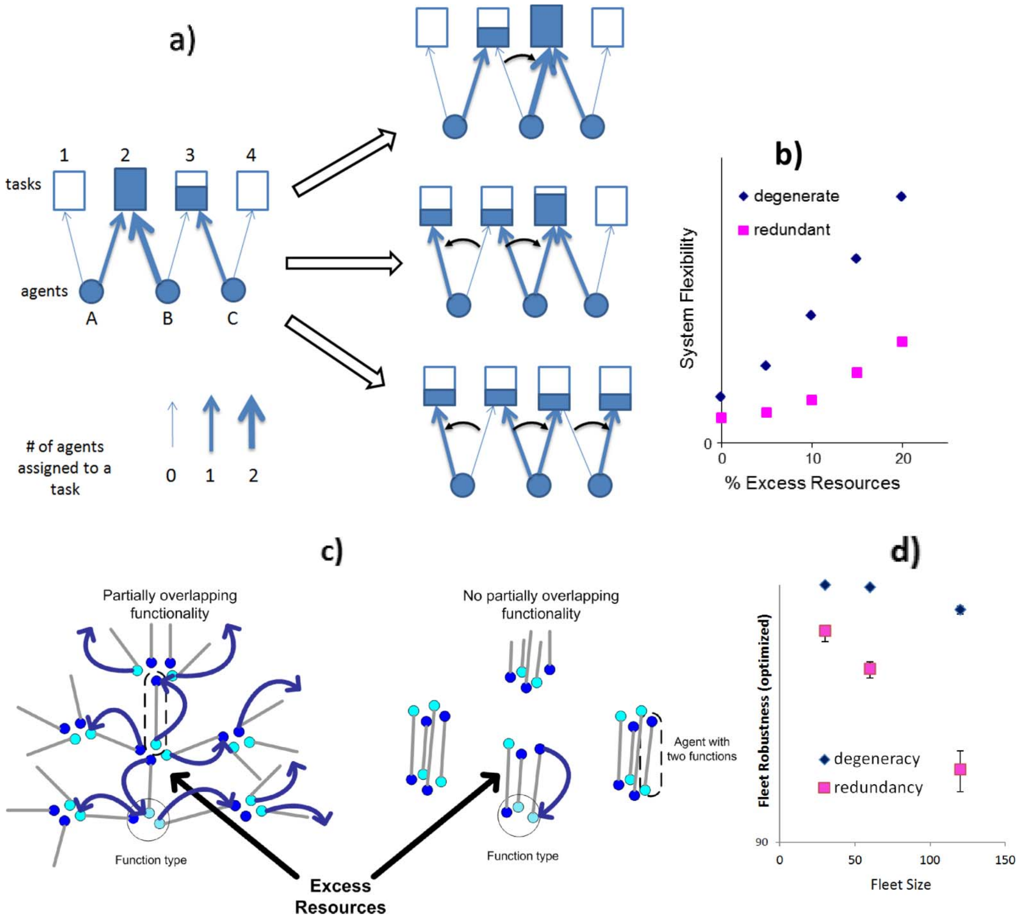

Figure 5. (a) The presence of unutilized type B agents can directly buffer fluctuations for task types 2 and 3. Type B agents are partly interoperable with type A and C agents, and thereby can free up resources for changes in requirements for task types 1 and 4. (b) Flexibility conferred in random (non-optimized) protein affinity networks with increases in protein expression (excess resources) for protein affinity networks with and without degeneracy. Taken from [49]. (c) Bifunctional agents designed to be partially interoperable (degenerate) with high flexibility conferred through a degeneracy backbone (left diagram), and fully interoperable (redundant) agents that are operationally isolated within functional clusters (right diagram). Adapted from [48]. (d) Robustness in fleets of multifunctional vehicles that were optimized to be maximally flexible towards a variety of mission scenarios. Quantitative differences in scenario task requirements are scaled proportionally with fleet size. Adapted from [52].

the repertoire of response options towards different task requirements can become very large. This is qualitatively illustrated in Figure 5c and quantitatively validated in multi-agent simulations in Figure 5d. In contrast, when multifunctional components are designed/trained so that functions are clustered within organizational stovepipes, a multiplier effect is not observed (Figure 5b,d). The distributed robustness just described is arguably seen in electron transport within large proteins [32], in bipartite protein-ligand interaction networks [50], in the mutualistic and trophic interactions of ecosystems [50], and in human organizations and military field vehicle fleets [52].

For instance, in evolutionary computation simulations, we explored how degeneracy-based design affects a military field vehicle fleet's response to anticipated variability in mission requirements. When partial interoperability in vehicle task capabilities was permitted to arise in the fleet design, we recorded improvements in fleet performance in anticipated scenarios (Figure 5d) and greater robustness toward stresses that were not planned for during fleet development [52]. Interestingly, adaptive responses that were optimized to address specific mission variations were also found to support the emergence of pervasive flexibility options that could arise on demand and be coopted ad hoc to respond to conditions not previously encountered or planned for. We showed in [50] that these new adaptive capabilities are largely due to the NB effect (as illustrated in Figure 5c).

In [23] we proposed that similar principles can be applied to agile manufacturing processes involving semi-autonomous robots. Although the NB effect increases the complexity of resource allocation decisions, we also found that the pervasive flexibility from NB endows a system with many high-performance configurations. This desensitizes the system to suboptimal local decisions, thus alleviating conflicts between the system's adaptive response speed and the diversity offered by the response repertoire. In other words, NB partially resolves the robustness-efficiency and complexity-adaptation conflicts that plague engineered systems but are less prominent in biological systems.

One possible reason that the degeneracy-based design approach is rarely tested in man-made systems is that multifunctional agents must share their time across several related functions and, unlike unifunctional agents, are therefore not reliably available for a specified task function (Figure 4c). This may be considered disadvantageous in circumstances where it is desirable to tightly control the utilization of excess resources. Note, though, that groups of partly interoperable agents counterintuitively display improved functional reliability at the group level as a result of resource allocation flexibility. This flexibility and associated reliability have been found to grow rapidly (and at a faster rate than in purely redundant systems) with gradually increasing excess resources (Figure 5b).

While the distributed flexibility we have described so far only relates to quantitative changes in a system's designed functional outputs, it also becomes essential when exploiting exaptations or adaptations from response diversity. For instance, in systems optimized for efficiency, selecting an adaptive option from a repertoire of diverse resources can reduce the availability of these coopted resources for their original tasks, thereby creating internal stresses on the what/when/where of a system's other operational outputs. In such circumstances, task assignment flexibility becomes an important foundation for these higher-level adaptive capabilities.

6 Recommendations for Enabling Degeneracy

6.1 Summary

Shared protocols, agent versatility, and loose coupling constitute a set of quantifiable design principles for realizing degeneracy and the emergence of pervasive flexibility in living technologies.

6.2 Introduction

Degeneracy and network buffering architectures can be incorporated into living technologies through clearly definable system features. These features have evolved in biological systems over long periods of time through major evolutionary transitions and have become ubiquitous in present-day species, particularly in multicellular eukaryotes [28]. In contrast to their possibly fortuitous development in biological contexts, we contend that they can be intentionally selected, designed, and encouraged. By enabling the systematic development of degeneracy, these properties support the development of living technologies that flexibly respond to unplanned changes at all scales of a system—from operational environment and internal design to user preferences and competitive marketplaces.

As a design principle, degeneracy can be realized in technological systems that exhibit the following features:

1. Shared protocols
2. Agent versatility
3. Loose regulatory coupling

An overview of these design features is provided in Figure 6. In the following sections, we describe each feature while ignoring their biological and engineering motivations, their history, and the technical jargon that pervades the research informing this discussion.

6.3 Shared Protocols

Plug-and-play compatibility provides ample opportunities for communication and other interaction among technological artifacts. This supports the fortuitous discovery of new service combinations and the occasional reorganization of networked services to reveal novel capabilities.

Such combinatorial flexibility is achieved in part by requiring agents to adhere to *protocols*. Protocols are standard procedures or "rules of engagement" [18] that specify conditions that must be met in order to execute a particular task or elicit a particular behavior in other agents [14, 17, 18]. Protocols enable nontrivial interactions among agents with one agent needing to know little (or even nothing) about the internal operations of the others. Instead, small amounts of information sharing between agents can be used to inform and guide elaborate patterns of action.

In agent-agent interactions, protocols manage relationships by defining rules for how the behavior of an agent is influenced by others. Although protocols constrain how agents can interact and communicate, they also deconstrain the potential for new collaborations, because any agent can plug and play once it meets the protocol criteria. Within the context of the task-centered conceptual framework favored in this article, protocols can be viewed as an agreed standard for the successful completion of tasks. From this perspective, protocols provide a design principle by which diverse components can engage in functionally similar activities (functional redundancy) and thus help to manage a key requirement for the realization of degeneracy and subsequent networked buffering.

There are many examples of protocols that arise in different systems contexts. In biology, protocols have evolved and are seen in the energy currencies of metabolism, in activation-potential-enabled cell-cell communication of neurons, and in the universal usage of nucleotide sequence codons in gene transcription. In informal social systems, protocols emerge through acceptance of cultural norms that spread like viruses over socially connected and susceptible segments of a society. In technological systems, protocols are often explicitly established during systems design, for example, the Internet's TCP/IP protocol stack.

The role of protocols in agent-based collaborations is not restricted to direct interactions. For instance, the manipulation of shared environmental artifacts using standards of manipulation (known to biologists as stigmergy) can provide cues for actions that are taken by other agents in the environment.

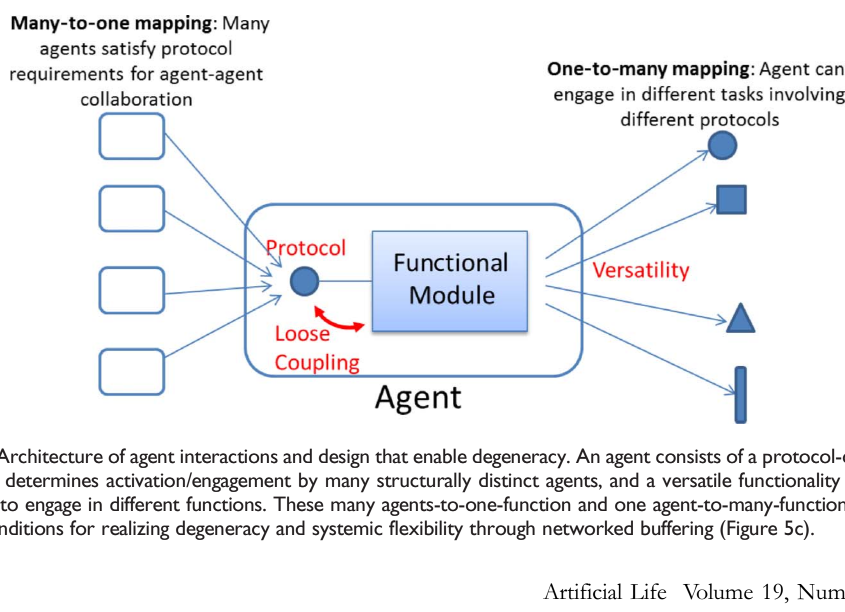

Figure 6. Architecture of agent interactions and design that enable degeneracy. An agent consists of a protocol-constrained input that determines activation/engagement by many structurally distinct agents, and a versatile functionality that allows the agent to engage in different functions. These many agents-to-one-function and one agent-to-many-functions mappings are preconditions for realizing degeneracy and systemic flexibility through networked buffering (Figure 5c).

6.4 Agent Versatility

Versatility means the ability to competently perform a range of partly related tasks or functions. When agents are functionally versatile, functions are invoked based on the agent's current state and cues from its local surroundings. Within a systems context, functional versatility might allow an agent to collaborate with or support ("plug into") a variety of different agents.

The behavior of versatile components is influenced by environmental cues and therefore can be responsive toward changing task requirements, for example, changed requirements in operational outputs, replacement of degraded units, or replacement of units that have been assigned to other tasks. Pairs of versatile agents may appear functionally interoperable for certain tasks but uniquely functionally qualified for others, thereby enabling defining attributes of degeneracy to arise; see Section 4. Conversely, degeneracy cannot be observed without functional versatility.

6.5 Loose Regulatory Coupling

Loose regulatory coupling refers to circumstances where the architectural design features that determine an agent's behavioral responses (input protocols) are encapsulated and independent from design features that influence an agent's functional capabilities; see Figure 6. With loose regulatory coupling, design changes to agent functions will rarely require new input specifications or changes to the external cues that motivate agent activity. In contrast, without loose regulatory coupling, agent design changes can alter an agent's protocols for engagement. This may in turn require collaborating agents to modify their behavior, or worse, may require changes in the design of collaborating agents that then propagate new requirements to still other agents of the system.

Loose regulatory coupling is a key innovation in gene regulation and some signaling pathways [7] and has played a major role in the evolvability of biological systems [28]. The lack of loose regulatory coupling increases fragility within interdependent systems toward even small changes in design, because of propagating change requirements in other interacting components. Conversely, loose regulatory coupling enables cooptation of degenerate elements for innovative use in new applications without the propagation of new design requirements, so that it is a vital facilitator of design exaptation by degeneracy.

7 Discussion

7.1 Understanding Complexity in Living Technologies

Complexity is a poorly understood concept in science, in part because it has been assigned multiple imprecise meanings. While the term generally relates to the interdependence of component behavior/actions/functions, it is an otherwise ambiguous term, and there is no consensus as to its meaning or measurement (e.g., [1, 11, 20, 24, 26, 35, 40]). In [3], Alderson and Doyle point out at least three different meanings that are commonly associated with system complexity:

- *Complexity of size:* A large number of diverse components in the system
- *Complexity of interconnection:* A large number of highly specific forms of engagement among components
- *Complexity of interaction:* System components are functionally versatile and interact with numerous other components in a manner that is tailored to the interaction context.

In engineering, complexity often refers to sophisticated services that require interdependent actions of single-purpose devices, each occurring in specific ways and at specific places and times. In other words, engineering complexity often relates to complexity of interconnection. In the absence of redundancy and diversity, interconnection complexity can reduce a system's adaptive potential in a manner that is easy to appreciate.

Starting with a single device, the number and exactness of operational constraints/specifications will restrict the proportion of operating conditions that will meet these requirements. Although the tradeoff between operating constraints and operational feasibility is not necessarily linear or monotonic, the reliability of many multi-device services becomes more fragile to novel internal and external conditions as more components are added that each co-specify the feasible operating conditions of others contributing to the service. In other words, the operating requirements placed on each device become more exacting as its function becomes more reliant on the actions/states/behaviors of others, for example, through direct interaction, through sharing or modifying the same local resources, or indirectly through failure propagation. Services with this interconnection complexity can become more fragile to uncharacteristic component behaviors and atypical events because a greater proportion of events will exceed the operational tolerance thresholds in at least one device, with the propagation characteristics of these threshold-crossing events determining the likelihood of subsystem and system-wide failure. To reduce the frequency of failures, a design approach is sometimes taken that assumes predictability and relies on carefully placed backup devices, the monitoring of intermediate system states, and the incorporation of a variety of fail-safe procedures. However, failures can still be prevalent in such systems, as discussed in detail for Xerox photocopiers [2] and for some large organizations that have experienced rapid change (e.g., DuPont [8]). While design principles such as modularity and loose coupling can help reduce the size and frequency of some failures, adaptation processes containing repertoires of system response options are still essential for achieving reliable performance under unexpected conditions.

The same conditions that limit operational robustness toward unanticipated events can also place limits on the adaptability of system design. When systems are designed from single-purpose devices that are each uniquely suitable for a system-critical function, a tight coupling is established between system performance, the reliability of a function, the continued normal operation of the device providing that function, and the continued compatibility of that device with other interacting devices.[1] Novel redesign of devices is thus constrained by a need for highly specific forms of engagement with other devices. With engineering driven to maximize efficiency and performance, small design adjustments are repeatedly made over time to improve efficiency under standard operating conditions. Thus, the system's design becomes well adapted to a specific and well-controlled environment. As a system's design matures (evolves) in this way, there are fewer alternative system configurations for achieving each given task and fewer degrees of freedom for modifying a system's design without compromising function.

If the environment or system priorities change substantially, a need for system redesign arises and a lack of useful redesign options can create tension that grows over time. Eventually a failure to meet system-level goals can make re-engineering unavoidable. However, given the redesign constraints just discussed, design modifications become necessary in many components simultaneously—a phenomenon that is ubiquitous in complex engineering artifacts and well illustrated by the dramatic redesign of complex software such as operating systems. A large re-engineering effort often runs roughshod over the accumulated and highly contextual knowledge that was built during the system's earlier maturation, causing many large re-engineering and change management projects to appear as failures when compared with prior system performance.

7.2 Complexity in Biology

Highly sophisticated services also exist in biological systems that require many different subfunctions and process pathways to be executed. However, the building blocks of biological systems are not single-purpose devices with predefined functionality and instead display considerable overlap in function, functional versatility, and degeneracy.

While the occasional slowdown in the tempo of adaptation is inevitable and occurs in biological evolution as well (e.g., under stabilizing selection), there is little evidence to suggest that biological systems experience the same built-up tension from gradual changes in the environment or the same sensitivity to incremental design changes. We believe this is because degeneracy affords a weaker

1 While functional redundancy is sometimes designed into a system, it is almost always treated as a backup device that is not utilized under standard operating conditions.

coupling between the functions performed and the components involved in achieving them [32]. Within an abstract design space or fitness landscape, one can argue that engineered systems are located on isolated peaks, and large movements in design space are needed so that the systems can find new feasible and viable design options. Biological systems, on the other hand, reside on highly connected neutral plateaus. Although many complexity science researchers have used the metaphor of rugged fitness landscape to advocate the necessity for disruptive and explorative search in the evolution of technologies, this is neither needed nor observed in biological evolution.

In biological evolution, continued species survival requires the discovery of design changes that are incrementally adaptive and do not lead to the propagation of redesign requirements in other components of the system. Thus, macromutation is a negligible contributor to the evolution of complex species. Instead, single heritable design changes lead to novel context-specific interaction opportunities for a component, the flexible reorganization of component interactions (that still maintain core functionalities), and in some cases a subsequent compounding of novel opportunities within the system [31]. In other words, the requirement is one of incremental changes in design and compartmentalized, but not necessarily incremental, changes in system behavior.

In [28], Kirchner and Gerhart present a number of biological examples where this flexible reorganization takes place at cellular and developmental levels. Degeneracy and its associated complexity of interaction (multifunctionality) can support this flexibility through redundancy and diversity and by providing redesign options through networked buffers. In short, based on the arguments outlined here, we propose that degeneracy could play a major role in the realization of pervasive flexibility within living technologies.

7.3 Complexity in Defense Capabilities

Albeit not at the same time scale as biological systems, defense capabilities have also evolved over many generations. The need and the want to defend, potentially with physical force, a society's material, cultural, religious, and moral (among other) values are as old as humankind itself and have resulted in the development of sophisticated warfare equipment, support services, command structures, and military organizations. Combat capabilities—those capabilities that need to operate in a wide spectrum of physical and social environments and are exposed to dramatic fluctuations in external conditions (e.g., when an adversary attacks)—show complexity of interaction that resembles that of biological systems. For instance, a deployed military force can generate indirect fire (e.g., to secure a beachhead during an amphibious landing) through the orchestration of land artillery, naval guns, and/or air strike fighters, involving the elaborate command and control of various systems that interoperate at the operational and tactical levels. This complexity of interaction is enabled through degeneracy-based design:

1. The combat agents (artillery, naval gun, strike fighter, etc.) are functionally versatile in that they can execute other tasks if the context so demands. For instance, the strike fighter can fly reconnaissance missions, engage in air close combat, or execute precise lethal strikes on identified targets (such as bunkers or weapons depots).
2. The interaction between agents is by and large governed by military doctrine, that is, protocols that are shared among all agents that participate in warfare and in other military operations. These shared protocols even allow for indirect agent-agent interaction through stigmergy—for example, when the actions of a subset of the military force are triggered once the environment has been appropriately shaped by the actions of other agents in the force.
3. The agents' behavioral responses are primarily encapsulated in standard operating procedures (supplemented by COA planning) and rules of engagement, which offer fallback options for actions when plan execution arrives at an impasse or, worse, when a plan needs to be abandoned and the deployed force needs to reorganize. Standard operating procedures and rules of engagement ensure that when an agent's design changes (e.g., through attrition of military equipment and/or loss of soldiers' lives), its protocols for engagement do not alter.

Defense sociotechnical systems use degeneracy-based design not only at the tactical level (when defense capability is exposed to potentially high levels of conflict), but also at the strategic level. Modern military forces have a plug-and-play capability that facilitates the generation of task-organized combined arms teams (CATs)—military units that have the resilience, flexibility, and robustness to deal as well as possible with the partially unknown conditions of their next deployment (be it a war-fighting, peace-keeping, or humanitarian assistance operation). The task organization of optimized CATs has to occur on short notice: a modern military force needs to have the inherent ability to assemble diverse CAT options from extant building blocks. In addition, modern forces nowadays have to shift quickly from war-fighting to other lines of operation, such as population protection and/or indigenous capacity building. It therefore comes as no surprise that CAT bricks[2] are degenerate building blocks, that is, they exhibit functional redundancy in some contexts (e.g., all bricks are composed of soldiers who are trained to fight with and without small arms) but functional diversity in others.

While degeneracy-based design is widespread in military sociotechnical systems, designers and decision makers of defense capability rarely reflect on the reasons for this ubiquity. Actually, apart from "it has always been this way," military force structure developers have few arguments at hand to justify force structure designs that many might consider inefficient and uneconomical. The principles described in this article provide a basis for the scientific evaluation of advantages and disadvantages of degeneracy-based design in military sociotechnical systems.

7.4 Swarm Robotics

Many of the aspirational properties of living technologies are also implicit goals in swarm robotics research. Swarms of simple cooperative robots are being studied for their ability to solve complex real-world tasks in which adaptive responses to unexpected conditions and component failure is essential to success. The swarm concept is inspired in part by social insects, for which local information, limited communication, and decentralized control can provide powerful emergent properties and robust problem-solving skills.

Many studies have only considered robots that are identical in form and function. Swarms of these unifunctional robots have been found to produce only a limited repertoire of system behaviors. An important exception is the Swarmanoid Project [16]. The Swarmanoid Project has constructed systems of heterogeneous and dynamically connected autonomous robots. Through functional collaboration, communication, and assembly, a swarmanoid system achieves multi-robot functions that are qualitatively distinct from the capabilities of the individual robots acting in isolation. Three types of robots have been created—eye-bots, hand-bots, and foot-bots (Figure 7). They cooperate in a swarm of dozens of robots to achieve goals involving multiple complex subtasks. Each robot in the swarmanoid system is functionally versatile and interacts with its environment in several nontrivial ways. For instance, the eye-bots are specialized for sensing and analyzing the environment, but can also fly and magnetically attach themselves to a ceiling, thereby expanding the robot's sensing and communication capabilities. Hand-bots are able to use their hands in multiple ways; they can climb, grab, and manipulate other robots. Foot-bots are specialized to move over rough terrain and can transport a variety of objects, including other robots.

The Swarminoid Project provides a proof of principle that swarms of functionally unique robots can cooperate to achieve complex tasks. Protocols for robot-robot engagement enable a plug-and-play architecture that, in principle, can be extended to integrate new robot designs into the swarm collective. It is also conceivable for novel functions to arise from swarm configurations that are guided by new patterns of environmental cues. In short, important forms of operational adaptation and design evolution are attainable in a swarmanoid system.

A video of the swarmanoid system in action won the 2011 AAAI video competition for exciting advances in artificial intelligence (http://www.youtube.com/watch?v=M2nn1X9Xlps [37]). While

2 A *brick* is a small military unit that can perform specific military functions and is operationally viable (i.e., self-sufficient for a certain period of time).

Figure 7. Swarmanoid eye-bot (left), hand-bot being carried by three foot-bots (center), and foot-bot (right). Courtesy of Marco Dorigo.

constituting a remarkable advance over existing swarm systems, the video suggests that the swarm relies on the presence of well-controlled environmental conditions in order to perform its tasks. Based on the principles outlined in this article, we propose a few changes to the swarmanoid system that could improve system flexibility and thereby improve swarm performance in heterogeneous and uncertain environments.

When decomposing the swarmanoid system into its simplest functions (e.g., sensing, moving, climbing, flying), it becomes clear that the components that make up the robots exhibit high levels of redundancy but virtually no degeneracy. The design of each functional component used in the construction of a robot is the result of decisions involving tradeoffs between energy efficiency, physical dimensions, functional range, strength, processor power, and durability. Designing each robot component thus corresponds to the selection of a solution from the Pareto optimal set (POS). As with any multi-objective optimization problem, improvements in one objective are likely to have a negative impact on at least one other objective. Each design option for a component satisfies a range of conditions in which the component can perform a given function (as in Figure 3a for military field vehicles). Comparisons between component design options residing on the POS may in some cases reveal a partial overlap in the conditions in which two designs are interchangeable; this would indicate the presence of degeneracy. Similar functional overlaps can also arise in comparisons between entire robots and multi-robot assemblies.

To realize benefits from degeneracy, multiple distinct components should be added that are functionally redundant under conditions where a function is most commonly needed, while providing unique functional competencies in less common but still important conditions. Unlike the wastefulness of simple redundancy, the components should only remain in the system if they are able to contribute regularly to tasks: No components should be allowed to sit idle and be retained for a low-chance contingency. For the pervasive flexibility afforded by degeneracy to prove its benefits, it would also be necessary for the swarm to operate in a complex environment where many partially related tasks are required of the system.

With a plug-and-play architecture introduced at the robot assembly level (e.g., plug-and-play of sensors, flying component, magnetic attachment, grabbing device, chassis, battery), and with design alternatives provided for each component, large combinations of new robot configurations would be available to expand the structural and functional diversity of the robot swarm. This would result in a system that could no longer be easily decomposed into distinct classes of foot-, hand-, and eye-bots. Assuming that protocols for collaboration and robot-robot assembly were maintained, and assuming that robots could explore their functional limitations in a safe environment, this diversity should allow the swarmanoid system to complete complex tasks in a greater variety of operational environments. Through the incremental testing and evolution of new robot designs, the degree of degeneracy would grow at the component, robot, and multi-robot assembly levels; the flexibility of the swarm to deal reliably with complex tasks under a range of conditions would increase; and the swarm's commercialization potential for complex and dangerous environments such as rescue operations, mining, and space exploration would improve.

Importantly, however, degeneracy provides more than just functional reliability in a volatile environment. It also expands the opportunities for the cooptation of existing components for new functions, because new multi-robot functions can be discovered through their assembly and usage within novel environments. While the new functions might not be optimized, they still provide useful information and guidance for the designer to create new component or robot designs that can eventually expand system operations into an important niche environment. This narrative of environment-revealed adaptations followed by design modification is paralleled in biological evolution in the form of exaptations (see Section 3) followed by mutation-driven genetic accommodation/assimilation, that is, the fine-tuning of the new functions.

8 Conclusions

Our research into distributed robustness suggests that the reductionist evaluation of resource redundancy can sometimes be highly inaccurate and lead to erroneous conclusions regarding the value of versatile assets that are deployed within a volatile and uncertain environment [52]. We claim that a better understanding of distributed robustness requires moving beyond heuristics that only attribute robustness from diversity to bet-hedging and portfolio-theoretic principles. What is sorely lacking is a principled approach to the systematic design of living technologies where forms of robustness and flexibility can emerge on demand [52].

The NB described in this article illustrates one form of emergent robustness that can be understood and designed without knowing precisely where flexibility will be needed or what perturbations will be experienced. To achieve NB, elements in the system must display a partial overlap in functional capabilities: in some contexts providing functional redundancy, in others providing response diversity. In biology this unique type of group behavior is known as degeneracy.

Previous investigations into the potential importance of degeneracy in multi-agent systems suggest that degeneracy could broadly contribute to the flexibility needed in living technologies because:

1. Degeneracy can lead to an emergent form of systemic flexibility through NB. The NB hypothesis [50] extends the value of diversity beyond that of portfolio theory toward a system-level relationship between diversity, performance, and resilience.

2. When organized into buffering networks, degeneracy can enhance a system's capacity to take advantage of opportunities originating from component changes; exploit novel environmental conditions; and mitigate the effects from unanticipated stresses [52]. Stated differently, systems with inbuilt degeneracy can develop an innate capacity to deal with unforeseen, new challenges.

3. Degenerate systems can sometimes violate robustness-efficiency tradeoffs that are widely perceived as hard constraints. The robustness-efficiency balance attainable in systems with innate degeneracy significantly outperforms that achievable in systems where degeneracy is absent [32].

4. The strategic advantage of highly degenerate multi-agent systems does not appear to degrade operational effectiveness. In particular, the benefits of NB appear to be attainable at small cost in management overheads when decision making is distributed across the system [52].

Degeneracy is a system property that can be clearly articulated and defined for any system composed of functionally versatile elements. Currently, several research programs are exploring how the degeneracy concept can be translated into design principles for the realization of more flexible and resilient systems in different disciplines [23, 38, 51, 52]. For instance, in defense capability simulations, we have shown that fleets of military field vehicles with high degeneracy in task capabilities can improve operational robustness within anticipated mission scenarios, yet at the strategic level provide exceptional design and organizational adaptability for responding to unanticipated challenges [6, 52]. We are also looking at how the degeneracy concept can be translated into the design of more flexible manufacturing and assembly systems [23], and can improve the performance of population-based dynamic optimization [51]. Other researchers are using the concept to understand the requisite conditions for embodied [21] and simulated artificial life [10, 27, 36].

In this article we have described how component-level functional versatility and network-level functional redundancy enable degenerate elements to facilitate exaptations. We have also described how the inclusion of protocols and loose coupling can enable the incremental evolution of highly degenerate systems. We looked at how these ideas can help realize some of the desirable features of living technologies and focused on swarm robotics as a case study. We evaluated a particular system known as the swarmanoid. In many ways, the Swarmanoid Project captures key design principles for the realization of behavioral adaptation, configurational flexibility, and design evolvability that are rare outside of natural systems. Swarmanoid exemplifies a system architecture with loose coupling in component design and a plug-and-play modularity that facilitates dynamic assembly guided by environmental context. The swarmanoid system supports self-assembly, distributed decision making, and component diversity: features that are absent from most robotics research.

However, to fully realize its potential as an applied living technology within environments that are volatile and uncertain, the swarmanoid system needs a greater degree of flexibility for responding to real-world conditions. The flexibility afforded by degeneracy could provide advantages in the design and operation of robotic swarms, while the swarmanoid's plug-and-play architecture should enable this diversity to arise from a relatively small set of core building blocks.

8.1 General Recommendations

Enhancing a system's adaptive capabilities involves difficult decisions. Tradeoffs between the cost of redundancy and the need for flexibility require careful choices that reflect expectations on the size and nature of future volatility. While degeneracy-based design is advantageous for many living technologies and sociotechnical systems, there are arguably circumstances in which degeneracy may not be needed. Deriving benefits from degeneracy requires both a top-down valuation of a system (and its configurational variations) under different scenarios (i.e., what aspects of diversity are likely to translate into a competitive advantage) and a bottom-up assessment of the opportunities for the deployment of existing assets, resources, and components under a variety of new conditions. Such evaluations need to be traded off against design, training, and management overheads that may come about with the introduction of degeneracy. To justify a change toward degeneracy-based design, decision makers will need to:

- assess the changes in capabilities that expanded component functionality confers;
- identify what cost containment is achievable from reusable training modules for skill development or from reusable physical modules in construction;
- estimate the expected returns on investment (for example, does the system naturally lend itself to the creation of networked buffers where robustness can be increased significantly with negligible decreases in efficiency?).

These decisions should be rational to individual stakeholders, but are difficult to make in view of the scarcity of information that relates to the largely intangible benefits of degeneracy, notably the potential for adaptation under unanticipated conditions.

Critics of nature-inspired design often claim that biological systems display costly levels of component complexity and gratuitous inefficiency. Degeneracy does indeed require component diversity, which can come with design costs and management overheads. Biological systems solve this problem by creating degeneracy through the versatile reuse of a small number of molecular building blocks. Modular technological and sociotechnical systems with a plug-and-play architecture should similarly be able to achieve high levels of degeneracy at relatively low cost.

References

1. Adami, C. (2002). Sequence complexity in Darwinian evolution. *Complexity*, *8*, 49–57.
2. Adler, P. S., & Borys, B. (1996). Two types of bureaucracy: Enabling and coercive. *Administrative Science Quarterly*, *41*(1), 61–89.
3. Alderson, D. L., & Doyle, J. C. (2010). Contrasting views of complexity and their implications for network-centric infrastructures. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, *40*, 839–852.
4. Barkai, N., & Leibler, S. (1997). Robustness in simple biochemical networks. *Nature*, *387*, 913–917.
5. Bedau, M. A., McCaskill, J., Packard, N., & Rasmussen, S. (2010). Living technology: Exploiting life's principles in technology. *Artificial Life*, *16*, 89–97.
6. Bender, A., & Whitacre, J. M. (2011). Emergent flexibility as a strategy for addressing long-term planning uncertainty. In *NATO Risk-Based Planning Conference*.
7. Bhattacharyya, R. P., Reményi, A., Yeh, B. J., & Lim, W. A. (2006). Domains, motifs, and scaffolds: The role of modular interactions in the evolution and wiring of cell signaling circuits. *Annual Review of Biochemistry*, *75*, 655–680.
8. Chandler, A. D. (1993). *Strategy and structure*. Cambridge, MA: MIT Press.
9. Chandra, F. A., Buzi, D., & Doyle, J. C. (2011). Glycolytic oscillations and limits on robust efficiency. *Science*, *333*, 187.
10. Clark, E., Nellis, A., Hickinbotham, S., Stepney, S., Clarke, T., Pay, M., & Young, P. (2011). Degeneracy enriches artificial chemistry binding systems. In *ECAL 2011* (pp. 133–140), Cambridge, MA: MIT Press.
11. Crutchfield, J. P., & Görnerup, O. (2006). Objects that make objects: The population dynamics of structural complexity. *Journal of the Royal Society Interface*, *3*, 345–349.
12. Csermely, P. (2004). Strong links are important, but weak links stabilize them. *Trends in Biochemical Sciences*, *29*, 331–334.
13. Csermely, P. (2006). *Weak links: Stabilizers of complex systems from proteins to social networks*. Berlin: Springer-Verlag.
14. Csete, M. E., & Doyle, J. C. (2002). Reverse engineering of biological complexity. *Science*, *295*, 1664–1669.
15. Csete, M., & Doyle, J. (2004). Bow ties, metabolism and disease. *Trends in Biotechnology*, *22*, 446–450.
16. Dorigo, M., Floreano, D., Gambardella, L. M., et al. (2012). Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics & Automation Magazine* (in press).
17. Doyle, J. C., Alderson, D. L., Li, L., Roughan, M., Shalunov, S., Tanaka, R., & Willinger, W. (2005). The "robust yet fragile" nature of the Internet. *Proceedings of the National Academy of Sciences of the U.S.A.*, *102*(41), 14497–14502.
18. Doyle, J., & Csete, M. (2007). Rules of engagement. *Nature*, *446*, 860.
19. Edelman, G. M., & Gally, J. A. (2001). Degeneracy and complexity in biological systems. *Proceedings of the National Academy of Sciences of the U.S.A.*, *98*, 13763–13768.
20. Edmonds, B. (2000). Complexity and scientific modelling. *Foundations of Science*, *5*, 379–390.
21. Fernandez-Leon, J. A. (2011). *Behavioural robustness and the distributed mechanisms hypothesis*. Ph.D. thesis, University of Sussex.
22. Figge, F. (2004). Bio-folio: Applying portfolio theory to biodiversity. *Biodiversity and Conservation*, *13*, 827–849.

23. Frei, R., & Whitacre, J. (2012). Degeneracy and networked buffering: Principles for supporting emergent evolvability in agile manufacturing systems. *Journal of Natural Computing*, *11*(3), 417–430.

24. Gell-Mann, M. (1995). What is complexity. *Complexity*, *1*, 1–9.

25. Gould, S. J., & Vrba, E. S. (1982). Exaptation: A missing term in the science of form. *Paleobiology*, *8*, 4–15.

26. Hazen, R. M., Griffin, P. L., Carothers, J. M., & Szostak, J. W. (2007). Functional information and the emergence of biocomplexity. *Proceedings of the National Academy of Sciences of the U.S.A.*, *104*, 8574–8581.

27. Kerkstra, S., & Scha, I. R. (2008). *Evolution and the genotype-phenotype map*. Master's thesis, University of Amsterdam.

28. Kirschner, M., & Gerhart, A. (1998). Evolvability. *Proceedings of the National Academy of Sciences of the U.S.A.*, *95*, 8420–8427.

29. Kitano, H., & Oda, K. (2006). Robustness trade-offs and host-microbial symbiosis in the immune system. *Molecular Systems Biology*, *2*, 2006.0022.

30. Kondoh, M. (2003). Foraging adaptation and the relationship between food-web complexity and stability. *Science*, *299*, 1388–1391.

31. Kurakin, A. (2009). Scale-free flow of life: On the biology, economics, and physics of the cell. *Theoretical Biology and Medical Modelling*, *6*, 6. doi:10.1186/1742-4682-6-6.

32. Kurakin, A. (2011). The self-organizing fractal theory as a universal discovery method: The phenomenon of life. *Theoretical Biology and Medical Modelling*, *8*, 4. doi:10.1186/1742-4682-8-4.

33. Levchenko, A., & Iglesias, P. (2002). Models of eukaryotic gradient sensing: Application to chemotaxis of amoebae and neutrophils. *Biophysical Journal*, *82*, 50–63.

34. Macia, J., & Solé, R. (2009). Distributed robustness in cellular networks: Insights from synthetic evolved circuits. *Royal Society Interface*, *6*, 393–400.

35. McShea, D. W. (1996). Perspective: Metazoan complexity and evolution: Is there a trend? *Evolution*, *50*, 477–492.

36. Mendao, M., Timmis, J., Andrews, P. S., & Davies, M. (2007). The immune system in pieces: Computational lessons from degeneracy in the immune system. In *Foundations of Computational Intelligence, 2007. FOCI 2007. IEEE Symposium on Foundations of Artificial Immune Systems* (pp. 394–400).

37. O'Grady, R., Birattari, M., & Dorigo, M. (2011). Swarmanoid, the movie. In *AAAI-11 Video Proceedings*.

38. Randles, M., Lamb, D., Odat, E., & Taleb-Bendiab, A. (2010). Distributed redundancy and robustness in complex systems. *Journal of Computer and System Sciences*, *77*, 293–304.

39. Reznick, D. N., & Ghalambor, C. K. (2001). The population ecology of contemporary adaptations: What empirical studies reveal about the conditions that promote adaptive evolution. *Genetica*, *112*, 183–198.

40. Senge, P. (1993). The fifth discipline: The art and practice of the learning organization. *Consulting Psychology Journal: Practice and Research*, *45*, 31–32.

41. Schindler, D. E., Hilborn, R., Chasco, B., Boatright, C. P., Quinn, T. P., Rogers, L. A., & Webster, M. S. (2010). Population diversity and the portfolio effect in an exploited species. *Nature*, *465*, 609–612.

42. Siegal, M. L., & Bergman, A. (2002). Waddington's canalization revisited: Developmental stability and evolution. *Proceedings of the National Academy of Sciences of the U.S.A.*, *99*, 10528–10532.

43. Stelling, J., Sauer, U., Szallasi, Z., Doyle, F. J., & Doyle, J. (2004). Robustness of cellular functions. *Cell*, *118*, 675–685.

44. Tilman, D. (1996). Biodiversity: Population versus ecosystem stability. *Ecology*, *77*, 350–363.

45. van Wageningen, S., Kemmeren, P., Lijnzaad, P., et al. (2010). Functional overlap and regulatory links shape genetic interactions between signaling pathways. *Cell*, *143*, 991–1004.

46. Wagner, A. (2000). Robustness against mutations in genetic networks of yeast. *Nature Genetics*, *24*, 355–362.

47. Wagner, A. (2005). Distributed robustness versus redundancy as causes of mutational robustness. *BioEssays*, *27*, 176–188.

48. Whitacre, J. M. (2010). Degeneracy: A link between evolvability, robustness and complexity in biological systems. *Theoretical Biology and Medical Modelling*, *7*(6). doi:10.1186/1742-4682-7-6.

49. Whitacre, J. M., & Bender, A. (2010). Degeneracy: A design principle for achieving robustness and evolvability. *Journal of Theoretical Biology*, *263*, 143–153.

50. Whitacre, J. M., & Bender, A. (2010). Networked buffering: A basic mechanism for distributed robustness in complex adaptive systems. *Theoretical Biology and Medical Modelling*, *7*(20). doi:10.1186/1742-4682-7-20

51. Whitacre, J. M., Rohlfshagen, P., Bender, A., & Yao, X. (2010). The role of degenerate robustness in the evolvability of multi-agent systems in dynamic environments. In *Proceedings of the 11th International Conference on Parallel Problem Solving from Nature: Part I* (pp. 284–293) Berlin: Springer-Verlag.

52. Whitacre, J. M., Rohlfshagen, P., Bender, A., & Yao, X. (2012). Evolutionary mechanics: New engineering principles for the emergence of flexibility in a dynamic and uncertain world. *Natural Computing*, *11*(3), 431–448.

53. Yi, T. M., Huang, Y., Simon, M. I., & Doyle, J. (2000). Robust perfect adaptation in bacterial chemotaxis through integral feedback control. *Proceedings of the National Academy of Sciences of the U.S.A.*, *97*, 4649.