Lints, T. (2010). How to facilitate variability. In *Artificial Life XII, Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, pages 569–576. The MIT Press.

T. Lints, "How to facilitate variability," in *Artificial Life XII, Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, pp. 569–576, The MIT Press, 2010.

```
@inproceedings{Lints10_Variability,
  author = {Taivo Lints},
  title = {How to Facilitate Variability},
  year = {2010},
  booktitle = {Artificial Life XII, Proceedings of the Twelfth
International Conference on the Synthesis and Simulation of
Living Systems},
  publisher = {The MIT Press},
  pages = {569-576}
}
```

How to Facilitate Variability

Taivo Lints

Research Laboratory for Proactive Technologies
Tallinn University of Technology, Ehitajate tee 5, 19086 Tallinn, Estonia
taivo@taivo.net

Abstract

The concepts of *life* and *intelligence* almost *require* the system to be adaptive. And adaptivity, in turn, is usually strongly dependent on the continual generation of variations in the system. The paper discusses various ways of producing the required variations, and how to support these production processes.

Introduction

The property of being alive seems to almost *require* (if not yet with scientific rigor, then at least intuitively) the existence of adaptational processes in the system – it is difficult to imagine a lifeform whose internal processes and behavior would not depend in any reasonable (fitness-linked) way on the situation the organism is in. Adaptivity, in turn, has strong, though less strict, ties with the generation of *variations* in / by the system.

The evolution theory inspired approaches to adaptation consider it to be a process where variations of existing individuals are being generated and where selection operates on those variants, probabilistically eliminating the less fit ones. The variation-selection loop is not a strict requirement for adaptation in general (because adaptive behavior can also be displayed by a system that is able to accurately enough estimate the required states and actions and generate them in "one shot"), but nevertheless a notable portion of adaptational processes can be described as having such a character.

In cybernetics, too, the importance of variety for a system's ability to cope is emphasized, though in a slightly different sense: "The larger the variety of actions available to a control system, the larger the variety of perturbations it is able to compensate." (Ashby's (1956) idea of requisite variety, as summarized by Heylighen and Joslyn, 2001). Here, the variants are not exactly competing with each other for survival, but rather form an operational repertoire the system can draw from as required by the circumstances.

The widespread usage of the concept of diversity in debates about sustainability and problem solving furthermore suggests that the existence of variations in a system may increase its adaptivity as well as robustness.

And, finally, the need for some kinds of variations in a system that is considered adaptive derives directly from the essence of adaptation itself, which can be defined as "changing something (itself, others, the environment) so that it would be more suitable or fit for some purpose than it would have otherwise been" (Lints, 2010) – the term 'change' is pretty much synonymous with 'variation in time', i.e., something is transformed from one state to another and there are different variants of it at different time points (which, in turn, may, or may not, depending on the system, be facilitated by the existence of multiple simultaneously present variations of system elements (components, processes, relations, etc.)).

All in all, then, it is of great import for adaptation research, and, consequently, for ALife research, to study the ways how variability can be stimulated. At least three issues can be identified. Firstly, the very generation itself – what are the ways to produce variations. Secondly, how to support that generation, i.e., how to make it easier for the generative processes to operate well in a system. And thirdly, how to trigger the production of new relevant variations when the mechanisms are already in place but latent or unguided. This paper explores the first two of these issues. It should be noted that the paper grew out of the author's untested pondering on the topic of adaptivity and does *not* attempt to survey the variability related research done so far (and, accordingly, the given references are *not* representative of the main research efforts of that direction; but, on the other hand, it is exactly because of that why the paper might potentially provide some perspectives, connections and summarizations interestingly divergent from the usual).

Ways of Generating Variations

There exist several perspectives from which to dissect the ways of producing variations. One might be called a "creativity perspective", which lists the possibilities in accordance with how (or if) the novelty is produced (surely, the terms *creativity* and *novelty* are somewhat difficult to define, but for our current purposes they serve mostly as referential labels and thus the lack of rigorous definitions is not particularly problematic). The baseline would be having no novelty at all, from the system's own perspective and (to keep the current discussion within reasonable limits) with regard to the set of variations, not the set of pairings of variations with the situations. This would be the case when, for example, all the possible variations already exist in some kind of an internal repository and the system merely draws them from this store.

Combinatorial novelty can be produced through, as the name suggests, producing novel combinations of existing elements, be they physical system parts or various signals, processes, arrangements, etc. In genetics, a typical example would be the crossover operation that basically takes some DNA strands from two individuals and swaps some of their sections with each other. But combinatorial novelty is not limited to preserving the sizes or numbers of inputs, of course, and may in principle use any kind of element pool to produce any other kind of element pool constructable from the (parts of the) initial material. If the arrangement of the elements is important for the system, then a mere rearrangement (permutation) can also be considered to produce a novel variant from existing parts. Another noteworthy possibility is the so-called *bootstrapping* where the products of one generational cycle are used as elementary building blocks in the next cycle (it is worth emphasizing, though, that bootstrapping is a powerful method not limited to combinatorial approach and can be used with most of the other techniques as well).

To produce new alterations in a possibly noncombinatorial way (though it can also be used with the combinatorial method), the first approach would be incremental tuning or modification of system's parameters and parts, i.e., moving around relatively smoothly in the space of modifiables. Whether this translates to the system moving around smoothly in its state space as well depends on the mappings from modifiables to system states and dynamics, as well as on the general complexity and nonlinearity of the system. In developmental systems the extent of the effect a modification has is usually also strongly dependent on how early in the development the modification was made – early changes often have strong effects (which helps to explain why, especially in biology, early development often remains relatively conservative in comparison to later development: the large impacts of early alterations render, in most cases, the system unfit (Bennett, 1997) and thus are selected against).

Moving up on the hypothetical creativity ladder we find the revolutionary, "truly creative" change, the existence of rigorous meaning and essence of which is somewhat questionable, but intuitively it implies the occurence of particularly noteworthy advances, strong originality and innovation, and large unexpected (but clever, at least in hindsight) changes in modifiables, as opposed to the more mundane step-by-step tuning. In practice, though, the line between incremental and revolutionary is blurry, and even more so with the occasional distinction between truly creative and "just" combinatorial, as it is actually common for the breakthrough ideas to stem from intensive work with extensive presence of both incremental and combinatorial methods. Also, in nonlinear systems the slight tuning of some system parameter can lead to substantial changes in other variables.

A classification somewhat orthogonal to the previously described one can be reached at when differentiating between the system being self-contained with regard to novelty creation versus it drawing some variants, or elements of them, from external sources. The most obvious situation would be using an external knowledge repository, the form of which can range from databases through helpful systems / agents up to the vast accumulated knowledge of the whole human, or other, culture. Another possibility is the incorporation of (or merging with) external components that supplement system's own capabilities. This might be done temporarily on the basis of need, or also permanently. In some cases even the temporary inclusion of a component (say, an employee) can permanently upgrade the system's abilities (say, in the form of idea exchange / extraction). Probably the most complex, but accordingly with the highest potential payoff, way of acquiring variations from external world is a (mutual, creative, constructive, temporally extended) interchange process between the system and various external agents.

Yet another perspective on producing variations can be constructed by focusing on the spectrum of possible uses of randomness and determinism in the system – whether the search for new variations (or the act of retrieving existing ones from some repository) is random or determined, guided by previous experience or not, and what characteristics the sources of randomness have.

A fully random search with a flat probability distribution samples the search space, by definition, uniformly and without any guidance from previous experience. A possibility to be noted, though, is that if the search space is not the same as the space of directly testable outcomes (e.g., genotypes are being varied but the selection is based on final organisms that develop under the guidance of those genotypes), the probability distribution may well become skewed somewhere in the mappings from modifiables to testables (the mappings can be very complex, involve generative rules, randomness, context-dependence, emergent behavior, self-organization, etc.). For the system this could be either a problem or an opportunity.

As the probability distributions become less and less flat, either through the changes in the aforementioned mappings or directly at the source of randomness, there will be more and more predictability (at least in principle) in the system, finally in the limit reaching full determinism. The shaping of distributions might be accidental, but a considerably more interesting case is when it is used as a way to store previous experience or externally acquired knowledge – those regions of search space that have become known to be more likely to contain good solutions are searched more thoroughly and preferentially earlier than other regions. One has to be careful, though, to take into account the possibility of the circumstances changing or of the existence of special cases, for both of which the solutions may lie in areas previously experienced as solution-poor.

In some cases the search can also be exhaustive, generating all the possible variations of the modifiable(s). While exhaustive search is typically prohibitively costly, and purely random search too unintelligent, the option to use them should not be totally forgotten or immediately discarded, as occasionally they may really turn out to be the most viable ways to find good solutions (e.g., Wolfram, 2002, page 393).

As a fair share of interesting systems could be classified as nonlinear and complex, there is one more potentially important source of variations: deterministic chaos. It can amplify minor fluctuations and deviations, both deliberate and accidental, deterministic and random, into major changes in system dynamics totally, and in practice quite unpredictably, altering the system's behavior in the long run.

For the probabilistic and deviation-amplifying methods to work properly, it is necessary to have a source of randomness. This can be located either inside or outside the system, and be truly random or pseudorandom. If the usage of the source is deliberate, the values of the random variable might be explicitly acquired from the source, but in most cases the randomness kind of "leaks in" as noise in imperfect sensors, signal channels, processing elements, actuators, etc., or in the form of perturbations of the "normal" system behavior, composition or organization.

One more informative way of classifying the variation-producing methods rests on the sequential-parallel scale, distinguishing between systems that create new variations one by one in a row (and, in extreme cases, only allow the existence of one variant at a time) and systems that either spawn multiple simultaneously active variety generators or just generate a number of alternatives more or less instantaneously (at least from the practical viewpoint).

While it is educative to be aware of all the described techniques, it should be kept in mind that they are not mutually exclusive – it can often be advantageous to combine various approaches instead of relying on a single mechanism. The partial orthogonality of the "perspectives" is relatively obvious, but even within a single perspective there are possibilities for diversity, e.g., having both random and deterministic, or both parallel and sequential variation generators present in the same system. The different mechanisms can be applied to altering different modifiables, be cooperating on the same ones, act as backups for each other, and so on.

Supporting the Generation of Variations

For the various aforementioned methods to have a possibility to work well, the system they operate in should provide some specific support in the form of having certain features and resources. Some of the most important ways of help are described in the following subsections.

Making the Modifiables Easy to Change

The job of a variation generator could be roughly described as producing altered versions of the system, usually based on the system's previous state(s) or on some template or seed. An alteration is basically a change of some modifiable features of the system, executed either in the very same system (component) or by fabricating a new altered copy instead. It is quite straightforward to deduce, then, that making the modifiables easy to change can make the job of the generator much easier.

The specifics of how the effortlessness can be achieved depend, obviously, on the particular system, but in general the following keywords might give the first hints on the directions to pursue: tunability, reconfigurability, rearrangeability, reroutability, flexibility, plasticity, elasticity, adjustability. The main connective idea here, almost by definition, is to reduce the resistance to change. This includes reducing the cost of adjustment actions, increasing responsiveness (the speed at which the changes can be made), relaxing constraints (except maybe the ones that directly support variation generation by keeping the corresponding mechanisms functional, e.g., in genetic systems "the extremely high internal correlations underlying the transcription and translation mechanisms allow for a large ensemble of variants" (Conrad, 1983, page 338)), removing various barriers, and also increasing the number of options for each modifiable feature (both by expanding the range and by upping the density of allowed positions in that range) as well as the number of modifiables themselves. In addition to reducing the cost of adjustment actions, the (meta-level) costs of *maintaining* the flexibility are also important to be paid attention to and reduced as much as possible or feasible.

As of increasing the number of options, an interesting concept is neutral variation on a flat plateau of fitness landscape, meaning that something can be varied a lot without affecting the measure of system's current successfulness much. In general this is *not* what we would like to have when enlarging the set of options, because by definition the added options on the same plateau give the same fitness result as those already existing there. However, there still exist potential ways to use it. One is to notice that although different spots on the same plateau do have the same elevation, their neighboring areas might not, thus the new options might provide better access to new interesting places on the fitness landscape while being easy to reach themselves due to neutrality (because of being similar to other variants there is likely to be less resistance against moving into them) (e.g., Lenski et al., 2006). Another possibility is to look at some kind of an "opposition to alterations" landscape instead (the construction of which is trickier, though, as the resistance to moving into a given point depends not only on the static paramater values of that point but also on dynamics, and is typically not the same for different origins of alteration), find plateaus there and define neutrality on such a basis. Then the areas of interest would be plains of low resistance but of useful variability in fitness-relevant dimensions.

A related concept, originating from physics, is referred to as the system having glassy properties and means, among other things, that there are "multiple low-energy minima in the energy landscape of the system" (Menashe et al., 2000). This, in turn, means that there is no uniquely predetermined state to which the system would always try to fall, but instead a variety of equi-energetic states to "choose" from. And that possibility of choice would increase the system's potential capacity to adapt, and would move it closer to (or further in) the domain of biology (Stec, 2004).

Yet another related idea for fostering variety is to keep the system sufficiently far from equilibrium so that it has plenty of stationary states to choose from (Heylighen, 2001).

Whereas neutral variation and ideas related to it definitely deserve further research about how to apply them for supporting variety generation, they are probably not the key concepts and were given a somewhat disproportionate attention here mainly due to their intellectual appeal. A considerably better studied and in all likelihood more important notion is that of modularity – something consisting of changeable pieces is typically a lot easier to modify than a monolithic structure. Although modularity promptly associates with some physical system or software being composed of distinct components, the idea has a lot wider applicability. To give a few examples, it is possible (and sometimes possibly enlightening) to talk about modularity in time, modularity of search space, state space, action space, or some more exotic space, modularity of representations, behaviors, signals, protocols, functionality, resources, and much more.

Linking the concepts of tunability and modularity, we can arrive at the idea of having tunable and exchangeable components. In general this is a thought too obvious maybe to even mention, but in some areas it does not necessarily come to mind that easily, yet is exceedingly useful nevertheless. An example would be for a system to have switchable sets of tunable behaviors where tuning improves the currently active set and changeovers are triggered by context changes, as opposed to having only a single tunable set that can slowly become another (distant) one as is common in simpler artificial learning systems (Moorman and Ram, 1992).

An additional option for supporting variation generation is to make the modifiables polyadjustable, that is, to have the same feature be adjustable by a variety of different mechanisms (Knoll and Järvenpää, 1994). Depending on the specific circumstances this can provide the system with the possibility to choose the most efficient change mechanism for given situation, to have backup if some of the mechanisms fail, to more effortlessly generate interesting and complicated variations by playing around with several interacting mechanisms, and so forth. But, assuredly, polyadjustability may also make it more difficult to tune something if the various mechanisms interact in a particularly intricate way. Polygenic control is an example of natural use of polyadjustability, where some characteristic of a biological organism is controlled by more than one gene.

Looking at the problem of reducing resistance to change from the viewpoint of psychology adds yet another perspective to the discussion, one that is concerned with systems being deliberate agents, or collections of them. In this view, the topic is more commonly referred to as *openness to new*, where "new" includes both the easier case of novel input that agrees well with agent's current worldview and the more challenging situation of input that does not.

The main problem with regard to variation generation (and to adaptivity in general) is that people and social groups have a tendency, after initial developmental period, to become quite fixed in their ways of thinking and doing. We have cognitive predispositions to confirmation bias, fallacy of centrality, hubris, normalization, typification, and bottom-up salience of cues, as well as to lock-in and fixation (Weick, 2005). Similarly, in social groups and institutions various behaviors and beliefs more or less spontaneusly emerge and form the "culture of the organization", which will then create a great deal of inertia to change (Grisogono, 2005). To allow for novel variations to be introduced into such systems it is thus necessary to offset those cognitive predispositions (Weick, 2005), to induce openness to conflicting inputs (Harvey et al. 1961, page 333, as referred to by Hunt, 1966), to break the addiction to listen and accept only perspectives similar to one's own (Holley, 2005), etc. Whereas the common approach is to just inform people about how it would be better to act and then expect or require them to follow the guidelines, it would be considerably more effective to take the time and really help people (or whoever / whatever the deliberate agents are in the system of interest) break old behavioral habits in combination with establishing new ones. Also, enough psychological safety should be provided in order to combat the urge for closure and certainty. This means it should be assured that "it is much more important to be prepared to be wrong in order to learn, than to always be right (and therefore either or both risk-averse or in denial) and conversely, being prepared to 'decriminalise' others being wrong" (Grisogono and Ryan, 2007), as well as made sure that the group or organization is safe for interpersonal risk taking (speaking up, offering suggestions, critiques, expertise, advise) (Stagl et al., 2006). The habits of constantly challenging one's own thinking and being prepared to look for both confirming and contradictory evidence (Grisogono and Ryan, 2007), making explicit (even vocalizing, for particularly critical processes and decisions) the situation reviews, alternative diagnoses and plans (Weick, 2007), and being tolerant of uncertainty and responsibility (Ku, 1995, page 316) should be encouraged.

Also aimed mainly at deliberate agents is the suggestion to avoid various plans becoming too prescriptive (Holmqvist and Pessi, 2006). By using an extended understanding of what a plan is (can include blueprints, generative codes, various evolvable constraints, guidelines learned from experience, and much more) it can well be applied to most other adaptive systems, too. Having plans is, assuredly, very often beneficial, and the very process of planning itself can help a lot with understanding and solving the problem at hand. But if the plans are followed through rigidly, the adaptivity of the system in general and the generation of (unplanned) variations in particular may suffer a lot. Multiple ways of achieving plan flexibility exist. One is to just keep revising the plan dynamically, taking into account the new situational information (Burke et al., 2006). Another is to make the plans themselves somewhat loose, for example to have strategies suggesting boundaries on behavioral parameters rather than precise values (Ram and Santamaría, 1997). And finally there is a possibility to plain discard parts of the plan, or the whole of it, as deemed necessary. In group situations the latter option can be made easier by avoiding strongly binding contracts and building an ability to replace some of the planning with on-time communication (Andersen, 2003).

Ending the current list of the ways of making the modifiables easy to change, but certainly not closing the set of all possibilities, is the option of adding some form of redundancy to the system. Having multiple copies of the same components not only can increase the reliability of the whole system, but also facilitates transformability and mutability (Conrad, 1983, page 337): in addition to the straightforward potential benefit of having more elements to target with altering actions, the workings of the system do not depend critically on single components anymore and thus the unsuccessful variants of the elements do not immediately render the whole system inoperative (except in some particularly unfortunate cases of highly disruptive variants), which encourages more aggressive varying. A possibly even safer approach would be to decouple the exploration architecturally and functionally from the rest of the system. The better variants could then either directly and forcefully substitute the ones currently in effect in the main part of the system or, as suggested by Grisogono and Ryan (2007), "to work provisionally alongside established ways of doing things, without relying on them, but using the parallel system enough to identify and fix flaws with it until confidence in it grows sufficiently that users start transferring to it in preference to the previous system". Finally, taking this direction of adding redundancy and separating it from the main operational part to its logical conclusion, we reach virtual variation generation that is executed in models and simulations and thus potentially allows for particularly rapid alteration production and testing. But, surely, the use of models has various possible drawbacks as well, e.g., a less than ideal match with reality might lead to erroneous results and decisions.

Making the System Tolerant to Errors

In real life, variation generation almost inevitably produces a significant number of unfit alterations along with the acceptable ones. If those mistakes have a strong negative effect on the system, either real or imaginary (e.g., psychological problems), then the whole variation generating process may be considered undesirable and its activity reduced to minimum, with potentially dire consequences to system's adaptivity. Thus, making the system tolerant to errors is an important factor in supporting the generation of novel variations. For deliberate agents with psychological problems that might involve making them aware of the near unavoidability, or even desirability, of mistakes on the path of success, but in general it is mostly about increasing robustness, redundancy, reversibility and / or repairs, and actually also adaptivity (regardless of the slight touch of circularity that it seems to bring into our discussion) which would allow for incorporating some of the errors in a way that transforms them from mistakes into neutral or even useful features.

Robustness, as understood here, is the capacity to withstand various perturbations without needing an active, adaptive, response. It can come about in multiple ways, mostly by having the important functionality being just plain insensitive to disturbances (as in neutral variation discussed earlier), by making the critical parameters very difficult to change, or by having enough redundancy in the system so that single failures cannot eliminate important functionality. Redundancy can provide even more safety if it is implemented not by simply having multiple copies of the very same element, but by having different components with partially overlapping functionalities, because this protects better against systemic errors that affect all instances of some element type (e.g., Edelman and Gally, 2001).

The more active side of error tolerance – reversing, repairing, or adapting to mistakes – either tries to restore the pre-mistake state of the component or reorganize the system to now use what was previously considered a problem as a useful feature instead. Reversibility can be fostered, for example, by representing the targets of modification so that each modification would be a simple flip of some bit (or a switch between few alternatives), the undoing of which is relatively straightforward (except only when the rest of the system has already changed too much due to the unfit alteration and will not restore itself appropriately after reverse modification). Or, in some cases, the so-called system restore points can be occasionally created by saving the system state in a recoverable way, up to producing full back-ups every once in a while (especially before potentially dangerous modifications). Usually this would require the implementation of several special reversibility-related mechanisms, but sometimes there may also exist possibilities to achieve similar effects with less effort. An example would be to have the new variant just functionally override the previous one without actually removing it from the system immediately, so that for recovery it would be enough to withdraw the new element and thus allow the previous one to function again. *Genic occlusion* is a natural instance of such a method: a gene is suppressed through addition of a further "upstream" gene to the epistatic set (a set of interacting genes), with no actual change to the original locus itself (Brock, 2000, page 245). As of repairing and readjusting, some options (in a military context) are listed by Unewisse and Grisogono (2007): shifting of essential tasks from damaged to undamaged elements, exploiting redundancy within system; redistribution of tasks within system, exploiting multiroled or multifunctioned elements; repair of damage, which requires the capacity to detect damage, assess and repair it, exploiting capacity for frontline repair and rapid mobilisation of logistic chains; redistributing tasks so that essential ones are done vice non-essential; compensate for the damage by changing the resources available to the system.

One has to be careful, though, with using the error tolerance increasing methods for supporting variability, because more often than not the system will also be less sensitive to the variations themselves, somewhat counteracting the expected positive effect. Occasionally the very opposite action would be beneficial instead, as illustrated by yet another example from genetics where one way to increase mutation rate (in conditions calling for higher adaptivity) is through inhibition of DNA repair processes (Hersh et al., 2004; Denamur and Matic, 2006). The latter option is particularly suitable for harsh situations where the survival of the system (usually a population) is put into considerable danger and the normal adaptational mechanisms are unlikely to be of enough help – then the high occurence of (totally) unfit variants is outweighed by the increased probability of also finding some new viable forms because the alternative would likely be an irreversible extinction of the whole system.

Choosing Suitable Representations

A large share of nontrivial systems make use of various internal representations in order to process information and store knowledge. In principle there can be a near infinite number of different representations that refer to the same "real" entities, and furthermore a near infinite number of mappings both from the referenceable set to representations and back. While equal in some ultimate respect, those alternative representations and mappings may present different practical opportunities and constraints for the system, including to the variation generation mechanisms. If the modifications executed in an adapting system target the very representations themselves, then the influence of the choice of representations on the variation generation is often obvious. But even if they do not, the representations may be important intermediaries in the chains from introduced modifications to systemic results and thus can still have a significant impact on how easy it is to produce relevant variations.

When representations are looked at as yet another kind of modifiables, then the general ideas discussed in current paper apply to them just as well as to other modifiables and are thus not repeated here. One problem worth a separate mentioning is about whether to use distributed and possibly implicit representations or not. Having "an ecology of cooperating and competing models, each partially representing some aspects" (Ryan, 2006) may help variation generation both by providing a large set of different combinable elements and possibly by making variations *emerge* even in the course of "normal" system behavior without any explicit generators in place. On the other hand, implicit, distributed and inscrutable internal representations make it difficult to use bootstrap learning processes (Provost, 2007, page 5), so the variations may remain to be generated on a very low level where it rarely leads to very complex solutions due to the vastness of search space down there. Thus some balance suitable for a given system should be searched for.

Regarding the mappings between entities and their representations, there are several issues to be paid attention to. If variation mechanisms are applied to representations (e.g., the genotype), but fitness is mainly dependent on the "real" features deriving from those representations (e.g., the phenotype), then one of the main concerns is the question of whether the representations and mappings allow the mechanisms to properly explore the phenotype space.

The first problem is coverage – which and how big parts of the phenotype are in principle derivable from the genotype. If no representations exist that lead to high-fitness phenotypes, then the variation generator cannot possibly reach them. If, on the contrary, most of the representations lead to only good solutions, then the generator is without much effort very good at producing fit variants, but only as long as the fitness landscape does not change radically with regard to what is covered. Thus in the longer perspective it would make sense to either have full coverage or, possibly even better, to have adaptive representation (or mapping) structure that keeps the coverage on high fitness areas.

Secondly, in addition to the static correspondence between genotype space and phenotype space there is also correspondence of dynamics – how does a movement in one space get reflected in the other. If the mapping is relatively straightforward (e.g., small movements of the modifiable in a certain "direction" generally produce small movements of the testable also in some certain "direction"), then variation generating mechanisms will have the possibility to guide the search in a systematic way. On the other hand, if the mapping is complicated and small changes in genotype space cause significant and difficult to predict jumps in phenotype, then the production of high diversity and large amount of novelty is made easy. Which of these is preferred depends on the particular system and / or situation. Similarly, there is a trade-off involved in the amplification factor: small movements in one space corresponding to small movements in the other makes fine-tuning easy, but small movements corresponding to large ones helps with the rate of exploration in the phenotype space, especially if representations are for some reason difficult to change in large steps.

The third point to be considered is somewhat related to both previous ones: would it be a good idea to have the representations together with mappings form nontrivial generative rules that produce the phenotype in a developmental, step-by-step fashion (as opposed to providing a fully detailed blueprint from which the structures are directly "copied" into reality)? If yes, then should they be deterministic or probabilistic, and context-sensitive or not (or to what extent)? The usage of generative rules can surely make the correspondence between modifiables and testables more complex and thus difficult to guide, but accordingly it can facilitate the production of novel interesting variants that would have been burdensome to explicitly encode in all detail. Then again, if the generative rules make good use of contextual information during execution, and possibly utilize self-organization, they can in principle provide valuable support in channeling the variants into high-fitness regions of solution space, with the almost inseparable flip side of reducing solution diversity. In less fortunate cases the channeling might also occur into low-fitness regions.

And the fourth interesting issue with representations is their abstractness. For example, psychology has observed that the ability to generalize (i.e., to abstract) and transfer knowledge and skills supports (or reflects) system's ability to adapt (Ployhart and Bliese, 2006), and that "greater abstractness is associated with lower stereotypy and greater flexibility in the face of complex and changing problem situations, toward greater creativity, exploration behavior, tolerance of stress, etc." (Harvey and Schroder, 1963, page 134, as referred to by Hunt, 1966). As of variability, the abstractness could be viewed as increasing the scope, or applicability, of each variant and thus reducing the number of different internal alternatives required to cover the areas of interest in phenotype and interaction space. On the other hand, though, abstract representations may be more difficult to interpret, therefore being better suited for advanced systems that possess enough processing capacity and knowledge for transforming between abstract and specific.

Providing Various Internal and External Resources

The generation of variations can also be supported by providing the corresponding mechanisms with an adequate supply of all the necessary and helpful resources. Particularly noteworthy among them are reservoirs of elements that can be used for combinatorial purposes, of prefabricated variants, of ideas, and of accumulated knowledge and experience to be used either directly or more loosely in the form of inspiration. These can be set up as, for example, repositories that can store the components or knowledge either in an explicit and ready-to-use state or also in some more implicit fashion where the full content is not readily extractable but usable nevertheless. The resource pools can also exist as secondary functions of some other subsystems, as well as be totally external. The various ways of using external resources for variation generation include obtaining / copying knowledge and ideas only, acquiring by incorporation of or by merging with external objects, and executing a more interactive process where there exists at least two-way communication between the system and external entities. The lines between these can occasionally be somewhat fuzzy, but the first one is generally thought of as taking place through system's sensory channels, while the second is likely to involve some special intake mechanism and the third can be a combination of the first two with the addition of outward communication. An obvious precondition for using external resources is the very existence of these resources in combination with them being accessible to the system. The latter could be supported by giving the system the necessary interfacing mechanisms, by having some external transportation and communication infrastructure in place, and by other, more elaborate, supportive systems.

Conclusion

Generating variations efficiently and wisely can sometimes be the key for making a system adaptive enough with regard to the goal at hand. And adaptivity, in turn, is one of the key ingredients of life and intelligence. As described in this paper, there are a lot of aspects to be paid attention to in this seemingly simple process of variation generation, and thus both further research of these issues and inventive application of the found ideas can be considered an important part of the fields of ALife and AI, as well as of the studies of Complex Adaptive Systems in general.

Acknowledgements

Thanks to Leo Mõtus for providing me the possibility to do highly unconstrained research. This work was supported in part by Research Laboratory for Proactive Technologies and Department of Computer Control in Tallinn University of Technology, Estonian Doctoral School in ICT, Estonian Information Technology Foundation, and Estonian Ministry of Education and Research (grant SF0140113As08). Conference participation was financed by ESF DoRa 8 via Archimedes Foundation.

References

Andersen, E. (2003). Genesis of an anthill: Wireless technology and self-organizing systems. *Ubiquity*, 3(49).

Ashby, W. R. (1956). *Introduction to Cybernetics*. Chapman & Hall.

Bennett, A. F. (1997). Adaptation and the evolution of physiological characters. In Dantzler, W. H., editor, *Handbook of Comparative Physiology, Vol. 12*, pages 3–16. Oxford University Press.

Brock, J. P. (2000). *The Evolution of Adaptive Systems: The General Theory of Evolution*. Academic Press.

Burke, C. S., Hess, K. P., and Salas, E. (2006). Building the adaptive capacity to lead multi-cultural teams. In Burke, C. S., Pierce, L. G., and Salas, E., editors, *Understanding Adaptability, A Prerequisite for Effective Performance within Complex Environments*, pages 175–211. Elsevier.

Conrad, M. (1983). *Adaptability. The Significance of Variability from Molecule to Ecosystem.* Plenum Press.

Denamur, E. and Matic, I. (2006). Evolution of mutation rates in bacteria. *Molecular Microbiology*, 60(4):820–827.

Edelman, G. M. and Gally, J. A. (2001). Degeneracy and complexity in biological systems. *PNAS (Proceedings of the National Academy of Sciences of the United States of America)*, 98(24):13763–13768.

Grisogono, A.-M. (2005). Co-adaptation. In Bender, A., editor, *SPIE Symposium on Microelectronics, MEMS and Nanotechnology*, pages 23–37.

Grisogono, A.-M. and Ryan, A. (2007). Operationalising adaptive campaigning. In *International Command and Control Research and Technology Symposium, 12th ICCRTS*.

Harvey, O. J., Hunt, D. E., and Schroder, H. M. (1961). *Conceptual Systems and Personality Organization*. Wiley.

Harvey, O. J. and Schroder, H. M. (1963). Cognitive aspects of self and motivation. In Harvey, O. J., editor, *Motivation and Social Interaction: Cognitive Determinants*, pages 95–133. Ronald.

Hersh, M. N., Ponder, R. G., Hastings, P. J., and Rosenberg, S. M. (2004). Adaptive mutation and amplification in Escherichia coli: two pathways of genome adaptation under stress. *Research in Microbiology*, 155(5):352–359.

Heylighen, F. (2001). The science of self-organization and adaptivity. In *The Encyclopedia of Life Support Systems*. Eolss Publishers, Oxford.

Heylighen, F. and Joslyn, C. (2001). The law of requisite variety. `http://pespmc1.vub.ac.be/REQVAR.HTML`.

Holley, J. (2005). Transforming your regional economy through uncertainty and surprise: Learning from complexity science, network theory and the field. In McDaniel Jr., R. R. and Driebe, D. J., editors, *Uncertainty and Surprise in Complex Systems*, pages 165–176. Springer.

Holmqvist, M. and Pessi, K. (2006). Agility through scenario development and continuous implementation: a global aftermarket logistics case. *European Journal of Information Systems*, 15:146–158.

Hunt, D. E. (1966). A conceptual system change model and its application to education. In Harvey, O. J., editor, *Experience, Structure & Adaptability*, pages 277–302. Springer.

Knoll, K. and Järvenpää, S. L. (1994). Information technology alignment or "fit" in highly turbulent environments: the concept of flexibility. In *SIGCPR '94: Proceedings of the 1994 computer personnel research conference on Reinventing IS: managing information technology in changing organizations*. ACM.

Ku, A. (1995). Modelling uncertainty in electricity capacity planning. Doctoral dissertation. London Business School.

Lenski, R. E., Barrick, J. E., and Ofria, C. (2006). Balancing robustness and evolvability. *PLoS Biology*, 4(12):2190–2192.

Lints, T. (2010). The essentials of defining adaptation. In *Systems Conference, 2010 4th Annual IEEE*, pages 113–116.

Menashe, D., Biham, O., Laikhtman, B. D., and Efros, A. L. (2000). Glassy properties and fluctuations of interacting electrons in two-dimensional systems. *Europhysics Letters*, 52(1):94–100.

Moorman, K. and Ram, A. (1992). A case-based approach to reactive control for autonomous robots. In *Proceedings of the AAAI Fall Symposium on AI for Real-World Autonomous Mobile Robots*.

Ployhart, R. E. and Bliese, P. D. (2006). Individual adaptability (I-ADAPT) theory: Conceptualizing the antecedents, consequences, and measurement of individual differences in adaptability. In Burke, C. S., Pierce, L. G., and Salas, E., editors, *Understanding Adaptability, A Prerequisite for Effective Performance within Complex Environments*, pages 3–39. Elsevier.

Provost, J. (2007). Reinforcement learning in high-diameter, continuous environments. Doctoral dissertation. The University of Texas at Austin.

Ram, A. and Santamaría, J. C. (1997). Continuous case-based reasoning. *Artificial Intelligence*, 90:25–77.

Ryan, A. (2006). About the bears and the bees: Adaptive responses to asymmetric warfare. In Minai, A., Braha, D., and Bar-Yam, Y., editors, *Proceedings of the Sixth International Conference on Complex Systems*. New England Complex Systems Institute.

Stagl, K. C., Burke, C. S., Salas, E., and Pierce, L. (2006). Team adaptation: realizing team synergy. In Burke, C. S., Pierce, L. G., and Salas, E., editors, *Understanding Adaptability, A Prerequisite for Effective Performance within Complex Environments*, pages 117–141. Elsevier.

Stec, B. (2004). Living and nonliving matter. *Science*, 305:41. Letter to the editor.

Unewisse, M. and Grisogono, A.-M. (2007). Adaptivity led networked force capability. In *International Command and Control Research and Technology Symposium, 12th ICCRTS*.

Weick, K. E. (2005). Managing the unexpected: Complexity as distributed sensemaking. In McDaniel Jr., R. R. and Driebe, D. J., editors, *Uncertainty and Surprise in Complex Systems*, pages 51–65. Springer.

Weick, K. E. (2007). Drop your tools: On reconfiguring management education. *Journal of Management Education*, 31(1):5–16.

Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media, Champaign, IL.