

Response to comments on my article "Genetic Programming and Emergence"

Wolfgang Banzhaf

Received: date / Accepted: date

At the outset, I would like to thank all colleagues who took the time to read my opinion piece "Genetic Programming and Emergence", thought about it, and finally wrote their comments. I am very happy about these comments, as they provide clear evidence that my contribution has been well received by the community. I shall next take up these comments in alphabetic order of their authors, briefly summarize the gist of their argument and provide for each one what I hope is a useful response.

Lee Altenberg [1] emphasizes that emergent phenomena in Genetic Programming are produced in a formal system. To bolster this argument, he uses the mathematical structures of a matrix and a vector as an even simpler example of an interacting structure, whose formal manipulation leads to emergent phenomena in the mathematical world examined by "Linear Algebra". As examples he mentions eigenvalues, spectral radii, etc. Thus, he argues, that an entire field of Mathematics occupies itself with emergent phenomena in a formal system, and he rightly is convinced that other areas of Mathematics do the same.

Having thus fortified the notion of emergent phenomena and their respectability in science, he states that emergent phenomena in Genetic Programming could also be treated in formal or mathematical terms, and possibly lead to interesting discoveries. He concedes that such an undertaking might be more difficult, due to the complexity of interactions in Genetic Programming that give rise to emergence, but he points to one example from the last decade where something along these lines has been achieved [10]. In his own words: "But this unruly space is no less formal or no less observable than the simple matrix, and their mathematical properties are simply awaiting our discovery". Well spoken.

Where Altenberg notices a slight disagreement with me is in regard to the role of selection and variation operators. I am not sure whether there really is much of a disagreement, but his comment sheds some light on an interesting aspect of evolution, that has been passionately discussed [6]. The argument states that "selection and the genetic operators are co-equal partners in the creation of these emergent phenomena" [1]. I agree with this statement, as Figure 2 of my contribution demonstrates. The diversity generator is of constituting importance in generating the variety of struc-

ture and function from which a selection device can then select and amplify the most appropriate.

To clearly state it again: Selection (top-down causation) can only work on variation if the latter is offered, otherwise it will become powerless. But I do see a difference between variation and causation. Variation can be said to provide "input" to the mechanisms of bottom-up causation which is constituted by the mapping from genotypes to fitness.

Anika Ekart [4] agrees with the usefulness of emergence in Genetic Programming but is not satisfied with the current state of affairs in GP and asks what we can do to get to what she calls "perpetual emergence". What she has in mind here, I believe, is what others have called "open-ended evolution", i.e. a continuous process that would allow a GP system to adapt to new and unforeseen tasks, a system that permanently learns.

She points out that systems equipped with some sort of self-reflection are necessary to arrive at more interesting and more powerful GP models and goes on to cite a few ideas from the literature that might lead the way. Overall she is in agreement with my observation of emergent phenomena in GP, and calls to exploit them.

In response, I would agree that some sort of internal or self-reflecting mechanism is key to systems that are open-ended and continue to produce novelty. We cannot expect perpetual emergence if we define a fixed fitness function. While that will allow adaptation of a given structure, it does not allow the invention of higher level structures with correspondingly more complex functions. We might be able to build a system equipped with a number of lower levels before it comes to fitness evaluation (the implementation of genotype-phenotype maps, perhaps even involving artificial regulatory networks, is an example), but we cannot expect the system to do anything outside its once defined fitness function. The term "open-ended evolution" can be easily misunderstood, in particular in artificial evolution because we are so accustomed to using an explicit fitness function in those approaches.

Krzysztof Krawiec [7] discusses in his contribution the role of interaction in emergent phenomena and zooms in on the interaction between program code and input of a GP system as the key to the emergence of solutions. He remarks that most of what GPs have been doing to date was concerned with structure rather than with semantic content/functions. "But a GP system cannot be fully understood without them [semantic aspects] ...".

He then argues that the semantics of a program cannot be seen only on the level of the program as a whole, but must be assigned at least partially, to parts of the program, for instance, to modules or larger building blocks gradually emerging from randomness. He states: "Nevertheless - and this is the pivot of my argument - there is room for emergence *during* program execution." Thus, Krawiec wants to (re-)use intermediate results. I agree that this approach is fruitful. However, the premise that semantics cannot emerge on the level of the program as a whole seems to be a misunderstanding of what I said. While in Section 4 of my contribution I indeed refer to various phenomena of emergence I did never state that the entire solution did not emerge. The presence of selection that selects for high fitness does not rule out to state that highly fit solutions *emerge*. The point is that fitness selection does not prescribe ("instruct") what a program has to look like. There is considerable freedom to realize high fitness with different program code. Agreed, the fitness itself does not emerge, but the programmatic solution that is able to produce this fitness can be said to emerge.

I fully agree that it is no surprise that selection for high fitness results in highly fit individuals. Yet, the actual way these individuals realize their high fitness is emergent, and often surprising, especially when, upon running a GP system we realize that it has found a way to wiggle out of our tight descriptions and found a shortcut to high fitness we deem "illegal". In fact, emergence is the explanation for this tendency of evolutionary systems.

Andre Leier [8] comments that it is the multi-scale nature of evolution in time and space that makes evolution such a powerful paradigm: "From genes to proteins, to organelles, to cells, to tissue, to organs, and complete organisms: Between genome and phoneme there are many spatial and temporal scales, and functional levels and emergent phenomena can be observed at each of these." He then asks whether additional levels in living organisms are a means to increase adaptability and success of survival. I fully agree with this point of view, but think that multi-level systems in GP have so far not garnered the attention they deserve.

How can we allow GP systems to foster emergent levels of organization, both at higher levels from where we currently have defined "fitness" and at lower levels where we need to learn to compose solutions in different ways? How can we open up the scales of time and space for evolution? As *Bonner* points out [3], a key discovery of natural evolution was that greater size of an organism would allow it to occupy empty ecological niches. This, in the end, brought about the success of multicellularity. Clearly, a change in spatial scale of evolution is involved. But it is interesting to note that multi-cellularity also involves an increase in the range of control over temporal scales: Development requires epigenetic interactions which control dynamic, time-dependent switches to provide for structural control [5].

Leier also asks whether we need to start studying what he calls evolutionary risk management, perhaps in the form of repair mechanisms that are itself under the control of the artificial evolutionary process. This is an interesting idea, however, one might argue that what is studied in general under the heading of "parameter self-adaptation" in evolutionary algorithms is an abstraction of this idea.

David Montana's comment [9] is one of the most critical. He is skeptical about the claim that we find emergence in artificial systems such as a GP system. *Montana* speaks of "top-down emergence", a term that I did not use nor understand. Rather than interpreting what *Montana* means, I would respond to the strongest statement he makes: "... the emergent system must be self-contained. When an external driver causes emergence in a system, this should not be considered an example of top-down emergence."

Thus, *Montana* starts his argument by making a distinction between the system in which emergence occurs, and the "driver" of that system. My response would be that every system must have a boundary, otherwise it would not be a system, separated from its environment. So we can agree that there will always be insides and outsides to a system. I am, however, unclear what *Montana* means by the driver. He alludes to it with the example of the development of flocking behaviour, but I am afraid this confuses me even more.

I guess that if we accepted *Montana's* argument, we would have to deny that breeding foxes for tamability [12], i.e. the artificial evolution of mammals for a particular trait, an experiment which produced a number of emergent side-effects, would also not count as emergent. Yet in this particular example we clearly see that both intended and unintended features are emergent. This happens despite of the fact that "the driver"

is outside the system. A case in point is that Darwin based a substantial portion of his writing about evolution by natural selection on the knowledge of breeders.

In summary, let me refute David Montana's argument by stating:

1. There is no such process as "top-down emergence". There is only emergence, a process by which "global" phenomena appear, produced by the interplay of bottom-up and top-down causation.
2. In a variation - selection mechanism, top-down causation is realized by selection, regardless of what or who is doing the selection.
3. The actual material of what is under emergence is not relevant, as long as that material has the potential to provide a rich set of variations.

The only point I am willing to concede is that emergence might be richer in systems with access to the level structure of itself, something Ekart calls "internal emergence". I have spoken of this in the context of multi-level systems and consider it empowering, so it is not really a concession. The difference is a matter of degree but not of principle.

Moshe Sipper [11] points out his work on a test for emergence which involves a scientist in a simulation-based experiment who would (i) design the system, (ii) observe the system's behaviour and (iii) compare the observations with his expectations. Emergence in this view can be tied to the fact that "there is a cognitive dissonance between the observer's mental image of the system's design stated in [a language] \mathcal{L}_1 and his contemporaneous observation of the system's behaviour in [a language] \mathcal{L}_2 ."

Sipper states that this approach to emergence could be considered complementary to my approach.

Let's for a moment delve into the inner workings of his idea. A gut reaction to trying to define a "mental image" is that it might be difficult to come up with something objectively measurable in such a context. However, we are assured by the existence of languages that at least something formal can be stated about this mental image. The key, however, to Sipper's emergence test is the existence of a second language, with the second one - sic! - describing global behaviour, most likely over a longer time-scale than individual events which are the subject of the first. By changing scope (both in space and time), this second language is able to capture other patterns than the former. So a difference ("cognitive dissonance", in Sipper's words) between these descriptions is plausible to exist.

Sipper's test raises a number of questions:

1. If the languages are constructed as described, isn't there always a difference ("cognitive dissonance") in the description of some phenomenon, by the definition of these languages?
2. If there are bound to be differences anyway, the test hinges on "results are not obvious". But that is in the eye of the beholder, and isn't it therefore subjective? So are we not back to where we started without the formal description in \mathcal{L}_1 and \mathcal{L}_2 ?
3. Suppose, the test can be done. Is it only valid for simulation environments? We hear claims that emergence only happens in Nature (I dispute such a claim) [2], but shouldn't the test work in natural systems in particular? What about the "formalizability" of these natural systems, is it reasonable to expect that at all?

Peter Whigham [13] was inspired to demonstrate by simulation experiment that the notions in my opinion piece are only part of the story. The gist of his argument is that selection is not fundamental to the process of emergence. He states: "Therefore

selection is just an amplifier of patterns and that the emergence of such patterns is a fundamental and inevitable consequence of recombination acting on variable-length representations.” He goes on to demonstrate this by studying the development of the number of repeated patterns, an emergent phenomenon claimed by my essay to be the result of its mechanisms.

In particular, he studies the emergence of repeated patterns under drift and recombination, but in the absence of selection. He uses a variable length representation, and corrects for the expected number of repeated sequences that might appear randomly. The result of his study is that repeated patterns emerge even without selection. As a proxy for a selection process, he uses a system with a smaller population. His conclusion is that (i) either repeated patterns are not emergent, or (ii) the mechanism by which they appear has nothing to do with selection.

My response is that (a) a smaller population is only an approximation for selection. The experiment would better be repeated using a population of the same size as the original population. Most likely, the distribution will look different under those conditions; (b) I am happy to concede that even under drift and recombination only (without selection) an effect might be visible. Perhaps this can be termed ”weak emergence”. This effect, however, is usually orders of magnitude smaller than an effect under selection. In addition, repeated patterns are only under weak selection pressure in Genetic Programming. Finally, there are more mechanisms that lead to emergent phenomena. I have not stated that selection is the only way to produce emergence.

Acknowledgements

I am indebted to Lee Spector, the Editor-in-Chief of Genetic Programming and Evolvable Machines, for proposing this contribution format and for handling the process of submission so diligently.

References

1. ALTENBERG, L. Commentary on ”Genetic Programming and Emergence” by Wolfgang Banzhaf. *Genetic Programming and Evolvable Machines 14* (2013), tba.
2. BERSINI, H. Emergent phenomena belong only to biology. *Synthese 185* (2010), 257 – 272.
3. BONNER, J. *First Signals: The evolution of multicellular development*. Princeton University Press, 2001.
4. EKART, A. Emergence in Genetic Programming - Let’s exploit it! *Genetic Programming and Evolvable Machines 14* (2013), tba.
5. JABLONKA, E., AND LAMB, M. *Evolution in four dimensions*. MIT Press, 2005.
6. KAUFFMAN, S. *The Origins of Order*. Oxford University Press, 1993.
7. KRAWIEC, K. Genetic Programming: Where Meaning Emerges from Program Code. *Genetic Programming and Evolvable Machines 14* (2013), tba.
8. LEIER, A. Emergence in simulated evolution. *Genetic Programming and Evolvable Machines 14* (2013), tba.
9. MONTANA, D. A Response to ”Genetic Programming and Emergence”. *Genetic Programming and Evolvable Machines 14* (2013), tba.
10. POLI, R., LANGDON, W., AND DIGNUM, S. On the limiting distribution of program sizes in tree-based genetic programming. In *Proceedings EuroGP 2007* (2007), LNCS, Springer.
11. SIPPER, M. Commentary on ”Genetic Programming and Emergence”. *Genetic Programming and Evolvable Machines 14* (2013), tba.
12. TRUT, L. Early Canid Domestication: The Farm-Fox Experiment. *American Scientist 87* (1999), 160.
13. WHIGHAM, P. Response to ”Genetic Programming and Emergence” by Wolfgang Banzhaf. *Genetic Programming and Evolvable Machines 14* (2013), tba.