Quantifying Viability

Matthew Egbert[1,2,3], Juan Pérez-Mercader[1,4]

[1]Harvard University, Dept. of Earth and Planetary Sciences, Cambridge, 02142, USA
[2]University of Auckland, Dept. of Computer Science, Auckland, 1142, New Zealand
[3]EON/ELSI Tokyo Institute of Technology, Ookayama, Meguro-ku, Tokyo, 152-8550, Japan
[4]Santa Fe Institute, Santa Fe, USA
mde@matthewegbert.com, jperezmercader@fas.harvard.edu

Living and life-like systems vary in *viability*. They are alive or dead, healthy or unhealthy, getting better or worse, or dying. Despite the ease of applying these descriptions informally, there do not yet exist general methods for richly quantifying viability or health in such systems —even when every aspect of the system is available for experimental variation and measurement. Nevertheless, for a given system of interest, it is sometimes possible to distinguish between states where the system will persist for the foreseeable future (there are termed *viable states*) and those where it will not. This is perhaps the most basic, binary classification of states in terms of viability and it can be used to identify different regions in 'viability space' (see Figure 1 and Barandiaran and Egbert, 2013). An improved measure would make it possible to not just categorize systems but to compare the relative viability of two states that are in the same category, i. e. that are both expected to persist or both expected to die. This type of measure would make it possible to identify whether a system is becoming more or less viable, or to evaluate the influence of a given external perturbation upon viability, thereby enhancing our ability to understand and influence the viability of complex life-like systems.

One way to formulate such a measure is to assume that the system of interest is subjected to unpredictable fluctuations that perturb its autonomous dynamics. If this is the case, the argument can be made that the farther away a viable system is from the viability-interface (the surface between the viable and non-viable regions of viability space), the *more* viable it is, as there is a smaller set of perturbations that will cause the system to become non-viable. (A similar argument can be used to describe non-viable systems as being more and more non-viable as their distance from the viability boundary increases.) There is a problem however: the dimensions of viability space (i. e. the essential variables) are almost always measured in entirely different units, and these units have no relation to viability. As an example, an organism might require a specific range of temperature to survive and a specific range of atmospheric pressure. It should be clear that the units for measuring these phenomena do not relate to viability and nor do they relate to each other. A perturbation of 3 atmospheres will in general not have the same influence on viability as a change of 3 degrees! Further work is needed if we are to develop a meaningful measure of distance in viability space. [1]

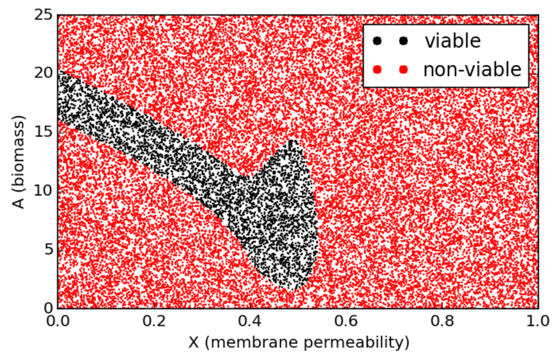

Figure 1: Viability class for various initial conditions in a simple two-dimensional model of a bio-reactor. Randomly sampled initial conditions plotted in red do not survive, whereas those plotted in black do. Details of the model are not relevant and are not presented in this abstract.

In a soon-to-be submitted paper, we have proposed a method that uses the shape of the viability interface to rescale the system's essential variables so as to define a normalized viability space, where a perturbation of a given magnitude has the same likelihood of crossing the viability regardless of the direction of the perturbation. The method works by calculating the extent to which the viability-interface "faces" each dimension and then scaling the values in that dimension by this amount. More formally, for each dimension of viability space, X, we identify I_X, the average magnitude of the X-component of the viability-interface surface normals:

$$I_X = \frac{\iint_I ||\hat{\mathbf{n}} \cdot \hat{\mathbf{e}}_{\mathbf{X}}||\, dI}{I}, \tag{1}$$

[1] This problem was first brought to Egbert's attention in a seminar given by Nathaniel Virgo and Simon McGregor at the University of Sussex in or around 2009.

and use this value to rescale values into normalized units, thus: $\hat{x} = I_X x$. In the above Equation, $\hat{\mathbf{e}}_\mathbf{X}$ is the basis unit-vector for dimension X, and $\hat{\mathbf{n}}$ is the surface normal of I, the viability-interface.

In normalized viability space, there is a meaningful minimal distance between any given state and the viability interface: on average over initial-conditions, a perturbation of a given magnitude will have an equal chance of crossing the viability-interface *regardless of the angle of the perturbation vector*. In other words, a perturbation of given magnitude in normalized viability space has the same chance of transforming a randomly selected viable state into a non-viable state (or *vice versa*) whether it is a perturbation of one essential variable (e. g. pressure) or another (temperature), or a combination thereof. Figure 2 shows the same system as Figure 1, but plotted in normalized viability space, with shading used to indicate a viability gradient based upon the minimum distance to the viability interface.

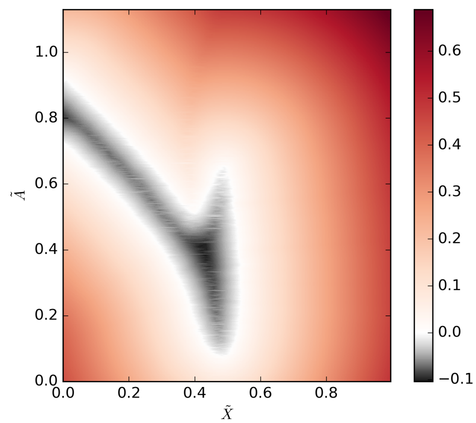

Figure 2: The same system, plotted in a normalized viability space with signed-distances to the viability interface indicated. The lower the value, the healthier the system, with negative values corresponding to viable states, i. e. states that in the absence of external perturbation, are expected to persist for the foreseeable future.

Normalizing viability space in this way allows us to compare states in terms of their relative viability. This in turn allows us to describe how a system's viability is changing over time. When additional information is available concerning the system's autonomous dynamics, and/or the cost/difficulty of influencing the system's essential variables, it is possible to make additional observations relevant to the system's viability, such as to identify the future state from which the minimum perturbation is necessary to cross the viability interface.

Using information theoretical analysis, it is also possible to identify correlation between variables and these measures of viability. This allows us to identify and evaluate the quality of *viability indicators*, variables that are good at predicting a system's viability. This connects with some of our previous work, where we have shown how an organism can respond to their own viability-indicators, and in so doing become capable of (i) adapting to phenomena neither it nor its ancestors have ever previously experienced (Egbert et al., 2010); and (ii) adapting to changes in its own needs and abilities, resulting in a more evolvable organism (Egbert et al., 2011; Egbert and Pérez-Mercader, 2016).

Within the enactive approach (Stewart et al., 2010), the concept of viability has been used to naturalize concepts of adaptivity, agency and normativity. In particular, Di Paolo (2005) compares trajectories in terms of their dynamics relative to the viability boundary to formulate a definition of adaptivity. In a previous publication, we presented an argument showing how an organism's viability can be used to develop a naturalized concept of normativity (Barandiaran and Egbert, 2013). The research presented herein extends these works, providing a way to normalize viability space and compare states in terms of viability and to measure distance from the viability boundary.

More broadly, identifying viability-indicators in natural systems could improve our ability to predict or influence their viability, and similarly identifying high quality viability-indicators in synthesized protocells will allow us to better understand how to create artificial life-forms that are capable of surviving in the diverse conditions found outside of tightly controlled laboratory environments.

Acknowledgments. This publication was supported by the ELSI Origins Network (EON), which is supported by a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of EON or the John Templeton Foundation.

References

Barandiaran, X. E. and Egbert, M. D. (2013). Norm-Establishing and Norm-Following in Autonomous Agency. *Artificial Life*, 20(1):5–28.

Di Paolo, E. A. (2005). Autopoiesis, Adaptivity, Teleology, Agency. *Phenomenology and the Cognitive Sciences*, 4(4):429–452.

Egbert, M. D., Barandiaran, X. E., and Di Paolo, E. A. (2010). A Minimal Model of Metabolism-Based Chemotaxis. *PLoS Comput Biol*, 6(12):e1001004.

Egbert, M. D., Barandiaran, X. E., and Di Paolo, E. A. (2011). Behavioral Metabolution: The Adaptive and Evolutionary Potential of Metabolism-Based Chemotaxis. *Artificial Life*, 18(1):1–25.

Egbert, M. D. and Pérez-Mercader, J. (2016). Adapting to Adaptations: Behavioural Strategies that are Robust to Mutations and Other Organisational-Transformations. *Scientific Reports*, 6:18963.

Stewart, J., Gapenne, O., and Di Paolo, E. A. (2010). *Enaction: Toward a New Paradigm for Cognitive Science*. The MIT Press.