sible experiments or observations ("choice space"). If we look more deeply into the machinery of this complex process, we see a large number of agents, i.e., scientists, working on their individual proposals for the correct $\mathcal{L}$. In that sense, they play the role of the walkers for the physics community agent. Based on the models already developed, there is a strong preselection for $\mathcal{L}$, similar to the distinction between green and red walkers in Figure 2 of the target article. However, they are following FSX themselves. Their individual number of future choices is determined by academic success, e.g., by the number of citations, or by the number of interviews offered later for academic positions. This space is, again, entirely different from the text space in which the Lagrangians ($\mathcal{L}$) dwell. Each of these walker-Lagrangians is intended by the author to be published, which means that its position in text space must differ sufficiently from the positions of all other walker-Lagrangians. This is very similar to the agents in Figure 1, above, where no two agents can come too close. At the same time, it is favorable in view of the number of citations one might expect to keep in close proximity to similar published walker-Lagrangians. As a result, scientists tend to assemble in swarms studying similar lines of attack to a problem, and this can be empirically observed (Shifman 2012). We can conclude that the development of physical theory follows FSX on several levels.

« 9 » As the search for the correct Lagrangian continues, it has acquired more and more terms over time. As of today, the Lagrangian that tries to unify all interaction forces (except gravity, however) consists of more than 150 terms and takes about a full page to be written down legibly (Shivni 2016). Since each term (or group of terms) enables completely new classes of experiments to be accounted for, the total number of choices in choice space is multiplied by the number of new experiments a term allows for. Hence, choice space increases roughly exponentially with the number n of terms. Hence, predictive density scales as e^n/n. As this is monotone and strongly increasing, it tends to the formation of larger and larger formulas for $\mathcal{L}$, such that there is no convergence to a "beautiful" (i.e., simple) theory in sight. The search for mathematical "beauty" is a major driving force in this field of physics (Hossenfelder 2018) and rests on the idea that the beauty we see in theories of nature can come only from the inherent beauty of nature itself. This reasoning, however, is logically stringent only from a strictly positivist point of view. A constructivist approach would rather appreciate that the theory can only describe a version of "nature" constructed by an observer. This version may then well (or, rather, is bound to) have inherited structure from the observer through the process of symbolization. Hence the "beauty" may well reside exclusively "in the eye of the beholder," to use a proverbial phrase.

References

Foerster H. von (1972) Perception of the future and the future of perception. Instructional Science 1: 31–43. ▸ https://cepa.info/1647

Halzen F. & Martin A. D. (1984) Quarks and leptons. John Wiley & Sons, New York.

Hornischer H., Herminghaus S. & Mazza G. M. (2019) Structural transition in the collective behavior of cognitive agents. Scientific Reports 9: 12477.

Hossenfelder S. (2018) Lost in math: How beauty leads physics astray. Basic Books, New York.

Shifman M. (2012) Frontiers beyond the standard model: Reflections and impressionistic portrait of the conference. Modern Physics Letters A 27(40): 1230043.

Shivni R. (2016) The deconstructed Standard Model Equation. Symmetry Magazine: A Joint Fermilab/SLAC Publication: 07/28/16. https://www.symmetrymagazine.org/article/the-deconstructed-standard-model-equation

Stephan Herminghaus obtained his PhD in Physics at the University of Mainz in 1989. In 1998 he became full professor at the University of Ulm. Today he is a director at the Max Planck Institute for Dynamics and Self-Organization in Göttingen. In his research, he performs experiments, simulations, and analytic calculation to gain insight into phase transitions and other collective phenomena in complex and active matter. This includes systems as diverse as granular flows, biological matter, and traffic systems. His overarching interest is how a sustainable anthropocene may be organized.

Received: 16 October 2020
Accepted: 31 October 2020

Control and Behavior – Spoiled for Options? Coping with Uncertainty

Elena Rovenskaya
International Institute for Applied Systems Analysis, Laxenburg, Austria and Lomonosov Moscow State University, Russia
rovenska/at/iiasa.ac.at

Nikita Strelkovskii
International Institute for Applied Systems Analysis, Laxenburg, Austria • strelkon/at/iiasa.ac.at

> Abstract • While we agree with Hornischer et al. that there is an apparent similarity between the Future State Maximization principle and von Foerster's Ethical Imperative, we suggest that the former can be seen as a computational overlap with the latter rather than its computational interpretation. Further, we reflect on the connection between these two concepts and the theory of robust control.

Handling Editor • Alexander Riegler

« 1 » In their target article, Hannes Hornischer et al. reflect on the Future State Maximization (FSX) approach for describing the behavior of multi-agent systems developed earlier. The authors highlight the difference between FSX and traditional machine learning (ML) methods as FSX is forward-looking and explores possible future options rather than being backward-looking and relying on past data, which is common in ML. In this way, FSX follows the spirit of reinforcement learning (RL). The main claim of the target article is that FSX is a computational interpretation of Heinz von Foerster's Ethical Imperative (HF-EI) that reads: "Act always so as to increase the number of choices." However, we suggest a more nuanced approach in this regard. We suggest that FSX is not merely a computational interpretation of HF-EI. Rather, both FSX and HF-EI can be seen as overlapping projections from different realms onto the space of agents' actions. Agents can be either artificial agents in a computer or entities in a natural setting. In the next paragraphs, we elaborate on this proposition in detail.

« 2 » First, both FSX and HF-EI focus on the intrinsic motivation of agents to have a larger number of options available. HF-EI is indifferent with respect to other possible objectives that the agent might have, which may be in contradiction with it. In other words, as an ethical imperative, HF-EI can ignore the other agent's objectives. It is not evident that FSX can or should be positioned in the same way. In §§19f of the target article, an example is presented where the agent's behavior is being determined solely by FSX. In this example, FSX and HF-EI conform with each other. However, in several other examples, the agent or a group of agents act so as to achieve a certain practical objective. In Example 1, the agent aims to collect a higher reward. In Example 2, robots aim to collect information about the environment.

« 3 » As an imperative, HF-EI supposedly dominates over any other possible objective if the agent has one. In contrast, the relations between FSX and any other objective(s) that an agent might have can be different. In Examples 1 and 2, for instance, the future options included in FSX seem to be only those that are *compatible* with the other objective of the system, rather than *any* options. It is a very important aspect that challenges the perception of FSX

« 4 » Furthermore, while HF-EI calls the agent to consider any options, in the examples of the FSX application presented in the target article, these options were constrained by the model design. For instance, in the example considered in §§19f, available options are going North/East/South/West or staying put. We would expect that an agent that had an intrinsic motivation according to HF-EI would "discover" also the possibility of diagonal moves (e.g., going North-East).

« 5 » The next issue to which we would like to draw attention is that HF-EI does not specify whose number of options should be maximized – the agent's own options of those of other agents in its environment. In many cases, there is a tradeoff between the number of one's own options and the number of options of others. The FSX implementations, however, all seem to focus on the number of options of the considered agent only (Charlesworth & Turner 2019: 15363).

Figure 1 • A simple homing guidance problem for two points in one-dimensional space.

« 6 » HF-EI emerged in the 1970s in a certain cultural, economic, and social context. As any ethical concept, it cannot pretend to be absolute. Ethics is a product of a choice of humans that is, at least to some degree, conscious. For example, in the 21st century, humanity broadly embraced the idea of sustainability. The prominent milestone of that is the adoption of seventeen Sustainable Development Goals (SDGs) by the UN by all countries in 2015. According to the definition given by the Brundtland's Commission, sustainable development meets the needs of the current generation without compromising the next generations with regard to the possibility of meeting their needs.[1] Formal computational models are considered as one of the ways to better grasp uncertainties of the future (Riegler 2015; Gash 2020: §33). The commitment of the current generation to leave to the next generations as many possibilities as possible to be able to satisfy their needs seems to be compatible with both EI and FSX.

« 7 » On the other hand, there are examples in nature where organisms seem to follow the FSX principle without relating themselves to any ethical considerations. For example, stem or tumor cells maximize the number of future options available to them – i.e., to stay undifferentiated as long as possible; differentiation to a specific cell type implies reducing the number of future options. In both cases, maximizing the number of options seems to be a survival strategy rather than ethics.

1 | Cf. "Report of the world commission on environment and development: Our common future," United Nations General Assembly, 1987. http://www.un-documents.net/our-common-future.pdf

« 8 » In practice, maintaining a large number of options is costly. Cost should be understood in the broad sense in relation to searching for, storing, assessing, and deciding on these options. Hornischer et al. touch upon this issue in §4, mentioning the costs of maintaining a model; we suggest that this issue requires greater attention. It is a significant challenge both for living creatures and artificial agents. For example, it is known that humans who have to choose from many options are prone to decision fatigue (Vohs et al. 2008). Artificial agents would also require more power, memory, and time to process a higher number of options. Thus, we suggest introducing some kind of a "flexibility penalty" for a higher number of options to be processed.

« 9 » Furthermore, maximizing the number of options for oneself can incur costs for the other agents operating in the same environment. As we discussed in §6 above, it can also put constraints on the number of options of the others, which raises broader ethical considerations.

« 10 » In practical implementations of FSX, it may not always be necessary or even desirable – if costs are taken into consideration – to maximize the number of options. We argue that the preferred multiplicity of options is tightly connected with the amount of information available to the agent who makes a decision. The less information she has, the broader the space of options she may need to keep open. However, when the agent is close (in some sense) to achieving her goal, she may not need to keep many options – given that doing so implies high costs. Moreover, trying to maximize the number of options, she can even miss the target. As an illustration, in the next two paragraphs, we refer to two examples.

« 11 » The first example is from chess. In chess, when making a move, players typical-

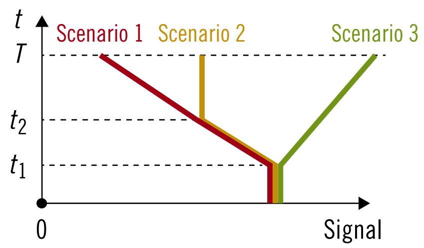

Figure 2 • A tree-like information structure that drives an agent's decisions. Different scenarios are schematically denoted by different colors.

ly strive to maximize the number of available options. A larger set of available options makes winning the game more likely, and vice versa (Charlesworth & Turner 2019: 15362). An empty set of available options indicates a checkmate or a stalemate for a player. However, when a player clearly sees a way to checkmate her opponent unavoidably, there is no reason for her to maximize the number of future options. Rather she should merely execute the order of the winning moves. If, on the contrary, she wished to make a move that would increase the number of available options, she might lose the advantage of the game.

« 12 » The second example comes from the theory of robust control. Consider a simple homing guidance problem. The agent operating in a one-dimensional space aims to reach a target by a given time *T* by steering its acceleration. However, the agent does not know from which of the two different points, *A* and *B*, it has started. Points *A* and *B* are equally remote from the target (see Figure 1). To be able to select a successful control, the agent needs to observe a signal from the environment to identify its position in relation to the target. Assume that until time moment *T/2*, the agent is not able to receive any information. From *T/2* on, the agent receives a signal that allows it to identify the initial point without ambiguity. The optimal course of action in this case is to stay put while no information is available and, once the information starts to arrive and the initial point has been identified, to accelerate left or right, accordingly. Acting thus, the agent leaves herself the maximum number of the options available on the "no information" stretch and chooses the only option on the "full information" stretch. Any deviation from such a strategy, e.g., to accelerate in any direction on the "no information" stretch or stay put once the full information is known, would lead to a suboptimal resource use or even to a failure to achieve the target set. This example is similar to the Buridan's ass paradox, illustrating a problem of decision making between equally attractive options and a necessity to make a choice to achieve the objective (Rescher 1960). The presence of information about the options plays a decisive role in this example.

« 13 » In the agent homing guidance problem described above and many other, more complex problems of this kind, before executing any action, the agent needs to create a portfolio of controls allowing it to achieve its objective in each possible scenario of the uncertain environment and other agents' actions. This means that relying on its model of the problem, the agent explores the solution space corresponding to these scenarios. This exploration yields a tree-like information structure that drives decisions (Figure 2). Bifurcation points correspond to time moments in which new information arrives (i.e., t_1, t_2). In the bottom part of the tree, which has only few branches, the agent has little information. Consequently, she has to choose an action from a narrow set that contains actions leading to the target in all plausible scenarios under this available information. As time goes on and more information becomes available, the agent progresses through the bifurcation points to branches corresponding to the unfolding scenario (Strelkovskii 2015). This exploration procedure is very similar to the concept of "walkers" from the target article (§21). Note that in both FSX and the theory of robust control, we do not use any knowledge regarding other agents and their objectives. We only rely on the knowledge of how the actions of the other agents impact the objective function of the considered agent.

« 14 » As the last remark, we wish to support the authors' reflection that FSX "implies mitigating uncertainty through the increase of uncertainty" (§5). Indeed, keeping more options available often allows agents to reduce risk. Portfolio diversification of investment is a prominent example (Wagner & Lau 1971).

References

Charlesworth H. J. & Turner M. S. (2019) Intrinsically motivated collective motion. Proceedings of the National Academy of Sciences 116(31): 15362–15367. https://www.pnas.org/content/pnas/116/31/15362.full.pdf

Gash H. (2020) Constructivism, fast thinking, heuristics and sustainable development. Constructivist Foundations 16(1): 1–12 (this issue). ► https://constructivist.info/16/1/001

Rescher N. (1960) Choice without preference: A study of the history and of the logic of the problem of "Buridan's ass." Kant-Studien 51(1–4): 142–175.

Riegler A. (2015) Knowledge and belief: Some clarifications. Cybernetics and Systems 46: 484–509.

Strelkovskii N. (2015) Constructing a strategy for the guaranteed positioning guidance of a linear controlled system with incomplete data. Moscow University Computational Mathematics and Cybernetics 39(3): 126–134.

Vohs K. D., Baumeister R. F., Schmeichel B. J., Twenge J. M., Nelson N. M., & Tice D. M. (2008) Making choices impairs subsequent self-control: A limited-resource account of decision making, self-regulation, and active initiative. Journal of Personality and Social Psychology 94(5): 883–898.

Wagner W. H. & Lau S. C. (1971) The effect of diversification on risk. Financial Analysts Journal 27(6): 48–53.

Elena Rovenskaya is Program Director of the Advanced Systems Analysis (ASA) Program, International Institute for Applied Systems Analysis (IIASA) and Research Scholar at the Faculty of Computational Mathematics and Cybernetics, Lomonosov Moscow State University, Russia. Her research interests lie in the areas of socio-environmental modeling, decision support, and optimization.

Nikita Strelkovskii is a Research Scholar in the Advanced Systems Analysis (ASA) Program, International Institute for Applied Systems Analysis (IIASA). He has a background in control theory. He develops agent-based models and qualitative systems analysis methods.

Received: 24 October 2020
Accepted: 29 October 2020