PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Paths to non-deterministic autonomy: practical and qualitative considerations towards a Hawking-Musk-esque nightmare

Wolfgang Fink

Wolfgang Fink, "Paths to non-deterministic autonomy: practical and qualitative considerations towards a Hawking-Musk-esque nightmare," Proc. SPIE 10982, Micro- and Nanotechnology Sensors, Systems, and Applications XI, 1098225 (17 May 2019); doi: 10.1117/12.2518335

Event: SPIE Defense + Commercial Sensing, 2019, Baltimore, Maryland, United States

Invited Paper

Paths to Non-Deterministic Autonomy:
Practical & Qualitative Considerations towards a Hawking-Musk-esque Nightmare

Wolfgang Fink*[a]

[a]Visual and Autonomous Exploration Systems Research Laboratory, College of Engineering, University of Arizona, 1230 E Speedway Blvd, Tucson, AZ 85721, USA

ABSTRACT

Current computational approaches, such as Artificial Intelligence, artificial neural networks, expert systems, fuzzy logic, fuzzy-cognitive maps, other rule-based approaches, etc., fundamentally do not lend themselves to building non-deterministic autonomous reasoning systems. Especially for AI, high hopes were raised more than 50 years ago, but AI has largely failed to deliver on its promises and still does. As such, the paper discusses different ingredients and approaches towards completely non-deterministic autonomous systems that are based on and exhibit critical capabilities, such as, but not limited to, self-organization, self-configuration, and self-adaptation. As such, any two initially identical autonomous systems will exhibit diverging and ultimately completely unpredictable developmental trajectories over time, once exposed to the same or similar environment, and even more so, once exposed to different environments.

Keywords: Non-deterministic autonomy, true randomness, objective global feature extraction and analysis, stochastic optimization framework, self-organization, self-configuration, self-adaptation, working hypothesis generation

1. INTRODUCTION AND MOTIVATION

Current computational approaches, such as Artificial Intelligence (AI), artificial neural networks (ANNs), expert systems, fuzzy logic, fuzzy-cognitive maps, other rule-based approaches, etc., fundamentally do not lend themselves to building non-deterministic autonomous reasoning systems. Especially for AI, high hopes were raised more than 50 years ago, but AI has largely failed to deliver on its promises and still does. In the following we will discuss ingredients and approaches towards completely non-deterministic autonomous systems that are based on and exhibit critical capabilities, such as, but not limited to, self-organization, self-configuration, and self-adaptation.

Although by far not the only area of application, most of the following will be put in the context of current and future autonomous robotic space exploration missions (Fig. 1) that will be necessary as we venture out into deep space.

Figure 1. Required degrees of operational autonomy for past, current, and future robotic space exploration missions [Mission Images credit: NASA, CMU, UT Austin, Colorado School of Mines, SRI, Breakthrough].

* e-mail: wfink@email.arizona.edu; phone: +1-520-621-8734; website: http://autonomy.arizona.edu

Micro- and Nanotechnology Sensors, Systems, and Applications XI, edited by Thomas George, M. Saif Islam, Proc. of SPIE Vol. 10982, 1098225 · © 2019 SPIE
CCC code: 0277-786X/19/$18 · doi: 10.1117/12.2518335

2. NSF'S DEFINITION OF "SMART" SERVICE SYSTEMS – A USEFUL BASELINE DEFINITION OF NON-DETERMINISTIC AUTONOMOUS SYSTEMS

The National Science Foundation (NSF) defines a "smart" service system as follows [NSF PFI:BIC 2015[1]]: *"A "smart" service system is a system capable of learning, dynamic adaptation, and decision making based upon data received, transmitted, and/or processed to improve its response to a future situation. […] These capabilities are the result of the incorporation of technologies for sensing, actuation, coordination, communication, control, etc. The system may exhibit a sequence of features such as detection, classification, and localization that lead to an outcome occurring within a reasonable time."* Especially the first sentence of this definition provides a useful baseline for non-deterministic autonomous systems as it names the key characteristics of *learning*, *dynamic adaptation*, and *decision making* – all based on/influenced by received, transmitted, and/or processed data.

Along similar considerations we postulate that a *non-deterministic autonomous system*, for the purposes of this publication, exhibits the following characteristics (note: non-exhaustive list and in no particular order):

- Non-reproducible, i.e., unique development over time;
- Can sense/assess environment through onboard sensors;
- Assesses environment objectively in sensor feature-space only rather than based on user-bias;
- Can be influenced by environment;
- Can make decisions based on sensor input and potential overall mission goal(s) and/or behavior(s);
- Can effectuate change to its environment and to itself (e.g., mobility);
- Only provided with atomic law(s), behavior(s), and tool(s) a priori;
- May be provided with mission goal(s) a priori or not;
- Can self-modify its behavior(s) and decision making and thereby "evolve";
- May exhibit forms of "self-awareness" but not expected to develop it.

3. TECHNICAL CORE INGREDIENTS FOR NON-DETERMINISTIC AUTONOMY

Rather than designing non-deterministic autonomous systems, which would constitute a borderline oxymoron, we propose to provide/instill basic atomic governing law(s), behavior(s), and tool(s) or process(es) as *technical core ingredients for non-deterministic autonomy to develop subsequently and in consequence as a result*, such as, but not limited to:

- Basic laws of physics, e.g., principle of least action or Hamilton's principle (Fig. 2);[2]
- Atomic behavior motifs, e.g., for Lidar-supported mobility;[3]
- Stochastic optimization as a tool, e.g., Stochastic Optimization Framework (SOF; Fig. 3);[4]
- Mathematical framework for manipulating/analyzing heterogeneous sensor data sets;[5,6]

as well as capabilities to:

- Sense/assess the environment (i.e., through sensors and/or instruments);
- Self-modify (especially the decision making aspect);
- Actuate/influence/move about the environment (e.g., through actuators).

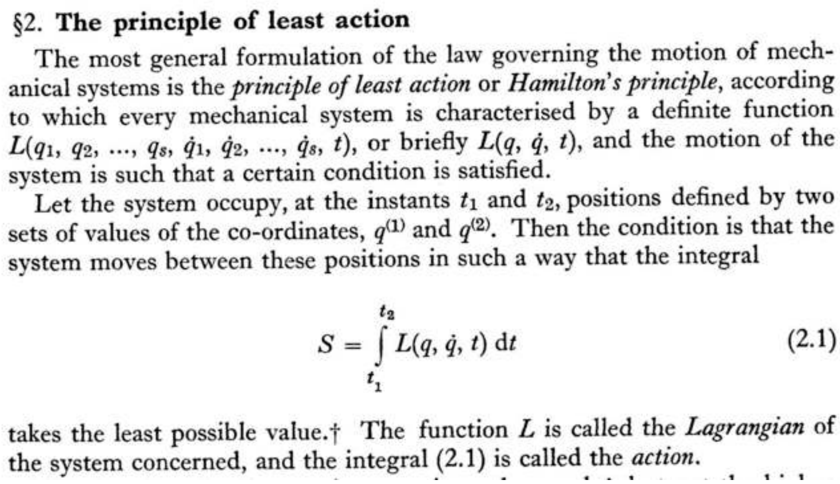

§2. **The principle of least action**

The most general formulation of the law governing the motion of mechanical systems is the *principle of least action* or *Hamilton's principle*, according to which every mechanical system is characterised by a definite function $L(q_1, q_2, \ldots, q_s, \dot{q}_1, \dot{q}_2, \ldots, \dot{q}_s, t)$, or briefly $L(q, \dot{q}, t)$, and the motion of the system is such that a certain condition is satisfied.

Let the system occupy, at the instants t_1 and t_2, positions defined by two sets of values of the co-ordinates, $q^{(1)}$ and $q^{(2)}$. Then the condition is that the system moves between these positions in such a way that the integral

$$S = \int_{t_1}^{t_2} L(q, \dot{q}, t)\, dt \tag{2.1}$$

takes the least possible value.† The function L is called the *Lagrangian* of the system concerned, and the integral (2.1) is called the *action*.

Figure 2. Basic law of physics example: principle of least action or Hamilton's principle as the most general formulation of the law governing the motion of mechanical systems, leading to Lagrange's and Hamilton's equations of motion [Excerpt from Landau-Lifshitz "Mechanics", 1976[2]].

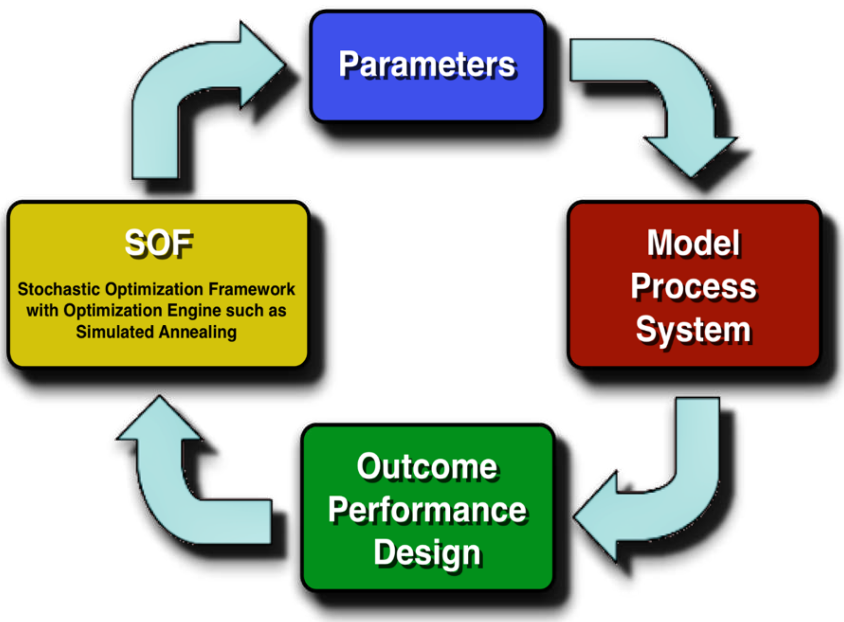

Figure 3. Functional schematic of a Stochastic Optimization Framework (SOF; after Fink, 2008[4]). SOFs allow for efficient sampling of the entire model/process/system-intrinsic parameter space of the model/process/system to be optimized by repeatedly running the respective model/process/system forward, comparing the outcomes against the desired outcome, which results in a fitness measure to be optimized, and by modifying the parameters via stochastic optimization algorithms, such as Simulated Annealing (SA),[7,8] Genetic Algorithms (GA),[9,10] Evolutionary Algorithms (EA), and Genetic Programming (GP),[11] etc.

These ingredients enable the self-development of higher order, more complex behavior(s), such as, but not limited to:

- "Self-awareness":
  - For example, through a combination of basic laws of physics and onboard sensor feedback, such that motion and reconfiguration experiments conducted onboard the autonomous robotic entity can yield the formulation of an "emergent self-model" about its makeup,[12] e.g., determining the number of

robotic legs and their configuration/arrangement with respect to the body of the robot without a priori knowledge;

- Advanced Mobility:
  - Robotic limb actuation without prior knowledge, e.g., learning how to walk from scratch;[12]
- Self-preservation:
  - Collision avoidance;[3]
- Unlimited/enhanced mobility, e.g., NASA's Space Technology Grand Challenge: *"All Access Mobility"*:[13]
  - Multi-objective path/traverse planning and optimization;[14-16]
  - Maximized exploration, e.g., deepest path exploration;[3]
  - Autonomous tele-commanding of robotic agents towards target areas or to avoid obstacles;[16-25]
  - Autonomous robotic agent redeployment/reconfiguration;[18-25]
- Answer (science) questions, e.g., through detection/identification of (feature-based) anomalies and/or regions of interest, such as, but not limited to, heat sources, locales of methane outgassing, subsurface water ice deposits, etc.[26-28,5,6]
  - Autonomous characterization of and anomaly detection in an operational area;[26,27,5,6]
  - Autonomous target prioritization;[28]
  - Autonomous robotic limb/actuator deployment towards target areas, e.g., for sampling or probing.[29]

3.1 Examples of Non-designed, Non-architected Locomotion Behavior

In the following, several examples of non-designed, non-architected locomotion behavior are briefly listed (with respective references for details), which were developed at the *Visual and Autonomous Exploration Systems Research Laboratory (autonomy.arizona.edu)*:

- Robotic limb deployment via SOF;[29]
- Adaptive hierarchical Lidar-based autonomous robotic navigation through atomic behavior motifs;[3]
- Robotic multi-objective path/traverse planning and optimization via SOF;[14-16]
- Round-robin overhead autonomous tele-commanding[18-25,16,17] for tier-scalable reconnaissance missions;[18-25]
- Design and optimization of low-thrust orbit transfers and associated mission trade studies for ion-engine propelled spacecraft via SOF.[30,31]

3.2 Approaches for True Randomness vs. Pseudo Random Number Generation for Stochastic Optimization

Stochastic optimization in general and Simulated Annealing (SA),[7,8] Genetic Algorithms (GA),[9,10] Evolutionary Algorithms (EA), and Genetic Programming (GP)[11] in particular, make heavy/heaviest use of random numbers. All computer-based *pseudo random number generators (PRNGs)* suffer from several challenges, such as, but not limited to:

(a) *Finite periodicity* of random number generators, i.e., pseudo random number sequences of finite length (albeit sometimes relatively long, e.g., *ran2()* random number generator:[32] $\sim 2 \times 10^{18}$);

(b) Only *pseudorandom* number sequences, i.e., algorithm-based and therefore completely deterministic;

(c) *Fully reproducible* random number sequences given same random number seed.

These limitations steer away from a possible uniqueness of the temporal trajectory development of an autonomous system that uses stochastic optimization. A partial solution is the use of timestamps as random number seeds. However, limitations (a) and (b) still remain in effect. Hence, overall, PRNGs generate random numbers that are entirely predictable/deterministic, and, given the same random number seed, even repeatable exactly. As such, no unexpected behavior can result from these PRNGs. Thus, for the development of non-deterministic autonomous systems, true randomness as a key requirement is postulated.

As a potential solution we propose to replace pseudo random number generators with processes that are "truly random," i.e., at least as far as we currently know. One such example is *radioactive decay*, which is currently still considered "truly random" since we do not know the exact physics laws that determine when a particular nucleus is going to decay. On the other hand, a single radioactive sample may not deliver random events fast and plentiful enough to sustain the huge demands of a Stochastic Optimization Framework. Half-lives can range from 10^{-24} seconds to 10^{30} seconds [Source: Wikipedia]. For example, Polonium-215 has a half-life of 0.0018 seconds and Uranium-238 has a half-life of 4.5 billion years [Source: Wikipedia]. As a possible solution one can use radioisotopes with very short half-lives, and/or multiple, i.e., N samples of the same radioisotope simultaneously. The proposed procedure would then be to measure inter-decay-event time via, e.g., a Geiger-Mueller-counter or similar devices, and to convert this time to a random number, or, to use the timestamp for each decay event as a random number. The detection rate, i.e., the temporal resolution of the detector will determine, which radioisotopes may be useable.

Alternatively, *quantum random number generators (QRNGs)* have become increasingly popular, exploiting *quantum indeterminism*, such as, but not limited to:

- ID Quantique's Quantis Random Number Generator:[33] exploiting the fundamentally random nature of quantum optics;
- Australian National University (ANU) Quantum Random Numbers Server:[34] measuring the quantum fluctuations of the vacuum;
- Joint R&D effort of PicoQuant GmbH and the Nano-Optics groups at the Department of Physics of Humboldt University:[35] QRNG based on quantum randomness of photon arrival times;
- National Institute of Standards and Technology (NIST) Randomness Beacon:[36,37] An intense laser hits a special crystal that converts laser light into pairs of photons that are entangled, a quantum phenomenon that links their properties. These photons are then measured to produce a string of truly random numbers.

3.3 Approaches for Objective Global Feature Extraction, Analysis, Anomaly Detection, and Target Prioritization

Systems such as the Caltech-patented *Automated Global Feature Analyzer™ (AGFA™)*[26-28,5,6] can be employed as a driver for *operational autonomy*, comprising, but not limited to, the following fundamental characteristics:[25]

(1) Automatic characterization of operational areas from different vantages (e.g., space, air, ground, subsurface);

(2) Automatic sensor deployment and data gathering;

(3) Automatic feature extraction, anomaly detection, and region-of-interest or target identification;

(4) Automatic region-of-interest or target prioritization;

(5) Subsequent automatic redeployment and navigation of robotic agents to regions or targets of interest.

AGFA is an extensible analysis and classification framework to perform automated target/region-of-interest identification and unbiased anomaly detection. AGFA performs superpixel-based image segmentation and analysis on images of operational areas under investigation with an array of heterogeneous sensors and/or instruments (including cameras), extracts features, and generates feature vectors for all identified targets. AGFA uses these feature vectors subsequently to (a) classify individual targets according to their extracted features; (b) summarize the operational area numerically for each extracted feature category; and (c) automatically flag significant anomalies within the feature space using a patented *multi-stage normalization cascade*[5,6] followed by, e.g., agglomerative clustering and Principal Component Analysis (PCA). All this is done in an objective, i.e., exclusively feature-space-driven and thus unbiased manner, as opposed to a biased, human hypothesis-driven manner. As such, it lends itself as a technological core ingredient for an autonomous (space exploration) system, to accommodate the "maxim" of planetary exploration: "to expect the unexpected."

Moreover, AGFA is equipped with a mathematical prioritization framework[28] that allows the determination of the most promising target(s) for in-situ follow-up investigation in an operational area that has previously been visited/examined/mapped from a distance. The underlying process, termed "hypothetical probing"[28], seeks to minimize the *clustering quality* expressed as an *objective function E(t)*:

$$E(t) = \sum_{k=1}^{K} \sum_{i=1}^{N} M_{ki}(t) \left(\left\| \mathbf{c}_i(t) - \mathbf{cc}_k(t) \right\|^2 - \mu \right),$$

where K is the number of clusters, N the number of feature vectors, $c_i(t)$ the feature vector i, $cc_k(t)$ the centroid vector of cluster k, $M_{ki}(t)$ the membership strength of feature vector i with respect to cluster k (accommodating both hard and soft clustering) – and all of that at time t. μ is a constant reward/penalty term. The minimization of $E(t)$ is accomplished by hypothetically changing the respective membership strengths at the next time step, i.e., $M_{ki}(t+1)$, and checking whether $\Delta E = E(t+1) - E(t)$ is negative or positive (or zero).[28]

3.4 Approaches for Self-Organization, Self-Configuration, and Self-Adaptation of Artificial Neural Networks in Support of Synthetic Reasoning

Artificial Neural Networks (ANNs) are powerful methods for (1) the classification and analysis of multi-dimensional data, (2) the learning and generalization of rules underlying data, as well as (3) for the control of (highly non-linear) dynamic systems. ANNs are at the core of Machine Learning, Deep Learning, and Artificial Intelligence techniques, etc. To overcome the inherent challenge of choosing/designing a suitable architecture (e.g., number of neural layers and number of neurons for each layer), connectivity between neural layers, and associated neural coupling strengths for ANNs to solve complex classification or control tasks, a Stochastic Optimization Framework[4] can be employed instead to auto-generate Artificial Neural Networks appropriate/sufficient for the task at hand.[38] Hereby, the objective or fitness function to be minimized by the SOF is a measure of misclassification or misbehavior of the current version of the ANN against a training set or expected/desired behavior of a dynamic system, e.g., a legged robot relocating from point A to B by any means necessary. This SOF-based auto-generation mechanism for ANNs paves the way towards dynamic "self-(re-)design", especially when the ANN architecture, connectivity, and neural coupling strengths are subject to constant change over time while interacting with the environment, i.e., while influencing/being influenced by the environment.

3.5 Approaches for Working Hypothesis Generation through General-Purpose Machine Learning Systems as Enablers of Higher Autonomy

General-Purpose Machine Learning Systems (GPMLSs) that are designed to operate in real world environments need to exhibit the ability to adapt, i.e., to modify themselves based on feedback or successful assessment of situations or actions they take or cause, e.g., by interacting with the environment. A GPMLS, e.g., a *Learning Classifier System (LCS; Figs. 4, 5),*[39,40] can start from a set of rules or classifiers that is being used to assess situations and/or the environment (Figs. 4, 5). Based on the feedback or successful assessment of situations/environment, those rules that successfully contributed to the assessment are being rewarded (i.e., gain in strength; Fig. 5, top; for details see Farmer et al., 1986[41]), and those that did not contribute successfully are being penalized (i.e., lose strength; Fig. 5, middle; for details see Farmer et al., 1986[41]). A similar process can be found in living brains – so-called *Bonhoeffer-effect* – where synaptic strengths increase or decrease over time depending on inter-neural "traffic." In addition, an overall taxation scheme can be implemented as a Darwinian pressure term after a certain number of iterations of the GPMLS (Fig. 5, bottom; for details see Farmer et al., 1986[41]). This can lead to the elimination of some potentially unimportant or repeatedly unproductive rules. Conversely, new rules can be introduced into the system on a random, occasional basis that may improve the overall performance of the GPMLS over time, by introducing a new "perspective." Moreover, rules can be modified in ways very similar to those used in Genetic Algorithms,[9,42,10] such as *point mutation*, *crossover*, *inversion*, etc. – so-called *genetic operators*. As such, the rules underlying the GPMLS will evolve over time, and also change in number if not kept constant or increase if not capped, to assess situations and resulting decisions more successfully. Equipped with these capabilities, general-purpose machine learning systems are adaptive, and, as opposed to Artificial Intelligence-based systems that are mostly or entirely based on a *fixed* set of *unchangeable rules*, e.g., Mamdani-type IF-THEN-rules,[43,44] GPMLSs are capable of continuously evolving and modifying themselves. Therefore, GPMLSs are promising candidates for the generation of *working hypotheses*,[45,46] a concept used in predominantly *abductive* disciplines, such as Geology, Biology, and Medicine, as opposed to *deductive-inductive* disciplines, such as Physics and Mathematics.

Hereby, the equivalent of "working hypotheses" would be represented by the currently active set of rules/classifiers in the GPMLS at any point in time. By the nature of the dynamic update process underlying a GPMLS, some "working hypotheses" will be momentary (i.e., short-term), temporary (i.e., mid-term), or long-term (potentially permanent).

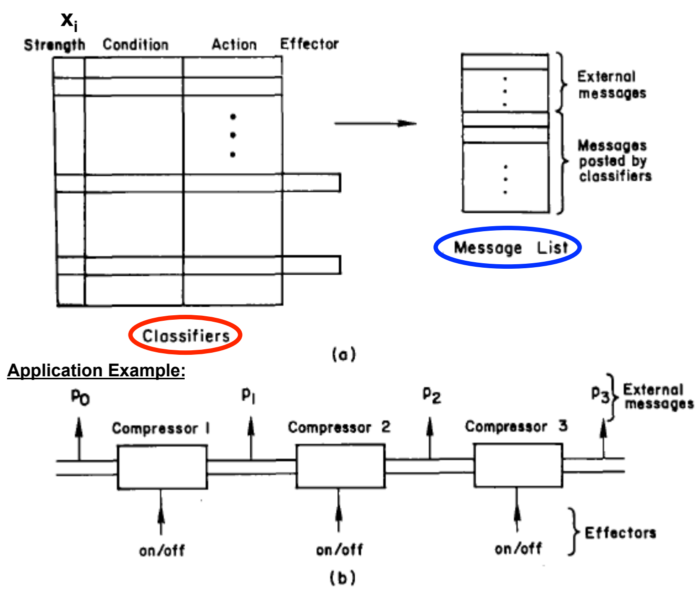

Message list is interface between outside world and system; can be potentially unlimited

Condition, Action, and Effector are binary strings that match the length of messages.

Condition is used to determine whether or how much posted messages by other classifier rules pertain to a classifier rule, i.e., the mutual overlap

Actions and Effectors are message postings, where Action is internal and Effector is external

Figure 4. General schematic overview of a Learning Classifier System (a),[47,41] applied to a dynamic control problem (b), i.e., pressure control within a gas pipeline using pressure sensors and compressors[47] [Schematic on the left from Goldberg, 1983[47] and Farmer et al., 1986[41]].

Strength increase (pay back) due to own messages being read successfully by other rules

$$\Delta x_i = c \sum_{j}^{N} m_{ji} x_j f_i H(m_{ji} x_j - T).$$

Strength decrease (bid) due to successful posting of own messages in response to messages of other rules

$$\Delta x_i = -c \sum_{j=1}^{N+n} m_{ij} x_i f_j H(m_{ij} x_i - T).$$

Total change in strength taking decrease and increase, as well as overall global payoff and taxation terms into account

$$\begin{aligned} \Delta x_i &= x_i^{t+1} - x_i^t \\ &= c\left[\sum_{j=1}^{N} m_{ji} x_j f_i H(m_{ji} x_j - T) - \sum_{j=1}^{N} m_{ij} x_i f_j H(m_{ij} x_i - T)\right] \\ &\quad + f_i P^t - k_2 x_i. \end{aligned}$$

Figure 5. Inner workings of Learning Classifier Systems: Governing equations for system trajectory,[41] i.e., change of strengths of rules/classifiers over time/per iteration t [Equations on the right from Farmer et al., 1986[41]]: x_i = strength of rule i, m_{ij} = match specificity between condition part of classifier i and message posted by classifier j, f_i = index function whether classifier i posted action on message list, $H()$ = Heaviside function, T = bid threshold, P = performance function, k_2 = taxation, c = fractional constant.

4. CONCLUSIONS – DISCUSSIONS – OUTLOOK

A set of technical core ingredients were introduced that would enable a dynamic system to develop non-deterministic autonomy over time – as opposed to a priori design – through learning, dynamic adaptation, and decision making – all based on/influenced by received, transmitted, and/or processed data. In other words, rather than designing the "finished cake" a priori, only the ingredients and the "baking process" are provided to a dynamic system – both embedded in or guided by a Stochastic Optimization Framework, and constantly influenced by and influencing the environment – to leave it up to the system to dynamically "bake the cake" appropriate for the task(s) and operational environment(s) at hand, and, moreover, to adjust it on a need be basis to accommodate changes in the task(s) and operational environment(s) over time.

Several approaches were briefly discussed to further support this developmental process, such as, but not limited to (in random order): (1) "true" randomness; (2) objective global feature extraction, analysis, and anomaly detection; (3) self-organization, self-configuration, and self-adaptation of artificial neural networks in support of synthetic reasoning, and (4) working hypothesis generation through Learning Classifier Systems. In contrast to Artificial Intelligence (AI) schemes, these core ingredients and approaches outlined above neither depend on, nor operate within the confines of expert-defined rule sets. As such, any two initially identical autonomous systems, which embrace at least some of these technical core ingredients, will exhibit diverging and ultimately completely unpredictable developmental trajectories over time, once exposed to the same or similar environment, and even more so, once exposed to different environments. Especially through the use of uncapped, i.e., size-unlimited LCSs can the environment influence the autonomous system and vice versa. Using genetic operators, the "bucket brigade" and taxation framework inherent to LCSs,[39-41] coupled with environmental factors and "true" randomness will provide for unique developmental and, most importantly, non-deterministic trajectories for autonomous systems, which was the declared goal of this paper.

Several of the above discussed ingredients and approaches have been implemented by the Visual and Autonomous Exploration Systems Research Laboratory, first at Caltech and now at UofA. In the context and for the purposes of autonomous robotic space exploration, these elements are being integrated in a robotic testbed (at present including aerial and surface based robotic agents) for *Tier-Scalable Reconnaissance*[18-25] at the University of Arizona, and, paired with the integration of automated objective global feature analysis and anomaly detection,[26,27,5,6] automated target prioritization,[28] stochastic and/or globally optimal multi-objective traverse/path planning,[14-16] round-robin overhead autonomous tele-commanding,[18-25,16,17] Lidar behavior motifs,[3] robotic limb deployment,[29] etc., can now be fielded, studied, and validated with respect to operational autonomy for planetary exploration.

Future work at the Visual and Autonomous Exploration Systems Research Laboratory is targeted towards the development of an entirely SOF-governed robotic agent as a steppingstone towards a non-deterministic autonomous system. Similar to, or as an extension of robotic limb deployment and multi-objective path planning, more complex behaviors would be formulated or cast as optimization tasks, and solved, i.e., executed through a Stochastic Optimization Framework.[4] For example, if a robotic swarm pursued the goal of reaching a certain destination, the path to which is blocked by an obstacle (e.g., a wall), some of the robotic agents could be commanded to slam into the obstacle to destroy it and/or to clear a path for the remaining robotic agents of the swarm to reach their destination. On the outside looking in, this seemingly deliberate and selfless act could be construed as a remarkable gambit or self-sacrificial "behavior" for the "well-being" or success of the greater good of the collective. In reality, it is nothing more than the maximization of an objective function through multi-objective optimization, i.e., the reaching of the destination with as many robotic agents of the swarm as possible, taking the shortest distance possible, and in the shortest time possible. In a similar fashion, other complex, even more human-like behaviors can be emulated, i.e., cast as multi-objective optimization tasks, and solved, i.e., executed through a Stochastic Optimization Framework.[4]

5. ACKNOWLEDGEMENTS

This publication was created at the University of Arizona with partial support from the Edward & Maria Keonjian Endowment at the University of Arizona.

6. AUTHOR DISCLOSURE STATEMENT

Author WF may have a financial interest in some of the underlying technologies presented here, as he is named as an inventor on several issued or pending patents (see References) on behalf of the California Institute of Technology (Caltech) and the University of Arizona (UofA).

REFERENCES

[1] https://www.nsf.gov/pubs/2014/nsf14610/nsf14610.htm

[2] Landau LD, Lifshitz EM (1976) Mechanics; Volume 1 of Course of Theoretical Physics, Transl. by J.B. Sykes and J.S. Bell, 3rd Edition, Butterworth-Heinemann, Reed Educational and Professional Publishing Ltd. 1981.

[3] Brooks A J-W, Fink W, Tarbell MA (2018) An adaptive hierarchical approach to Lidar-based autonomous robotic navigation; Proc. SPIE, Micro- and Nanotechnology Sensors, Systems, and Applications X, SPIE DCS Conference 2018, Orlando, FL

[4] Fink W (2008) Stochastic Optimization Framework (SOF) for Computer-Optimized Design, Engineering, and Performance of Multi-Dimensional Systems and Processes; Proc. SPIE, Vol. 6960, 69600N (2008); DOI:10.1117/12.784440 (invited paper).

[5] US 9,122,956 and international progenies: "Automated feature analysis, comparison, and anomaly detection" (Fink; Caltech)

[6] US 9,424,489 and international progenies: "Automated feature analysis, comparison, and anomaly detection" (Fink; Caltech)

[7] Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E, Equation of State Calculation by Fast Computing Machines, *J. of Chem. Phys.*, 21, 1087 – 1091, 1953.

[8] Kirkpatrick S, Gelat CD, Vecchi MP, Optimization by Simulated Annealing, *Science*, 220, 671 – 680, 1983.

[9] Holland JH, *Adaptation in Natural and Artificial Systems*, The University of Michigan Press, Ann Arbor, Michigan, 1975.

[10] Goldberg DE, Genetic Algorithms in Search, Optimization and Machine Learning. Addison-Wesley, 1989.

[11] Koza JR (1992) *Genetic Programming: On the Programming of Computers by Means of Natural Selection,* Cambridge, MA: The MIT Press, 1992.

[12] "A Robot Teaches Itself How to Walk" by Cornell University professor Hod Lipson (2012): https://www.youtube.com/watch?v=iNL5-0_T1D0

[13] http://www.nasa.gov/pdf/503466main_space_tech_grand_challenges_12_02_10.pdf

[14] Fink W, Baker VR, Flammia M, Tarbell MA (2015) Rover Traverse-Optimizing Planner For Multi-Objective Deployment Scenarios; IEEE Aerospace Conference Proceedings, paper #2722, Big Sky, Montana

[15] Fink W, Baker VR, Flammia M, Tarbell MA (2015) Avoiding Planetary Rover Damage by Multi-Objective Rover Traverse Optimization [abstract 1353]. In 46th Lunar and Planetary Science Conference Abstracts [CD-ROM], Lunar and Planetary Institute, The Woodlands, TX

[16] US20170254657: "TRAVERSE AND TRAJECTORY OPTIMIZATION AND MULTI-PURPOSE TRACKING" (Fink; UofA)

[17] Fink W, Mahmood Q (2019) Template-based Overhead Detection, Identification, and Tracking of Multiple Surface-based Exploration Vehicles; IEEE Aerospace Conference Proceedings, paper #2853, Big Sky, Montana

[18] Fink W, Dohm JM, Tarbell MA, Hare TM, Baker VR (2005) Next-Generation Robotic Planetary Reconnaissance Missions: A Paradigm Shift; Planetary and Space Science, 53, 1419-1426

[19] US 6,990,406: "MULTI-AGENT AUTONOMOUS SYSTEM" (Fink, Dohm, Tarbell; Caltech)

[20] US 7,734,063: "MULTI-AGENT AUTONOMOUS SYSTEM" (Fink, Dohm, Tarbell; Caltech)

[21] US 7,742,845: "MULTI-AGENT AUTONOMOUS SYSTEM AND METHOD" (Fink, Dohm, Tarbell; Caltech)

[22] Fink W, Tarbell MA (2007) Tier-Scalable Reconnaissance© - A Paradigm Shift in Autonomous Robotic Planetary Exploration; 7th IAA International Conf. on Low-Cost Planetary Missions 2007, Pasadena, CA, abstract and poster

[23] Fink W, Tarbell MA, Jobling FM (2008) Tier-Scalable Reconnaissance - A Paradigm Shift in Autonomous Remote Planetary Exploration of Mars and Beyond; Chapter 1 in "Planet Mars Research Focus", Ed. L. A. Costas. Nova Science Publishers, Hauppauge, NY. ISBN: 1-60021-826-1

[24] Kean S (2010) Making Smarter, Savvier Robots; Science, 329, 508-509

[25] Fink W, Brooks A J-W, Tarbell MA, Dohm JM (2017) Tier-Scalable Reconnaissance: The Future in Autonomous C^4ISR Systems has arrived – Progress towards an Outdoor Testbed; Proc. SPIE 10194, Micro- and Nanotechnology Sensors, Systems, and Applications IX, 1019422 (18 May 2017); doi: 10.1117/12.2257333
[26] Fink W, Brooks A J-W, Tarbell MA (2018) Anomaly detection and target prioritization in planetary imagery via the Automated Global Feature Analyzer (AGFA): Progress towards a driver for autonomous C^4ISR missions; Proc. SPIE, Micro- and Nanotechnology Sensors, Systems, and Applications X, SPIE DCS Conference 2018, Orlando, FL
[27] Fink W, Datta A, Dohm JM, Tarbell MA, Jobling FM, Furfaro R, Kargel JS, Schulze-Makuch D, Baker VR (2008) Automated Global Feature Analyzer (AGFA) – A Driver for Tier-Scalable Reconnaissance; IEEE Aerospace Conference Proceedings, paper #1273; DOI: 10.1109/AERO.2008.4526422
[28] Fink W (2006) Generic Prioritization Framework for Target Selection and Instrument Usage for Reconnaissance Mission Autonomy, Proc. IEEE, WCCI 2006, Vancouver, Canada, 11116-11119
[29] Fink W, Huntsberger TL, Aghazarian H (2010) Dynamic Optimization of N-Joint Robotic Limb Deployments; Journal of Field Robotics, Volume 27, Issue 3, p. 268-280, DOI: 10.1002/rob.20323
[30] Lee S, von Allmen P, Fink W, Petropoulos AE, Terrile RJ (2004) Design and Optimization of Low-thrust Orbit Transfers; IEEE Aerospace Conference Proceedings; DOI: 10.1109/AERO.2005.1559377
[31] Lee S, Russell RP, Fink W, Terrile RJ, Petropoulos AE, von Allmen P (2006) Low-Thrust Mission Trade Studies with Parallel, Evolutionary Computing; IEEE Aerospace Conference Proceedings, paper #1042, DOI: 10.1109/AERO.2006.1656038
[32] Press WH, Flannery BP, Teukolsky SA, Vetterling WT. Numerical recipes in C: the art of scientific computing. Cambridge, NY: Cambridge University Press; 1991. p. 286–9.
[33] ID Quantique's Quantis Random Number Generator: https://www.idquantique.com/random-number-generation/products/quantis-random-number-generator/
[34] Australian National University (ANU) Quantum Random Numbers Server: https://qrng.anu.edu.au/
[35] QRNG through Joint R&D effort of PicoQuant GmbH and the Nano-Optics groups at the Department of Physics of Humboldt University: https://qrng.physik.hu-berlin.de/
[36] NIST Randomness Beacon: https://www.nist.gov/programs-projects/nist-randomness-beacon
[37] NIST's New Quantum Method Generates Really Random Numbers: https://www.nist.gov/news-events/news/2018/04/nists-new-quantum-method-generates-really-random-numbers
[38] Fink W (2009) Autonomous Self-Configuration of Artificial Neural Networks for Data Classification or System Control; Proc. SPIE, Vol. 7331, 733105 (2009); DOI:10.1117/12.821836
[39] Holland JH, Reitman JS (1978) Cognitive systems based on adaptive algorithms; reprinted in: Evolutionary computation. The fossil record. In: David BF (ed) IEEE Press, New York 1998. ISBN 0-7803-3481-7.
[40] Holland JH (1986) Escaping brittleness: The possibilities of general purpose learning algorithms applied to parallel rule-base systems; in: Machine Learning 2, R.S. Michalski, J.G. Carbonell and T.M. Mitchell, eds. (Morgan Kauffman, Los Altos, in press), chap. 20.
[41] Farmer JD, Packard NH, Perelson AS (1986) The Immune System, Adaptation, and Machine Learning; Physica 22D, 187-204.
[42] Holland JH (1981) "Genetic Algorithms and Adaptation", Technical Report #34, Univ. Michigan, Cognitive Science Department.
[43] Mamdani EH, Assilian S (1975) An experiment in linguistic synthesis with a fuzzy logic controller; International Journal of Man-Machine Studies, 7(1), 1-13. http://dx.doi.org/10.1016/S0020-7373(75)80002-2
[44] Mamdani EH (1977) Applications of fuzzy logic to approximate reasoning using linguistic synthesis; IEEE Trans. Comput. 26 (No. 12) (1977) 1182–1191.
[45] Chamberlin TC (1890) The method of multiple working hypotheses; Science 15:92–96. (reprinted in Science 148: 754–759 [1965]).
[46] Baker VR (2017) Debates—Hypothesis testing in hydrology: pursuing certainty versus pursuing uberty; Water Resources Research, v. 53, p. 1770-1778.
[47] Goldberg S (1983) Computer Aided Gas Pipeline Operation Using Genetic Algorithms and Rule Learning; Dissertation, University of Michigan.