*Article*

# Can A Robot Have Free Will?

**Keith Farnsworth**

School of Biological Sciences, Queen's University Belfast; Northern Ireland, UK; k.farnsworth@qub.ac.uk

**Abstract:** Using insights from cybernetics and an information-based understanding of biological systems, a precise, scientifically inspired, definition of free-will is offered and the essential requirements for an agent to possess it in principle are set out. These are: a) there must be a self to self-determine; b) there must be a non-zero probability of more than one option being enacted; c) there must be an internal means of choosing among options (which is not merely random, since randomness is not a choice). For (a) to be fulfilled, the agent of self-determination must be organisationally closed (a 'Kantian whole'). For (c) to be fulfilled: d) options must be generated from an internal model of the self which can calculate future states contingent on possible responses; e) choosing among these options requires their evaluation using an internally generated goal defined on an objective function representing the overall 'master function' of the agent and f) for 'deep free-will', at least two nested levels of choice and goal (d-e) must be enacted by the agent. The agent must also be able to enact its choice in physical reality. The only systems known to meet all these criteria are living organisms, not just humans, but a wide range of organisms. The main impediment to free-will in present-day artificial robots, is their lack of being a Kantian whole. Consciousness does not seem to be a requirement and the minimum complexity for a free-will system may be quite low and include relatively simple life-forms that are at least able to learn.

**Keywords:** self-organization; downward causation; autocatalytic set; goal-oriented behaviour; autopoiesis; biological computing

---

## 1. Introduction

Why do things do what they do? We have a hierarchy of explanations that roughly reflects a gradient in complexity, matched by the epistemic hierarchy which starts with the physics of Hamiltonian mechanics (and Shrödinger's equation), extends through statistical mechanics and complex systems theory, but then declines in power as we try to account for the behaviour of living systems and finally of the human condition, for which we have no satisfactory scientific explanation. One of the most persistent open questions, at the far end of the complexity gradient, is whether we as humans have free will. Here, I attempt to address this question with respect to a broader category of active agent (*sensu* Sharov [1]): anything that can make decisions and act in the physical world. The premise of this paper is that an understanding of the interaction and dynamics among patterns – of the distribution of matter and energy in space and time – may bring such high-level phenomena into resolution. That means a focus on information and its interactions using cybernetics and computation theory, but it also requires a broad concept of information addressing the relationship among patterns (as data) in general, rather than just the statistics of data transmission. If the emergence of material reality (as we experience it) is the assembly of stable configurations (of matter), undergoing transformations and combining as stable composites (e.g. molecules forming materials, forming structures [2]), then a deep understanding of it requires a mathematically precise account of the physics of patterns that are simultaneously the product of material structure and the

cause of it [3]. For this investigation, we must focus on the kind of 'information' that is embodied by, and processed by natural systems. This concept of intrinsic structural 'information' (surely it ought to have a word of its own), is not necessarily ontological (as in the theories of Weizsäcker [4] and Stonier [5]), but refers to at least observable patterns having observable effects [6], is objectively quantifiable [7] and useful in understanding biological processes in terms of cybernetic systems [1], and functions [8,9]. To avoid ambiguity (and conflict) I will refer to this kind of 'information' as pattern and the 'information' which reduces 'uncertainty' in a receiver as Shannon information.

I will start by defining what I mean by 'free-will' and then give a very brief overview of the philosophical debate about free-will, identifying some of the problems. Then I will introduce ways of thinking about these problems based on cybernetic/computation theory and use these to identify the necessary and sufficient conditions for free-will according to my definition. The class of systems for which these conditions are fulfilled will then be identified and with this, the minimum complexity compatible with autonomous action will be implied. The term 'robot' is specifically used here to mean any cybernetic system coupled to a physical system that allows it to independently act in the physical world; this includes living systems (see e.g. the requirements specified for a molecular robot in Hagiya *et al.* [10]) and artificial systems. I am going to conclude that only certain living systems (including humans, together with their anthropic extensions) so far achieve the requirements for organisational autonomy and therefore free-will.

### *1.1. A definition of free-will*

Free-will is defined here as the condition in which all of the following are jointly true: a) there exists a definite entity to which free-will may (or may not) be attributed, b) there are viable alternative actions for the entity to select from, c) it is not constrained in the exercising of two or more of the alternatives, d) its 'will' is generated by non-random process internal to it, e) in similar circumstances, it may act otherwise according to a different internally generated 'will'. In this definition the term 'will' means an intentional plan which is jointly determined by a 'goal' and information about the state (including its history) of the entity (internal) and (usually, but not necessarily) its environment. The term 'goal' here means a definite objective that is set and maintained internally. The terms used here will be explained and justified in what follows.

### *1.2. The Philosophical background*

In philosophy, freedom 'to act as one wills' is often referred to as 'superficial freedom' and the freedom 'to determine what one wills' as 'deep freedom' (Kane [11] provides a good introduction to the subject, Westen [12] gives a deeper criticism of it). For example, one may be free to drink a bottle of whisky in one sitting and as an alcoholic, or in an irresponsible mood, or emotional turmoil one may will it, but knowing the consequences, one may master the desire and will otherwise: rejecting the opportunity of this poisonous pleasure. If the alcoholism had taken over, or one was under duress, the 'freedom to choose otherwise' might be denied and deep freedom would be lost with it, although the superficial freedom to drink would remain. The classical philosophical argument on free-will consists of a) whether it is compatible or not with determinism and b) whether at least some agents (usually people) are at least in part the ultimate cause of their actions.

Determinism is the idea that there is, at any instant, exactly one physically possible future [12,13], summarised in the slogan 'same past: same future' (see List [14] for a more rigorous analysis). Cybernetics captures determinism in the definition of a determinate machine (DM) as a series of closed, single valued transformations (for example describing a finite state automaton (FSA)). Superficial freedom is often seen as the absence of constraint, leading to the (relatively trivial) conclusion that it is compatible with determinism. But deep freedom needs more than an absence of constraints, it requires alternative paths into the future to provide the 'freedom to do otherwise' [13]. The cybernetic model of a system with this capacity is of course the non-determinate machine (NDM). However, the NDM is usually conceived as a probabilistic process in which a set of possible

states $S\{s_1...s_n\}$ of the system, given the present conditions $C$ (in general including the previous history), may occur at random, with probability set $P\{p_1...p_n\}$ where $\sum p_i = 1$ and each $p_i$ is the probability of each possible state $s_i$. Most philosophers agree that randomness is not compatible with self-determination, indeed it seems to be the opposite, so they reject random spontaneity as a means of achieving deep freedom (and so do I). They reason that if an agent's action were ultimately caused by e.g. a quantum fluctuation or thermal noise, then we could not reasonably hold the agent responsible for it. This indicates that philosophers supporting the existence of deep free-will are searching for an ultimate cause of actions within the agent of those actions.

Unfortunately for them, a paradox arises: since any agent is the product of its composition and previous experiences (the making of it) and these are beyond its control. If it did not make itself and select its own experiences, then its behaviour must be determined by things other than itself. An agent is not free in the deep sense unless it has control over all the events that led to it's choice of action. Recognising that all events in the universe belong to a chain of cause and effect that extends back before the existence of the agent, some philosophers conclude that either a) this deep freedom cannot exist and is considered an illusion (e.g. Van Inwagen [13], reiterated in [15]) , or b) the agent is indeterminate so that we get 'same past: different futures'. If they also rule out randomness, then (b) suggests that an agent which could act in more than one possible way from exactly the same state and history (i.e. it is indeterminate) must act without cause. They conclude that this is self-contradictory, hence deep freedom cannot exist. This line of thinking is closely related to Strawson's [16] 'Basic Argument' against free-will, which starts from the premise that for an agent to have free-will it must be the cause of itself and shows, via infinite regress, that this is not possible. Axiomatic to these positions is that the action of an agent can only be either a) random, b) exogenous or c) of itself, with only c) being compatible with free-will.

## 2. Systems with Identity: the *closure condition*

If I am talking of my own free-will, what exactly am I? Only by answering this question do I reach the position of being able to determine if I have free-will or not (rarely, if ever, has this question been posed in the philosophy of free-will). Free-will requires a definite boundary between the internal and external, not only physically (as supplied by e.g. a casing or skin), but more profoundly in organisation and control. For example a computer controlled robot must have all the necessary provisions for physical independence (as in the extraterrestrial exploration robots), but this still leaves them organisationally linked to humanity because their existence is entirely dependent on our gathering and processing the materials for their 'bodies' and assembling these, implicitly embodying them with functional information [1] and programming their control computer (including the goal for operation). For these reasons, they remain extensions of ourselves: tools just as sophisticated hammers. In general, for free-will, the control information of an agent must be independent of anything beyond a cybernetically meaningful boundary. Put the other way round, for the identification of free-will, we must first identify the boundary of the agent, which is defined by independence of control. Given this, the Mars Rover coupled with its human design team seems to meet the closure condition, but the Mars Rover alone does not.

This idea of a boundary surrounding a system, such that whatever is within the boundary has the property of organisational independence from what lies without, was encapsulated by the concept of the 'Kantian whole' by Kauffman [17] and can be formally described in cybernetic terms as *organisational closure*.

### *2.1. The Kantian whole*

A system composed of parts, each of whose existence depends on that of the whole system is here termed a 'Kantian whole', the archetypal example being a bacterial cell [18]. The origin of this terminology lies in Emanuel Kant's definition of an organised whole [19]. For the present purpose, the closure condition is best explained in terms of self-construction, since this implies the

embodiment of self with the pattern-information that will then produce the agents behaviour. In other words, we are to consider a cybernetic system that, by constructing itself materially, determines its transition rules, by and for itself (material self-construction may not be the only way to ensure this self-determination,but assuming the cybernetic relations it embodies are, we may proceed without loss of generality). An autocatalytic chemical reaction network with organisational closure (and this is also what Kauffman [17] considered a Kantian whole) is an anabolic system able to construct itself [20].

Hordijk and Steel [20] and Hordijk *et al.* [21] define their chemical reaction system by a tuple $Q = \{X, \mathcal{R}, C\}$, in which $X$ is a set of molecular types, $\mathcal{R}$ a set of reactions and $C$ a set of catalytic relations specifying which molecular types catalyse each member of $\mathcal{R}$. The system is also provided with a set of resource molecules $F \subseteq X$, freely available in the environment, to serve as raw materials for anabolism (noting that whilst we are defining an organisational closure, we may (and indeed must) permit the system to be materially and thermodynamically open). The autocatalytic set is that subset of reactions $\mathcal{R}' \subseteq \mathcal{R}$, strictly involving the subset $X' \subseteq X$, which is:

- *reflexively autocatalytic*: every reaction $r \in \mathcal{R}'$ is catalysed by at least one of molecular type $x' \in X'$ and
- composed of $F$ by $\mathcal{R}'$: all members of $X'$ are created by the actions of $\mathcal{R}'$ on $F \cup X'$ .

This definition of an autocatalytic set is an application of relational closure (in computation theory terms, the molecule types are symbol strings for which the Kleene closure is readily defined) and it has been implemented in experiments for exploring aspects of the origin of life (e.g. the GARD system simulating 'lipid world' [22]). Clearly with the two conditions above met, everything in the system is made by the system, but there is a more important consequence. The system is made from the parts (only) and can only exist if they do. Organisational closure of this kind has been identified as a general property of individual organisms [23], many biochemical sub-systems of life [17] and embryonic development [24]).

As it is defined above, living systems fulfil the closure condition, but can we conceive of a non-living system also reaching this milestone? Von Neumann's [25] self-replicating automata show that some purely informational (algorithm) systems have the capacity to reproduce within their non-material domain, but they cannot yet assemble the material parts necessary, nor can they build themselves from basic algorithmic components (they rely on a human programmer to make the first copy). What is needed for the physical implementation of a Kantian whole is the ability to 'boot-strap' from the assembly of simple physical components to reach the point of autonomous replication (i.e. the system must be autopoietic [23,26]. This is necessary to answer Strawson's [16] 'Basic Argument': that for deep free-will an entity must be responsible for shaping its own form and it provides a motivation for rejecting dualism (the idea that the 'mind' is not created from the material universe).

### *2.2. Emergence and downward causation*

Considering the forgoing, we might ask what it is that is making an autopoietic system (e.g. an organism) - is it the components, or the system itself, and in any case, what really is the 'system'. Cybernetics provides an answer to the second question, in that the system is the organisational pattern-information embodied in a particular configuration of interactions among the component parts. Because it is abstract of its material embodiment, it is 'multiply realisable', i.e. composed of members of functionally equivalent parts (see Auletta *et al.* [27]; and Jaeger and Calkins [28] for biological examples). It is not the identity of the components that matters, rather it is the functions they perform (e.g. a digital computer may be 'embodied by' semiconductor junctions, or water pipes and mechanical valves, without changing its identity). Crucially, 'function' is defined by a relationship between a component and the system of which it is a part. According to Cummins [29], 'function' is an objective account of the contribution made by a system's component to the 'capacity' of the whole system. At least one process performed by the component/s is necessary for a process

performed by the whole system. This implies that the function of a component is predicated on the function of the whole. This definition was recently modified to more precisely specify the meaning of 'capacity' and of whole system, thus: "A function is a process enacted by a system A at organisational level $L$ which influences one or more processes of a system B at level $L+1$, of which A is a component part" [30].

In this context, organisational level means a structure of organisation that is categorically different from those above and below in the hierarchy because it embodies novel functional information (levels may be ontological or merely epistemic in meaning: that is an open debate in philosophy). The self-organisation of modular hierarchy has been described as a form of symmetry-breaking phase transition [31], so the categories either side are quantitatively and qualitatively different. Organisational levels were defined precisely in terms of meshing between macro and micro dynamics (from partitioning the state-space of a dynamic system) by Butterfield [32] and also using category theory to specify supervenence relations and multiple-realisability among levels by List [33]. Neither definition, though, deals specifically with the phenomenon of new pattern-information 'emerging' from the organisation of level $L$ components at level $L+1$, which is responsible for the emergence of 'new phenomena'.

Ellis [34] shows that a multiply realisable network of functions, self-organised into a functional whole, emerges to (apparently) exercise 'downward causation' upon its component parts [34,35]. The organisational structure is selecting components from which to construct itself, even though it is materially composed of only the selected components. Since it is purely cybernetic (informational) in nature, the downward control is by pattern-information [28,34] which transcends the components from which it is composed. The pattern-information arises from, and is embodied by, the interactions among the components, and for these reasons it was termed a 'transcendent complex' by Farnsworth *et al.* [36]. Examples are to be found in embryonic development, where a growing cluster of cells self-organises using environmental signals created by the cells taking part [37] and the collective decision making of self-organising swarms (e.g. honey bees in which the hive acts as a unity [38]).

There is something significant here for those who conflate determinism with causation. All causal paths traced back would be expected to lead to the early universe. Despite the appearance of near maximum entropy from the uniformity of background microwave radiation, there is broad agreement that the entropy of the early universe was low and its embodied pattern-information (complexity) could not account for the present complexity, including living systems [39]. Novel pattern-information has been introduced by selection processes, especially in living systems, for which Adami *et al.* [40] draw the analogy with Maxwell's demon. Selection is equivalent to pattern matching, i.e. correlation, and is accompanied by an increase of information. Since its beginning, the entropy of the universe has been increasing [39] and some of this has been used as a raw material for transformation into pattern-information. This is achieved by creating the 'order' of spatial correlation through physical self-assembly (atoms into molecules into molecular networks into living systems). This self-assembly embodies new information in the pattern of a higher level structure through the mutual provision of context among the component parts [3]. The process of self-assembly is autonomous and follows a boot-strap dynamic, so it provides a basis for answering Strawson's [16] 'Basic Argument' in which the putative agent of free-will is an informational (pattern) structure of self-assembly.

### *2.3. Purpose and will*

Much of the literature on downward causation uses the idea of 'purpose', though many are uncomfortable with its teleological implication. The aim of this section is to form a non-teleological account of purpose and its connection with will in non-human agents.

Cause creates correlation (usually, but not necessarily, in a time-series): the pattern of any action having a cause is correlated with its cause. An action without cause is uncorrelated with anything in the universe and accordingly considered random. If an action is fully constrained, then its cause is the

constraint. Thus, freedom from at least one constraint allows the cause to be one of either: random, or exogenous control or agent control (in which 'control' means non-random cause). By definition the cause is only taken to be the agent's will if it originated in agent control. Correlation alone, between some outcome variable $x$ and some attribute $a$ of the agent, is not sufficient to establish will: a) because correlation has no direction (but metrics such as 'integrated information' [41] can resolve direction) and b) because $a$ may itself be random in origin and thereby not of the agent's making. Marshall *et al.* [42] showed that cause can be established at the 'macro' level of agent (as opposed to the 'micro' level of its components) using an elaboration of integrated information, so the pattern in $x$ can be attributed to agent-cause. Because the agent-based cause could be random ((b) above), we must form and test a hypothesis about the effect of $x$ on the agent before we can attribute the cause to the alternative of agent-will. The hypothesis is that the effect of $a$ on $x$ is to increase the overall functioning $F$ of the agent. If this were true, then to act wilfully is to reduce the entropy of $x$, by increasing the probability of an outcome $x'$ where $x' \Rightarrow F' > \bar{F}$ (and $\bar{F}$ is the average $F$). That means that the mutual information between a wilful action $a(t)$ at time $t$ and the resulting function $F(t+\tau)$, $\tau \geq 0$ is greater than zero. This mutual information between action and future functioning is taken to imply a 'purpose' for the action, so purpose is identified by the observations that:

$$H(a) + H(F) > H(a|F) \text{ and } F|a > F|r, \tag{1}$$

where $r$ is a random (comparator) variable. This is clearly an observational definition and is in some way analogous to the Turing test, but it is a test for purpose rather than 'intelligence'. It represents our intuition that if a behaviour repeatedly produces an objectively beneficial outcome for its actor, then it is probably deliberate (repeatedly harmful behaviour is also possibly deliberate, but all such actions are regarded as pathological and thereby a subject beyond the present scope).

To recap, for attributing the action to the agent's will, we must at least identify a purpose so that Eq. 1 is true. The purpose is a pattern embodied in the agent, which acts as a template for actions of the agent that cause a change in future states (of the agent, its environment, or both). We may call this pattern a 'plan' to attain an objective that has been previously set, where the objective is some future state to which the plan directs action. Specifically, let the objective be a state X (of the system or the world, etc.), which can be arrived at through a process P from the current state Y, then the purpose is a 'plan' to transform $\mathrm{Y} \rightarrow \mathrm{X}$ by the effect of at least one P and at least one function F is necessary for the process P to complete.

The homeostatic response to a perturbation, for example, has maintenance at the set-point as its purpose. Y is the perturbed state, F is some function of the internal system having the effect of causing a process P, i.e. some transition $\mathrm{Y} \rightarrow \mathrm{X}$. In general there is more than one P and more than one F for achieving each. This results in a choice of which to use: it is a choice for the agent described by the system. To make a choice requires a criterion for choosing (else the outcome is random and therefore not a choice). The criterion for choosing is a 'goal' G, consisting of one or more rules, which identify a location in a function describing the outcome (which we may call the *objective function*). In general, this location could be any and it is essential for freedom of will that it be determined by the agent of action alone. However, in practice it is most likely to be an optimisation point (in living systems, this is implied by Darwinian evolution and in designed systems, it is the basis of rational design). So, narrowing the scope, but with justification, let us take the criterion for choosing to be a 'goal' G, consisting of one or more optimisation rules. For example, of all the possible systems performing homeostasis, the purposeful one is defined as enacting $\mathrm{P}'$ such that $\mathrm{Y} \rightarrow \mathrm{X}$, with $\mathrm{P}' \in \mathrm{P}|\max(\mathcal{G})$, where $\mathcal{G}$ is an *objective function* for which the optimisation goal $\mathrm{G}(\mathcal{G})$ is satisfied, contingent upon the options (e.g. P proceeds as quickly as possible, or with minimum energy expenditure, etc.). Accordingly, 'will' is defined by a purpose which is a plan to enact a process causing a transition in state, 'as well as possible' (according to $\mathrm{G}(\mathcal{G})$) – notwithstanding the earlier comment about pathological purposes. A free-will agent has a choice of transition and a goal which

identifies the most desirable transition and the best way to enact it, from those available. These two choices can be united (by intersection), without loss of generality, to one choice of best transition.

One of the reasons for objecting to teleological terms such as 'purpose' , 'plan' and 'goal' in relation to natural systems has been the belief that a plan implies a 'designer' - the concept at the centre of the most famous battles between science and religion. This implication is not necessary and is rejected here (following the argument of Mayr [43]) . A plan is merely a pre-set program of steps taking the system from Y to X; it is the concept for which computation theory was developed. It may be designed (the work of an engineer), but also may have evolved by natural selection (which also supplies the goal - in which case it is a teleonomic system (*sensu* Mayr [43]).

A plan, as an ordered sequence of transformations, is an abstraction of information from the physical system, which for free-will must be embodied within the system. A more subtle implication of 'plan' is that as a path leading from Y to X, it is one among several possible paths: different plans may be possible, perhaps leading to different outcomes. There is a fundamental difference between this and the inevitability of a dynamic system which follows the only path it may, other than by the introduction of randomness. The reason is that for a dynamical system all the information defining its trajectory is pre-determined in the initial (including boundary) conditions and the laws of physics. The initial conditions constitute its one and only 'plan'. If a system embodies pattern-information (by its structure) which constitutes a developed plan, then this pattern-information may direct the dynamics of the system along a path other than that set by the exogenous initial conditions (though we may consider the structure of the system to be a kind of initial condition). The point is that the embodied plan gives freedom to the system, since it 'might be otherwise' - there could be a different plan and a different outcome. We see this in the variety of life-forms: each follows its own algorithm of development, life-history and behaviour at the level of the individual organism. The existence of a plan as abstract pattern-information is a pre-requisite for options and therefore freedom of action.

### *2.4. Goal, master function and will nestedness*

Now let us complete the connections between will, goal and function. The previous argument reveals an important difference between downward and any other kind of causation (considered important by Walker [44]): the former must always be directed by a purpose, for which we need to identify a goal (upward and same level causation are satisfactorily explained by initial conditions [34]). Viewing entities and actions both as the consequence of information constraining (filtering) entropic systems, then the role that is taken by initial conditions in upward causation, is taken by system-level pattern-information (the transcendent complex [36]) in downward causation. Since the goal $G$ is a fixed point in an objective function $\mathcal{G}$, it constitutes information (e.g. a homeostatic set-point) that must be embodied in the agent's internal organisation. Since the objective function $\mathcal{G}$ represents the overall functioning of the system (at its highest level), it matches the definition given by Cummins [29] and Farnsworth *et al.* [30]. The highest level function from which we identify the purpose of a system was termed the 'master function' by Jaeger and Calkins [28] , so the will of an agent is instantiated in the master function. This then identifies $\mathcal{G}$ with 'master function' and 'will' with $G(\mathcal{G})$.

Ellis [34] identifies 5 types of downward causation, the second being 'non-adaptive information control', where he says "higher level entities influence lower level entities so as to attain specific fixed goals through the existence of feedback control loops..." in which "the outcome is not determined by the boundary or initial conditions; rather it is determined by the goals". Butterfield [32] gives a more mathematically precise account of this, but without elaborating on the meaning or origin of 'goals'. Indeed, as both Ellis [34] and Butterfield [32] proceed with the third type: downward causation "via adaptive selection" they refer to fitness criteria as "meta-goals" and it is clear that these originate before and beyond the existence of the agent in question. Ellis [34] describes meta-goal as "the higher level 'purpose' that guides the dynamics" and explains that "the goals are established through the process of natural selection and genetically embodied, in the case of biological systems,

or are embodied via the engineering design and subsequent user choice, in the case of manufactured systems".

This suggests a nested hierarchy of goal-driven systems and for each, the 'goal' is the source of *causal power* and as such may be identified as the 'will' (free or otherwise). We may interpret the definition of 'deep freedom' [11] as meaning that an agent has at least two nested levels of causal power, the higher of which, at least, is embodied within the agent (as causal pattern-information). This concept may be formalised after introducing the discrete variable 'will-nestedness' $\mathcal{N}$ which counts the number of levels of causal power exercised over a system, *from within the agent as a whole* (i.e. at the level of master function), the $\mathcal{N}$th level being the highest-level internal cause of its actions.

Among organisms in general, the master function specifies the criteria by which the organism is to assess its possible future reactions to the environment. It is so much an integral part of the organism that without it, the organism would not exist. However, it was not chosen by the organism (in the sense of deep free-will) because it was created by evolutionary filtering and inherited from its parent(s) - as all known life has been created by the previous generation copying itself. For single celled organisms the biological master function is to maximise their cell count by survival and reproduction, but in multicellular organisms, this master function exists, by definition, at the level of the whole organism (the unconstrained drive to proliferate a single cell line leads to cancer). Organisms with a central nervous system, regulated by neuro-hormone systems, with their corresponding emotions, can implement more complex (information rich) and adaptable (internally branched) algorithms for the master function, which may include will-nestedness $\mathcal{N} > 1$. In humans, this is taken to such an extent that the biological master function may seem to have been superseded (but the weight of socio-biological evidence may suggest otherwise [45,46]).

### *2.5. The possibility of choice and alternative futures*

So far I have identified organisational-closure and the internal generation of a goal-based plan as prerequisites for free-will, but have not yet addressed the strict possibility of freedom for an agent constructed from elemental components that necessarily obey physical determinism. List [14] provides a philosophical argument for meeting this requirement, constructed from supervenence and multiple realisation of an agent in relation to its underlying (micro) physical level: "an agent-state is consistent with every sequence of events that is supported by at least one of its physical realizations" [14]. He shows that this may apply not only to multiple micro-histories up to $t$, but in principle includes subsequent $(t + \tau)$ sequences at the micro-level, which may map to different agent states and therefore permit different courses of action at the (macro) agent-level. To explain: for any given time $t$, the macro-state $Q_i(t)$ is consistent with a set of micro-states $\mathbf{s}$, at least one of which $s_i \in \mathbf{s}$ may lead (deterministically) to a new state $s_j \in \mathbf{s}$ at $t + 1$, with which a different macro-state $Q_j$ is consistent, thus giving the agent a choice of which micro-state history to 'ride' into $t + 1$ (this idea is developed with rigour by List [14], and illustrated with 'real-world' examples; it is the basis on which he concludes that agents may be 'free to do otherwise', despite supervening on deterministic physical processes).

The possibility of choice may exist at multiple levels of system organisation within a hierarchical structure, re-applying the same principles as identified by List [14] for each level of macro-micro relations. For any system level $L$ to have this potential, it must have the attributes of an agent-level: supervenence and multiple realisation such that pattern-information with causal power emerges at level $L$ from $L - 1$: i.e. a transcendent complex exists at level $L$. But this does not necessarily give *free will* to a system of that level, since for that, it must be organisationally closed. If it were not so, we would not be able to identify the system at level $L$ as an entity to which free-will could be ascribed. Thus will-nestedness cannot be attributed to levels of organisation below that of the Kantian whole. Since the Kantian whole is, by definition, the highest level of organisation to which free-will may be ascribed (any causal power beyond it rules out its free will), then will-nestedness can only apply at the level of the Kantian whole. Given this, the will-nestedness must be constructed from purely

organisational, i.e. pattern-informational and therefore be purely computational in nature. This is an important deduction: free-will can only be an attribute of a Kantian whole and it can only result from the cybernetic structure at the level of the Kantian whole.

## 3. Choosing possible futures: The *computational condition*

We see that for free-will, an agent must have an independent and internally generated purpose for action and that this requires it to be organisationally closed. Free-will further requires the agent to use this purpose to choose among options. To do so, it needs an internal representation of *possible* futures from which to choose and an internally generated means of choosing. We now turn to the conditions which enable these essentially computational facilities.

### *3.1. Information abstraction*

Information abstraction at the organisational boundary is crucial to achieving autonomy. It strips off the physical effect of the external environment to take only the abstract information as a signal. Causes are transformed into signals, their effects being rendered responses (which thereby may become optional). It is this separation of information (as signals) from the physical force of cause and effect that releases the agent from attachment to the cause-effect determinism of its environment. The material apparatus for performing this task is a transducer: the tegumental membranes of bacteria and other cells contain a wide variety of transducers (receptors) and we expect an artificial robot to be well equipped with them too. This is not a merely technical point. The closure condition gives the system a degree of causal independence and the boundary transducers give it a sensitivity to its environment whilst preserving this independence. The boundary is the place where the inevitability of cause and effect of the environment meets that of the internal processes of the system and the transducers are the interface between these causal chains. Internally to the system, the environment is 'reduced' to abstract, representational, information by the transducers [47]. Now the question is, what must the agent do with this information in order to exercise free-will?

Of course the answer is to compute: more specifically, to perform transformations on the data as a result of a sequence of physical changes in the physical structure of the agent. Such changes are described by automata theory, for which the most basic automaton has two states and can potentially change state on receiving a signal to which it responds: it is a switch (e.g. a protein molecule with two conformations is an 'acceptor' of all strings from the alphabet $\{0,1\}$, where these symbols may represent the presence / absence of e.g. another molecule or a level above a threshold (e.g. temperature)). Obviously, the switch is the elemental component for generating *discrete options*. A less obvious, but crucial property of the switch for the physical embodiment of computation is 'thermodynamic indifference'. Walker and Davies [48] focus on computation in explaining the origin of life, referring to genetics-first theories as 'digital-first', emphasising the need for 'programmability' and its provision by informational polymers (the genetic oligomers RNA, DNA etc.). By programmability they meant that components of a system are configured so that the system state can change reversibly, approximately independent of energy flow: i.e. changes of state are not accompanied by substantial changes in potential (stored) energy. If they were, then switching would always be biased by the difference in energetic cost between e.g. switching on and off. Energetically unbiased switching is the physical underlying mechanism of information abstraction referred to by Walker and Davies [48]. In reality, switching (and state-changes in general) always have energetic consequences (more deeply, there is always an exchange of entropy between the system and an external energy source), which is one of the reasons an autonomous agent must complete work cycles as Kauffman [49] specifies. What makes the informational polymers (e.g. DNA) of life special is the fact that they are reversible in a way that is thermodynamically indifferent (or very nearly so - see Ptashne [50]). Any ordered set of $n$ switch positions (e.g. 1,1,0,0) has very nearly the same potential energy as any other ordered set of $n$ (e.g. 0,1,0,1).

For free-will, both the self-assembly of autocatalytic systems and the computational requirements (switches and memory) are jointly necessary. Walker and Davies [48] and Walker [44] proposed that the autonomy of living systems arises from the combination of 'analogue' chemical networks and 'digital information processing'.

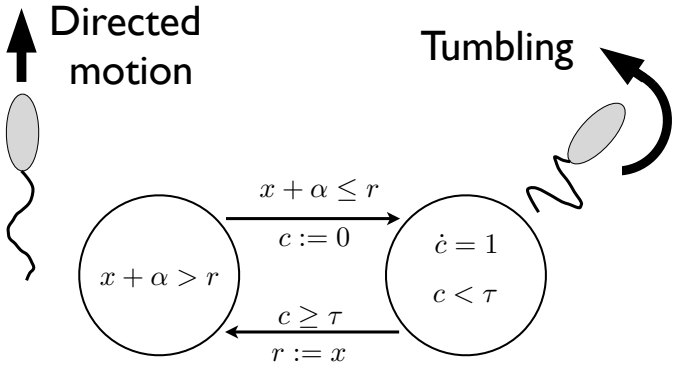


**Figure 1.** A bacterial chemotaxis controller described as a hybrid automaton, realised in practice by e.g. *E-coli* species, but also as an engineering design for a molecular robot, using DNA-based components by Hagiya *et al.* [10], from which the figure, slightly modified, is taken. Note that $\alpha$ and $\tau$ are internal set-points, which constitute pattern-information embodied in the molecular robot's structure and are ultimately determined by Darwinian evolution (natural robots) or intentional design (human artefacts). The bacterium responds to the environmental concentration (e.g. of glucose) - its objective function $x$, which is detected by a membrane transducer which generates the internal variable $r$. It is searching for a higher concentration when $x$ is below a set threshold (directed motion) and tumbles randomly for a time set internally by $\tau$ using an internal 'clock' signal $c$, whenever the concentration is at least equal to the threshold. The swimming and tumbling are modes of action of its flagella. The objective function of this system is the experienced concentration $x$ and the goal is $\max(x)$. The chemotaxis controller exhibits will-nestedness of 1.

A hybrid automaton is a good model for the construction of a free-will agent. Fig. 1 (from Hagiya *et al.* [10]) shows an example of a hybrid automaton which combines discrete-state with continuous dynamical systems, such that the discrete states (as modes of functioning) determine the system's responses to dynamic variables and these responses potentially influence the trajectory of the dynamics. In this pedagogical example from Hagiya *et al.* [10], there are two state variables $\alpha$ and $\tau$ which determine the set points for autonomous chemotaxis behaviour in the bacterium represented. This system can freely maximise its (experienced) environmental concentration of $x$ (e.g. glucose), so its goal $\mathrm{G}(\mathcal{G})$ is defined, but it cannot choose how (hence $\mathcal{N} = 1$), so it cannot express free-will in the deep sense. However, if it had independent control over $\alpha$ and $\tau$, with the ability to adjust these values according to a plan of its own making, then it would fulfil the condition of having willed its own behaviour ($\mathcal{N} = 2$), at least in the sense defined in the previous section. As part of a living bacterium, the system depicted would be a component of a Kantian whole (the free-living organism), so all that is missing for free-will is a plan for determining $\alpha$ and $\tau$ according to an internally generated goal (a master function) and some means of computing this. The computational requirements for free-will are identified as follows.

### *3.2. The representation of self*

Firstly, a free-will agent must maintain an internal representation of itself, and also the effect of its environment on its internal state, to enable it to assess each of its options for action.

At the root of automata theory lies an attempt to fully describe (and therefore predict) the behaviour of a system without a detailed mechanistic account of its internal workings (the black box approach). The system is captured in the mapping between environmental stimuli and responses:

$$R(t+1) = W(\mathbf{S}(t), S(t)) \tag{2}$$

where $R(t)$ and $S(t)$ are the response and stimulus, and $\mathbf{S}(t)$ is the history of stimuli experienced by the system, *from the beginning of its formation* up to $t$. This presents an immediate problem, since in general $\mathbf{S}(t)$ is arbitrary and infinite in range (it is instructive to think of Strawson's [16] 'Basic Argument' in terms of Eq. 2: the response of a system, in general, depends on its environment from before the system came into existence). The solution to this indefinite $\mathbf{S}(t)$ problem (provided by Moore [51]) is to assume that the infinite set of $\mathbf{S}(t)$ may be partitioned into a finite number of disjoint equivalence sets, each containing the histories that are equivalent in their effect on $R(t+\tau), \forall\tau$, where $\tau$ is in the interval $[0, \infty]$. These equivalence sets are represented as states $Q(t)$ of the system, so that:

$$R(t+1) = W(Q(t), S(t)) \text{ and } Q(t+1) = U(Q(t), S(t)), \tag{3}$$

where $U(Q(t), S(t))$ is the transformation function dictating the transition of system state given its present state and that of the stimulus (see Minsky [52]. p16-17). Thus $Q(t)$ represents how the system is now, *given* its previous history of experiences. $Q(t)$ corresponds to the agent-state of List [14], which is multiply-realisable and which is, at least in principle, free to take more than one value at some point in the future $t+\tau$, despite supervening on a a wholly determined set of micro-histories. List [14] showed the possibility of choice at the agent-level, but this does not necessarily mean that $Q(t)$ is indeterminate. Specifically, for free-choice (of $Q(t+1)$ and implied $R(t+1)$), the direction taken at the branching point $t$ must be determined by a process internal to the agent that represents its 'purpose' (as defined earlier). For the choice to be purposful, it must be based on an assessment of the outcomes that would arise from choosing each of the options. This entails a prediction of possible futures, for which a free-will system must have a model of itself in its environmental context.

The question now is, how can a system create such a representation by and for itself - not 'programmed' by some exogenous source of information? The answer seems to be as it is with material self-assembly: a boot-strap, step by step, gathering of pattern-information embodied in form, such that as the form grows, it increases in complexity. In the particular case of building a model of self and environment, this process is one of learning, for which the field of 'machine learning' provides our understanding. Well known advances in this field have already led to sophisticated learning among pre-existing (i.e. not self-assembled) computation systems such as deep neural networks etc.. The difference here is that the learning is not merely a statistical problem, but one of simultaneous self-construction, which must begin with simple systems, so in the remainder of this section, only basic and simple systems capable of unsupervised learning are discussed.

#### 3.2.1. Learning in a constant environment

The most basic form of learning is operant conditioning (reinforcement learning), described mathematically by Zhang [53] as follows. Let $R_i$ be one of a set of $N$ possible responses ($R_i \in \mathbf{R}$, $i \in [1, N]$) in a constant environment, occurring with probability $r_i$. For each response there is a 'reward' $\rho$ (which coincides with the objective function $\mathcal{G}$ that defines the goal of the system: $\rho \rightarrow f(\mathcal{G})$ ), so that the incremental change in probability of the $i$th response is:

$$\Delta r_i = a\rho_k(\kappa_{k,i} - r_i), \tag{4}$$

where $\kappa_{k,i}$ is the Kronecker delta function (equal to 1 if $i = k$, else equal to zero) and $a$ is a learning rate constant. Since the average change in response over the ensemble of possible responses is the frequency-weighted sum: $\Delta\widehat{r} = \sum_k^N r_k \Delta r_k$, the result is that the frequency of the $i$th response

incrementally increases in proportion to the difference between its reward and the average over all rewards:

$$\Delta r_i = a r_i(\rho_i - \widehat{\rho}) \text{ , where } \widehat{\rho} = \sum_k^N r_k \rho_k. \tag{5}$$

The dynamic quantified in Eq. 5 describes learning by maximising the reward experienced. Such learning is equivalent to making an increasingly accurate model of the (static) relationship between the agent's internal state and the environment, via (Bayesian) 'trial and error' sampling of responses. Given a constant environment, the solution to Eq. 5 yields a single, reward maximising response: $R^* = R_i$ such that $\rho_i = \widehat{\rho}$ and $r_i = \kappa_{k,i}$.

To achieve this in practice the system must at least keep a record of the reward for the last response made and the average reward, for which an automaton with an external memory is required (e.g. a push-down automaton, though this is still essentially a DFA). Quantifying the complexity of such a system is probably best achieved through a programme (algorithmic) complexity measure since the information instantiated by such an automaton is almost all in its transition mapping and there are robust methods for reducing this to the minimum description, leading directly to the Kolmogorov complexity. The process of learning can be interpreted in information terms: if the starting probability distribution of $\mathbf{R}$ is $\tilde{r}$, then the initial Shannon entropy of the system is $H = -\sum_i^N \tilde{r}_i \log(\tilde{r}_i)$ and the final entropy is zero: having completed its learning, the system has no uncertainty about the best way to respond to this environment. In this state, the automaton is a complete representation of its interaction with its environment (i.e. the distribution of rewards over its repertoire of responses) and it embodies exactly $H$ units of information: a quantity which should match the algorithmic complexity measure (though not tested here).

### 3.2.2. Extension to a variable environment

Generalising to a variable environment, for which a set of finite states $\mathbf{S}$ is an adequate representation, there would exist a reward maximising response for each state: $R_i^* \rightarrow (S_i)$, such that $R_i^*$ solves Eq. 5. The agent may choose to maximise its reward over all $\mathbf{S}$, and to enable this, it must learn the best response for every $S_i \in \mathbf{S}$. In information terms, first let $H(\mathbf{s})$ be the entropy of the environment having probability distribution $\mathbf{s}$ and $H_t(\mathbf{r_t})$ be that of the responses, given their probability distribution $\mathbf{r_t}$ at time $t$. The Shannon information the agent has about its environment (in terms of its rewards) is:

$$\begin{aligned} I_t(\mathbf{s} : \mathbf{r_t}) &= H(\mathbf{s}) + H(\mathbf{r_t}) - H(\mathbf{s}, \mathbf{r_t}) \\ &= H(\mathbf{s}) + H(\mathbf{r_t}) - H(\mathbf{s}|\mathbf{r_t}), \end{aligned} \tag{6}$$

meaning that the agent is learning both the distribution $\mathbf{s}$ and the reward associated with each $S_i \in \mathbf{S}$. This *mutual information* is embodied in the structure of the agent and can be used as a measure of its complexity. The structure of the agent may be too simple to embody as much as the maximum mutual information, in which case its learning will be limited and it will not make optimal responses, so Eq. 6 is a measure of the minimum complexity required for optimal behaviour from the agent. It would be possible for an agent to implement this learning system by 'growing' multiple copies of the DFA with memory (one for each $S_i \in \mathbf{S}$) that is used for a constant environment. The output of each of these would then be the input to a further DFA which it uses to maximise the reward across all of them. This 'growth' would be enacting meta-learning: the agent would increase its complexity in response to rewards. The number of states in $\mathbf{S}$ is not known by the agent a-priori, so the number of DFAs needing to be 'grown' is indeterminate. Further, account should be taken of the extended time needed to perform such laborious learning and the consequences of the agent being wrong in so many trials. It seems that for practical reasons there comes a point when a more powerful kind of

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 February 2017 doi:10.20944/preprints201702.0105.v1

Peer-reviewed version available at Entropy 2017, 19, , 237; doi:10.3390/e19050237



computation becomes necessary. In computation terms, such a learning problem requires at least a finite and non-volatile memory, effectively to store multiple instances of the single learning problem encountered in a constant environment. For this reason a Turing Machine would be a more realistic option.

### 3.3. The Free-Will Machine

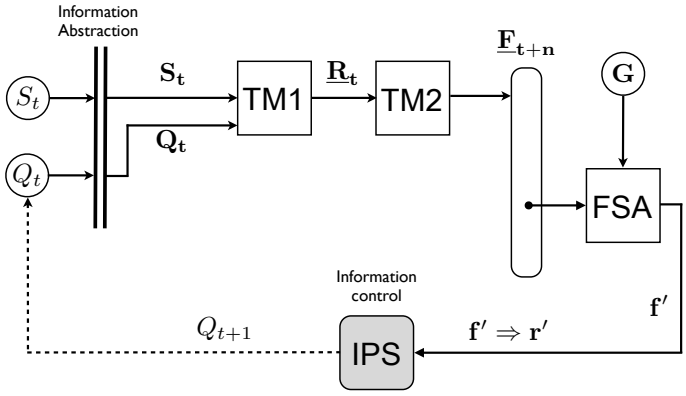


**Figure 2.** A conceptual 'free-will' machine, which generates predictions of its state in alternative futures $\underline{F}_{t+n}$using an internal representation of itself interacting with its environment, selecting the optimal from among these, using a goal-bases criterion G, in which the goal is internally determined (further explanation in the text).

These computational requirements for free-will to become possible are brought together in the hypothetical 'free-will machine' of Fig. 2. This information processing must be implemented by the agent to which free-will is ascribed and that agent must, further, be a Kantian whole for the requirements of free-will to met. The current state of the environment (external) and of the agent (internal) are derived by information abstraction from the physical world: the array of receptor molecules in the cell membrane, the nervous senses of an animal, or the transducers of a human artefact all perform this task. The first Turing machine implementation TM1 constitutes a representation of the agent in the present (relevant aspects of its environment are included) and is informed (updated) by the state information $Q_t$ and $S_t$. The function of this representation is to identify the set of possible responses $\underline{R}_t$ that the agent can make, given $Q_t, S_t$ (underline notation denotes a set).

TM2 uses these hypothetical responses to compute the set of possible futures at a time $t + n$ ($n$ may take any positive value) $\underline{F}_t$ for which in a simple case TM2 : $\underline{R}_t \rightarrow \underline{F}_t$ is a map of responses onto possible futures (simple, because there is not necessarily a 1:1 relation between $\underline{R}_t$ and $\underline{F}_t$, but we need not be concerned with that at present). There is no limit on the number of possible futures that TM2 may compute, but it must be at least 2 for a choice to exist. These possible futures are each represented by a set of states $\underline{f}$, each member $f_n$ being equivalent to a prediction of a possible $Q_{t+n}$. The FSA chooses from among these, using a selection criterion based on the objective function defined by a goal G, which is generated by the agent (not exogenously). This goal is the maximisation of the master function, (e.g. for a living agent this is life-time reproductive success). The goal enables the optimal possible future state to be recognised (it is the one which maximises the master function) and this future state $f'$ implies an optimal response $r'$ (in general there could be more than one, in which case the agent will be indifferent among them). Having selected an optimal response, the agent then must implement it in the physical realm. Since the computation of $r'$ has been conducted in the realm of information, this step appears to involve the control of material by information. In practice all the computation and indeed all the information is instantiated by material and energy acting in the

physical realm, so our cybernetic model is merely an abstract representation of the organisation of the physical processes which lead to the implementation in the physical system and this return to physical reality is represent by the action (IPS). Such implementation inevitably results in a transformation of the agent into a new state $Q_{t+1}$, together with $S_{t+1}$ and this restarts the cycle. It may be noted that Von Neumann's self-replicating automata are proven universal Turing machines [25] and Turing machines are thought to be common among living systems, so this computational arrangement is not beyond the bounds of possibility.

## 4. Discussion and Synthesis

To summarise, the essential requirements for free will are:

- R1) There must be a self to self-determine.
- R2) There must be a non-zero probability of more than one option being enacted.
- R3) There must be an internal means of choosing among options (which is not merely random, since randomness is not a choice).

for R3 to be fulfilled:

- R4) Options must be generated from an internal model of the self which can calculate future states contingent on possible responses.
- R5) Choosing among these options requires their evaluation using an internally generated goal.
- R6) For 'deep free-will', at least two nested levels of choice and goal (R4-R5) must be enacted by the agent.

R1 and references to 'internally generated' are fulfilled by organisational closure. For R2, the possibility of options, which implies 'multiple futures' has been established for the level of the agent by List [14]. R3 and its predicates R4-R5 imply a minimum level of computational power, which in principle can be met by a small set of Turing machines, which may in principle be implemented by a Von Neuman architecture computer, a network or cellular automaton-based or any other sort of computer, including a biomolecular system such as found in higher animal life, but it seems to be beyond the power of a single living cell (though that last point is not yet established). R4 in particular seems to require a finite memory (the size depends on the complexity of the agent and its environment) and R6, the qualifier for deep free-will, adds a little more to the computing power necessary, but it is important to note that this extra is not a step-change: it is not qualitatively more demanding than the automated decision making required by R3.

The question of free-will is not one of whether an agent's actions are caused, since all actions ultimately have a cause. The ultimate cause of any action can be understood as resulting from selection over random actions by a pattern, which leaves a correlation with the pattern that caused the selection (instantiating pattern-information). All living organisms, including people, were produced by information-pattern filtering, proximally by molecular replication (creating inheritance) and ultimately by Darwinian selection. All human artefacts were created by following a design pattern (though it may not have been completed before artefact construction), so they correlate with their design. Even inanimate objects, such as stars, lakes and sand grains, owe their form to the information-pattern of underlying physical laws, Pauli's exclusion principle and the distribution of matter and energy in space following the big-bang. To this, we must add randomness which has been entering as 'informational raw material' into the universe, disrupting the original patterns and opening opportunities for novelty (evolutionary for life) and more widely directing the course of the universe in unexpected ways as its history tracks a course in the highly ergodic space of possibilities.

Taking Strawson's [16] Basic Argument seriously, this pattern-correlation and the injection of randomness both deny free will. From them, we obtain a model in which the identity of all things, including human beings, is an illusion: as if the universe was all one complex manifestation which only *appears* to include separate agents. Closer inspection shows how the nested-hierarchical

construction of this complexity entails the creation of genuinely new pattern-information, caused by and embodied in the interactions among component parts of putative agents. This novel pattern-information transcends its component parts and can exert downward causation upon them. Some structures (such as autocatalytic sets) created this way are organisationally closed (though materially and thermodynamically remain open systems). Because of this, their internal dynamics are, at least partially, separated from the external dynamics of their environment and this gives them an organisational boundary, enabling internal to be defined against external. At this boundary, external and internal chains of cause and effect interact through transducers which transform physical determinism into stimulus-response relations. Systems with these properties are essentially cybernetic and although their low-level processes are continuous with the rest of the universe, List [14] has shown that in principle they may have options for their next state and response: they are freed from physical determinism. To translate this freedom into free-will, requires that the (partially) independent agent chooses from among its options and this entails an internal computation of possible futures and their evaluation against a goal representing the fulfilment of the agents' 'master function' (i.e. its purpose). This goal is a fixed point in an objective function which may be simple (as in a homeostatic system), but also arbitrarily complex and multi-layered, taking account of multiple time-scales and interactions with other agents. If the objective function is at least two-layered (will-nestedness $\mathcal{N} \geq 2$), then it effectively has a choice of what to choose and thereby could fulfil the established definition of (deep) free-will [11,12]. This calls into question the idea that free-will is an all or nothing capacity, instead, it suggests free-will to be a discrete quantity and even something we could in principle measure as a trait of a system.

The reason is that deep free-will has so far been defined as the freedom to choose ones will, but the analysis presented here shows that to be wilful, a choice must be purposeful, which means optimising an objective function. Freedom is in the choice of objective function. Since therefore, the core of will is the objective function, deep freedom is the freedom to choose this, but to be wilful, this choice in turn must optimise a hierarchically superior objective function, which must have been determined by something. We can conceive of a large but finite nested set of such objective functions, but *ultimately* the highest of them all must be provided either arbitrarily (e.g. at random) or by natural selection (or its unnatural equivalent), or by design: in all cases, not the free choice of the agent. This applies even to human beings, who are still subject to 'design' by inheritance, selected by evolution. The depth of free-will is therefore a discrete, finite variable ('will-nestedness' $\mathcal{N}$) and we may speak of one kind of agent being more deeply free than another, but no kind of agent can be *ultimately* free-willed in the sense meant by deep free-will (presumably, humans attain the highest will-nestedness of known agents).

If interpreting the philosophers' definition of deep free-will as $\mathcal{N} \geq 2$ is correct, then it is not hard to achieve in principle. The impediment to a non-life robot acquiring that sort of free-will is not computational, it is the closure constraint with its requirement for bootstrapping self-assembly (especially since this includes the 'growth' of the sort of computational apparatus indicated in Fig. 2). That is a problem already solved by life. Quite likely $\mathcal{N} \geq 2$ in most or all organisms having at least a limbic system, so free-will defined this way can be attributed to most or all vertebrates [54]. The computational requirements for exercising purposeful choices are not very challenging for artificial computers. Among human artefacts, including what most people would define as robots, the organisational closure condition is the major hurdle not yet leapt. For the time-being, it seems free-will, as defined here, is a unique property of living things.

**Acknowledgments:** This work was unfunded, but was inspired by attendance at the workshop "Information, Causality and the Origin of Life" in Arizona State University at Tempe, AZ Sept. 30th - Oct. 2nd 2014, funded by the Templeton World Charity Foundation. The cost of open access publication was met by the Queen's University Belfast.

**Author Contributions:** Just one author

**Conflicts of Interest:** The author declares no conflict of interest. The founding sponsors had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, and in the decision to publish the results.

## Abbreviations

The following abbreviations are used in this manuscript:

| | |
|---|---|
| DFA | Derterminate Finite State Automaton |
| DM | Determinate Machine |
| FSA | Finite State Automaton |
| NDM | Non-Determinate Machine |
| TM | Turing Machine |
| UTM | Universal Turing Machine |

## Bibliography

1. Sharov, A.A. Functional Information: Towards Synthesis of Biosemiotics and Cybernetics. *Entropy* **2010**, *12*, 1050–1070.
2. Hazen, R.M. The emergence of patterning in life's origin and evolution. *Int. J. Develop. Biol.* **2009**, *53*, 683–692.
3. Farnsworth, K.; Nelson, J.; Gershenson, C. Living is Information Processing: From Molecules to Global Systems. *Acta Biotheor.* **2013**, *61*, 203–222.
4. Weizsäcker, C.v. *Die Einheit der Natur [The unity of nature]*; Deutscher Taschenbuch Verlag: Munich, Germany, 1974.
5. Stonier, T. Information as a basic property of the universe. *Bio Systems* **1996**, *38*, 135–140.
6. Devlin, K.J. *Logic and information*; Cambridge University Press: Cambridge, U.K., 1992.
7. Floridi, L. Information. In *The Blackwell Guide to the Philosophy of Computing and Information*; Floridi, L., Ed.; Blackwell Publishing Ltd, 2003; pp. 40–61.
8. Szostak, J.W. Functional information: Molecular messages. *Nature* **2003**, *423*, 689–689.
9. Hazen, R.M.; Griffin, P.L.; Carothers, J.M.; Szostak, J.W. Functional information and the emergence of biocomplexity. *Proc. Natl. Acad. Sci. U. S. A.* **2007**, *104*, 8574–8581.
10. Hagiya, M.; Aubert-Kato, N.; Wang, S.; Kobayashi, S. Molecular computers for molecular robots as hybrid systems. *Theor. Comp. Sci.* **2016**, *632*, 4 – 20. Verification of Engineered Molecular Devices and Programs.
11. Kane, R. *A contemporary Introduction to free will.*; Oxford University Press.: Oxford, UK, 2005.
12. Westen, P. Getting the Fly out of the Bottle: The False Problem of Free Will and Determinism. *Buffalo Criminal Law Review.* **2005**, *8*, 599–652.
13. Van Inwagen, P. *An Essay on Free Will.*; Oxford University Press.: Oxford, UK, 1983.
14. List, C. Free will, determinism, and the possibility of doing otherwise. *Noûs* **2014**, *48*, 156–178.
15. Van Inwagen, P. Some Thoughts on An Essay on Free Will. *Harvard Rev. Phil.* **2015**, *22*, 16–30.
16. Strawson, G. *Freedom and Belief*; Oxford University Press: Oxford, UK, 1986.
17. Kauffman, S.A. Autocatalytic sets of proteins. *J. Theor. Biol.* **1986**, *119*, 1–24.
18. Kauffman, S.; Clayton, P. On emergence, agency, and organization. *Biol. Philos.* **2006**, *21*, 501–521.
19. Ginsborg, H., A Companion to Kant.; Wiley-Blackwell, 2006; chapter Kant's biological teleology and its philosophical significance.
20. Hordijk, W.; Steel, M. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *J. Theor. Biol.* **2004**, *227*, 451–461.
21. Hordijk, W.; Hein, J.; Steel, M. Autocatalytic Sets and the Origin of Life. *Entropy* **2010**, *12*, 1733–1742.
22. Segré, D.; Ben-Eli, D.; Lancet, D. Compositional genomes: Prebiotic information transfer in mutually catalytic noncovalent assemblies. *Proc. Natl. Acad. Sci. U. S. A.* **2000**, *97*, 4112–4117.
23. Luisi, P. Autopoiesis: a review and a reappraisal. *Naturwissenschaften* **2003**, *90*, 49–59.
24. Davies, J.A. *Life Unfolding: how the human body creates itself.*; Oxford University Press: Oxford, UK, 2014.
25. Von Neumann, J.; Burks, A. *Theory of self-reproducing automata*; Illinois University Press: Chicago, IL.,USA., 1966.

26. Varela, F.; Maturana, H.; Uribe, R. Autopoiesis: the organization of living systems, its characterization and a model. *Curr. Mod. Biol.* **1974**, *5*, 187–96.
27. Auletta, G.; Ellis, G.; Jaeger, L. Top-down causation by information control: from a philosophical problem to a scientific research programme. *J. R. Soc. Interface* **2008**, *5*, 1159–1172, [http://rsif.royalsocietypublishing.org/content/5/27/1159.full.pdf].
28. Jaeger, L.; Calkins, E.R. Downward causation by information control in micro-organisms. *Interface Focus* **2012**, *2*, 26–41.
29. Cummins, R. Functional Analysis. *J. Philos.* **1975**, *72*, 741–765.
30. Farnsworth, K.D.; Albantakis, L.; Caruso, T. Unifying concepts of biological function from molecules to ecosystems. Oikos - In review.
31. Lorenz, D.; Jeng, A.; Deem, M. The emergence of modularity in biological systems. *Phys. Life Rev.* **2011**, *8*, 129–160.
32. Butterfield, J. Laws, causation and dynamics at different levels. *Interface Focus* **2012**, *2*, 101–114.
33. List, C. Levels: descriptive, explanatory, and ontological.
34. Ellis, G. Top-down causation and emergence: some comments on mechanisms. *Interface Focus* **2012**, *2*, 126–140.
35. Ellis, G. On the nature of causation in complex systems. *Trans. R. Soc. S.Afr.* **2008**, *63*, 1–16.
36. Farnsworth, K.D.; Ellis, G.F.; Jaeger, L., Living through Downward Causation; Cambridge University Press, 2017; chapter 13, pp. 303–333.
37. Gilbert, S. *Developmental Biology*; Sinauer Associates: Sunderland, MA. USA., 2013.
38. Seeley, T. *Honeybee Democracy*; Princeton University Press: Princeton, USA., 2010.
39. Lineweaver, C.H.; Egan, C. Life, gravity and the second law of thermodynamics. *Phys. Life Rev.* **2008**, *5*, 225–242.
40. Adami, C.; Ofria, C.; Collier, T. Evolution of biological complexity. *Proc. Natl. Acad. Sci. U. S. A.* **2000**, *97*, 4463–4468.
41. Hoel, E.; Albantakis, L.; Marshall, W. Can the macro beat the micro? Integrated information across spatiotemporal scales. *J. Conscious. Sci.* **2016**, *1*, niw012.
42. Marshall, W.; Albantakis, L.; Tononi, G. Black-boxing and cause-effect power. *ArXiv e-prints* **2016**, [arXiv:q-bio.NC/1608.03461].
43. Mayr, E. Teleological and Teleonomic: A New Analysis. *Boston Studies Phil. Sci.* **1974**, pp. 91–117.
44. Walker, S.I. Top-down causation and the rise of information in the emergence of life. *Information* **2014**, *5*, 424–439.
45. Wilson, D.; Wilson, E.O. Rethinking the Theoretical Foundation of Sociobiology. *Q. Rev. Biol.* **2007**, *82*, 327–348, [http://dx.doi.org/10.1086/522809]. PMID: 18217526.
46. Wilson, E.O. *Sociobiology: The New Synthesis.*; Harvard University Press: Cambridge (MA), USA., 1975.
47. Danchin, A. Bacteria as computers making computers. *FEMS Microbiol. Rev.* **2009**, *33*, 3–26.
48. Walker, S.; Davies, P. The algorithmic origins of life. *J. R. Soc. Interface* **2013**, *10*.
49. Kauffman, S.A. *Investigations*; Oxford University Press, 2000.
50. Ptashne, M. Principles of a switch. *Nat. Chem. Biol.* **2011**, *7*, 484–487.
51. Moore, E.F., Automata Studies; Princeton University Press., 1956; chapter Gedanken-Experiments on sequential machines., pp. 129–153.
52. Minsky, M.L. *Computation – Finite and Infinite Machines.*; Prentice Hall: Englewood Cliffs, NJ. U.S.A., 1967.
53. Zhang, J. Adaptive learning via selectionism and Bayesianism, Part 1: Connection between the two. *Neural Networks.* **2009**, *22*, 220228.
54. Bruce, L.L.; Neary, T.J. The Limbic System of Tetrapods: A Comparative Analysis of Cortical and Amygdalar Populations. *Brain Behav. Evol.* **1995**, *46*, 224–234.