

A Model for Artificial General Intelligence

Andy E. Williams, Nobeah Foundation, Nairobi, Kenya

Abstract

A Functional Modeling Framework (FMF) for defining and comparing models of consciousness and cognition has recently been developed. This framework proposes to have the capacity to represent the complete set of the functionality of human consciousness and cognition, which if true, would suggest that all models of consciousness and cognition can be represented within the framework. The framework also proposes to define the criteria for a model of cognition to have the potential for the general problem solving ability commonly recognized as true human intelligence. The FMF provides a single framework for defining models of consciousness and cognition that is human-centric in that the functions can be validated through experiments that can be performed within innate human self-awareness rather than being dependent on assumptions made by any specific model. This human-centric functional modeling approach is intended to enable different models of AGI to be more easily compared so research can reliably converge on a single understanding, enabling the possibility of massively collaborative interdisciplinary projects to research, and implement models of consciousness or cognition where such massive collaboration has not proved possible before. The FMF defines requirements for all the functional components defined by the framework, but leaves specific models to define their own implementations. This paper summarizes a model of cognition developed within this framework that is proposed to meet the criteria of an AGI as defined within this framework. This description is expanded in a number of other papers.

Background – The Problem of Cognition

Problem definitions don't exist in a vacuum. What is a problem to one individual might not be a problem to another. In this way, problems are defined by the perspective of the problem solver. In this sense, general problem solving ability requires a general perspective, that is, a general problem to solve. All living organisms have the capacity to adapt to any state of their environments that can be accommodated within the domains of adaptability they have evolved. And in this sense all living organisms have general problem solving ability. Human beings are the one species known to have the capacity for abstract reason as an adaptive domain. The one perspective shared by all humans, that is, the one general problem shared by all humans, is the problem of well-being. Where current AI models might lack a definition of well-being, and therefore lack a general problem to solve, the model of cognition described in this paper aims to provide the capacity for general problem solving ability through providing a quantifiable definition of well-being [4], through providing a mechanism for detecting the state of well-being and for executing a library of reasoning processes to navigate to a different state of well-being, and through governing that navigation with a mathematical function of these reasoning processes that is globally stable in well-being, though the reasoning processes themselves may follow a potentially chaotic (unstable) path in the reasoning space defined for them.

Introduction

The Functional Modeling Framework (FMF) defines living systems in terms of functional components, each of which has a set of discrete functions. The FMF then describes the systems in terms of the states accessible through those functions. These states form a functional state space through which the system navigates a path. In this way, the FMF can be used to represent and compare models of living systems in terms of how they implement those functions, and in terms of how those implementations govern the dynamics of the system through that functional state space. In the FMF the human system is defined in terms of a hierarchy of adaptive processes, a set of functional systems that each arise from a different adaptive process, and as a set of functional components that also each arise from an adaptive process,

where these functional components might be present in multiple different functional systems as well. This paper focus on only one adaptive functional component (adapting through reasoning).

In the FMF each adaptive process in the hierarchy of adaptive processes might be implemented by an entire system such as the consciousness system that implements the ability to adapt through directing awareness, the cognitive system that implements the ability to adapt through directing reason, or the emotional system that implements the ability to adapt through directing motivation. These systems themselves consist of a hierarchy of functions. The hierarchy of functions in the cognitive system is described in the sections of this paper to follow. The importance of providing adaptive processes that adapt functions to become more fit is that all constraints against optimizing problem definitions or constraints against optimizing choice of solution to that problem, are forces that act to reduce problem solving fitness. Where narrow AI requires enough problem solving fitness to select an optimal solution, general problem solving systems such as AGI require fitness at a broader range of problems, including deciding the optimal problem to solve. If some of the constraints against problem definition and problem solving might exist in each adaptive domain, and if each adaptive domain must exist in a hierarchy of other adaptive domains if it is itself to have the capacity for adaptation, then removing these constraints must require a hierarchy of adaptive processes that enable the system to adapt to do so.

The approach to functional modeling used in this paper may be a radical departure in that it attempts to create a bridge between approaches for understanding the human system in terms of functions that can be observed in the individual's own self awareness, and scientific approaches that aim to do the same through measurement that can be verified externally. Where the vast tradition of such observations has not before been readily accessible to the sciences, this human-centric approach formalizes the process of representing systems in terms of their functions that human beings already use intuitively, so that it is possible to leverage that vast understanding. Furthermore, rather than introducing jargon that forces researchers to adjust to an individual researcher's way of framing solutions, this human-centric formalization attempts to frame solutions in a way that can be intuitively understood in natural language by anyone with a deep understanding of the problem.

The Functional Modeling Framework (FMF) developed using this approach decomposes systems into a set of functions through an exercise in "functional decomposition" that leverages the most commonly perceived understanding of what humans observe the system as functioning to achieve. These functions are used to define the states the system can enter through executing those functions, thereby creating a representation that unambiguously defines each state of the system as a unique point in a "functional state space", and that unambiguously defines each function of the system as a transform between two points in that functional state space.

As an example, in a number of traditions for observing the human system in innate human self-awareness the human system is defined in terms of the functional components: body, emotions, mind, and consciousness. The mind (the system of cognition) is defined in this human-centric functional modeling approach as consisting of a number of functional units that process neural signals into concepts, and a number of functional units that process concepts to provide the functions involved in cognition.

In addition, breaking systems down to a set of discrete, objectively defined functions that can be independently implemented by people from different disciplines may facilitate massive interdisciplinary cooperation to do so, where such cooperation has not proven possible before. Functional modeling approaches are commonly used in systems and software engineering to facilitate

cooperation in the design of complex systems by removing the need for individuals in interdisciplinary teams to understand each other's approaches. A functional modeling approach that is also human-centric enables functional modeling to be extended to systems like consciousness or cognition for which functions can be observed within our innate human awareness, but for which the mechanisms of operation are unknown, and being unknown with no universally agreed upon models, researchers may propose models of those mechanisms from mathematics, neurology, physics, or a wide variety of other backgrounds that aren't shared. Human-centric functional modeling creates the potential for massive interdisciplinary collaboration across these disciplines, and between projects to implement functional components of poorly understood human functions like consciousness or cognition, where that collaboration has not proved possible before. The proliferation of models of cognition may tend to remain in silos, and their lessons remain unexplored wherever the complexity of translating between them remains too great to permit more than a tiny minority of models to be readily understood by people in different fields. This functional modeling approach enables all work to be combined in a way that has the potential to reliably converge on the functions of a working model.

The Functional State Spaces

The cognitive system is represented here functionally as a system that uses reasoning processes and other cognitive functions to map one set of concepts to another. Each set of concepts is itself a concept, so in this sense all reasoning processes and other cognitive functions map one concept to another. Because all these concepts are represented as being located in a "conceptual space", all reasoning processes and other cognitive functions that map one concept to another concept are represented as transforms that trace a path through this conceptual space. Within the FMF this conceptual space forms the functional state space of the cognitive system. This representation of the state of any system of cognition as a point in a functional state space, and this representation of any reasoning process as a path through that state space, enables the functions of AI solutions to be unambiguously modeled as such paths, so they can potentially be added to a library that a system of artificial general intelligence can use to increase its problem solving ability [15]. It also enables potentially subjective terms like "general problem solving ability" to be unambiguously defined as is done later in this paper.

From a functional perspective, various traditions for observing the functions of the human system in our innate self-awareness, including yoga, have defined the human system as consisting of four sub-systems. These are the body, the emotions, the mind, and the consciousness. Assume we can define a functional state space for each of these sub-systems, where this functional state space uniquely identifies the state of the sub-system [14]. For the cognitive system to have the ability to perceive of each of these states as concepts it must have the ability to detect these functional states. And for the cognitive system to have the ability to direct reasoning to any of these functional states (to conceive of them) it must have the ability to direct perception to them (to find them in conceptual space).

In the FMF the body is defined as moving through one functional state space (the sensory space), the emotional system is defined as moving through a second functional state space (the emotional space), the cognitive system is defined as moving through a third functional state space (the conceptual space), and the consciousness is defined as moving through a fourth functional state space (the awareness space). This paper proposes that in order to understand sensations, emotions, and awarenesses, the cognitive system must also have functional components for detecting perceptions located in those functional state spaces, and that this requires detecting position, sequence, and patterns of perceptions in the sensory, emotional, or conscious awareness state spaces. In this way the cognitive system receives input as understanding in analogy with the way the body receives sensations as input.

The body not only receives sensory input however. It also directs sensory output to the motor functions. By further analogy with the body system, in the FMF the cognitive system not only receives input as understanding, but also directs output of reasoning towards concepts (conception). Within the cognitive system this action is proposed to consist of locating concepts representing states in the sensory, emotional, or awareness state spaces so that the cognitive system may direct conception to concepts of physical states, emotional states, or conscious awarenesses.

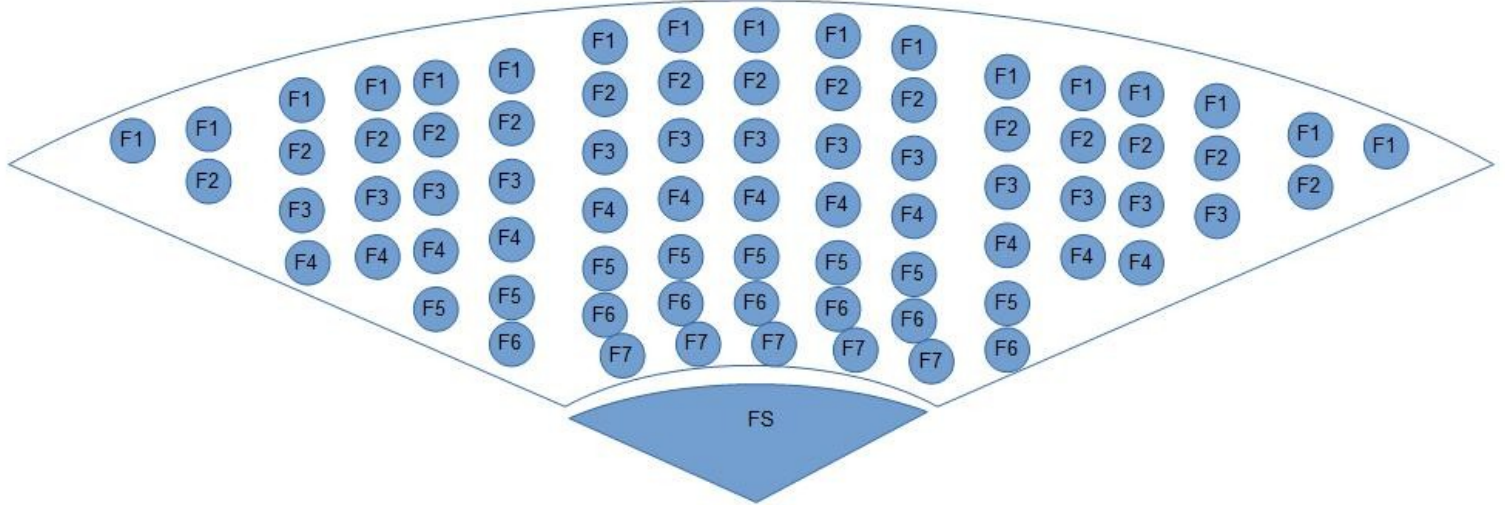
The Components of an AGI in the Functional Modeling Framework

The FMF proposes that the individual mind's cognitive functions consist of a number of functional units that process neural signals into concepts, and a number of functional units that process concepts according to the functions involved in cognition. Three lower order cognitive functions represented by the functional units F1 to F3 map to and from signal space to the conceptual space. And the higher order cognitive functions F4 to F7 and FS consisting of storage (memory), recollection, recognition, etc., receive concepts from the functional state space of the cognitive system ("conceptual space") as input, and produce other concepts as output to that "conceptual space". By executing reasoning processes defined in terms of these functions the cognitive system navigates this conceptual space.

The FMF defines the complete set of cognitive functions that it proposes occur on the input processing path (cognition of sensory or other input). The set of these input cognitive functions are proposed to act to receive understanding in terms of concepts (understanding meaning the process that enables comprehension of the sentence "the quick brown fox jumped over the lazy dog"). On the output path (using cognition to drive reason towards conclusions) these cognitive functions are proposed to direct understanding through reasoning (reasoning meaning the process that enables answering the question "what fox jumped over what dog?").

Functional Units in Systems of Cognition		
Functional Unit	Input Function	Output Function
F1 to F3	Create Concept	Create Signals from Concept
F4	STORE (Store Concept)	DECOMPOSE STORAGE (Determine Concept in Storage Function)
F5	RECALL (Recall Concept)	DECOMPOSE RECALL (Determine Concept in Recall Function)
F6	DETECT PATTERN (Detect Pattern in Concept)	DECOMPOSE PATTERN (Detect Concept in Pattern)
F7	DETECT SEQUENCE (Detect Sequence of Patterns in Concept)	DECOMPOSE SEQUENCE (Detect Concept in Sequence of Patterns)
FS	COGNITIVE AWARENESS (Choose Problem to Solve and Reasoning Process to Solve it)	

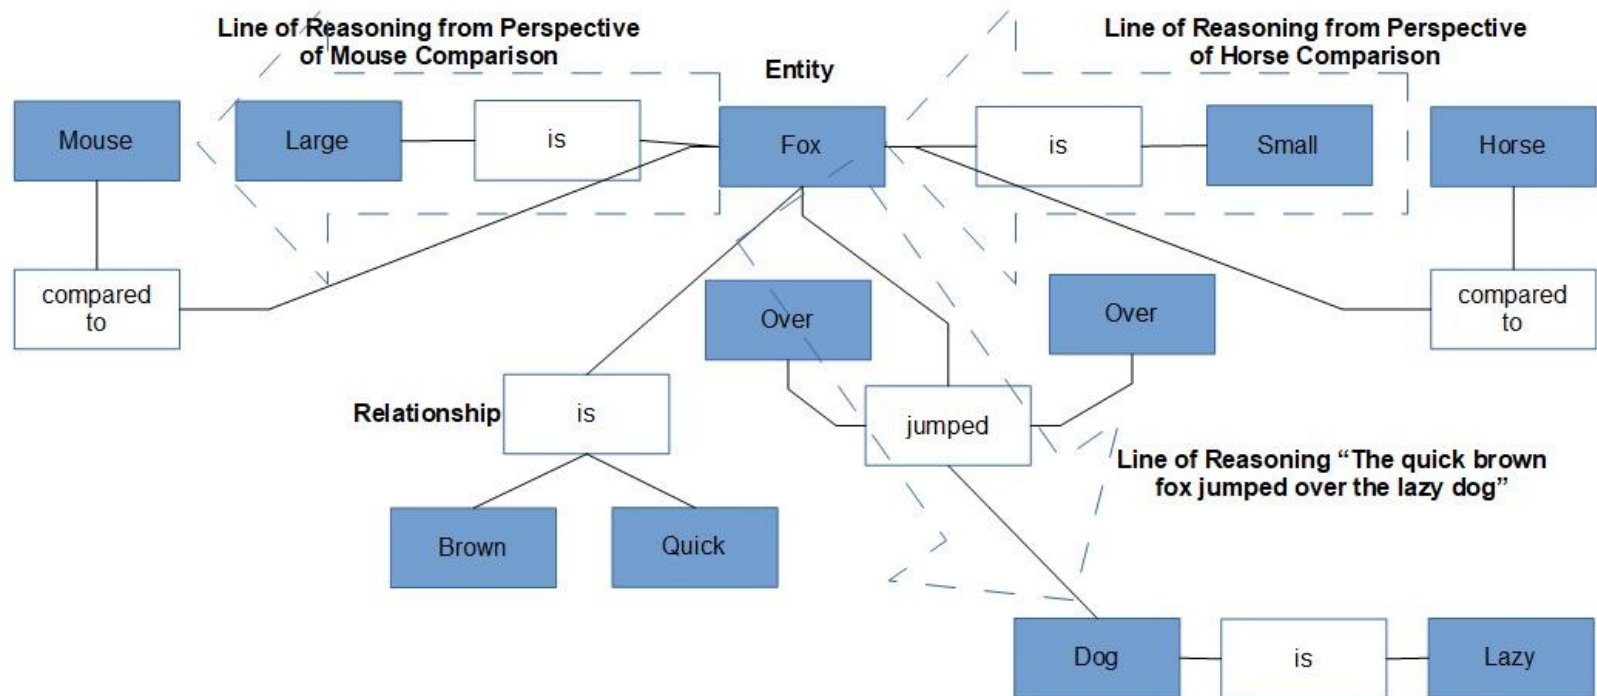
Multiple instances of each of these functions in a conceptual model of the brain that connects these functional units into paths through some mechanism not yet specified, are proposed to have the capacity to represent any intuitive or rational methodical reasoning process.



This paper proposes that representation of reasoning processes in this way is possible because any intuitive reasoning can be represented in a functional model as a form of pattern detection in concepts (F6). And since the set of functions AND, OR, as well as NOT can represent all logic and is therefore Turing complete, this paper proposes that any logic, and therefore the logic in any rational methodical reasoning process, can be represented in a functional model in terms of a function to detect patterns representing a Turing complete set of logical operations on concepts, whether or not those operations are the functions AND, OR, and NOT, and in terms of a sequence of those patterns (F7).

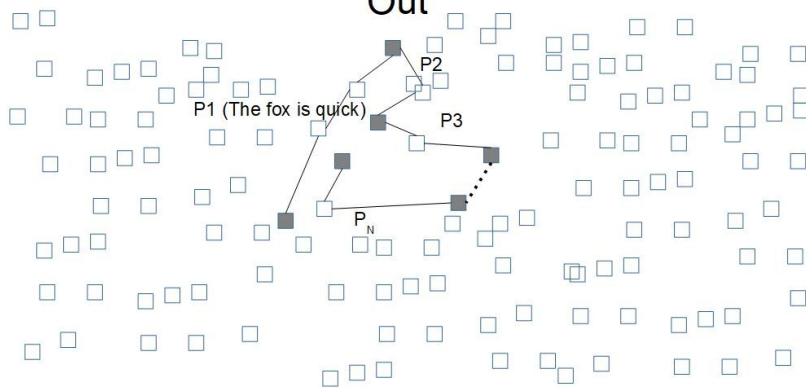
As an example, consider how the following sentence might be represented with the set of cognitive functions and other functional components defined by the framework: "The quick brown fox jumped over the lazy dog". The relationships between concepts from a given perspective are proposed to define the position of concepts in the conceptual space that is defined by the Functional Modeling Framework.

Constructs in Conceptual Space – Zoom In



The reasoning process that produces this natural language sequence can potentially be modeled in this case as beginning at a position on the diagram above representing a given perspective on the entity "fox", and then executing the RECALL function on the properties "quick" and "brown" and the DETECT PATTERN function to associate them with the "fox" to produce "the fox is quick" and "the fox is brown". The process might then execute the DETECT SEQUENCE function to group "quick", "brown", and "fox" into "quick brown fox". The process might then execute the RECALL function on the relationship "jumped". And then might execute the RECALL function on the modifier "over". Finally, it might execute the RECALL function on "lazy dog", and then execute the DETECT SEQUENCE function to group "the quick brown fox", "jumped over", and "the lazy dog". Reasoning processes, such as those required to construct text or speech in natural language, then become a sequence of paths through the conceptual space. In this case, the first path P1 is "the fox is quick".

Constructs in Conceptual Space – Zoom Out



As noted in the first diagram, there are a multitude of relationships connecting the fox to entities that define other of its properties. For example, from the perspective of a comparison with a "mouse" the fox is "large". From the perspective of a comparison with a "horse" the fox is "small". In order to be able to retrieve all the relationships relevant to a given perspective, the representation of the conceptual space must be complete enough to store such perspectives.

Adaptive Processes

As mentioned, the FMF also defines a system of cognition to require a hierarchy of functional components with which it can adapt all other functional components to be more fit (the basic life processes L1 to L9). The FMF defines requirements for the basic life processes and all other components, but leaves models to define implementations. This model of AGI described in this paper includes a hierarchy of adaptive processes to ensure that the most fit component at executing any given required functionality can be taken from models suggested by different researchers, while ensuring the overall model continues to become more fit at representing the functionality of cognition. The adaptive processes provided in the FMF are listed in the table below:

Adaptive Processes (Basic Life Processes)	
Life Process	Name
L1	Homeostasis
L2	Autopoiesis (Reproduction)
L3	Growth (Maximize Outcomes)
L4	Evolve
L5	Act
L6	Cooperation (Between Functional Units)
L7	Learning
L8	Motivation
L9	Reason

General Problem Solving Ability

Reasoning processes are represented in the FMF as tracing a path from one concept in the conceptual space to another. Problems are represented as the gap between one concept in the conceptual space and another. Solutions are represented as the set of reasoning process which bridge that gap. General problem solving ability is then the ability to sustainably navigate the conceptual space so that it is potentially possible to navigate from any problem that can be defined within that conceptual space to any solution that can be formulated within that conceptual space.

Where a non-intelligent system such as current computer programs solves the problem it's designers have chosen for it, a system with general problem solving ability or true intelligence, chooses which problem to solve. The model of cognition described within this paper chooses which problem to solve through maintaining global stability in the dynamics with which it executes all reasoning processes, where that stability exists within a fitness space related to cognitive well-being.

The system of cognition is modeled as projecting the cognitive value minus cost of each activity being executed (its "fitness" in achieving its targeted outcome in terms of cognitive well-being) and either investing resources into the current reasoning activity until complete, or discontinuing the current reasoning activity to invest resources into the next (choosing to solve another problem) in a way that maintains stability in fitness to continue to execute these cognitive functions. The dynamics of the system that are required for stability in fitness space then helps determine the cognitive activity executed.

The fitness space is defined in this model by the projected, actual, and targeted outcomes of the currently executing reasoning process and the impact of those outcomes on cognitive well-being, or the impact of the actual outcomes of selecting a new reasoning process or problem definition, and the projected, and targeted outcomes of all the reasoning process being considered for execution next and the impact of those outcomes on cognitive well-being.

As a reasoning process executes, or as new reasoning processes are considered for execution, their actual impact on outcomes as well as the impact of those outcomes on well-being are updated, and those actual impacts are used to update the projected outcomes and impact of those outcomes on cognitive well-being. If with the continual updates to these projections the actual outcome of the current reasoning process is not projected to reach its target, then other reasoning processes targeting the same outcome, or reasoning processes that choose other problem definitions that target different outcomes, might be selected by the GCI according to their projected impact on the targeted state of cognitive well-being. The targeted outcomes required to keep the cognitive system within a stable range of well-being are also updated as the cognitive system changes its state in functional state space (conceptual space) with the ongoing execution of the reasoning process. The cognitive system then traces a path through this fitness space as the current reasoning process executes and traces a path through conceptual space.

However, when the current reasoning process either completes without achieving its targeted outcome, or when it is not projected to achieve the targeted outcome before exiting the boundary of the stable region in fitness space, then the cognitive system might search for and select the a different reasoning process that now is projected to achieve the targeted outcome, or that redefines the problem in a way that is projected to make the targeted outcome achievable. The cognitive system then continues its path through this fitness space as the next selected reasoning process executes. As it continues its path through this fitness space it must stay within that bounded region of fitness if the cognitive functions are to be stable. It might also follow a potentially chaotic path through the unbounded conceptual space.

In this way, the cognitive system must invest sufficient resources for a given reasoning process to be initiated, but then ranges between investing insufficient resources for the current reasoning process to continue, and investing too many resources to be able to continue any subsequent reasoning in the current reasoning process. This is proposed to form a kind of convection that is reflected in the motion of the cognitive system through fitness space. To implement this model, a system of equations capable of demonstrating this convection throughout a three dimensional fitness space (the Lorenz equations for convection) can then be used to define forces of selection of reasoning processes according to projected, targeted, and actual impact on cognitive well-being so that the path through fitness space might form this stable convection, despite the path through the conceptual space being potentially chaotic.

Having defined the equations governing this relationship, the sequence of reasoning activities selected to be executed by the cognitive process is expected to function in a way that maximizes stability in terms of keeping the cognitive well-being as close as possible to the center of a stable range, and not allowing that stable range to be exited. By executing reasoning activities in a sequence that keeps the state of cognitive well-being within a stable range, the cognitive system is proposed to gain the capacity to navigate the conceptual space as well as to gain the capacity to navigate the state space of the environment. In this way, reasoning in the cognitive system is an adaptive process that enables the entity to find stability in greater regions of the external environment (to understand and reason about the external world).

The parameters of the Lorenz equations can be chosen to form a globally stable dynamics (a strange attractor) in the well-being space, despite a chaotic path through the conceptual space. The same Lorenz equations are used to implement all the other functional components in the model so that their dynamics within their fitness spaces and state spaces obeys the same global stability despite local instability. More detail regarding the application of these equations is provided elsewhere [2], [3].

Conditions for an AGI to be Valid in the FMF

In the FMF the ability to solve a specific problem, such as accomplished by narrow AI, is represented as a path from one concept (the problem) to another concept (the solution). General problem solving ability is the ability to sustainably execute a library of reasoning processes, including reasoning processes that generate new reasoning, so that the cognitive system navigates the conceptual space in a sustainable way that creates the potential to navigate the entire cognitive space. That is, so it is potentially possible to navigate from any one concept (any problem) to any other concept (the solution). In order to be a valid model of AGI, the FMF then requires this global stability in dynamics despite following a potentially locally chaotic path through the conceptual space.

Implementation

In order to implement AGI models the FMF dictates that a number of functional components must be implemented. However having defined these functional components and their requirements, implementations can proceed independently of each other, and in fact may already exist and might just need to be identified. Functional unit F3 performs pattern detection, and since some form of pattern detection is general to all neural networks this has been demonstrated. In the case of position as in F1, sequence detection as in F2, storage as in F4, and the generalization involved in learning as in F7, we can show that each of these functions has been implemented as a neural network (position [5],[6], sequence detection [7],[8], storage [9],[10],[11], and the generalization [12],[13]) and therefore that each mechanism has been explored in an actual implementation.

Conclusions

A model suggested to represent an AGI has been presented. But it may not be clear from the perspective of other approaches why this approach has the potential for general problem solving ability. The answer is that from the point of view of the FMF general problem solving ability is the ability to sustainably navigate the conceptual field with all available reasoning processes from one concept representing any problem, to another concept representing its solution. This sustainability requires stability in some cognitive well-being (cognitive fitness) space, where fitness is defined as the capacity to execute available cognitive functions. Without this stability in fitness a computer program might, for example, merrily go about solving a decryption problem that it projects will take 50,000 years to solve, not taking into consideration that it certainly will not be operational by then, and that no one will be around to care about the solution in any case. With this stability, and with the general problem solving ability it makes possible, the computer program would simply choose to solve a different problem.

Another point to clarify is what exactly is new about this approach, since far more sophisticated neural nets already exist. The difference is that the organization by the same set of functional units enables logic to be constructed the same way for every implementation, so that fitness of every implementation can steadily improve in achieving the functionality required for consciousness and cognition. In addition, with organization by the same set of functional units, reasoning and awareness processes can always be implemented the same way so that the library of such processes can steadily grow.

Finally, the meaning of the adaptive processes in the context of an AGI must be clarified. Do implementations have to implement the adaptive processes such as reproduction themselves? In other words do the functional units in an AGI need to reproduce in this model in order for the AGI to operate? In a word, no. Adaptation from the perspective of the FMF requires a single domain be defined for each adaptive process so every function in that adaptive domain can be changed to become more fit. Without well-defined adaptive domains governing functions through which systems can adapt in each domain, the compatibility of adaptations might quickly be broken.

Finally, to reiterate, by formalizing the process of representing systems in terms of their functions that human beings already do intuitively, this functional modeling approach enables AI researchers to access the vast tradition in which the functions of human cognition have been observed, where these observations have not been readily accessible to the sciences before. This formalization allows processes of observation to work within well-defined state spaces so that processes of observation can be seen as attempting to transmit truthful information, with the result that the ability of processes of observation to reliably transmit truth is governed by well-understood information theory. Where before much of such self-observation had to be discarded as “anecdotal evidence”, this formalism is proposed to make external verification of self-observation reliably achievable [16].

References

- [1] Andy E. Williams, The Functional Modeling Framework (FMF) as a Tool for Modeling Individual and Collective Consciousness and Cognition, under review
- [2] Model for Human, Artificial & Collective Consciousness (Part I), Journal of Consciousness Exploration & Research, Andy E. Williams, Vol 10, No 4 (2019), pages 250-269
- [3] Model for Human, Andy E. Williams, Artificial & Collective Consciousness (Part II), Journal of Consciousness Exploration & Research, Vol 10, No 4 (2019), pages 270-293
- [4] Andy E. Williams, The Semantic Metrics Framework, working paper
- [5] P Bruyndonckx, S Léonard, S Tavernier, Neural network-based position estimators for PET detectors using monolithic LSO blocks, IEEE Transactions on Nuclear Science, 2004 - ieeexplore.ieee.org

- [6] IE Ebong, P Mazumder, CMOS and Memristor-Based Neural Network Design for Position Detection. Proceedings of the IEEE, 2012 – ieeexplore.ieee.org
- [7] I Sutskever, O Vinyals, Sequence to sequence learning with neural networks, QV Le - Advances in Neural Information Processing Systems 27 (NIPS 2014) - papers.nips.cc
- [8] G Houghton, The problem of serial order: A neural network model of sequence learning and recall - Current research in natural language generation, 1990 – dl.acm.org
- [9] S Nara, P Davis, H Totsuji, Memory search using complex dynamics in a recurrent neural network model, Neural Networks, 1993 - Elsevier
- [10] MA Cohen, S Grossberg, Absolute stability of global pattern formation and parallel memory storage by competitive neural networks - IEEE Transactions on Systems, Man, and Cybernetics, 1983 - ieeexplore.ieee.org
- [11] K Yao, B Peng, Y Zhang, D Yu, Spoken language understanding using long short-term memory neural networks, 2014 IEEE Spoken Language Technology Workshop (SLT) –ieeexplore.ieee.org
- [12] IS Aranson, A Pikovsky, NF Rulkov, Advances in Dynamics, Patterns, Cognition, LS Tsimring - 2017 – Springer
- [13] J Sietsma, RJF Dow, Creating artificial neural networks that generalize, Neural networks, 1991 – Elsevier
- [14] Andy E. Williams, Defining the Sensory State Space as a Model for the Consciousness and Cognitive State Spaces, working paper
- [15] Andy E. Williams, Defining Functional Models of Artificial Intelligence Solutions to Create a Library that an Artificial General Intelligence can use to Increase General Problem Solving Ability, under review.
- [16] Andy E. Williams, A Mathematical Model for Identifying Truth in Observations Made within Individual Human Self-Awareness, working paper.