

Artificial Intelligence and Brain Simulation Probes for Interstellar Expeditions

Eray Özkural

Celestial Intellect Cybernetics
celestialintellect.com

November 23, 2018

Abstract

¹ We introduce a mission design for an interstellar expedition to nearby earth-like exoplanets, which our analysis determined to be Tau Ceti and Gliese 667C, at the time of analysis in 2013. We review the research problems in propulsion and AGI that must be addressed to launch an AI guided interstellar probe within 100 years. We propose a new semi-autonomous agent approach for intelligent control of the spacecraft. We introduce the concept of a semi-autonomous agent as having built-in safety guarantees that constrain operation. An autonomous agent case study is presented formulating the objective, constraints, AGI implementation of the agent based on Solomonoff's Alpha architecture, and adding specificity. We discuss the training required to reach human-level and trans-sapient levels of intelligence which corresponds to an entire crew of AI experts specialized in fields such as astrophysics, astromechanics, astrobiology, quantum physics, computer science, molecular biology, and so forth. We project the feasibility of the human-level AI technology based on empirical findings in neuroscience, and find that it should be feasible by 2030. We analyze Solomonoff's infinity point hypothesis in light of Koomey's law about energy efficiency of computing and find that the trends in 2013 indicated an early singularity by 2035, which implies that we might encounter physical bottlenecks which will decelerate computing technology improvements significantly. We recommend thus year 2040 for launching the probe by which date other required technologies will have been developed. We discuss the scenario of a virtual crew made of brain simulations, which is a bio-information based AI approach. We detail the subsystems of command and control, communication, scientific instrumentation, power, propulsion, navigation, and shielding. We propose a variation of ICAN-II/AIMStar propulsion which uses a positron source

¹This paper is an extended version of the original report published at 100 Year Starship Symposium 2013, in Houston, Texas, USA.

instead of anti-protons to initiate micro-fusion reactions. We combine the positron initiated fusion pulse propulsion scheme with a miniaturized version of the Daedalus fusion thruster obtaining high performance. We derive two mission profiles one for fusion pulse propulsion, and the yet hypothetical Q-Thruster. Fusion thruster requires 132.2 years for Tau Ceti, and 233.2 years for Gliese 667C, while Q-Thruster takes only 42.3 years for Tau Ceti, and 62.5 years for Gliese 667C. We also discuss extended roles for intelligent interstellar probes such as self-reproduction via nanotechnology, refueling, construction, robotic bodies, and transmission of brain simulations.

1 Introduction

We are witnessing the dawn of a technological society, in which we will not be merely surrounded by technology, but the very members of the society will have technological components or be comprised of technology altogether. We have ascertained early on that interstellar travel shall be a primary application of human-level AI technology, along with such intelligence requiring applications as unification of general relativity and quantum mechanics, and curing cancer. The renewed interest in interstellar research confirms that boundless amounts of creativity are required to feasibly achieve interstellar flight in a short time. We anticipate that human-level AI technology will be immensely useful both for accelerating interstellar research and for controlling the spacecraft. The present study addresses the latter problem.

We present a high-level design for an interstellar probe built around artificial general intelligence and brain simulation technologies. Due to the communication delay in interstellar missions, an intelligent control system is required for navigation, scientific tasks and emergencies. While biological astronauts fit the bill, the astronauts and life support systems as well as crew space and tools with human interfaces increase the mission mass significantly. Moore's law suggests that around 2025, a laptop computer will have enough processing speed to simulate a human brain at real time. Therefore, fast and small computers that can operate human-level artificial general intelligence agents or brain simulations may be used for interstellar missions in the near future. We outline a one-way mission design for expeditions to Tau Ceti and Gliese 667C systems which house earth-like planets with a hypothetical timeline beginning in 2040. We present a design sketch for the probe, detailing its control, scientific instrumentation, communication, propulsion and shielding components. We detail how a goal-following or reinforcement-learning agent can be designed which will navigate the probe safely to a destination star system, survey the planets, perform experiments and establish communications autonomously. We discuss how the intelligent agent can be trained with sufficient scientific knowledge, and mission simulations. The brain simulation option allows us to create a virtual crew which may rank in hundreds; we detail scenarios of how the virtual crew can be effectively employed inside the interstellar probe. We show the computing and energy costs for such artificial crews. The secondary significant cost would be the mass

and energy requirements for establishing an interstellar communication link at the destination star system, as well as any scientific equipment. We assume laser based communication for two-way interstellar communication and calculate the costs. We base the scientific instrumentation costs based on existing interplanetary probes, and take into account fuel requirements to survey the entire system. We consider the following extended roles for interstellar probes. We consider whether the probe can refuel or manufacture any equipment such as habitats and robots at the star system using future nanotechnology assembly methods, and we also discuss replicating probes. We consider the cost of carrying a lander module for exploring earth-like planets. For propulsion, we assume advanced propulsion systems such as fusion pulse propulsion to reach a significant percentage of light speed, discuss shielding and navigation requirements for ultra-high speed probes, and discuss the benefits of artificial/virtual crews over human crews in terms of fuel, mass and construction regarding our hypothetical mission.

The recurring themes in our proposal are miniaturization, efficiency and autonomy. A design which optimizes these features is especially fitting for an interstellar probe, whereas miniaturization allows us to use as little propellant as possible, and efficiency lets us get by with little energy and with maximum speed, and autonomy is the cornerstone of our proposal, eventually allowing trans-sapient levels of adaptability and error-recovery to be achieved for spacecraft control.

2 Background

2.1 Target selection

Name	Distance (ly)	Planets	Planets in HZ	Max. ESI	Confirmed
Gliese 667C [1]	22	6 (7)	3	0.82	Yes
Gliese 581	20.2	5	1	0.81	Yes
Tau Ceti [43]	11.9	5	2	0.77	No
Gliese 163	48.8	3	1	0.74	Yes
HD 40307	41.7	3	1	0.72	No

Table 1: Promising nearby star systems with earth-like planets.

Recent surveys of extrasolar planets have revealed 743 star systems which contain 973 planets, 504 of which are closer than 100 light-years [45]. Most habitable exoplanets are expected to be found around M-class stars in recent studies that refines the concept of habitable zone; [18] revises the habitable zone according to atmospheric data, and [3] develops an empirical habitable zone definition which suggests prevalence of M stars with habitable planets (36.5%) followed by K, A, G, F in order of decreasing percentage. Some of the most promising nearby planetary systems are displayed in Table 2.1, the columns of which indicate the name, the distance from solar system in light-years, the total number of planets,

number of planets in the habitable zone, and maximum Earth Similarity Index (ESI), taken from the as of current list in Habitable Exoplanets Catalog [20] with high ESI values which measures the similarity of the planet to Earth incorporating many parameters [5]. Most of the exoplanets discovered in the habitable zone are either gas giants, or super-Earths, i.e., rocky planets that are more massive than Earth. All the habitable planets in are super-Earths, with masses less than 10 Earths, and they have greater ESI values than Mars. Extra-terrestrial life depends on more than Earth likeness naturally, such as stability, surface temperature variations, and chemical composition. Some may even turn out to be gaseous and have habitable moons. Also notable is the idea that astrobiology may be fundamentally different from ours, alleviating the need for similarity to Earth. However, for colonization purposes, a similar planet would be most convenient.

For our intelligent probe, we have chosen Tau Ceti and Gliese 667C as the target systems among the candidate star systems in , due to the close proximity of the former, and the high ESI values of the latter's three planets. Although the habitable super-Earths in Tau Ceti are yet unconfirmed, there is reasonable speculation in support. Tau Ceti e has a high ESI value, suggesting its atmosphere is similar to Earth. Although Tau Ceti f is said to be a cold planet, due to greenhouse effect, its surface temperature may vary, providing liquid water for complex life. Gliese 667C is naturally an attractive target as it is both close enough for a 21st century mission, and has three habitable planets. There is insufficient data about the habitability of these planets yet, however, there is ample time for future observations.

2.2 The challenges of propulsion

It is common knowledge that chemical rockets do not have enough energy density to support interstellar travel within reasonable time frames, which suggests using nuclear and anti-matter rockets, and alternative propulsion concepts [24]. The first application of nuclear power to interstellar propulsion was the Orion project [8], which used nuclear pulse propulsion by periodically detonating small nuclear explosives (fusion-enhanced fission reactions) against a pusher plate. Orion project was abandoned due to Partial Test Ban Treaty and issues of nuclear fallout.

Numerous credible design studies succeeded Orion, some of which replaced fission with smaller yield fusion reactions, as spacecraft design is limited by the size of smallest explosive package possible. The most comprehensive of such studies to date is the Daedalus project of the British Interplanetary Society which used $D-He_3$ pellets for electron beam driven inertial confinement fusion, that directly transformed fusion products into thrust via a magnetic pusher plate [13]. Daedalus propulsion combines kinetic shock absorption with magnetic impulse from a superconductive shell. Each pellet of the Daedalus propulsion system is a miniature explosive package composed of a $D-T$ trigger and $D-He_3$ fuel encapsulated in a superconductive shell, which after magnetic injection into a hemispherical reaction chamber, burns into fusion and yields thrust, repeatedly detonated by the relativistic electron beam at 250 Hz.

The primary disadvantage of fusion pulse propulsion is the inordinate mass of the driver, cooling, conductance and plasma containment systems. The Daedalus mission weight for Barnard's Star was about 53000 tonnes which may present orbital construction challenges in the near-future. Project Icarus proposed replacing the electron beam driver with laser and decreasing the nuclear pulse to 150 Hz with smaller pellets [25]. The fusion-driven rocket partially solves the problem by using metal liners to both contain the plasma and conversion to thrust [34]. Light-sails work without any propellant, and laser-driven light-sails can reach up to 0.1c easily with around half GW power, making light-sails an ideal candidate for fly-by missions [21]. Anti-matter ignited fusion pulse propulsion proposals solve the mass problem by substituting the lasers with antiprotons. Antiproton annihilation ignites a small fission reaction which then powers fusion in the pellets. Anti-matter production and storage present research challenges of its own, however. Current research indicates that positron annihilation driven rockets is a feasible concept, as pursued by Positronics Research LLC. Simultaneously, groups like Positron Dynamics LLC have been reporting advances in positron storage. Likewise, small cylindrical inertial electrostatic confinement fusion devices pursued might turn out to be practical [28]. McGuire's follow-up work at Lockheed aims supporting 100MW reactors in 2 metric cubes. Such a small fusion reactors can henceforth power efficient plasma rockets with high specific impulse, partially solving the dreaded mass problem with fusion rockets.

2.3 Artificial General Intelligence

Artificial General Intelligence (AGI) field seeks to build a human-level, general-purpose Artificial Intelligence (AI) system, which may be defined as a computer system that can solve any problem that humans can. The most fundamental feature of human intelligence is its apparent universal character, that it can in principle solve any problem. Furthermore, it has been observed that the problems of intelligence are essentially *prediction* problems, i.e., the most basic intelligence problem is that of extrapolating from data, which corresponds to inductive inference in scientific inquiry. That is to say, the task of formulating a general law from given observations forms the core of our intelligence. As such, Einstein's theory of general relativity is a fine accomplishment of induction, as well any scientific discovery. Ray Solomonoff's universal induction theory [37] addresses the universal inductive inference problem, based on a quite benign assumption: that the probability distribution of the universe is computable. This assumption apparently holds for the observable universe due to Bekenstein bound, since the observable universe has finite entropy, and thus has a computable probability distribution, even if quantum events are indeed random. When this assumption is satisfied, it has been proven that predictions have very small error dependent only on the universal computer chosen to represent information, and not the data. The predictions converge very rapidly (convergence theorem). Solomonoff's induction method is incomputable, however, it may be approximated arbitrarily, and many practical universal induction algorithms already exist. Solomonoff's theory has unveiled a beautiful concept, which is the probability of an object

to be produced by a random program. The probability of bitstring x , which may be anything on a computer, is known as universal distribution and it is defined as:

$$P(x) = \sum_{U(\pi_i)=x^*} 2^{-|\pi_i|} \quad (1)$$

where U is a universal computer (a computer that can interpret a programming language), x^* is x and any extension of x , and $|\pi_i|$ is the length of program π_i that generates x as a prefix. In application to prediction, it can solve arbitrary sequence, set, and operator induction problems [40]. Solomonoff's induction method fully formalizes Occam's razor, the distribution explicitly shows that probability of a program decreases exponentially with increasing program length, yet it allows multiple explanations at once (Epicurus's law of multiple explanations). The benefit of the universal distribution may be most clearly understood in the context of Bayesian inference. In Bayesian inference, if there is little data, the inference does not work well. There is also the problem of often not knowing the a priori probability distribution. The universal distribution acts as a universal prior for any application of Bayes theorem, alleviating these major problems. In our experiments, we have seen that universal induction works extremely well with small amounts of data.

Let us try to give a better sense of how an inductive inference engine works. Given a set $\{(x_1, y_1), (x_2, y_2), \dots\}$ of input, output pairs, operator induction can learn the generalized conditional pdf $P(Y|X)$, and it allows us to predict $\operatorname{argmax}_Y P(Y|X)$. Operator induction solves any classical ML problem; x_i 's may be faces and y_i 's names, and it is perfectly general. It has been argued in the literature that we can apply induction to solve any AI problem [16].

Some approximation methods exist for Solomonoff's universal inductive inference method, which are basically Levin's universal search [22] and evolutionary programming (which has also produced our intelligence). The major application of inductive inference is building a cognitive architecture that uses inductive inference module, an early example of which is the OpenCog system [11].

Some forthcoming challenges of AGI research are

1. Choosing the right universal computer (variants of LISP, FORTH, machine language, MATLAB, etc.)
2. Designing an efficient approximation algorithm as the search space is exponential in the number of bits and parallelism is required
3. Memory: designing a transfer learning method and making memory work with induction
4. Modularity and scalability: making a scalable cognitive architecture that can scale up to human-level complex problems, comprised of many kinds and instances of cognitive modules

Examples of AGI systems we may recount here are

- Alpha [39]: assimilates narrow-AI and universal AI systems, it is modular, it has preprogrammed higher-order cognitive procedures that perform analysis and synthesis, and it can solve free-form time-limited optimization problems

- Gödel Machine [35]: Self-reflective AI agent design, which can use and improve any reaction policy
- AIXI [15]: extends universal sequence induction and it is an optimal reinforcement learning agent model which uses input, actions, rewards to make an optimal plan, and it can optimize expected future rewards.
- Teramachine [32]: a universal inductive inference engine that partially solves the memory and parallelism problems.

3 Semi-autonomous Agent for Intelligent Control

3.1 Rationale

Intelligent probes are required mainly due to communication latency with interstellar spacecraft, however, there are other factors which may be as significant in practice. There are multiple intelligence-requiring tasks onboard an interstellar spacecraft, the most immediate of which are:

1. Navigation: Neither is the interstellar environment void of hazards, nor do we know much about it. While future missions plan to map the heliopause, there will still be much that is unknown in the interstellar space, especially nearer the target system. In addition to this, the voyage shall take several years, which implies variability in system performance and operation, all of which point out that an intelligent navigation system would be most useful.
2. Scientific tasks: Without an adequate scientific payload and an array of informative scientific tasks, an interstellar expedition would merely be an engineering demonstration. Our lack of knowledge about the host system, however, implies much adaptivity in the execution of the scientific tasks, which casts doubt on pre-programmed scientific instruments. An AI system may also perform open-ended queries on a target system, employing as much creativity as a human, in finding answers to questions such as the presence of exobiology in the system.
3. Emergencies and fault-tolerance: Due to mission longevity requirements, we may expect many emergencies and faults during the mission. Intelligent recovery from such erroneous conditions will improve the robustness of the system. An AI system may also respond to such emergencies much quicker than any human can, and it may even be able to resolve emergencies which were not anticipated.

An AI control module thus affords several advantages that would be unwise to neglect. It first eliminates significant mass for the habitat module and extra shielding, while retaining every advantage of a human crew, a point which was also well taken by the Daedalus Project. It offers improved longevity, as AI may do well in conserving resources, and carrying out mission tasks as efficiently as possible. In contrast to a human crew, AI may have better rationality, as we may program it to be free of human

emotions in its decision making. AI may have improved reaction times to all manners of hazards and faults during the voyage, and this might matter much, as for instance, one may imagine an unanticipated hazard in the fusion propulsion system which may be prevented with the right command within a matter of microseconds. A successful AI control system may also have superior error recovery abilities, as it may have command of every engineering method ever devised. For instance, whereas a human may start writing a MATLAB program, an AI may already have the answer from an obscure engineering journal and have solved an optimization problem prior to the accident in case any error of the variety might happen. Possibly, the most interesting point of an AI system however is increased intelligence, as the AI might be able to solve error recovery problems that would take a large community of human scientists to solve.

3.2 Agent Designs

An AI agent is an abstraction of an intelligent animal, which acts in an environment in an intelligent manner. Two common agent designs are relevant to interstellar probes: the goal-following agent, and the utility-maximizing agent. Both agent designs interpret sensor signals to infer the state of the world, and manage actuators intelligently so as to influence the environment. Both agents know how their actions will influence the world state. The goal-following agent chooses present (and future) actions in order to satisfy pre-determined goals. Likewise, the utility-maximizing agent chooses in actions that will maximize its (future, cumulative, or expected, or defined otherwise plausibly) utility. The latter kind of agent has received an unusually intense treatment in the Artificial General Intelligence community, the best known example of which is the AIXI reinforcement-learning agent which solves the problem in a universal setting [15]. A reinforcement-learning agent assigns positive or negative reward signals, much like an animal does, to solve the utility maximization problem. We shall show that both kinds of agent designs are fitting for the AI probe.

3.3 Agent sensors and actuators

Sensory input corresponds to spacecraft's sensor instruments and computer's instruments, such as its ability to read its RAM. The sensory input may be provided to AI in many ways. It may be provided as raw data to the AI, which then has to run appropriate perceptual processing algorithms to extract relevant information, which requires the highest level of intelligence and learning rate for the agent. It may be supplied in the form of calculations over raw data, such as location of pulsars from the X-ray telescope, or other relevant astrophysics calculations such as a power spectrum, which would accelerate perceptual processing. One may also run perceptual-specific narrow AI's and astrophysics calculations to directly supply the AI with such data structures like an astronomical map, much like the data structures autonomous driving programs work with.

Actuators are the outputs of the control system, and they correspond to every device which manipulates the environment and the spacecraft

itself, such as the thrusters. Similar to sensors, there are many ways to deal with actuators; the AI can either drive the control signals of every subsystem directly, or it can issue high-level computer commands to actuator subsystems, for instance it can send the command “FUSION THRUSTER THROTTLE 40” to the propulsion subsystem. The distributed design seems more manageable and akin to large-scale project development.

3.4 Primary mission objectives

To better comprehend the problems that the agent will have to solve, let us enlist the most essential mission objectives in temporal order. Recall that for the AI to be specified in the right manner, they ought to be expressed in the most general manner of conceptualization to allow it freedom.

1. Navigate safely to the target star system: Travel to the host star system in the shortest time possible, using as few resources as possible, and incurring as little damage as possible, maintaining as many subsystems as fully operational.
2. Survey the planets: Chart the star system by determining the orbits and astronomical properties of all bodies with significant mass, then determine a short list of planets which may embody life, and then travel to the orbits of these planets and map the planets’ topography and run the sensory instruments on the planets such that resources are not depleted for the mission duration.
3. Perform experiments: Carry out intelligent sub-plans corresponding to each scientific experiment in order of decreasing priority.
4. Establish communications: Determine the exact location of a FOCAL probe, deploy an inflatable RF antenna, and establish communication with the FOCAL probe such that bandwidth is maximized. Transmit as much relevant mission information as possible, in order of decreasing priority to the FOCAL probe, such that energy is preserved during transmission.

Note that any/all of these objectives may be carried out concurrently.

3.5 Goal specification for goal-following agents

The goal-following agent may work with explicit statement of goal-states in a logical language. For instance, it may have a form such as:

`DistanceLessThan(CurrentPos, TauCeti, 1-AU).`

It may have natural language commands, as in OpenCog Bot [12], which can answer natural language queries, and accept natural language commands. Thus, the above goal may be stated as “Travel to Tau Ceti system and stop around an orbit of 1 AU distance to Tau Ceti star”, which the agent would first translate to an internal knowledge representation and carry out a plan which corresponds to the commands. The goals may also be expressed in a logical natural language such as LojBan [10]. Thus, a

succession of goal-states may be either pre-programmed or transmitted to the spacecraft on-the-fly.

Another means to implement a goal-following agent would be to express the goal as a free-form optimization problem. For instance, the above goal may be expressed as an optimization problem of the following form.

$$\begin{aligned} & \min.d(\text{current}, \text{tauceti}) - 1AU \\ & s.t. \text{spent-fuel}(t_{end}) < \text{total-fuel} \end{aligned}$$

In the context of AI probe, the free-form optimizer must be general enough to account for multiple objectives and multiple constraints, and must be adaptable to real-time operation, wherein the values of input variables may rapidly change.

3.6 Reinforcement learning agent

Discrete goals, such as the arrival to a star system of the preceding examples, may be expressed as 0/1 rewards to the reinforcement-learning (RL) agent. Numerical objectives may be mapped to utilities. For instance, a reward may consider the percentage of a planet's surface that has been mapped by the probe.

For reinforcement-learning agents, an important problem may be the question of how to combine the various rewards, as this may create unforeseen strange behaviors in bizarre, unanticipated circumstances. A simple method is a linear combination of rewards, however, coefficients must be then tried out in simulation to see if whether this results in stable behavior. Other more sophisticated approaches may use reward vectors, other possible reward structures (a tree may denote reward priorities) or statistical combination of various rewards. Yet more sophisticated approaches may have an optimizer that considers multiple rewards.

3.7 Autonomy

An important trade-off for the intelligent probe is specific vs. unspecific goals. A specific goal would contain numerically precise model states. For instance, as relevant to our goal example, it might be "Orbit around target star at 1.0 AU +/- 10m". Such numerically precise goals run risk of failure, due to tight constraints. In some circumstances, it might be too difficult or costly to achieve them, leading the spacecraft to sub-optimal behavior. On the other hand, unspecific goals allow the probe the freedom to plan. For instance, we may instruct the probe to gather as much information as possible about the system. We call this kind of maximally general objective a "universal goal". Universal goals are likely more appropriate for generally intelligent agents, as they exploit the unbounded creativity of these agents.

Another relevant trade-off is fully autonomous vs. semi-autonomous intelligent agent. It is well known that fully autonomous agents may diverge and have drives of their own [30], which may well be detrimental to (real) mission goals. We propose a simple solution to the problem of AI drives:

1. One must use physical constraints. There must be interpreted sensor inputs which account for basic physical measurements, such as location, velocity, energy spent, and so forth. Thereafter, the agent may be instructed to respect the physical constraints, such as a certain time-space bound is respected, and only a certain amount of energy is spent.
2. One must specify the goals generally, but at the same time sufficiently precisely. This is achieved by striking the right balance between specificity and unspecificity, as discussed before. The goal must be precise enough such that the agent *knows* what we wish it to achieve, and at the same time it must be imprecise enough such that the agent may conceive as many ways as possible to satisfy the goal.
3. One must avoid open-ended goals. Open-ended goal roughly means an optimization problem with an unbounded objective. For instance, maximizing the lifetime of the agent is such an objective. Plugging such objectives into the agent may interfere with other essential tasks. For instance, the survivalist agent given as an example would likely try to conserve resources and avoid making any scientific experiments eventually (as much the utility combination allows).
4. One may use priorities suitably. The use of weights/priorities corresponding to probe goals may be useful, in that, the agent knows which goal is the most important. For instance, arriving at the target system may be more important for a mission than achieving a stable orbit. In that case, the probe may decide to switch to a fly-by mission, since a propulsion component was damaged beyond repair.

4 Autonomous Agent Example

A quite unspecific agent may be expressed as an optimization problem of the following form.

4.1 Objective

The objective is to maximize amount of (true/verified) information discovered about the target system and transmitted back to Earth. The amount of information may be expressed as the expected cumulative information for the entire mission duration, thus it encompasses all future actions, resulting in an integrative plan. The amount of information may be calculated with respect to all questions that may be asked of the system, which may be formalized as improving the prediction accuracy for any induction problem that contains the system as a variable. The questions thus may be generated in order of increasing a priori probability of the question, assigning much more weight to simpler/shorter questions, and they may be weighed accordingly. Thus, it is more important for the probe to answer a question whether there are any rocky planets, than the number of mountains on a particular planet that have an elevation of more than 12000 meters. The list of these potential questions would

be examined before mission, however, the importance of this approach is that as the probe increases its knowledge, the order of questions change, giving it an intelligent inquisitive character.

Note that an RL formulation is also possible. In this case, the idea is loosely equivalent to the Bayesian knowledge-seeking agent [31], which can gather as much information about an unknown environment as possible.

4.2 Constraints

As the objective is general, so must the constraints.

1. Spacecraft resources are not depleted within maximum mission duration
2. Spacecraft's critical subsystems remain operational for specified period
3. Spacecraft remains within given space-time region and energy limits.
4. Generalized non-interference clause: spacecraft does not interfere with any intelligent agents encountered.

4.3 Advanced Formulation of AI Probe

An extremely advanced formulation of the optimizing/goal-following agents is possible due to Solomonoff's general-purpose AI system Alpha, stage 2 [39] which can solve time-limited free-form optimization problems. Alpha is of particular interest to us because it is the culmination of our research program, as teramachine was conceived as a candidate for Alpha, stage 1 [32]. Shortly, merely collecting as much information as possible is not enough. One must also optimize the rate of information extraction, which Solomonoff briefly explains in [39]. Also, real-time operation is of paramount importance for a spacecraft.

The advanced formulation optimizes the main objective function within a time-limit we call a planning quanta, which may be set as 1 millisecond. Thus, an action plan is continually updated. The plan optimizes the efficiency of knowledge extraction, in terms of bits/J.sec, that is, it optimizes the number of bits extracted per unit of energy spent within unit time. For this approach to work, Alpha must be made fully incremental, so that its batch processing approach can be converted to real-time. Another approach is to divide the plan into a top-level plan which is updated slowly and a reactive plan which is updated rapidly. Such hierarchy allows separation of reflective thinking and reactive thinking as in the Gödel Machine [35], and as many levels as useful may be created for various temporal requirements.

4.4 Adding Specificity

We may modify the objective such that problems are restricted to astrophysics and astrobiology domains. This may be achieved by conditioning the problem generator on the corpuses that correspond to these domains.

We may also modify the objective such that prediction accuracy in answering given questions increase. Following are examples of such questions:

- Is there life in the host system?
- Is there intelligent life in the system?
- Are there technological artifacts in the system?
- What are the astronomical properties of bodies in the system?
- What are the chemical compositions of the bodies?
- What are the topographies of the bodies?

which are examples of questions we are truly curious about. It must not discourage the reader that these are most general questions. General approaches to solution of some of these questions are already known. For instance, in the absence of any evolution, we expect the proportions of chemicals in a planet to be much more uniform. Therefore, looking for any sort of skewed distribution of elements which may not be easily explained by usual geological interactions would be a general method to detect life. It is expected that the human-level AI knows and *understands* all such general methods sufficiently well, i.e., as well as any human scientist, therefore it should not be prohibitive to ask these questions in the simplest fashion possible. Furthermore, the set of pre-specified questions may be weighed, causing the probe to spend more energy on questions we are more wondrous about.

The unspecific agent in general may be much more flexible. Assume that the second question in the preceding text was answered positively, and an RF-capable civilization was detected on a planet. The AI could learn to communicate with the extra-terrestrial intelligence and directly ask them the questions it could not answer itself, translate them, and send us its findings. On the other hand, the specific agent of the sort outlined here may be more efficient due to the exploitation vs. exploration trade-off; note that there are many suitable proposals for information acquisition that may handle this trade-off automatically [41]. The correct application of curiosity in an AGI controlled interstellar probe seems, therefore, a matter of importance, and thus general theories of curiosity are relevant [36, 42]. The probe must be curious, but must not be so curious as to diverge to random walks.

5 Training

The training is likely to start in 2030 using supercomputers, and should be regarded as the most expensive part of an AGI application. Since by 2030, we assume human-level energy efficiency will be achieved, training the AI should not take any more than training a human by that time, in fact, it should be much faster, since we will be using faster computers than human.

5.1 Modes of intelligence

Various modes of intelligence may be sought for the AI agent. Following are some scenarios from the simplest mode of intelligence to the most sophisticated.

- A high-level gating system that acts like an animal. Assume that there are several fixed programs that solve various subsystem functions in the probe. One way to visualize this is to consider all the sections in a Mission Control Room. Each section in mission control may be automated by a narrow AI system of its own, optimizing the behavior of given subsystem, using many fixed programs supplied to the probe, which perform the complex calculations required for orbit insertion and so forth. The various Mission Control AI's provide a narrow but useful degree of adaptivity to each command section. While, on the other hand, the general-purpose AI acts as a gating system, which knows enough to give the proper commands to the mission control AI's, but does not know in detail on its own how to perform them, i.e., it acts like a manager, or a mission directorate.
- A human-level AI that has expertise in related fields. The AI agent is trained, much like a university student, so that it is proficient enough to obtain a university degree in the relevant fields, for instance, astronautics, astrophysics, astrobiology, terrestrial sciences. The AI henceforth is a good scholar, and it can use applied knowledge it learned during its training to perform the required mission objectives.
- A self-reflective, self-improving trans-sapient agent (artificial star-fleet captain). This AI has accumulated a good portion of the knowledge generated by humans, and can process this information and invent new solutions at a much higher rate than human. It has demonstrated ability to learn new languages quickly and assimilate and apply new fields of inquiry as it has needed.

5.2 Background knowledge and testing

Standard material in any human-level AI encompasses:

- Training in pattern recognition problems (visual, audio, etc.)
- Extensive training in math, physics, English
- Includes much of high-school/college curriculum
- Capability to learn and apply human knowledge assessed (can pass SAT/GRE)

We estimate that the training of the standard human-level AI will take at most \$100 M and 5 years once the algorithms are stabilized. Thus, if in 2025 we have the first human-level AI program, then in 2030 we may train it to a level where it can comprehend arbitrary text.

Interstellar mission knowledge would be required in addition:

- Expert-level training in physics, chemistry, astrophysics, astronomy, astrobiology

- Spacecraft specific training for use and control of every subsystem
- Mission training involving hand-crafted and random mission simulations
- Mission readiness determined by an Interstellar Turing Test

In 10 years, with an additional budget of \$200 M, expert level training and realistic mission simulations may be achieved, in time for mission launch in 2040.

6 Human-Level AI Projection

We calculate a new upper bound for relevant human brain computational capacity to determine a worst-case estimation for when human-level AI will be feasible. To our knowledge, this figure was not achieved in this manner before. Number of neurons in the brain is about 10^{11} , and the number of synapses in the neocortex, the region which is thought to include higher-level cognition, is about 1.64×10^{14} , while total number of synapses is less than 5×10^{14} . Synaptic bandwidth is about 1500 bits/sec maximum [6]. Therefore, cortical speed is approximately 2.5×10^{17} bits/sec, which is about 3.8 petaflop/sec. Contrast with Moravec's estimate, which was 100 Teraflop/sec [29]; a large gap which suggests that Moravec was optimistic.

Energy efficiency of cortical computation is thus 3.8 petaflop/sec at 20W, which is 192 Teraflop/sec.W. Note that tighter estimates are possible, but we are merely making a rough worst-case estimate. While it seems that Moore's law in its original form has been trumped, Koomey's law suggests that energy efficiency of computing doubles every 18 months, and it has been observed to be more stable than Moore's law. Our prior estimate using Moravec's optimistic estimate of human-level energy efficiency was 2026 [33] based on NVIDIA chips' energy efficiency. We now base our calculation on Adapteva's multi-core chips, which obtain 32 GigaFlops for \$100 currently, at 72 GigaFlops/Watt energy efficiency for their 64-core CPU. When we extrapolate using Koomey's law for upper bound, we find that in 17 years human-level energy efficiency would be achieved, which sets our worst-case prediction as 2030. By then, we anticipate that human-level AI will already be available, and the interstellar probe will merely be an application.

6.1 Infinity Point

Infinity Point is a striking consequence of Solomonoff's mathematical analysis of social effects of AI [38], and it is also known as singularity by authors who popularized this discovery consequent to Solomonoff. Solomonoff has observed that human-level AI accelerates Moore's law, as AI may run on hardware to which Moore's law is applicable, and it can accelerate computer hardware research, resulting in a positive feedback loop, improving the already exponential improvement of Moore's law. As such, it is a macro-economic hypothesis about proliferation of human-level AI technology.

Moore's law may be summarized as "number of transistors placed on a microprocessor at a fixed cost doubles every two years". Current doubling time is about three years, which means that Moore's law has already tapered off. However, let us first present the theory and show how it will be salvaged. According to the theory, we have two assumptions:

- Computer Science community size \sim rate of improvement
- Computer Science community size \sim rate of logarithm of computing efficiency

That is, the larger the community, the smaller the doubling time gets. Solomonoff's macro-economical theory suggests that if fixed amount of money is invested in AI every year, we thus obtain infinite improvement in finite time, which he describes as an impossible mathematical singularity.

In practice we have only finite improvement, we improve from 600 years from unaltered Moore's law to a few decades. However, Moore's law is already slowing down so it is inapplicable. We can replace Moore's law with Koomey's law and thus still predict an infinity point. By 2030, it will be a negligible amount of money to augment each computer with an equally efficient AI, which sets $R = 1$ in Solomonoff's theory, his calculations show the infinity point to arrive in less than 5 years with this rate of investment [38]. Therefore, infinity point is expected to occur by 2035 *in the worst case*, which turns out to be quite surprising. Note that this estimate rests on the assumptions stated above, and analysis of them is part of ongoing research.

6.2 Computer Technology by 2040

We extrapolate energy efficiency of computation by 2040, which is the launch date of our hypothetical mission. Without infinity point, we have 195 Petaflop/sec.W efficiency (101.5 x human); one brain simulation costs about 0.2W. With infinity point, we assume human-level AI by 2035. At 20W per brain simulation, we double CS community every year. This incurs a negligible cost with respect to global economy by then. Double-exponential improvement in energy efficiency occurs extremely rapidly, and infinity point then is reached in 4.62 years, which implies that computational resources will be abundant by 2040, and efficiency will be much closer to physical limits than present. For our design, however, we do not assume infinity point, and thus our projections may be interpreted as a worst-case scenario.

7 Brain Simulations for Virtual Crew

Brain scanning technology improves very rapidly. Human Brain Project will complete by 2025 and it aims to build the first complete brain simulation by then. By 2040, 20W can simulate 100 virtual crew, and thus we can build a particular spaceship of the imagination that houses a virtual crew and take our best minds on a trip across the stars.

As Tipler and others have suggested brain simulations (sims) may interface with body images in VR and they can occupy a virtual bridge

and a virtual city. Sims can carry a copy of the internet and all libraries with them and thus have access to vast knowledge resources.

Even more interestingly, we can build a trans-sapient artificial astrophysicist which will possess an extended neocortex in simulation, that will be able to directly perceive astrophysics modalities, and can invoke computer tools with no delay, allowing it to drive the ship's controls as it were its appendages.

Sims will be trained before voyage on virtual missions, allowing us to explore every mission variation and hazard. As with the AI, sims may use a sleep mode to conserve energy, and run at different speeds to accommodate various spaceship tasks. During the voyage, sims may continue collecting and interpreting data, and conduct scientific work, as well as preparing a rich log of the voyage for analyzing back at Earth. Sims may also engage in all manners of recreational and creative activities, as well as inventing and producing new technology. A major advantage of sims is that they have human experience and thus might more directly apprehend what we would prefer in case of unprecedented situations.

8 Subsystems

8.1 Command and control

By 2040, we will have ample speed to simulate 100 humans. If we allow for five times the resources required for virtual reality and auxiliary programs, or instead deploy 1 trans-sapient AGI agent (at $500\times$ human speed), we require 19.5 exaflop/sec computing speed. We may also assume 10 exabyte storage. The entire power budget thus is about 100 W, and with three-way redundancy for the computer system, we require 300 W power, and only 1 kg payload (which fits into a 1U unit). The AI can scale down to 2W on-route (which is $10\times$ human speed), since not all tasks need to be given minute attention during a long voyage.

8.2 Communication

Two reasonable approaches are laser communication which requires only 20W and a 3 meter telescope like Hubble, and gravitational-lensing amplification for RF (FOCAL mission) [27]. The FOCAL proposal exploits gravitational-lensing of Sun (and host system), and it requires a mission to 550 AU distance from Sun. We find that gravitational-lensing may be preferable since it might be lighter than optical communication, and since RF is not obstructed by interstellar medium. The analysis of radio communication using the FOCAL probe suggests that 40W power would be enough and we may use only an inflatable RF antenna. We assume an inflatable 8m RF antenna shaped as a parabolic dish, that fits into a 3U unit with 5kg mass, requiring 30W. The material currently used is mylar, but we estimate that lighter materials will be available, and larger dishes can be constructed for the same mass.

Name	Form	Mass	Power
Long-Range Imaging	3U	3kg	10W
Infrared Imaging Spectrometer	1U	1kg	10W
Imaging Spectrometer	1U	1kg	10W
Mass Spectrometer	3U	3kg	20W
LIDAR	1U	1kg	20W
Dust detector	1U	1kg	10W

Table 2: Scientific Instruments

8.3 Scientific instrumentation

We assume a standard array of scientific instrumentation similar to those on interplanetary probes such as the New Horizons probe [9]. We propose a modular cubesat design and we have extrapolated approximate figures for various scientific instrumentation modules from current cubesat technology. The form, mass and power requirements of cubesat modules thus estimated are shown in Table 2.

8.4 Power

Power is vital in a long space mission. We think that the spacecraft must be supplied with various power systems so that during emergencies there is sufficient power to resume operation.

We propose using the following power modules:

1. Thermoelectric converter for main power: the converters will transform excess heat from the fusion chamber to electric power. We assume Quantum Well Film converters, which will provide $\sim 1.5\text{kW}$ for $30 \times 50\text{W}$ modules, with 600 gr mass and a $\Delta T = 200$. We assume $\sim 1.5\text{kg}$ coolant, which may fit in a 2U cubesat, with only 2 kg mass.
2. Radioisotope Thermal Generator (RTG) for auxiliary power: we propose plutonium for the longer half-life required during the interstellar mission. The unit is supposed to provide at most 50W for a 3U cubesat, with 3 kg mass.
3. Solar panels for auxiliary power: these will be most useful at the target system, possibly giving the probe indefinite lifespan. We propose using a deployable cubesat solar panel, which will fit into a 2U, 2kg unit, providing up to 50W power.
4. Battery (emergency power): useful when auxiliary power is not available and thruster has to be restarted, 2U, 2kg cubesat unit with 300Wh capacity.

8.5 Propulsion

As described in Section 2.2, fission, fusion, anti-matter and sail concepts are generally thought to be feasible, while a host of technological devel-

opments are necessary to realize them. Nuclear pulse propulsion seems too massive for small probes due to minimum pellet size limitation and massive containment. Laser/beam driver is massive as well, and requires auxiliary power and inter-dependent mission structure. Light-sail seems inappropriate for rendezvous missions. Pure antimatter rocket requires much antimatter which may not be available by 2040. We thus recall the anti-proton ignited nuclear pulse propulsion proposal (ICAN-II) [23], in which the basic idea is to employ antimatter annihilation to ignite micro-fission/fusion detonations, overcoming the massive overhead of usual nuclear pulse propulsion designs and taking a large payload to Mars. Likewise, [4] suggests combining antiproton induced fission and magnetically insulated inertial confinement fusion. Anti-protons to burn fusion pellets (Cassenti et al) A beam of anti-protons/positrons ignites fusion pellets which consist of a fissile core and fusion fuel encapsulated in a tungsten shell for containment, with an inner uranium shell.

On the other hand, AIMStar uses antiproton annihilation for burning fusion fuel droplets for a cubesat mission to the Oort cloud [19]. We have found that extrapolating the AIMStar mission is appropriate for our interstellar mission. AIM stands for Antimatter Initiated Micro-fusion; the crucial component is an antiproton cloud confined in a Penning trap. The trap is about 10 kg and has dimensions of $0.1m \times 0.1m \times 0.3$ which is the size of a 3U cubesat. Fusion fuel is magnetically injected into the reaction chamber. Antimatter is supplied at the rate of 10^{11} anti-protons/sec. The fusion pulse is provided by a 42ng $D-He3$ droplet (5×10^{15} pairs), which is 45nm in diameter. 5×10^8 anti-protons burn 2% molar mixture of fission fuel (e.g., U238) which fully ionizes the fusion fuel droplet, while the fission fragments are not radioactive. Thereafter, the droplet is magnetically compressed for a fusion ignition. A 50Hz cycle is maintained yielding 0.75MW continuous power in the form of protons and alpha particles. A chamber henceforth may transfer fusion power to hydrogen propellant, and an analysis of 100% efficiency is presented. We assume for sake of simplicity that such a chamber (not designed in AIMStar proposal) is possible and may be scaled down from the Daedalus chamber. We propose to maintain AIMStar's advantages, as micro-fusion is more cost-effective than Daedalus pellets, whereby propellant mass decreases by many orders of magnitude. If we assume a chamber made of molybdenum that is 1 inches thick and has a radius of 0.1m, then we estimate a 32kg chamber. Such a configuration would be feasible for interstellar probes, however note that additional components such as magnets would likely be required to harness the fusion products as thrust. In particular, one may wish to use the alpha particles as thrust and deflect the protons with a magnetic pusher plate.

Let us also note that Positronics Research LLC is currently developing a new positron rocket design that uses an attenuation matrix to transfer the energy of positron annihilation to propellant, while Positron Dynamics LLC is working on positron storage technology. We foresee another propulsion design for probes which annihilates positrons to ignite micro-fusion as we may expect sufficient amount of positrons to be storable by 2040, although it is not yet known if such a design is possible, either.

Another recent development is the quantum vacuum plasma thruster

experiments developed by Eagleworks laboratory of NASA [44]. If the work in question bears fruit, it would instantly solve interstellar propulsion problems for it is a propulsion system that uses the quantum vacuum as propellant, requiring energy only for the generated fields.

8.6 Navigation

Four pulsars can be used to compute location [7], however for that a proper X-ray telescope is required. Traditional X-ray telescopes are too massive for a probe, while newly developed lobster eye optics improves efficiency up to 1000 times [14]. A mini lobster X-ray imaging module for picosatellites has been proposed by Hudec et. al and a 10cm x 10cm x 30cm design is underway (3U cubesat). We estimate a 3kg device with 10W power.

We also require attitude control and auxiliary thrust for navigation. We assume 3×3 U Cubesat Ambipolar Thruster (CAT) plasma thruster modules, which spend 20W, and have 5 kg mass each, with the performance of 2mN continuous thrust for 10W, and 20mN thrust for pulsed 100W, as recently proposed in a crowdfunding project by Benjamin Longmier [26]. 2 modules shall have 4×2 small thrusters, 1 module shall have 3-axis 5×2 small thruster.

A sensor module provides navigation sensors, which we take to be accelerometers, gyros, and magnetometers for 3-axis. We also require an attitude-meter, a star-sensor and radiation monitor. The navigation sensors ought to fit into a 1U cubesat, with 1kg mass, using 10W power.

8.7 Shielding

Artificial mini-magnetosphere is an idea to shield spacecraft from charged particles by generating a magnetic field imitating earth's magnetosphere [2]. The laboratory experiment used a 0.5T natural magnet, and in various publications Bamford recounts that 50nT sufficient for space applications with a $\sim 100m$ radius. We estimate 500W, produced from fusion burn and a 1kg additional weight for mini-shield. We also require a carbon nanotube EMI shield, which has been previously tested on cubesat. We have not yet considered neutron shields, although Project Daedalus has, and neutral particles must be considered in any interstellar voyage.

9 Mission Profile

Table 3 summarizes the subsystems in terms of mass and power requirements. Note that this is merely a preliminary design, which requires many additional components to be realistic. We attempted to show that a reaction mass on the order of 100kg is realistic, although it may require many technological developments to achieve such mass efficiency. In particular, an antimatter ignited fusion propulsion system is likely going to be much bigger in practice, and also propulsion tanks are not accounted for in this calculation which is a major shortcoming. However, we still give a full mission profile in Table 4 assuming that all fusion energy is converted to

Name	Mass (kg)	Power (kg)	Name	Mass (kg)	Power (kg)
			<i>Science</i>		
<i>Propulsion</i>			Long-Range	3	10
Reaction Trap	10	100	IR Imaging	2	10
Chamber	32		Imaging Spectr.	2	10
<i>Power</i>			Mass Spectr.	3	20
Thermoelectric	2		LIDAR	1	20
RTG	3		Dust Detector	1	10
Solar Panel	2		<i>Control</i>		
Battery	2		AI	1	2-300
<i>Navigation</i>			Shielding	1	0-500
Mini-Lobster	3	10	<i>Communication</i>		
Sensors	1	10	Inflatable RF	5	30
Attitude Ctrl.	15	60	Spaceship Infr.	2	10
			Total	91	302-1102

Table 3: Subsystems

	Tau Ceti	Gliese 667C
Distance	11.9 light-years	22 light-years
Cruise Speed	0.1c	–
Burnout time	9.183 years	–
Burnout distance	0.258 light-years	–
Antimatter required	0.97 milligrams	–
Fusion fuel required	36469 kg	–
Cruise Time	113.8 years	214.8 years
Total Voyage Time	132.2 years	233.2 years

Table 4: Mission Profile for AIM propulsion

thrust, because it corresponds to a best-case analysis for this propulsion technology. As it is quite likely that by 2040, technological progress will have been much accelerated compared to present rate.

Table 5 shows the mission profile assuming quantum vacuum plasma thruster technology. A SAFE-400 reactor with 512kg mass and 400kW thermal output, and 100kW electrical output is assumed. We assume 4N/kW thruster performance, and a thruster with 1000kg mass. It is observed that the quantum thruster might achieve extraordinary performance.

	Tau Ceti	Gliese 667C
Distance	11.9 light years	22 light years
Mass	1561 kg	–
Cruise Speed	0.5c	–
Burnout time	18.5 years	–
Burnout distance	4.63 light years	–
Cruise Time	5.26 years	25.4 years
Total Voyage Time	42.3 years	62.5 years

Table 5: Mission Profile for Q-Thruster

10 Extended Roles for Intelligent Interstellar Probes

Intelligent interstellar probes may be supplemented with many technological tools for performing extended roles at the target system. Self-reproducing interstellar probes is a well-known concept, and valid extension of the Daedalus probe [17]. Likewise, our probe may be extended for this most intriguing extended role, however fashioned for high technology and the miniaturized design of our probe.

By 2040, we expect nano-assembly technology to be widely available. Such a machine is best characterized by the ability to construct nanotechnological machinery by assembling atoms. If all the components in the above text are designed in this manner, a probe can assemble parts of itself for repairs and additional construction. An amount of material may be brought along, considering possible repairs, which is the most immediate application of a nano-tech assembler. However, if material can be mined at the host system, and designs only require materials that are easily mined, it may well be possible to 3d print the entire probe at the destination.

Refueling is almost as significant a role as reproduction and it is required for reproduction to be useful. It may be conceived that the probe may mine D and He_3 from gas giants and water might be mined from carbonaceous asteroids, however, much additional hardware would have to be added to the probe for this role, defeating the point of miniaturization. To mine the gas giants, the probe may print balloon supported factories. Instead, we may try to design propulsion systems that use widely available fuel so that refueling task requires less effort.

Construction at the host system is an appealing role, and asteroid material might be sufficient for many products, given how versatile the carbon atom is. At the destination, it would be worthwhile to construct a station for enhanced solar (stellar) power, computation, and communication facilities. It may be possible to bootstrap essential parts of a technological civilization at the host system via this route. If mini-probes, robots, and landers may be printed, they can be used to extend the survey and increase the capabilities. For instance, if miner robots and more nano assemblers may be printed, miners can gather much more material,

assemblers can produce required technologies, and an industrial base may thus be established.

Robotic bodies can be built for sims at the target system. Differences of sims can be communicated back to earth, so that sims at Sol may merge. Likewise, sims from Sol may be uploaded to the host system for visiting the facility.

Acknowledgments

Thanks to participants of 100YSS symposium for making several suggestions which improved the paper. Thanks also to Centauri Dreams blog author Paul Gilster for keeping us up to date on interstellar research. Thanks to Zeynep Acuner for sending the reference on pulsar based navigation.

References

- [1] G. Anglada-Escudé, M. Tuomi, E. Gerlach, R. Barnes, R. Heller, J. S. Jenkins, S. Wende, S. S. Vogt, R. P. Butler, A. Reiners, and H. R. A. Jones. A dynamically-packed planetary system around GJ 667C with three super-Earths in its habitable zone. *Astr. Ap.*, 556:A126, August 2013.
- [2] R Bamford, K J Gibson, A J Thornton, J Bradford, R Bingham, L Gargate, L O Silva, R A Fonseca, M Hapgood, C Norberg, T Todd, and R Stamper. The interaction of a flowing plasma with a dipole magnetic field: measurements and modelling of a diamagnetic cavity relevant to spacecraft protection. *Plasma Physics and Controlled Fusion*, 50(12):124025, 2008.
- [3] Justin R. Cantrell, Todd J. Henry, and Russel J. White. The solar neighborhood xxix: The habitable real estate of our nearest stellar neighbors. *The Astronomical Journal*, 146(4):99, 2013.
- [4] B. N. Cassenti, T. Kammash, and D. L. Galbraith. Antiproton catalyzed fusion propulsion for interplanetary missions. *Journal of Propulsion and Power*, 1997.
- [5] Schulze-Makuch D, Mndez A, Fairn AG, von Paris P, Turse C, Boyer G, Davila AF, Antnio MR, Catling D, and Irwin LN. A two-tiered approach to assessing the habitability of exoplanets. *Astrobiology*, 11(10), December 2011.
- [6] RR de Ruyter van Steveninck and SB Laughlin. The rate of information transfer at graded potential synapses. *Nature*, 379:642–645, 1996.
- [7] X. P. Deng, G. Hobbs, X. P. You, M. T. Li, M. J. Keith, R. M. Shannon, W. Coles, R. N. Manchester, J. H. Zheng, X. Z. Yu, D. Gao, X. Wu, and D. Chen. Interplanetary spacecraft navigation using pulsars. *Advances in Space Research*, 52:1602–1621, November 2013.

- [8] Ulam S.M. Everett, C.J. On a method of propulsion of projectiles by means of external nuclear explosions. Technical report, University of California, Los Alamos Scientific Laboratory, August 1955.
- [9] G. H. Fountain, D. Y. Kusnierkiewicz, C. B. Hersman, T. S. Herder, T. B. Coughlin, W. C. Gibson, D. A. Clancy, C. C. Deboy, T. A. Hill, J. D. Kinnison, D. S. Mehoke, G. K. Ottman, G. D. Rogers, S. A. Stern, J. M. Stratton, S. R. Vernon, and S. P. Williams. The New Horizons Spacecraft. *Space Sci. Rev.*, 140:23–47, October 2008.
- [10] Ben Goertzel. Lojban++: An interlingua for communication between humans and agis. In *AGI 2013, Beijing*.
- [11] Ben Goertzel. Opencogprime: A cognitive synergy based architecture for artificial general intelligence. In George Baciu, Yingxu Wang, Yiyu Yao, Witold Kinsner, Keith Chan, and Lotfi A. Zadeh, editors, *IEEE ICCL*, pages 60–68. IEEE Computer Society, 2009.
- [12] Ben Goertzel, Hugo de Garis, Cassio Pennachin, Nil Geisweiller, Samir Araujo¹, Joel Pitt¹, Shuo Chen, Ruiting Lian, Min Jiang, Ye Yang, and Deheng Huang². Opencog bot: Achieving generally intelligent virtual agent control and humanoid robotics via cognitive synergy. In *Proceedings of ICAI 2010 Beijing*.
- [13] Project Daedalus Study Group. Project daedalus. the final report on the bis starship study. *Journal of the British Interplanetary Society*, 31, 1978.
- [14] R. Hudec, L. Pina, A. Inneman, and L. Sveda. Lobster - astrophysics with lobster eye telescopes. In Andrei P. Lobanov, J. Anton Zensus, Catherine Cesarsky, and Phillip J. Diamond, editors, *Exploring the Cosmic Frontier*, ESO Astrophysics Symposia European Southern Observatory, pages 73–74. Springer Berlin Heidelberg, 2007.
- [15] Marcus Hutter. Towards a universal theory of artificial intelligence based on algorithmic probability and sequential decisions. *Proceedings of the 12th European Conference on Machine Learning (ECML-2001)*, (IDSIA-14-00):226–238, September 2001.
- [16] Marcus Hutter. *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability*. Springer, 2005.
- [17] Robert A. Freitas Jr. A self-reproducing interstellar probe. *Journal of the British Interplanetary Society*, 33, 1980.
- [18] R. K. Kopparapu. A Revised Estimate of the Occurrence Rate of Terrestrial Planets in the Habitable Zones around Kepler M-dwarfs. *Ap. J. Letters*, 767:L8, April 2013.
- [19] Kevin J. Kramer, Raymond A. Lewis, Kirby J. Meyer, Gerald A. Smith, and Steven D. Howe. Aimstar: Antimatter initiated microfusion for pre-cursor interstellar missions. *AIP Conference Proceedings*, 504(1):1412–1419, 2000.
- [20] Planetary Habitability Laboratory. Habitable exoplanets catalog. Retrieved from <http://phl.upr.edu/projects/habitable-exoplanets-catalog>, September 2013.

- [21] Geoffrey A. Landis. Advanced solar- and laser-pushed lightsail concepts. Technical report, Ohio Aerospace Institute, May 1999.
- [22] L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.
- [23] R. A. Lewis, G. A. Smith, E. Cardiff, B. Dundore, J. Fulmer, B. J. Watson, and S. Chakrabarti. Antiproton-catalyzed microfission/fusion propulsion systems for exploration of the outer solar system and beyond. *AIP Conference Proceedings*, 387(1):1499–1504, 1997.
- [24] K. F. Long, R. K. Obousy, A. C. Tziolas, A. Mann, R. Osborne, A. Presby, and M. Fogg. Project icarus: Son of daedalus flying closer to another star. In *British Interplanetary Society Daedalus after 30 years symposium*, September 2009.
- [25] K.F. Long, R.K. Obousy, and A. Hein. Project icarus: Optimisation of nuclear fusion propulsion for interstellar missions. *Acta Astronautica*, 68(1112):1820 – 1829, 2011.
- [26] Benjamin W. Longmier, Leonard D. Cassady, Maxwell G. Ballenger, Mark D. Carter, Franklin R. Chang-Diaz, Tim W. Glover, Andrew V. Ilin, Greg E. McCaskill, Chris S. Olsen, Jared P. Squire, and Edgar A. Bering. Vx-200 magnetoplasma thruster performance results exceeding fifty-percent thruster efficiency. *Journal of Propulsion and Power*, 27(4), 2011.
- [27] Claudio Maccone. Focal mission to 1,000au as an interstellar precursor. In *Mathematical SETI*, Springer Praxis Books, pages 349–360. Springer Berlin Heidelberg, 2012.
- [28] Thomas J. McGuire. *Improved Lifetimes and Synchronization Behavior in Multi-grid Inertial Electrostatic Confinement Fusion Devices*. PhD thesis, MIT Department of Aeronautics and Astronautics, 2007.
- [29] Hans Moravec. When will computer hardware match the human brain. *Journal of Transhumanism*, 1, 1998.
- [30] Stephen M. Omohundro. The basic ai drives. In *Proceedings of the First AGI Conference*, 2008.
- [31] Laurent Orseau, Tor Lattimore, and Marcus Hutter. Universal knowledge-seeking agents for stochastic environments. In Sanjay Jain, Rmi Munos, Frank Stephan, and Thomas Zeugmann, editors, *Algorithmic Learning Theory*, volume 8139 of *Lecture Notes in Computer Science*, pages 158–172. Springer Berlin Heidelberg, 2013.
- [32] Eray Özkural. Towards heuristic algorithmic memory. In Jürgen Schmidhuber, Kristinn R. Thórisson, and Moshe Looks, editors, *AGI*, volume 6830 of *Lecture Notes in Computer Science*, pages 382–387. Springer, 2011.
- [33] Eray Özkural. Diverse consequences of algorithmic probability. In David L. Dowe, editor, *Algorithmic Probability and Friends. Bayesian Prediction and Artificial Intelligence*, volume 7070 of *Lecture Notes in Computer Science*, pages 285–298. Springer Berlin Heidelberg, 2013.

- [34] Anthony Pancotti, John Slough, David Kirtley, Michael Pfaff, Christopher Pihl, and George Votroubek. Mission design architecture for the fusion driven rocket. In *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. American Institute of Aeronautics and Astronautics, 2012.
- [35] J. Schmidhuber. Gödel machines: Fully self-referential optimal universal self-improvers. In B. Goertzel and C. Pennachin, editors, *Artificial General Intelligence*, pages 199–226. Springer Verlag, 2006. Variant available as arXiv:cs.LO/0309048.
- [36] Jrgen Schmidhuber. Powerplay: Training an increasingly general problem solver by continually searching for the simplest still unsolvable problem. *Frontiers in Psychology*, 4(313), 2013.
- [37] Ray J. Solomonoff. A formal theory of inductive inference, part i. *Information and Control*, 7(1):1–22, March 1964.
- [38] Ray J. Solomonoff. The time scale of artificial intelligence: Reflections on social effects. *Human Systems Management*, 5:149–153, 1985.
- [39] Ray J. Solomonoff. Progress in incremental machine learning. Technical Report IDSIA-16-03, IDSIA, Lugano, Switzerland, 2003.
- [40] Ray J. Solomonoff. Three kinds of probabilistic induction: Universal distributions and convergence theorems. *The Computer Journal*, 51(5):566–570, 2008. Christopher Stewart Wallace (1933-2004) memorial special issue.
- [41] J. Storck, S. Hochreiter, and J. Schmidhuber. Reinforcement driven information acquisition in non-deterministic environments. In *Proceedings of the International Conference on Artificial Neural Networks, Paris*, volume 2, pages 159–164. EC2 & Cie, 1995.
- [42] Yi Sun, Faustino Gomez, and Jrgen Schmidhuber. Planning to be surprised: Optimal bayesian exploration in dynamic environments. In Jrgen Schmidhuber, KristinnR. Thrisson, and Moshe Looks, editors, *Artificial General Intelligence*, volume 6830 of *Lecture Notes in Computer Science*, pages 41–51. Springer Berlin Heidelberg, 2011.
- [43] M. Tuomi, H. R. A. Jones, J. S. Jenkins, C. G. Tinney, R. P. Butler, S. S. Vogt, J. R. Banes, R. A. Wittenmyer, S. O’Toole, J. Horner, J. Bailey, B. D. Carter, D. J. Wright, G. S. Salter, and D. Pinfield. Signals embedded in the radial velocity noise. Periodic variations in the τ Ceti velocities. *Astr. Ap.*, 551:A79, March 2013.
- [44] Dr. Harold ‘Sonny’ White, Paul March, Nehemiah Williams, and William ON. Eagleworks laboratories: Advanced propulsion physics research. Technical report, NASA, 2011.
- [45] Ivan Zolotukhin. The extrasolar planets encyclopaedia. Retrieved from exoplanet.eu, September 2013.