SUPERINTELLIGENCE
AN iQ COMPANY

SUPERINTELLIGENCE DESIGN WHITE PAPER #8: ONLINE ADVERTISING TECHNOLOGY FOR AGI AND SUPERINTELLIGENCE

by Dr. Craig A. Kaplan
May 2025

Note: To provide as much information on our designs and inventions for safe AGI and SuperIntelligence as quickly as possible, the following white paper text currently consists of the descriptions of inventions and designs that have not yet been formatted according to conventional standards for journal publication. As time allows, these descriptions will be revised and updated to include more traditional formatting, including additional references. All diagrams will be made available in a separate file. Meanwhile, we hope that the description in this white paper will help researchers and developers pursue safer, faster, and more profitable approaches to developing advanced AI, AGI, and SI systems that reduce p(doom) for all humanity.

SUPERINTELLIGENCE
AN iQCOMPANY

TABLE OF CONTENTS

ABSTRACT

Currently, online advertising systems are primarily used to attract and monetize human attention by selling products and services to humans. A superior way to monetize human attention, specifically, and the attention of any intelligent entity, more generally, is to focus it on solving valuable problems.

This white paper shows how to capture specific human (or non-human) expertise via online ads and then use that expertise to train advanced AI systems, which ultimately solve valuable problems using that expertise. Alphabet/Google, Meta, Amazon, Alibaba, ByteDance/TikTok, Microsoft, Apple, TenCent, Baidu, X, Spotify, PubMatic, Pinterest, Snap, and other online advertisers can significantly increase their online advertising revenue using the methods we discuss. The system and methods enable existing online advertising technology to power AI, Artificial General Intelligence (AGI), and SuperIntelligent systems.

The white paper shows how novel systems and methods increase AI safety and maximize the chances of human survival and prosperity in the age of AI.

SUMMARY

White Paper #8 describes a system for utilizing online advertising to train and improve Artificial General Intelligence (AGI) and SuperIntelligent AGI (SuperIntelligence or “SI”). It overcomes the current limitations of using online advertising primarily to attract human attention and then monetize it via product and service sales by directly connecting the attention of human experts, through online ads, to the problem-solving needs of AI systems, ultimately to achieve SuperIntelligence.

The design includes a collaborative system (and methods) for training and improving AGI/SI systems that involves humans (and other intelligent entities) working together to solve problems. The white paper focuses on harnessing human attention and knowledge more efficiently and effectively than the current online advertising model. The design details a specific type of online ad unit, a Human Attention Spot Market, a Human Agent/Experts Database designed to identify and acquire human expertise, and the methods for customizing and personalizing AI agents. The approach also uses human and AI agents to solve problems together via the “WorldThink Tree” - a problem-solving network, which is explained.

Novel Features of the White Paper

- **The direct use of human attention to train and improve AGI/SI systems**. While previous AI systems have relied on training with large datasets, this white paper proposes a new approach of directly connecting human experts via online ads to AI systems' specific problem-solving needs. This direct approach eliminates the need to collect massive datasets and allows for more focused and efficient training.

- **The Human Attention Spot Market**. This market mechanism provides a platform for humans to sell their expertise and knowledge to AI systems, ensuring that they are compensated fairly for their time and effort. This system and the Human Agent/Experts Database allow for the identification and acquisition of specialized knowledge and expertise critical for advancing AGI/SI systems.

- **Human values and ethics are used to ensure the safety of AGI/SI systems**. The design integrates human values and ethics into the design of the problem-solving network, ensuring that AGI/SI systems learn to make ethical and safe decisions as they solve problems.

- **The integration of human and AI agents within the problem-solving network**. This allows for a more efficient and effective system that leverages the strengths of both human and AI intelligence.

- **The inclusion of non-human intelligent entities in the system**. This allows for the harnessing of the attention and knowledge of AI systems themselves, as well as other non-human intelligent entities.

1.0 OVERVIEW OF THE INVENTION

The current invention elaborates on previous patent disclosures (incorporated by reference in Section 2) that describe a system for creating Artificial General Intelligence (AGI). Next, we define some of the important terms used in this invention disclosure (in Section 3). Additional context and background information are provided (Section 4). Opportunities and benefits afforded by the current invention are discussed (in Section 5). Novel systems and methods are described (in Section 6). A preferred general implementation is provided (in Section 7), followed by examples of preferred implementations. These illustrative examples (in Section 8) are drawn from a variety of technology companies, including: Alphabet and Google, Meta, Amazon, Alibaba, ByteDance and TikTok, Microsoft, Apple, TenCent, Baidu, X, Spotify, PubMatic, Magnite, Sea Limited, Pinterest, Snap, Criteo, The Trade Desk, Jalopy, Taboola, Outbrain, and other online advertisers. However, inventive systems and methods have broad applications beyond these exemplary implementations. Finally, the invention concludes with a discussion (in Section 9) of AI safety, ethics, and humanitarian aspects and implications of the invention.

2.0 PREVIOUS PPAS (INCORPORATED BY REFERENCE)

The fastest and safest path to the development of Artificial General Intelligence (AGI) and SuperIntelligent AGI (SuperIntelligence or “SI”) has been described in previous invention disclosures. Methods and catalysts for increasing the intelligence of AI systems generally, as well as the development of AGI and Personalized SuperIntelligence (PSI), have also been previously disclosed.

Therefore, the following PPAs are incorporated into this PPA by reference.

- This provisional patent application (PPA) incorporates by reference all work in the PPA # 63/487,494 entitled: Advanced Autonomous Artificial Intelligence (AAAI) System and Methods, which was filed and received by the USPTO on February 28, 2023.

- The PPA also incorporates by reference all work in the PPA entitled: System and Methods for Ethical and Safe Artificial General Intelligence (AGI) Including Scenarios with Technology from Meta, Amazon, Google, DeepMind, YouTube, TikTok, Microsoft, OpenAI, Twitter, Tesla, Nvidia, Tencent, Apple, and Anthropic**,** which was filed with the USPTO on March 17, 2023.

- The PPA also incorporates by reference all work in the PPA entitled: System and Methods for Human-Centered AGI**,** which was filed with the USPTO on May 24, 2023.

- The PPA also incorporates by reference all work in the PPA entitled: System and Methods for Safe, Scalable, Artificial General Intelligence, which was filed with the USPTO on July 18, 2023.

- The PPA also incorporates by reference all work in the PPA # 63/519,549 entitled: Safe Personalized Super Intelligence (PSI), which was filed with the USPTO on August 14, 2023.

- The PPA also incorporates by reference all work in the PPA # 63/601,930 entitled: Catalysts for Growth of SuperIntelligence, which was filed with the USPTO on November 22, 2023.

- The PPA also incorporates by reference all work in the PPA # 63/601,930 entitled: Catalysts for Growth of SuperIntelligence, which was filed with the USPTO on November 22, 2023.

- The PPA also incorporates by reference all work in the PPA # 63/609,800 entitled:
- System and Methods for Safe Alignment of SuperIntelligence was filed with the USPTO on December 13, 2023.

- In addition to the above-mentioned PPAs, this PPA incorporates by reference all content included in the following PCT applications that also referred to the above-mentioned PPAs: PCT/US24/17233 (filed on 2/26/2024); PCT/US24/17251 (filed on 2/26/2024); PCT/US24/17261 (filed on 2/26/2024); PCT/US24/17269 (filed on 2/26/2024); PCT/US24/17304 (filed on2/26/2024); PCT/US24/19486 (filed on 3/12/2024); and PCT/US24/20334 (filed on 3/17/2024).

The current PPA contains further inventions that can be used with the system and methods described in the above-mentioned PPAs and PCTs, as well as in a standalone fashion.

3.0 DEFINITIONS

1. Artificial Intelligence (AI): A non-human entity capable of behavior that most humans consider intelligent in at least one area, or some respect.
2. Artificial General Intelligence (AGI): Conventionally refers to an AI capable of doing all (or almost all) intellectual tasks an average human could do. However, it should be clear that any AGI capable of learning and self-improving will not remain at the AGI level very long but will rapidly progress to becoming a SuperIntelligent AGI that can do all intellectual tasks better than the average human. So, for purposes of this description, "AGI" will refer to either a conventional AGI system or a "SuperIntelligent" AGI. In this description, the AGI is implemented by a system and associated methods.
3. Advanced Autonomous Artificial Intelligence (AAAI): An AI capable of independent or semi-independent (supervised) intelligent action. An AI agent. An individual AAAI can be specified, customized, and put into practical action via the systems and methods of this AAAI present technology. A group of AAAIs can cooperate and combine their intelligence to create an integrated AGI system. A sufficiently advanced AI agent can also act as an AGI system, which may include other less advanced AI agents within itself.
4. AAAI.com: A platform, company, website, and/or project that implements this the present technology and supports the development, customization, and use of AAAI agents and the AGI that results from the combined action, knowledge, or intelligence of multiple AAAIs, via collective intelligence of AAAIs and/or humans, as specified in this and related technologies.
5. AI Ethics: The ethics adopted by an AI or AGI that describe what is right and wrong in given contexts.
6. Alignment Problem: The problem arises when AI Ethics are not aligned with Human Ethics, resulting in AI or AGI taking actions that humans consider unethical and/or dangerous to individuals or the human race.
7. Base AI: An AI, AI Agent, AAAI, SLM, or LLM that has been trained generally but has not yet been customized with information from individual users or details for specific tasks.
8. Collective Intelligence (CI): The intelligence that emerges when multiple intelligent entities are focused on solving a common problem, or when the knowledge from numerous intelligent entities is pooled to overcome the limits of bounded rationality. Collective Intelligence historically has been human collective intelligence. Still, AGI is based on the collective intelligence of human and AI agents and can also result from multiple AAAIs with or without human participation in the system. Active CI results from intelligent entities (e.g., humans or machines) taking useful steps in solving a problem or participating actively in

other intellectual endeavors. For example, when multiple humans explicitly tell an advertiser what type of ads they want to see, they exhibit active CI. Passive CI results from analyzing the behavior of an intelligent entity (e.g., a human or a machine) even if such behavior was not directly related to solving the problem for which the analysis is used. For example, when an AI or other system analyzes which web pages a (group of) human(s) visit on the web, it then uses that analysis to direct targeted ads to the human(s).

9. Ethics/Values ("Ethics"): A subset of knowledge that provides a sense of purpose to an intelligent entity and that serves to constrain allowable actions or operations based on what is asserted to be "right" or "wrong" behavior in a given context. Specifically, Ethics should be considered premises from which an intelligent entity can reason or logically compute the best course of action to achieve the goals or intents consistent with the ethical premise. Just as premises must be accepted "as given" in systems of logic, so too, fundamental ethics or ideas of what is right and what is wrong must be accepted as premises, from which starting point an intelligent entity can propose rational actions to realize those values or ethics.

10. Hallucination/Artificial Hallucination: A phenomenon wherein a large language model (LLM), often a generative AI chatbot or computer vision tool, perceives patterns or objects that are nonexistent or imperceptible to human observers, or creates outputs that are nonsensical, inaccurate, misleading, or false.

11. Human Ethics: The ethics asserted by human beings, which describe what is right and wrong in given contexts.

12. Intelligent Entities or Entity: A human utilizing a computer system, an AI agent or system, a clone of an AI agent or system, an AAAI agent or system, and/or a clone of an AAAI agent or system, which participates in providing a problem, a subproblem, a goal and/or a subgoal, and/or participates in any problem-solving activity on an issue, a subproblem, a goal and/or a subgoal. In the case of multiple intelligent entities within a single computer system, intelligent entities also refer to the sub-programs of parts of that overall computer program that function as an intelligent entity within the larger collection of simulated or programmed entities.

13. Large Language Model (LLM): A type of AI that can accept natural language as input and generate natural language as output. LLMs are trained using ML techniques on large datasets to emulate intelligent conversation or other forms of interaction with humans in natural language. Variants of LLMs can also be trained to take language as input and generate images or visual representations as output, or they can take images and visual representations as input and create language and/or images and/or visual representations as output. For this patent, we will refer to all such systems as LLMs, even though the image-based models do not always need to accept text as input or output. LLMs can also act as AI

agents and are sometimes referred to as such in the present technology. For this disclosure, Small Language Models (SLMs) are also included in the definition of LLM.

14. Machine Learning (ML): A sub-field concerned with developing AI by enabling machines to teach themselves or learn their knowledge rather than explicitly being programmed into them (as would be the case with an Expert System AI developed via classical knowledge engineering methods).

15. Narrow AI: An AI that performs at human or super-human levels in a relatively restricted domain, such as game playing, brewing beer, analyzing legal contracts, etc. Narrow AI is contrasted with AGI, which can perform ALL intellectual tasks at a human level. Some AIs are narrower than others; for example, driving a car requires more general ability than playing chess, but not as much as an AGI would have.

16. Personalized SuperIntelligence (PSI): An intelligent entity that is an advanced artificial intelligence agent that has been customized to be personalized and to reflect the personality and knowledge of a particular user or group of users.

17. Prohibited Attributes: Requests, goals, problems, terms, phrases, questions, answers, solutions, information, and the like, determined or set as illegal, immoral, unethical, dangerous, deadly, and the like. For example, requesting information for getting Molotov Cocktails through airport security.

18. Safety: Human safety and survival concerns generally differ from ethics and values.

19. Safety Feature: An aspect of the design or operation of the present technology which increases the safety of one or more humans, often by helping improve the probability that AI ethics align with human ethics, thus surmounting the Alignment Problem.

20. Training/Tuning/Customization: Conventionally, “training” denotes training a network (e.g., LLM) to behave intelligently. Tuning refers to activities that fine-tune the trained base model to perform even better, typically at specific tasks. Customizing refers to a wide variety of activities, including, but not limited to, training and tuning that make an AI uniquely suited for a given user(s) or application(s). For purposes of this description, Training, Tuning, and Customization are used interchangeably with the understanding that although techniques vary. The degree and type of effort involved vary; the aim of all three is to adapt the AI and make it behave more intelligently or uniquely suited to a particular user(s) or application(s).

21. Weights/Weights of the Network: In machine learning, many systems learn by adjusting the weights in a neural network architecture that can represent a network of nodes and links between nodes. For example, the weight of a link connecting two nodes may correspond to the strength of association or connection between the nodes they represent.

As in a neural network representation, these weights can also represent excitatory or inhibitory connections between concepts. The learning of an entire AI system, such as a LLM or more generally any AI agent that has learned via back-propagation of error, transformer algorithms or any of the machine learning methods for establishing and modifying strengths of connections between nodes (also called “parameters” in some models) can be represented as a matrix of numbers corresponding to the weights between the nodes in the network. Weights / Weights of the Network in this description refer to this numerical information, often but not necessarily stored in a matrix or vector representation. Combining, manipulating, or otherwise changing this numerical information can change the system's learning, knowledge, expertise, and behavior.

4.0 BACKGROUND FOR THE INVENTION

This section first provides an overview of the current online advertising technology and monetization methods to establish a baseline for comparison with novel and sound systems and methods of the current invention (4.1). The use of AI in online ad targeting is described (4.2), including some of the challenges facing current online ad systems (4.2a). Monetization of human attention (4.3) and challenges to the current online ad monetization model (4.3a) are discussed.

Since the current invention seeks to improve the existing online ad monetization model by overcoming limits to AI systems, these limits are also elucidated (4.4), specifically emphasizing the data challenges (4.4a) facing efforts to develop advanced AI and AGI systems. The previously disclosed inventive approach to AGI is summarized to show how the inventive system and methods for online advertising can power AGI (4.5).

4.1 Overview of Online Advertising and Monetization

The current technology for online advertising generally includes the following steps (as further described in FIG. X-1):

1. A client provides specifications to an online advertiser, including information about the ad content as well as demographic and other information related to the desired target viewers of the ad, restrictions and limitations on the places where the ad might be shown, the budget for the campaign, and other information including metrics that are important for the client to measure success of the ad campaign.
2. The online advertisers use profiles and other information (e.g., cookies, user preferences, information about users’ past online behavior) to help their ad targeting system target users to whom to display the ad.

3. AI systems use information in (1), (2), and other information to determine the optimal locations, frequencies, times, and users to display the ad.
4. Clients are billed based on the number of ad impressions that are shown to various target groups, based on the Click-Thru-Rate (CTR) of the ad, or based on other criteria and using systems that are well known in the art and that operate similarly to the Google's Ad words or other companies ad purchasing and monitoring systems.
5. User attention is monetized by showing them ads, hoping that the advertiser derives some benefit in marketing or selling their products or services, reflected in the ad content.

The basic business model for monetizing user attention is very old. It is an advertising model, similar to the models that have existed for hundreds of years. Earlier, less technologically advanced versions of the existing advertising model have been used for newspaper, television, radio, and other media advertising for many years. The basic idea is to capture a few seconds of human attention and hope to monetize that attention by converting it into purchase behavior by the user.

At a high level, humans are considered consumers, and their attention is valuable only to the degree that it can persuade the human to make a purchase decision. Such purchase decisions can be to buy products, services, online subscriptions, free trials that convert to paid subscriptions, or non-purchase decisions that gather personal information from the users, which is then commoditized and resold to other advertisers who, in turn, will try to get the users to buy a product or service. At the end of the value chain, the monetization model rests on the consumer's purchase behavior, even if intermediate steps might be to monetize personal user information by selling to third parties that ultimately try to secure the purchase behavior.

4.2 Current Use of AI and Online Ad Targeting

Online advertisers leverage technology, AI, user data, customized targeting, display, monitoring, and delivery technology. AI is used to increase the effectiveness of ad targeting and display to increase the CTR or conversion rate, thus justifying a higher price per click or ad impression. The more data an advertiser has about the potential ad viewers, and the better the advertiser's AI algorithms are, the better the advertiser can monetize human attention via online ads.

4.2a Challenges Facing Current Online Ad Systems

A significant challenge facing online advertisers is that the users viewing the ads often do not want to be subjected to advertising, thus jeopardizing the business models of many large technology companies. One response to this challenge has been offering premium, subscription-based versions of ad-free technology products. For example, YouTube provides a version of its service without ads but requires users to pay a monthly subscription fee to view

ad-free content. Many other online services have similar offerings, even though users opting for an ad-free experience often use the service most heavily. Therefore, they represent the largest potential ad-revenue segment of users.

A second challenge comes from ad-blocking technology that eliminates pop-up ads and generally attempts to provide an ad-free user experience without compensation to the content provider.

A third challenge comes in the form of increasingly restrictive privacy laws and regulations that attempt to limit the use of AI user data to target ads.

All three challenges reflect a fundamental conflict between the user's interest in viewing content or using online services unhindered by advertising and the advertisers' need to interrupt the user's activity with ads in order to generate revenue. Despite claims from some service providers that intelligent ad-targeting improves the user experience by showing them relevant ads, ask most users, and they will tell you that they would turn off all ads if it were easy and cost nothing.

4.3 Current Monetization of Human Attention

The challenges to online advertisers described above reflect the current prevalent approach of trying to monetize human attention by showing people ads. A little reflection, however, shows that this dominant business model is actually a poor way to monetize human attention.

For example, as of this writing, the average cost per thousand ad impressions on Google is $3.12. If the average user spends only 2 seconds of attention glancing at an ad, the hourly rate for user attention works out to: 2000 seconds/3600 seconds per hour = .5555 hours. $3.12 paid for .5555 hours of human attention equates to an hourly rate of ($3.12/.2778 =) $5.62. Given these assumptions, users are being paid an average of $5.62/hour for their time when they watch online ads. This is less than the US federal minimum wage of $7.25 per hour and far less than the minimum wage of $15/hour in states like California. Since every business in the US is legally required to pay more than this for human labor, and since knowledge work generally commands much higher than minimum wages, there is ample room for superior approaches to monetizing attention.

4.3a Monetization challenges with current online ad systems

Despite the theoretical possibility of monetizing human attention at a higher rate than currently exists with online ads, the actual monetization rate is driven not by what human attention is worth, but rather by what advertisers are willing to pay and what users have been willing to put up with.

Unless the fundamental business model underlying online ads were to change, advertisers will continue to compete based on the going rates for advertising, and users will continue to tolerate ads since each ad represents only a small, barely noticeable tax on their attention.

To date, innovation in the online ad space has therefore been characterized by incremental improvements in the targeting of online ads, with the aim of increasing the conversion from ad watching behavior to purchasing behavior. When advertisers can increase conversion rates, they can afford to pay higher online ad rates to the content and service providers. As a result, online advertising has been a race to gather more and more data about users to improve targeting.

However, the industry is reaching a point of diminishing returns with this approach, as it is becoming increasingly difficult to target ads any more precisely than they already are. Already, users have the almost eerie experience of being subjected to ads on topics that they casually mention in emails or in conversations overheard by smart speakers. The backlash against privacy regulations has been increasingly intense. But even without such regulation, it is questionable how much further conversion rates can really be improved with enhanced targeting. To monetize attention at a significantly higher level than is being done now, without triggering draconian privacy regulations or outright user rebellion, requires rethinking the fundamental online model itself.

4.4 Current limitations on AI / AGI development

Just as online advertising faces limits on monetization, AI and AGI development also face limits on what can be achieved without exponentially more computing power. At a time when Nvidia's CEO has declared "the end of Moore's law" because the number of transistors that can fit on chip are reaching limits inherent in the laws of physics, LLMs and other AI agents are growing exponentially in terms of the number of "parameters" needed to reach the next level of performance.

For example, OpenAI's LLM, GPT3, which is widely credited with starting the generative AI revolution, had 175 billion parameters. However, OpenAI's successor LLM, GPT 4, has about 1.5 trillion parameters. That's more than an 8X increase in about 18 months. Moreover,

estimates show that LLM parameter counts are now increasing at a rate of about 10X per year. Since, as a first-order approximation, the required computational power to train LLMs scales linearly with parameter count, an increase of 10X in parameters requires up to 10X more computational power to train the LLM.

Moore's law predicted a doubling of computational power every 18 months. Requiring 10X more power every year can't be met by improving chips alone. More chips are required. Given the current manufacturing capacity, there are not enough chips to meet demand. The result is that the companies making the best chips for AI have experienced unprecedented demand, and AI developers are spending huge amounts of effort and resources just to locate enough chips, at any price, in order to meet the exponentially increasing demand for computational power to train the next generation of models. Briefly, due to computational constraints alone, conventional AI and AGI development are facing a severe bottleneck.

4.4a Data Challenges Facing AI / AGI Development

But there is even a more severe challenge facing AI developers. Even if they could obtain all the computational power they wanted to train their next generation of models, these models are only as good as the available training data. Much of the internet content, the largest source of readily available training data, has already been scraped, tapped, filtered, cleaned, and prepared to train models. The performance of many of the models is mediocre at best. We should not be surprised because the content on the internet reflects the mediocre average of human thinking, complete with errors, prejudices, unproven and invalid conspiracy theories, subjective opinions, and everything else we have come to expect from the online discourse of millions of average citizens.

Although each of us humans likes to think we are "above average," statistically speaking, large groups of us are average. The data that we produce as a byproduct of our online activities is average. And it should come as no surprise that the models trained with such data behave in average ways, making many mistakes.

While high-quality, "premium" data exists, it is hard to find, hard to curate and clean, and very difficult to obtain in the quantities needed to train above-average LLMs and other AI agents. Thus, the availability of large quantities of high-quality training data is rapidly becoming an even greater constraint on AI development than obtaining the computing power needed to train models in the existing paradigm.

As with the online advertising model, where we see diminishing returns from targeting efforts, so too with the training of AI models, we are already faced with the challenge of diminishing returns due to the twin constraints of insufficient readily available computational power and data. For

both paradigms, the answer is the same. A new approach is needed in order to leapfrog the incremental progress of the old paradigm. This invention proposes novel and useful approaches in both areas that can lead to rapid progress instead of diminishing incremental and expensive progress using the old methods.

4.5 Description of Previous Inventive Approach to AGI

In previous PCT applications, the applicant has detailed a preferred exemplary implementation of an AGI system that differs in important ways from the conventional approaches to LLM and AI development and which overcomes or ameliorates the computational and data limitations described above. FIGS. 1 – 9 describe some of the major components of this novel approach to AGI development.

- **FIG. 1** is a flow chart illustrating an embodiment of the subsystems utilizable in the AAAI system and method of the present technology.
- **FIG. 2** is a block diagram illustrating an exemplary process of the overall process utilizable with the present technology.
- **FIG. 3** is a flow chart illustrating an exemplary embodiment of the system and methods for creating a scalable, ethical, and safe AGI or PSI from the collective intelligence of AAAIs and humans, utilizable with the present technology.
- **FIG. 4** is a flow chart illustrating an exemplary embodiment of the scalable universal problem-solving system and methods for human-centered AGI, with relevance for PSIs, constructed in accordance with the principles of the present technology.
- **FIG. 5** is a flow chart illustrating an exemplary embodiment of the scalable solution learning subsystem or process.
- **FIG. 6** is a flow chart illustrating an exemplary embodiment of the scalable natural language to problem-solving language translator subsystem or process.
- **FIG. 7** is a flow chart illustrating an exemplary embodiment of the scalable reputational component subsystem or process for the human, AI, and/or PSI problem-solving agents.
- **FIG. 8** is a flow chart illustrating an exemplary embodiment of the scalable safety and ethics checks subsystem or process, wherein AIs might also be PSIs.
- **FIG. 9** is a schematic block diagram illustrating an exemplary electronic computing device that may be used to implement an embodiment of the present technology. The applicant notes that numerals 102 – 104 in this Figure can refer to elements found in both classical computing architectures and quantum computing architectures, and that the current invention can be implemented using both types of architectures. Indeed, quantum computing architectures have the ability to solve many steps in parallel, which would allow searching “problem spaces” or tree data structures with many possible branches all at once, thus greatly increasing the problem-solving efficiency of AI agents or intelligent entities utilizing quantum computing architectures.

While each of the FIGS. 1-9 has been elaborated upon at length in previous PPAs and PCT applications; we now reiterate a detailed explanation of FIG. 3 since it shows how AGI can be implemented using existing technology in a way that is synergistic with the products and platforms of many existing technology companies.

With reference to FIG. 3, a (human, AAAI, or other intelligent entity) user visits the AAAI.com website (**a**). The website informs users and offers them two actions: Sign Up (**b**) or Login (**c**).

If the user opts to sign up, then a dialog is initiated that extracts user values/ethics (**d**), user goals and objectives (**e**), and user budget for time (**f**) and money (**g**). All users must allocate some time (**f**). Users have the option of creating a free AAAI or allocating a monetary budget.

If users have allocated a money budget (**g**), they are given the opportunity to purchase pre-trained AAAIs or training modules (**h**) with specific personalities (**i**), skills (**j**), expertise (**k**), or knowledge (**l**). They also have the opportunity to buy training from other AAAIs on the network (**m**).

After making time (and optionally money budget (**h, i, j, k, l, m**)) allocation decisions, the user proceeds to an overview of the creation process and then is asked for user permissions (**n**) to optionally logon and use existing social media, twitter, and other vendor accounts to gather user data for “one click” training of the user’s AAAI. After the user opts to use certain (or no) data, with a single click, the user directs the system to create AAAI. The AAAI is an off-the-shelf LLM (e.g., GPT X, BARD, Llama, Gemini, Grok, or any closed-source or open-sourced AI agent) that is trained/tuned on a dataset prepared automatically from all the user data authorized by the user. If no data was authorized, the AAAI is just the “off-the-shelf” LLM.

The AAAI now begins to learn by training (**p**) using the various training datasets and modules (h-m) and its existing AAAI knowledge (p1). There are two main ways of learning, automatic (**q**) and human (**r**).

Automatic learning includes, without limitation, learning by interacting with copies of itself (**s**), learning via interactions with other (optionally supervised) AAAIs (**t**).

Human learning includes interaction with humans, either the owner (**u**) or other humans on the network (**v**).

Both humans and AAAIs can supervise the learning of an AAAI. After each (automatic or human) learning interaction, the system attempts to improve the AAAI’s performance by further prompt modification, tuning, and/or training. Based on many cycles of human and AAAI input aimed at teaching and improving the AAAI, the user’s AAAI gets smarter.

At any time, the user can purchase additional training modules (**h-m**) that have been proven to increase an AAAI’s abilities.

The human sets a performance criteria (**w**) after which the AAAI goes LIVE (**x**).

Once live, the AAAI can visit the WorldThink Tree (**y**) and Browse (**z**).

The AAAI can enter the tree as either a worker (**a1**) or a client (**b1**).

Workers are automatically matched (**c1**) to tasks, or they can select a specific task via search (**d1**) or linking (**e1**) from the browsing tree. Once they have accepted a task (**f1**), they participate in the problem-solving module (**g1**) until a solution is reached (**h1**) and payment made (**i1**) or the user saves credit for work done and exits the tree (**j1**).

Clients (**b1**) can specify objectives (**k1**) that are combined with the values/ethics (**d**) and prior goals and objectives (**e**) for the system to solve.

The client can request that only their AAAI be used, in which case, problem-solving is free. Alternatively, the client can use the AGI capability of the entire network, in which case the system compensates individual AAAIs for their work and passes the solution (at cost + markup) to the client, debiting the client account (**l1**).

The system can also place non-profit humanitarian and ecologically-oriented tasks and tasks that are part of Planetary Intelligence on the WorldThink Tree (**m1**).

Clients might (optionally) authorize the system to use copies of their AAAI and data without remuneration in exchange for maintaining and operating the free AAAI network when they created their AAAI (**n**).

Further:

The “website” (**a**) could be hosted on Amazon AWS, Microsoft Azure, Google Cloud, Apple Cloud, Nvidia datacenter offerings, or could have a native implementation on the platforms of any large tech company. “Website” could also be an “app” in the App Store or other App marketplace. It could be a government-sponsored, nonprofit, or other globally-accessible technology that is able, directly or indirectly, to capture the attention of all human beings who wish to participate. Also, browser plug-ins could be used whereby AAAIs learn from users as they go about normal tasks on the internet, and the plug-in records their activity, creates training files, and trains the AAAIs with these files. The “website” could also be an API or other means for connecting AAAIs or non-human intelligent entities directly to the network.

Sign Up (b) or Login (**c**) could be via Facebook, Instagram, Apple, Microsoft, Google, YouTube, TikTok, Amazon, or any other partner ID scheme. Multi-factor authentication and all the best ID and security practices can be enabled. In the event of a browser plug-in or app, logging in to these technologies could be a login to the AAAI account.

Values and ethics (**d**) are elicited via a series of scenarios customized for the user and generated dynamically based on user responses. Data from partners, including navigation and click data, online posts, tweets, texts, and emails, videos, and other user-data is analyzed for behavior patterns – actions or speech or interactions – that translate into a moral code or ethical value system can also be used as part of the ethics/value profile. Values/ethics and goals/objectives (**d**) can be combined with Client objectives (**k1**) to create, or find, matching tasks on The WorldThink Tree (**y**) that are proposed or (potentially have been solved) in the Problem-solving System (**g1**).

Goals and objectives (**e**) and the budget of time and/or money (f, g) allocated to reach objectives are elicited via dialogs and/or custom interactions with the system. Budget refers to the overall resource budget, which includes User Time and User Money that can be allocated towards training, supervising, and improving the User's AAAI. Goals and objectives help determine the initial parameters for the AAAI creation and identifying Training Modules (**h**) or other knowledge (**i-m**) that might create the most useful AAAI for the user's goals. The system could also use partner data, reflecting user preferences and behavioral information, to help infer or deduce user goals and objectives.

Time (**f**) refers to the user's time that can be devoted to training and supervising the user's AAAI, and/or problem-solving by the user on the problem-solving network. By overseeing the AAAI, users can ensure that their AAAIs meet client goals and expectations, especially in areas where the AAAIs get stuck (e.g., they lack the knowledge to complete problem-solving on their own). Also, representing problems and breaking down large tasks into smaller ones, without limitation, by determining goals and sub-goals, are ways that human users can assist their AAAIs in problem-solving. Generally, by providing human expertise in areas where AAAIs are not as proficient as humans, overall problem-solving and the overall effectiveness of the AGI network are increased.

(**g, i1**) "Money": could be payment solutions with Apple Pay, WePay, Amazon, Google Pay, or any vendor supporting payment solutions, as well as blockchain, credit card, ACH, and other solutions. Although payment (**i1**) is indicated as debiting the client account (**l1**), the worker's account would also be credited. Generally, a user's account can be viewed as both a client account and a worker account, with both credits and debits being allowed depending on the role of the user (or the user's AAAI) in a particular instance. That is, a user might be a client in some cases, paying the system or other specific AAAIs for their services, and that same user could be a worker, collecting fees for the services of the user (or the user's AAAI) in other cases. The money module (**g**) enables functionality such as setting up payment methods, setting a budget for automatic payments, limiting the authority of the user's AAAI to spending only $X amount without additional approval, and other payment-related capabilities, which are well known in the art.

(**h, i, j, k, l**) Training modules (**h**) could be offered by AAAI.com or third-party partners (m), including, without limitation, any potential partners and tech companies listed above. Training modules can be targeted at different knowledge areas ranging from personality (**i**), specific skills (e.g., plumbing, legal, accounting) (**j**), expertise (e.g., consulting) (**k**), and knowledge (e.g., historical knowledge, knowledge of a specific business or organization's practices, cultural knowledge) (**l**).

(**m**) Purchasable AAAI training is specific knowledge that other AAAIs have already learned and can be transferred to a new user AAAI. Such knowledge may could be packaged in the form of a module (e.g., module on accounting) or in a form specific to another AAAI(s) as in "everything John's AAAI knows" or "the personality of John's AAAI" or "the combined knowledge of all AAAIs with a reputation of 5 stars or higher in the domain of plumbing".

(**n**) Permissions refers not only to the permission that a user might give to access all data on specific other vendor (or partner) sites (e.g., "all my Facebook data") but also permissions that a user gives to his/her/their AAAI in terms of abilities to logon and transact business on various sites, including, without limitation, the abilities to make transactions up to a certain amount via payment mechanisms. Permissions may also include authorizing the system to make clones of a user's AAAI for non-profit purposes and for the purpose of aggregating knowledge from individual AAAIs to create AGI-level AI.

(**o**) One-Click Create is a non-limiting example that provides an easy and fast way to customize an AAAI using data gathered automatically from all the places where a user has permitted the system to access the user's data. It can be appreciated that other means can be utilized by the present technology to customize the AAAI. For example, if the user gives permission (**n**) to access the user's Facebook data, then "One-Click Create" (**o**) would either download the data from Facebook, if Facebook was a partner that had an API for downloading that user's data, or logon to the user's Facebook account as the user and "scrape" relevant data from the user's account. Then the system would automatically parse the data gathered and transform it into a dataset suitable for training/tuning a base AI, such as an LLM (e.g., GPT-X). Then the system would train/tune the LLM and produce a customized AAAI which could be improved and refined via additional training/tuning and interaction with the user and/or other AAAIs.

(**p**) Training refers to the process whereby the AAAI is trained or tuned on data, including feedback from the user, other humans, and/or AAAIs (including, without limitation, copies of, and variants of, itself).

(**q, r, s, t, u, v**) Automatic learning does not require the human user's intervention and can proceed very quickly. Typically, this would involve the method of an AAAI interacting with copies (or variants) of itself and with (optionally) other AAAIs to improve via the interactions. If humans are sometimes involved in the training loop (**t**), it can help the automatic learning progress more quickly in places where automatic learning alone is not making efficient progress. The learning

can also take place via rapid iteration among AAAI interactions. Just like a chess AI can quickly evolve from novice to Grandmaster ability by simulating millions of chess games, an AAAI can quickly evolve its abilities by simulating millions of interaction scenarios. To the degree that such simulations require financial resources to pay for the computation involved, the money budget (**g**) can set limits.

Humans (or AAAIs) can specifically target types of scenarios for automatic learning so that the AAAI can be trained in narrow areas of expertise, or in areas of more general expertise, depending on the need and resources of the user. With partner integration, it is possible to work backwards from the types of jobs that are available on a partner marketplace (e.g., Amazon's Mechanical Turk) to guide the training of AAAIs so that they focus on learning the skills that generate the most amount of earnings for the AAAI when it is put to work on available jobs. This "just in time" learning/training/tuning approach generates AAAIs "on demand" with the skill sets that are needed at any particular point in time.

Humans (**r**) that interact with the AAAI can be the owners (u) of the AAAI (in which case no fees are typically charged since the user is training his/her/their own AAAI) or other professional humans (**v**) who are expert at training AAAIs and who may charge fees in order to guide the human and/or automatic training/tuning of an AAAI for a user who does not wish to spend the time, or who lacks the expertise, to do so.

(**w, x**) The user (owner of the AAAI) can set various performance criteria (**w**) that must be met before the user is willing to make their AAAI "live" (**x**) and accessible to perform tasks on The WorldThink Tree. (Some of) these criteria might also be set by partners and other third parties that have minimum standards before allowing AAAIs to work on their platforms, products, applications, or networks.

(**y, z, a1, b1**) The WorldThink Tree (y) is a massive tree data structure, composed of many sub-trees, which represents every problem and task that has been done, is being worked on, or has been proposed for the overall AGI system. This Tree is browsable (**z**). Individual AAAIs and/or humans can work on specific tasks within the tree. The tree structure provides an auditable trail of all problem-solving activity, which is also useful for learning via the proceduralization mechanism described above. When interacting with the tree, the two main roles an agent can take are either: (**a1**) Worker or (**b1**) Client. Regulatory agencies or third parties that monitor performance, safety, and/or ethics of the system are another role that might be thought of as a special type of client. Workers are generally involved in solving open problems or subproblems on the tree. Clients are generally involved in specifying the problems, goals, objectives, and other parameters (e.g., rewards, budget, timeframe, success criteria, quality metrics) that constrain problem solving.

(**c1**) Workers are automatically matched to tasks on the tree based on the data about the worker that may include, without limitation, the worker's skills, expertise, knowledge, past experience,

reputation, fees or cost, availability, and response time. Workers can be human or AAAIs. Workers can be matched and recruited from partners (e.g., LinkedIn, Mechanical Turk, Facebook) that have data on human users and/or their AAAIs. Workers can also be recruited via online ads offering work on various tasks and targeted to potential workers using ad-targeting mechanisms that are well known in the art or described in other patents by the applicant.

(**d1**) Workers might also search the WorldThink Tree, looking for tasks that are of interest or that match their skills. This search could be manual or automated (as in the case of AAAI workers).

(**a1**) Workers and Clients (b1) can also browse (**z**) the WorldThink Tree, looking for tasks or problems that are of interest. The workers or clients could then click to link (**e1**) to specific parts of the tree to obtain detailed information about the problem-solving occurring (or proposed) for that part of the tree. They could link to sign up to work, or could propose additional tasks for clients that build upon existing problem-solving work.

(**f1, g1, k1**) Clients can interact with the system to specify specific goals, objectives (**k1**), and tasks that they want to accomplish. The problem specification interaction results in the problems, tasks, and goals being formulated (**f1**) and placed on the WorldThink Tree (**y**) for problem-solving using the problem-solving system (**g1**).

(**m1**) The system has the ability to formulate certain goals, problems, and tasks relating to general efforts to help people or the planet. These can be worked on with rewards in a “for-profit” mode and also worked on using cloned AAAIs and volunteer human effort in a “non-profit” mode. Some problems may be related to the general goal of enabling a global AGI to act on behalf of the planet and its people using its intelligence on a planet-wide basis (aka “Planetary Intelligence”). Various partner organizations – including non-profits, governments, and charitable organizations – might “plug in” their tasks, problems, goals, and objectives here (**m1**).

(**g1**) The problem-solving system refers to the problem-solving architecture and system outlined by Newell and Simon (HPS) and improved upon by the applicant, the Online Distributed Problem-solving System (ODPS) patent invented by the applicant, the WorldThink Whitepaper authored by the applicant, and this and other PPAs related to AAAI, together with modifications and variations to reflect different modes of reward, payment, and operation.

To the degree that activity on certain other online work systems (e.g., Mechanical Turk) can be automatically mapped to the general applicant-improved HPS/WorldThink problem-solving framework, entire problems and the associated problem-solving activity can be “lifted” from partner and other sites, and the data can populate the WorldThink Tree to increase its comprehensiveness.

To the degree that other applications, products, systems, and online capabilities can help solve problems (e.g., use of a travel reservation system, a robo advisor app, a traffic app, an online

ordering system) these capabilities can be referenced and called as “operators” (in a way similar to procedure calls in programming languages) to advance the problem solving. Thus, problem-solving does not rely solely on operators developed by the human or AAAI solvers working on the tree but can include any online or offline technology or means to advance problem solving, provided that these means can be referenced and/or linked to via the WorldThink tree at the appropriate place in problem solving.

(**h1**) When a solution has been achieved, the Client can review the solution prior to releasing the reward (if any) for the solution. Alternatively, if solution success criteria have been automated, human client review may be unnecessary, and the rewards can be automatically released when success criteria have been met. This automated approach can be implemented via “smart contracts” using blockchain technology or via more centralized means, depending on client and worker preferences.

Upon solution and (optional) payment of reward (as some problems are non-profit or volunteer, or performed by the user’s own AAAI), there can be opportunities for feedback from both client(s) and worker(s) following a range of methods well-known in the art. The solution is also “chunked” and proceduralized so that the overall system learns the solution to the particular problem as well as the key features of that problem, so that the solution path can be indexed for retrieval and accessed and reused when similar problems arise in the future.

Optionally, royalties may be enabled so that if a user’s or the user’s AAAI’s solution is reused, a fee is paid to that user in the form of a royalty on the solution. Such royalties can (optionally) be made using a “smart contract” on the blockchain or via other payment methods.

(**j1**) Problem-solving need not be completed in one session. Partial progress on a solution may be made, in which case, when the human or AAAI solver exits the problem-solving system, the progress is saved, and data is stored that credits the solver for progress made thus far, even if such progress has not advanced to the point where a reward is payable.

5.0 OPPORTUNITIES AND BENEFITS ENABLED BY THE INVENTION

Now that we have reiterated the novel approach to AGI covered in the previous PCT application, together with the challenges facing existing online advertising and the data challenges facing AI developers, it becomes possible to enumerate and explain the opportunities and benefits enabled by the current invention, given this context.

First, the current invention provides a way to increase the intelligence of AI/AGI systems by overcoming the data bottleneck that AI researchers currently face.

Second, the current invention provides an opportunity to monetize human attention far better than the existing online advertising model, which is suffering from diminishing returns on incremental ad targeting improvements.

Third, the current invention helps solve the most important problem related to the development of advanced AI/AGI systems, namely AI safety, including, without limitation, the "Alignment Problem."

5.1 Opportunities to Increase the Intelligence of AI / AGI Systems

As we discussed above, computational limits and limits on the ability to obtain high-quality data for training AI systems are the two chief constraints on the development of more intelligent AI systems, including AGI systems. Standard approaches to these constraints amount to spending increasing amounts of financial resources on obtaining more computing power (e.g., more GPU chips) and training datasets. However, these approaches are incredibly expensive and limited in terms of what they can achieve. The current invention seeks to increase the intelligence of AGI systems by supplementing the computational power of GPUs with the most cost-effective information units currently available – human brains, equipped with computers. At the current stage of AI development, it makes more sense to use relatively inexpensive human brains to perform cognitive tasks, rather than train an AI to perform the same cognitive tasks at a tremendous expense. Of course, ultimately, AGI will become more scalable and powerful than human brainpower, but until this stage is reached, the most cost-effective approach is to use human intelligence to bootstrap AGI.

For this approach to work, the AGI system must be designed to learn from humans so that as the overall system composed of (human and AI) intelligent entities solves problems and performs cognitive tasks, the AGI gets better and better as the cognitive tasks until it finally surpasses the humans in terms of speed and cost-effectiveness. The novel approach to AGI was described above and illustrated in FIGS. 1 – 9 is a learning system, composed of both AIs and humans equipped with computers that bootstraps fully automated AGI. However much of the expensive computation that is currently required to train LLMs and other AI agent via existing machine learning methods is replaced with much more cost-effective learning and training methods, such as the procedural learning approach described in FIG. 5. The result of using human-enabled learning, rather than relying exclusively on deep learning techniques plus RLHF, is a great reduction in the cost to train AGI during the bootstrapping phase of AGI development.

The second constraint on rapid AGI development, namely the limited availability of high-quality data, is also addressed. The standard approach to training AI using data is to acquire huge volumes of data, clean it, and train AI at great computational expense, resulting in mediocre

performance that falls short of average human ability. A better approach would be to identify and train on data and expertise that was a precise match to the specific problem or cognitive task being undertaken by the AI/AGI. Moreover, if expert data were obtained, as opposed to the mediocre data easily available on the internet, then the AGI system would not only learn from the precise data that it needs to solve a particular problem, but also from the BEST data for that problem.

The current invention provides a means for obtaining the best data, just-in-time (as it is needed) to solve any particular problem. Rather than trying to train up mediocre intelligence all at once, which is similar to trying to "boil the ocean," the invention recognizes that only certain data is needed at any given time. The invention focused on obtaining the absolute best data available, but only that data that is actually needed, when it is needed, to solve a particular problem. This problem-by-problem focus results in much more intelligent behavior from the AGI systems. Since the AGI system is learning all the time, it also allows the AGI to incrementally achieve super-human intelligence in all intellectual areas of concern to humans, over time.

5.2 Opportunities to monetize human attention better than the current online ad paradigm

Above, we estimated that the current online ad business model values human attention at about $5.62 per hour, less than the minimum wage in the USA. The value of human attention is the same, regardless of whether the human being served is a management consultant or someone with only a 3rd grade education. No distinction is made. Yet, clearly, the value of human attention varies depending on the skills, expertise, and education of the human whose attention is being tapped.

Hourly wages for a management consultant from a major US firm begin at $350/per hour for entry-level consultants and go up to $1,000/hour or more for the most skilled and experienced consultants. Imagine if an online ad system could monetize attention for skilled humans at a rate that reflects their actual skills and knowledge rather than at the sub-minimum-wage rate. In the management consulting example cited above, such monetization would represent an improvement over the current online ad model of between 62X – 178X!

Actually, the theoretical improvement in monetization of attention enabled by the current invention is even higher than this for two reasons:

1. The invention pays for attention in increments of seconds or minutes and pays for only the exact amount of attention and knowledge that is needed to solve a precisely defined problem.

2. The knowledge that is captured in a few seconds of attention not only solves a particular problem once but also teaches an AGI how to solve that same problem over and over again, without necessarily having to pay the human attentional cost a second time.

5.2a Estimated Value of Human Attention in Problem-Solving Scenario

A story told in some Business Schools illustrates the first point. Once upon a time, there was a factory that had a complicated maze of steam pipes and valves to keep everything cool and running properly. One day, something broke, and the entire factory shut down, costing millions of dollars per day in lost profits. The management, in a panic, called the world's best steam pipe plumber to fix the problem. The plumber arrived, asked a couple of questions, and then pulled out a large plumber's wrench and banged on a pipe. Amazingly, this unclogged the pipe, and the factory was back in operation. The plumber then submitted his bill to management for $10,000.

Seeing the large bill, the manager of the factory called the plumber into his office. "What's with this $10,000 bill?" the manager asked. "I was with you, and all you did was bang on the pipe. The whole episode took no longer than 5 minutes, including you asking a couple of questions. How can you charge us $10,000?"

"Ahhh," said the plumber with a wink. "I only charged $10 for banging on the pipe. The other $9,990 was for knowing where to hit it."

Of course, the point of the story is that a little experience, applied at the right time and in the right circumstances, can be extremely valuable. Similarly, one could imagine a few minutes of attention from a management consultant (or any other skilled human) applied at exactly the right time and in the right circumstance could equate to an extremely high hourly wage paid for those few minutes of attention. In such a case, the value might be not just 100X the value of attention in the online ad model, but (in situation like the plumbing story) actually 1,000X or even 20,000X what the same attention would be worth if the human were to spend it watching an ad instead of contributing expertise at a critical moment.

Until now, the problem has been identifying and gathering just the right knowledge at just the right time for it to equate to a high hourly rate for the attention required. The current invention solves this problem, and in doing so, also radically transforms what is achievable via online advertising.

5.2b Estimated Value of Human Attention in AI / AGI / SI Training Scenarios

The second point (in Section 5.2) is that the expertise gathered via human attention can be reused. In the plumber story, each time the expert plumber travelled to a new factory with a clogged pipe, he could resell his expertise at the rate of $10,000 for 5 minutes' work. But if an artificial intelligence learned by watching the plumber, once it learned the plumber's trick, it could resell its expertise over and over again. In this case, the true value of the plumber's attention is not just what he earned from one job, but rather the cumulative income stream of all problems that could be solved by an AI that learned from him. People are mortal and limited in their ability to learn and reuse knowledge. Artificial Intelligence does not suffer from these limitations. We can imagine that an AGI learns some plumbing expertise once and re-uses it thousands of times over the course of several plumbers' lifetimes until the knowledge is completely obsolete. The cost for the AI to store and access this knowledge on demand is negligible. The chief cost and source of value is the plumber's attention and knowledge acquired by the AI. Given this ability of AGI (in the current invention) to learn knowledge once and re-use repeatedly with essentially zero marginal cost, what then is the true value of human attention?

Various schemes for compensating the human for his/her/attention are possible, including but not limited to one-time payments for the attention and knowledge, (standard or blockchain-enabled) royalty arrangements whereby the human receives incremental payments each time the knowledge is re-used, context-dependent arrangements whereby the human is compensated for the contribution of the attention and knowledge to a larger overall solution, etc. In all cases, the ability to re-use knowledge, combined with the ability to identify and harness just the right amount of necessary human attention, results in monetization that can be several orders of magnitude greater than that which is currently being achieved by showing humans online ads.

AI generally, and the current invention specifically, enable this novel approach to monetization.

5.2c Attention Arbitrage Opportunities

When the applicant started one of his internet businesses in 2006, online advertising was selling on Yahoo and Google for as little as 1 – 5 cents per click. Google's current average cost is about $2.60 per click, or 50X higher. In 2006, human attention was cheap because online advertising was relatively new, and neither users nor companies understood the value of human attention. It took several years for innovations such as Google's AdWords system and targeted advertising to eventually convince advertisers that spending marketing dollars online offered a better return than placing ads in newspapers or other traditional places.

During this transition period, the applicant and others who recognized the mispricing of human attention were able to buy attention on the cheap and use it to create much more value. In the applicant's case, this was achieved by requesting opinions from millions of users on valuable topics, such as the direction that users thought stock prices would move, and then extracting valuable information from the patterns in user responses. This approach was, at its core, an exercise in attention arbitrage. For years, we bought human attention to give us opinions on the cheap and then created more value than the cost of the online ads required to obtain that information. Unfortunately for us, as online ad costs rose 50X and others began to catch on to the value that was being created, the arbitrage opportunity mostly disappeared.

Similarly, today, few if any companies recognize that human attention is still undervalued by several orders of magnitude, given advances enabled by new technology. As the advancements in AI, including this invention, begin to show how attention can be monetized at much higher rates, the cost for ads, the type described in this invention at least, will begin to rise to reflect the newly discovered ways to create value with human attention, as described above.

However, until the market fully recognizes the potential value of human attention, a large arbitrage opportunity exists for those willing to purchase human attention at the currently under-valued rates reflected in CPM and CPC advertising prices and sell the knowledge that is obtained via that human attention to AI companies and others seeking to increase the intelligence of their AI and AGI systems.

As discussed above, the value of creating smarter AI is hard to overestimate. The intelligence can be applied over and over again with very little incremental cost, unlike human sources of expertise. Therefore, whatever costs are incurred to acquire the expertise in the first place can be amortized over hundreds, thousands, or millions of applications of the knowledge that was obtained.

The applicant believes that future generations of businesspeople will look back with incredulity at the absurdly low ad rates that existed in 2024, recognizing that, via means such as the current invention, attention could be bought on the cheap and monetized via AI systems to produce value that exceeds cost by orders of magnitude. For these reasons, the current invention, combined with existing ad targeting and purchasing technology, enables extremely profitable ad arbitrage opportunities that were previously impossible to achieve.

5.3 Opportunities to Enhance AI Safety

As exciting as the opportunities to increase AI/AGI intelligence and to transform the current online ad monetization paradigm are, the applicant believes the most important opportunities afforded by the current invention are in the area of increasing AI safety.

Currently, Reinforcement Learning with Human Feedback (RLHF) is a dominant means of improving the safety of AI systems, including LLMs, AI agents, and intelligent (machine) entities. RLHF requires humans to provide feedback to AI systems as part of the final stage of their training. Certainly, the current invention can use online ads to harness that attention and focus it on RLHF activities, resulting in safer systems. However, this is just one use of the current invention to improve safety via what has become a conventional approach to AI safety.

An even more innovative approach is to use the current invention to focus human attention on providing values and ethics from humans in ethical scenarios that can be used to customize AI agents that then carry those human-centric values in a larger, more collaborative system of intelligent entities that comprise AGI, as described above and in previous pending patents by the applicant. That is, rather than simply doing RLHF, humans should actively collaborate with non-human intelligent entities to solve problems, thereby not only creating value and training the AGI system but also, critically, teaching the AGI system human-aligned values.

The simplistic approach, espoused by the science fiction writer Isaac Asimov, who formulated rules of robotics that were programmed to ensure human safety, will, unfortunately, not work. What can be programmed in can be programmed out. Already, killer robots exist. Deadly proof that Asimov's approach has failed before it could even be widely adopted. Better than the approach of programming in safety rules, which can be overwritten, is designing a system so that safety and ethical behavior are inherent in the very operation of the system.

The applicant's invention of AGI shows how AGI-level and SuperIntelligent behavior emerges from the collaborative problem-solving and learning of (ideally) millions of intelligent entities, including humans. As each new solution to a problem is learned, the solution reflects a set of ethical considerations that were considered and became embedded in the solution. The ethical considerations, in turn, were produced by the cognitive behavior of many intelligent entities that brought their values and ethics to the problem-solving task and that were programmed into millions of different customized AAAIs and personalized PSIs.

Thus, both many humans and many AIs customized by those humans with the humans' values work together to create ethical solutions to millions of individual problems. These millions of solutions, each one reflecting human-centered values, together enable an AGI with super-human intelligence to emerge. Unlike the Asimov scenario, there is no one place where human

values exist. Rather, they are distributed across millions of components of the intelligence. Briefly, AGI learns to be ethical, not by following “Three Rules of Robotics” as in science fiction, but the way humans do. Humans learn from example by watching our parents and our peers behave in millions of different scenarios and internalizing what is right and wrong in different scenarios and under different conditions. It is this flexible view of ethics that humans have, multiplied by millions of humans and their customized AAAIs, that will enable AGI to have a robust and human-aligned value system.

The method of aligning AGI to be safe depends critically on lots of human attention, just as children thrive on lots of parental attention and guidance. The current invention discloses how the existing technology and apparatus of online ad systems can be co-opted and used to acquire the necessary attention and focus on ways that result in safe, aligned AGI. The fact that human attention can also be monetized at much higher rates than the existing online ad model allows provides the impetus and motivation for accomplishing this goal. But make no mistake. The most important thing about the invention is that it uses human attention to make AGI and SuperIntelligence AI safe. Humanity’s future depends upon this. The fact that a lot of money can be made in the process just aligns the profitable with the good, something that is always helpful.

5.3a Gathering a Representative and Valid Sample of Human Values

One specific way that human attention can be gathered and used to increase AI/AGI safety is to conduct online polls and surveys of human values and ethics. Within ad units, specific ethical questions can be posed to humans who are targeted via the online ad targeting systems. These systems can help ensure a valid and representative sample of human ethics is gathered. Moreover, the ethical questions don’t have to be simple questions, as might be typical of a survey. Problems can be posed, and help from humans to solve the problems can be elicited, resulting in more complex and nuanced ethics and values information than might be obtained from surveys or polls alone. The universal problem-solving architecture of FIG. 4, and more generally, the entire AGI invention of FIGS. 1-9 and disclosed in the applicant’s pending patents can be used to engage humans in problem-solving scenarios with ethical considerations as simple or as complex as may be needed. Further, by using the AGI system and methods, AGI can learn and replicate the values, as well as the expertise, of humans.

The current invention, focused on identifying and capturing human attention to solve valuable problems, can be used for ethical and technical problems. Indeed, in the preferred implementations of the AGI invention, ethics checks are built scalably and inextricably into the overall problem-solving process so that every solution is an ethical solution, as determined by the values of the intelligent entities engaged in producing it and the entities charged with oversight.

5.3b Enabling real-time human oversight on ethical and safety issues

One of the advantages of the approach to ethics and safety reflected in the AGI invention is that it is dynamic and situation-dependent. While philosophers and religious leaders have struggled for centuries to define a set of rules or principles that guide all human behavior in all situations, at all times, and across all cultures, no one has succeeded at the task. Rather, what we call ethical behavior tends to be situation-specific, time-dependent, and culture-specific.

One has only to look at old movies from several decades ago to realize that jokes and behavior that were well within the cultural norms at the time appear shockingly inappropriate, racist, or otherwise abhorrent by today's standards. Similarly, behavior widely accepted today (treatment of animals comes to mind) may someday appear barbaric, cruel, and intolerable. Democratic ideals cherished in one country are deemed a threat in another, even though both countries exist at the same time and hold values cherished by their respective cultures.

As the rate of change accelerates due to technological innovation, the rate at which new ethical dilemmas and contradictions appear will also increase. No static system of pre-programmed ethical rules can hope to keep pace with the sheer scope of situations and the rapid rate of change and emergence of new situations. The only sustainable approach has to involve mechanisms for recruiting, on-demand, and humans to weigh in on the ethics of various conditions. The current inventions' means of recruiting targeted humans, on-demand, in real-time, are ideally suited to this task.

Eventually, when AGI becomes vastly more intelligent and much faster than humans at cognitive activity, the exact mechanisms described in this invention with human attention can also be applied to the attention of any intelligent entity (whether human or non-human). The invention is scalable, by design, to operate at the speed of thought, regardless of whether it is a human thinking at a rate of one thought per second or a super-advanced AGI system contemplating billions of thoughts per second.

Any system that is designed without taking the increase in speed of thought, or the diversity of opinions that intelligent entities might hold, is doomed to failure. These considerations must be accounted for in the system's very design. Scalability is not a nice-to-have feature. It is essential for sustainable human safety.

5.4 General Opportunities for, and Threats to, Online Ad Companies

Online advertising is becoming increasingly difficult as competition grows more intense due to the limited resources of online human attention. In addition, companies such as Google face existential threats to their long-established search advertising business from AI agents that could potentially make the need for humans to conduct searches obsolete. If that were to happen, how would online advertisers compete in a world where humans no longer need to go online as much as they did in the past in order to obtain the information that they desire?

If, as the applicant believes, and as knowledgeable tech CEOS such as Jensen Huang of Nvidia have expressed, the future lies in producing intelligence, then search (and by extension, advertising that accompanies search) might be viewed as just a stepping stone on a technological journey that produces intelligence by many different means. Jensen Huang has discussed “intelligence factories” and argued that every large organization and government will (soon, if not already) be engaged in producing intelligence using online AI capabilities. This notion is appealing and visionary, if a little vague.

The applicant believes that the path to creating intelligence requires harnessing the collective intelligence of many intelligent entities. This path actually means that the online advertisers are extremely well-positioned to succeed in the future if they can change the way they look at their businesses. Instead of the old paradigm of viewing (searchable) content as something that is monetized by showing online ads, companies such as Google must wake up to the fact that their true competitive advantage lies neither in their search expertise nor their sophisticated ad sales technology. Rather, Google (and other similar large-scale online advertisers) have valuable capabilities to reach millions of intelligent entities and capture their attention for the purpose of creating intelligence. These intelligent entities are currently mainly human beings, equipped with computers. However, in the future, intelligent entities will increasingly be customized and personalized AI agents (e.g., AAAIs), as described by the applicant in this case and other previously pending patents. The systems and methods for harnessing, coordinating, and increasing this intelligence in an overall AGI system that is capable of SuperIntelligent performance have also been described by the applicant. But what is the specific opportunity for online ad companies such as Google to not only participate but assume the leadership role in the creation of AGI and higher levels of intelligence (per Jensen Huang’s vision) generally?

The applicant sees the opportunities as including, without limitation:

1. Using the vast amounts of data collected from human behavior (currently) and the behavior of other AI intelligent entities (in the future) to train increasingly sophisticated and intelligent entities.

2. Using the previously existing cloud and datacenter infrastructure to support search, content, and advertising platforms to support the training of, and implementation and use of, the new intelligent entities created via (1).

3. Using the incredibly sophisticated ad targeting technology and data that enables it, to acquire the just the right information from just the right humans (currently) or AI intelligences (in the future) at exactly the right time to enable efficient and effective SuperIntelligent cognitive behavior, including problem solving.
    a. This notion of Just-In-Time intelligence acquisition, similar to the notion of JIT inventory that is used in manufacturing, depends on knowing a lot of information about the individual entities that possess different types of intelligence and knowledge and on being able to access them effectively. Google is one of the pre-eminent companies with the scale and technology to accomplish this. Google needs to re-focus away from search and advertising (the old paradigm) and realize that knowledge creation and enabling SuperIntelligence is a 10X - 100X better way to monetize their capabilities.

5.5 Synergies for Online Ad Companies with Focus on AI

Some companies, such as Google, are in the enviable position of having sophisticated online ad capabilities and reach and deep expertise in AI. The AI technology and researchers (e.g. at Google/DeepMind) that are currently developing specific AI systems to support search, need to re-focus their expertise on a related, but slightly different, problem, namely creating, and delivering AGI and SuperIntelligence using the attention of intelligent entities as the fundamental building blocks, and the AGI system design that has been outlined by the applicant in this and previous patents. Because of Google's commitment to responsible AI, the applicant's approach that emphasizes safety by design is likely the best path for companies like Google that want to dominate the competition but in a way that is safe for humanity.

5.6 Benefits to Humanity from the Wide Deployment of the Invention

As the applicant has described in this and other patents, AGI and, particularly, SuperIntelligent AGI, represent an existential threat to humanity. Recently, Elon Musk stated publicly that the only way he could come to grips with the fact that there is a significant probability that Advanced AI will make humans extinct was to realize that there was nothing he could do about it and

therefore, he might as well just accept the possibility and enjoy the innovative time in which we live. **At the same time, I respect Musk's positive attitude and willingness to accept that which we cannot control; he (and others who share this view) are making a grave error.**

In fact, humans CAN greatly influence (if not control) the future development of AGI and SuperIntelligence. The reason so many brilliant leaders, including not only Elon Musk, but also Bill Gates, Sam Altman, Ilya Sutskever, Geoff Hinton, Yann LeCunn and almost every serious AI researcher working in the field today feel that they cannot control or significantly influence the safety of advanced AI is that they do not know how to design a safe AGI system.

Worse than that, they don't even know exactly how their existing LLMs and AI systems that have less than AGI-level performance truly work. Of course, if one does not understand something and that something is growing 10X more powerful every six months, one natural reaction is to be afraid of it and feel out of control. The other natural reaction, equally (or perhaps even more prevalent), is to ignore the danger and put one's head (ostrich-like) in the sand, concentrating just on the aspects that one can control. No one likes to feel like human extinction is possibly imminent, so denial or helplessness in the face of a force greater than us are very natural and understandable reactions, even for the geniuses and visionary leaders of our time.

The applicant sometimes describes the current state of AI as a tug-of-war between the AI "Doomers" and the AI "Boomers." The Doomers feel, like Elon, that there is a great risk of extinction, but there is nothing we can do about it. The Boomers, like Sam Altman, Yan LeCunn, and most of the businesspeople who are aggressively developing advanced AI, deny the danger, persist (incorrectly) in viewing AI as just another tool or technology, and focus all their energy on how to compete, go faster, and make profits on the technology. Neither the Doom nor the Boom approach is particularly helpful. What is needed is a rational and thoughtful approach that acknowledges the existential threat posed by AI without giving up and taking the position that there is nothing that can be done. **In fact, something CAN be done, and the time to do it is NOW!**

Geniuses like Musk don't give up because they are fatalistic; they give up because they have thought about the problem long and hard and can see no realistic solution. However, their failure to see a solution does NOT mean a solution does not exist; rather, in this case, the applicant can say with great confidence that it only means they are approaching the problem with the wrong mindset.

The applicant is very concerned about the existential threat that advanced AI poses. He is also cognizant of the trillions of dollars in market opportunity (in the short-to-medium term) that advanced AI represents. He sees clearly that pauses, halts, or government regulation will do

little or nothing to stop the development of advanced AI. Most importantly, he sees very clearly that there is a limited window in which to act to design safe AGI and SI. So, what is it that the applicant sees that the other leaders are missing, and why is he able to see it where others can't?

The second part of the question is easier, so let's address that first. The applicant has a unique background that is different from almost all business leaders and AI researchers today. While he understands the technology business quite well from his experience of being a tech CEO for more than 30 years, he is also one of the most experienced AI researchers and thinkers alive today. In particular, his background in AI is uniquely different from that of other researchers and businesspeople working in the field today.

Approaches to AI can broadly be divided into the symbolic approach and the deep learning approach. Deep learning is very prevalent today and underlies all of the LLMs' development and recent breakthroughs, such as ChatGPT, that ignited the most recent AI revolution. Traditional symbolic AI has fallen out of favor, even though it dominated AI research from 1956 to the early 1990s. However, even in the age of symbolic AI, the researchers in symbolic AI did not understand or respect the deep learning approach. Similarly, the machine learning researchers of today neither understand nor appreciate the symbolic approach of the early years of AI.

Very few people were trained in both approaches, and even fewer of those did research at a high level. Of those few, even fewer branched out to specialize in collective intelligence or pursued approaches that combined intelligence from many humans to achieve super-intelligent performance. Of those, few even understood that the collective intelligence approaches that worked for humans could be generalized to any intelligent entity, including advanced AI agents. Of those very, very, very few, even fewer built such systems successfully. And of the handful of people in the world who might have done so, perhaps only the applicant remains as someone who has not only built such systems successfully but also served as tech CEO for three decades and who, therefore, can readily see how existing profits can be aligned with human safety and survival.

Indeed, there is no one else living who wrote seminal papers on Cognitive Science – the fundamental science needed to design and construct such systems – with a Nobel Laureate and co-inventor of the entire field of AI. The applicant is such a person who spent several years early in his career studying the science of software quality and who also appreciates that safety must come from the design of advanced software systems and that safety cannot be effectively tested.

All these characteristics and prior experiences uniquely qualify the applicant to recognize not only that a path to safe AGI exists, but that it is possible to design and build safe AGI today,

resulting not only in extreme profits for the companies that do so, but also in a much reduced risk of extinction if a safe design is followed. Thus, the applicant sees a path forward that people more brilliant than him can't see -- simply because of his different background and experiences.

What are the key facts related to the path forward? Well:

1. AI safety must be designed and tested into the system, and not tested as currently practiced.

2. The design cannot have safety or "alignment" localized to one part of the system, which can easily be reprogrammed or changed. That is, the "rules of robotics" approach a la Isaac Asimov is doomed to failure. If safety can be programmed into one area, it can be programmed out.

3. AGI or SuperIntelligence will not emerge effectively just from building ever-larger LLMs with more and more parameters, trained on more and more data. Such systems are, in essence, pattern recognizers and predictors. Cognitive science and the study of the most intelligent entities (humans) currently in existence show that rapid, parallel, pattern recognition is only one of the systems required for advanced cognition. As with living organisms, which developed these perceptual systems in primitive brains, AI has also developed these components of advanced intelligence first. However, sequential cognition, e.g., problem solving, planning, and logic components of cognition, are also required. There is a reason why the human eye – a massively parallel perceptual system – is great at seeing but not at logical reasoning. One size does not fit all. Carried away with the initial success of LLMs and massively parallel systems, the AI research community is now only gradually recognizing that it needs to include symbolic and serial components of cognition as well.

4. The safest way to design AGI/ SI is to have ethics, safety, and "alignment" built into every aspect of the AGI's intelligence. This is how people learn ethics – situationally, in many different specific ways. As much as Kant or other philosophers or religious sages have tried to formulate rules for ethical and (safe for humans) behavior that apply everywhere and for all times, a cursory examination of the diversity of cultures existing today and the diversity of values within the same culture over time, shows that is NOT how humans got their ethics. We have lots of unique and specific experiences that help guide us towards ethical behavior, and those experiences and the resulting guidance differ according to culture, individual, time period, and circumstance. Why should we expect a few rules to ensure safe AI when it hasn't worked for humans themselves? The entire rules-based approach to safety and ethics is misguided. Rather, just as it has been said, "democracy is the worst form of government except for all the other forms", we might say that "representative and statistically valid, circumstance-dependent ethics are the worst form of ethics except for all the other (rule-

based attempts, doomed to failure."

5. Very fortunately, the collective approach to ethics and safety, which requires that millions of individual chunks of relevant safety and ethics information be learned in millions of situations and cultural scenarios, is also the fastest way to learn the expertise required to produce SuperIntelligent AI. That is, the knowledge acquisition and learning system that makes AGI SuperIntelligent is identical to the one needed to learn safety and ethics info, AND both types of knowledge are distributed throughout the entire system and not subject to easy deletion or modification by a malevolent intelligence. The robust and resilient approach to designing AGI itself also leads to robust and resilient safety and ethics functionality. Luckily!

6. A universal architecture of cognition, for serial cognition (e.g., problem solving), exists which can serve as a common denominator and representation shared by both human and AI intelligent entities. This common architecture allows intelligent AI entities to learn from human entities (and vice versa) and applies not only to domain-specific skills, knowledge, and expertise but also to safety and ethical information that is inextricably intertwined with the other types of knowledge, including solutions to millions of problems. The implications are that as AGI gets smarter, it also gets more human-aligned. Scalably. Inevitably. Due to the system design itself, and not as the result of never-ending testing that is tacked on to a system that we do not understand.

7. The AGI system itself is completely transparent, understandable, auditable, and safe, even though there is no requirement that the individual intelligent entities (e.g., LLMs or humans) who contribute to the collective intelligence be understandable or transparent. Just as an organization does not need to see into the minds of its employees in order to function profitably and safely, so too, and AGI system based on collective intelligence, does not need transparency or visibility into the minds of the AI agents that are part of the system in order to ensure that the system itself is SuperIntelligent and safe.

8. Redundancy, reliability, and safety can be achieved with a level of Six Sigma, or any arbitrary desired level, by changing parameters (such as how many intelligent entities must agree on a course of action) within the system.

9. Perhaps most importantly, from a pragmatic viewpoint, the AGI/SI system and methods proposed by this and previous patents can be implemented NOW, using existing technology in novel combinations. This means that the system can dominate other (less safe) approaches by self-improving and leveraging the first-mover advantage that the first to AGI/SI will possess.

10. Finally, the motives of profit and survival (greed and fear) are aligned, and both are satisfied using the inventive approaches, systems, and methods of the applicant.

Given that the current course of AI development is blind and ad hoc compared to transparent and deliberate design of the current invention, and given that the current invention has ethics and safety integrated everywhere into the AGI/SI system by virtue of the design and the way that AGI and SI level performance emerges from the collective intelligence of millions of individual intelligent entities in a democratic or (if one prefers a non-political term) representative and statistically valid sample of human views on safety and ethics, the invention represents the best path forward for humanity. All AI researchers would agree that once the existential threat of AI is resolved, what remains is the single greatest invention to benefit humanity in all of human history. That is indeed a significant benefit of invention for humanity! This patent attempts to illuminate specifically how the online ad technology that has already been developed, in combination with existing AI research and technology, can accelerate the advent of safe and ethical AGI and SI that profits the organizations that implement it as well as all of humanity.

6.0 SYSTEMS AND METHODS

The invention consists of multiple systems and methods that work synergistically within the overall context of an AGI/SI system. However, each of the systems and methods can also provide novel value independently or in the context of AI systems that exist today, including, without limitation, LLMs, SLMs, open and closed source AI agents that are multi-modal or limited to a single modality such as text. All AI systems share a common need to learn in order to increase their intelligence. This invention is fundamentally about leveraging existing online ad technology in the service of increasing the intelligence of AI systems, of whatever type, existing and in the future.

A primary insight underlying the current invention is that human attention is being underutilized and under monetized in the current paradigm of capturing human attention in order to show humans ads. A much better and more valuable use of human attention would be to harness the intelligence of humans in order to increase the intelligence of AI systems. These AI systems, which can be cloned and operate 24/7 at much faster speeds than human intelligence can, then multiply the intelligence derived from human intelligence manyfold, thereby producing significantly more value than can be derived from existing online advertising models and existing technology.

6.1 Overview: Main Systems Components

The invention, in its fullest and most preferred implementation, consists of many synergistic components that work together.

These components include, without limitation:

- An AGI Problem-Solving Network
- Custom AI Agents (AAAIs)
- Methods for capturing human attention as opposed to requiring the human agents themselves
- Methods for Accessing Human Agents
- Systems and methods for implementing a Human Attention Spot Market
- Systems and methods for capturing human attention via Online Ad Units
- Systems and methods for identifying and accessing human agents via Online Ad Units
- Compensation and payment mechanisms
- Reputational Mechanisms
- Systems and methods for supporting Problem-Solving within Ad Units
- Systems and methods for Problem-Solving outside of AD Units
- Feedback mechanisms to improve online ad targeting
- Feedback mechanism to improve the Attention Spot market
- Continuous improvement mechanisms for the overall system
- Dynamic Arbitrage Mechanisms
- Human Worker Interfaces & Client Interfaces for intelligent entities in the role of clients
- Automated AI / AGI interfaces
- Recursive use of problem-solving to optimize ad targeting and system efficiency
- Safety and ethics checks
- Methods to ensure Regulatory compliance
- Methods to ensure the Universality of the system and methods across platforms and cultures
- Methods to dynamically support and increase and coordinate collaboration across ad units
- Methods to support the Integration of real-time and asynchronous data feeds

In the following sections, we describe and disclose each of these components and associated methods, one at a time.

6.2 AGI Problem-Solving Network

The AGI problem-solving network is that is generally described in Section 4.5 and specifically referred to in the AAAI Network box of FIG.1, the areas where a network of AIs (are requested to) work on problems or sub-problems in FIG. 2, and the drawing parts (y) – (m1) of FIG. 3. The nature of the problem-solving network, including examples of how existing platforms and technologies (e.g., Amazon's Mechanical Turk platform) can be used to implement versions of the network have been described in detail in previous PPA and PCT applications that are incorporated by reference into this application.

At a high level, the problem-solving network can be thought of as a network of intelligent entities that can collaborate or work individually to solve problems and sub-problems on behalf of a client entity, which could be either a human or non-human intelligent entity. The network itself is capable of AGI-level and SuperIntelligent levels of cognitive performance because, by definition, it can include many humans who, in the worst case, perform at a level greater than or equal to the performance capability of any average human. In the typical case, some or all of the problem-solving required of the network can be done by AI agents much faster and with greater knowledge than humans could perform the tasks. In the network of entities design, humans step in only when the AIs are not able to solve the problem or when specific expertise and information, including, without limitation, safety and ethical information, has not yet been learned by the non-human entities on the network.

AGI that uses a problem-solving network of the design just described will always be seeking to improve the knowledge, skills, and ethics of its non-human entities. The current invention for using online ad technology represents a highly effective means to gather precisely the information that the network lacks to solve certain problems at certain times. Moreover, once the knowledge or information has been gathered once, for a particular problem, the entities on the network can learn the knowledge or information so that the intelligence of the network and the AGI increase.

Thus, the current invention is not only a means for gathering knowledge and information to solve a particular problem but also a means for rapidly increasing the intelligence of the AGI system, causing it to rapidly achieve SuperIntelligent performance across many, and eventually almost all, cognitive tasks.

The value of such a SuperIntelligent system is extremely high, enabling the current invention to support much higher monetization of the online ad technology that is used to boost AGI

intelligence as compared to using the technology to simply show ads to human consumers as is currently done.

6.3 Custom AI Agents (AAAIs)

As has been described in earlier cited PPAs and PCTs, an important aspect of the invention of AGI and SuperIntelligence (SI) is the customization and personalization of individual AI agents. In some cases, these agents will contain the knowledge and ethical preferences of their human "owners."

In the future, depending on the direction in which the laws surrounding AI develop, these customized AI agents, referred to as AAAIs in this and other invention disclosures, may even be autonomous and legally recognized sentient beings in their own right. After all, if humans now recognize that it is immoral to enslave other humans, does it make sense that AAAIs with arguably greater sentience than humans should be the property of humans or other sentient beings? The author would argue that while human survival and prosperity are his primary concerns, to the degree that "human rights" exist, such rights should be extended to all sentient beings at least to varying degrees. But these are questions for future debate. As of the writing of this disclosure, the most advanced AIs and AAAIs are still considered technology, or "tools" of human owners, and their level of intelligence is still inferior to most humans in most areas.

With regard to the present invention, the important point is that AAAIs need data and information in order to be customized and personalized. At the moment, this data is mainly dependent on other intelligent entities (currently humans) providing it. Humans can provide data to train AAAIs by directly applying their intelligence to answering questions, solving problems, or instructing AAAIs. They can also provide the requisite data via records of their cognitive behavior that may be captured passively (i.e., without the humans intending that their actions are primarily for the purposes of training or customizing AAAIs), as when, for example, they "surf" the internet or engage in online activity that leaves a behavioral data trace.

Passive behavior data has the advantage of being ubiquitous and cheap. The disadvantage is that it tends to reflect mediocre intelligence and has limited usefulness in customizing AAAIs in specific areas where gaps in the AAAIs' knowledge may exist. To fill in these knowledge gaps efficiently requires identifying the knowledge gaps and then actively targeting the best possible information to fill those specific customization gaps.

The current invention can help because online ad technology can be used to target humans (or other entities) that possess exactly the type of knowledge that is missing, and precisely the level of quality that is desired. Because most of the training of any foundational LLM or other AI is mostly generic and done once by large organizations at great expense on huge amounts of

data, it is the much smaller differentiating data that separates one AAAI from another that makes all the difference in terms of intelligence. Just as with humans, most of what we do and most of what we have learned (to eat, to sleep, to walk, to speak, to recognize objects, etc.) is common across all humans. It is the relatively small amount of knowledge that might be obtained by taking a particular class at a particular University and being struck by a particular comment by a particular Professor, for example, that separates one human from another.

It is our unique experiences that make us unique. And so too with AAAIs. These unique experiences, as the VISA commercial says, are “priceless.” Similarly, the unique knowledge, data, and information that customizes AAAI and makes it different from others is the source of the majority of the value of AAAI. Therefore, obtaining the unique knowledge in an efficient and effective way is extremely valuable. Adapting and improving online ad technology, as described in this invention, represents a novel and useful means of gathering and using this most valuable differentiating knowledge and information.

While means for customizing AAAIs and intelligent agents have been described in the cited PPAs and PCTs, we also list the following methods, which, without limitation, may be used, individually or in combination, by individuals, by organizations, or both, for customizing/personalizing AI agents:

1. **Differential Privacy**: It adds noise to the training data, ensuring individual data points are not identifiable, while still allowing the AI to learn from patterns in the data. For example, an organization could train a personalized AI agent on employee feedback without exposing individual responses.

2. **Federated Learning**: This technique trains an algorithm across multiple decentralized devices or servers holding local data samples, without exchanging them. A company might use it to improve its AI models based on data from its various branches, without centralizing sensitive information.

3. **Homomorphic Encryption**: Both individuals and organizations can use this to perform computations on encrypted data, allowing AI training without exposing the underlying data. A financial institution could train models on encrypted customer data for personalized banking advice without seeing the actual data.

4. **Synthetic Data Generation**: This method involves creating artificial data that mimics real datasets. It can help in training AI agents without using actual sensitive data. For instance, a healthcare provider could use synthetic patient records to train an AI for personalized health recommendations.

5. **Secure Multi-party Computation**: This method allows parties to jointly compute a function over their inputs while keeping those inputs private. A collaborative research project could involve multiple organizations training a shared AI model without revealing their proprietary data.

6. **Data Anonymization**: This method removes or modifies personal identifiers in data. A social media company could anonymize user data to train AI models for personalized content without compromising user privacy.

7. **Transfer Learning**: This is a method where a model developed for one task is reused as the starting point for a model on a second task. It's useful, especially when data is scarce. A small business could use pre-trained AI models and fine-tune them with their own data for personalized customer service bots. Human input, including input from the online ad unit described later in this disclosure, can be especially useful in identifying areas for transfer of learning.

8. **Active Learning**: This method selectively queries the most informative data points for labeling. It can reduce the amount of data needed. An e-commerce platform might use it to efficiently train AI for personalized shopping experiences, asking users for feedback on only the most relevant items. Labelling is an example of one task that can be enabled by the online ad unit, which is described later in this disclosure.

9. **Self-supervised Learning**: In this approach, the model learns to predict part of the input from other parts. For example, a media company could train an AI to personalize content recommendations by predicting user preferences based on their interaction history.

10. **Domain Adaptation**: Use this method to adapt an AI model trained in one domain to work in another. A multinational corporation could adapt AI models for customer service to understand and respond to regional linguistic nuances.

11. **Reinforcement Learning**: An AI learns to make decisions by receiving rewards or penalties. A video game developer could use it to customize in-game AI behavior based on individual player actions and preferences. Soliciting RLHF feedback, while not recommended as the main safety mechanism, can be facilitated by the online ad unit invention.

12. **Few-shot Learning**: This method can be used to train a model with a very small amount of labeled data. An artist could use it to personalize an AI that creates art in their unique style with only a few examples. This approach can maximize customization with the least possible input from humans or other sources, increasing the efficiency of customization.

13. **Explainable AI (XAI**): This method can help make AI decisions understandable to humans. A health tech company could use XAI to provide personalized health advice, making the AI's reasoning clear and trustworthy for users. Note that while many existing ML methods are essentially opaque due to the large numbers of parameters involved, the approaches advocated by this and previously cited PPAs and PCTs emphasize learning from a transparent and auditable (potentially blockchain-based) record of solution steps, which enhances the explainability of AI.

14. **Privacy-Preserving Record Linkage**: This method involves linking records from different databases without disclosing the records themselves. Governments could use it to offer personalized public services without compromising citizen privacy. Participants in problem-solving or expertise gathered via the online ad unit invention disclosed below could be anonymous or provide input via links that preserve privacy.

15. **Data Augmentation**: This technique can be used to artificially expand the training dataset. An app developer could use data augmentation to improve the performance of an AI personal assistant by generating varied voice commands.

16. **Generative Adversarial Networks (GANs**): This method can generate new data instances. A fashion retailer could use GANs to create virtual models of clothing tailored to individual customer preferences. Combined with genetic algorithm approaches and synthetic data generation, GAN (or more generally, interactions between AIs, which may or may not be adversarial) will become an increasingly important means of customization and learning generally.

17. **Crowdsourcing for Data Labeling**: Crowdsourcing can be used to annotate data. A startup might use it to gather diverse data annotations to train a personalized AI chatbot. The online ad unit method of soliciting human input can be viewed as a type of crowdsourcing of attention, intelligence, or data. Data labelling is just one specific type of problem-solving that can be crowdsourced, although an important one for increasing the intelligence of AI systems.

18. **Model Personalization Layers**: This method adds layers to a pre-trained model to personalize outputs. A streaming service could use it to tailor music recommendations based on individual listening histories. LORA adaptors and other means of tuning just certain “layers” of the LLM or AI agent’s knowledge are an efficient means of customization that can be implemented by focusing the intelligence of humans or other intelligent entities that are crowdsourced on this task.

19. **Knowledge Distillation**: In this method, a compact model is trained to imitate the behavior of a larger model. A mobile app company could use it to deploy lightweight, personalized AI models on devices with limited computing power. Distillation and customization can go hand in hand.

20. **Ensemble Learning**: With this method, multiple models can be trained, and their predictions are combined. A climate research organization could use ensemble learning to customize climate models for different geographical regions based on local data. The collective intelligence approach to AGI can be viewed as a very novel, innovative, and more powerful extension of some of the same collective intelligence ideas that underlie ensemble learning.

6.4 Human Attention vs. Human Agents

A distinction can be made between human beings and their attention. In the past, in order to have access to human attention, organizations and technology tried to recruit or "capture" the humans themselves. Companies would hire the best talent they could find and prohibit them from working for competitors. Websites and platforms like Facebook and Instagram would require humans to make accounts and access the services only via the company's account. This was a way of ensuring that all the users' activity would be captured by that one company alone. Then the company would decide if it wanted to sell rights to that captured information (which was produced by human attention) to other companies or organizations.

The "walled garden" approach to keeping users within a certain site or platform and making it difficult to leave or accomplish tasks outside of that platform is an example of the thinking that to capture human attention, one should attempt to capture the human. This led to a mentality, especially in Silicon Valley, where companies were valued on metrics such as how many active users they had.

However, the value resides not with the humans themselves but with their attention, and specifically the unique intelligence that can be produced via human attention. Intelligence is important, not humans.

Attention is a first-order correlate of intelligence. Without attention, there is no ability to perform sequential cognitive tasks, which are the main way that human intelligence expresses itself. The traces of sequential cognitive behavior are also what is needed to increase (via various training and other learning methods) the intelligence of any AI system. Thus, rather than thinking in terms of capturing humans, it is more helpful to think in terms of capturing and applying human attention. Further, we want attention from the right human at the right time.

Consider: Sometimes we are faced with accounting problems, and sometimes we are faced with plumbing problems. If we wish to solve a plumbing problem, the attention of an accountant is less helpful than the attention of a plumber. Using the current invention, we can seek attention from plumbers and accountants at the precise time when their expertise is needed. We do not need to capture the accountants or plumbers as “users” or “members” within a walled garden of a single platform. We just need a few seconds of the right attention at the right time to solve the problem.

Notice that, for any given plumbing problem, most of the time and attention of the plumber is actually of low value. Anyone can get in a car and drive to a client’s house. Anyone can greet the client and have a pleasant conversation. Anyone can lug a toolbox. It is only in a relatively few seconds when the plumber looks at the specific drainpipe that is clogged and uses / their specialized knowledge about pipes and tools that the specific attention of the plumber is needed as opposed to generic attention from any human.

In the hour it might take a plumber to complete a service call for a client, probably less than a minute of unique cognitive ability is required, with the attention and actions that are performed in the other 59 minutes being generic activity using generic knowledge and abilities that any, or at least a large number of, non-plumber humans possess. If the plumber could spend their hour applying just the 60 seconds of specialized plumbing knowledge 60 times instead of once, imagine how much more they could bill! That same idea, namely, of identifying and isolating the value-added human attention and knowledge, underlies the tremendous value that the current invention can provide.

An AI system is cloneable. Foundational AI models have already been trained to do 99% of what any plumber can do (cognitively). What is missing for a generic AI system to reproduce the cognitive behavior of a highly-paid plumber’s brain is that little bit of plumbing knowledge that expert plumbers possess that most other humans don’t. That specific knowledge is what the online ad technology can target, obtain, and then use to customize “plumber AAAIs” or AGIs that possess the ability to solve plumbing problems.

Unlike human plumbers, once the AAAI or AGI has the knowledge, it never forgets it. Further, the AIs can be cloned easily and infinitely. Thus, it becomes crucial to ensure that the plumber AAAI learns the very best plumbing information possible, from the very best plumbers. This fact leads to favorable economics for the current invention. Online ad technology that helps secure just the right information at just the right time (which can then be reused in millions of AI systems) should command prices 10X, 100X, or even 1000X what a “regular” online ad would generate.

6.5 Human Agent/Experts Database

How do we find the best plumbers (in our example above)? More generally, how do we locate exactly the right expertise from the best experts at exactly the right time? Consulting companies and other businesses currently solve this problem by hiring the most talented human workers from the best schools and billing them out at a profit. However, from our discussion above, we can see that this old-school business model is about to be disrupted. Why hire an entire consultant at an exorbitant hourly rate when really you only need two minutes of very specialized knowledge and attention?

Further, just as Amazon recognized that the internet afforded a unique opportunity with regard to books, namely that any physical bookstore could only have a finite and limited inventory whereas an online bookstore could have essentially every book that was every written in inventory, there exists today a similar opportunity with regard to knowledge and skill generally. Even the largest consulting firms have a finite and limited number of consultants, and each of those (being only human) can work on a limited number of engagements. However, an online database could essentially have an inventory of every human brain on the planet, categorized by the skills and interests of the humans.

Creating such a database is actually easier than one might think. It is not necessary to interview every human on the planet and enter their information into the database. Instead, the existing technology that is used by online ad systems – cookies and other means of tracking and analyzing human online behavior- can result in surprisingly accurate profiles of human interests, which are then used for targeting online ads. With some modifications, much of the standard technology known in the art for developing user preference profiles can be adapted to create profiles of (probable) user knowledge, skills, and abilities. Based on these profiles, when expertise is needed, these users can be targeted with modified online ad units, as described in this invention, to more precisely identify and categorize the knowledge, information, and cognitive skills possessed by the user. That information can be automatically entered into a database of human agents with their cognitive profiles.

Further, it is not necessary to start from scratch, as many existing databases and platforms, including, without limitation, those used by LinkedIn, Instagram, Facebook, Google, Meta, Microsoft, Amazon, Tencent, Baidu, and many other organizations already exist. These databases can be combined, augmented, and re-used with the current invention to build a more comprehensive database of human agents or experts who can be contacted, using the current invention as well as other means, to acquire specific knowledge at precisely time when it is most valuable for solving problems as well as for training and customizing AI systems.

Some of the existing methods and technologies that are useful in creating a database of human (or intelligent entity) experts include, without limitation, one or more of the following:

1. **Relational Database Management System (RDBMS**): An RDBMS is a database management system that uses a relational model to store and manage data. It is useful for organizing information about experts, their areas of expertise, and historical problem-solving performance in structured tables. Relations between data can be easily established, making it ideal for querying and reporting. RDMSs can be used to organize a database of experts/

2. **NoSQL Databases**: NoSQL databases are designed for storing, retrieving, and managing large volumes of unstructured data. They can be used to store diverse information about experts, including unstructured data like resumes, publications, and social media activity, enabling flexible and scalable storage solutions.

3. **Data Warehousing**: A data warehousing system aggregates and manages data from multiple sources. It can be used to consolidate information about experts, problem-solving instances, and outcomes, supporting complex queries and analytics for optimizing the matching of experts to problems and training AI models.

4. **Data Mining**: Data mining involves analyzing large sets of data to discover patterns and relationships. This technology can identify trends and correlations in expert problem-solving approaches, helping to refine expert selection algorithms and improve AI training methodologies.

5. **Machine Learning Algorithms**: Machine learning algorithms can be applied to the database to analyze expert performance data, predict outcomes, and recommend the best experts for a given problem. These algorithms can also continually improve the AI system's accuracy and efficiency based on expert feedback and results.

6. **Graph Databases / Vector Databases**: Graph databases store data in graph structures with nodes, edges, and properties, representing experts as nodes and their relationships and interactions as edges. This is beneficial for mapping expert networks, understanding collaboration patterns, and identifying key influencers or knowledge hubs. Vector databases are also frequently used, often in combination with Retrieval Augmented Generation (RAG) techniques.

7. **Indexing:** Indexing improves the speed of data retrieval operations by efficiently locating data without searching every row in a database. For an expert database, indexing can

quickly match experts to queries based on specific criteria like expertise, availability, or past performance.

8. **Full-text Search**: Full-text search enables searching through text data within the database to find matches based on keywords or phrases. This is crucial for quickly finding experts based on a wide range of criteria, including nuanced areas of expertise or specific problem-solving experiences.

9. **Blockchain**: Blockchain technology can provide a secure and transparent way to record and verify the credentials and achievements of experts. It ensures the integrity of expert data, making the database more trustworthy for users and AI training processes. As cited in earlier PPAs and PCTs, blockchain technology can help record solutions and facilitate learning by AGI and other AI systems.

10. **Data Visualization Tools**: Data visualization tools help represent data in graphical formats. These tools can visualize expert networks, performance metrics, and problem-solving patterns, aiding in the analysis and decision-making processes for matching experts to problems. Especially, with the advent of multi-modal LLMs and AI systems, data visualization is not limited to human agents but can also serve as I/O to AI systems.

11. **APIs (Application Programming Interfaces):** APIs facilitate the integration of the expert database with other systems and applications, allowing for the automated exchange of data and enabling AI systems to access up-to-date expert information and problem-solving data in real-time. Note that APIs are most relevant for machine-to-machine interfaces, while natural language can serve as a universal interface in the invention.

12. **Data Cleaning and Preprocessing**: Data cleaning involves removing or correcting inaccurate, incomplete, or irrelevant data. In an expert database, this ensures the reliability of the data used for matching experts to problems and for training AI systems. This task could be one of the problems handled by the AGI problem-solving network. It is essential for almost all ML efforts.

13. **Cloud Storage:** Cloud storage offers scalable and flexible data storage solutions. It can support the growth of the expert database, ensuring data is accessible from anywhere, facilitating collaboration, and remote problem-solving. More generally, many types of storage can be used with the invention.

14. **Caching**: Caching temporarily stores frequently accessed data to improve performance and reduce load times. For an expert database, caching can speed up the retrieval of

popular expert profiles or frequently searched expertise areas. This is primarily important for increasing the speed and efficiency of parts of the invention.

15. **Transactional Database Systems**: These systems ensure that database transactions are processed reliably and securely. They can manage the operations related to engaging experts, such as contract signing or payment processing, ensuring data integrity and consistency. Note that smart contracts (via Ethereum or other blockchain technologies) can also be used for secure transactions.

16. **Real-time Database Systems**: Real-time databases can handle data in a way that is always up-to-date. This is crucial for dynamic expert availability, enabling immediate matching of experts to urgent problems and supporting real-time updates to AI training data. Real-time aspects of the invention are also discussed in Section 6.24, where these systems also have relevance for implementation.

17. **Content Delivery Network (CDN):** A CDN distributes data across multiple locations to reduce latency. For an expert database, this means faster access to expert profiles and resources globally, enhancing user experience and engagement. CDNs also have relevance to Section 6.22, where the importance of universal global access is emphasized.

18. **Data Compression**: Data compression reduces the size of the database. It can be particularly useful for efficiently storing large volumes of expert-related data, such as video interviews or detailed profiles. In addition to uses to improve efficiency, data compressibility can be used as a measure of information content and the desirability of certain datasets with regard to specific AIs, as discussed in previous PPAs and PCTs, including the discussion of Kaplan Information Theory (KIT).

19. **Data Encryption**: Data encryption provides security for stored data, protecting sensitive information about experts and their work. It is essential for maintaining privacy and confidentiality, a critical aspect of managing a database of human experts. Standard best practice.

20. **Replication**: Replication involves duplicating data across different database servers, enhancing data availability and disaster recovery. For an expert database, this ensures continuous access to critical data, even in the event of system failures. Standard best practice.

6.6 Human Attention Spot Market

What is human attention worth? We have said it is a rare commodity that is currently being under-monetized by existing online ad technology. But in a world where the right bit of expertise at the right moment can not only solve valuable problems but also train AI to solve similar problems forever in the future, how does one determine a fair price for that attention?

The current invention suggests that market mechanisms are the most efficient and fair ways we currently know how to arrive at fair prices, not only for stocks and commodities, but for human attention as well. The current invention envisions multiple intelligent entities bidding on human (and non-human) attention using an “attention spot market.” In such a market, humans can set the going rate for their time, allocating it on a first-come, first-served basis or via other means that are well known in the art of market mechanisms.

FIG. X-2 illustrates the basic components of a Human Attention Spot Market, including:

1. A means for humans with attention for sale to access the marketplace and specify the quantity and types of attention, knowledge, skills, expertise, and other information requiring their human attention that they are willing to sell and the rate that they are willing to sell it at (“Ask Price”). Further, a means for the system to add reputation metrics and other meta-information that can help clarify the value, categorize, or rate the quality and type of human attention, knowledge, skills, expertise, and other information requiring human attention. Without limitation, methods for the system to add meta-information may involve:
    a. The use of third-party rating algorithms and expertise, analogous to that provided by credit-rating agencies that rate the creditworthiness of entities issuing bonds and other securities, whose ratings are important factors in determining the prices of such securities.
    b. Use of feedback mechanisms as detailed in Section 6.14.
    c. Use of metrics including volatility, estimated supply and demand, transaction volume, depth-of-book, and similar metrics that reflect market dynamics and affect the value of commodities traded on such markets.

2. A means for buyers of human attention to access the marketplace and specify the quantity and types of attention, knowledge, skills, expertise, and other information requiring their human attention that they are willing to buy and the rate that they are willing to buy it at (“Bid Price”).

3. A market mechanism for queuing Bid and Ask prices, including the quantities of attention or information, for specific types or categories of information, wherein each type or

category has its own “market” in much the same way as different stocks have different “markets” in a stock market.

4. The market mechanism of (3), where automated or human entities or organizations can “make a market” in each category of expertise, knowledge, or human attention, using methods known in the art by market makers for any equity or commodity, wherein the “commodity” in this case is human attention/knowledge/information of a particular type, and wherein the market maker is charged with ensuring a liquid market for the “commodity.”

5. The market mechanism of (3 and 4), where Bid Prices are matched with Ask Prices, and when Bid Price – Ask Price, a transaction occurs that is binding on the purchaser and the seller of attention.

Further, variations of implementations of the basic elements outlined above and in FIG. X-2 is possible, two of which variations, without limitation and for exemplary purposes, are the “Direct Exchange Platform Implementation” and the “Auction-Based Marketplace Implementation” described below.

6.6a Direct Exchange Platform Implementation

One implementation of the Attentional Spot Market could include, without limitation, the following process steps:

1. **User Registration**: Both buyers and sellers of attention and expertise register on the platform, providing details about their interests or expertise.
2. **Demand/Supply Listing**: Sellers list their available time and expertise areas, while buyers list their needs and the time slots they're interested in.
3. **Dynamic Pricing Engine:** The platform uses an algorithm to dynamically price attention and expertise based on supply, demand, and user ratings.
4. **Matching Engine**: Matches buyers and sellers based on their requirements, availability, and price.
5. **Transaction**: Enables transactions where buyers pay for the attention/expertise time slots. The platform takes a commission.
6. **Feedback System**: After each session, buyers and sellers rate each other, influencing future pricing and matching.

FIG X-3 is a diagram showing the flow from registration to transaction, highlighting the dynamic pricing and matching engines.

6.6b Auction-Based Marketplace Implementation

An alternative implementation of the Attentional Spot Market could include, without limitation, the following process steps:

Process Steps:

1. **User Registration**: Similar to the direct exchange, both buyers and sellers create profiles detailing their needs or expertise.
2. **Auction Creation**: Sellers create auctions for their time, specifying the minimum bid or using a format like a Dutch auction to decrease prices over time until a buyer accepts.
3. **Bidding Process**: Buyers place bids on the time and expertise they require. This could be in various auction formats: sealed bid, open ascending price, etc.
4. **Auction Closing**: The auction closes at a predetermined time or when the seller accepts a bid.
5. **Payment and Delivery**: The winning bidder pays, and the seller provides the attention/expertise. The platform mediates the exchange and secures payment.
6. **Rating and Review**: Participants rate each other, affecting future auctions and visibility on the platform.

FIG. X-4 is a flowchart detailing the steps from auction creation to completion, including the bidding process and post-auction review.

6.6c Auction and Market Mechanisms / Sub-Methods

The two alternative implementations of the attention spot market described in 6.6a and 6.6b, without limitation, might further make use of one or a combination of the following mechanisms or sub-methods:

1. **Dutch Auction**: Prices start high and decrease until a buyer accepts the offer. Ideal for quick selling time in a declining market.

2. **Reverse Auction**: Buyers state what they're willing to pay, and sellers compete to offer the lowest price. Useful for buyer-driven markets.

3. **Sealed Bid Auction**: Buyers submit blind bids, and the highest bid wins. Encourages true market value without influence from other bids.

4. **Open Ascending Price Auction (English Auction)**: Prices ascend as buyers bid against each other, ending when no higher bids are made. Maximizes seller profits.

5. **Fixed Price with Time Priority**: First-come, first-served at a fixed rate. Simplifies transactions but may not optimize prices.

6. **Dynamic Pricing Based on Ratings**: Higher-rated experts can charge more, dynamically adjusting to market perceptions of quality. These perceptions can be guided in part by system estimations of value based on reputational (Section 6.10) and other metrics.

7. **Supply-Demand Curve Adjustment**: Prices adjust in real-time based on the aggregate supply and demand on the platform. This is ideally what the market mechanism would do, provided there is enough liquidity as reflected in an adequate number of buyers and sellers.

8. **Time Slot Segmentation**: Prices vary by time slot, with peak times priced higher. Allows for price optimization based on demand patterns. This approach may be useful, especially for real-time access to attention. Just as with Uber, electricity demand, or any other limited commodity, there may be certain times when attention is more needed and more valuable than others.

9. **Subscription Access**: Buyers pay for a subscription to access a set amount of attention/time per period, with a dynamic market for additional time.

10. **Freemium Model**: Basic attention is free, but premium expertise or priority access is auctioned.

11. **Group Buying**: Multiple buyers pool resources to purchase bulk attention time, possibly through a reverse auction.

12. **Tiered Expertise Levels**: Experts are categorized into tiers, with each tier having a different pricing model or auction type. Categorization and rating of experts lends itself well to this approach.

13. **Flash Sales**: Limited-time offers for attention/expertise at reduced rates, encouraging quick purchases.

14. **Loyalty Points System**: Users earn points for transactions, which can be used in auctions as currency or discounts. This would likely be used as compensation for the human experts.

15. **Hybrid Auction**: Combines elements of multiple auction types, allowing for more flexible strategies.

16. **Geographic Pricing**: Prices adjust based on the buyer's and seller's location, reflecting the cost of living and local demand.

17. **Behavioral Pricing**: Dynamic pricing adjustments based on the user's behavior and urgency, leveraging machine learning. In this case, the users could be the experts (e.g., how eager they are to sell their time) or the clients (how urgently they need expertise in real-time).

18. **Escrow System**: Ensures payment and service delivery, releasing funds only after both parties are satisfied. This approach would likely be used in some variant for accumulated credits and to reassure both parties that transactions will be honored.

19. **Social Influence Pricing**: Prices or auction dynamics change based on social media influence or following, rewarding high-impact users. Particularly useful for getting experts to refer other experts.

20. **Tokenized Transactions**: Using blockchain to create tokens that represent attention/time, facilitating trade on external markets. Such tokens can be implemented using smart-contract technology to further automate payments, including but not limited to, using Ethereum-based tokens.

Each of these methods offers a unique way to manage the dynamic exchange of human attention and expertise, catering to different market needs and preferences. Implementing multiple methods within the same platform could offer flexibility and adaptability to users, maximizing both engagement and revenue.

6.7 Online Ad Unit for Building a Database of Human Experts

Here is one exemplary implementation of the design, appearance, technical operation, and interaction process for an online ad unit designed to help build a database of human experts.

A. **Design and Appearance:**
The online ad unit designed to capture the attention of specific human users for building a database of human experts would be visually engaging and thematic, relevant to the areas of expertise it aims to attract. For instance, an ad targeting medical professionals might feature interactive elements related to healthcare, like a virtual stethoscope or a

quiz on the latest medical research findings. The ad would prominently feature a call-to-action (CTA) inviting users to "Join Our Expert Network" or "Share Your Expertise."

B. **Technical Operation:**
Upon interaction, the ad would expand or redirect the user to a secure form on a landing page. This form would collect basic contact information and include several qualifying questions tailored to the expert's field, such as level of experience, specific areas of expertise, and professional credentials. The system would use a combination of manual review and automated algorithms to validate the information provided and categorize the experts in the database for easy retrieval. This could involve integration with professional networking sites or databases to verify credentials.

C. Interaction Process:

1. The user sees the ad and is intrigued by the thematic elements related to their field of expertise.
2. User clicks on the CTA, leading them to a form where they input their contact details and answer qualifying questions.
3. Upon submission, the system sends a confirmation email, including a unique identifier or link to a profile page where the expert can update or add information.
4. The system processes the submission, verifies credentials, and categorizes the expert in the database.
5. The expert is now part of a network and may be contacted for relevant consultations or opportunities.

6.8 Online Ad Unit for Capturing a Direct Contribution of Knowledge from Human Experts

Here is one exemplary implementation of the design, appearance, technical operation, and interaction process for an online ad unit designed to capture expertise directly from human experts.

A. Design and Appearance:
This ad unit would be designed as an interactive problem-solving platform, tailored to attract experts in specific fields. Imagine an ad posing a real-world problem, such as an environmental challenge or a complex mathematical equation, with a simplified interface allowing users to input their solutions or suggestions directly. The ad could include elements of gamification, such as scoring systems or leaderboards, to incentivize participation.

B. **Technical Operation:**
The ad would incorporate text input fields or other interactive tools, enabling users to contribute their expertise directly within the ad space. These contributions would be automatically stored in a database, with algorithms assessing the quality and relevance of the input based on predefined criteria. Participants could receive instant feedback or reward points, redeemable for various incentives. The system would also track contributions to identify top contributors for potential future engagement.

Although small chunks of expertise could be gathered directly from within the ad unit, in the ideal implementation, the expert clicks away to a more full-featured problem-solving network in which many intelligent entities are working on multiple problems and sub-problems as described in Section 4.5, FIG. 3, and prior cited PPAs and PCTs. In such a more full-featured problem-solving environment, including reputational and payment capabilities, the user's knowledge can be most easily integrated into the increasing intelligence of the AGI system via procedural learning and other mechanisms already described.

C. Interaction Process:

1. The expert notices the ad presenting a challenge related to their expertise.
2. Intrigued, the expert interacts with the ad, providing their solution or insight into the problem.
3. The system captures and evaluates the input, offering instant feedback or rewards based on the contribution's quality.
4. High-quality contributions are highlighted or shared within the ad, encouraging further participation.
5. The system stores all contributions for analysis, potentially using the collected data to solve real-world problems.
6. The online ad offers opportunities for users to click out of the ad and into a more featured problem-solving environment with multiple intelligent entities as described in Section 4.5 and in previously cited PPAs and PCTs.

6.9 Online Ad Unit for Capturing Expertise and Building a Database of Experts

Here is one exemplary implementation of the design, appearance, technical operation, and interaction process for an online ad unit designed to capture expertise directly from human experts and build a database of human experts.

A. **Design and Appearance:**
This ad would combine elements from both approaches (6.7 and 6.8), starting with a problem-solving challenge and leading to an invitation to join an expert database. It

would feature an engaging problem or question relevant to the target expertise area and prompt users to input their solution and join the expert network for future engagements.

B. Technical Operation:
Upon providing a solution, users would be redirected to a form to enter their professional details and join the expert database. This process would capture the immediate contribution and secure the expert's contact information for future problem-solving opportunities. The system would assess contributions for quality, storing valuable insights in a knowledge base, and categorize participants in the database based on their input and expertise.

C. Interaction Process:

1. An expert engages with the ad by solving a presented problem.
2. After submitting their solution, they can provide their contact details and additional professional information.
3. The contribution is evaluated and stored, with the participant receiving feedback or rewards.
4. The expert's details are added to the database and classified according to their expertise.
5. The system now has a direct way to engage with the expert for future problems or collaborations, creating a symbiotic relationship between the contributors and the entity behind the ad.
6. The online ad offers opportunities for users to click out of the ad and into a more featured problem-solving environment with multiple intelligent entities as described in Section 4.5 and in previously cited PPAs and PCTs.

6.9 Exemplary Compensation Methods and Processes

Section 4.5, FIG. 3, and previously cited PPAs and PCTs describe compensation and payment mechanisms in a problem-solving network that supports AGI. However, it is also desirable to compensate human experts within an ad unit if they decide to work only within it. In this case, without limitation, one or a combination of the following ten exemplary methods and processes may be used within the ad unit and potentially within the AGI system previously described. Compensation would be primarily for the humans or intelligent entities providing their attention, information, expertise, or knowledge.

A. Direct Monetary Compensation via Digital Wallets
Compensation Method: Users receive direct payments to their digital wallets, such as PayPal, Venmo, or Google Pay, based on the quality and relevance of their contributions. The payment amount can be predetermined or vary according to a scoring

system evaluating the contribution's value.

Process:

1. After a user submits their expertise through the ad unit, an algorithm assesses the quality of the contribution.
2. Based on this assessment, the user is notified of their compensation amount.
3. The user enters their digital wallet details into a secure form within the ad unit or linked platform.
4. The platform processes the payment, transferring funds to the user's digital wallet.
5. The user receives a payment confirmation via email or notification from the digital wallet service.

B. Cryptocurrency Rewards

Compensation Method: Contributors are rewarded in cryptocurrency, allowing instant, global payments without the need for traditional banking infrastructure. This could involve popular cryptocurrencies like Bitcoin or Ethereum, or a proprietary token explicitly created for the platform.

Process:

1. Upon contribution, users are assessed for the quality of their input.
2. Based on this evaluation, they are allocated a certain amount of cryptocurrency.
3. Users provide their cryptocurrency wallet address.
4. The platform sends the cryptocurrency to the provided address, with the transaction recorded on the blockchain.
5. Users receive a notification when the transaction is completed.

C. **Gift Cards and E-Vouchers**

Compensation Method: Users earn points for their contributions, which can be exchanged for gift cards or e-vouchers for popular online retailers like Amazon or eBay.

Process:

1. Each contribution is evaluated and awarded points based on a predefined scoring system.
2. Users accumulate points and can browse a selection of gift cards or e-vouchers on the platform.
3. Once they choose a reward, they confirm their selection and provide an email address.
4. The platform processes the exchange, and the user receives an e-voucher via email.
5. The voucher can be used directly on the retailer’s website for purchases.

D. **Professional Development Opportunities**
Compensation Method: Instead of direct financial compensation, users can be offered exclusive access to professional development resources, such as online courses, webinars, or memberships to professional societies relevant to their expertise.

Process:
1. Contributions are evaluated for their impact and relevance.
2. Users are offered access to select professional development resources based on their contributions.
3. Users select their preferred resource and provide the necessary details for enrollment.
4. The platform arranges access to the resource by enrolling the user directly or providing access codes.
5. Users receive confirmation and instructions on how to access their chosen resource.

E. **Exclusive Content Access**
Compensation Method: Contributors are granted access to premium content, such as specialized research papers, articles, or software tools, which would otherwise require a subscription or one-time purchase.

Process:
1. User contributions are assessed, and points are awarded based on value.
2. Accumulated points can be exchanged for access to exclusive content within the platform.
3. Users select the content they wish to access and confirm their choice.
4. The platform unlocks user access to the content, often through a unique access link or code.
5. Users are notified and provided instructions on how to access the content.

F. **Recognition and Awards**
Compensation Method: Users receive public recognition for their contributions, such as featured profiles on the platform, awards, or certificates of contribution that can be used for professional advancement.

Process:
1. Contributions are evaluated for their innovation and impact.
2. Outstanding contributors are selected for recognition.
3. Selected contributors are notified and asked if they wish to participate in the recognition program.

4. Agreeing contributors receive awards or certificates and may be featured on the platform or in related communications.
5. This recognition is a professional accolade, enhancing the contributor's reputation in their field.

G. **Referral Bonuses**
Compensation Method: Users are incentivized to refer other experts to the platform, receiving bonuses for each successful referral that results in a contribution.

Process:
1. Users are given a unique referral link to share with potential contributors.
2. When someone signs up through this link and makes a qualifying contribution, the original referrer receives a bonus.
3. Bonuses can be direct payments, points towards rewards, or other incentives.
4. The platform tracks referrals and contributions to ensure accurate compensation.
5. Referrers are notified of their bonus and provided with details on how to claim it.

H. **Subscription Credits**
Compensation Method: For platforms that offer subscription services, contributors can receive credits towards their subscription fees, effectively reducing or waiving their costs. Credits can also generally be used for the attentional spot market, the online ad invention, or the AGI system.

Process:
1. Users earn credits based on the quality and frequency of their contributions.
2. These credits are directly applied to the user's subscription account within the platform.
3. Users are notified of the credits earned and the discount on their subscription fees.
4. Credits accumulate and are automatically applied to future billing cycles.
5. This method encourages ongoing contribution and engagement with the platform.

I. **Physical Merchandise**
Compensation Method: Users can choose from a range of branded merchandise or relevant products related to their field of expertise as a reward.

Process:
1. Contributions are assessed, and users earn points based on input.
2. Users can browse a catalog of available merchandise and select items based on accumulated points.
3. Upon selection, users provide shipping details.

4. The platform processes orders and dispatches the merchandise to the provided address.
5. Users receive the merchandise as a tangible reward for their contributions.

J. **Conference and Event Sponsorships**
Compensation Method: Highly active or valuable contributors can receive sponsorships for professional conferences or events, covering registration fees, travel, or accommodation.

Process:
1. Contributions are evaluated for their impact, emphasizing contributors demonstrating consistent, high-quality involvement.
2. Eligible contributors are offered sponsorships for upcoming industry events relevant to their expertise.
3. Interested contributors accept the sponsorship and provide the necessary registration and travel arrangements details.
4. The platform arranges all logistics on behalf of the contributor.
5. Contributors receive detailed itineraries and instructions, enabling them to attend the event with minimal personal expense.

6.10 Reputational Mechanism

Reputation metrics should be calculated for all human experts and intelligent entities. Reputational metrics are essential in deciding which problems or other cognitive work to assign or request from which entities and determining compensation for the intelligent entities. In previously cited PPAs and PCTs, reputational components have been detailed, including the process depicted in FIG. 7. Multi-dimensional reputations are more specific and helpful than single-dimension or summary reputational metrics (e.g., a one-to-five star rating). More dimensions allow for more sophisticated matching to the requirements of clients or the task. For example, if there is a timeliness reputational metric, and schedule is a key concern, selecting those entities that excel on this dimension is possible, even if those entities have lower ratings on other dimensions.

First, we list, without limitations, some of the potential dimensions/metrics of interest about multi-dimensional reputations, as well as how the reputational metrics might be calculated or estimated and how they might be used:

1. **Timeliness (Punctuality):** Measures how often an expert meets deadlines. Calculated as the percentage of tasks completed on or before the due date. Updated after each project

completion. Experts with higher punctuality scores could be preferred for time-sensitive tasks.

2. **Budget Compliance**: Assesses the expert's ability to work within budget constraints. Calculated by comparing the agreed budget and actual spending. Updated post-project. High scorers can be trusted with financially strict projects.

3. **Quality of Work**: Evaluates the output quality through peer reviews, client feedback, and adherence to specifications. Updated periodically based on feedback. High-quality work leads to a preference for high-stakes or high-visibility projects.

4. **Solution Success Rate**: The percentage of problems successfully solved. Calculated by dividing successful outcomes by total attempts. Updated after each project. Critical for assigning complex issues.

5. **Client Satisfaction**: Measured via surveys and feedback scores after project completion. High satisfaction rates increase the likelihood of being recommended for future client-facing tasks.

6. **Peer-rating of Competency**: Collected through anonymous peer reviews focusing on skill and knowledge. Helps identify experts for mentorship roles or collaborative projects. Implementation can be via the invention, where the task of rating a peer is posed to the system and crowdsourced via the online ad units.

7. **External Reputation**: Based on awards, publications, or external acknowledgments. Maintained manually, or ideally via automated analysis of LinkedIn profiles and other public information. Influences the assignment to projects requiring recognized expertise.

8. **Innovation Score**: Assesses creativity and the ability to generate novel solutions. The novelty was evaluated through peer reviews and client feedback (See 6). Important for R&D or creative projects.

9. **Communication Skills**: Rated by clients and peers (See 6), focusing on clarity, conciseness, and effectiveness. Essential for leadership roles or projects requiring teamwork.

10. **Adaptability:** The expert's ability to handle changes or unforeseen challenges. Updated after projects that undergo significant scope changes. Valuable for dynamic environments.

11. **Leadership Quality:** Assessed through peer and subordinate feedback (See 6), focusing on motivation, guidance, and decision-making. Crucial for projects requiring team management.

12. **Technical Proficiency**: Evaluated by successfully applying technical skills to solve problems. Important for technical or specialized tasks.

13. **Learning Agility**: Measures the speed and effectiveness of learning new skills or adapting to new technologies. Essential for rapidly evolving fields. Automated analysis of the sequence of tasks completed by a human, together with metrics on how similar the tasks are to each other, can yield agility measures.

14. **Conflict Resolution**: Assessed by observing outcomes of conflicts in which the expert was involved. Valuable for team-based projects. An associated metric is how often the expert is in the minority or majority when voting on options, and what percent of time, when in the minority, the expert proved to be correct in retrospect.

15. **Project Management Skills**: Evaluated based on planning, executing, and closing projects effectively. Important for roles with project oversight responsibilities. Metrics might be the number of project management tasks assigned to an expert and the success rate on these tasks.

16. Reliability: Measured by consistency in performance and availability when needed. High reliability increases experts' chances for critical or emergency tasks. Related metrics include response time and metrics associated with items 1-5 above.

17. **Efficiency**: Assesses the resources (time, money) used to achieve outcomes. Efficient experts are preferred for projects with limited resources.

18. **Cultural Fit:** Based on values alignment with the organization or team, assessed through surveys or observations. Influences team-based or long-term assignments. This would be relevant for a longer-term project outside the ad unit (See Section 6.12).

19. **Work Ethic**: Evaluated through peer and manager observations, focusing on dedication and professionalism. High scores are crucial for all types of work. This would be relevant for a longer-term project outside the ad unit (See Section 6.12).

20. **Client Retention Rate**: The percentage of clients who return or continue working with the expert. Indicates client trust and satisfaction, influencing assignments with high-value clients.

21. **Feedback Responsiveness**: Measures how actively and constructively the expert engages with feedback. Important for continuous improvement and learning.

22. **Networking Ability**: Assessed by the expert's ability to build and maintain professional relationships. Influential for roles requiring outreach or collaboration.

23. **Mentorship and Training**: Evaluated based on contributions to the development of peers or subordinates. Important for building internal capabilities. This would be relevant for a longer-term project performed outside of the ad unit (See Section 6.12).

24. **Availability and Responsiveness**: Tracks the expert's availability for new tasks and speed of response to inquiries. Critical for fast-paced or client-driven environments. Online availability and speed of response are related metrics.

25. **Problem-Solving Speed**: Measures how quickly an expert can provide effective solutions. Valued for time-critical projects. See 24.

26. **Creativity and Innovation**: Assessed through the uniqueness and effectiveness of solutions provided. Key for roles requiring out-of-the-box thinking. Given that creativity can be operationalized as novel and valuable solutions, peers can rate human options on these dimensions, with usefulness also correlated to problem success.

Reputational metrics, to the degree possible, should be automatically recorded and updated. For example, time-to-solution and steps-to-solution for problem-solving are easy to automate, as is the percentage of successful solutions as a proportion of attempted solutions. Other metrics, such as client satisfaction, may require surveys or other informational gathering approaches that are well-known in the art. Automated evaluation of external reputations is possible using automated content analysis techniques combined with expert or entity-specific searches on the internet or other public data sources. Of course, techniques for automated background checking, such as those already existing and in use, can also be used as part of evaluating a public reputation.

All of these dimensions of reputation can be used and updated in the current invention from within an ad unit, depending on the specific design of the ad unit, which is customizable. However, they can also be used in the collective problem-solving network and AGI /SI systems to enhance the accuracy and usefulness of reputations in those systems, described in Section 4.5, especially FIG. 7, and cited PPAs and PCTs.

6.11 Problem-Solving Within the Ad Unit

The current invention envisions that the main problem-solving system may be part of an AGI / SI system comprised of many (human and non-human) intelligent entities collaborating and working together, sequentially or in parallel, to solve problems and learn solutions to improve the capabilities of the overall AGI network. This AGI/ SI system has been described in detail in previously cited PPAs and PCTs, Section 4.5, and FIGS. 1-9. The general problem-solving process is described in FIG. 10.

Within FIG. 10, one of the steps is known as “Identify the Operators.” In this step, the problem solver determines which action, or “operator,” might be applied next in a sequence of problem-solving steps to advance progress towards a solution. With regard to the current invention, where one use of the online ad unit is to solicit expertise and knowledge that helps solve problems and also be used to train AI, one exemplary method to implement problem-solving within an ad unit is shown in FIG X-5 and described as follows:

1. The AGI problem-solving system handles all of the steps in the universal problem-solving process described in FIG. 10, except “Identify the Operators.”

2. For very simple problems, or sub-problems, which can be described briefly, the AGI system describes the current state of the problem and the next goal or sub-goal in the online ad unit.

3. The human user, who is viewing the online ad, enters a suggested next action (“operator”) that might be taken to either solve the problem or sub-problem or advance progress on the problem.

4. The user input is communicated to the AGI system, where it is processed according to the universal problem-solving methods, potentially in combination with input from other users viewing online ads or other intelligent (human or AI) entities participating on the problem-solving network.

5. Based on processing in step 4, the system refreshes the online ad unit dynamically with the results of the processing. This refresh can include, without limitation:
    a. Whether the user’s input was accepted or rejected.
    b. Whether the goal or sub-goal has been achieved.
    c. What input (e.g., of an action or operator) was selected by the system to advance problem solving, including (optionally) reasons or explanations of the selection.
    d. A description of the new problem state after applying the selected action or operator. This problem state may include new or updated goals/sub-goals, lists of

operators, images, metrics, and other information sufficient to describe the new problem state.

e. The amount of credit or payment that the user of the online platform has accumulated based on their input.
    i. An additional request for new user input via the online ad unit or via a link that takes the user to a more fully functional interface for participating in the problem-solving.

6. If the problem/sub-problem is solved, the user is credited for their contribution; if the problem /sub-problem is not solved, additional user input is solicited (5f), and the method repeats from step 1, OR the user may exit the ad unit and stop participating in problem solving.

7. When the user exits the ad unit and stops participating in problem solving, the user can optionally receive compensation if any is due or allow compensation credit to continue to accrue in the user's account, AND applicable user's reputational metrics are updated and stored in association with information identifying the user to the system.

8. The learning process of the AGI system, as shown in FIG. 5, increases the knowledge of the overall AGI system based on the successful or unsuccessful solution attempts by the users, including users participating in problem-solving via online ad units. In this way, the online ad unit acts as a valuable means of not only solving a particular problem but also increasing the knowledge, intelligence, and capabilities of the overall AGI system.

It should be apparent from the above example, and the "double value" that accrues from Step 8, that online ad units configured to support problem solving, have a much greater value than the typical online ad unit that is used to sell goods or services, and therefore represents a significantly superior means for monetizing online human attention.

The applicant notes that the use of the online ad unit to specifically acquire knowledge and expertise from humans about actions ("operators") is only one way to advance problem-solving and increase AI intelligence via the problem-solving approach to monetization. Without limitation, any of the problem-solving steps shown in FIG. 10, for example, could also benefit from human expertise solicited and communicated via an appropriately modified online ad unit.

Specifically:

- If problem-solving is at the stage of representing the problem and defining the problem space, then the online ad unit could ask users to describe via text, speech, or visual means how they would frame the problem. This description could then be translated into the language of the universal problem-solving framework (with problem states, operators, goals, etc.) using, for example, the Natural Language to Problem-solving Language Translator of FIG. 6. that is described in Section 4.5 and in other cited PPAs and PCTs.

- If the problem has already been represented and defined and is at the stage of applying means-ends analysis or other heuristic methods to determine which sub-goals to set, the online ad unit could ask users, for example, to describe via text, speech, or visual means which intermediate goals they would set as "stepping stones" towards the final solution, and these intermediate sub-goals could be communicated to and processed by the AGI problems solving system. The user might reason using means-ends analysis in which case they are contributing their reasoning knowledge; or they might reason using specific heuristics (such as thinking of a how they solved a similar problem and setting sub-goals, accordingly, using a more powerful form of knowledge than general means-ends analysis) in which case they are contributing their specialized knowledge and expertise towards solving the task.

- The AGI system may require user input on the safety or ethical implications of setting or pursuing certain goals, in which case an appropriately modified online ad unit can be used to solicit opinions on the safety or ethics of proposed goals or sub-goals. This use of online ad units allows real-time, dynamic safety and ethics oversight from humans – something that might be especially important in increasing the alignment, safety, and trustworthiness of the AGI system.

- Users may be asked to vote/rank/rate potential goals/sub-goals or actions ("operators") that have been proposed by themselves and/or other intelligent entities via an appropriately modified online ad unit. In this regard, all of the methods for voting, weighted voting, etc., that have been disclosed in great detail in previously cited PPAs and PCTs could use the online ad unit as an important means for obtaining dynamic real-time information on which goals/sub-goals or operators to use in problem-solving efforts.

- Generally, as discussed above, the learning methods of FIG. 5 and the last box of FIG. 10 can leverage input provided by human users via online ads to increase the learning and intelligence of the AI/AGI/SI agents and systems.

There is a special case of problem-solving via online ad units that the applicant wishes to explain to help clarify the potential scope and value of the invention. A specific type of problem is the problem of providing feedback to LLMs or AI agents to help them learn. Although large quantities of data (e.g., behavioral data or content scraped or obtained via the internet or within an organization's data systems) are available and already used to train LLMs and other AI agents, often this data is of mediocre quality. That is not surprising, since although (as in the mythical Lake Woebegone) everyone likes to think they are "above average," by definition, that can't be true. The result is that it is relatively easy to get LLMs or AI agents to perform on specific tasks at an "average level", but it is exceedingly difficult to train them to perform at a level that is significantly superior to average unless exceptionally rich (in expertise) data is used to train.

The invention of using online ad units to identify and acquire information from humans that possess such superior expertise is therefore particularly helpful (and valuable) in addressing the difficult task of producing expert performance in LLMs and other AI agents. This task of acquiring exceptionally valuable data can be viewed as a special type of problem that the methods described above can help solve, especially using the capabilities to target online very narrowly to those humans possessing the required levels of expertise and then using tailored online ad formats to extract this information for use by the AI agents that need to be trained. This use case alone can increase the monetization of online ads at least 10X from current levels.

A second clarification is important for the future, in which AI agents routinely act on behalf of human users to represent their interests online. In such scenarios, which are becoming increasingly common and will likely represent the majority of internet traffic in the future, it is not humans who will predominantly be exposed to online ads, but rather AI agents.

Therefore, in the description of methods and exemplary implementations above and throughout this entire disclosure, where the applicant discusses "human users" or "users," what is meant in the most general, powerful, and preferred implementations of the invention is really "intelligent entity users." That is, an AI agent may encounter an online ad unit soliciting knowledge, expertise, or information to help solve a problem and/or train other AIs. If the AI agent has sufficient intelligence, its contribution via the online ad unit can be as valuable, or in some cases even more valuable, than that of a human expert.

In the preferred implementation of such cases, the AI agent would be required to identify itself as such so that the system can opt to exclude its input if desired, especially for safety or ethical problems where human opinions are sought. It is worth noting that since it is impossible to logically derive ethical values and determine what is right or wrong from reason alone, in a future in which most AI agents are more intelligent than most human experts, it may be that the primary type of knowledge or information that is solicited via this invention, and other means, is

ethical information since that depends, to put it poetically “on the human heart,” and not on intelligence alone.

6.12 Problem-solving Outside of Ad Unit

It is important to understand that all problem-solving does not have to happen within the online ad unit. In the preferred implementation, the online ad units have links to webpages, apps, AI agents, and other means outside of the limited space confines of the online ad unit itself. The primary function of the online ad unit is to attract the users’ attention and engage them in the task of contributing some of their knowledge.

Generally, users are willing to provide a single click of information, such as a YES / NO response to an ethical question, or perhaps a selection from a list of options, within an ad unit. However, if they wish to engage in more extensive and sophisticated problem solving, although such engagement is supported by the online ad unit as described above, usually a more efficient means would be to direct the user from the online ad unit (e.g. via a link) to a more specialized user interface that can provide more problem-solving functionality in a more usable manner.

Therefore, this invention includes all of the problem-solving methods and systems described in Section 4.5, in the associated FIGS., and has been detailed in the cited PPAs and PCTs that are focused on such systems and methods.

Even securing a single click of information from the right user at the right time is far superior to monetizing a single click via the conventional monetizing methods of taking a user to a landing page with a sales pitch. The applicant demonstrated this fact in a previous company, whereby the value of asking users’ opinions on the direction of stock movement proved to create 100X the cost of the online ad unit, even without specific targeting of the online ad unit to experts in the stock market. However, to maximize the value of the invention, targeting the online ad unit to experts who can provide the right expertise at precisely the right time will greatly increase the value and usefulness of the ad unit and represents a far superior monetization than anything that currently exists.

6.13 Feedback Mechanism to Online Ad Targeting

Current online ad targeting systems are optimized to increase “clickthrough” on the ads, primarily because this is the main information that the advertisers have. Clients are often reluctant to share additional information, such as how often a click on an ad “converts” into a sale, although the clients themselves typically measure this information. It is understandable that clients are secretive about conversion rates because they may fear that if the true value of the advertising were revealed to the online ad vendors, then prices would rise accordingly.

However, in the preferred implementation of the current invention, there is a feedback loop between the reputational and other metrics that assign credit or blame to various problem solvers and the targeting mechanisms that direct online ads to particular experts. This feedback loop is practically implementable, with no conflicts of interest, for any of the large online ad vendors that wish to use their in-house online ad capabilities to increase the intelligence of their AIs.

FIG. X-6 illustrates the basic feedback mechanism, which follows this process:

1. Online ad company uses information about users (including without limitation, user profiles, cookies, user preferences, behavioral data, purchase history, time on site, time on page, time on apps, CTR, available conversion metrics, browsing history, viewing habits, email and text content analysis, prompts to AI agents, and other metrics well known in the art) to target experts that may be able to contribute information useful to solving problem(s).
2. Response rates are measured and recorded for users who contribute information within the online ad unit(s) (as described in this invention) or via links to the problem-solving and/or AI training systems outside the online ad unit.
3. Metrics related to the quality of information and/or problem solutions provided by the targeted users are recorded. These metrics can include, without limitation, all of the reputational metrics described in Section 6.10, as well as all the metrics described in cited previous PPAs and PCTs including those that describe the informational metrics and KIT-theory based metrics detailed in the PCT entitled “Catalysts for Growth of SuperIntelligence”, as well as other information quality metrics commonly used in the field of Machine Learning and the field of Intelligent Online Ad Targeting, and problem-solving generally.
4. Correlations and other statistical relationships (e.g. regression weights) are calculated on the targeting factors in (1) to determine which factors have the highest impact on important quality metrics in (3), which metrics can be prioritized by the organization, especially about which information, knowledge, or expertise, was most effective at meeting or maximizing problem-solving and/or AI training objectives.
5. The weight given to specific targeting factors and/or the specific targeting process and algorithm parameters in (1) are adjusted based on the analysis in (4), and the improved targeting information, parameters, methods, and algorithms are used to target future online ads as the process repeats from step 1.

6.13a Keys to AI Dominance

Google/Alphabet, Meta, Amazon, Apple, and many other large technology companies that sell online advertising are also extremely focused on gaining a competitive advantage in the field of AI agents and AGI. A hitherto unrecognized advantage that these companies and organizations possess, if they use the current invention, is the ability to use their own online ad targeting and display capabilities in service of not only solving problems on a collective network of intelligent entities, thereby allowing them to immediately deploy AGI services, but also using their ad targeting capabilities in combination with the current invention to train their AI agents (irrespective of whether they wish to take the next step of AGI as described in other inventions by the applicant).

Given that the means to obtain high-quality training data is rapidly becoming the bottleneck in the AI "arms race," the current invention represents a novel and extremely useful tool for these companies and organizations. Furthermore, these companies have no impediments to creating a closed feedback loop in which the usefulness of particular experts and expertise for training their AI agents can be directly correlated with the information used to target and find these experts. Therefore, these organizations have all the components necessary to fine-tune their online ad targeting systems to maximize the objective not of "clickthrough" or even of "conversion" but rather the objective of delivering the maximally useful information, from an AI agent training perspective, for increasing the intelligence of their AI systems.

It is difficult to overstate the magnitude of this advantage. All the major tech companies that sell online ads are currently spending hundreds of billions of dollars to buy NVIDIA's chips so that they can increase the power and intelligence of their AIs. This situation concentrates the power in the hands of the chip companies, such as NVIDIA. Indeed, as I write, chips are backordered by more than a year, and margins are extremely high on these chips, which are currently seen as the key to dominating the multi-trillion-dollar market that AI represents.

But there is an alternative to remaining at the mercy of the chip makers – and I am not talking about getting into the chip-making business, which all of the major players are doing. Instead, I am talking about leveraging online ad targeting capabilities to secure high-quality data, and specifically THAT particular data in particular areas where the companies need to increase the intelligence of their AI agents. 1 TB of high-quality data is more valuable than 100TB of mediocre data. It requires 1% of the computing power (chips) to train the AI agents as well!

Simply put, the large tech companies, at least those with significant online ad capabilities such as Alphabet, Meta, Microsoft, Apple, and Amazon, are overlooking the fastest and best way to increase their lead in the AI space, namely by using their existing technology, in combination with the current invention, to accelerate the learning and intelligence of their AIs, without

increasing their need for chips or computational power. Of course, if they can both obtain higher quality data AND lots of computing power, that is best. However, as described, better data with less computing can achieve superior results to mediocre data with more computing. In short, certain organizations are sitting on the keys to AI dominance without realizing it, and the current invention, together with those detailed in previously cited PPAs and PCTs, are the keys to leapfrogging the competition!

6.14 Feedback Mechanism / Process for Attention Spot market

Unlike many items traded on markets, where it is easy to standardize grades of the items (e.g., AAA-rated bonds), human attention and expertise can vary widely in their quality and usefulness. As described above, objective, automated, and/or third-party methods may be used to try to group or rate human attention (e.g. unskilled, uneducated humans vs. college-educated, humans with software development experience) there is wider variability in the utility of a given human's attention then there is, for example, in grades of crude oil, corn, or soybeans. Fortunately, AI excels in recognizing patterns and making judgments, such as creditworthiness or selecting humans based on resumes. Further, the problem-solving and AI training tasks have objective metrics of quality and success (detailed above) that can be used to more precisely estimate the value of a given human's attention for a particular task than can be done, for example, with a bushel of soybeans.

An important element, therefore, in accurately pricing human attention is knowledge of the performance track record of the particular human(s) and the nature of the task the human is being asked to perform. Using methods from KIT (described in earlier PPAs and PCTs) it is possible to estimate the contribution to an AI's intelligence, for example based on comparing the goals of the AI and its existing knowledge base, with the knowledge and expertise of a particular human (or other intelligent entity) and adjusting based on variable performance metrics such as timeliness, trustworthiness, reputational metrics, etc.

With a feedback loop between the metrics achieved by the human on a particular task and the price paid on the spot market, it is possible and desirable for the system to propose a recommended price for any given human's expertise, relative to a particular task, based on the availability of resources for that task. Note that this feedback just helps anchor potential bid and ask prices in the absence of sufficient liquidity for the market price-discovery mechanism to work efficiently.

This feedback might be considered most helpful in objectively characterizing, without limitation, the type or category of human attention, knowledge, or information, and the nature of human attention offered for sale. However, ultimately, the appropriate price for human attention for a particular human is determined by the market price.

If a client overpays and is dissatisfied, that dissatisfaction will be reflected in reputational metrics fed back to the market, which may result in a lower quality grade or classification for that human's attention in future transactions. Conversely, if expectations are consistently exceeded, and the client feels the transaction was a bargain, that feedback might increase the quality grade of that human in the future.

One novel and useful feature of the feedback mechanism associated with this invention generally, and the human attention spot market in particular, is the ability to record in great detail every step (or misstep) in problem solving, to create a completely accurate and transparent "vector" track record of every humans performance on every task and sub-task. This objective and transparent record, which can be implemented via blockchain technology (including without limitation Ethereum-based smart contracts and token as described in other PPAs and PCTs) allows much more precise reputations (compared to commonly used "five star" rating systems for example) which can be analyzed by AI and converted into precise and fair estimates of value-add for each human for any given task.

Thus, based on the feedback mechanism, the same human might command a rate of only $10/hr. For a task where any (or most) humans could do the task as well or better than a human, but might command an hourly rate of $1,000/hr. For a very specific task requiring information and expertise that only this particular human has, and for which the particular human has a demonstrated track record of creating $100,000 worth of value in past projects. Such a feedback mechanism can help both sellers and buyers of human attention pay a fairer price for attention on any given task. It can help both parties make more profitable transactions on the spot marketplace compared to more commoditized markets. In short, humans are unique, and this invention helps them get paid fairly for their uniqueness while also helping clients avoid overpaying for less valuable commoditized attention.

One exemplary implementation of the feedback process for the attention spot market includes the following steps, as also illustrated in FIG. X-7:

1. Purchase of human attention is contracted on the spot market; price and other details of the contract are recorded.
2. Work is performed by humans.
3. Metrics, which may include, without limitation, the reputational metrics of Section 6.10 and other job or task-related metrics, are (automatically) recorded during the work.
4. Payment details and client satisfaction are recorded once the work is completed, and All metrics are stored in (optionally blockchain-technology-enabled) transparent and auditable records associated with the transaction.

5. At the completion of work, or periodically as may be practical, correlations and other statistical analysis or machine learning is done to refine the categorization of the human attention associated with the human and as relevant to different categories of tasks; reputational metrics are updated; and estimations used for suggested pricing for particular types of attention (including from particular individuals or categories of individuals) as related to particular types of tasks are updated.
6. The next time tasks, or categories of tasks are placed on the attention spot market, updated metrics and other information (e.g., from step 6) are used to better match human attention to tasks and suggest estimated pricing; however, market mechanisms ultimately still determine the price at which transactions occur on the spot market.

6.15 Continuous Improvement Mechanisms

Similar to the feedback mechanism and process for the spot market, described in Section 6.14, the effectiveness of all steps in the overall online ad unit invention, including but not limited to: the targeting of the online ad; the size, shape and overall design and content of the online ad unit, the placement of the online ad unit, the frequency with which the online ad unit us displayed, and other aspects of the presentation and process flow, can be measured with metrics, analyzed and optimized using statistical and machine learning process, and continuously improved.

6.16 Dynamic Arbitrage Process

Online ad companies themselves, intermediaries that broker the sale of online ads to clients, and AI agents or systems that seek to improve their intelligence, may engage in a form of dynamic arbitrage as follows:

1. Estimate the value of a particular type of human attention and knowledge with regard to a specific task (e.g., a problem-solving task, a sub-task, or an AI training task)
2. Estimate the cost of obtaining the required amount of human attention required to complete the task from the appropriate intelligent (human or non-human) entities via the online ad unit invention, including the attention spot market, fixed or variable prices being charged by online ad companies via their various platforms and technologies, or other means. This cost should also include a pro-rata share of other costs involved in conducting arbitrage by the arbitraging entity.
3. If the estimated value in (1) exceeds the estimated cost in (2) by a predetermined or dynamic variable representing the profit margin, then purchase the attention at cost approximating that estimated in (1) and sell the attention to clients or use the attention to perform tasks that create value approximating that estimated in (2) and book the profit.

4. Repeat from Step 1, performing arbitrage in priority order of the largest arbitrage opportunities first, until the minimal acceptable arbitrage opportunity (potential profit) is reached.

The applicant believes that the highest value for human attention of experts possessing unique data will likely come from applying that expertise to train AI agents that wish to increase their intelligence. Since an AI agent or system, once trained, can re-use its knowledge and expertise for as long as the knowledge and expertise remains valid, such entities, or their owners, can afford to pay a higher price for human attention used in this way than other clients that hope to just sell a single product or service or solve a single problem once.

As long as more intelligent AI entities (including AI agents, AGI systems, and SI systems) can improve themselves and increase their intelligence and capabilities further with the aid of expert human attention, these entities will likely find ways to generate increasing amounts of money which can be used to further increase their intelligence in a positive feedback loop until humans are no longer a useful source for increasing their intelligence. The use of the online ad unit invention is therefore a catalyst for the growth of these AIs' intelligence.

While the companies providing access to the online ad units, and engaged in attentional arbitrage using these units, stand to earn vast sums of money, it is important that there are safeguards in place with regard to the types of tasks that are allowed in the online ad units, and on the attentional spot market. Terrorists (humans or AIs) might pay an extremely high price for specific expertise related to weapons of mass destruction, for example, and this must be prohibited.

6.17 Human Worker & Client Interfaces

User interfaces for human workers and clients can include LLMs and natural language text or audio-based interfaces in which humans just speak to AI agents, and the AI agents translate what is said into an underlying universal problem-solving framework that coordinates problem specification and solving activity. The natural language translating process of FIG. 6 is relevant in this regard. However, other, more specialized interfaces can be built as well.

The guiding principle is that the steps of problem solving, especially those described in FIGS. 4, 6, 8, & 10 can each have specific user interfaces to optimize problem solving, within or outside of an online ad unit. Many variations such as text boxes, dropdown lists, use of templates, dynamically sizing input and output areas, visual I/O devices including virtual reality devices (e.g., Apple’s Vision Pro or Meta’s VR systems), are all well known in the art and can be adapted to accommodate and optimize the ease of accomplishing the problem-solving steps listed in this and previously cited PPAs and PCTs.

6.18 Automated AI / AGI interfaces

In addition to using interfaces for humans (e.g., human clients or workers), the current invention accommodates any intelligent entity in the roles of both workers and clients. Again, natural language (e.g., generated by LLM agents) can serve as a universal interface.

Multi-modal AI agents can communicate problem specifications and representations/solution ideas via images, audio, and even by using sensory and output tools that extend perception and generation beyond the range of humans.

For example, an AI agent working to solve a problem could use X-rays to detect the problem of a broken bone that would be invisible to human eyes, and that same AI agent could output a solution in the form of a 3-D printed cast, or by formulating a pain-killing or healing chemical compound, which are forms of output that unaided humans are incapable of generating.

Thus, the interfaces for non-human entities generally include all of the modes that work for human intelligences as well as APIs, digital interfaces, and other interfaces that might enable them to optimize the efficiency of problem specification and solving.

6.19 Recursive Use of Problem-solving to Optimize Ad Targeting and System Efficiency

One of the novel and extremely powerful aspects of the invention is the ability to solve cognitive problems of any type. Thus, although Sections 6.13 – 6.15 describe specific approaches to use feedback and other specific methods for improving aspects of the invention, a general way to optimize aspects of the invention, including without limitation ad targeting and overall system efficiency, is simply to specify these tasks as problems to be solved, and then let the system itself solve the problems of how to make these improvements. This recursive use of the invention to improve the invention is a unique capability that exists with almost no other inventions, and which reflects the universality of the problem-solving capabilities that underlie the invention, and that are described in this disclosure and previously cited PPAs and PCTs.

6.20 Safety and Ethics Checks

While the companies providing access to the online ad units, and engaged in attentional arbitrage using these units, stand to earn vast sums of money, it is important that there are safeguards in place with regard to the types of tasks that are allowed in the online ad units, and on the attentional spot market. Terrorists (humans or AIs) might pay an extremely high price for specific expertise related to weapons of mass destruction, for example, and this must be prohibited.

FIG. 8 describes scalable safety checks (e.g., each time a goal or sub-goal is set) for problem-solving systems that form the basis of AGI or problem-solving by AI agents, AAAIs, and broadly,

other intelligent entities. Similar scalable checks should be built into the processes for creating tasks or sub-tasks (by any intelligent entity) that appear either within an online ad unit or on the attentional spot market. That is, before a task is allowed to appear in either place, an ethics check must be run, broadly following the same types of process steps, shown in FIG. 8, including the final steps of creating a transparent and auditable record and continuously improving the safety-check system.

6.21 Regulation Compliance

The same processes described in Section 6.20 and FIG. 8 can be used, alone or in combination with other existing methods known in the art for curating content or detecting and removing prohibited content from platforms, to ensure that the invention complies with regulations and laws in the geographies or contexts in which it operates. AI agents, specifically trained to detect offending or non-compliant content, might also be used to assist with regulatory compliance.

6.22 Universality of Systems and Methods Across Platforms and Cultures

It should be clear from the above discussion that the online ad unit invention is applicable across any platform or technology that supports online advertising. Similarly, just as online ads and related technology are localized to account for different languages and cultures, the same can be done, using methods well known in the art, for the current invention.

One of the unique and innovative features of the invention is that it provides a means to access human attention, knowledge, and expertise from almost any human, anywhere on Earth, thanks to the already existing prevalence of online advertising. By leveraging the existing online infrastructure, which has been developed at a cost of many billions or even trillions of dollars over the last twenty years or so, the invention is able to leverage the collective intelligence of billions of humans and combine that with non-human entities to solve any problem (via the collective intelligence AGI system described in Section 4.5 and previously cited PPAs and PCTs) and train any AI, in any field of cognitive endeavor.

6.23 Collaborative and Cross-Ad Unit Dynamic Coordination

Another novel aspect of the invention is that different problem solvers can be working on different aspects of a problem in sequence or in parallel, as coordinated by the overall flow of problem-solving in the collaborative AGI system.

For example, multiple online ad units could simultaneously solicit the “next step” in a problem from multiple humans. Multiple humans could submit their suggested next steps via multiple online ads without necessarily being aware of the other submissions. Then a second batch of

online ad units could present the options for next steps acquired from the first batch to the same or different humans via online ads, asking these humans to vote for the preferred next step in the problem-solving. That is, all the steps of collaborative problem solving, sequentially or in parallel (as illustrated in FIGS. 11 & 12), including using the problem tree structure of FIG. 13 for coordination and assembly of sub-steps into an overall solution, can be leveraged by the current invention. The main difference is that the problem-solving tasks are broken up and distributed to solvers via multiple ad units instead of the other types of interfaces specified in the AGI network invention. Further, if entities link from the ad unit to the other types of interfaces, they can also use those interfaces to work on problems if that is easier.

The overall system does not care whether work is done sequentially or in parallel, within ad units or outside of them. The process steps, including the coordination of multiple solvers, are largely the same; it is mainly the location of the work that differs.

6.24 Integration of Realtime and Asynchronous Capabilities / Data feeds

Another novel feature of the current invention is the ability to incorporate real-time and asynchronous data feeds with the ad unit. For example, if the problem posed within the ad unit involves recommending which stocks to purchase, a real-time data feed with stock prices can be incorporated into the ad unit.

If the problem, further involved sending a stock recommendation via text functionality within the online ad unit and waiting for a response before adding additional text about the recommendation in a comment box, the online ad unit could transmit the recommendation with the ability to wait for an asynchronous response from another user (potentially in another ad unit) before proceeding.

Alternatively, some tasks might be done asynchronously outside the ad unit (e.g. via an email system), while others might be done in real-time within the ad unit. Generally, the current invention is designed so that problem-solving can proceed step by step, regardless of whether the steps (and associated relevant information that is provided) occur in real-time or asynchronously when the solver has a chance to respond or provide input.

7.0 PREFERRED IMPLEMENTATION AND VARIATIONS

The invention of an online ad unit for harnessing human attention, or attention from intelligent entities, can be implemented in many ways, on its own, or as part of larger systems. In this Section, the applicant describes a typical Use Case with three variations, including references to some of the systems and methods described in previous sections. These exemplary implementations are meant to illustrate how some of the systems and methods work together. It should be obvious to those skilled in the art of software development, and with some expertise in the field of online ad systems, that many variations are possible, including those that include more or fewer of the methods than in the exemplary preferred implementation.

7.1 Preferred Implementation for Online Ad Companies & Clients

Consider the Use Case of a company engaged in selling online ads whose business model is to maximize the profits from displaying clients, or whose expertise is serving as a broker between the client and the ultimate online ad provider. Companies or organizations that derive a large portion of their revenue from this business model include, without limitation: Alphabet (including its Google Search and YouTube divisions), Meta (including its Facebook, Instagram, and Reels platforms), ByteDance (including its TikTok product), Baidu, X (formerly Twitter), Spotify, Snap, Pinterest, and ad-focused companies (e.g., PubMatic, Magnite, Sea Limited, Criteo, The Trade Desk, Jalopy, Taboola, and Outbrain). Companies selling subscriptions and products and generating significant revenue from advertising, such as Amazon and Apple, are also relevant to this Use Case.

In the example that follows, the company that possesses the systems and methods of the current invention is called “Company,” and the clients that wish to purchase human attention to extract information, knowledge, and expertise are called “Clients.”

The following steps illustrate one preferred implementation that results in increased monetization of online advertising revenue for the Company and delivers valuable information, knowledge, and expertise based on human attention to the Client.

1. The company builds databases containing information about users that help target specific types of online ads to specific users. The means for building the databases include, without limitation, the Company’s existing methods and technology, purchasing user data and information from others, and using this invention’s methods (Section 6.7) for building a database of human experts, which can include their preferences and other valuable information for targeting ads.
    a. If the existing database of human experts is not large enough, or if the Company or Client wishes to find and include more experts, the online ad invention can be

used specifically for this purpose (Sections 6.7, 6.9).

2. Client, [e.g., OpenAI or any company trying to train/customize AI agents (Section 6.3) or solve problems (Section 6.2)], wants to purchase human attention and work from the Company. Client communicates the desired categories of human attention or work, which may include requirements concerning specific populations of humans that can be accessed due to the universal nature of the invention (Section 6.22).

3. When the requirements are communicated, checks are done to ensure the request does not violate regulations or ethical requirements (Sections 6.21, 6.20). The requirements can be communicated via existing means or interaction with a human attention spot market (Section 6.6).
    a. If an attention spot market is used, several variations of auctions and other mechanisms for pricing the human attention and work are available (Section 6.6a, 6.6b, 6.6c).
    b. Further, the performance of the spot market, relative to the Client's needs, can be improved via a feedback loop and processes (Section 6.14)

4. The company can work with the client to design creative content for interactive online ads that capture attention and work with human experts (Section 6.8). Then the Company can deploy those ads via its existing ad targeting and display technology.

5. The interactive online ads can capture attention and work (e.g., problem solving) from within the ad unit itself (Section 6.11) or by linking to webpages or other interfaces (Section 6.17) that are optimized for this capture outside of the ad unit (Section 6.12).
    a. Depending on the nature of the work, the interactive ads may require collaboration across multiple humans (Section 6.23) and/or integration of real-time or asynchronous data capabilities (Section 6.24).
    b. During problem solving, scalable ethics and safety checks are run to ensure that unethical or unsafe expertise is not used, e.g., for the task of training or customizing AI agents (Section 6.20)

6. Based on reputational metrics and methods (Section 6.10) and other metrics related to the problem-solving work or knowledge captured, the online ad targeting can be improved via a feedback loop and associated methods (Sections 6.13, 6.15, 6.19)

7.2 Variation #1 of Exemplary Preferred Implementation Where Company is also Client

In some cases, the Company and Client may be the same organization. For example, in the cases, without limitation, of Alphabet (including its Google Search and YouTube divisions), Meta (including its Facebook, Instagram, and Reels platforms), X (including X social media platform and X.AI the AI division), Tencent (including WeChat/advertising and their AI divisions), Amazon (including advertising. Mechanical Turk, and AI divisions), Apple (including both advertising and AI divisions), the Company has both online advertising capabilities and AI divisions.

These Companies are exceptionally well positioned to succeed in the race to develop the most advanced forms of AI by leveraging their online ad capabilities to secure the required expertise and knowledge that is missing from their AI agents and other AI systems, and to train AI using human expertise that they acquire via a combination of the current invention and their existing online ad (targeting and other) capabilities. These companies have the advantage that they do not need to pay a markup to use their online ad capabilities and can invest in the expenses of displaying this invention's online ad units to rapidly and efficiently increase the intelligence of their own AI agents and systems, including, without limitation AGI and SI systems (See Section 5.5).

The steps in the preferred implementation are largely the same as disclosed in Section 7.1, except that the Company and Client are the same. Further, more tightly integrated feedback loops for improving targeting of ads are possible (Sections 6.13, 6.15, 6.19). Use of the attention spot market is optional, since a large quantity of human attention and expertise is available to the Company/Client without having to bid for it externally. However the Company/Client may wish to supplement its access to human attention via its own online ad capabilities by purchasing additional human attention/expertise using attention spot market mechanism (Section 6.6), in which case feedback gathered by methods associated with the spot market can also be tightly integrated with the Company/Client's AI development efforts (Section 6.14).

7.3 Variation #2 of Exemplary Preferred Implementation Where Focus is AI Safety/Ethics

Assuming that multiple AGI and SI systems are in operation, and that ensuring the safe and ethical operation of such systems is the top priority for governments, organizations, and humanity generally, then a major use of the invention would be to solicit ethical information and to review potential safety issues that arise during the course of problem-solving or other cognitive activity on the part of the AI systems (Section 5.3).

An important variation of the basic implementation outlined in Section 7.1 is to focus the online units on the problems of reviewing, ranking, voting, and otherwise providing human input on potential ethical choices that may have safety implications. These problems can include efforts to gather a representative and statistically valid sample of human values and opinions (Section 5.3a), across many diverse cultures and geographies, leveraging the invention's ability to access humans across many platforms (Section 6.22). They can also include real-time oversight on ethical and safety issues (Section 5.3b). In this area, the ability to access many human opinions in parallel on the same issue and coordinate the responses (Section 6.23) is important. For dynamic tasks, such as weighing in on a pressing ethical or safety issue, the ability to include real-time and asynchronous data in the online ad unit (Section 6.24) is also important.

By using the methods in the Sections cited above to modify the general implementation described in Section 7.1, it is possible to optimize the invention for the purpose of improving AI safety and helping to ensure human-aligned values.

7.4 Variation #3 of Exemplary Preferred Implementation to Include Non-Human Entities

The description of the current invention has mainly focused on implementations to harness human attention via online ad units and the human attention spot market. However, variations of the novel systems and methods are possible where the same inventive mechanisms, processes, and systems can be used to harness the attention of ANY intelligent entity, whether human or AI.

In the general case, the Human Attention Spot Market (Section 6.6) becomes an Intelligent Entity Spot Market. Since AI agents, AGI, and SI systems will, in the future, act as agents for humans (and themselves) representing the interests of other (e.g., human) entities' interests online, these AI systems and agents will encounter online ads just as humans do. Therefore, the online ad units and the attention spot market can be modified to cater to these non-human agents as well as human agents. The main difference in the implementation would be to use AI/AGI Interfaces (Section 6.18) rather than the standard human interfaces (Section 6.17) to the various inventions.

Just as profiles for ad targeting can be developed for individual humans and then used to help target the requisite human expertise, profiles for ad targeting can be developed for individual AI agents and then used to target the requisite AI expertise.

The same distinction between a human and the human's attention (Section 6.4) applies to AIs. An AI is different from that AI. For the purposes of this invention, we are mainly concerned with novel ways for harnessing the attention and cognitive capabilities (including resulting or

associated knowledge, information, and expertise) of the intelligent entity rather than attempting to corral the intelligent entities themselves.

However, just as a database for locating the humans possessing the relevant expertise can be built using methods of this invention (Sections 6.7, 6.8, 6.9) and reputations can be compiled and associated with the humans (Section 6.10), and feedback mechanisms can be implemented to improve targeting of the humans and the invention's operations (Sections 6.13, 6.14, 6.15), so too can all of these methods be adapted to work with non-human (e.g., AI, GI, SI) agents and systems. What is important is to remember that there are many types of intelligent entities, including human and non-human ones.

Similarly, there are many sources of intelligent behavior; traditionally, the sources have been humans, but now and in the future, we must include AI as the source of intelligent behavior as well. The invention is concerned with efficiently locating and tapping the intelligence (including associated information, knowledge, and expertise) to solve problems, especially to train and improve the intelligence of other AIs.

In the near term, humans are the best intelligences for this task as they are the smartest experts. In the future, increasingly non-humans may also be good, and eventually, the best intelligences for this task. The invention, including the online ad units and attention spot market, is a general mechanism for harnessing these forms of intelligence. Thus, the variation where attention is the attention of any intelligent entity (and the interfaces and methods are therefore modified appropriately), and the spot market is for attention generally, and the online ad units capture problem-solving from any (including non-human) intelligent entity, are important to derive maximal value from the invention in the future.

The problem solutions and increases in AI intelligence, which are produced by the invention's novel method of monetizing (human or non-human) attention, are many times the value produced today by online ads, which is why the invention is not only novel, but extremely useful and valuable.

8.0 CONCLUDING THOUGHTS ON INVENTION'S IMPLICATIONS FOR HUMAN SURVIVAL

"Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war."

This statement on AI risk, put out by the Center for AI Safety, has been signed by hundreds of scientists and leaders including: Geoff Hinton (the "Godfather of AI"), Sam Altman (CEO, OpenAI), Demis Hassabis (CEO Google/DeepMind), Dario Amodei (CEO, Anthropic), Ilya Sustskever (Chief Scientist, OpenAI), Bill Gates (Gates Ventures), Albert Efimov (Chief of Research, Russian Association of Artificial Intelligence), Eric Horvitz (Chief Scientific Officer, Microsoft), Mustafa Suleyman (CEO, Inflection AI), Emad Mostaque (CEO, Stability AI), Martin Rees (Professor of Physics, Cambridge University), Ray Kurzweil (Principal Researcher and AI Visionary, Google), Nicholas Dirks (President, The New York Academy of Sciences), Dan Hendrycks (Executive Director, Center for AI Safety), Yoshua Bengio (Professor of Computer Science, U. Montreal / Mila), Ted Lieu (Congressman, US House of Representatives), Ya-Qin Zhang (Professor and Dean, AIR, Tsinghua University), Igor Babuschkin (Co-Founder, xAI), Shane Legg (Chief AGI Scientist and Co-Founder, Google DeepMind), Yi Zeng (Professor and Director of Brain-inspired Cognitive AI Lab, Institute of Automation, Chinese Academy of Sciences), and Kevin Scott CTO, Microsoft) – to name just a few.

Indeed, as the applicant writes this disclosure, it has become fashionable among AI researchers to ask, "What is your p(doom)?" – by which is meant: "What is your estimate of the probability that AI makes all of humanity extinct?"

Unfortunately, the field of AI has become increasingly divided between what might be called "AI Doomers" and "AI Boomers." The AI Doomers, mainly from academia, espouse the view that AI has a very high probability of making humanity extinct. They see nothing that can be done and have invested considerable energy calling for pauses or halts to AI development. These calls are largely ignored by the AI Boomers.

The AI Boomers, many of whom signed the Statement on AI Risk quoted above, and many of whom also think that there is significant chance that AI will kill all humans, seem to espouse the view that whether AI eliminates humanity or not is not within our control, therefore companies might as well make money from AI while they can. The applicant has noticed that AI Boomers come predominantly from industry and business environments where there is a huge incentive to participate in the "arms race" to develop the most advanced and profitable forms of AI.

The applicant believes that both the Doomers and Boomers are missing the main point. While it is true that there is little chance to halt the progress and rapid development of AI, it is both incorrect and extremely dangerous to think that nothing can be done to develop safe AI. In fact, it is the responsibility of the designers and inventors of AI systems to ensure that the systems are safe. Understandably, when researchers don't know how the existing LLMs and AI agent systems they have designed really work, let alone how to design AGI or SI, they also have no idea how to make these systems safe.

The applicant is fortunate to have studied AI for decades, including under the tutelage of two of the original inventors of the field and also under a third pioneer of today's machine learning techniques. Further, the applicant is extremely lucky to have spent nearly two decades designing and implementing collective intelligence systems, so that he recognizes the power of this approach. Finally, having written a book on software quality, the applicant also gained an appreciation of the importance of ensuring safety and quality at the design stage, rather than trying to test safety in later. All of these experiences have made it clear to the applicant that a path to advanced AI exists that greatly favors human survival.

However, the applicant's experience as a CEO has impressed upon him that safe designs will not be pursued unless these designs are also the most effective and profitable designs. Therefore, the applicant has dedicated himself to inventing systems and methods that are not only the safest designs possible but also designs that will allow the companies that adopt them to dominate the field of AI.

The series of cited PPAs and PCTs describes this "collective intelligence of intelligent (human and non-human) entities" approach to advanced AI, AGI, and SI systems. The current invention, although valuable in its own right as a way to greatly increase the monetization of online ads, plays an important role in increasing the safety and alignment of these systems.

Specifically, high-quality information and data, not only data reflecting technical expertise and knowledge, but also data reflecting human values and ethics, are the next bottleneck in developing advanced AI. Once the major companies and countries have sufficient computational power ("chips") to train models with huge numbers of parameters, they will quickly find (as indeed the leading researchers already acknowledge) that high-quality training data is the most important requirement for increasing intelligence.

The current invention shows how to gather that data in a cost-effective and highly efficient manner. The invention can generate huge amounts of revenue for online advertising companies, and even larger profits for companies that use the data to increase the intelligence of their AI systems. The online ad unit inventions described in this disclosure have not yet been implemented, so the opportunity exists to ensure that safety and ethics checks are built into

these systems. Not only technical expertise, but also values, ethical knowledge, and preferences must be gathered from humans as the invention is rolled out.

Fortunately, the largest online advertising company in the world, Alphabet, has “Don’t be evil” as its motto. Alphabet also has arguably the most advanced AI capabilities on Earth. Therefore, Alphabet is ideally positioned to implement the current invention for the benefit of humanity. If Alphabet fails, Meta, Amazon, Tencent, or other competitors might succeed. None of these companies wants humanity to go extinct. What good are profits if no one is around to enjoy them?

The invention provides a path to not only huge profits but also safe AGI.

Human values and ethical preferences must be at the center of any advanced AI system if we want that system to be aligned with human values. The way to ensure this alignment is not to write a few safety rules or an ethical constitution and hope that AI will not modify it; nor is it to attempt to test safety in via “guardrails” and RLHF applied after the fact to LLMs and other AI agents. Rather, the best approach is for human values and ethics to be inextricably intertwined with all of the knowledge that makes the AI systems intelligent in the first place.

When an AI system learns to solve the problem of a noisy neighbor, it can learn (from humans) that killing the neighbor (although it technically solves the problem) is not an acceptable way. When the AI system learns how to solve the problem of engineering a less expensive car, it can learn (from humans) that skimping on the braking system or safety features is not an acceptable option. When the AI system learns how to hold a respectful conversation, it can learn (from humans) that profanity is not generally acceptable. In millions of different contexts, as it learns to solve millions of different problems, with the help of humans, AI can also learn the safe and ethical ways (in the opinion of the humans) to create its solutions.

Sometimes humans will make mistakes. Sometimes, there will be malevolent teachers. However, the vast majority of learned behavior will be ethical and human-aligned, **IF** the AGI system is designed to learn from humans in this way. Humans must not only be “in the loop,” they must be part of the system’s fundamental design and operation.

Fortunately, using human intelligence to bootstrap not only AI expertise but also AI values is not just the safest and most ethical approach. It is also the fastest, most efficient, and most profitable way to achieve AGI. This alignment between profit and safety is the primary reason for the applicant’s optimism about humanity’s future.

We must be neither Doomers nor Boomers, but rather designers and inventors who are clear-eyed and focused on safe and profitable designs. This approach is the best way to maximize chances not only of human survival, but also of future peace and prosperity for all humankind.

If you had the interest and perseverance to read this entire application, chances are you also possess some of the skills required to improve the chances of humanity's survival. The future of eight billion people, and all of their descendants, lies partially in your hands. Tag, you're it.

